IACR Communications in Cryptology
ISSN 3006-5496, Vol. 3, No. 3, 36 pages.

https://doi.org/10.62056/aebnjbhdj

Check for updates

Forensic categories: a framework for SQIsign-like primitives

Andrea Basso[1], Luca De Feo[1], Sikhar Patranabis[2], Ilinca Radulescu[3] and Benjamin Wesolowski[3]

[1] IBM Research Europe, Zurich, Switzerland
[2] IBM Research India, Bangalore, India
[3] ENS de Lyon, CNRS, UMPA, UMR 5669, Lyon, France

Abstract. Using the language of categories, we introduce a novel framework abstracting the key algorithmic features of the Deuring correspondence between supersingular elliptic curves and quaternion orders and of the post-quantum signature scheme SQIsign based on it. We then show how to construct an interactive identification scheme and digital signature within this framework, and also instantiate more advanced primitives, such as a chameleon hash function. We present two distinct instantiations of the framework based on isogenies of supersingular elliptic curves: the first recovers (one-dimensional) SQIsign, while the second leads to the recently published SQInstructor.

Keywords: Isogeny-based cryptography · Applications of categories · SQIsign · Chameleon hash function

1 Introduction

From pairing groups [BF03, Jou04] to LWE [Reg05], abstractions have proven to be a powerful tool in cryptography to design advanced protocols without the burden of a heavy mathematical language. They also help unify language and transfer techniques from one subdomain to another.

Isogeny-based cryptography is reputed for its hardcore mathematical background and steep entry barrier; thus, it is no surprise that abstractions of it have become a popular topic. The most successful one is the framework of *cryptographic group actions* [BY91, Cou06, ADMP20], which abstracts the theory of *complex multiplication* and that of *supersingular oriented curves* [CK20] to construct a wide array of schemes ranging from key exchange [RS06, Cou06] to signatures [Sto12], threshold schemes [DM20] and much more. More recently the same framework has also been applied outside of isogeny-based cryptography to construct signatures and other advanced schemes based on error-correcting codes [BMPS20, BBMP24, KLP26].

There is a dichotomy in post-quantum isogeny-based cryptography, however, and cryptographic group actions only cover half of the story. First introduced in [CLG09] as a cryptographic tool, isogenies of (non-oriented) supersingular curves have been used to construct key exchange [JD11, FMP23, BF23], encryption [BMP23, BM25], signature schemes [GPS17, DKL+20, BBC+25] and more [Ler25], none of which can be appropriately modeled by cryptographic group actions. Among these, SQIsign holds a place of regard, owing to it being the only isogeny-based candidate still running in NIST's post-quantum selection process [AAA+24], but also because its security is based on one of the most fundamental problems in the field, the endomorphism ring problem [Koh96, Wes22].

E-mail: ilinca.radulescu@ens-lyon.fr (Ilinca Radulescu)

This work is licensed under a "CC BY 4.0" license.
Received: 2026-05-05 Accepted: 2026-09-25

Our contributions. In this work we introduce a framework that abstracts and generalizes the fundamental properties of SQIsign, and we also construct more advanced primitives based on it: a signature which mimics SQIsign, and a novel chameleon hash function (not previously known in isogeny-based cryptography). We present two instantiations of the framework: one which recovers SQIsign, and one that is new, albeit also based on isogenies of supersingular curves. We hope this new framework will help design more advanced protocols based on SQIsign, but also that it will find other instantiations that are not based on isogenies of elliptic curves, enabling a transfer of knowledge between fields.

Technical overview. SQIsign is based on the Deuring correspondence, an equivalence of categories between supersingular elliptic curves and left ideal classes of maximal orders of a quaternion algebra. We discuss this in more detail in Section 2.2.1. Only one word in the previous sentence is relevant to cryptography: "category". Indeed, most of SQIsign may be explained in terms of objects and morphisms, forming a sort of generalized graph. We want to reassure the reader: although categories have a reputation of being technical and excessively abstract, here we only need the very basic definitions from the theory, and they are all summarized in Section 2.2.

We start by defining a fairly straightforward list of computational axioms a cryptographically useful category should satisfy: representations of objects and morphisms, composition of morphisms, and similar. We then add the requirement that the category can be "walked": it must have an *origin* object where walks can start, and it must be feasible to select a random morphism originating at one object and *step* to its destination, similarly to a walk in a graph.

These axioms alone do not characterize SQIsign particularly well, and indeed they could also be taken as a (pedantic) starting point for cryptographic group actions. The key insight of this work is adding to the category something we call a *fingerprint*: a finite set of labels attached to the morphisms of the category. Fingerprints must be efficiently computable, and they capture the essence of a morphism: it must be generically difficult to construct two morphisms between the same two objects with distinct fingerprints, unless one possesses some sort of trapdoor. The trapdoor in question will be the knowledge of a path from the origin object to the source of the morphism.

Fingerprints will serve as challenges in the interactive identification protocol underlying the signature. Because a fingerprint "incriminates" the signer, we call these categories *forensic*.

Fingerprints are a defining feature of SQIsign: the larger the fingerprint space, the lower the soundness error of the protocol, the more compact the signature. The availability of an exponentially-sized fingerprint space for SQIsign is responsible for its best-in-class public key and signature sizes, a crucial feature that cryptographic group actions do not share with forensic categories.

Remark 1. The reader versed in categories may wonder how we got from an equivalence of categories to a single category with some axioms. We give a detailed explanation in Section 3.

Besides our main contribution of defining forensic categories, we also achieve the following results:

- We build some primitives on forensic categories: a signature and a chameleon hash (which in turn implies chameleon signatures [KR98] and trapdoor collision-resistant hash functions). We highlight that the latter is a novel construction not previously known to be achievable in isogeny-based cryptography.

- We give two instantiations of forensic categories. The first, based on kernels of smooth-degree isogenies, leads to the original instantiation of SQIsign. The latter, based on level structures, is entirely novel and gives a first example of a *homomorphic*

forensic category. Although we could not leverage the homomorphic property in this work, we hope that future work will build upon it.

Related work. Our second instantiation based on level structures makes an appearance in the recent work SQInstructor [BFLS26]. While our focus is on formalizing an abstract framework and using it to build new protocols, this concurrent work only considers signature schemes and focuses on the algorithmic aspect.

Paper outline. After some background material in Section 2, and an in-depth discussion of the Deuring correspondence in Section 3, we define our forensic categories in Section 4 and describe the primitives built on them in Sections 5 and 6. Finally we discuss two instantiations of forensic categories in Section 7.

1.0.1 Acknowledgments.

Radulescu and Wesolowski are supported by the Agence Nationale de la Recherche under grant ANR-22-PETQ-0008 (PQ-TLS) and the European Research Council under grant No. 101116169 (AGATHA CRYPTY). Basso and De Feo are supported by SNSF Consolidator Grant CryptonIs 213766.

2 Preliminaries

Definition 1 (Negligible function [Bel02])**.** A function $f\colon \mathbb{N} \to \mathbb{R}$ is negligible if, for every $c \in \mathbb{N}$, there exists an integer x_0 such that, for any $x \geq x_0$, $f(x) \leq x^{-c}$.

In other words, a function $f\colon \mathbb{N} \to \mathbb{R}$ is negligible if it approaches zero faster than the reciprocal of any polynomial.

2.1 Elliptic curves and isogenies

The following preliminaries on elliptic curves are not necessary to the understanding of the framework: Sections 4, 5 and 6 are perfectly understandable without them. We will need them for Section 7, where we instantiate the framework. For a complete reference on elliptic curves see [Sil09].

Let p be a prime number. To instantiate our framework we exclusively work with supersingular elliptic curves over $\mathbb{F}_{p^2}$, a finite field of size p^2. We denote by $\overline{\mathbb{F}}_p$ the algebraic closure of $\mathbb{F}_{p^2}$.

Definition 2 (Isogeny)**.** Let E_1, E_2 be two elliptic curves over a field $\mathbb{F}_{p^2}$. An isogeny from E_1 to E_2 is a non-constant algebraic morphism $\varphi\colon E_1 \to E_2$ preserving the point at infinity.

An isogeny induces a group homomorphism $E_1(\overline{\mathbb{F}}_p) \to E_2(\overline{\mathbb{F}}_p)$. The degree of an isogeny is its degree as a rational map. When an isogeny has degree $\deg(\varphi) = d$ coprime to the characteristic of the field, p, the isogeny is *separable*. The degree of isogenies is multiplicative under composition, meaning that the degree of $\psi \circ \varphi$ is $\deg(\psi \circ \varphi) = \deg(\psi)\deg(\varphi)$. Any isogeny can be factored as a composition of isogenies of prime degree.

For any non-zero $m \in \mathbb{Z}$, the multiplication-by-m map $[m]_E\colon E \to E\colon P \mapsto mP$ is an isogeny. We simply write $[m]$ when the curve E is clear from context.

Every isogeny $\varphi\colon E_1 \to E_2$ has a unique dual isogeny $\hat{\varphi}\colon E_2 \to E_1$ such that $\varphi \circ \hat{\varphi} = [\deg(\varphi)]$, the multiplication-by-$\deg(\varphi)$ map on E_2. Similarly, $\hat{\varphi} \circ \varphi$ is the multiplication-by-$\deg(\varphi)$ map on E_1. Furthermore, an isogeny φ has the same degree as its dual, $\hat{\varphi}$.

If φ is a separable isogeny, the kernel of φ, $\ker(\varphi)$ is a group of order $\deg(\varphi)$. The kernel of $[m]_E$ is the m-torsion subgroup $E[m]$. Separable isogenies are uniquely determined by

their kernel. That is, there is a one-to-one correspondence between separable isogenies (up to an isomorphism of the codomain curve) and finite subgroups of $E(\overline{K})$. Let G be the kernel of an isogeny φ, the isomorphism class of the codomain of φ is denoted by E/G. Then, Vélu's formulas [Vél71] can be used to compute the isogeny $\varphi\colon E \to E/G$ from its kernel G.

If φ, ψ are two isogenies, we write $\psi \succ \varphi$ and say that ψ *factors through* φ if there exists χ such that $\psi = \chi \circ \varphi$. If Ψ is a (potentially infinite) set of isogenies from the same domain, we write

$$\bigwedge \Psi = \{\varphi \mid \psi \succ \varphi \text{ for all } \psi \in \Psi\}. \tag{1}$$

This set has a maximum, i.e. an element $\varphi \in \bigwedge \Psi$ such that $\varphi \succ \chi$ for every $\chi \in \bigwedge \Psi$, and this maximum is unique up to post-composition with an isomorphism. We call any such maximum a *greatest common factor* of Ψ, and we denote it by $\gcf(\Psi)$. Note that if Ψ contains a separable isogeny, then $\gcf(\Psi)$ is separable and $\ker \gcf(\Psi) = \bigcap_{\psi \in \Psi} \ker \psi$.

2.1.1 Endomorphism ring

An isogeny that maps a curve E to itself is called an *endomorphism*. An example is $[m]$, the multiplication-by-m map for m an integer. An endomorphism of this form is known as a *scalar endomorphism* of E. Let $\mathrm{End}(E)$ denote the set of all endomorphisms of E. Then $\mathrm{End}(E)$ is a ring under addition and composition. The unit group is $\mathrm{Aut}(E)$, the group of endomorphisms of degree 1. An in-depth explanation of the endomorphism ring and its structure can be found in [Sil09, Chapter 5, Section 3].

A key concept for the instantiation of the framework is the *efficient representation* of isogenies. As defined in [Wes24, Definition 1.3]:

Definition 3 (Efficient Representation)**.** Let $\mathcal{A}$ be a polynomial time algorithm. It is an *efficient isogeny evaluator* if for any $D \in \{0,1\}^*$ such that $\mathcal{A}(\mathsf{validity}, D)$ outputs T, there exists an isogeny $\varphi\colon E \to E'$ (defined over some finite field $\mathbb{F}_q$, where q is the power of a prime p) such that:

1. on input $(\mathsf{curves}, \mathsf{D})$, $\mathcal{A}$ returns (E, E')
2. on input $(\mathsf{degree}, \mathsf{D})$, $\mathcal{A}$ returns $\deg(\varphi)$,
3. on input (eval, D, P) with $P \in E(\mathbb{F}_{q^k})$, $\mathcal{A}$ returns $\varphi(P)$.

If furthermore D is of polynomial size in $\log(\deg \varphi)$ and $\log q$, then D is an **efficient representation** of φ (with respect to $\mathcal{A}$).

We now introduce two of the hard problems in isogeny based cryptography, the ONEEND and ISOGENY problems. As seen in [DKL+20, Section 3.2, Problem 1],

Definition 4 (ONEEND)**.** Given a prime p and a supersingular elliptic curve E over $\mathbb{F}_{p^2}$, find a non-trivial cyclic endomorphism of E (in efficient representation).

Defined in [RS06, Cou06], we now introduce the ISOGENY problem with respect to efficient representation,

Definition 5 (ISOGENY)**.** Given E, E' two supersingular elliptic curves defined over $\mathbb{F}_{p^2}$ find an efficient representation of $\varphi\colon E \to E'$.

2.1.2 The ℓ-isogeny graph

The *supersingular ℓ-isogeny graph* is the graph whose vertices are the isomorphism classes of supersingular elliptic curves over $\mathbb{F}_{p^2}$, namely the j-invariant in $\mathbb{F}_{p^2}$, and whose edges are the ℓ-isogenies between them. [Piz80] and [Koh96] have shown that the supersingular ℓ-isogeny graph is connected. Recall that every ℓ-isogeny has a unique dual ℓ-isogeny, and hence, the graph is also undirected (excluding $j = 0, 1728$, which are special cases). From each j-invariant there are exactly $\ell + 1$ outgoing edges, so the graph is also $(\ell + 1)$-regular. Lastly, [Piz90] offers a detailed proof for the supersingular ℓ-isogeny graph satisfying the Ramanujan property.

2.2 Categories

For our purposes, categories can be seen as a generalization of multi-graphs. A (small) category $\mathcal{C}$ is defined by a set of *objects* $\mathrm{obj}(\mathcal{C})$ and a set of *morphisms* $\mathrm{mor}(\mathcal{C})$ with the following properties:

1. Every morphism φ has a *domain* object A and a *codomain* object B; we denote this relation by $\varphi\colon A \to B$.
2. The set of all morphisms with domain A and codomain B is denoted by $\mathrm{Hom}(A, B)$; every such set is called a *homset*.
3. There exists a *composition law*, written $\circ$, that maps a morphism $\varphi\colon A \to B$ and a morphism $\psi\colon B \to C$ to a morphism $\psi \circ \varphi\colon A \to C$, and that is associative: $(\varphi \circ \psi) \circ \chi = \varphi \circ (\psi \circ \chi)$.
4. For every object A there exists a morphism $1_A\colon A \to A$ such that $1_A \circ \varphi = \varphi$ and $\psi \circ 1_A = \psi$ for every $\varphi\colon Z \to A$ and every $\psi\colon A \to B$.

A *functor* $F : \mathcal{C} \to \mathcal{D}$ is a structure-preserving mapping between two categories $\mathcal{C}, \mathcal{D}$. It is defined by a mapping on objects and a mapping on arrows, both denoted by $F()$, with the following properties:

1. For every object $A \in \mathcal{C}$, the object $F(A) \in \mathcal{D}$ is well defined.
2. For every morphism $f : A \to B$ of $\mathcal{C}$, the morphism $F(f) : F(A) \to F(B)$ is well defined.
3. It is compatible with composition: $F(f \circ g) = F(f) \circ F(g)$.
4. It preserves identities: $F(1_A) = 1_{F(A)}$.

Every category has an *identity functor* mapping every object and every morphism to themselves. If functors $F : \mathcal{C} \to \mathcal{D}$ and $G : \mathcal{D} \to \mathcal{C}$ are such that $F \circ G$ and $G \circ F$ are the respective identity functors, we say that they are the (strict) inverse of one another and that the pair forms an *isomorphism of categories*.

From a category $\mathcal{C}$ we construct the *opposite category* $\mathcal{C}^{\mathrm{op}}$ by taking the same set of objects and defining for every morphism $f : A \to B$ in $\mathcal{C}$ the opposite morphism $f^{\mathrm{op}} : B \to A$. A *contravariant* functor from $\mathcal{C}$ to $\mathcal{D}$ is a functor $\mathcal{C}^{\mathrm{op}} \to \mathcal{D}$.

2.2.1 The Deuring functor

Fix a prime p, and let $\mathcal{O}$ be a maximal order of the quaternion algebra over $\mathbb{Q}$ ramified at p and ∞. Two categories will be of special importance to us. The first is $\mathbf{SSing}_p$, the category of supersingular elliptic curves in characteristic p, whose objects are $\overline{\mathbb{F}}_p$-isomorphism classes of supersingular elliptic curves, and whose morphisms are isogenies

between them. The second is $\mathcal{O}$-**Cls**, the category of left $\mathcal{O}$-ideal classes with non-zero $\mathcal{O}$-module homomorphisms.

The *Deuring correspondence* is an isomorphism of categories $D : \mathcal{O}\text{-}\mathbf{Cls}^{\mathrm{op}} \to \mathbf{SSing}_p$.

Let us first understand the correspondence between objects of these categories: D induces a bijection between (isomorphism classes of) supersingular elliptic curves over $\overline{\mathbb{F}}_p$, and left $\mathcal{O}$-ideal classes. Explicitly, let $E_0 \in \mathbf{SSing}_p$ be such that $\mathrm{End}(E_0) \simeq \mathcal{O}$. Let us fix such an isomorphism, and, abusing notation, identify $\mathrm{End}(E_0)$ with $\mathcal{O}$. Naturally, the trivial ideal $\mathcal{O}$ corresponds to this "reference" curve E_0 (i.e., $D(\mathcal{O}) = E_0$). Now let I be any integral left $\mathcal{O}$-ideal. Via the identification $\mathcal{O} = \mathrm{End}(E_0)$, the ideal I is actually a set of endomorphisms, and we can form the isogeny

$$\varphi_I = \mathrm{gcf}\, I,$$

as defined in Section 2.1 (the greatest common factor of I). Write E_I for the codomain of $\varphi_I : E_0 \to E_I$. The ideal I then corresponds to the curve E_I, i.e., $D(I) = E_I$. It is easy to check that the isomorphism class of E_I only depends on the left ideal class $[I]$ of I, and that $E_I \simeq E_0$ if and only if I is principal. In particular, D indeed maps *ideal classes* to *isomorphism classes* of elliptic curves.

It is a bijection, and the inverse can be defined as follows. Given any elliptic curve E, the ring $\mathrm{End}(E)$ is isomorphic to a maximal order in the quaternion algebra $\mathcal{O} \otimes \mathbb{Q}$. Write $\iota_E : \mathrm{End}(E) \to \mathcal{O} \otimes \mathbb{Q}$ for one such embedding. There exists an (essentially unique) left $\mathcal{O}$-ideal I_E which is also a $\iota_E(\mathrm{End}(E))$-ideal on the right. This ideal I_E depends on the choice of ι_E, but its ideal class $[I_E]$ does not, and satisfies $D([I_E]) \simeq E$. In other words, $E \mapsto I_E$ is the inverse of D.

Now, D being a functor, let us also clarify how it maps morphisms. Let I and J be left ideals and $h : J \hookrightarrow I$ a non-zero homomorphism of modules. We wish to construct $D(h)$, an isogeny from $D(I) = E_I$ to $D(J) = E_J$ (note that the domain and codomain are swapped: this explains the op in $\mathcal{O}$-$\mathbf{Cls}^{\mathrm{op}}$). Since $\bigwedge I \subseteq \bigwedge h(J)$, the isogeny $\mathrm{gcf}(h(J))$ factors through $\mathrm{gcf}(I) = \varphi_I$; we can thus write $\mathrm{gcf}(h(J)) = \varphi_h \circ \varphi_I$. It is clear that φ_h does not depend on the choice of representatives of the classes $[I], [J]$, and a simple computation shows that $\varphi_h : E_I \to E_J$. We thus define $D(h) = \varphi_h$.

Deuring proved [Deu41] that every supersingular curve over $\overline{\mathbb{F}}_p$ is isomorphic to some E_I, and that every isogeny can be uniquely written as some φ_h, i.e. that the functor D is an isomorphism of categories.

3 In search of the "one-way functor"

Before giving technical definitions, we explain how the concept of a forensic category arises naturally from attempting to axiomatize the cryptographic content of the Deuring correspondence. This section is heavier in algebraic content than the following ones and can be safely skipped by the reader not versed in isogeny-based cryptography.

It is intuitively clear to anyone familiar with SQIsign that what enables it is a sort of *one-way property* of the Deuring functor: going from ideal classes to elliptic curves is easy, going back from curves to ideal classes is hard. There is no established definition of a "one-way functor", though the concept was previously explored by Montgomery and Patranabis [MP23], who used functors to model cryptographic primitives. Following their lead, we may define a one-way functor as one that can be evaluated in polynomial time on both objects and arrows, but such that antecedents, to both objects and arrows, are generically hard to find. To fix ideas, we informally state the "forward" direction of this definition, though we defer a more formal treatment to Section 4.

Definition 6 (informal)**.** A functor $F : \mathcal{C} \to \mathcal{D}$ is *effective* if there exists an efficient algorithm that on input objects A, B of $\mathcal{C}$ and a morphism $\varphi : A \to B$ outputs $F(A)$, $F(B)$ and $F(\varphi)$.

Then any one-way function $f : X \to Y$ can be trivially *categorified* by viewing X and Y as discrete categories (i.e. categories with no morphisms other than the identities) and f as a functor. We refrain, however, from immediately stating the "backward" part of the definition, as we are going to need a definition of hardness that is considerably more involved than just "hard to invert".

Hardness and fingerprints. The Deuring functor is clearly effective: given ideals I, J and a homomorphism $h : J \hookrightarrow I$, we have polynomial-time algorithms [EHL+18, BDD+24] to compute E_I, E_J and the isogeny $\varphi_h : E_I \to E_J$.

But the functor is not as hard to invert as one may expect: given isogenies $\varphi_I : E_0 \to E_I$ and $\varphi_h : E_I \to E_J$, in many cases[1] we have polynomial-time algorithms to compute I, J and $h : J \hookrightarrow I$. What appears to be genuinely hard, instead, is the problem of computing I given only E_I, which is equivalent to the ISOGENY problem.

This dichotomy points to the actual cryptographic use of the Deuring correspondence: ideal classes I, J can serve as secret data associated to the public curves E_I, E_J. The owner of the secrets can do computations on I and J and evaluate the Deuring functor to produce isogenies $\varphi_h : E_I \to E_J$, whereas someone lacking the knowledge of I or J cannot. This already suggests an interactive identification scheme along the lines of [GPS17, GPS20]. But there is more to it.

Isogenies $\varphi_h : E_I \to E_J$ between public curves come equipped with associated publicly computable data. Examples are: the degree of φ_h, whether $\varphi_h \succ \psi$ for some given ψ, how φ_h acts on torsion points of E_I, E_J, ... Every single one of these has been used in some cryptographic protocol [BBC+25, DKL+20, BFLS26]. Through the Deuring correspondence, these data have analogues in $\mathcal{O}$-**Cls**: degrees correspond to norms, isogeny factorizations correspond to ideal filtrations, actions on torsion points correspond to families of sublattices [BFLS26]. The power of the Deuring correspondence resides in that these data can be computed and manipulated on the $\mathcal{O}$-**Cls** side, before evaluating the functor and letting the other party check them on the $\mathbf{SSing}_p$ side. Crucially, these data are not malleable in $\mathbf{SSing}_p$: from an isogeny $\varphi_h : E_I \to E_J$, we do not know how to produce a new isogeny $\varphi' : E_I \to E_J$ with different associated data in any non-trivial way. We abstract this crucial property of the Deuring functor by defining *fingerprint maps*: maps from morphisms to a finite set of observable data.

Definition 7 (informal)**.** Let $F : \mathcal{C} \to \mathcal{D}$ be an effective functor and let $\mathcal{M}$ be a finite set. A *fingerprint map* is a pair of efficiently-computable partial mappings

$$\mathrm{fp}_{\mathcal{C}} : \mathrm{mor}(\mathcal{C}) \to \mathcal{M}, \qquad \mathrm{fp}_{\mathcal{D}} : \mathrm{mor}(\mathcal{D}) \to \mathcal{M},$$

such that $\mathrm{fp}_{\mathcal{D}}(F(\varphi)) = \mathrm{fp}_{\mathcal{C}}(\varphi)$. Additionally, a fingerprint map comes with an efficient algorithm that on input objects A, B of $\mathcal{C}$ and a valid fingerprint $m \in \mathcal{M}$ outputs a morphism $\varphi : A \to B$ such that $\mathrm{fp}_{\mathcal{C}}(\varphi) = m$.

Example 1. Intuitively, fingerprints capture partial information on morphisms that can be easily verified both in $\mathcal{C}$ and $\mathcal{D}$, but only constructed in $\mathcal{C}$. Smooth isogeny kernels provide one such example in $\mathbf{SSing}_p$: fix a smooth integer N (e.g., $N = 2^n$) such that $E_I[N] \subset E_I(\mathbb{F}_{p^2})$ and fix a basis (P, Q) of $E_I[N]$. Given any cyclic isogeny $\varphi : E_I \to E_J$ such that $N | \deg(\varphi)$, we can easily find the unique projective pair $(a : b) \in \mathbb{P}(\mathbb{Z}/N\mathbb{Z})$ such that $\varphi([a]P + [b]Q) = 0$. For such isogenies we define $\mathrm{fp}(\varphi) = (a : b)$.

[1] The case where the isogenies φ_I and φ_h have smooth degree is among these. See [CII+23] for a more in-depth overview of the "isogeny-to-ideal problem".

Defining the corresponding fingerprint in $\mathcal{O}$-**Cls** is technical, and we will not need it in the rest of this work, so we skip the details. Suffice to say that not only can $\mathrm{fp}(h)$ be computed for any morphism $h : J \to I$ of left $\mathcal{O}$-ideal classes, but also given any $(a : b)$, I and J, we can efficiently compute $h : J \to I$ such that $\mathrm{fp}(h) = (a : b)$. The fact that N is smooth is crucial for the efficiency claim.

Going back to $\mathbf{SSing}_p$, note that, given E_I and $(a : b)$, we can easily compute an isogeny with fingerprint $(a : b)$: it suffices to take the isogeny whose kernel is exactly $\langle [a]P + [b]Q \rangle$, which can be easily computed using Vélu's formulas. What is really difficult is to compute $\varphi : E_I \to E_J$ such that $\mathrm{fp}(\varphi) = (a : b)$ when E_I and E_J are given as inputs. This is the key property at the heart of SQIsign: in the interactive identification scheme, the prover sets up a pair of curves E_I and E_J, the "public key" and the "commitment" curve, and is challenged to produce an isogeny $\varphi : E_I \to E_J$ such that $\mathrm{fp}(\varphi) = (a : b)$ for some random fingerprint chosen by the verifier. Only the legitimate prover, thanks to its knowledge of the corresponding ideal classes I and J, can produce such an isogeny.

We give formal definitions for this *kernel fingerprint* in Section 7.2.

Example 2. A related fingerprint is encountered in SQInstructor [BFLS26]. Fixing again bases (P_I, Q_I) of $E_I[N]$ and (P_J, Q_J) of $E_J[N]$, we let $\mathrm{fp}(\varphi : E_I \to E_J)$ be the matrix $\left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$ such that $\varphi(P_I) = [a]P_J + [c]Q_J$ and $\varphi(Q_I) = [b]P_J + [d]Q_J$. For reasons that will soon become clear, these fingerprints must be taken up to equivalence in $\mathrm{GL}_2(\mathbb{Z}/N\mathbb{Z})/\Gamma$, where Γ is a subgroup of matrices containing the scalar matrices. Hence, the fingerprint of an isogeny φ can be interpreted as its action on Γ-level structures. We give formal definitions for this *level-structure fingerprint* in Section 7.4.

For maximum flexibility, we put very few restrictions on fingerprint maps, and thus our definition is not very "categorical". However, in some cases fingerprint maps may be functors to a category with one object. We can now state an assumption that models the hardness of the Deuring functor with respect to fingerprint maps.

Definition 8 (informal)**.** An effective functor $F : \mathcal{C} \to \mathcal{D}$ and a fingerprint map $(\mathrm{fp}_{\mathcal{C}}, \mathrm{fp}_{\mathcal{D}})$ form a *forensic pair* if, given a random object A of $\mathcal{D}$ it is hard on average to find two morphisms $\varphi, \psi : A \to B$ such that $\mathrm{fp}_{\mathcal{D}}(\varphi), \mathrm{fp}_{\mathcal{D}}(\psi)$ are defined and $\mathrm{fp}_{\mathcal{D}}(\varphi) \neq \mathrm{fp}_{\mathcal{D}}(\psi)$.

We will show in Section 4.3 that this assumption implies the hardness of finding morphisms between objects of $\mathcal{D}$, i.e., the one-way property on morphisms.

Example 3. Going back to Example 1, we illustrate how the hardness of finding two morphisms of distinct fingerprints is the abstraction of the ONEEND problem, described in Definition 4. By the definition of the kernel fingerprint, breaking security of the forensic pair is equivalent to finding two cyclic isogenies, $\varphi, \psi : E_I \to E_J$ such that $\mathrm{fp}(\varphi) \neq \mathrm{fp}(\psi)$. This implies that $\ker(\varphi) \cap \ker(\psi) \cap E_I[N] = \{0\}$, and hence, $\hat{\psi} \circ \varphi : E_I \to E_I$ is a non-scalar endomorphism of E_I. Therefore, finding two isogenies of distinct fingerprints is equivalent to finding a solution to the ONEEND problem.

Similarly for the level-structure fingerprint of Example 2, suppose that φ, ψ have distinct fingerprints. Recalling that fingerprints are matrices describing the action of the isogenies on the N-torsion, we have $\mathrm{fp}(\hat{\psi} \circ \varphi) = \deg(\psi)\mathrm{fp}(\psi)^{-1}\mathrm{fp}(\varphi)$. This alone is not sufficient to show that $\hat{\psi} \circ \varphi$ is a non-scalar endomorphism, however if $\mathrm{fp}(\psi)$ and $\mathrm{fp}(\varphi)$ are distinct in PGL_2, then $\mathrm{fp}(\hat{\psi} \circ \varphi)$ is a non-scalar matrix and hence $\hat{\psi} \circ \varphi$ is a solution to ONEEND.

For formal proofs of these reductions, we point the reader to Theorem 8 and Theorem 15.

Forensic representable functors. Now we could go on and fill the gaps in our informal definitions by stating what operations can be efficiently computed on $\mathcal{C}$ and $\mathcal{D}$. Understanding $\mathcal{C}$ as the "easy category" acting on the "hard category" $\mathcal{D}$, this would have a very

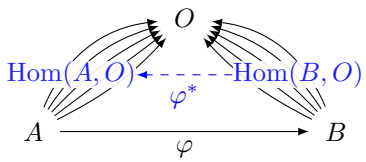


Figure 1: The functor $\mathrm{Hom}(-, O)$. Source category in black, image of the functor in blue.

similar flavor to Cryptographic Group Actions [ADMP20], albeit with a different set of axioms. However doing so would be slightly verbose, as we would need to duplicate many definitions for $\mathcal{C}$ and $\mathcal{D}$. It turns out that the properties of the Deuring functor are entirely contained within the image category $\mathcal{D}$, a consequence of it being a *representable functor*.

Let O be an object of $\mathcal{D}$. The *hom-functor* $\mathrm{Hom}(-, O) : \mathcal{D}^{\mathrm{op}} \to \mathbf{Set}$ maps an object A of $\mathcal{D}$ to the homset $\mathrm{Hom}(A, O)$ and a morphism $\varphi : A \to B$ to the pullback morphism $\varphi^* : g \mapsto g \circ \varphi$, as illustrated in Figure 1. A functor is *representable* if it is "of the form"[2] $\mathrm{Hom}(-, O)$ for some object O, called the *representing object*. The inverse Deuring functor D^{-1} is a representable functor represented by E_0, via the (natural) isomorphism

$$\begin{aligned} \mathrm{Hom}(E_I, \mathcal{O}) &\longrightarrow I \\ \psi &\longmapsto \psi \circ \varphi_I, \end{aligned}$$

where for each ideal class $[I]$ we choose an arbitrary representative I.

Computationally, the datum of the isogeny $\varphi_I : E_0 \to E_I$ realizes $\mathrm{Hom}(E_I, \mathcal{O})$ as an ideal $I \subset \mathcal{O}$. If we assume the existence of an efficient algorithm[3] that on input φ_I computes I, we may then use φ_I in place of I as input to any algorithm that requires I.

This idea applies generally to any representable functor, and we can play a game of restating the informal definitions above internally to the "hard" category $\mathcal{D}$. Definition 6 and Definition 7 combine to the existence of an efficient algorithm that on input a pair of morphisms $O \to A$ and $O \to B$ and a valid fingerprint m, outputs a morphism $A \to B$ matching the fingerprint. This will be formally stated in the "Triangularizable" axiom of Definition 13. The hardness assumption in Definition 8 is essentially unchanged, simply dropping the mentions of the functor F and the "easy" category $\mathcal{C}$.

Although this "one-category" formalism loses a bit of expressiveness compared to the one-way-functor formalism, we believe it is sufficiently general to express all cryptographic constructions of interest, and has the clear benefit of being more succinct and requiring less background in category theory. This is why we chose to develop it in the next sections.

4 Forensic categories

We introduce the main object of our study. We enrich the classical notion of category with a set of computational axioms defining what can be efficiently computed and what problems are assumed to be cryptographically hard. The determining feature of our categories is the *fingerprint*, a finite set of labels attributed to the morphisms of the category. Looking ahead, fingerprints will serve as challenge material in the interactive protocols of Sections 5 and 6.

[2] Technically, *naturally isomorphic*. For brevity we skip the definition of natural isomorphism.

[3] We already mentioned such an algorithm is not known for arbitrary isogenies, however we do not need φ_I to be arbitrary, and it is not a stretch to think that such an algorithm may eventually be found.

4.1 Effective categories

We introduce the concept of an effective category, a category equipped with efficient algorithms to compute basic functionalities, such as composing morphisms or generating new ones.

Definition 9. A **pre-effective category family** is a sequence $(\mathcal{C}_n)_{n\in\mathbb{N}}$ where each $\mathcal{C}_n$ is a category, together with algorithms such that the following holds. For each n, we write $\lambda(n) = \log(n)$. All objects and morphisms in $\mathcal{C}_n$ have encodings as binary strings, and there exist (probabilistic) polynomial time algorithms (in the length of the input) for the following tasks:

1. An algorithm that on input n and a binary string A, decides whether A is a valid encoding of an object in $\mathcal{C}_n$. The binary encoding of every object is unique and of length polynomial in $\lambda(n)$.

2. An algorithm that on input n and a binary string φ, decides whether φ is a valid encoding of a morphism in $\mathcal{C}_n$.

3. An algorithm that on input n and the encoding of a morphism $\varphi\colon D \to C$ in $\mathcal{C}_n$, outputs the unique encodings of domain D and codomain C.

4. An algorithm that on input n and the encodings of two morphisms $\varphi\colon A \to B$ and $\psi\colon B \to C$ in $\mathcal{C}_n$, returns an encoding of the composition $\psi \circ \varphi\colon A \to C$.

5. An algorithm that on input n and the encoding of an object A in $\mathcal{C}_n$, outputs an encoding of the identity morphism $1_A\colon A \to A$.

Remark 2. In the rest of this article, we abuse terminology and notation and say *categories* instead of *families of categories*. The index n is always an input of the algorithms, which will remain implicit from now on. Asymptotic statements such as "there exists a polynomial time algorithm" are with respect to the input length, hence in particular, depend on the implicit quantity $\lambda = \lambda(n)$. This implicit λ should be understood as a "security parameter", as is common in the cryptographic literature.

Our definitions are pretty liberal: besides forcing the object set to be finite (of size at most quasi-polynomial in n), they impose very few constraints on the categories. Homsets may be infinite, morphisms may have multiple representations and their equality does not need to be decidable. In a pre-effective category, even producing the encoding of a single object does not need to be efficient. Inspired by the examples in isogeny-based cryptography, we opt for a minimal definition of how objects and morphisms may be discovered. We only assume that a single object O of $\mathcal{C}_n$, named the *origin*, is known at the start. Every other object and morphism is discovered by a process akin to randomly *walking* in a graph.

Definition 10. A **walk** in a pre-effective category $(\mathcal{C}_n)_{n\in\mathbb{N}}$ is defined as a pair (Walk, μ), where

- Walk is a deterministic polynomial time algorithm which takes as input an index n, an object A in $\mathcal{C}_n$, and a sequence of (random) coins $r \in R^*$ (for some finite set R, e.g., $R = \{0, 1\}$) and produces a pair (B, ψ), where ψ is a morphism $\psi : A \to B$, and

- $\mu = (\mu_n)_{n\in\mathbb{N}}$ is a sequence of distributions, where μ_n is fully supported on the objects of $\mathcal{C}_n$, with the following property. There is a bound $X = n^{O(1)}$ such that for any $x > X$, if r is a sequence of x uniformly random coins, the statistical distance between μ_n and the distribution of the codomain B in the output of $\mathsf{Walk}(n, A; r)$ is a negligible function of n.

As previously, the index n will stay implicit, and we simply write $\mathsf{Walk}(A; r)$ and μ in place of $\mathsf{Walk}(n, A; r)$ and μ_n. Furthermore, the input r will often be treated as implicit

random coins, and we simply write $\mathsf{Walk}(A) \coloneqq \mathsf{Walk}(n, A; r)$ for the randomized algorithm using X uniformly random coins (or $\mathsf{Walk}(A, x)$ to emphasize a number x of random coins used).

Definition 11. An **effective category** is a pre-effective category together with a walk Walk and an origin object O whose representation can be computed in polynomial time.

4.2 Defining the fingerprint map

We now define the notion of a *fingerprint map* on an effective category.

Definition 12. A **fingerprint map** fp on a category $\mathcal{C}$ is a collection of maps

$$\mathrm{fp}\colon \mathrm{Hom}(-,-) \to \mathcal{M} \cup \{\perp\},$$

such that $\mathcal{M}$ is finite and contains at least two elements, i.e. $|\mathcal{M}| \geq 2$. The special symbol $\perp$ is used to denote those cases when fp is undefined. An output m of fp is a **fingerprint**, which we call **valid** if it is defined, i.e. $m \in \mathcal{M}$.

We define a *forensic category* as an effective category with a fingerprint map that satisfies a list of properties, as follows.

Definition 13. A **forensic category** is an effective category where the fingerprint map satisfies the following properties:

Effective: Given φ, computing $\mathrm{fp}(\varphi)$ is efficient.

Hard (on average): The language

$$\mathcal{L} = \left\{ \big(A, (\varphi, \psi)\big) \;\middle|\; \exists B \text{ such that } \begin{array}{l} \varphi, \psi\colon A \to B, \\ \mathrm{fp}(\varphi) \neq \mathrm{fp}(\psi), \\ \mathrm{fp}(\varphi), \mathrm{fp}(\psi) \neq \perp \end{array} \right\} \tag{2}$$

is hard on average for A following distribution μ. More explicitly, for any PPT adversary $\mathcal{A}$, given as input an object A sampled from distribution μ, the probability $\mathsf{Adv}^{\mathsf{hard}}(\mathcal{A})$ that $\mathcal{A}(A)$ returns a pair (φ, ψ) such that $(A, (\varphi, \psi)) \in \mathcal{L}$ is negligible.

Triangularizable: There exists a PPT algorithm, $\mathsf{Triangle}$, that on inputs two morphisms $\varphi\colon O \to A, \psi\colon O \to B$ and a valid fingerprint $m \in \mathcal{M}$, returns $\chi\colon A \to B$ such that $\mathrm{fp}(\chi) = m$. Furthermore, the distribution of χ only depends on A, B and m.

(Computationally) Walk-Indistinguishable: There exists a PPT algorithm $\mathsf{SimWalk}$ which takes an object A as input and returns a morphism $\chi : A \to B$ with a valid fingerprint, indistinguishable from the output of $\mathsf{Triangle}$. Formally, for any PPT adversary $\mathcal{A}$ with access to an oracle $\mathcal{O}$, and for a $\varphi : O \to A$ sampled as the output of $\mathsf{Walk}(O)$, the following advantage is negligible

$$\mathsf{Adv}^{\mathsf{Walk\text{-}Ind}}(\mathcal{A}) = \left| \Pr\left[\mathcal{A}\left(\varphi; \mathcal{O}^{\varphi}_{\mathsf{SimWalk}}\right) = 1\right] - \Pr\left[\mathcal{A}\left(\varphi; \mathcal{O}^{\varphi}_{\mathsf{Triangle}}\right) = 1\right] \right|,$$

where $\mathcal{O}^{\varphi}_{\mathsf{SimWalk}}$ and $\mathcal{O}^{\varphi}_{\mathsf{Triangle}}$ are defined in Algorithm 1.

Remark 3. The language $\mathcal{L}$ in Definition 13 is a *total search problem*: every instance A has a solution (φ, ψ) such that $(A, (\varphi, \psi)) \in \mathcal{L}$. This follows from the Triangularizable property and the assumptions that $|\mathcal{M}| \geq 2$ and that the support of μ is the set of all objects.

Remark 4. We could have defined several other variants of indistinguishability. Stronger variants may let the adversary choose φ, or even let it be computationally unbounded; weaker variants may omit φ in the inputs to the adversary. We chose the present one as an acceptable middle ground: it is satisfied by the instantiations of Section 7 and is strong enough to build all the primitives of Sections 5 and 6.

Algorithm 1 Oracles of the walk-distinguishing experiment

oracle $\mathcal{O}^{\varphi:O\to A}_{\mathsf{SimWalk}}$
  $\chi \leftarrow \mathsf{SimWalk}(A)$
  return χ

oracle $\mathcal{O}^{\varphi:O\to A}_{\mathsf{Triangle}}$
  $(B,\psi) \leftarrow \mathsf{Walk}(O)$
  $m \leftarrow_{\$} \mathcal{M}$
  return $\chi \leftarrow \mathsf{Triangle}(\varphi,\psi,m)$

4.3 Finding a morphism between two objects is hard

In this section, we prove that finding morphisms between objects in a forensic category is hard, even in the average case.

Problem 1 (The Morphism Problem)**.** *Let $\mathcal{C}$ be a forensic category. Given two objects A and B, find a morphism $\varphi\colon A \to B$ (or decide that none exists).*

We define the average case of Problem 1 as the case where the input objects A and B are independent random objects with distribution μ. For any PPT algorithm $\mathcal{A}$, we write $\mathsf{Adv}^{\mathsf{morph}}(\mathcal{A})$ for its probability of success in solving the morphism problem.

Proposition 1. *In any forensic category $\mathcal{C}$, the Morphism Problem is hard on average. That is, for any PPT adversary $\mathcal{A}$ that solves the average-case Morphism Problem with non-negligible advantage $\mathsf{Adv}^{\mathsf{morph}}(\mathcal{A})$, there exists a PPT algorithm $\mathcal{B}$ that breaks the hardness (on average) property of $\mathcal{C}$ with non-negligible advantage, i.e.*

$$\mathsf{Adv}^{\mathsf{morph}}(\mathcal{A}) \leq \mathsf{Adv}^{\mathsf{hard}}(\mathcal{B}).$$

Proof. Consider the algorithm $\mathcal{B}$, as described in Algorithm 2. Since Walk, $\mathsf{Triangle}$, and sampling are polynomial time algorithms, and the adversary $\mathcal{A}$ is efficient, Algorithm 2 is also polynomial time. Then, the output of $\mathcal{B}$ is two morphisms χ, χ' from A to the same object with distinct valid fingerprints. The success probability of $\mathcal{B}$ is the same as that of $\mathcal{A}$. □

Algorithm 2 The algorithm $\mathcal{B}$ in the proof of Proposition 1

Input: An object A
Input: An adversary $\mathcal{A}$ that solves the (average-case) Morphism Problem
Output: $\chi\colon A \to B, \chi'\colon A \to B$ for some object B such that $\mathrm{fp}(\chi) \neq \mathrm{fp}(\chi')$ and $\mathrm{fp}(\chi), \mathrm{fp}(\chi') \in \mathcal{M}$

Run $(C,\gamma) \leftarrow \mathsf{Walk}(O)$
Compute $(\psi\colon C \to A) \leftarrow \mathcal{A}(C,A)$
Run $(B,\varphi) \leftarrow \mathsf{Walk}(O)$
Sample $m \leftarrow \mathcal{M}$
Run $\chi \leftarrow \mathsf{Triangle}(\psi\circ\gamma,\varphi,m)$ ▷ $\chi\colon A \to B$, $\mathrm{fp}(\chi) = m$
Sample $m' \leftarrow \mathcal{M}\setminus\{m\}$ ▷ $\mathcal{M}\setminus\{m\} \neq \emptyset$ because $|\mathcal{M}| \geq 2$
Run $\chi' \leftarrow \mathsf{Triangle}(\psi\circ\gamma,\varphi,m')$ ▷ $\chi'\colon A \to B$, $\mathrm{fp}(\chi') = m'$
return χ, χ'

4.4 Fingerprint stability

We now explore three additional axioms that, while not strictly necessary, enhance a forensic category making it more expressive. They all have a similar flavor, expressing the fact that the space of valid fingerprints $\mathcal{M}$ is stable under some algebraic transformation.

We start with *undisturbing walks*: walks naturally act on homsets by *pullback*; when this action induces a permutation on valid fingerprints, we say the category has undisturbing walks. This condition tends to be easy to satisfy, and is indeed satisfied in the instantiations presented in Section 7. We use it in Theorem 1 to prove a worst-case to average-case reduction.

Definition 14. Let fp be a fingerprint map. An algorithm UWalk is an *undisturbing walk* for fp if it is a walk as described in Definition 10 such that the following holds: for any objects A, B, whenever $\mathsf{UWalk}(A)$ returns a morphism $\chi\colon A \to C$, letting $F_\chi : \psi \mapsto \psi \circ \chi$ be the pullback map, there exists a bijection $\alpha\colon \mathcal{M} \cup \{\perp\} \to \mathcal{M} \cup \{\perp\}$ fixing $\perp$ such that the diagram below commutes.

$$\begin{array}{ccc} \mathrm{Hom}(C,B) & \xrightarrow{F_\chi} & \mathrm{Hom}(A,B) \\ {\scriptstyle \mathrm{fp}}\downarrow & \circlearrowright & \downarrow{\scriptstyle \mathrm{fp}} \\ \mathcal{M}\cup\{\perp\} & \xrightarrow[\alpha]{} & \mathcal{M}\cup\{\perp\} \end{array} \tag{3}$$

The next property is *camouflage*: we postulate the existence of a free group action on valid fingerprints.

Definition 15. Let G be a group[4] and let $\mathcal{M}$ be the set of valid fingerprints. We say that the fingerprint map is *camouflaged* by G if there is an efficiently computable free group action of G on $\mathcal{M}$.

This property, too, tends to be easily satisfied. For example, if $\mathcal{M}$ is in bijection with binary strings $\{0,1\}^n$, we may take XOR as group operation on $G = \mathcal{M}$ acting on itself. In Section 6 we show how to construct several primitives from camouflage.

Finally, we introduce a stronger but very natural property that implies both of the above.

Definition 16. We say that a fingerprint map fp is *homomorphic* if the following holds. $\mathcal{M}$ is a group whose group operation can be efficiently computed and $\mathrm{fp}(\psi \circ \varphi) = \mathrm{fp}(\psi)\mathrm{fp}(\varphi)$, for any ψ, φ such that their fingerprints are valid (i.e. $\mathrm{fp}(\psi), \mathrm{fp}(\varphi) \neq \perp$). We also require $\mathrm{fp}(\psi \circ \varphi) = \perp$ whenever $\mathrm{fp}(\psi)$ or $\mathrm{fp}(\varphi)$ are invalid.

Lemma 1. *For any homomorphic fingerprint map* fp *and any walk* Walk *whose outputs* ψ *satisfy* $\mathrm{fp}(\psi) \neq \perp$*, the algorithm* Walk *is an undisturbing walk for* fp*.*

Proof. Let A be any object, and $\chi\colon A \to C$ an output of $\mathsf{Walk}(A)$. Let $\alpha\colon \mathcal{M} \cup \{\perp\} \to \mathcal{M} \cup \{\perp\}$ be defined as $\alpha(\perp) = \perp$ and $\alpha(m) = m \cdot \mathrm{fp}(\chi)$. Using the notation of Equation (3), we obtain that for any $\psi : C \to B$ with a valid fingerprint,

$$\alpha(\mathrm{fp}(\psi)) = \mathrm{fp}(\psi) \cdot \mathrm{fp}(\chi) = \mathrm{fp}(\psi \circ \chi) = \mathrm{fp}(F_\chi(\psi)).$$

If $\mathrm{fp}(\psi) = \perp$, we have $\mathrm{fp}(\psi \circ \chi) = \perp = \alpha(\mathrm{fp}(\psi))$. In other words, the diagram of Equation (3) is commutative. □

Lemma 2. *Let* fp *be a homomorphic fingerprint map with valid fingerprint set* $\mathcal{M}$*, then* $\mathcal{M}$ *camouflages* fp*.*

Proof. Obviously $\mathcal{M}$ is a group and thus acts on itself freely and transitively. □

[4] A magma with divisibility and an identity, i.e. a *loop*, would be sufficient for our needs. However loops are rather exotic objects, so we stick to the more restrictive definition.

4.5 Worst-case to average-case reduction of $\mathcal{L}$

Consider the language

$$\mathcal{L} = \left\{ (A, (\varphi, \psi)) \;\middle|\; \exists B \text{ such that } \begin{array}{l} \varphi, \psi \colon A \to B, \\ \mathrm{fp}(\varphi) \neq \mathrm{fp}(\psi), \\ \mathrm{fp}(\varphi), \mathrm{fp}(\psi) \neq \perp \end{array} \right\} \tag{4}$$

For forensic categories, Definition 13 requires this language to be hard on average. Here we prove that if the category has an undisturbing walk, this average hardness follows from the worst-case hardness of the language: it benefits from a worst-case to average-case reduction.

Theorem 1 (Worst-case to average-case reduction of $\mathcal{L}$)**.** *Consider any fingerprint map* fp *for which there exists an undisturbing walk* UWalk*. If the associated language* $\mathcal{L}$ *is hard, then it is hard on average.*

Proof. Let μ be the limit distribution of UWalk. Suppose that there exists a PPT algorithm $\mathcal{A}$ that, given as input a statement A' distributed according to μ, outputs $(A', (\psi, \psi')) \in \mathcal{L}$ with non-negligible probability. Consider the following algorithm, $\mathcal{B}$:

1. Run $(C, \chi) \leftarrow \mathsf{UWalk}(A)$.
2. Since C follows distribution μ, the algorithm $\mathcal{A}(C)$ outputs $\psi, \psi' \colon C \to B$ such that $\mathrm{fp}(\psi) \neq \mathrm{fp}(\psi')$ and $\mathrm{fp}(\psi), \mathrm{fp}(\psi') \neq \perp$ with non-negligible probability.
3. Return $(A, (\psi \circ \chi, \psi' \circ \chi))$.

Note that algorithm $\mathcal{B}$ is PPT since UWalk and $\mathcal{A}$ are PPT algorithms. In the following analysis assume we are in the (non-negligible) event that $\mathrm{fp}(\psi) \neq \mathrm{fp}(\psi')$ and $\mathrm{fp}(\psi), \mathrm{fp}(\psi') \neq \perp$, and we prove that the output satisfies $(A, (\psi \circ \chi, \psi' \circ \chi)) \in \mathcal{L}$, proving the result.

Now, let us use the property of the undisturbing walk UWalk. With the same notation as in Equation (3), we have that

$$\big(A, (\psi \circ \chi, \psi' \circ \chi)\big) = \big(A, (F_\chi(\psi), F_\chi(\psi'))\big).$$

Moreover, by the commutativity of the square in Equation (3), we have

$$\mathrm{fp}(F_\chi(\psi)) = \alpha(\mathrm{fp}(\psi)) \quad \text{and} \quad \mathrm{fp}(F_\chi(\psi')) = \alpha(\mathrm{fp}(\psi')),$$

where α is the bijection from $\mathcal{M} \cup \{\perp\}$ to $\mathcal{M} \cup \{\perp\}$ fixing $\perp$. Hence, we have that $\alpha(\mathrm{fp}(\psi)) \neq \alpha(\mathrm{fp}(\psi'))$ and $\alpha(\mathrm{fp}(\psi)), \alpha(\mathrm{fp}(\psi')) \in \mathcal{M}$, because $\mathrm{fp}(\psi) \neq \mathrm{fp}(\psi')$, $\mathrm{fp}(\psi), \mathrm{fp}(\psi') \in \mathcal{M}$ and α is a bijection. In particular, we have

$$\mathrm{fp}(\psi \circ \chi), \mathrm{fp}(\psi' \circ \chi) \neq \perp \quad \text{and} \quad \mathrm{fp}(\psi \circ \chi) \neq \mathrm{fp}(\psi' \circ \chi).$$

In other words, $(A, (\psi \circ \chi, \psi' \circ \chi)) \in \mathcal{L}$, finishing the proof. □

5 Signing in forensic categories

Let $\mathcal{C}$ be a forensic category. We show how to construct a *digital signature* within our framework, and denote it by $\mathrm{FCSignature}_{\mathcal{C}}$. This digital signature is designed to mirror the underlying structure of SQIsign. Note that different instantiations of a forensic category $\mathcal{C}$ will lead to distinct variants of the SQIsign protocol [DKL+20, BFLS26].

We start by recalling some basic definitions for Sigma-protocols. A Sigma-protocol for a relation $\mathcal{R}$ is a public-coin three-move interactive protocol, where a prover P tries to convince a verifier V that it knows the witness w to some statement $(x, w) \in \mathcal{R}$. A sequence of messages (m, c, r) between P and V that ends in V accepting the proof is called an *accepting transcript*. Important properties for a Sigma-protocol are:

Completeness: For any $(x, w) \in \mathcal{R}$, if $P(x, w)$ and $V(x)$ follow the protocol, V always accepts.

Special soundness: From any x and any pair of accepting transcripts on input x, $(m, c, r), (m, c', r')$ such that $c \neq c'$, one can efficiently compute w such that $(x, w) \in \mathcal{R}$.

Honest-verifier zero-knowledge (HVZK). A Sigma-protocol for a relation $\mathcal{R}$ is honest-verifier zero-knowledge (HVZK) if there exists a PPT simulator $\mathcal{S}$ which, given only a statement x, outputs a transcript (m, c, r) such that, for an honestly generated (uniformly random) pair $(x, w) \in R$, the advantage of any PPT adversary $\mathcal{A}$ given x (but not w) in distinguishing the output of $S(x)$ from the (accepting) transcript of an honest execution on (x, w) is negligible.

Strong honest-verifier zero-knowledge (strong HVZK). A Sigma-protocol is *strong* HVZK if there exists a simulator as above for which indistinguishability between real and simulated transcripts continues to hold even when the PPT adversary $\mathcal{A}$ is given the witness w in addition to x. The simulator still receives only x, and not w.

We clarify that the term "strong" above *does not* refer to strong unforgeability of a signature scheme obtained from the Sigma-protocol.

We now present a Sigma-protocol $\mathrm{FCSigma}_{\mathcal{C}}$ for the hard language $\mathcal{L}$ of Equation (2).

InstanceGen. Run $(A, \varphi) \leftarrow \mathsf{Walk}(O)$. A will serve as public key, while φ will be the secret key used to produce tuples $(A, (\psi, \chi)) \in \mathcal{L}$.

Commitment. The prover generates $(B, \psi) \leftarrow \mathsf{Walk}(O)$, and sends B as the commitment object.

Challenge. The verifier samples a valid fingerprint $m \in \mathcal{M}$ uniformly at random and sends it to the prover.

Response. The prover rejects if $m \notin \mathcal{M}$. The prover runs $\psi_{\mathrm{rsp}} \leftarrow \mathsf{Triangle}(\varphi, \psi, m)$ to obtain a morphism ψ_{rsp} such that $\mathrm{fp}(\psi_{\mathrm{rsp}}) = m$. It sends ψ_{rsp} to the verifier.

Verification. The verifier accepts if ψ_{rsp} has A for domain, B for codomain, and if $\mathrm{fp}(\psi_{\mathrm{rsp}}) = m \in \mathcal{M}$.

It is immediate to see that the protocol is complete: given the secret key φ, a prover can successfully respond to any challenge $m \in \mathcal{M}$. To begin the analysis, observe that recovering φ is an instance of the Morphism Problem, which is hard thanks to Proposition 1. We now prove that the protocol is special-sound and zero-knowledge.

Proposition 2. *Let $\mathcal{C}$ be a forensic category. Then,* $\mathrm{FCSigma}_{\mathcal{C}}$ *is a special-sound proof of knowledge for the Hard language of Definition 13.*

Proof. We say that $\mathrm{FCSigma}_{\mathcal{C}}$ is special-sound if, given two accepting transcripts $(B, m, \psi_{\mathrm{rsp}})$ and $(B, m', \psi'_{\mathrm{rsp}})$, with the same commitment B and distinct challenges $m \neq m'$, one can find a witness (χ, χ') for A in $\mathcal{L}$ (Equation (2)). By construction, $(\psi_{\mathrm{rsp}}, \psi'_{\mathrm{rsp}})$ is indeed a witness for A in $\mathcal{L}$, and we are done. □

Observe that in the proof above we only recover a witness instance for the instance A in $\mathcal{L}$. Because $\mathcal{L}$ is a hard language, this is sufficient to prove security, but it does not show how to recover the secret key φ, and indeed in instantiations of SQIsign several such witnesses may be needed for forging signatures.

Also note that the soundness error of the interactive proof is the inverse of the size of the challenge space, i.e. $1/|\mathcal{M}|$, as explained in [Dam10, Theorem 1]. Hence, the forensic categories with the best security / bandwidth ratio are those where $|\mathcal{M}|$ is of exponential size.

Proposition 3. *Let $\mathcal{C}$ be a forensic category. Then,* $\mathrm{FCSigma}_{\mathcal{C}}$ *is strong honest-verifier zero-knowledge.*

Proof. Let $\varphi : O \to A$ be a secret key. We define the simulator as follows. On input the public key A:

1. Run $(B, \chi) \leftarrow \mathsf{SimWalk}(A)$;
2. Compute $m \leftarrow \mathrm{fp}(\chi)$;
3. Return (B, m, χ).

Comparing the simulator with the oracle $\mathcal{O}^{\varphi}_{\mathsf{SimWalk}}$ in the Walk-distinguishing game (see Algorithm 1), we see their outputs are identically distributed for each A, independently of φ. Similarly, comparing accepting transcripts to $\mathcal{O}^{\varphi}_{\mathsf{Triangle}}$, we see again that outputs are identically distributed. We deduce that the advantage of an adversary for the HVZK game is bounded by the advantage of an adversary for the walk-distinguishing game of the forensic category. □

It is a well-known result that a complete, special sound, and strong[5] HVZK Sigma-protocol, is also a canonical identification scheme secure against impersonation under passive attacks (see, e.g., [CSD+23, Theorem 9.1.4]). Then, applying the generalized Fiat–Shamir transform of [AABN02] we derive $\mathrm{FCSignature}_{\mathcal{C}}$, a digital signature scheme unforgeable under chosen message attacks in the random oracle model.

Remark 5. This construction does not quite capture the HD variants of SQIsign [DLRW24, BDD+24, NOC+24, DF24], which deviate from the classical notion of honest-verifier zero-knowledge: they require a hint-assisted simulator (see [ABD+25]). The axioms could be adapted to cover these variants, but we do not explore this direction in this paper.

6 Towards more cryptographic primitives in forensic categories

In this section, we explore further cryptographic primitives that can be constructed in the framework of forensic categories. The first, a *preimage-resistant hash*, follows immediately from the axioms of a forensic category, while the subsequent primitives, namely a *ring signature* and a *chameleon-hash function*, require additional structure on the fingerprint map.

6.1 Preimage-resistant hash from forensic categories

We study the hash function family induced by the walk of an effective category (Definition 10), inspired by the CGL hash-function [CLG09]. Let X be the set of objects of the category. Recall that the family of functions, indexed by $A \in X$, is defined as

$$H_A : R^* \longrightarrow X : r \longmapsto \mathrm{codomain}(\mathsf{Walk}(A; r)).$$

We consider the probability distribution μ on the key space X (the limit distribution of the walk). In particular, one can sample a random function from the family by running the randomized algorithm $(A, \varphi) \leftarrow \mathsf{Walk}(O)$.

This construction does not rely on fingerprints, but fingerprints unlock its analysis. Indeed, we prove in this section that in a forensic category, this family is preimage-resistant.

Theorem 2. *In a forensic category, the family $(H_A)_{A \in X}$ is* ***preimage resistant.***

[5] In fact, a weaker HVZK notion suffices.

Proof. To show that the hash function satisfies preimage resistance, we show that for any PPT adversary $\mathcal{A}$, we must have that the quantity

$$\mathbf{Adv}_{H_A}^{\mathrm{Pre}[m]}(\mathcal{A}) = \Pr\left[H_A(x_2) = B \;\middle|\; \begin{array}{c} \mathsf{Walk}(O); x_1 \xleftarrow{\$} R^m; \\ B \leftarrow H_A(x_1); x_2 \leftarrow \mathcal{A}(A, B) \end{array} \right]$$

is negligible.

The objects $A \leftarrow \mathsf{Walk}(O)$ and $B \leftarrow H_A(x_1)$ sampled in the preimage game are both statistically indistinguishable from μ (the distance is a negligible function of the implicit security parameter). Also, by construction, an output $x_2 \leftarrow \mathcal{A}(A, B)$ which satisfies $H_A(x_2) = B$ actually encodes a morphism $(B, \psi_{A,x_2} : A \to B) = \mathsf{Walk}(A; x_2)$. Therefore, $\mathcal{A}$ solves the average case of the morphism problem (Problem 1) with probability $\mathbf{Adv}_{H_A}^{\mathrm{Pre}[m]}(\mathcal{A})$, up to a negligible error. By Proposition 1, this morphism problem is hard on average in a forensic category, so this probability $\mathbf{Adv}_{H_A}^{\mathrm{Pre}[m]}(\mathcal{A})$ must be negligible. □

We now move to primitives that do not immediately follow from the axioms of a forensic category, but can be instantiated provided some extra structure on the fingerprint map. Throughout the rest of the section we let $\mathcal{C}$ be a forensic category with fingerprint map camouflaged by a group G (see Definition 15), and denote by $\star$ the action of G on the valid fingerprints $\mathcal{M}$.

6.2 Ring signatures

A ring signature was already derived from SQIsign in [BLL24] using the generic framework of [AOS02], which constructs a linear-size ring signature from any identification protocol. It would be tempting to apply the same construction to the protocol of Section 5, but such an attempt is bound to fail because of one missing property: the identification protocol needs to satisfy *special HVZK*.

Definition 17. A sigma-protocol is *special HVZK* if there exists a PPT simulator $\mathcal{S}$ which on input a statement x and a challenge c outputs a transcript of the form (m, c, r). For any PPT adversary $\mathcal{A}$ choosing challenges c, and for a randomly sampled $(x, w) \in \mathcal{R}$, the advantage of $\mathcal{A}(x)$ in distinguishing between the outputs of $\mathcal{S}(x, c)$ and accepting transcripts conditioned on c is negligible.

Definition 18. A sigma-protocol is *special strong HVZK* if the same prescribed-challenge simulation requirement as in Definition 17 holds even when the adversary $\mathcal{A}$ is given w in addition to x, both when choosing the challenge c and when distinguishing the transcripts.

Note that the above requirements play distinct roles in the ring-signature construction framework of [AOS02]. Special HVZK allows the signer to simulate each non-signing member's transcript for a challenge determined by the preceding hash computation, while the identity of the member supplying the real response remains anonymous. Special strong HVZK supports this anonymity guarantee even after the members' secret keys are exposed.

It turns out that none of the published versions of SQIsign satisfy special and strong HVZK. The original SQIsign simulator samples a response isogeny and then constructs the challenge and commitment around it, rather than accepting a prescribed challenge. Hence, it does not satisfy special HVZK. Further, [BLL24, Theorem 3] gives an efficient distinguisher against the SQIsign simulator when given access to both x and w, thus violating strong HVZK. On the other hand, the HD variants discussed in Remark 5 instead use hint-assisted simulation [ABD+25], which does not yield a public simulator of the form $\mathcal{S}(x, c)$ required for special HVZK.

[BLL24] had to work around this problem by, in the language of this paper, changing the fingerprint map. We present here a fix for achieving special HVZK, distinct from the one in [BLL24], which leverages a camouflaged fingerprint map. This requires only a small modification to the Sigma-protocol of Section 5.

InstanceGen. Run $(A, \varphi) \leftarrow \mathsf{Walk}(O)$.

Commitment. The prover generates $(B, \psi) \leftarrow \mathsf{Walk}(O)$, samples $m \leftarrow \mathcal{M}$ uniformly, and sends (B, m) as the commitment object.

Challenge. The verifier first rejects if $m \notin \mathcal{M}$. The verifier samples $g \in G$ uniformly and sends it to the prover.

Response. The prover runs $\psi_{\mathrm{rsp}} \leftarrow \mathsf{Triangle}(\varphi, \psi, g \star m)$ to obtain a morphism ψ_{rsp} such that $\mathrm{fp}(\psi_{\mathrm{rsp}}) = g \star m$. It sends ψ_{rsp} to the verifier.

Verification. The verifier accepts if $m \in \mathcal{M}$, ψ_{rsp} has A for domain, B for codomain, and if $\mathrm{fp}(\psi_{\mathrm{rsp}}) = g \star m$.

Proposition 4. *The sigma-protocol described above is complete, special-sound and special strong HVZK.*

Proof. Completeness is clear. For soundness we proceed like in Proposition 2: two transcripts $\big((B, m), g, \psi_{\mathrm{rsp}}\big)$ and $\big((B, m), g', \psi'_{\mathrm{rsp}}\big)$ with distinct challenges $g \neq g'$, yield two morphisms $\psi_{\mathrm{rsp}}, \psi'_{\mathrm{rsp}} : A \to B$ with respective valid fingerprints $g \star m$ and $g' \star m$, which must be distinct because the action of G is free.

To prove special strong HVZK, we define the simulator as follows. On input a public key A and a challenge $g \in G$:

1. Run $(B, \chi) \leftarrow \mathsf{SimWalk}(A)$;
2. Compute $m \leftarrow \mathrm{fp}(\chi)$;
3. Return $\big((B, g^{-1} \star m), g, \chi\big)$.

It is clear that this is a valid transcript.

Observe that, for every prescribed g, the map $u \mapsto g^{-1} \star u$ is a permutation of $\mathcal{M}$. Applying the transformation $(B, u, \chi) \mapsto ((B, g^{-1} \star u), g, \chi)$ to the real and simulated distributions of Proposition 3 yields the corresponding distributions here. Note that walk-indistinguishability holds even when the distinguisher knows φ. This establishes special strong HVZK. □

Finally, we can apply the transform of [AOS02] to the protocol above to obtain a linear-size ring signature. The construction uses the same building blocks as the identification scheme and is similar to a Fiat-Shamir signature: A signer generates a commitment, then hashes it to generate a challenge meant for *the next signer in the ring*. But, instead of having the actual signer respond to the challenge, their commitment and response are generated by the special HVZK simulator, then the commitment is hashed to generate a challenge for the next signer, and so on. After the last signer has been simulated, the hash of their commitment becomes the challenge for the first, real, signer, who finally generates their response and closes the ring. The signature is the concatenation of all transcripts. We defer to [BLL24] for details.

6.3 Chameleon hash functions

Chameleon hash functions (aka *trapdoor commitments*) are trapdoor collision-resistant hash functions parametrized by a public key. We show how a forensic category with camouflaged fingerprint map can be used to upgrade a collision-resistant hash function into a chameleon one. Instantiating the framework with any of the forensic categories in Section 7 yields the first chameleon hash in the SQIsign framework, and one of the most compact post-quantum chameleon hashes in general.

We defer to [DSS20] for a thorough formalisation of chameleon hash functions, and only reproduce here a minimal set of definitions adapted to our construction.

Definition 19 (Chameleon Hash)**.** A chameleon hash function is a tuple of five PPT algorithms $(\mathsf{CHPG}, \mathsf{CHKG}, \mathsf{CHash}, \mathsf{CHeck}, \mathsf{CHange})$ with the following interface:

$\mathsf{CHPG} : (1^\lambda) \to (\mathsf{pp})$**.** On input a security parameter λ, outputs public parameters pp.

$\mathsf{CHKG} : (\mathsf{pp}) \to (\mathsf{sk}, \mathsf{pk})$**.** On input the public parameters pp outputs a private-public key pair $(\mathsf{sk}, \mathsf{pk})$.

$\mathsf{CHash} : (\mathsf{pk}, \mathsf{msg}) \to (c, h)$**.** On input a public key pk and a message msg, outputs a hash h and a *check value* (aka *randomness*) c.

$\mathsf{CHeck} : (\mathsf{pk}, \mathsf{msg}, c, h) \to b$**.** On input a public key pk, a message msg, a check value c, and a hash h, outputs a bit b, indicating whether the hash h is valid.

$\mathsf{CHange} : (\mathsf{sk}, \mathsf{pk}, \mathsf{msg}, \mathsf{msg}', c, h) \to c'$**.** On input a key pair $(\mathsf{sk}, \mathsf{pk})$, a pair of messages $\mathsf{msg}, \mathsf{msg}'$, a check value c, and a hash h, outputs a new check value c'.

Definition 20 (Correctness)**.** A chameleon hash function is correct if for all $\lambda \in \mathbb{N}$, for all public parameters $\mathsf{pp} \leftarrow \mathsf{CHPG}(1^\lambda)$, for all $(\mathsf{sk}, \mathsf{pk}) \leftarrow \mathsf{CHKG}(\mathsf{pp})$, for all $\mathsf{msg}, \mathsf{msg}' \in \{0,1\}^*$, for all $(c, h) \leftarrow \mathsf{CHash}(\mathsf{pk}, \mathsf{msg})$ and for all $c' \leftarrow \mathsf{CHange}(\mathsf{sk}, \mathsf{pk}, \mathsf{msg}, \mathsf{msg}', c, h)$

$$\mathsf{CHeck}(\mathsf{pk}, \mathsf{msg}, c, h) = \mathsf{CHeck}(\mathsf{pk}, \mathsf{msg}', c', h) = 1.$$

A chameleon hash should satisfy a form of collision-resistance, informally stating that, without knowledge of the secret key, it is hard to find pairs (msg, c) and (msg', c') with $\mathsf{msg} \neq \mathsf{msg}'$ both passing CHeck for the same hash h. Several forms of collision-resistance with access to increasingly powerful oracles are formalized in [DSS20]. Here we only analyze the weakest of these notions, with no access to oracles.

Definition 21 (Weak Collision-Resistance)**.** A chameleon hash function is weakly collision-resistant (CHwCR) if for any public parameters $\mathsf{pp} \leftarrow \mathsf{CHPG}(1^\lambda)$, any $(\mathsf{sk}, \mathsf{pk}) \leftarrow \mathsf{CHKG}(\mathsf{pp})$ and any PPT adversary $\mathcal{A}$, the probability that $\mathcal{A}(\mathsf{pk})$ returns a tuple $(\mathsf{msg}, c, \mathsf{msg}', c', h)$ such that

$$\mathsf{CHeck}(\mathsf{pk}, \mathsf{msg}, c, h) = \mathsf{CHeck}(\mathsf{pk}, \mathsf{msg}', c', h) = 1$$

and $\mathsf{msg} \neq \mathsf{msg}'$ is negligible.

Finally, an optional but common requirement for chameleon hashes is that it is infeasible to distinguish between check values produced by CHash and those produced by CHange.

Definition 22 (Strong Indistinguishability)**.** A chameleon hash function is strongly indistinguishable ($\mathsf{CH\text{-}fInd}$) if for any $\lambda \in \mathbb{N}$, for any public parameters $\mathsf{pp} \leftarrow \mathsf{CHPG}(1^\lambda)$, for any PPT adversary $\mathcal{A}$ with access to an oracle $\mathcal{O}$, and for $(\mathsf{sk}, \mathsf{pk}) \leftarrow_\$ \mathsf{CHKG}(\mathsf{pp})$, the following advantage is negligible

$$\mathsf{Adv}^{\mathsf{CH\text{-}fInd}}(\mathcal{A}) = \left| \Pr\left[\mathcal{A}\left(\mathsf{sk}, \mathsf{pk}; \mathcal{O}^{\mathsf{sk},\mathsf{pk}}_{\mathsf{CHash}}\right) = 1\right] - \Pr\left[\mathcal{A}\left(\mathsf{sk}, \mathsf{pk}; \mathcal{O}^{\mathsf{sk},\mathsf{pk}}_{\mathsf{CHange}}\right) = 1\right] \right|,$$

where $\mathcal{O}^{\mathsf{sk},\mathsf{pk}}_{\mathsf{CHash}}$ and $\mathcal{O}^{\mathsf{sk},\mathsf{pk}}_{\mathsf{CHange}}$ are defined in Algorithm 3.

Algorithm 3 Oracles in the CH-fInd game

oracle $\mathcal{O}^{\mathsf{sk},\mathsf{pk}}_{\mathsf{CHash}}(\mathsf{msg}, \mathsf{msg}')$
  $(c, h) \leftarrow \mathsf{CHash}(\mathsf{pk}, \mathsf{msg}')$
  return (c, h)

oracle $\mathcal{O}^{\mathsf{sk},\mathsf{pk}}_{\mathsf{CHange}}(\mathsf{msg}, \mathsf{msg}')$
  $(c, h) \leftarrow \mathsf{CHash}(\mathsf{pk}, \mathsf{msg})$
  $c' \leftarrow \mathsf{CHange}(\mathsf{sk}, \mathsf{pk}, \mathsf{msg}, \mathsf{msg}', c, h)$
  return (c', h)

We now let $\mathcal{C}$ be a forensic category, G be the group camouflaging its fingerprint map, and $H : \{0,1\}^* \to G$ a collision-resistant hash function. Note that H need not be surjective. Algorithm 4 describes our instantiation of a chameleon hash function. In a nutshell, a public key will be an object A in $\mathcal{C}$, with corresponding trapdoor a morphism $\varphi\colon O \to A$. To hash a message msg we hash it to $H(\mathsf{msg})$ and use it to camouflage the fingerprint of a random morphism $\psi\colon A \to B$. To adapt a hash we use the knowledge of φ to invoke $\mathsf{Triangle}$ and produce a new camouflage $H(\mathsf{msg}')$ of the fingerprint.

Algorithm 4 Chameleon hash function from a camouflaged fingerprint map

function $\mathsf{CHKG}(\mathsf{pp})$
  $(A_{\mathsf{pk}}, \varphi_{\mathsf{sk}}) \leftarrow \mathsf{Walk}(O)$
  return $\varphi_{\mathsf{sk}},\ A_{\mathsf{pk}}$

function $\mathsf{CHash}(A_{\mathsf{pk}}, \mathsf{msg})$
  $(B_h, \psi_c) \leftarrow \mathsf{SimWalk}(A_{\mathsf{pk}})$
  $m_h \leftarrow H(\mathsf{msg}) \star \mathrm{fp}(\psi_c)$
  return $\psi_c,\ (B_h, m_h)$

function $\mathsf{CHeck}(A_{\mathsf{pk}}, \mathsf{msg}, \psi_c, (B_h, m_h))$
  Parse domain and codomain of $\psi_c : A \to B$
  If $\mathrm{fp}(\psi_c) = \bot$, return 0
  return $H(\mathsf{msg}) \star \mathrm{fp}(\psi_c) = m_h$ **and** $A = A_{\mathsf{pk}}$ **and** $B = B_h$

function $\mathsf{CHange}(\varphi_{\mathsf{sk}}, A_{\mathsf{pk}}, \mathsf{msg}, \mathsf{msg}', \psi_c, (B_h, m_h))$
  if not $\mathsf{CHeck}(A_{\mathsf{pk}}, \mathsf{msg}, \psi_c, (B_h, m_h))$ **then return** $\bot$
  return $\psi'_c \leftarrow \mathsf{Triangle}(\varphi_{\mathsf{sk}}, \psi_c \circ \varphi_{\mathsf{sk}}, H(\mathsf{msg}')^{-1} \star m_h)$

Theorem 3. *If $\mathcal{C}$ is a forensic category and* fp *is camouflaged, the chameleon hash function in Algorithm 4 is correct.*

Proof. Let $A_{\mathsf{pk}}, \varphi_{\mathsf{sk}}$ be a key pair and $\mathsf{msg}, \mathsf{msg}'$ two messages. Let

$$\psi_c, (B_h, m_h) \leftarrow \mathsf{CHash}(A_{\mathsf{pk}}, \mathsf{msg}),$$
$$\psi'_c \leftarrow \mathsf{CHange}(\varphi_{\mathsf{sk}}, A_{\mathsf{pk}}, \mathsf{msg}, \mathsf{msg}', \psi_c, (B_h, m_h)).$$

It is immediate that the output of CHash is accepted by CHeck.

The only technicality to check is that the call to $\mathsf{Triangle}$ in CHange succeeds. The morphisms φ_{sk} and $\psi_c \circ \varphi_{\mathsf{sk}}$ both have O as domain, and the value $H(\mathsf{msg}')^{-1} \star m_h = H(\mathsf{msg}')^{-1} H(\mathsf{msg}) \star \mathrm{fp}(\psi_c)$ is a valid fingerprint because $\mathcal{M}$ is closed under the action $\star$. Thus CHange outputs $\psi'_c : A_{\mathsf{pk}} \to B_h$ such that $H(\mathsf{msg}') \star \mathrm{fp}(\psi'_c) = m_h$, hence $\mathsf{CHeck}(A_{\mathsf{pk}}, \mathsf{msg}', \psi'_c, (B_h, m_h))$ accepts. □

Theorem 4. *If $\mathcal{C}$ is a forensic category,* fp *is camouflaged, and H is collision-resistant, the chameleon hash function in Algorithm 4 is weakly collision-resistant.*

Proof. Let A_{pk} be a public key. An adversary $\mathcal{A}$ that wins the weak collision-resistance game computes $\mathsf{msg}, \psi_c, \mathsf{msg}', \psi'_c, (B_h, m_h)$ such that

$$\mathsf{CHeck}(A_{\mathsf{pk}}, \mathsf{msg}, \psi_c, (B_h, m_h)) = \mathsf{CHeck}(A_{\mathsf{pk}}, \mathsf{msg}', \psi'_c, (B_h, m_h)) = 1.$$

Then by definition ψ_c and ψ'_c both have domain A_{pk} and codomain B_h, their fingerprints lie in $\mathcal{M}$, and

$$H(\mathsf{msg}) \star \mathrm{fp}(\psi_c) = m_h = H(\mathsf{msg}') \star \mathrm{fp}(\psi'_c).$$

Now we have two possibilities:

1. $H(\mathsf{msg}) = H(\mathsf{msg}')$, but $\mathsf{msg} \neq \mathsf{msg}'$, thus $\mathcal{A}$ has produced a collision for H;
2. $H(\mathsf{msg}) \neq H(\mathsf{msg}')$, but the group action is free, thus

$$\mathrm{fp}(\psi_c) = H(\mathsf{msg})^{-1} \star H(\mathsf{msg}') \star \mathrm{fp}(\psi'_c) \neq \mathrm{fp}(\psi'_c),$$

thus $\mathcal{A}$ has produced a witness for the hard language of the category.

Either way the probability that $\mathcal{A}$ wins the game is negligible. □

Theorem 5. *If $\mathcal{C}$ is a forensic category and* fp *is camouflaged, the chameleon hash function in Algorithm 4 is strongly indistinguishable.*

Proof. We bound the advantage of an adversary in the strong indistinguishability game by using the game hopping strategy outlined in Algorithm 5.

Algorithm 5 Game hops for the CH-fInd game

oracle $\mathcal{O}^{\varphi_{\mathsf{sk}}, A_{\mathsf{pk}}}_{\mathsf{CHange}}(\mathsf{msg}, \mathsf{msg}')$
$(B_h, \psi_c) \leftarrow \mathsf{SimWalk}(A_{\mathsf{pk}})$
$m_h \leftarrow H(\mathsf{msg}) \star \mathrm{fp}(\psi_c)$
$m' \leftarrow H(\mathsf{msg}')^{-1} \star m_h$
$\psi'_c \leftarrow \mathsf{Triangle}(\varphi_{\mathsf{sk}}, \psi_c \circ \varphi_{\mathsf{sk}}, m')$
return $\psi'_c, (B_h, m_h)$

oracle $\mathcal{O}^{\varphi_{\mathsf{sk}}, A_{\mathsf{pk}}}_{1}(\mathsf{msg}, \mathsf{msg}')$
$(B_h, \psi) \leftarrow \mathsf{Walk}(O)$
$m_h \leftarrow_{\$} \mathcal{M}$
$\psi'_c \leftarrow \mathsf{Triangle}(\varphi_{\mathsf{sk}}, \psi, H(\mathsf{msg}')^{-1} \star m_h)$
return $\psi'_c, (B_h, m_h)$

oracle $\mathcal{O}^{\varphi_{\mathsf{sk}}, A_{\mathsf{pk}}}_{2}(\mathsf{msg}, \mathsf{msg}')$
$(B_h, \psi) \leftarrow \mathsf{Walk}(O)$
$m \leftarrow_{\$} \mathcal{M}$
$\psi'_c \leftarrow \mathsf{Triangle}(\varphi_{\mathsf{sk}}, \psi, m)$
$m_h \leftarrow H(\mathsf{msg}') \star m$
return $\psi'_c, (B_h, m_h)$

oracle $\mathcal{O}^{\varphi_{\mathsf{sk}}, A_{\mathsf{pk}}}_{\mathsf{CHash}}(\mathsf{msg}, \mathsf{msg}')$
$(B_h, \psi'_c) \leftarrow \mathsf{SimWalk}(A_{\mathsf{pk}})$
$m \leftarrow \mathrm{fp}(\psi'_c)$
$m_h \leftarrow H(\mathsf{msg}') \star m$
return $\psi'_c, (B_h, m_h)$

In the first hop we replace $\mathcal{O}_{\mathsf{CHange}}$ with $\mathcal{O}_1$. The output distribution is unchanged, up to a negligible statistical distance:

- B_h is distributed according to μ;
- $\mathrm{fp}(\psi_c)$ is uniformly distributed and so is m_h, because $m \mapsto H(\mathsf{msg}) \star m$ is a permutation on $\mathcal{M}$;
- $\psi_c \circ \varphi_{\mathsf{sk}}$ and ψ both have domain O and codomain B_h, thus the output of $\mathsf{Triangle}$ has the same distribution in both oracles.

When moving from $\mathcal{O}_1$ to $\mathcal{O}_2$, fingerprints m and m_h are uniformly distributed in both cases, and we ensure that $\mathrm{fp}(\psi'_c) = H(\mathsf{msg}') \star m$, thus the two games are indistinguishable.

Steps 2–4 in $\mathcal{O}_2$ correspond to $\mathcal{O}^{\varphi_{\mathsf{sk}}}_{\mathsf{Triangle}}$ in Algorithm 1, and we replace them with steps 2–3 in $\mathcal{O}_{\mathsf{CHash}}$, which correspond to $\mathcal{O}^{\varphi_{\mathsf{sk}}}_{\mathsf{SimWalk}}$. Thus the advantage in distinguishing between $\mathcal{O}_2$ and $\mathcal{O}_{\mathsf{CHash}}$ is bounded by $\mathsf{Adv}^{\mathsf{Walk\text{-}Ind}}$.

We have thus successfully replaced $\mathcal{O}_{\mathsf{CHange}}$ with $\mathcal{O}_{\mathsf{CHash}}$ in the CH-fInd game, bounding the advantage in distinguishing between the two by $2\mathsf{Adv}^{\mathsf{Walk\text{-}Ind}}$ plus negligible factors. □

7 Instantiations of forensic categories

We now instantiate our framework in the setting of elliptic curves. For the rest of this article, we fix a prime p.

7.1 The category $\mathbf{SSing}_p$

Theorem 6. *The category $\mathbf{SSing}_p$ of supersingular elliptic curves is an effective category: under the Generalized Riemann Hypothesis when $p \equiv 1 \bmod 8$, unconditionally otherwise. Furthermore, the endomorphism ring of the origin object can be computed in polynomial time.*

Proof. We show that $\mathbf{SSing}_p$ satisfies the fundamental properties of a pre-effective category as in Definition 9, that it has an origin object E_0 whose representation is known, and it is equipped with a Walk algorithm as per Definition 10.

Claim. *The objects and morphisms have unique representation.*

There are several possibilities to uniquely represent isomorphism classes of supersingular curves. The j-invariant classifies elliptic curves up to isomorphism and would be a natural candidate, however in what follows we need to work with explicit coordinates, so we will choose a uniquely defined plane model instead.

A common choice in practice is the Montgomery model $y^2 = x^3 + Ax^2 + x$. Every supersingular curve is isomorphic to a Montgomery curve with $A \in \mathbb{F}_{p^2}$, and for every isomorphism class there are at most 6 distinct Montgomery coefficients A, it is thus sufficient to take the lexicographically least among them.

We now sketch the method for how to obtain a unique representation for an isogeny $\varphi\colon E \to E'$. A detailed discussion can be found in [Rob25, Section 6].

1. For each $i = 1, \dots, k$, let ℓ_i be the i-th prime not dividing $\deg(\varphi)$, with k the smallest integer such that $\prod_{i=1}^{k} \ell_i > \deg(\varphi)$.

2. For each $i = 1, \dots, k$, compute a deterministic (i.e. depending only on E and ℓ_i) basis (P_i, Q_i) of $E[\ell_i]$.

3. For each $i = 1, \dots, k$, compute the images $(\varphi(P_i), \varphi(Q_i))$.

4. Then, as shown in [Rob25], an isogeny interpolation algorithm makes the data $(E, E', \deg(\varphi), (P_i, Q_i, \varphi(P_i), \varphi(Q_i))_{i=1}^{k})$ an efficient representation of φ, in the sense of Definition 3. It is a unique representation thanks to the deterministic choice of bases.

Claim. $\mathbf{SSing}_p$ *is a pre-effective category.*

1. *Deciding validity of encodings.* To decide if an equation encodes a supersingular curve, it is sufficient to apply a polynomial-time point-counting algorithm [Sch85]. To decide if interpolation data encodes an isogeny, one can run the interpolation algorithm [Rob25], which will succeed if and only if the interpolation data is valid.

2. *Computing domain and codomain.* These are part of the interpolation data.

3. *Composing morphisms.* To compose morphisms $\varphi\colon A \to B$ and $\psi\colon B \to C$, following [Rob25, Section 2.5], we use the efficient representations of φ and ψ to evaluate φ and ψ on arbitrary points, and thus construct an efficient representation of $\psi \circ \varphi$.

4. *Computing identity morphisms.* We return the unique efficient representation of the identity map $1_E\colon E \to E$.

Claim. *There exists an algorithm to find an "origin curve" E_0 and to compute its endomorphism ring.*

As shown in [EHL+18, Proposition 3], given a prime p, there exists an algorithm to find a supersingular curve E_0 of known endomorphism ring. The algorithm terminates in polynomial time if $p \not\equiv 1 \bmod 8$, or for any p if the Generalized Riemann Hypothesis is true.

Claim. *There exists a walk algorithm* (Walk, μ) *as described in Definition 10.*

Fix any prime number $\ell = O(1)$ (e.g., $\ell = 2$). The walking algorithm Walk is instantiated by following a (non-backtracking) random walk of length x from the input curve in the ℓ-isogeny graph. The stationary distribution of this walk is given by[6] $\mu(E) = \frac{24}{(p-1)\#\mathrm{Aut}(E)}$. Since this graph is Ramanujan [Piz90], the distance between the distribution of the output of the walk and the stationary distribution is a negligible function of the walk's length x. For non-backtracking random walks, the convergence to the stationary distribution is quantified in [BCC+23, Theorem 11]. In particular, there is a bound $X = O(\log_\ell(p))$ such that a random walk of length $x > X$ is at statistical distance $1/p$ from the stationary distribution. In particular, taking $X = \lceil (\log p)^2 \rceil$, the statistical distance from the stationary distribution is negligible in p for every walk length $x > X$.

□

Definition 23. For any choice of a prime ℓ, the outputs of Walk described above are called the (non-backtracking) ℓ-isogeny walks. We write Walk_ℓ when the prime is not clear from context. Note that its stationary distribution μ does not depend on ℓ.

7.1.1 Torsion bases.

We now detail one way of instantiating a fingerprint map for the category $\mathbf{SSing}_p$, which relies on torsion subgroups of the elliptic curves. First, we will restrict the category $\mathbf{SSing}_p$ to morphisms of smooth degree, to facilitate their manipulation. Recall that an integer is S-smooth if all of its prime factors are smaller than S.

Definition 24. Let $S > 0$. The category $\mathbf{SSing}_p^S$ is the category with the same objects as $\mathbf{SSing}_p$, and with all the morphisms of $\mathbf{SSing}_p$ with S-smooth degree.

For the rest of this paper, we fix an arbitrary bound $S = (\log p)^{O(1)}$, with $\ell \leq S$, and work in the effective category $\mathbf{SSing}_p^S$.

We fix a positive integer N not divisible by p nor ℓ, so we have $E[N] \cong \mathbb{Z}/N\mathbb{Z} \times \mathbb{Z}/N\mathbb{Z}$.

We also require that N is S-smooth, so that discrete logarithms in $E[N]$ are efficiently computable, and that $N \mid (p^n - 1)$ for some $n \in \mathrm{polylog}(p)$, implying that elements of $E[N]$ have coordinates in $\mathbb{F}_{p^n}$, so that the algorithms we describe next can run in polynomial time.

To define fingerprint maps, it will be necessary to attach to each supersingular curve E a *deterministic basis* (P_E, Q_E) which generates the whole $E[N]$, chosen deterministically among the possible bases of $E[N]$ (see, e.g., [AAA+24]).

7.2 The kernel fingerprint

This first instantiation aims to emulate (classical) SQIsign [DKL+20]. Let $\psi\colon E \to E'$ be an isogeny and define $\ker(\psi)[N] \coloneqq \ker(\psi) \cap E[N]$. Let us define the map

$$\mathsf{prefp}(\psi) = \begin{cases} \ker(\psi)[N], & \text{if } \ker(\psi)[N] \cong \mathbb{Z}/N\mathbb{Z} \\ \bot, & \text{otherwise.} \end{cases}$$

[6] since $\#\mathrm{Aut}(E) = 2$ for all curves except at most 2, and $\#\mathrm{Aut}(E) \leq 6$ for these exceptions, this distribution is essentially the same as the uniform distribution

In essence, this is the fingerprint of ψ. However, as written it depends on the domain curve of ψ, whereas we need a single set of valid fingerprints for all curves. We call prefp a *prefingerprint map*, and to define the actual fingerprint map we use the well known bijection between cyclic subgroups of $(\mathbb{Z}/N\mathbb{Z})^2$ and the projective line $\mathbb{P}(\mathbb{Z}/N\mathbb{Z})$. Let (P_E, Q_E) be the deterministic basis for $E[N]$, the set of pairs (a, b) such that $aP_E + bQ_E \in \mathsf{prefp}(\psi)$ forms an additive subgroup of $(\mathbb{Z}/N\mathbb{Z})^2$. We choose a canonical generator of this subgroup as fingerprint for ψ as follows.

Definition 25 (Kernel Fingerprint)**.** The *N-kernel fingerprint* for an isogeny $\psi\colon E \to E'$ is

$$\mathrm{fp}_{\mathcal{K}}(\psi) = \begin{cases} (1, b), & \text{if } \ker(\psi)[N] = \langle P_E + bQ_E \rangle \text{ for } b \in (\mathbb{Z}/N\mathbb{Z}) \\ (a, 1), & \text{if } \ker(\psi)[N] = \langle aP_E + Q_E \rangle \text{ for } a \in (\mathbb{Z}/N\mathbb{Z}) \setminus (\mathbb{Z}/N\mathbb{Z})^\times \\ \bot, & \text{otherwise,} \end{cases}$$

where (P_E, Q_E) is the deterministic basis of $E[N]$.

Theorem 7 (Effectivity)**.** *The fingerprint map $\mathrm{fp}_{\mathcal{K}}$ is effective.*

Proof. Given an efficient representation for ψ, we first compute $\mathsf{prefp}(\psi)$ by evaluating it on $E[N]$. If $\mathsf{prefp}(\psi) = \bot$, we return $\bot$. Then we compute $\mathrm{fp}_{\mathcal{K}}(\psi)$ by a generalized discrete logarithm computation, taking time polynomial in $\log(N)$. □

Theorem 8 (Hardness)**.** *If the* ISOGENY *problem is hard (in the worst case), the fingerprint map $\mathrm{fp}_{\mathcal{K}}$ satisfies the Hard on average property.*

Proof. By the result [MW25, Theorem 1.2], if the ISOGENY problem is hard in the worst case, then ONEEND is hard on average.

Therefore, to prove the result, we will actually prove that if $\mathrm{fp}_{\mathcal{K}}$ does not satisfy the Hard on average property, then ONEEND is not hard on average either (hence ISOGENY is not hard in the worst case).

Since $\mathrm{fp}_{\mathcal{K}}$ is just an encoding of prefp, for the sake of simplicity we work with the prefp definition. We now assume that $\mathrm{fp}_{\mathcal{K}}$ is not hard on average: there exists a polynomial time algorithm $\mathcal{A}$ which, on input a uniformly random E_1, returns $\varphi, \psi\colon E_1 \to E_2$ such that $(E_1, (\varphi, \psi)) \in \mathcal{L}$ with non-negligible probability ε.

Let E_1 be uniformly random, and let us solve ONEEND: run $\mathcal{A}(E_1)$, which produces a pair (φ, ψ), and return the composition $\hat{\psi} \circ \varphi$. We need to prove that with good probability, $\hat{\psi} \circ \varphi$ is a solution of ONEEND, i.e. $\hat{\psi} \circ \varphi \in \mathrm{End}(E_1) \setminus \mathbb{Z}$.

Indeed, with probability ε the isogenies φ, ψ have domain E_1 and a common codomain E_2, and satisfy $\mathrm{fp}_{\mathcal{K}}(\varphi) \neq \mathrm{fp}_{\mathcal{K}}(\psi)$ and $\mathrm{fp}_{\mathcal{K}}(\varphi), \mathrm{fp}_{\mathcal{K}}(\psi) \neq \bot$. In that situation, $\ker(\varphi)[N]$ and $\ker(\psi)[N]$ are *distinct* cyclic subgroups of order N.

Assume, for the sake of contradiction, that $\hat{\psi} \circ \varphi \in \mathbb{Z}$, so there exist $a, b \in \mathbb{Z}$ such that $[a]\varphi = [b]\psi$. Without loss of generality, $\gcd(a, N) = 1$, and we get

$$\ker(\varphi)[N] = \ker([a]\varphi)[N] = \ker([b]\psi)[N].$$

In particular, we must have $\gcd(b, N) = 1$ (otherwise the non-cyclic $E_1[N]$ would have a non-trivial intersection with the cyclic $\ker(\varphi)[N]$). Therefore, $\ker([b]\psi)[N] = \ker(\psi)[N]$, and we deduce the contradiction $\ker(\varphi)[N] = \ker(\psi)[N]$. □

Theorem 9 (Triangularizability)**.** *The fingerprint map $\mathrm{fp}_{\mathcal{K}}$ satisfies Triangularizability.*

Proof. Let $\varphi\colon E_0 \to E_1, \psi\colon E_0 \to E_2$ be two isogenies from the origin curve E_0 (of known endomorphism ring), and $(a, b) \in \mathcal{M}$ a given valid fingerprint. Algorithm 6 instantiates $\mathsf{Triangle}(\varphi, \psi, (a, b))$. Its workhorse is the algorithm RSigningKLPT described in [BLL24, Algorithm 3] (where we have adjusted inputs for consistency with our notation[7]).

[7] Note that we provide isogenies as input to RSigningKLPT, instead of the corresponding ideals — since all our isogenies are smooth, the conversion from isogenies to ideals is efficient.

Algorithm 6 The Triangle algorithm for the kernel fingerprint map

Input: $\varphi\colon E_0 \to E_1, \psi\colon E_0 \to E_2, (a,b) \in \mathcal{M}$

1: Let (P_1, Q_1) be the chosen basis of $E_1[N]$
2: Let $P_{(a,b)} = aP_1 + bQ_1$
3: Compute the quotient isogeny $\varphi_{\mathrm{Velu}}\colon E_1 \to E_1/\langle P_{(a,b)}\rangle$ using Vélu's formulas
4: Use $\mathsf{RSigningKLPT}_\ell(\varphi_{\mathrm{Velu}} \circ \varphi, \psi)$ to compute $\psi_{\mathsf{KLPT}}\colon E_1/\langle P_{(a,b)}\rangle \to E_2$ of degree a power of ℓ

Output: $\psi_{\mathrm{resp}} = \psi_{\mathsf{KLPT}} \circ \varphi_{\mathrm{Velu}}$, where $\psi_{\mathrm{resp}}\colon E_1 \to E_2$

We show that ψ_{resp} is indeed an isogeny from E_1 to E_2 such that $\mathrm{fp}_{\mathcal{K}}(\psi_{\mathrm{resp}}) = (a,b)$. By construction, ψ_{resp} has domain E_1 and codomain E_2.

Now, we have

$$\langle P_{(a,b)}\rangle = \ker(\varphi_{\mathrm{Velu}}) \subseteq \ker(\psi_{\mathsf{KLPT}} \circ \varphi_{\mathrm{Velu}}) = \ker(\psi_{\mathrm{resp}}).$$

Since ψ_{resp} has degree $N\ell^m$ (for some integer m), and $\langle P_{(a,b)}\rangle$ is a cyclic subgroup of order N, we deduce that

$$\langle P_{(a,b)}\rangle = \ker(\psi_{\mathrm{resp}})[N] = \mathsf{prefp}(\psi_{\mathrm{resp}}).$$

Therefore, $\mathrm{fp}_{\mathcal{K}}(\psi_{\mathrm{resp}}) = (a,b)$, finishing the proof. □

7.2.1 Walk indistinguishability.

Like zero-knowledge in the original SQIsign [DKL+20] and in the ring signature [BLL24], walk-indistinguishability of the kernel fingerprint map relies on a computational assumption that we state below in simplified form (see [DKL+20, DLLW23, BLL24] for details).

Assumption 1. *The output of* RSigningKLPT *[BLL24, Algorithm 3] is computationally indistinguishable from uniformly random cyclic isogenies of the same degree and with the same domain.*

Theorem 10. *Under Assumption 1, the fingerprint map* $\mathrm{fp}_{\mathcal{K}}$ *is walk-indistinguishable.*

Proof. SimWalk is identical to Algorithm 6, except that we replace line 4 with

$$(E_2, \psi_{\mathsf{KLPT}}) \leftarrow \mathsf{Walk}_\ell(E_1/\langle P_{(a,b)}\rangle).$$

Note that this algorithm does not make use of the "secret" inputs φ, ψ, and only reads E_1 and (a,b) instead, so is a valid instantiation of SimWalk. Clearly, under Assumption 1, the output of this algorithm is computationally indistinguishable from the output of Triangle conditioned on the fingerprint being (a,b). □

7.2.2 Undisturbing walk.

We now show that the kernel fingerprint map, $\mathrm{fp}_{\mathcal{K}}$, also satisfies the additional Undisturbing Walk property. This also justifies our use of the Hard on average property.

Theorem 11. Walk_ℓ *is an Undisturbing Walk for the fingerprint map* $\mathrm{fp}_{\mathcal{K}}$.

Proof. We want to show that $\mathrm{fp}_{\mathcal{K}}$ and Walk_ℓ satisfy the commutative square in Equation (3). Like before, it will be equivalent and simpler to argue about prefp instead of $\mathrm{fp}_{\mathcal{K}}$.

Let $\chi : A \to C$ be an output of Walk_ℓ and let $\psi \in \mathrm{Hom}(C, B)$. The degree of χ is coprime to N, thus χ maps $A[N]$ to $C[N]$ bijectively. Hence

$$\ker(\psi \circ \chi)[N] \simeq \ker(\psi)[N],$$

thus $\mathsf{prefp}(\psi) = \perp$ if and only if $\mathsf{prefp}(\psi \circ \chi) = \perp$.

Now suppose $\mathsf{prefp}(\psi \circ \chi) = \langle P\rangle \neq \perp$, then $\mathsf{prefp}(\psi) = \chi(\langle P\rangle) = \langle \chi(P)\rangle$. But χ maps N-torsion points bijectively, so it also maps cyclic groups of order N bijectively, hence it induces a bijection on valid fingerprints. □

7.2.3 Camouflage.

We now prove that there is a group G that camouflages the kernel fingerprint map $\mathrm{fp}_{\mathcal{K}}$, in the sense of Definition 15: there is an efficiently computable free group action of G on $\mathcal{M}$. For simplicity we assume that $N = q^n$ is a power of a prime q, but the construction generalizes easily. We actually prove a stronger property by defining a free *and transitive* group action. Let us identify $\mathbb{Z}/q^n\mathbb{Z}$ with $\{0, 1, \ldots, q^n - 1\}$, so the set $\mathcal{M}$ is the set of pairs of integers $\mathcal{M} = \mathcal{M}' \sqcup \mathcal{M}_q$ (where $\sqcup$ denotes the disjoint union) with

$$\mathcal{M}' = \{1\} \times (\mathbb{Z}/q^n\mathbb{Z}) \text{ and } \mathcal{M}_q = \left((\mathbb{Z}/q^n\mathbb{Z}) \setminus (\mathbb{Z}/q^n\mathbb{Z})^\times\right) \times \{1\}.$$

Let $X = \{0, 1, \ldots, q^{n-1} - 1\}$. We now prove that $\mathcal{M}$ can naturally be decomposed as a disjoint union of $q + 1$ copies of X. First, we have $\mathcal{M}_q = \{qx \mid x \in X\} \times \{1\}$, and a bijection $\iota_q : X \to \mathcal{M}_q : x \mapsto (qx, 1)$. Then, we have $\mathcal{M}' = \bigsqcup_{i=0}^{q-1} \mathcal{M}_i$ with $\mathcal{M}_i = \{1\} \times \{i + qx \mid x \in X\}$, and for each i a bijection $\iota_i : X \to \mathcal{M}_i : x \mapsto (1, i + qx)$.

In summary, $\mathcal{M}$ is the disjoint union of the $q + 1$ sets $(\mathcal{M}_i)_{i=0}^q$, where each $\mathcal{M}_i$ is in bijection with X. We consider the group $G = (\mathbb{Z}/q^{n-1}\mathbb{Z}) \times (\mathbb{Z}/(q+1)\mathbb{Z})$, and define the group action by combining the natural actions of $\mathbb{Z}/q^{n-1}\mathbb{Z}$ on X with the action of $\mathbb{Z}/(q+1)\mathbb{Z}$ on the indices of the list $(\mathcal{M}_i)_{i=0}^q$. More explicitly, for $(x, j) \in G$, and $m \in \mathcal{M}_i$, we define

$$(x, j) \star m = \iota_{i+j \bmod (q+1)}(\iota_i^{-1}(m) + x) \in \mathcal{M}_{i+j \bmod (q+1)}. \tag{5}$$

Theorem 12. *The action of $G = (\mathbb{Z}/q^{n-1}\mathbb{Z}) \times (\mathbb{Z}/(q+1)\mathbb{Z})$ on the set of valid fingerprints $\mathcal{M}$ defined in (5) is free and transitive. In particular, G camouflages the kernel fingerprint map $\mathrm{fp}_{\mathcal{K}}$.*

Proof. Let $m \in \mathcal{M}_i$, and $(x, j), (x', j') \in G$ such that $(x, j) \star m = (x', j') \star m$. Since $(x, j) \star m \in M_{i+j \bmod (q+1)}$ and $(x', j') \star m \in M_{i+j' \bmod (q+1)}$, we must have $j \equiv j' \bmod (q + 1)$, hence $j = j'$. Now, we have

$$\iota_{i+j \bmod (q+1)}(\iota_i^{-1}(m) + x) = (x, j) \star m = (x', j) \star m = \iota_{i+j \bmod (q+1)}(\iota_i^{-1}(m) + x').$$

Since $\iota_{i+j \bmod (q+1)}$ is a bijection, we deduce $\iota_i^{-1}(m) + x \equiv \iota_i^{-1}(m) + x' \bmod q^{n-1}$, hence $x \equiv x' \bmod q^{n-1}$. Therefore $(x, j) = (x', j')$ in G, which proves that the action is free. Since the action is free and $\#G = \#\mathcal{M} = (q + 1)q^{n-1}$, it must also be transitive. □

7.3 The category $\mathbf{SSing}_{p,N}^S$

Our second instantiation uses a different effective category of supersingular elliptic curves.

Definition 26 (Action of $M_{2\times 2}(\mathbb{Z})$ on points)**.** Given any elliptic curve E, any pair of points $B = (P, Q) \in E^2$, and a matrix $M = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$, we write

$$B \cdot M = (P, Q) \cdot M = (aP + cQ, bP + dQ).$$

Note that for any $M, M' \in M_{2\times 2}(\mathbb{Z})$, we have $(B \cdot M) \cdot M' = B \cdot (MM')$.

Definition 27. Fix a prime p, a positive number $S = (\log p)^{O(1)}$, an S-smooth integer $N = p^{O(1)}$. We write $\mathbf{SSing}_{p,N}^S$ the category where

- **objects** are pairs (E, B), where E is a supersingular elliptic curve over $\mathbb{F}_{p^2}$, and B is any basis of $E[N]$ (i.e. curves with full N-level structure), and
- **morphisms** from (E, B) to (E', B') are pairs (φ, γ), where
  - $\varphi : E \to E'$ is an isogeny of S-smooth degree coprime to N, and
  - $\gamma \in \mathrm{GL}_2(\mathbb{Z}/N\mathbb{Z})$ is the matrix such that $\varphi(B) = B' \cdot \gamma$, where if $B = (P, Q)$ we define $\varphi(B) = (\varphi(P), \varphi(Q))$.

Remark 6. The category $\mathbf{SSing}^S_{p,N}$ is a *product category*, an extension of the concept of the Cartesian product of two sets. We are interested in this kind of categories as they offer a natural way to characterize SQInstructor [BFLS26].

Theorem 13. *The category of supersingular elliptic curves $\mathbf{SSing}^S_{p,N}$ is an effective category: under the Generalized Riemann Hypothesis when $p \equiv 1 \bmod 8$, unconditionally otherwise. Furthermore, the endomorphism ring of the origin object can be computed in polynomial time.*

Proof. We show that $\mathbf{SSing}^S_{p,N}$ satisfies the fundamental properties of a pre-effective category as in Definition 9, that it has an origin object (E_0, B_0) whose representation is known, and it is equipped with a Walk algorithm as per Definition 10.

Claim. *The objects and morphisms have unique representation.*

As in Theorem 6, to uniquely represent isomorphism classes of supersingular curves, we use the Montgomery model $y^2 = x^3 + Ax^2 + x$. This is detailed in [Rob25, Section 6]. This gives us the unique representation of E. Once we have such a canonical model for E, the basis $B_E = (P_E, Q_E)$ is uniquely expressed in the induced coordinate system. Hence we can uniquely represent the objects (E, B_E).

We now show how to obtain a unique representation of any morphism (φ, γ). To obtain a unique representation for an isogeny $\varphi\colon E \to E'$, we use the same method as in Section 7.1. A detailed discussion can be found in [Rob25, Section 6]. As for $\gamma \in \mathrm{GL}_2(\mathbb{Z}/N\mathbb{Z})$, we may employ any standard representation of matrices with modular integer coefficients.

Claim. $\mathbf{SSing}^S_{p,N}$ *is a pre-effective category.*

The proof is essentially the same as in Theorem 6. To test whether $B = (P, Q)$ is indeed a basis of $E[N]$, it is sufficient to verify that $e_N(P, Q)$ is of multiplicative order N, where e_N is the N-th Weil pairing.

Claim. *There exists an algorithm to find an "origin curve" (E_0, B_0), and to compute its endomorphism ring.*

This, again, can be done as in Theorem 6 — plus any algorithm to find a (deterministic) basis of $E_0[N]$.

Algorithm 7 The walking algorithm Walk_ℓ

Input: (E, B) and a walk length x
Output: An object (E', B') and a morphism $\tau : (E, B) \to (E', B')$

1: Compute a non-backtracking random isogeny $\varphi\colon E \to E'$ of degree ℓ^x, corresponding to a random walk of length x in the ℓ-isogeny graph
2: Sample $\delta \in \mathrm{GL}_2(\mathbb{Z}/N\mathbb{Z})$ uniformly at random
3: Parse $B = (P, Q)$
4: Compute $B' = (\varphi(P), \varphi(Q)) \cdot \delta$
5: Define the morphism $\tau := (\varphi, \delta^{-1})\colon (E, B) \to (E', B')$
6: **return** $(E', B'), \tau$

Claim. *There exists a walk algorithm (Walk, μ) as described in Definition 10.*

Fix any prime number $\ell = O(\log N)$ coprime to N. The walking algorithm Walk_ℓ is instantiated in Algorithm 7.

As in the proof of Theorem 6, the Ramanujan property of the ℓ-isogeny graph ensures that the distance between the distribution of E and the stationary distribution decays exponentially in the length x of the walk. Since $\delta \in \mathrm{GL}_2(\mathbb{Z}/N\mathbb{Z})$ is chosen uniformly at random, and $\varphi(B)$ is a basis of $E'[N]$, we obtain that $B' = \varphi(B) \cdot \delta$ is a uniformly random basis of $E'[N]$. □

7.4 A fingerprint from the Borel level structure

This fingerprint map is defined using *level structures*, a tool increasingly used in isogeny-based cryptography [CL26, BCC+23, PW24, DFP24]. Interestingly, it provides examples of homomorphic fingerprint maps. Informally a level structure of level N on E is a basis of $E[N]$ *up to change of basis* by some subgroup of invertible matrices. Formally, we follow [CL26, Definition 1.1].

Definition 28 (Level structure on elliptic curves)**.** Fix a positive integer N (of known factorization), and a subgroup Γ of $\mathrm{GL}_2(\mathbb{Z}/N\mathbb{Z}) = \mathrm{Aut}((\mathbb{Z}/N\mathbb{Z})^2)$. Let k be a field where N is invertible. For an elliptic curve E/k, a level Γ structure on E is an isomorphism $\iota\colon (\mathbb{Z}/N\mathbb{Z})^2 \to E[N]$ considered up to composition with an element of Γ, i.e. two isomorphisms ι, ι' are equivalent if there exists an element h in Γ such that $\iota = \iota' \circ h$.

Level structures for certain choices of Γ can be interpreted in more explicit ways. We list some of the relevant examples below:

- *Full level structure:* When $\Gamma = \{\mathrm{Id}\}$, a level Γ structure is equivalent to the choice of a basis of $E[N]$. In particular, any object in the category $\mathbf{SSing}^S_{p,N}$ is an elliptic curve with a full level structure.
- *Scalar level structure:* When $\Gamma = \{\lambda\mathrm{Id} \mid \lambda \in (\mathbb{Z}/N\mathbb{Z})^\times\}$, the subgroup of scalar matrices, a level Γ structure is equivalent to the choice of a basis up to multiplication by a scalar in $(\mathbb{Z}/N\mathbb{Z})^\times$.
- *Borel level structure:* When $\Gamma = \{\left(\begin{smallmatrix} * & * \\ 0 & * \end{smallmatrix}\right)\}$ is the subgroup of upper-triangular matrices, a level Γ structure is equivalent to the choice of a cyclic subgroup of order N in $E[N]$. Explicitly, the cyclic subgroup $\langle(1,0)\rangle \subset (\mathbb{Z}/N\mathbb{Z})^2$ is stable under the action of Γ, and a level structure ι is uniquely determined by the image $\iota(\langle(1,0)\rangle) \subset E[N]$.

The Borel level structure fingerprint map expresses how an isogeny "acts" on the level structure. For the rest of this section, we consider the subgroup $\Gamma = \{\left(\begin{smallmatrix} * & * \\ 0 & * \end{smallmatrix}\right)\}$ of upper-triangular matrices.

Definition 29 (Borel Level Structure Fingerprint Map)**.** Let (φ, γ) be a morphism from (E_1, B_1) to (E_2, B_2), as in Definition 27. Then, define the Borel level structure fingerprint map $\mathrm{fp}_{\mathrm{lvl}}$ as follows

$$\mathrm{fp}_{\mathrm{lvl}}(\varphi, \gamma) = \gamma\Gamma \subset \mathrm{GL}_2.$$

Note that every morphism (φ, γ) has a valid fingerprint, so there are no undefined cases. Hence, the set of valid fingerprints is just the quotient set $\mathcal{M} = \mathrm{GL}_2(\mathbb{Z}/N\mathbb{Z})/\Gamma$.

Remark 7. While we focus our efforts on the Borel level structure, other choices of Γ could be used to define fingerprint maps. Note however that not every choice of Γ will lead to a well-defined fingerprint map. In order to satisfy the *Hard* property, it is crucial that Γ contains the scalar matrices.

Moreover, considering different choices of Γ allows the fingerprint map to satisfy additional properties, such as those detailed in Section 4.4. For example, *the scalar level structure*, $\Gamma = \{\lambda\mathrm{Id} \mid \lambda \in (\mathbb{Z}/N\mathbb{Z})^\times\}$, also defines a forensic category, and additionally satisfies the *homomorphic property* (Definition 16).

Lemma 3 (Unique representation)**.** *Any value* $\mathrm{fp}_{\mathrm{lvl}}(\psi, \gamma) \in \mathcal{M}$ *can be represented uniquely, and this representation can be computed in polynomial time.*

Proof. Writing down $\gamma\Gamma$ is not sufficient to uniquely represent the fingerprint, as different isogenies with the same fingerprint may have different matrices in the same coset. We need to define a unique representative for the coset $\gamma\Gamma$.

Recall that $\Gamma = \left\{ \left(\begin{smallmatrix} a & b \\ 0 & d \end{smallmatrix}\right) | ad \in (\mathbb{Z}/N\mathbb{Z})^\times \right\}$. Let $\gamma = \left(\begin{smallmatrix} x & y \\ u & v \end{smallmatrix}\right)$ and assume that x and $v - uyx^{-1}$ are invertible. Then, Algorithm 8 can be used to uniquely represent $\mathrm{fp}_{\mathrm{lvl}}(\psi, \gamma) = \gamma\Gamma$:

Algorithm 8 Unique representation for the Borel level structure fingerprint

Input: $\begin{pmatrix} x & y \\ u & v \end{pmatrix}$

1: $\begin{pmatrix} x & y \\ u & v \end{pmatrix} \begin{pmatrix} x^{-1} & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & -y \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & (v - uyx^{-1})^{-1} \end{pmatrix}$

Output: $\begin{pmatrix} 1 & 0 \\ ux^{-1} & 1 \end{pmatrix}$

Evaluating the complexity of Algorithm 8 is straightforward. Since

$$\left(\begin{smallmatrix} x^{-1} & 0 \\ 0 & 1 \end{smallmatrix}\right) \left(\begin{smallmatrix} 1 & -y \\ 0 & 1 \end{smallmatrix}\right) \left(\begin{smallmatrix} 1 & 0 \\ 0 & (v-uyx^{-1})^{-1} \end{smallmatrix}\right) \in \Gamma,$$

it remains to show that $\left(\begin{smallmatrix} 1 & 0 \\ ux^{-1} & 1 \end{smallmatrix}\right)$ is the unique element of this form in the orbit.

For the sake of contradiction, suppose that in the same orbit as $\left(\begin{smallmatrix} 1 & 0 \\ ux^{-1} & 1 \end{smallmatrix}\right)$, there exists a second element of the same form, $\left(\begin{smallmatrix} 1 & 0 \\ u'x^{-1} & 1 \end{smallmatrix}\right)$, such that $u \neq u'$ in $\mathbb{Z}/N\mathbb{Z}$. Then, $\left(\begin{smallmatrix} 1 & 0 \\ ux^{-1} & 1 \end{smallmatrix}\right) = \left(\begin{smallmatrix} 1 & 0 \\ u'x^{-1} & 1 \end{smallmatrix}\right) h$, for $h \in \Gamma$. Hence, $\left(\begin{smallmatrix} 1 & 0 \\ u'x^{-1} & 1 \end{smallmatrix}\right)^{-1} \left(\begin{smallmatrix} 1 & 0 \\ ux^{-1} & 1 \end{smallmatrix}\right) = h$, so $\left(\begin{smallmatrix} 1 & 0 \\ (u-u')x^{-1} & 1 \end{smallmatrix}\right) = h$. Since $h \in \Gamma$, h is a matrix of the form $\left(\begin{smallmatrix} * & * \\ 0 & * \end{smallmatrix}\right)$, so $(u - u')x^{-1} = 0$. As x^{-1} is invertible, $u - u' = 0$, leading to a contradiction.

Comment: If x and $v - uyx^{-1}$ are not invertible, the reader can follow a similar argument to handle other cases.

□

Theorem 14 (Effectivity)**.** *The fingerprint map* $\mathrm{fp}_{\mathrm{lvl}}$ *is effective.*

Proof. Recall that $\mathrm{fp}_{\mathrm{lvl}}(\varphi, \gamma) = \gamma\Gamma$. Hence, to compute the fingerprint, we simply need to compute the unique coset representative of $\gamma\Gamma$. To do this we apply Lemma 3. □

Theorem 15 (Hardness)**.** *If the* ISOGENY *problem is hard in the worst case, the fingerprint map* $\mathrm{fp}_{\mathrm{lvl}}$ *satisfies the Hard on average property.*

Proof. As in the proof of Theorem 8, we use the result in [MW25, Theorem 1.2], stating that if the ISOGENY problem is hard in the worst case, then the ONEEND problem is hard on average.

Therefore, to prove the result, we will actually prove that if $\mathrm{fp}_{\mathrm{lvl}}$ does not satisfy the Hard on average property, then ONEEND is not hard on average either (hence ISOGENY is not hard in the worst case).

Suppose that $\mathrm{fp}_{\mathrm{lvl}}$ does not satisfy the Hard on average property. That is there exists a polynomial time algorithm $\mathcal{A}$, such that on input (E_1, B_1), where (E_1, B_1) follows the uniform distribution, $\mathcal{A}(E_1, B_1)$ returns $(\varphi, \gamma), (\psi, \delta)$: $(E_1, B_1) \to (E_2, B_2)$ such that $((E_1, B_1), ((\varphi, \gamma), (\psi, \delta))) \in \mathcal{L}$ with non-negligible probability.

Let E_1 be a uniformly random supersingular elliptic curve, and let us solve $\mathsf{ONEEND}(E_1)$. First, we sample a uniformly random basis B_1 of $E_1[N]$. Run $\mathcal{A}(E_1, B_1)$ which produces a pair of morphisms $((\varphi, \gamma), (\psi, \delta))$; return the composition $\hat{\psi} \circ \varphi$.

With non-negligible probability, the output of $\mathcal{A}$ satisfies $((E_1, B_1), ((\varphi, \gamma), (\psi, \delta))) \in \mathcal{L}$ (and in particular, $\mathrm{fp}_{\mathrm{lvl}}(\varphi, \gamma) \neq \mathrm{fp}_{\mathrm{lvl}}(\psi, \delta)$). Assume this is the case, and we finish the proof by proving that $\hat{\psi} \circ \varphi$ is a solution of ONEEND (i.e. $\hat{\psi} \circ \varphi \in \mathrm{End}(E_1) \setminus \mathbb{Z}$). Clearly,

we have $\hat{\psi} \circ \varphi \in \mathrm{End}(E_1)$. Assume for the sake of contradiction that $\hat{\psi} \circ \varphi$ is a scalar endomorphism, then $\delta^{-1}\gamma$ is a scalar matrix. This implies that $\gamma\Gamma = \delta\Gamma$. But we have that $\mathrm{fp}_{\mathrm{lvl}}((\varphi, \gamma)) \neq \mathrm{fp}_{\mathrm{lvl}}((\psi, \delta))$, meaning that $\gamma\Gamma$ and $\delta\Gamma$ are distinct cosets of Γ, leading to a contradiction.

□

7.4.1 Triangularizability.

The article [BFLS26] describes algorithms which, reframed in our language, amount to triangularizing fingerprints for several types of level structures. When specializing to Borel level structures, the algorithm uses the same building blocks as Algorithm 6. We sketch it below, and refer to [BFLS26] for details.

Theorem 16. *The fingerprint map* $\mathrm{fp}_{\mathrm{lvl}}$ *satisfies the Triangularizability property.*

Sketch. Let $(\varphi, \gamma)\colon (E_0, B_0) \to (E_1, B_1)$ and $(\psi, \delta)\colon (E_0, B_0) \to (E_2, B_2)$ be two morphisms from the origin curve (E_0, B_0), as in Definition 27. Let $M \in \mathcal{M}$ be a given valid fingerprint. Recall that $\mathrm{End}(E_0)/N\,\mathrm{End}(E_0) \simeq M_{2\times 2}(\mathbb{Z}/N\mathbb{Z})$. Furthermore, we know a basis of $\mathrm{End}(E_0)$ and the factorization of N, so by solving a linear system modulo N, we can efficiently find an endomorphism $\alpha \in \mathrm{End}(E_0)$ such that $\mathrm{fp}_{\mathrm{lvl}}(\alpha, \zeta) = (\delta\Gamma)^{-1}M\gamma\Gamma$. Let us write $\mathrm{fp}_{\mathrm{lvl}}(\alpha, \zeta) = \zeta\Gamma$. Then, letting $\chi = \psi \circ \alpha \circ \hat{\varphi}$, and $\sigma = \delta\zeta\gamma^{-1}$ we find that $(\chi, \sigma) : (E_1, B_1) \to (E_2, B_2)$ is a morphism with $\mathrm{fp}_{\mathrm{lvl}}(\chi, \sigma) = M$.

Now (χ, σ) is already a valid triangularization, but let us make it "look" like a walk, in preparation for the *Walk-indistinguishable* property.

We now apply BorelSigningKLPT [BFLS26] with inputs I and I_N, where I is the ideal of χ and I_N the ideal of the isogeny $E_1 \to E_1/\langle Q_1\rangle = E_1/\iota_1(\langle(1, 0)\rangle)$. The output is an ℓ-walk $\chi'\colon E_1 \to E_2$ with same fingerprint σ as χ — i.e., a solution to $\mathsf{Triangle}((\varphi, \gamma), (\psi, \delta), M)$.

□

Before we prove Walk-Indistinguishability for the Borel level structure fingerprint map we introduce the following Assumption 2. It is similar to Assumption 1, and a more detailed discussion can be found in [BFLS26, Section 5.3].

Assumption 2. *The output of* BorelSigningKLPT *[BFLS26, Algorithm 7] is computationally indistinguishable from uniformly random cyclic isogenies with the same level structure, same degree and same domain.*

Theorem 17 (Walk indistinguishability)**.** *Under Assumption 2, the fingerprint map* $\mathrm{fp}_{\mathrm{lvl}}$ *is Walk-indistinguishable.*

Proof. The algorithm SimWalk is defined as Walk_ℓ (Algorithm 7) with the same length parameter as e in BorelSigningKLPT [BFLS26, Algorithm 7]. The indistinguishability then follows immediately by construction, thanks to the *ad hoc* Assumption 2. □

8 Conclusion

This article initiates the study of abstractions of the SQIsign digital signature scheme, stripping away the algebraic geometry and isolating core properties via the very general language of categories. The key notion of our framework is the fingerprint map, which captures the singular way in which SQIsign uses challenges.

We do not claim that our definition of a forensic category is the final answer to abstracting SQIsign. On one hand, it captures well the original SQIsign design, and it is expressive enough to derive several more cryptographic primitives. On the other hand, several SQIsign-adjacent schemes remain out of its scope. Most notably, SQIsignHD [DLRW24] and its successors [BDD+24, NOC+24] deviate from the classical notion of honest-verifier

zero-knowledge in a way that is not compatible with our choice of axioms. We believe that the definition of a forensic category could be adjusted to accommodate these constructions, and encourage further research in this direction.

Different instantiations (e.g., $\mathbf{SSing}_p^S$ with the kernel fingerprint map and $\mathbf{SSing}_{p,N}^S$ with the Borel level structure fingerprint map) each have specific properties. If we were to express these properties as axioms of the framework, then the instantiations would satisfy distinct sets of axioms. Therefore, for maximum expressivity, the framework might need some modularity. Such an approach, with increased flexibility, might allow the framework to include cousin constructions like SQIPrime [DF24] or Leroux' VRF [Ler25].

The instantiations presented in Section 7 are not the end of their own story either. Our framework highlights the grey zones of the literature, particularly regarding walk indistinguishability. This property still relies on *ad hoc* assumptions, and the strength of the indistinguishability statement could still be improved by further refining the approach of [BLL24].

We hope that our approach will encourage research into making the instantiations more general and robust, just as the formalisation of (un)restricted effective group actions [BY91, Cou06, ADMP20] encouraged the line of research that led to the unrestricted instantiations Clapoti [PR23] and PEGASIS [DEF+25].

We hope that forensic categories can be used to construct more advanced primitives. They offer protocol designers a new opportunity to create isogeny-based protocols without the burden of algebraic geometry. In particular we believe that, in the right hands, the *homomorphic* property (Definition 16) has more to offer.

Finally, we wish to emphasize that forensic categories do not require any isogeny, until we attempt to instantiate them. In principle, they could be realized using very different objects. We encourage the search for instantiations that do not rely on algebraic geometry.

References

[AAA+24] Marius A. Aardal, Gora Adj, Diego F. Aranha, Andrea Basso, Isaac Andrés Canales Martínez, Jorge Chávez-Saab, Maria Corte-Real Santos, Pierrick Dartois, Luca De Feo, Max Duparc, Jonathan Komada Eriksen, Tako Boris Fouotsa, Décio Luiz Gazzoni Filho, Basil Hess, David Kohel, Antonin Leroux, Patrick Longa, Luciano Maino, Michael Meyer, Kohei Nakagawa, Hiroshi Onuki, Lorenz Panny, Sikhar Patranabis, Christophe Petit, Giacomo Pope, Krijn Reijnders, Damien Robert, Francisco Rodríguez Henríquez, Sina Schaeffler, and Benjamin Wesolowski. SQIsign. Technical report, National Institute of Standards and Technology, 2024. available at https://csrc.nist.gov/Projects/pqc-dig-sig/round-2-additional-signatures.

[AABN02] Michel Abdalla, Jee Hea An, Mihir Bellare, and Chanathip Namprempre. From identification to signatures via the Fiat-Shamir transform: Minimizing assumptions for security and forward-security. In Lars R. Knudsen, editor, *EUROCRYPT 2002*, volume 2332 of *LNCS*, pages 418–433. Springer, Berlin, Heidelberg, April / May 2002. doi:10.1007/3-540-46035-7_28.

[ABD+25] Marius A. Aardal, Andrea Basso, Luca De Feo, Sikhar Patranabis, and Benjamin Wesolowski. A complete security proof of SQIsign. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part VI*, volume 16005 of *LNCS*, pages 190–222. Springer, Cham, August 2025. doi:10.1007/978-3-032-01887-8_7.

[ADMP20] Navid Alamati, Luca De Feo, Hart Montgomery, and Sikhar Patranabis. Cryptographic group actions and applications. In Shiho Moriai and Huaxiong

Wang, editors, *ASIACRYPT 2020, Part II*, volume 12492 of *LNCS*, pages 411–439. Springer, Cham, December 2020. doi:10.1007/978-3-030-64834-3_14.

[AOS02] Masayuki Abe, Miyako Ohkubo, and Koutarou Suzuki. 1-out-of-n signatures from a variety of keys. In Yuliang Zheng, editor, *ASIACRYPT 2002*, volume 2501 of *LNCS*, pages 415–432. Springer, Berlin, Heidelberg, December 2002. doi:10.1007/3-540-36178-2_26.

[BBC+25] Andrea Basso, Giacomo Borin, Wouter Castryck, Maria Corte-Real Santos, Riccardo Invernizzi, Antonin Leroux, Luciano Maino, Frederik Vercauteren, and Benjamin Wesolowski. PRISM: Simple and compact identification and signatures from large prime degree isogenies. In Tibor Jager and Jiaxin Pan, editors, *PKC 2025, Part III*, volume 15676 of *LNCS*, pages 300–332. Springer, Cham, May 2025. doi:10.1007/978-3-031-91826-1_10.

[BBMP24] Michele Battagliola, Giacomo Borin, Alessio Meneghetti, and Edoardo Persichetti. Cutting the GRASS: Threshold GRoup action signature schemes. In Elisabeth Oswald, editor, *CT-RSA 2024*, volume 14643 of *LNCS*, pages 460–489. Springer, Cham, May 2024. doi:10.1007/978-3-031-58868-6_18.

[BCC+23] Andrea Basso, Giulio Codogni, Deirdre Connolly, Luca De Feo, Tako Boris Fouotsa, Guido Maria Lido, Travis Morrison, Lorenz Panny, Sikhar Patranabis, and Benjamin Wesolowski. Supersingular curves you can trust. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part II*, volume 14005 of *LNCS*, pages 405–437. Springer, Cham, April 2023. doi:10.1007/978-3-031-30617-4_14.

[BDD+24] Andrea Basso, Pierrick Dartois, Luca De Feo, Antonin Leroux, Luciano Maino, Giacomo Pope, Damien Robert, and Benjamin Wesolowski. SQIsign2D-West - the fast, the small, and the safer. In Kai-Min Chung and Yu Sasaki, editors, *ASIACRYPT 2024, Part III*, volume 15486 of *LNCS*, pages 339–370. Springer, Singapore, December 2024. doi:10.1007/978-981-96-0891-1_11.

[Bel02] Mihir Bellare. A note on negligible functions. *Journal of Cryptology*, 15(4):271–284, September 2002. doi:10.1007/s00145-002-0116-x.

[BF03] Dan Boneh and Matthew Franklin. Identity-based encryption from the weil pairing. *SIAM Journal on Computing*, 32(3):586–615, 2003. doi:10.1137/S0097539701398521.

[BF23] Andrea Basso and Tako Boris Fouotsa. New SIDH countermeasures for a more efficient key exchange. In Jian Guo and Ron Steinfeld, editors, *ASIACRYPT 2023, Part VIII*, volume 14445 of *LNCS*, pages 208–233. Springer, Singapore, December 2023. doi:10.1007/978-981-99-8742-9_7.

[BFLS26] Giacomo Borin, Luca De Feo, Guido Lido, and Sina Schaeffler. The sqinstructor: a guide to sqisign and the deuring correspondence with level structures. In Nadia Heninger and Mike Rosulek, editors, *Advances in Cryptology – CRYPTO 2026*, pages 537–569, Cham, 2026. Springer Nature Switzerland. doi:10.1007/978-3-032-35398-6_18.

[BLL24] Giacomo Borin, Yi-Fu Lai, and Antonin Leroux. Erebor and durian: Full anonymous ring signatures from quaternions and isogenies. *CiC*, 1(4):4, 2024. doi:10.62056/ava3zivrzn.

[BM25] Andrea Basso and Luciano Maino. POKÉ: A compact and efficient PKE from higher-dimensional isogenies. In Serge Fehr and Pierre-Alain Fouque, editors, *EUROCRYPT 2025, Part II*, volume 15602 of *LNCS*, pages 94–123. Springer, Cham, May 2025. doi:10.1007/978-3-031-91124-8_4.

[BMP23] Andrea Basso, Luciano Maino, and Giacomo Pope. FESTA: Fast encryption from supersingular torsion attacks. In Jian Guo and Ron Steinfeld, editors, *ASIACRYPT 2023, Part VII*, volume 14444 of *LNCS*, pages 98–126. Springer, Singapore, December 2023. doi:10.1007/978-981-99-8739-9_4.

[BMPS20] Jean-François Biasse, Giacomo Micheli, Edoardo Persichetti, and Paolo Santini. LESS is more: Code-based signatures without syndromes. In Abderrahmane Nitaj and Amr M. Youssef, editors, *AFRICACRYPT 20*, volume 12174 of *LNCS*, pages 45–65. Springer, Cham, July 2020. doi:10.1007/978-3-030-51938-4_3.

[BY91] Gilles Brassard and Moti Yung. One-way group actions. In Alfred J. Menezes and Scott A. Vanstone, editors, *CRYPTO'90*, volume 537 of *LNCS*, pages 94–107. Springer, Berlin, Heidelberg, August 1991. doi:10.1007/3-540-38424-3_7.

[CII+23] Mingjie Chen, Muhammad Imran, Gábor Ivanyos, Péter Kutas, Antonin Leroux, and Christophe Petit. Hidden stabilizers, the isogeny to endomorphism ring problem and the cryptanalysis of pSIDH. In Jian Guo and Ron Steinfeld, editors, *ASIACRYPT 2023, Part III*, volume 14440 of *LNCS*, pages 99–130. Springer, Singapore, December 2023. doi:10.1007/978-981-99-8727-6_4.

[CK20] Leonardo Colò and David Kohel. Orienting supersingular isogeny graphs. *Journal of Mathematical Cryptology*, 14(1):414–437, 2020. doi:10.1515/jmc-2019-0034.

[CL26] Giulio Codogni and Guido Maria Lido. Spectral theory of isogeny graphs. *Journal of Number Theory*, 286:131–184, 2026. URL: https://www.sciencedirect.com/science/article/pii/S0022314X26000594, doi:10.1016/j.jnt.2026.02.006.

[CLG09] Denis Xavier Charles, Kristin E. Lauter, and Eyal Z. Goren. Cryptographic hash functions from expander graphs. *Journal of Cryptology*, 22(1):93–113, January 2009. doi:10.1007/s00145-007-9002-x.

[Cou06] Jean-Marc Couveignes. Hard homogeneous spaces. Cryptology ePrint Archive, Report 2006/291, 2006. URL: https://eprint.iacr.org/2006/291.

[CSD+23] Jorge Chavez-Saab, Maria Corte-Real Santos, Luca De Feo, Jonathan Komada Eriksen, Basil Hess, David Kohel, Antonin Leroux, Patrick Longa, Michael Meyer, Lorenz Panny, Sikhar Patranabis, Christophe Petit, Francisco Rodríguez Henríquez, Sina Schaeffler, and Benjamin Wesolowski. SQIsign. Technical report, National Institute of Standards and Technology, 2023. available at https://csrc.nist.gov/Projects/pqc-dig-sig/round-1-additional-signatures.

[Dam10] Ivan Damgård. On σ-protocols. http://www.cs.au.dk/~ivan/Sigma.pdf, 2010.

[DEF+25] Pierrick Dartois, Jonathan Komada Eriksen, Tako Boris Fouotsa, Arthur Herlédan Le Merdy, Riccardo Invernizzi, Damien Robert, Ryan Rueger, Frederik Vercauteren, and Benjamin Wesolowski. PEGASIS: Practical effective class group action using 4-dimensional isogenies. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part I*, volume 16000 of *LNCS*, pages 67–99. Springer, Cham, August 2025. doi:10.1007/978-3-032-01855-7_3.

[Deu41] Max Deuring. Die Typen der Multiplikatorenringe elliptischer Funktionenkörper. *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, 14(1):197–272, December 1941. doi:10.1007/BF02940746.

[DF24] Max Duparc and Tako Boris Fouotsa. SQIPrime: A dimension 2 variant of SQISignHD with non-smooth challenge isogenies. In Kai-Min Chung and Yu Sasaki, editors, *ASIACRYPT 2024, Part III*, volume 15486 of *LNCS*, pages 396–429. Springer, Singapore, December 2024. doi:10.1007/978-981-96-0891-1_13.

[DFP24] Luca De Feo, Tako Boris Fouotsa, and Lorenz Panny. Isogeny problems with level structure. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part VII*, volume 14657 of *LNCS*, pages 181–204. Springer, Cham, May 2024. doi:10.1007/978-3-031-58754-2_7.

[DKL+20] Luca De Feo, David Kohel, Antonin Leroux, Christophe Petit, and Benjamin Wesolowski. SQISign: Compact post-quantum signatures from quaternions and isogenies. In Shiho Moriai and Huaxiong Wang, editors, *ASIACRYPT 2020, Part I*, volume 12491 of *LNCS*, pages 64–93. Springer, Cham, December 2020. doi:10.1007/978-3-030-64837-4_3.

[DLLW23] Luca De Feo, Antonin Leroux, Patrick Longa, and Benjamin Wesolowski. New algorithms for the Deuring correspondence - towards practical and secure SQISign signatures. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part V*, volume 14008 of *LNCS*, pages 659–690. Springer, Cham, April 2023. doi:10.1007/978-3-031-30589-4_23.

[DLRW24] Pierrick Dartois, Antonin Leroux, Damien Robert, and Benjamin Wesolowski. SQIsignHD: New dimensions in cryptography. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part I*, volume 14651 of *LNCS*, pages 3–32. Springer, Cham, May 2024. doi:10.1007/978-3-031-58716-0_1.

[DM20] Luca De Feo and Michael Meyer. Threshold schemes from isogeny assumptions. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, *PKC 2020, Part II*, volume 12111 of *LNCS*, pages 187–212. Springer, Cham, May 2020. doi:10.1007/978-3-030-45388-6_7.

[DSS20] David Derler, Kai Samelin, and Daniel Slamanig. Bringing order to chaos: The case of collision-resistant chameleon-hashes. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, *PKC 2020, Part I*, volume 12110 of *LNCS*, pages 462–492. Springer, Cham, May 2020. doi:10.1007/978-3-030-45374-9_16.

[EHL+18] Kirsten Eisenträger, Sean Hallgren, Kristin E. Lauter, Travis Morrison, and Christophe Petit. Supersingular isogeny graphs and endomorphism rings: Reductions and solutions. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part III*, volume 10822 of *LNCS*, pages 329–368. Springer, Cham, April / May 2018. doi:10.1007/978-3-319-78372-7_11.

[FMP23] Tako Boris Fouotsa, Tomoki Moriya, and Christophe Petit. M-SIDH and MD-SIDH: Countering SIDH attacks by masking information. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part V*, volume 14008 of *LNCS*, pages 282–309. Springer, Cham, April 2023. doi:10.1007/978-3-031-30589-4_10.

[GPS17] Steven D. Galbraith, Christophe Petit, and Javier Silva. Identification protocols and signature schemes based on supersingular isogeny problems. In Tsuyoshi Takagi and Thomas Peyrin, editors, *ASIACRYPT 2017, Part I*, volume 10624 of *LNCS*, pages 3–33. Springer, Cham, December 2017. doi:10.1007/978-3-319-70694-8_1.

[GPS20] Steven D. Galbraith, Christophe Petit, and Javier Silva. Identification protocols and signature schemes based on supersingular isogeny problems. *Journal of Cryptology*, 33(1):130–175, January 2020. doi:10.1007/s00145-019-09316-0.

[JD11] David Jao and Luca De Feo. Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. In Bo-Yin Yang, editor, *PQCrypto 2011*, pages 19–34. Springer, Berlin, Heidelberg, November / December 2011. doi:10.1007/978-3-642-25405-5_2.

[Jou04] Antoine Joux. A one round protocol for tripartite Diffie–Hellman. *Journal of Cryptology*, 17(4):263–276, September 2004. doi:10.1007/s00145-004-0312-y.

[KLP26] Veronika Kuchta, Jason T. LeGrow, and Edoardo Persichetti. Post-quantum (partially) blind signatures from matrix code equivalence. *Cryptography and Communications*, 18(5):1669–1721, April 2026. URL: http://dx.doi.org/10.1007/s12095-026-00879-x, doi:10.1007/s12095-026-00879-x.

[Koh96] David Russell Kohel. *Endomorphism rings of elliptic curves over finite fields*. PhD thesis, University of California, Berkeley, 1996.

[KR98] Hugo Krawczyk and Tal Rabin. Chameleon hashing and signatures. Cryptology ePrint Archive, Report 1998/010, 1998. URL: https://eprint.iacr.org/1998/010.

[Ler25] Antonin Leroux. Verifiable random function from the deuring correspondence and higher dimensional isogenies. In Serge Fehr and Pierre-Alain Fouque, editors, *EUROCRYPT 2025, Part VII*, volume 15607 of *LNCS*, pages 167–194. Springer, Cham, May 2025. doi:10.1007/978-3-031-91098-2_7.

[MP23] Hart Montgomery and Sikhar Patranabis. A computational category-theoretic approach to cryptography and average-case complexity. *Transactions on Mathematical Cryptology*, 3(2):24–52, Oct. 2023. URL: https://journals.flvc.org/mathcryptology/article/view/134668.

[MW25] Arthur Herlédan Le Merdy and Benjamin Wesolowski. Unconditional foundations for supersingular isogeny-based cryptography. In Benny Applebaum and Huijia (Rachel) Lin, editors, *TCC 2025, Part III*, volume 16270 of *LNCS*, pages 266–297. Springer, Cham, December 2025. doi:10.1007/978-3-032-12296-4_9.

[NOC+24] Kohei Nakagawa, Hiroshi Onuki, Wouter Castryck, Mingjie Chen, Riccardo Invernizzi, Gioella Lorenzon, and Frederik Vercauteren. SQIsign2D-East: A new signature scheme using 2-dimensional isogenies. In Kai-Min Chung and Yu Sasaki, editors, *ASIACRYPT 2024, Part III*, volume 15486 of *LNCS*, pages 272–303. Springer, Singapore, December 2024. doi:10.1007/978-981-96-0891-1_9.

[Piz80] Arnold Pizer. An algorithm for computing modular forms on $\gamma_0(n)$. *Journal of Algebra*, 64(2):340–390, 1980. doi:10.1016/0021-8693(80)90151-9.

[Piz90] Arnold K Pizer. Ramanujan graphs and hecke operators. *Bulletin of the American Mathematical Society*, 23(1):127–137, 1990.

[PR23] Aurel Page and Damien Robert. Introducing clapoti(s): Evaluating the isogeny class group action in polynomial time. Cryptology ePrint Archive, Paper 2023/1766, 2023. URL: https://eprint.iacr.org/2023/1766.

[PW24] Aurel Page and Benjamin Wesolowski. The supersingular endomorphism ring and one endomorphism problems are equivalent. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part VI*, volume 14656 of *LNCS*, pages 388–417. Springer, Cham, May 2024. doi:10.1007/978-3-031-58751-1_14.

[Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 84–93. ACM Press, May 2005. doi:10.1145/1060590.1060603.

[Rob25] Damien Robert. On theÂ efficient representation ofÂ isogenies. In Andrzej Dąbrowski, Josef Pieprzyk, and Jacek Pomykała, editors, *Number-Theoretic Methods in Cryptology*, pages 3–84, Cham, 2025. Springer Nature Switzerland.

[RS06] Alexander Rostovtsev and Anton Stolbunov. Public-Key Cryptosystem Based On Isogenies. Cryptology ePrint Archive, Report 2006/145, 2006. URL: https://eprint.iacr.org/2006/145.

[Sch85] René Schoof. Elliptic curves over finite fields and the computation of square roots mod p. *Mathematics of Computation*, 44(170):483–494, 1985. doi:10.2307/2007968.

[Sil09] Joseph H Silverman. *The arithmetic of elliptic curves*, volume 106 of *Graduate Texts in Mathematics*. Springer, 2009.

[Sto12] Anton Stolbunov. *Cryptographic schemes based on isogenies*. PhD thesis, Norges teknisk-naturvitenskapelige universitet, Fakultet for informasjonsteknologi, matematikk og elektroteknikk, Institutt for telematikk, 2012.

[Vél71] Jacques Vélu. Isogénies entre courbes elliptiques. *Comptes-Rendus de l'Académie des Sciences*, 273:238–241, 1971.

[Wes22] Benjamin Wesolowski. The supersingular isogeny path and endomorphism ring problems are equivalent. In *62nd FOCS*, pages 1100–1111. IEEE Computer Society Press, February 2022. doi:10.1109/FOCS52979.2021.00109.

[Wes24] Benjamin Wesolowski. Random walks in number-theoretic cryptology. HDR Thesis Manuscript, École Normale Supérieure de Lyon (ENS Lyon), August 2024. Presented on 29 August 2024.