Isogeny problems with level structure

Luca De Feo[1][0000−0002−9321−0773], Tako Boris Fouotsa[2][0000−0003−1821−8406], and Lorenz Panny[3]

[1] IBM Research Europe, Zürich, Switzerland; eurocrypt24@defeo.lu
[2] EPFL, Lausanne, Switzerland; tako.fouotsa@epfl.ch
[3] Technische Universität München, Germany; lorenz@yx7.cc

Abstract. Given two elliptic curves and the degree of an isogeny between them, finding the isogeny is believed to be a difficult problem—upon which rests the security of nearly any isogeny-based scheme.
If, however, to the data above we add information about the behavior of the isogeny on a large enough subgroup, the problem can become easy, as recent cryptanalyses on SIDH have shown.
Between the restriction of the isogeny to a full N-torsion subgroup and no "torsion information" at all lies a spectrum of interesting intermediate problems, raising the question of how easy or hard each of them is. Here we explore *modular isogeny problems* where the torsion information is masked by the action of a group of 2×2 matrices. We give reductions between these problems, classify them by their difficulty, and link them to security assumptions found in the literature.

Keywords: Isogenies · Post-quantum · Security reductions.

1 Introduction

Isogeny-based cryptography is a fast-changing field, with new schemes and assumptions appearing at a sustained pace and, recently, a series of powerful attacks shaking its foundations. It may be difficult for an outsider to make sense of the scores of different assumptions, and understand what level of security they actually offer. Luckily, in parallel with the accumulation of new assumptions, works on security reductions have helped somewhat systematize the landscape and reduce the amount of hypotheses to keep track of [41,67,66,59]. Also worth mentioning is the project "Is SIKE broken yet?"[1], which tries to collect most isogeny assumptions and track the best reductions and attacks known on them.

The goal of this work is to add another layer to our understanding of isogeny-based cryptography by giving a framework that encompasses several seemingly unrelated assumptions and proving reductions between them. We start from two well-known problems: on one hand SIDH [50], also known as the Computational Supersingular Isogeny (CSSI) problem, which was recently solved quite

* Author list in alphabetical order; see https://ams.org/profession/leaders/CultureStatement04.pdf. Date of this document: 2024-03-18.

[1] https://issikebrokenyet.github.io/

efficiently and in fair generality [15,56,63]; on the other hand the generic (fixed-degree) isogeny problem, for which no general classical or quantum algorithm better than exponential is known, and which is the foundation of all isogeny-based cryptography. SIDH has sometimes been described as an isogeny problem *with torsion-point information* [60], but what is known as "torsion-point information" in the cryptographic community has long been known as a special type of *level structure* among mathematicians. By generalizing the SIDH problem to other types of level structures, we obtain a family of problems, some easy, some hard, which happen to have reductions to/from isogeny problems that had previously appeared in the literature without apparent connection. Along the way, we extend the SIDH attacks to a more general setting, we prove that, somewhat ironically, some proofs of knowledge based on SIDH are *at least as hard* as CSIDH [18], and we improve on the best generic algorithms to compute isogenies of ordinary curves [36].

Torsion-point information, a.k.a. level structures. In the generic fixed-degree isogeny problem, one is given a pair of curves E, E' isogenous of exponentially large degree d, and is tasked with finding a d-isogeny $\phi : E \to E'$, for instance by exhibiting a generator for $\ker \phi$. In SIDH and variants, next to E, E' we add an (ordered) basis (P, Q) of $E[N]$ for some fixed parameter $N \approx d$ coprime to d (typically $N = 2^a$ or $N = 3^b$) and its image $(\phi(P), \phi(Q))$ under the secret isogeny. The goal is again to find ϕ.

The extra information provided by $(P, Q, \phi(P), \phi(Q))$ has been called *torsion-point information* in [60] and follow-ups. It is precisely this information that is exploited by the attacks on SIDH [15,56,63]; its absence in the generic isogeny problem is the reason why the rest of isogeny-based cryptography still stands.

A curve E together with a basis (P, Q) of $E[N]$ is called a *full level structure of level* N in the literature on modular curves. More generally, a level structure of level N is a basis of $E[N]$ *up to change of basis* by some group of matrices $\Gamma \leq \mathrm{GL}_2(\mathbb{Z}/N\mathbb{Z})$. So, for example, when Γ is the group of diagonal matrices a Γ-level structure is the set of bases (aP, bQ) for all $a, b \in \mathbb{Z}/N\mathbb{Z}$ such that ab is invertible, and when $\Gamma = \mathrm{GL}_2$ the associated level structure is the set of all possible bases.

Once we interpret SIDH as a generic isogeny problem between curves with full level structure, it becomes natural to define isogeny problems with Γ-level structures for arbitrary subgroups Γ. The interpretation is that we are given tuples (E, P, Q) and (E', P', Q'), with the promise that there exists an isogeny $\phi : E \to E'$ mapping (P, Q) to one of the bases in the orbit of (P', Q') by Γ. Thus, when Γ is the diagonal group, the Γ-SIDH problem is to find ϕ knowing that $\phi(P) = aP'$ and $\phi(Q) = bQ'$ for some unknown $a, b \in \mathbb{Z}/N\mathbb{Z}$ such that ab is invertible. The GL_2-SIDH problem is simply the generic isogeny problem.

Related work. Level structures were first considered in the context of isogeny-based cryptography by Arpin [3], who studied the relation between supersingular isogeny graphs with level structure and Eichler orders in a quaternion algebra.

In [6] it is proved that supersingular isogeny graphs with *Borel level structure* have the Ramanujan property, which is then used to construct a proof of isogeny knowledge with statistical zero-knowledge. The follow-up work [26] proves similar expansion properties for graphs with other level structures.

Level structures in disguise appeared in isogeny schemes quickly after SIDH was broken: the key exchange M-SIDH [43] is a variant of SIDH using the group of scalar matrices to mask the torsion point information, thus blocking the attacks; the trapdoor one-way function FESTA [8] uses a diagonal matrix as a trapdoor to mask a standard SIDH problem. Recent attacks against special instances of M-SIDH and FESTA [21] function by, essentially, reducing the Γ-SIDH problem to a plain SIDH problem.

This work was prompted by a question raised at the Leuven Isogeny Days 2022: is it possible to solve the "SIDH with only one point" problem in polynomial time, i.e., the problem where one is only given (E, P) and $(E', \phi(P))$ with P of order N? We answered in the affirmative when N contains a large smooth square factor, thus, in particular, when $N = 2^a$ or $N = 3^b$, by giving a reduction to the standard SIDH problem (see Corollary 12). This result having been circulated privately for more than a year, it has already been used to break some *ad hoc* instances of class group actions on ordinary curves [16, § 6], and to extend the attacks on overstretched FESTA [21, Remark 5].

Contributions. We generalize the SIDH problem to isogeny problems with arbitrary level structures. In doing so, we:

- Identify several interesting types of level structures which correspond to problems related to M-SIDH, FESTA, CSIDH, proofs of isogeny knowledge [37,34], etc., which had not previously been known to be connected;
- Give our main technical contribution: a polynomial-time reduction between different level structures (Corollary 10);
- As a special case, show an attack against "SIDH with only one point" (Corollary 12), which has already been weaponized in [16,21];
- As another special case, prove that breaking SIDH-based proofs of isogeny knowledge [34] is at least as hard as breaking CSIDH (Corollary 13 and Lemma 14);
- Improve upon the best generic algorithm [36] to compute isogenies between ordinary curves (Section 5.6).

Limitations. We stress that it is rare that the security of a cryptographic scheme reduces to a Γ-SIDH problem as stated here. For example, the security of key exchange schemes typically depends on DDH- or CDH-like assumptions which are usually stronger than the corresponding Γ-SIDH one. In the interest of conciseness, we also avoid discussing decisional variants of Γ-SIDH, with the only exception of Section 5.5.

Some high profile schemes that fit quite badly in our framework are SQIsign and its variants [38,23,31], whose security reduces to a distinguishing problem on isogeny walks generated according to an *ad-hoc* distribution. Pre-quantum

schemes such as verifiable delay functions [39] and delay encryption [13] are also out of the scope of this work.

Despite all this, it is often the case that the best known attack against any isogeny-based scheme consists in solving an instance of a Γ-SIDH problem, thus our classification is especially valuable for cryptanalysts. In what follows whenever we mention "breaking a cryptosystem", we mean finding its secrets, rather than just breaking the assumption.

Ultimately, we hope that our framework will help better systematize and assess the landscape of assumptions in isogeny-based cryptography.

Outline. We formally define level structures in Section 2. In Section 3 we define the isogeny problem with level structure, and discuss how the inputs and the outputs of the problem are represented. We give our main technical contribution in Section 4: a polynomial-time reduction between isogeny problems with different level structures. We conclude in Section 5 with a review of isogeny problems with level structure that have previously appeared in the literature, and spell out the consequences of our reduction.

Notation. We will work with several groups of 2×2 matrices. We will use asterisks $*$ to denote entries of matrices that can take arbitrary values, and leave zero entries blank, thus $\left(\begin{smallmatrix} * & *_\ell \\ & * \end{smallmatrix}\right)$ represents any upper-triangular matrix whose upper right coefficient is divisible by ℓ. We will use the same notation to mean the group of all matrices having a certain form.

GL_2 denotes the group of invertible 2×2 matrices and SL_2 its subgroup of determinant-1 matrices. We write $\Gamma \leq \Delta$ to indicate that Γ is a subgroup of Δ.

Lower-case Greek letters ϕ, ψ, χ will be used to denote isogenies. To reduce clutter, we will write ϕP instead of $\phi(P)$ for the value of ϕ at a point P.

Throughout the document, p is a prime, q a power of p, N an integer coprime to p, and d coprime to N. We write $\tilde{O}(x)$ as a shorthand for $O(x \operatorname{polylog}(x))$.

2 Level structures

In this section we consider an elliptic curve E defined over some finite field $\mathbb{F}_q$. In the cases of cryptographic interest which motivated this work, E is supersingular and the finite field is either a prime field $\mathbb{F}_p$ or a quadratic extension $\mathbb{F}_{p^2}$, but the main results of this work apply in general.

Let N be a positive integer coprime to q. The torsion subgroup $E[N]$, i.e., the group of points of order dividing N, taken over the algebraic closure of $\mathbb{F}_q$, is isomorphic to $(\mathbb{Z}/N\mathbb{Z})^2$. We call a *basis* of $E[N]$ any ordered pair of points (P, Q) that generate $E[N]$. Denote by $\mathcal{B}_E(N)$ the set of all bases of $E[N]$.

The group $\mathrm{GL}_2(\mathbb{Z}/N\mathbb{Z})$ of 2×2 invertible matrices with coefficients modulo N acts on $\mathcal{B}_E(N)$ on the left by

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \cdot (P, Q) = (aP + bQ, cP + dQ).$$

Consider a subgroup $\Gamma \leq \mathrm{GL}_2$ (from now on all matrix groups will implicitly have coefficients in $\mathbb{Z}/N\mathbb{Z}$). If (P, Q) is a basis, we write $\Gamma \cdot (P, Q)$ for its *Γ-orbit*, that is the set

$$\{\gamma \cdot (P, Q) \mid \gamma \in \Gamma\}.$$

The *Γ-level structures (of level N)* on E are precisely the Γ-orbits of the bases of $E[N]$, forming a partition of $\mathcal{B}_E(N)$. We write $\mathcal{B}_E(\Gamma)$ for the set of Γ-level structures on E.

Said otherwise, a Γ-level structure is a basis of $E[N]$, up to transformation by elements of $\Gamma \leq \mathrm{GL}_2$. So, for example, when $\Gamma = \{\left(\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}\right)\}$ the Γ-level structures are just the bases of $E[N]$.

A pair (P, Q) of points in $E[N]$ forms a basis of $E[N]$ if and only if the Weil pairing $e_N(P, Q)$ is a primitive N-th root of unity. If $\gamma \in \mathrm{GL}_2$, by the alternating property of the Weil pairing,

$$e_N\big(\gamma \cdot (P, Q)\big) = e_N(P, Q)^{\det \gamma}.$$

Hence, the action of SL_2 partitions $\mathcal{B}_E(N)$ into $\varphi(N)$ orbits, each corresponding to one value of the pairing. We will chiefly be interested in subgroups $\Gamma \leq \mathrm{SL}_2$, so that it makes sense to talk about the value of the Weil pairing on a Γ-level structure. If $S \in \mathcal{B}_E(\Gamma)$ is one such level structure, we write $e_N(S)$ for the value of the Weil pairing.

3 Modular isogeny problems

Let E, E' be isogenous elliptic curves over $\mathbb{F}_q$, so that $\#E(\mathbb{F}_q) = \#E'(\mathbb{F}_q)$. From now on we let $\phi : E \to E'$ be an isogeny of degree d defined over $\mathbb{F}_q$.

Suppose $\gcd(d, N) = 1$; then ϕ defines a bijection between $\mathcal{B}_E(N)$ and $\mathcal{B}_{E'}(N)$. Let $\Gamma \leq \mathrm{GL}_2$ and let $S \in \mathcal{B}_E(\Gamma)$ be a level structure. The image $\phi(S)$ of S under ϕ is a Γ-level structure on E'.

We can now define generalizations of the classic SIDH problem.

Definition 1 (Γ-SIDH). *Fix coprime integers d and N and a subgroup $\Gamma \leq \mathrm{GL}_2(\mathbb{Z}/N\mathbb{Z})$.*

Let E, E' be elliptic curves defined over $\mathbb{F}_q$ such that there exists an $\mathbb{F}_q$-rational isogeny $\phi : E \to E'$ of degree d. Assuming $\gcd(N, q) = 1$, let $S \in \mathcal{B}_E(\Gamma)$ be a Γ-level structure.

The (d, Γ)-modular isogeny problem (of level N) asks, given $(E, S, E', \phi(S))$ to compute ϕ. When d is clear from context, we call this the Γ-SIDH problem.

Although the S in SIDH stands for "supersingular", we consider these problems for ordinary and supersingular elliptic curves alike. For groups Γ of special interest, we will also give other names.

Remark 2. When d and N have common factors, the image $\phi(S)$ is not well defined, however the problem of computing ϕ given some information on how it behaves on $E[N]$ is still meaningful. This problem, though, is usually much

easier to solve: If ℓ is a prime dividing d, either $\phi(E[\ell])$ is trivial, in which case ϕ factors as $\phi' \circ [\ell]$ with $\deg \phi' = d/\ell^2$, or $\phi(E[\ell])$ is the kernel of an ℓ-isogeny $\psi : E' \to E''$ such that $\phi = \widehat{\psi} \circ \phi'$ with $\deg \phi' = d/\ell$. By repeatedly removing common factors of d and N in this way, we reduce to a Γ-SIDH problem.

Remark 3. The isogeny ϕ may not be unique. In this case we just ask for any isogeny that satisfies the statement.

Alternatively, one may be given a tuple (E, S, E', S') and be asked whether there exists a d-isogeny such that $E' = \phi(E)$ and $S' = \phi(S)$. Variations on this decisional problem occur in cryptography, and we will point to them where relevant.

We were purposefully vague on the data structures involved in the definition of Γ-SIDH. Indeed the meaning of "given S" and "compute ϕ" may vary depending on context. It shall be understood that elliptic curves are represented by some projective model (e.g., a Weierstrass equation), and points by their coordinates in the model. We shall assume that the factorization of N is known, and that the modular groups Γ have some simple description, e.g., through a set of generators, or implicitly such as in "Γ_0, the subgroup of triangular matrices of determinant 1". The representation of level structures and isogenies is subtler.

Representing isogenies. "Computing an isogeny" is usually understood as "computing a generator of its kernel". From this data, we can use Vélu's formulas to evaluate the isogeny at any point. However, in some cases of interest the points of the kernel may be defined over prohibitively large field extensions, or the isogeny may have large prime degree preventing the use of Vélu's formulas.

Instead, following [67,55], we will say that an algorithm (efficiently) *represents* an isogeny $\phi : E \to E'$ if, given any point $P \in E(\mathbb{F}_{q^k})$ as input, it outputs $\phi(P)$ in time $\mathrm{poly}(k \log(q))$. The goal of the Γ-SIDH problem will thus be to output a representation of ϕ, for instance as an arithmetic circuit.

Representing level structures. A level structure may be exponentially large, so representing it by the list of its bases is out of question. As a first attempt, we may represent $S \in \mathcal{B}_E(\Gamma)$ through an arbitrary basis in S and the group Γ, but even this representation may turn out to be prohibitively expensive. We now illustrate different ways of representing level structures through two examples.

Example 4. For supersingular curves, one never needs to look far to complete a basis. Indeed, if $P \in E(\mathbb{F}_q)$ is a point of order N, then a point Q such that (P, Q) forms a basis of $E[N]$ is always defined over an extension of $\mathbb{F}_q$ of degree at most 6 (or even 2, when $j(E) \neq 0, 1728$). It is thus possible to represent a level structure by an arbitrary basis (P, Q) and the group Γ at little extra cost.

Familiar examples are SIDH [50] and B-SIDH [28] public keys: these are triples (E, P, Q) where E is a supersingular curve over $\mathbb{F}_{p^2}$, and (P, Q) is a basis of $E[N]$, defined over $\mathbb{F}_{p^2}$ in SIDH's case, or over $\mathbb{F}_{p^4}$ in B-SIDH's case. In both cases Γ is the trivial group.

When N factors into primes as $N = \ell_1^{e_1} \cdots \ell_r^{e_r}$, it may be more advantageous to use the decomposition

$$E[N] = E[\ell_1^{e_1}] \oplus \cdots \oplus E[\ell_r^{e_r}]$$

and use a pair (P_i, Q_i) of generators of $E[\ell_i^{e_i}]$ for each factor. Indeed the fields of definition of each $E[\ell_i^{e_i}]$ will tend to be much smaller than their compositum, the field of definition of $E[N]$. The group Γ will then act on $E[\ell_i^{e_i}]$ in a similar way as $\Gamma \bmod \ell_i^{e_i}$.

Example 5. Ordinary curves behave differently: even when a point of order N is defined over some field $\mathbb{F}_q$, a full basis of $E[N]$ is, in general, only defined over an extension of degree $O(N)$. However some level structures may be appropriately described by a single point of order N. For example, given a basis (P, Q), its orbit under the group of upper-triangular matrices consists of all the bases $(aP + bQ, cQ)$, i.e., all the bases whose second generator lies in $\langle Q \rangle$. Such a level structure may thus simply be encoded by an arbitrary generator of $\langle Q \rangle$, or even by some implicit definition, such as "the unique subgroup of order N defined over the extension of degree n of $\mathbb{F}_q$".

The case where N is a power of a prime is well-known to be related to the theory of *isogeny volcanoes* [52,44]. These are graphs with elliptic curves for nodes and ℓ-isogenies for edges. The nodes are arranged into levels, corresponding to the endomorphism rings of the curves: At the top level, the *crater*, lie the curves with the largest endomorphism ring; at the bottom, the *floor* lie the curves with smallest endomorphism ring; from one level to the next the endomorphism ring grows or shrinks by a factor ℓ. ℓ-isogenies between curves on the same level are only possible on the crater, at every other level ℓ-isogenies can only go up or down one level. On the other hand, isogenies of degree d coprime to ℓ are only possible between curves on the same levels of the respective ℓ-isogeny volcanoes.

Miret, Sadornil, Tena, Tomàs, Rosana, and Valls [57] show that the structure of the ℓ-*Sylow* of $E(\mathbb{F}_q)$, i.e. of the largest subgroup of $E(\mathbb{F}_q)$ of order a power of ℓ, is controlled by the level in the volcano. At the floor of the volcano, the ℓ-Sylow is cyclic and is thus naturally identified with a $\left(\begin{smallmatrix} * & * \\ & * \end{smallmatrix}\right)$-level structure of level, say, ℓ^e; however a basis of the full $E[\ell^e]$ will only be defined over an extension of degree $O(\ell^e)$. At the other end, on the crater, the ℓ-Sylow may be isomorphic to $(\mathbb{Z}/\ell^{e/2}\mathbb{Z})^2$ and thus all bases of $E[\ell^{e/2}]$ would be defined, but we would have trouble representing any meaningful level structure of level $\ell^e \gg \ell^{e/2}$.

Because d is coprime to ℓ, the isogeny ϕ must be between curves on the same level of their respective ℓ-isogeny volcanoes, and thus preserve the structure of the ℓ-Sylows. Therefore, whichever representation of Γ-level structures works for E also works for E'.

As the examples show, there is not a single "good way" of representing level structures. In what follows we will enunciate algorithms assuming all points of $E[N]$ are always defined over the base field, for coherence with the most cryptographically relevant cases. The reader is left with the task of adjusting the algorithms to other representations where appropriate.

Restricting to $\mathbf{SL_2}$. In most cases it makes sense to restrict to Γ-SIDH problems for $\Gamma \leq \mathrm{SL}_2$. Indeed, provided we can solve discrete logarithms in the subgroup of N-th roots of unity of $\overline{\mathbb{F}}_q$, we can reduce any Γ-SIDH problem to a $(\Gamma \cap \mathrm{SL}_2)$-SIDH as follows.

Lemma 6. *Let $\Gamma \leq \mathrm{GL}_2(\mathbb{Z}/N\mathbb{Z})$ and denote by $\mu_N \subset \mathbb{F}_{q^r}^\times$ the subgroup of N-th roots of unity. Given an oracle to solve discrete logarithms in μ_N, there exists a reduction from Γ-SIDH to $(\Gamma \cap \mathrm{SL}_2)$-SIDH with complexity polynomial in r, $\log(q)$ and $\log(N)$.*

Proof. Let (E, S, E', S') be a Γ-SIDH problem. If $E(\mathbb{F}_q)$ has order divisible by N, then by assumption $E[N] \subset E(\mathbb{F}_{q^{2r}})$. We can thus choose a representative (P, Q) of S and compute its order-N Weil pairing $\zeta := e_N(P, Q) \in \mathbb{F}_{q^r}$. Let $\overline{S} = (\Gamma \cap \mathrm{SL}_2) \cdot (P, Q)$; in other words, $\overline{S}$ is the set of bases obtained by acting on the basis (P, Q) with matrices in Γ having determinant 1. Thanks to the compatibility of the Weil pairing with isogenies,

$$e_N(\phi(\overline{S})) = e_N(\overline{S})^{\deg \phi} = \zeta^d.$$

Choose now a representative (P', Q') of $\phi(S)$ and compute its Weil pairing $\xi := e_N(P', Q')$. Use the oracle to compute the discrete logarithm x of ζ^d to base ξ and find a matrix $\gamma \in \Gamma$ with $\det \gamma = x$. Define $\overline{S}' = (\Gamma \cap \mathrm{SL}_2) \cdot \gamma \cdot (P', Q')$; then $\overline{S}' = \phi(\overline{S})$. Hence, $(E, \overline{S}, E', \overline{S}')$ is an instance of $(\Gamma \cap \mathrm{SL}_2)$-SIDH and has the same solutions as the original problem. □

In the next sections we shall focus our attention on groups $\Gamma \leq \mathrm{SL}_2$, which happen to be the most common in cryptography.

4 A reduction

We come to the main technical result of this work: a reduction between the Γ-SIDH problems for different modular groups Γ.

For ℓ an integer dividing N, define the subgroup

$$\Gamma_0(\ell) = \left\{ \left(\begin{smallmatrix} a & b \\ c\cdot\ell & d \end{smallmatrix} \right) \right\} \leq \mathrm{SL}_2(\mathbb{Z}/N\mathbb{Z}).$$

For $\Gamma \leq \Gamma_0(\ell)$, define its *$\ell$-conjugate* as

$$\Gamma^* = \left\{ \left(\begin{smallmatrix} a & b\cdot\ell \\ c' & d \end{smallmatrix} \right) \;\middle|\; \left(\begin{smallmatrix} a & b \\ c\cdot\ell & d \end{smallmatrix} \right) \in \Gamma \text{ and } c' \equiv c \pmod{N/\ell} \right\};$$

it is easily verified that Γ^* is a subgroup of SL_2. The main subroutine of the reduction is Algorithm 1, which constructs from a curve E with Γ-level structure an ℓ-isogenous curve E' with Γ^*-level structure.

The correctness of Algorithm 1 follows from the following lemma:

Lemma 7. *Let N, ℓ be positive integers with $\ell \mid N$. Let E be an elliptic curve and (P, Q) a basis of $E[N]$. Let $\psi : E \to E'$ be an ℓ-isogeny with kernel $\langle (N/\ell) \cdot Q \rangle$ and let $\widehat{\psi}$ be its dual. A point $Q' \in E'[N]$ satisfies $\widehat{\psi} Q' = Q$ if and only if $e_N(\psi P, Q') = e_N(P, Q)$ and $\ell Q' = \psi Q$. In that case, the pair $(\psi P, Q')$ forms a basis of $E'[N]$.*

Algorithm 1: Changing from Γ-level structure to Γ^*-level structure.

Input: Integers N and $\ell \mid N$, a subgroup $\Gamma \leq \Gamma_0(\ell) \leq \mathrm{SL}_2(\mathbb{Z}/N\mathbb{Z})$, an elliptic curve E with Γ-level structure S.

Output: An elliptic curve E' with Γ^*-level structure S' and an ℓ-isogeny $\psi\colon E \to E'$ such that $\widehat{\psi}Q' \in \langle Q \rangle$ for all $(\_, Q) \in S$ and $(\_, Q') \in S'$.

1 Pick an arbitrary basis $(P, Q) \in S$;
2 Let $K := (N/\ell) \cdot Q$;
3 Compute an isogeny $\psi : E \to E'$ with kernel $\langle K \rangle$;
4 Let $P' := \psi P$ and $Q'' := \psi Q$;
5 Compute Q' such that $\ell Q' = Q'' = \psi Q$ and $e_N(P', Q') = e_N(P, Q)$;
6 Let $S' := \Gamma^* \cdot (P', Q')$ and **return** $\big(E', S', \psi\big)$.

Proof. Let Q' be such that $\widehat{\psi}Q' = Q$, then, by the properties of the Weil pairing

$$e_N(\psi P, Q') = e_N(P, \widehat{\psi}Q') = e_N(P, Q). \tag{1}$$

Moreover $\psi Q = \psi\widehat{\psi}Q' = \ell Q'$.

Conversely, let Q' be a point satisfying Eq. (1) and such that $\ell Q' = \psi Q$. By the non-degeneracy of the Weil pairing we have $\widehat{\psi}Q' \in Q + \langle P \rangle$. Writing $\widehat{\psi}Q' = Q + xP$ for some $x \in \mathbb{Z}/N\mathbb{Z}$, we have

$$\ell Q' = \psi\widehat{\psi}Q' = \psi Q + x\psi P = \ell Q' + x\psi P,$$

hence $x = 0$ because ψP has order N.

In either case, because the Weil pairing from Eq. (1) has maximal order, $(\psi P, Q')$ must be a basis of $E'[N]$. □

We are now ready to present a reduction from Γ-SIDH to Γ^*-SIDH. The idea underlying the algorithm is visualized in the diagram below, where the vertical arrows represent the ℓ-isogenies constructed by Algorithm 1: The key point is that all matrices in Γ are upper-triangular modulo ℓ, hence by construction the isogenies ψ, χ are parallel with respect to ϕ, i.e., $\ker \chi = \phi(\ker \psi)$, which implies the existence of ϕ'.

$$\begin{array}{ccc} E_0, \big(P_0, Q_0\big) & \xdashrightarrow{\quad\phi\quad} & E_1, \big(P_1, Q_1\big) \\ \psi \downarrow \quad \uparrow \widehat{\psi} & & \widehat{\chi} \uparrow \quad \downarrow \chi \\ \frac{E_0}{\langle N/\ell \cdot Q_0 \rangle}, \big(\psi P_0, \widehat{\psi}^{-1} Q_0\big) & \dashrightarrow^{\phi'} & \frac{E_1}{\langle N/\ell \cdot Q_1 \rangle}, \big(\chi P_1, \widehat{\chi}^{-1} Q_1\big) \end{array}$$

Theorem 8. *Let $\phi : E \to E'$ be an isogeny of degree d coprime to N, and let N be coprime to the characteristic. Let ℓ divide N and let $\Gamma \leq \Gamma_0(\ell)$. Given an instance of a Γ-SIDH problem with solution ϕ, Algorithm 2 outputs an instance of a Γ^*-SIDH problem with solution $\phi' = \chi\phi\widehat{\psi}/\ell$.*

Algorithm 2: Γ-SIDH to Γ^*-SIDH reduction.

Input: Integers N and $\ell \mid N$, a subgroup $\Gamma \leq \Gamma_0(\ell) \leq \mathrm{SL}_2(\mathbb{Z}/N\mathbb{Z})$, a Γ-SIDH instance $(E_0, S, E_1, \phi(S))$.

Output: A Γ^*-SIDH instance $(E_0', S', E_1', \phi'(S'))$ together with ℓ-isogenies $\psi\colon E_0 \to E_0'$, $\chi\colon E_1 \to E_1'$ such that $\chi \circ \phi = \phi' \circ \psi$.

1 Run Algorithm 1 with input (E_0, S) and let (E_0', S', ψ) be the result;
2 Run Algorithm 1 with input $(E_1, \phi(S))$ and let $(E_1', \phi'(S'), \chi)$ be the result;
3 **Return** $\big(E_0', S', E_1', \phi'(S')\big)$ *and* (ψ, χ).

Proof. Fix $(P_0, Q_0) \in S$, $(P_1, Q_1) \in \phi(S)$, $(P_0', Q_0') \in S'$, and $(P_1', Q_1') \in \phi'(S')$.

By definition of Γ-SIDH, there exists a matrix $\gamma = \left(\begin{smallmatrix} a & b \\ c\cdot\ell & d \end{smallmatrix}\right) \in \Gamma$ such that $\gamma \cdot (P_1, Q_1) = (\phi P_0, \phi Q_0)$. Thus we have $d(N/\ell) \cdot Q_1 = (N/\ell) \cdot \phi Q_0$, implying that $\ker \chi = \phi(\ker \psi)$ and therefore $\chi\phi\widehat{\psi}(E_0'[\ell]) = \chi\phi(\ker\psi) = \chi(\ker\chi) = \{0\}$. This proves that $\phi' := \chi\phi\widehat{\psi}/\ell$ is a well-defined isogeny.

Lemma 7 shows that (P_0', Q_0') and (P_1', Q_1') form bases of the respective N-torsion subgroups, and that $\widehat{\psi}Q_0' = Q_0$ and $\widehat{\chi}Q_1' = Q_1$. Using the properties of the dual isogeny, we readily see that $\widehat{\psi}P_0' = \ell P_0$ and $\widehat{\chi}P_1' = \ell P_1$.

To each of the isogenies in the diagram, we associate a matrix representing its action (on the left) on the N-torsion, with respect to the bases (P_0, Q_0), (P_1, Q_1), (P_0', Q_0') and (P_1', Q_1'). Thus, ϕ acts like γ^t, ψ and χ act like $\left(\begin{smallmatrix} 1 & \\ & \ell \end{smallmatrix}\right)$, and $\widehat{\psi}$ and $\widehat{\chi}$ act like $\left(\begin{smallmatrix} \ell & \\ & 1 \end{smallmatrix}\right)$. Writing $\left(\begin{smallmatrix} x & y \\ w & z \end{smallmatrix}\right)^t$ for the matrix of ϕ', we obtain the relations

$$\left(\begin{smallmatrix} 1 & \\ & \ell \end{smallmatrix}\right)\left(\begin{smallmatrix} a & c\cdot\ell \\ b & d \end{smallmatrix}\right) \equiv \left(\begin{smallmatrix} x & w \\ y & z \end{smallmatrix}\right)\left(\begin{smallmatrix} 1 & \\ & \ell \end{smallmatrix}\right) \pmod N; \qquad [\text{from } \chi\circ\phi = \phi'\circ\psi]$$

$$\left(\begin{smallmatrix} a & c\cdot\ell \\ b & d \end{smallmatrix}\right)\left(\begin{smallmatrix} \ell & \\ & 1 \end{smallmatrix}\right) \equiv \left(\begin{smallmatrix} \ell & \\ & 1 \end{smallmatrix}\right)\left(\begin{smallmatrix} x & w \\ y & z \end{smallmatrix}\right) \pmod N. \qquad [\text{from } \phi\circ\widehat{\psi} = \widehat{\chi}\circ\phi']$$

Whence

$$x \equiv a, \quad y \equiv b\ell, \quad w\ell \equiv c\ell, \quad z \equiv d \pmod N,$$

thus ϕ' acts like $\left(\begin{smallmatrix} a & b\cdot\ell \\ c' & d \end{smallmatrix}\right)^t$ for some $c' \equiv c \pmod{N/\ell}$.

Applying the same reasoning to all matrices $\gamma \in \Gamma$, we conclude that

$$\big(E_0', (P_0', Q_0'), E_1', (P_1', Q_1')\big)$$

is an instance of a Γ^*-SIDH problem with solution ϕ'. □

Example 9. Going back to Example 5, Algorithm 2 can be understood as moving up and down the volcano. For example, suppose one is given curves E_0, E_1 on the floor of their respective ℓ-volcanoes and seeks to compute an isogeny $\phi : E_0 \to E_1$. The ℓ-Sylows of E_0 and E_1 are cyclic, say of order ℓ^e, and are mapped one onto the other. Thus ϕ is solution to an instance of a Γ_0-SIDH problem of level ℓ^e.

There are unique rational ℓ-isogenies $\psi : E_0 \to E_0'$ and $\chi : E_1 \to E_1'$, both ascending, and their target curves E_0', E_1' have ℓ-Sylows isomorphic to $\mathbb{Z}/\ell\mathbb{Z} \times \mathbb{Z}/\ell^{e-1}\mathbb{Z}$. Algorithm 2 reduces the Γ_0-problem between E_0 and E_1 to a $\left(\begin{smallmatrix} * & *\ell \\ *\ell^{e-1} & * \end{smallmatrix}\right)$-problem between E_0' and E_1', matching the structure of the ℓ-Sylows.

To finish the reduction, we need to show that from a solution to the Γ^*-SIDH problem we can efficiently construct a solution to the original Γ-SIDH problem. This task essentially consists in evaluating the fractional isogeny $\phi = \widehat{\chi}\phi'\psi/\ell$; see the proof below for details.

A small difficulty arises when analyzing the complexity of Algorithm 1: Indeed during the computation of $\widehat{\psi}^{-1}Q$ in Step 5, it is possible that Q' is only defined over an extension field of $\mathbb{F}_q$. This cannot happen if we assume that $E[N]$ is defined over $\mathbb{F}_q$ for all curves in the isogeny class, as is the case for supersingular curves. We will only analyze the cost in this case, but note that the ordinary case can be handled by simply using Q'' as a representation for $\widehat{\psi}^{-1}Q$.

Corollary 10. *Γ-SIDH reduces to Γ^*-SIDH, with a polynomial overhead in ℓ, $\log(N)$ and $\log(q)$.*

Proof. Given an instance of Γ-SIDH, we apply Algorithm 2 and pass the result to a Γ^*-SIDH oracle, obtaining a representation of $\phi' : E'_0 \to E'_1$. We use ϕ' to build a representation of ϕ as follows. Let P be a point of E_0 of which we want to know the image ϕP. Compute R such that $\ell R = P$, then $\phi P = \ell\phi R$. Now observe that $\ell\phi = \widehat{\chi}\phi'\psi$, hence $\phi P = \widehat{\chi}\phi'\psi R$.

We now analyze the cost of Algorithm 1 in terms of field operations. The choice of the basis (P, Q) is assumed to be free. The scalar multiplication in Step 2 costs $O(\log(N))$. Evaluating ψ costs $O(\ell)$ operations using Vélu's formulas. We can compute Q' by finding a preimage to Q'', using a root finding algorithm such as Cantor–Zassenhaus [14], which costs $\tilde{O}(\ell \log(q))$.

The cost of Algorithm 2 is that of running Algorithm 1 twice. Finally the additional cost of evaluating ϕP amounts to a division by ℓ and a few ℓ-isogeny evaluations, thus the same cost as Algorithm 2. □

For an example of how the reduction in this section can be applied concretely, see Section 5.4 below.

5 Γ-SIDH problems in the wild

We finally review the occurrences of Γ-SIDH problems in the literature, and use the reduction of Section 4 to reveal some new connections. As previously stated, we only consider subgroups $\Gamma \leq \mathrm{SL}_2$. Figure 1 gives an overview of this section: The table on the left lists groups Γ, the best known attack against the generic Γ-SIDH problem, and some schemes whose security is "based" on it; the diagram on the right shows reductions between the problems.

5.1 The generic isogeny problem

When $\Gamma = \mathrm{SL}_2$, a level structure is just the set of all bases of $E[N]$ with a given value of the Weil pairing. Assuming we can solve discrete logarithms in μ_N, an arbitrary such basis can be computed from E using the technique described in

Γ	Best attack	Schemes
$\left(\begin{smallmatrix}1 & \\ & 1\end{smallmatrix}\right)$	poly	SIDH
$\left(\begin{smallmatrix}1 & * \\ & 1\end{smallmatrix}\right)$	poly	[16,21]
$\left(\begin{smallmatrix}\lambda & \\ & \lambda\end{smallmatrix}\right)$	exp	M-SIDH
$\left(\begin{smallmatrix}* & \\ & *\end{smallmatrix}\right)$	exp	FESTA, binSIDH, CSIDH, SCALLOP
$\left(\begin{smallmatrix}* & * \\ & *\end{smallmatrix}\right)$	exp	SIDH PoKs
SL_2	exp	generic

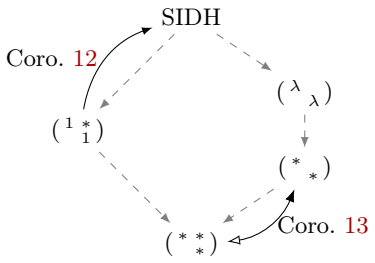

Fig. 1. Γ-SIDH problems, their difficulty and reductions. The *Schemes* column lists schemes that can be broken by solving the corresponding Γ-SIDH problem. Note these are not security reductions: some schemes may have better attacks (e.g., quantum subexponential time against CSIDH and SCALLOP). The reduction diagram uses dashed arrows to signify trivial inclusions and continuous arrows to represent reductions following from Theorem 10: The latter are between a Γ-SIDH problem of level n^2 and a Γ'-SIDH problem of level n; the filled arrow tip points towards the problem of level n.

Lemma 6. Thus SL_2-SIDH is simply the generic fixed-degree isogeny problem: given isogenous curves E, E', find $\phi : E \to E'$ of degree d.

The best generic algorithms to solve this problem have complexity polynomial in d, and are not much more advanced than plain exhaustive search. When $d = d_1 d_2$, a meet-in the middle approach [1,29] improves slightly over exhaustive search by intersecting the list of all curves d_1-isogenous to E with the list of all curves d_2-isogenous to E'.

When E and E' are supersingular, it is known that the generic isogeny problem is equivalent to the endomorphism ring problem (computing a quaternion representation of $\mathrm{End}(E)$ and $\mathrm{End}(E')$) [41,67][2]. These are considered to be the most fundamental problems in isogeny-based cryptography, and a solution to them would compromise almost all known schemes. The best algorithms known take $\tilde{O}(\sqrt{p})$ classical time [40,42,45] or $\tilde{O}(\sqrt[4]{p})$ quantum time [11]. The generic fixed-degree isogeny problem is only known to be equivalent to the endomorphism ring problem when d is at least $O(p^3)$ [5,9].

In the ordinary case, the theory of isogeny volcanoes applies and yields an improvement over generic algorithms when d is large enough. The first step is to ascend to the craters of the ℓ-isogeny volcanoes for each ℓ dividing the conductors of $\mathrm{End}(E)$ and $\mathrm{End}(E')$, which has cost polynomial in the largest such ℓ. The second step is a collision-search algorithm in the isogeny class of the maximal order, taking $O(\sqrt{C})$ classical operations [46], or a hidden shift algorithm taking $O(\exp(\sqrt{\log(C)}))$ quantum operations [53,61,54], where C is the size of the isogeny class.[3] The final step evaluates the actual isogeny of degree d and can be done in $O(\exp(\sqrt{\log(C)}))$ classical and quantum time [51,25]. The

[2] When d contains a large prime factor, the reduction is only quantum [24].

[3] The isogeny class of a random curve over $\mathbb{F}_q$ has size $O(\sqrt{q})$, however some curves (e.g., pairing-friendly curves) may be specially constructed with a small isogeny class.

same theory also applies to supersingular curves defined over a prime field as in CSIDH [18], or more generally to *oriented* supersingular curves [35]. We shall come back to these cases when discussing $\left(\begin{smallmatrix} * & \\ & * \end{smallmatrix}\right)$-SIDH in Section 5.6.

5.2 The SIDH problem

Appearance: SIDH key exchange and derivatives [50,49], B-SIDH [28], Séta [32].
Best attacks: polynomial time when N is smooth [15,56,63].

On the opposite end of the spectrum we have $\Gamma = \{\left(\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}\right)\}$, which is none else than (a minimal generalization of) the well-known SIDH problem: we are given d-isogenous curves E, E', points P, Q generating $E[N]$ and their images, and we want to find the isogeny.

The standard SIDH/SIKE setting [50,49] has $d = 2^n \approx 3^m = N$ and the (supersingular) curves are chosen so that $E[dN] \subset E(\mathbb{F}_{p^2})$. In some variants [28,32] d and N are coprime smooth integers, and $E[dN] \subset E(\mathbb{F}_{p^4})$. The attacks on SIDH [15,56,63] show that in all these instances the SIDH problem can be solved in polynomial time. More generally, Robert proves the following theorem.

Theorem 11. *Let E be an elliptic curve over a finite field $\mathbb{F}_q$. Let $N \in \mathrm{poly}(q)$ be a $\mathrm{polylog}(q)$-smooth integer, let $S \in \mathcal{B}_E(N)$, and suppose that S can be represented over an extension of $\mathbb{F}_q$ of size $\mathrm{polylog}(q)$.*

If $\phi : E \to E'$ is an isogeny of degree $d < N^2$, the instance $(E, S, E', \phi(S))$ of the SIDH problem can be solved in time $\mathrm{polylog}(q)$.

In more detail, in [62, Theorem 1.2] Robert provides an algorithm that, given $(E, S, E', \phi(S))$ and a point $R \in E(\mathbb{F}_{q^k})$, outputs ϕR. The only minor difference with the statement above is a stronger condition $d < N$, however this can be relaxed to $d < N^2$ using the technique described in [63, § 6.4].

Note that, when d is smooth too, one can recover a more traditional representation of ϕ as an isogeny walk by evaluating ϕ on $E[d]$.

In conclusion, the only instances where the (generalized) SIDH problem still appears to be hard are those where N contains a large prime factor, or where points of $E[N]$ cannot be represented over small extensions of $\mathbb{F}_q$.

5.3 $\left(\begin{smallmatrix} \lambda & \\ & \lambda \end{smallmatrix}\right)$-SIDH a.k.a. M-SIDH

Appearance: Masked-torsion SIDH [43].
Reduces to: SIDH, when N has few distinct prime factors.
Best attacks: exponential time [43], polynomial time for "special" supersingular curves [43] and for curves over $\mathbb{F}_p$ [21].

Acting by the group $\left\{\left(\begin{smallmatrix} \lambda & \\ & \lambda \end{smallmatrix}\right)\right\}$ of determinant-1 scalar matrices has been proposed as a countermeasure against the SIDH attacks. The key exchange scheme M-SIDH [43] (short for Masked-torsion SIDH) works exactly like SIDH, however,

before publishing the images of the basis (P, Q) of $E[N]$, it masks them by a random scalar λ. That is, an M-SIDH public key is a triple $(E', \lambda\phi P, \lambda\phi Q)$.

By definition $\lambda^2 = 1$, so a straightforward way to solve M-SIDH is to try and guess λ, thus reducing to an SIDH problem. To protect against this, M-SIDH chooses N to have r distinct prime factors, so that there are 2^r possible solutions for λ, rendering the attack infeasible.

In [43] it is shown how to solve M-SIDH when the starting curve E is supersingular and has small endomorphisms. In [21] it is shown how to solve it when E is defined over a prime field $\mathbb{F}_p$ or when E is connected to its Galois conjugate by a small-degree isogeny. In the general case, the best known attacks consist in guessing $\lambda \bmod N'$ where N' is a divisor of N with the minimal number of prime factors such that $d \leq N'^2$. As before, the correct guess for $\lambda \bmod N'$ allows a reduction to an SIDH problem. The degree d and the integer N are chosen such that any such N' has r distinct prime factors, where r depends on the desired security level ($r = 128, 192, 256$ for example), so that the complexity of this attack is at least 2^r.

5.4 $\left(\begin{smallmatrix}1 & *\\ & 1\end{smallmatrix}\right)$-SIDH a.k.a. Unipotent SIDH a.k.a. Γ_1-SIDH a.k.a. SIDH$_1$

Appearance: weak instances of class-group actions [16], flawed implementations of proofs of knowledge [34].
Reduces to: SIDH, when N contains a large smooth square factor.

The next group we consider is Γ_1, the group of unitriangular matrices $\left(\begin{smallmatrix}1 & *\\ & 1\end{smallmatrix}\right)$, a.k.a. the unipotent subgroup of SL_2. If $(P, Q) \in \mathcal{B}_E(N)$, its orbit $\Gamma_1 \cdot (P, Q)$ consists of all bases (R, Q) such that $e_N(R, Q) = e_N(P, Q)$, i.e., after fixing a value for the Weil pairing, Γ_1-level structures are in one-to-one correspondence with points of order N. Hence, Γ_1-SIDH is the variant of SIDH where, instead of the image of two generators of $E[N]$, only the image of a single point of order N is known. Because the notation of Γ_1 for the unipotent subgroup is standard in the theory of modular forms, we like to dub this the SIDH$_1$ problem.

Applying the reduction of Section 4, we prove a reduction from SIDH$_1$ to SIDH whenever N contains a large square smooth factor, e.g., when $N = \ell^e$ for some small prime ℓ.

Corollary 12. *Let n be an integer and ℓ its largest prime factor. The SIDH$_1$ problem of level n^2 reduces to the SIDH problem of level n in* poly(ℓ)*-time.*

Proof. Using Algorithm 2 repeatedly for each prime factor of n, we reduce SIDH$_1$ to Γ-SIDH with $\Gamma = \{\left(\begin{smallmatrix}1 & n*\\ n* & 1\end{smallmatrix}\right)\}$. But $\Gamma \bmod n$ is the trivial subgroup of $\mathrm{SL}_2(\mathbb{Z}/n\mathbb{Z})$, thus, restricting the level structure to the n-torsion (i.e., multiplying the basis generators by n), we obtain an instance of SIDH of level n. □

Even if N is not exactly a square, the corollary above gives a strategy to attack SIDH$_1$. Indeed, if $N = sn^2$ with a small squarefree factor s, we can simply ignore s and restrict to the SIDH$_1$ problem of level n^2, which reduces to

SIDH of level n. Should this not be sufficient, we can add a small exhaustive search on s to reduce to SIDH of level sn.

To the best of our knowledge, SIDH_1 has not appeared in the literature as a security assumption, however we are aware of at least two instances where it showed up somewhat unexpectedly.

The first instance is in variants of the Couveignes–Rostovtsev–Stolbunov group action [30,64] designed to be vulnerable precisely to Corollary 12 [16]. These isogeny classes are set up so that a large power $\ell^e > \sqrt{d}$ divides the discriminant. A self-pairing is then used to guess the image by ϕ of a point of order ℓ^e, leading to an SIDH_1 problem. Finally, our reduction is used to reduce to an SIDH problem that is solved using Robert's technique. Although this construction is artificial and not meant as a basis for cryptography, it shows that some isogeny-based group actions are less strong than originally thought.

The second instance is in proofs of knowledge of isogenies *à la SIDH* as seen in [37,34]. Both papers claim computational zero-knowledge based on the decisional variant of Γ_0-SIDH, however an implementation mistake makes them actually reliant on Γ_1-SIDH, and thus broken. We shall give more details on this in the next section.

5.5 $\left(\begin{smallmatrix} * & * \\ * & \end{smallmatrix}\right)$-SIDH a.k.a. Borel SIDH a.k.a. Γ_0-SIDH a.k.a. SIDH_0

Appearance: Proofs of Knowledge (decisional) [37,34,6].
Reduces to: $\left(\begin{smallmatrix} * & \\ * & \end{smallmatrix}\right)$-SIDH, when N contains a large smooth square factor.
Best attacks: generic.

Γ_0 is the group of determinant-1 upper-triangular matrices, the *Borel subgroup* of SL_2. The associated level structures correspond to cyclic subgroups of order N with a given value for the Weil pairing, and are sometimes called *Borel level structures* [6,26]. Again, we shorten Γ_0-SIDH into SIDH_0.

Using the reduction of Section 4, we prove that it is in fact equivalent to $\left(\begin{smallmatrix} * & \\ * & \end{smallmatrix}\right)$-SIDH. We are told this isomorphism between level structures is folklore among experts in modular curves (see [26, § 2.5]), but the algorithmic aspect appears to be new.

Corollary 13. *Let n be an integer and ℓ its largest prime factor. The SIDH_0 problem of level n^2 and the $\left(\begin{smallmatrix} * & \\ * & \end{smallmatrix}\right)$-SIDH problem of level n are $\mathrm{poly}(\ell)$-time-equivalent.*

Proof. Using Algorithm 2 repeatedly for each prime factor of n, we reduce SIDH_0 to Γ-SIDH with $\Gamma = \{\left(\begin{smallmatrix} * & n* \\ n* & * \end{smallmatrix}\right)\}$. Reducing modulo n we obtain a $\left(\begin{smallmatrix} * & \\ * & \end{smallmatrix}\right)$-SIDH problem of level n.

Conversely, let (E_0, S_0, E_1, S_1) be an instance of $\left(\begin{smallmatrix} * & \\ * & \end{smallmatrix}\right)$-SIDH. Pick arbitrary bases $(P_0, Q_0) \in S$ and $(P_1, Q_1) \in S_1$. Lift (P_0, Q_0) to a basis (P'_0, Q'_0) of $E_0[n^2]$ such that $nP'_0 = P_0$ and $nQ'_0 = Q_0$. Similarly lift (P_1, Q_1) to a basis (P'_1, Q'_1) of $E_1[n^2]$, with the additional constraint $e_{n^2}(P'_1, Q'_1) = e_{n^2}(P'_0, Q'_0)^d$, by solving a discrete logarithm for each factor ℓ of n, as in Lemma 6. Then

$(E_0, (P'_0, Q'_0), E_1, (P'_1, Q'_1))$ is an $\left(\begin{smallmatrix} * & n* \\ n* & * \end{smallmatrix}\right)$-SIDH instance of level n^2. Now apply Algorithm 2 to reduce to $\left(\begin{smallmatrix} * & * \\ * & * \end{smallmatrix}\right)$-SIDH and transpose to reduce to SIDH_0. □

SIDH_0 arises naturally in proofs of knowledge of isogenies *à la SIDH* [37,34]. To prove knowledge of an N-isogeny $\psi : E_0 \to E_1$ between supersingular curves, these schemes eventually produce an N-isogeny $\psi' : E_2 \to E_3$ such that $\ker \psi' = \phi(\ker \psi)$ for some secret isogeny ϕ of degree d. $\ker \psi = \langle Q \rangle$ and $\ker \psi' = \phi(\langle Q \rangle) = \langle Q' \rangle$ are cyclic groups of order N, thus, completing them to bases $(P, Q) = E_0[N]$ and $(P', Q') = E_1[N]$, we see that $\phi(P, Q) \in \Gamma_0 \cdot (P', Q')$. Hence, recovering ϕ from ψ and ψ' is naturally an SIDH_0 problem.

In fact, the zero-knowledge property of these schemes reduces to the decisional version of SIDH_0: the game is to distinguish $(E_0, \ker \psi, E_2, \phi(\ker \psi))$ from $(E_0, \ker \psi, E_2, G)$ where G is a random cyclic group of order N. This problem was named Decisional Supersingular Product (DSSP) in [37]. Recently, DSSP was proven undecidable when $\deg \phi \to \infty$, and even statistically undecidable as soon as $d \in O((pN)^c)$ for an explicit constant c depending on N, using the theory of *supersingular isogeny graphs with level structure* [6].

When implementing these schemes, it is a common mistake to encode $\ker \psi$ by some particular generator Q, and $\ker \psi'$ by $\phi(Q)$ rather than by some arbitrary generator of $\phi(\langle Q \rangle)$. However, in doing so the tuple $(E_0, Q, E_2, \phi(Q))$ becomes an instance of SIDH_1 rather than SIDH_0, and may thus be solved in polynomial time depending on the relative sizes of d and N. The fix is to multiply $\phi(Q)$ by a random scalar to hide the exact image of Q. The bug is present in [37] and [34], although the latter has been fixed in the online version [33].[4]

5.6 $\left(\begin{smallmatrix} * & \\ & * \end{smallmatrix}\right)$-SIDH a.k.a. Diagonal SIDH

Appearance: FESTA [8], CSIDH [18], SCALLOP [35], binSIDH [7].
Reduces to: $\left(\begin{smallmatrix} * & * \\ * & * \end{smallmatrix}\right)$-SIDH, when N contains a large smooth square factor.
Best attacks: generic.

The last group we consider is the diagonal group $\{\left(\begin{smallmatrix} * & \\ & * \end{smallmatrix}\right)\}$, whose associated level structures are pairs of cyclic subgroups of order N, sometimes called *split Cartan level structures* (see [26]). Corollary 13 shows that $\left(\begin{smallmatrix} * & \\ & * \end{smallmatrix}\right)$-SIDH is equivalent to SIDH_0 when N has a large square smooth factor.

The supersingular version of this problem appears in the security analysis of the FESTA encryption scheme [8], where it is called the Computational Isogeny with Scaled Torsion (CIST) problem. Solving CIST breaks FESTA, however the IND-CCA security of the scheme reduces to a "double" variant of CIST named CIST^2. The $\left(\begin{smallmatrix} * & \\ & * \end{smallmatrix}\right)$-SIDH problem also appears in binSIDH [7] where it is referred to as the Artificially Oriented Isogeny Problem.

[4] We heard rumors that the same bug was at some point also present in the source code for [6], but, looking at https://github.com/trusted-isogenies/SECUER-pok, it appears to have been fixed.

$(^*_*)$-SIDH also naturally appears in the theory of isogeny volcanoes, and thus in all "group action" schemes such as Couveignes–Rostovtsev–Stolbunov (CRS) [30,64], CSIDH [18] and SCALLOP [35].

When $\mathcal{O}$ is an imaginary-quadratic order, an *$\mathcal{O}$-oriented curve* [27] is an elliptic curve E together with an injection $\iota : \mathcal{O} \hookrightarrow \mathrm{End}(E)$. Thus, ordinary curves are $\mathrm{End}(E)$-oriented, curves in CSIDH are $\mathbb{Z}[\sqrt{-p}]$-oriented, and curves in SCALLOP are oriented by an order of large prime conductor inside $\mathbb{Z}[\sqrt{-1}]$.

To an ideal $\mathfrak{a}$ of $\mathcal{O}$ of norm coprime to q, we associate a subgroup

$$E[\mathfrak{a}] = \bigcap_{\alpha \in \mathfrak{a}} \{P \in E \mid \iota(\alpha)(P) = 0\}$$

and an isogeny $\phi_\mathfrak{a} : E \to E/E[\mathfrak{a}]$. A prime ℓ splits in $\mathcal{O}$ if the ideal $\ell\mathcal{O}$ factors as a product of two distinct ideals $\mathfrak{l}\bar{\mathfrak{l}}$, and in this case any $\mathcal{O}$-oriented curve has two distinguished cyclic groups of order ℓ, namely $E[\mathfrak{l}]$ and $E[\bar{\mathfrak{l}}]$.

Now suppose that one is given a pair of $\mathcal{O}$-oriented curves E and E', with the promise that there exists a d-isogeny $\phi : E \to E'$, then we can formulate the search for ϕ as a $(^*_*)$-SIDH problem as follows:

1. Let f be the conductor[5] of $\mathcal{O}$. For some bound B, take all the split primes $\ell_i < B$ not dividing df. Set $N = \prod \ell_i$.
2. For each ℓ_i factor $\ell_i\mathcal{O} = \mathfrak{l}_i\bar{\mathfrak{l}}_i$.
3. For each ℓ_i, compute $\langle P_i \rangle = E[\mathfrak{l}_i]$, $\langle Q_i \rangle = E[\bar{\mathfrak{l}}_i]$, $\langle P'_i \rangle = E'[\mathfrak{l}_i]$ and $\langle Q'_i \rangle = E'[\bar{\mathfrak{l}}_i]$. Each of the generators is defined in an extension of degree $O(B)$ of $\mathbb{F}_q$.
4. Then $\phi(\langle P_i \rangle) = \langle P'_i \rangle$ and $\phi(\langle Q_i \rangle) = \langle Q'_i \rangle$, defining a $(^*_*)$-SIDH problem of level N.

We stress that for any $\mathcal{O}$ there exists an infinity of split primes and that the asymptotic proportion of split primes is $1/2$. Hence N is bounded only by the largest extension of $\mathbb{F}_q$ we are willing to perform computations in.

Some instantiations of isogeny group actions, e.g. [22], do feature a known fixed degree d for all secrets, thus the strategy above reduces their key-recovery problem to $(^*_*)$-SIDH. However in the general key-recovery problem of isogeny group actions, the degree of ϕ is unknown: one is only given two $\mathcal{O}$-oriented curves $E, E'/\mathbb{F}_q$, and the goal is to find an ideal $\mathfrak{a}$ such that $\phi_\mathfrak{a} : E \to E'$. We now heuristically reduce this problem, known as the Group Action Inverse Problem (GAIP), to the $(^*_*)$-SIDH problem by arguing that there exists a polynomially-sized degree d that works for almost all pairs (E, E').

Lemma 14. *Let $\mathcal{O}$ be an imaginary-quadratic order of discriminant $-\Delta$ and suppose that all but $O(\log\log\Delta)$ prime factors of Δ are bounded by $\mathrm{polylog}(\Delta)$. Under heuristics on the distribution of ideal classes, one can find an integer d of size polynomial in $\log(q)$ and $\log(\Delta)$, such that for any polynomial g one can find $N \geq \exp(g(\log d))$ for which the GAIP with respect to $\mathcal{O}$ reduces in polynomial time to a (d, Γ)-SIDH problem with Γ the diagonal subgroup of $\mathrm{SL}_2(\mathbb{Z}/N\mathbb{Z})$.*

[5] The conductor of a quadratic order is its index in the ring of integers of its fraction field.

Proof. To apply the reduction discussed above, we sample a "default degree" d from a distribution which heuristically yields a valid guess for the degree at a $1/\operatorname{polylog}(\Delta)$ rate. The distribution is designed to work around two major obstructions: First, genus theory for binary quadratic forms implies that a suitable d must have the correct values for the *genus characters* of $\mathcal{O}$. We may recover the correct symbols for all characters of moduli lying in $\operatorname{polylog}(\Delta)$ using the DDH attack of [20,17] and randomly guess the remaining $O(\log\log\Delta)$ symbols; the chance of being correct is thus $1/\operatorname{polylog}(\Delta)$. Second, a much more elementary constraint is that if d has ω prime factors, then there are at most 2^ω distinct $\mathcal{O}$-ideals of norm d; hence, any d covering all ideal classes in some genus must necessarily have many prime factors. Absent any unforeseen further obstructions, it should therefore suffice to choose d as a product of a set of $\tilde{\Omega}(\log\Delta)$ random split primes of size $\operatorname{polylog}(\Delta)$, such that the chosen values for the genus characters are satisfied.[6]

Now let E and E' be two $\mathcal{O}$-oriented curves and assume there exists an ideal $\mathfrak{a} \subset \mathcal{O}$ of norm d such that $\phi_\mathfrak{a} : E \to E'$. Let g be a polynomial. To invoke the (d, Γ)-SIDH oracle, we proceed as sketched in the discussion before the lemma: Take enough small primes ℓ_i split in $\mathcal{O}$ such that $N = \prod \ell_i \geq \exp(g(\log d))$, compute the associated subgroups $E[\mathfrak{l}_i], E[\bar{\mathfrak{l}}_i], E'[\mathfrak{l}_i], E'[\bar{\mathfrak{l}}_i]$, and define the level structures $S = \left(\bigoplus E[\mathfrak{l}_i], \bigoplus E[\bar{\mathfrak{l}}_i]\right)$ and $S' = \left(\bigoplus E'[\mathfrak{l}_i], \bigoplus E'[\bar{\mathfrak{l}}_i]\right)$. Finally, pass (E, S, E', S') to the (d, Γ)-SIDH oracle.

There is one more caveat: In the supersingular setting, solving the constructed $(^*{}_*)$-SIDH problem may in fact produce a generic isogeny that fails to respect the orientation. In that case, to complete the reduction, we make use of known reductions between the supersingular endomorphism ring problem and the GAIP problem due to [19,66]: We first generate any $\mathcal{O}$-oriented supersingular curve E_0 of known endomorphism ring and apply the reduction above to both (E_0, E) and (E_0, E'). The resulting knowledge of smooth-degree isogenies $E_0 \to E$ and $E_0 \to E'$ thus reveals the endomorphism rings of E and E', respectively. This puts us into a position to invoke the reduction and finally recover an $\mathcal{O}$-ideal connecting E and E', as desired. □

In all existing isogeny group actions (CRS, CSIDH, SCALLOP), the discriminant of the order $\mathcal{O}$ in play has very few prime factors, thus satisfying the conditions of Lemma 14 on the factorization of Δ. Hence the respective GAIPs reduce to a $(^*{}_*)$-SIDH (and SIDH_0) problem of level N of arbitrary size.

Improved algorithm for isogenies between ordinary curves. The fact that computing isogenies between oriented curves reduces to an $(^*{}_*)$-SIDH

[6] Evidence (conditional on GRH) in support of this heuristic are (1) Bach's proof that ideals of norm up to $6\log(\Delta)^2$ generate the class group [4]; and (2) results in a line of work initiated by Landau and Bernays, see [10, part II, § 3] or for instance [58, § 5], which bound the density of integers represented by *all* quadratic forms in a given genus; however, those results are asymptotic and only become meaningful far beyond the sizes (relative to Δ) that are required for our purposes.

problem is well-known, and was already used in [36] in the context of the SEA point-counting algorithm [65]. The goal here is to compute an isogeny $\phi : E \to E'$ between ordinary curves, for a moderately large $d = \deg \phi$.

The algorithm in [36] finds a small split prime ℓ not dividing d and assumes that E and E' are on the craters of the respective ℓ-volcanoes (see Examples 5 and 9 for more background). Then it computes a *horizontal* basis B of $E[\ell^e]$ for $\ell^e > 2\sqrt{d}$, i.e., a basis (P, Q) such that $\langle P \rangle = \ker \phi_{\mathfrak{l}^e}$ and $\langle Q \rangle = \ker \phi_{\bar{\mathfrak{l}}^e}$. Then it computes a horizontal basis B' for $E'[\ell^e]$, so that $\phi(B) = \gamma \cdot B'$ for some matrix $\gamma \in \mathrm{GL}_2$. It finally finds γ by exhaustive search: for each choice it computes $\gamma \cdot B'$ and tries to compute ϕ by interpolation, halting when this succeeds. Since there are $O(\ell^{2e})$ diagonal matrices, and each interpolation step costs $\tilde{O}(d)$ operations, the total complexity is in $\tilde{O}(d^2)$.

We can give a first improvement using the technique of Lemma 6: choose B' so that $e_{\ell^e}(B') = e_{\ell^e}(B)^d$, then it is sufficient to go through the diagonal matrices in SL_2, which only number $O(\ell^e)$, leading to an algorithm with complexity $\tilde{O}(d^{1.5})$. Yet another improvement has recently become possible: instead of interpolation, we can use the algorithm of Theorem 11—whose dependency in d is only polylog(d)—to test whether ϕ maps B to $\gamma \cdot B'$. Hence, we can find γ in $\tilde{O}(\sqrt{d})$ operations. In principle, we could stop here and use the isogeny representation returned by Theorem 11, however the goal of [36] is to compute ϕ as a rational fraction, which we can now do by interpolation using $\tilde{O}(d)$ operations, which is quasi-optimal.

When E and E' are not on the crater, the paper reduces to this case by walking up the volcano. However one can again improve the algorithm as follows. Say the curves are at depth h, the ℓ-Sylows of E and E' will typically contain a unique cyclic subgroup of maximal order, at least ℓ^{2h} (see [57] for details). These two groups must be mapped one onto the other by ϕ, thus they define a $\left(\begin{smallmatrix} * & * \\ * & \end{smallmatrix}\right)$-SIDH problem. Using Corollary 13 we reduce this to a $\left(\begin{smallmatrix} * & \\ * & \end{smallmatrix}\right)$-SIDH problem of level (at least) ℓ^h, call it (F, S_1, F', S_1').

Now we proceed like before and compute horizontal bases of $F[\ell^h]$ and $F'[\ell^h]$, which define a second $\left(\begin{smallmatrix} * & \\ * & \end{smallmatrix}\right)$-SIDH problem, call it (F, S_2, F', S_2'). As argued above, S_2 (and S_2') consists of all possible bases (P, Q) such that $\langle P \rangle = \ker \phi_{\mathfrak{l}^h}$ and $\langle Q \rangle = \ker \phi_{\bar{\mathfrak{l}}^h}$. However S_1 (and S_1') consists of all bases (P, R) such that $\langle P \rangle = \ker \phi_{\mathfrak{l}^h}$ and $\langle R \rangle$ generates the kernel of the ℓ^h-isogeny descending back towards E.

Combining both constraints with the one coming from the Weil pairing, all degrees of freedom are removed and we are left with a pure SIDH problem. In the extreme case where $\ell^{2h} > 4d$, we can directly compute ϕ without trial-and-error; otherwise the number of trials will be divided by ℓ^h.

In conclusion, we have a generic algorithm with quasi-optimal complexity in d for computing isogenies of known degree between ordinary curves, albeit with large constants hidden inside the $O()$. This was previously only known in the case where the characteristic is larger than $2d$ and the curve models are *normalized* [12].

This new algorithm asymptotically beats all previously known approaches to the so called "Elkies step" of the SEA algorithm in the regime where characteristic and extension degree both grow polynomially, and may be useful in practice for solving some large point counting problems (likely outside cases of cryptographic interest). Note, however, that the overall complexity of SEA is dominated by other steps, thus we do not have an asymptotic improvement on point counting.

6 Conclusion

We introduced a new framework to study assumptions in isogeny-based cryptography. We hope that it will help classify and relate seemingly distant isogeny assumptions.

An important consequence of our main theorem is that the image of a single point of large-enough order is in general sufficient to recover an isogeny. Understanding how far these attacks can be pushed is a fundamental question for isogeny-based cryptography.

We also showed that CSIDH reduces to Diagonal SIDH, which is equivalent to SIDH_0, and is thus no harder than breaking SIDH-like proofs of knowledge. While this may lend more credibility to the security of these proofs of knowledge, we warn against using the reduction to set parameters. Indeed the best quantum attacks against CSIDH are subexponential, whereas the best quantum attacks against SIDH-like proofs of knowledge are exponential, thus the reduction is void if one uses the best possible parameters. Conversely, it is an interesting question whether or not SIDH_0 and Diagonal SIDH can be solved in quantum subexponential time.

Acknowledgments. Luca De Feo acknowledges support from the Swiss National Science Foundation through grant no. 213766, CryptonIs. This research project was initiated at the Isogeny Days 2022 workshop organized with support from the ERC grant no. 101020788, ISOCRYPT. At the time, Lorenz Panny was a postdoc at Academia Sinica, Taiwan, funded by Academia Sinica Investigator Award AS-IA-109-M01.

We would like to thank Wouter Castryck, Antonin Leroux, Christophe Petit, Fréderik Vercauteren and Benjamin Wesolowski for the fruitful discussions, and the anonymous referees for their useful suggestions.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Adj, G., Cervantes-Vázquez, D., Chi-Domínguez, J.J., Menezes, A., Rodríguez-Henríquez, F.: On the cost of computing isogenies between supersingular elliptic curves. In: Cid, C., Jacobson Jr:, M.J. (eds.) SAC 2018. LNCS, vol. 11349, pp. 322–343. Springer, Heidelberg (Aug 2019). https://doi.org/10.1007/978-3-030-10970-7_15

2. Agrawal, S., Lin, D. (eds.): ASIACRYPT 2022, Part II, LNCS, vol. 13792. Springer, Heidelberg (Dec 2022)
3. Arpin, S.: Adding level structure to supersingular elliptic curve isogeny graphs (2023). https://doi.org/10.48550/arXiv.2203.03531
4. Bach, E.: Analytic methods in the analysis and design of number-theoretic algorithms. MIT press Cambridge (1985)
5. Basso, A., Chen, M., Fouotsa, T.B., Kutas, P., Laval, A., Marco, L., Tchoffo Saah, G.: Exploring SIDH-based signature parameters. Cryptology ePrint Archive, Paper 2023/1906 (2023), https://eprint.iacr.org/2023/1906
6. Basso, A., Codogni, G., Connolly, D., De Feo, L., Fouotsa, T.B., Lido, G.M., Morrison, T., Panny, L., Patranabis, S., Wesolowski, B.: Supersingular curves you can trust. In: Hazay, C., Stam, M. (eds.) EUROCRYPT 2023, Part II. LNCS, vol. 14005, pp. 405–437. Springer, Heidelberg (Apr 2023). https://doi.org/10.1007/978-3-031-30617-4_14
7. Basso, A., Fouotsa, T.B.: New SIDH countermeasures for a more efficient key exchange. In: Guo, J., Steinfeld, R. (eds.) ASIACRYPT 2023, Part VIII. LNCS, vol. 14445, pp. 208–233. Springer, Heidelberg (Dec 2023). https://doi.org/10.1007/978-981-99-8742-9_7
8. Basso, A., Maino, L., Pope, G.: FESTA: Fast encryption from supersingular torsion attacks. In: Guo and Steinfeld [47], pp. 98–126. https://doi.org/10.1007/978-981-99-8739-9_4
9. Benčina, B., Kutas, P., Merz, S.P., Petit, C., Stopar, M., Weitkämper, C.: Improved algorithms for finding fixed-degree isogenies between supersingular elliptic curves. Cryptology ePrint Archive, Paper 2023/1618 (2023), https://eprint.iacr.org/2023/1618
10. Bernays, P.: Über die Darstellung von positiven, ganzen Zahlen durch die primitiven, binären quadratischen Formen einer nicht-quadratischen Diskriminante. Ph.D. thesis, Georg-August-Universität, Göttingen (1912)
11. Biasse, J.F., Jao, D., Sankar, A.: A quantum algorithm for computing isogenies between supersingular elliptic curves. In: Meier, W., Mukhopadhyay, D. (eds.) INDOCRYPT 2014. LNCS, vol. 8885, pp. 428–442. Springer, Heidelberg (Dec 2014). https://doi.org/10.1007/978-3-319-13039-2_25
12. Bostan, A., Morain, F., Salvy, B., Schost, E.: Fast algorithms for computing isogenies between elliptic curves. Mathematics of Computation **77**(263), 1755–1778 (Sep 2008). https://doi.org/10.1090/s0025-5718-08-02066-8, http://dx.doi.org/10.1090/S0025-5718-08-02066-8
13. Burdges, J., De Feo, L.: Delay encryption. In: Canteaut, A., Standaert, F.X. (eds.) EUROCRYPT 2021, Part I. LNCS, vol. 12696, pp. 302–326. Springer, Heidelberg (Oct 2021). https://doi.org/10.1007/978-3-030-77870-5_11
14. Cantor, D.G., Zassenhaus, H.: A new algorithm for factoring polynomials over finite fields. Mathematics of Computation **36**, 587–592 (1981)
15. Castryck, W., Decru, T.: An efficient key recovery attack on SIDH. In: Hazay and Stam [48], pp. 423–447. https://doi.org/10.1007/978-3-031-30589-4_15
16. Castryck, W., Houben, M., Merz, S.P., Mula, M., van Buuren, S., Vercauteren, F.: Weak instances of class group action based cryptography via self-pairings. Lecture Notes in Computer Science p. 762–792 (2023). https://doi.org/10.1007/978-3-031-38548-3_25, http://dx.doi.org/10.1007/978-3-031-38548-3_25
17. Castryck, W., Houben, M., Vercauteren, F., Wesolowski, B.: On the decisional Diffie–Hellman problem for class group actions on oriented elliptic curves. Research in Number Theory **8**(4), 99 (Nov 2022). https://doi.org/10.1007/s40993-022-00399-6

18. Castryck, W., Lange, T., Martindale, C., Panny, L., Renes, J.: CSIDH: An efficient post-quantum commutative group action. In: Peyrin, T., Galbraith, S. (eds.) ASIACRYPT 2018, Part III. LNCS, vol. 11274, pp. 395–427. Springer, Heidelberg (Dec 2018). https://doi.org/10.1007/978-3-030-03332-3_15
19. Castryck, W., Panny, L., Vercauteren, F.: Rational isogenies from irrational endomorphisms. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020, Part II. LNCS, vol. 12106, pp. 523–548. Springer, Heidelberg (May 2020). https://doi.org/10.1007/978-3-030-45724-2_18
20. Castryck, W., Sotáková, J., Vercauteren, F.: Breaking the decisional Diffie-Hellman problem for class group actions using genus theory. In: Micciancio, D., Ristenpart, T. (eds.) CRYPTO 2020, Part II. LNCS, vol. 12171, pp. 92–120. Springer, Heidelberg (Aug 2020). https://doi.org/10.1007/978-3-030-56880-1_4
21. Castryck, W., Vercauteren, F.: A polynomial time attack on instances of M-SIDH and FESTA. In: Guo and Steinfeld [47], pp. 127–156. https://doi.org/10.1007/978-981-99-8739-9_5
22. Chávez-Saab, J., Chi-Domínguez, J.J., Jaques, S., Rodríguez-Henríquez, F.: The SQALE of CSIDH: sublinear Vélu quantum-resistant isogeny action with low exponents. Journal of Cryptographic Engineering **12**(3), 349–368 (Sep 2022). https://doi.org/10.1007/s13389-021-00271-w
23. Chavez-Saab, J., Santos, M.C., De Feo, L., Eriksen, J.K., Hess, B., Kohel, D., Leroux, A., Longa, P., Meyer, M., Panny, L., Patranabis, S., Petit, C., Rodríguez Henríquez, F., Schaeffler, S., Wesolowski, B.: SQIsign. Tech. rep., National Institute of Standards and Technology (2023), available at https://csrc.nist.gov/Projects/pqc-dig-sig/round-1-additional-signatures
24. Chen, M., Imran, M., Ivanyos, G., Kutas, P., Leroux, A., Petit, C.: Hidden stabilizers, the isogeny to endomorphism ring problem and the cryptanalysis of pSIDH. In: Guo, J., Steinfeld, R. (eds.) ASIACRYPT 2023, Part III. LNCS, vol. 14440, pp. 99–130. Springer, Heidelberg (Dec 2023). https://doi.org/10.1007/978-981-99-8727-6_4
25. Childs, A., Jao, D., Soukharev, V.: Constructing elliptic curve isogenies in quantum subexponential time. Journal of Mathematical Cryptology **8**(1), 1–29 (2014). https://doi.org/10.1515/jmc-2012-0016
26. Codogni, G., Lido, G.: Spectral theory of isogeny graphs (2023). https://doi.org/10.48550/arXiv.2308.13913
27. Colò, L., Kohel, D.: Orienting supersingular isogeny graphs. Journal of Mathematical Cryptology **14**(1), 414–437 (2020). https://doi.org/10.1515/jmc-2019-0034, https://doi.org/10.1515/jmc-2019-0034
28. Costello, C.: B-SIDH: Supersingular isogeny Diffie-Hellman using twisted torsion. In: Moriai, S., Wang, H. (eds.) ASIACRYPT 2020, Part II. LNCS, vol. 12492, pp. 440–463. Springer, Heidelberg (Dec 2020). https://doi.org/10.1007/978-3-030-64834-3_15
29. Costello, C., Longa, P., Naehrig, M., Renes, J., Virdia, F.: Improved classical cryptanalysis of SIKE in practice. In: Kiayias, A., Kohlweiss, M., Wallden, P., Zikas, V. (eds.) PKC 2020, Part II. LNCS, vol. 12111, pp. 505–534. Springer, Heidelberg (May 2020). https://doi.org/10.1007/978-3-030-45388-6_18
30. Couveignes, J.M.: Hard homogeneous spaces. Cryptology ePrint Archive, Report 2006/291 (2006), https://eprint.iacr.org/2006/291
31. Dartois, P., Leroux, A., Robert, D., Wesolowski, B.: SQISignHD: new dimensions in cryptography. In: EUROCRYPT 2024. LNCS, Springer (2024), https://eprint.iacr.org/2023/436

32. De Feo, L., Delpech de Saint Guilhem, C., Fouotsa, T.B., Kutas, P., Leroux, A., Petit, C., Silva, J., Wesolowski, B.: Séta: Supersingular encryption from torsion attacks. In: Tibouchi, M., Wang, H. (eds.) ASIACRYPT 2021, Part IV. LNCS, vol. 13093, pp. 249–278. Springer, Heidelberg (Dec 2021). https://doi.org/10.1007/978-3-030-92068-5_9
33. De Feo, L., Dobson, S., Galbraith, S.D., Zobernig, L.: SIDH proof of knowledge. Cryptology ePrint Archive, Report 2021/1023 (2021), https://eprint.iacr.org/2021/1023
34. De Feo, L., Dobson, S., Galbraith, S.D., Zobernig, L.: SIDH proof of knowledge. In: Agrawal and Lin [2], pp. 310–339. https://doi.org/10.1007/978-3-031-22966-4_11
35. De Feo, L., Fouotsa, T.B., Kutas, P., Leroux, A., Merz, S.P., Panny, L., Wesolowski, B.: SCALLOP: Scaling the CSI-FiSh. In: Boldyreva, A., Kolesnikov, V. (eds.) PKC 2023, Part I. LNCS, vol. 13940, pp. 345–375. Springer, Heidelberg (May 2023). https://doi.org/10.1007/978-3-031-31368-4_13
36. De Feo, L., Hugounenq, C., Plût, J., Schost, E.: Explicit isogenies in quadratic time in any characteristic. LMS Journal of Computation and Mathematics **19**(A), 267–282 (2016). https://doi.org/10.1112/s146115701600036x, http://dx.doi.org/10.1112/S146115701600036X
37. De Feo, L., Jao, D., Plût, J.: Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. Journal of Mathematical Cryptology **8**(3), 209–247 (2014). https://doi.org/10.1515/jmc-2012-0015
38. De Feo, L., Kohel, D., Leroux, A., Petit, C., Wesolowski, B.: SQISign: Compact post-quantum signatures from quaternions and isogenies. In: Moriai, S., Wang, H. (eds.) ASIACRYPT 2020, Part I. LNCS, vol. 12491, pp. 64–93. Springer, Heidelberg (Dec 2020). https://doi.org/10.1007/978-3-030-64837-4_3
39. De Feo, L., Masson, S., Petit, C., Sanso, A.: Verifiable delay functions from supersingular isogenies and pairings. In: Galbraith, S.D., Moriai, S. (eds.) ASIACRYPT 2019, Part I. LNCS, vol. 11921, pp. 248–277. Springer, Heidelberg (Dec 2019). https://doi.org/10.1007/978-3-030-34578-5_10
40. Delfs, C., Galbraith, S.D.: Computing isogenies between supersingular elliptic curves over $\mathbb{F}_p$. Designs, Codes and Cryptography **78**(2), 425–440 (Feb 2016). https://doi.org/10.1007/s10623-014-0010-1
41. Eisenträger, K., Hallgren, S., Lauter, K.E., Morrison, T., Petit, C.: Supersingular isogeny graphs and endomorphism rings: Reductions and solutions. In: Nielsen, J.B., Rijmen, V. (eds.) EUROCRYPT 2018, Part III. LNCS, vol. 10822, pp. 329–368. Springer, Heidelberg (Apr / May 2018). https://doi.org/10.1007/978-3-319-78372-7_11
42. Eisenträger, K., Hallgren, S., Leonardi, C., Morrison, T., Park, J.: Computing endomorphism rings of supersingular elliptic curves and connections to path-finding in isogeny graphs. Open Book Series **4**(1), 215–232 (Dec 2020). https://doi.org/10.2140/obs.2020.4.215, http://dx.doi.org/10.2140/obs.2020.4.215
43. Fouotsa, T.B., Moriya, T., Petit, C.: M-SIDH and MD-SIDH: Countering SIDH attacks by masking information. In: Hazay and Stam [48], pp. 282–309. https://doi.org/10.1007/978-3-031-30589-4_10
44. Fouquet, M., Morain, F.: Isogeny volcanoes and the SEA algorithm. In: Fieker, C., Kohel, D.R. (eds.) Algorithmic Number Theory Symposium. Lecture Notes in Computer Science, vol. 2369, pp. 47–62. Springer Berlin / Heidelberg, Berlin, Heidelberg (2002). https://doi.org/10.1007/3-540-45455-1_23

45. Fuselier, J., Iezzi, A., Kozek, M., Morrison, T., Namoijam, C.: Computing supersingular endomorphism rings using inseparable endomorphisms (2023). https://doi.org/10.48550/arXiv.2306.03051
46. Galbraith, S.D., Hess, F., Smart, N.P.: Extending the GHS Weil descent attack. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 29–44. Springer, Heidelberg (Apr / May 2002). https://doi.org/10.1007/3-540-46035-7_3
47. Guo, J., Steinfeld, R. (eds.): ASIACRYPT 2023, Part VII, LNCS, vol. 14444. Springer, Heidelberg (Dec 2023)
48. Hazay, C., Stam, M. (eds.): EUROCRYPT 2023, Part V, LNCS, vol. 14008. Springer, Heidelberg (Apr 2023)
49. Jao, D., Azarderakhsh, R., Campagna, M., Costello, C., De Feo, L., Hess, B., Jalali, A., Koziel, B., LaMacchia, B., Longa, P., Naehrig, M., Renes, J., Soukharev, V., Urbanik, D., Pereira, G., Karabina, K., Hutchinson, A.: SIKE. Tech. rep., National Institute of Standards and Technology (2022), available at https://csrc.nist.gov/Projects/post-quantum-cryptography/round-4-submissions
50. Jao, D., De Feo, L.: Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. In: Yang, B.Y. (ed.) Post-Quantum Cryptography - 4th International Workshop, PQCrypto 2011. pp. 19–34. Springer, Heidelberg (Nov / Dec 2011). https://doi.org/10.1007/978-3-642-25405-5_2
51. Jao, D., Soukharev, V.: A subexponential algorithm for evaluating large degree isogenies. Algorithmic Number Theory p. 219–233 (2010). https://doi.org/10.1007/978-3-642-14518-6_19, http://dx.doi.org/10.1007/978-3-642-14518-6_19
52. Kohel, D.: Endomorphism rings of elliptic curves over finite fields. Ph.D. thesis, University of California at Berkeley (1996), https://i2m.univ-amu.fr/perso/david.kohel/pub/thesis.pdf
53. Kuperberg, G.: A subexponential-time quantum algorithm for the dihedral hidden subgroup problem. SIAM Journal of Computing **35**(1), 170–188 (2005). https://doi.org/10.1137/S0097539703436345
54. Kuperberg, G.: Another Subexponential-time Quantum Algorithm for the Dihedral Hidden Subgroup Problem. In: Severini, S., Brandao, F. (eds.) 8th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2013). Leibniz International Proceedings in Informatics (LIPIcs), vol. 22, pp. 20–34. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany (2013). https://doi.org/10.4230/LIPIcs.TQC.2013.20
55. Leroux, A.: A new isogeny representation and applications to cryptography. In: Agrawal and Lin [2], pp. 3–35. https://doi.org/10.1007/978-3-031-22966-4_1
56. Maino, L., Martindale, C., Panny, L., Pope, G., Wesolowski, B.: A direct key recovery attack on SIDH. In: Hazay and Stam [48], pp. 448–471. https://doi.org/10.1007/978-3-031-30589-4_16
57. Miret, J.M., Sadornil, D., Tena, J., Tomàs, R., Valls, M.: Volcanoes of ℓ-isogenies of elliptic curves. Publicacions Matemàtiques pp. 165–180 (2007), https://www.raco.cat/index.php/PublicacionsMatematiques/article/download/69987/387563
58. Odoni, R.: A new equidistribution property of norms of ideals in given classes. Acta Arithmetica **33**(1), 53–63 (1977)
59. Page, A., Wesolowski, B.: The supersingular endomorphism ring and one endomorphism problems are equivalent. Cryptology ePrint Archive, Paper 2023/1399 (2023), https://eprint.iacr.org/2023/1399

60. Petit, C.: Faster algorithms for isogeny problems using torsion point images. In: Takagi, T., Peyrin, T. (eds.) ASIACRYPT 2017, Part II. LNCS, vol. 10625, pp. 330–353. Springer, Heidelberg (Dec 2017). https://doi.org/10.1007/978-3-319-70697-9_12
61. Regev, O.: A subexponential time algorithm for the dihedral hidden subgroup problem with polynomial space. arXiv:quant-ph/0406151 (Jun 2004), http://arxiv.org/abs/quant-ph/0406151
62. Robert, D.: Evaluating isogenies in polylogarithmic time. Cryptology ePrint Archive, Report 2022/1068 (2022), https://eprint.iacr.org/2022/1068
63. Robert, D.: Breaking SIDH in polynomial time. In: Hazay and Stam [48], pp. 472–503. https://doi.org/10.1007/978-3-031-30589-4_17
64. Rostovtsev, A., Stolbunov, A.: Public-Key Cryptosystem Based On Isogenies. Cryptology ePrint Archive, Report 2006/145 (2006), https://eprint.iacr.org/2006/145
65. Schoof, R.: Counting points on elliptic curves over finite fields. Journal de théorie des nombres de Bordeaux **7**(1), 219–254 (1995), http://www.numdam.org/item/JTNB_1995__7_1_219_0/
66. Wesolowski, B.: Orientations and the supersingular endomorphism ring problem. In: Dunkelman, O., Dziembowski, S. (eds.) EUROCRYPT 2022, Part III. LNCS, vol. 13277, pp. 345–371. Springer, Heidelberg (May / Jun 2022). https://doi.org/10.1007/978-3-031-07082-2_13
67. Wesolowski, B.: The supersingular isogeny path and endomorphism ring problems are equivalent. In: 62nd FOCS. pp. 1100–1111. IEEE Computer Society Press (Feb 2022). https://doi.org/10.1109/FOCS52979.2021.00109