

Orientations and the supersingular endomorphism ring problem

Benjamin Wesolowski

Univ. Bordeaux, CNRS, Bordeaux INP, IMB, UMR 5251, F-33400, Talence, France
INRIA, IMB, UMR 5251, F-33400, Talence, France

Abstract. We study two important families of problems in isogeny-based cryptography and how they relate to each other: computing the endomorphism ring of supersingular elliptic curves, and inverting the action of class groups on oriented supersingular curves. We prove that these two families of problems are closely related through polynomial-time reductions, assuming the generalised Riemann hypothesis.

We identify two classes of essentially equivalent problems. The first class corresponds to the problem of computing the endomorphism ring of *oriented curves*. The security of a large family of cryptosystems (such as CSIDH) reduces to (and sometimes from) this class, for which there are heuristic quantum algorithms running in subexponential time. The second class corresponds to computing the endomorphism ring of *orientable curves*. The security of essentially all isogeny-based cryptosystems reduces to (and sometimes from) this second class, for which the best known algorithms are still exponential.

Some of our reductions not only generalise, but also strengthen previously known results. For instance, it was known that in the particular case of curves defined over \mathbf{F}_p , the security of CSIDH reduces to the endomorphism ring problem in subexponential time. Our reductions imply that the security of CSIDH is actually equivalent to the endomorphism ring problem, under polynomial time reductions (circumventing arguments that proved such reductions unlikely).

1 Introduction

We study two families of computational problems at the heart of isogeny-based cryptography, and how they relate to each other: computing the endomorphism ring of supersingular elliptic curves, and inverting the action of class groups on oriented supersingular curves. On one hand, the problem of computing endomorphism rings is of foundational importance to the field: its presumed hardness is necessary [GPST16,CPV20,FKM21] (and sometimes sufficient [CLG09,GPS20]) for the security of essentially all isogeny-based cryptosystems. On the other hand, the action of *ideal class groups* on sets of elliptic curves induces presumably hard inversion problems. This action, and the presumed hardness of its inversion, is the fertile ground upon which many cryptosystems have been built — from the early work of Couveignes [Cou06], to CSIDH [CLM⁺18] and its many variants [CD20,BKV19,CS21]. Thanks to the notion of *orientation* introduced by

Colò and Kohel [CK20], it has recently become clear that such actions play a ubiquitous role in isogeny-based cryptography. On one hand, orientations provide a framework that directly generalises the family of CSIDH-like cryptosystems. On the other hand, it has been identified [DDF⁺21] that even the security of cryptosystems where group actions were not expected, such as SIDH [JD11] and its variants, reduces to an action inversion problem, called the *Uber isogeny problem*, opening a new cryptanalytic avenue.

1.1 Oriented endomorphism ring problems

Isogenies are morphisms between elliptic curves, and endomorphisms of an elliptic curve E are isogenies from E to itself. They form a ring, written $\text{End}(E)$. Given a supersingular elliptic curve E over $\overline{\mathbb{F}}_p$, the endomorphism ring problem ENDRING consists in computing a basis of $\text{End}(E)$. This ENDRING problem was proved in [Wes22] (and heuristically since [EHL⁺18]) to be equivalent to the problem of finding isogenies between supersingular elliptic curves, assuming the generalised Riemann hypothesis. Let \mathfrak{D} be an order in a quadratic number field K . An orientation is an embedding

$$\iota : \mathfrak{D} \hookrightarrow \text{End}(E)$$

which cannot be extended to a superorder of \mathfrak{D} . We call (E, ι) an \mathfrak{D} -oriented elliptic curve, and E is \mathfrak{D} -orientable. We introduce three *oriented* variants of the endomorphism ring problem, in increasing order of hardness (precise definitions are provided in Section 4):

- \mathfrak{D} - ENDRING : given an \mathfrak{D} -oriented elliptic curve (E, ι) , compute a basis of $\text{End}(E)$. It is presumably easier than ENDRING since ι provides additional information.
- $\text{ENDRING}|_{\mathfrak{D}}$: given an \mathfrak{D} -orientable elliptic curve E , compute a basis of $\text{End}(E)$. It is simply the restriction of ENDRING to \mathfrak{D} -orientable inputs.
- \mathfrak{D} - ENDRING^* : given an \mathfrak{D} -orientable E , compute a basis of $\text{End}(E)$ together with an \mathfrak{D} -orientation expressed in this basis.

1.2 Class group action problems

A key feature of \mathfrak{D} -orientations is that they induce a group action. Given an \mathfrak{D} -oriented (E, ι) , and an invertible \mathfrak{D} -ideal \mathfrak{a} , one can construct another \mathfrak{D} -oriented elliptic curve $\mathfrak{a} \star (E, \iota) = (E^{\mathfrak{a}}, \iota^{\mathfrak{a}})$, and an isogeny $\varphi_{\mathfrak{a}} : E \rightarrow E^{\mathfrak{a}}$ connecting them. This construction induces a free action of the ideal class group $\text{Cl}(\mathfrak{D})$ on \mathfrak{D} -oriented curves up to isomorphism. We define four variants of the problem of inverting this group action (precise definitions are provided in Section 3):

- \mathfrak{D} - VECTORIZATION : given two \mathfrak{D} -oriented elliptic curves (E, ι) and (E', ι') , find an ideal \mathfrak{a} such that E' is isomorphic to $E^{\mathfrak{a}}$. The *effective* variant asks for

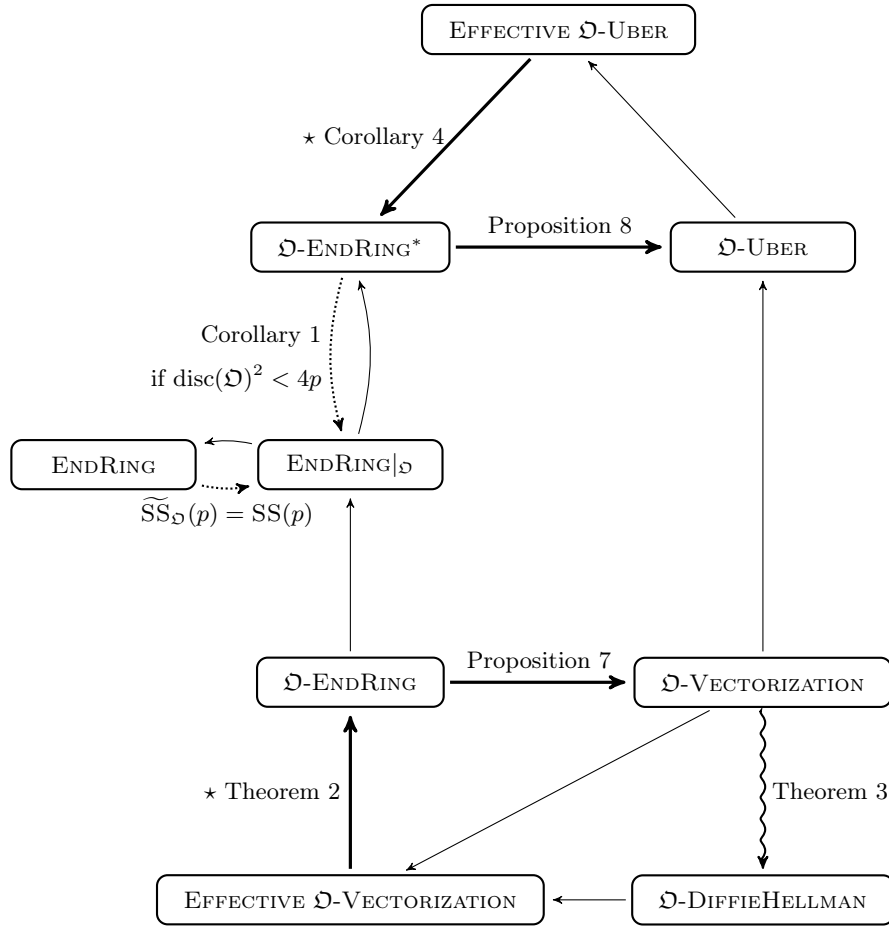


Fig. 1. Arrows represent probabilistic polynomial time reductions in the length of the instance (i.e., $\log p$, $\log |\text{disc}(\mathfrak{D})|$, and the length of the provided \mathfrak{D} -orientations), and for those marked by a star \star , in $\#(\text{Cl}(\mathfrak{D})[2])$, the size of the 2-torsion of the class group. Arrows with no reference (thin or dotted) are trivial reductions. Thick arrows assume that the factorisation of $\text{disc}(\mathfrak{D})$ is known. Dotted arrows assume the stated condition. The “snake” arrow is a quantum reduction. $\text{SS}(p)$ is the set of all supersingular elliptic curves over \mathbf{F}_{p^2} (up to isomorphism), and $\widetilde{\text{SS}}_{\mathfrak{D}}(p)$ is the subset of \mathfrak{D} -orientable curves. Non-trivial reductions assume the generalised Riemann hypothesis.

the isomorphism to preserve the orientation, and also requires a way to evaluate the action of \mathfrak{a} on any other \mathfrak{D} -oriented curve. The *vectorization* terminology comes from Couveignes’ work [Cou06]. The security of many cryptosystems reduces to this problem, such as CSIDH [CLM⁺18], CSI-FiSh [BKV19], CSURF [CD20], or generalisations to other orientations [CS21].

- \mathfrak{D} -UBER: given an \mathfrak{D} -oriented elliptic curve (E, ι) and an \mathfrak{D} -orientable E' , find an ideal \mathfrak{a} such that E' is isomorphic to $E^{\mathfrak{a}}$. The *effective* variant also requires a way to evaluate the action of \mathfrak{a} on any other \mathfrak{D} -oriented curve. This *Uber* terminology was introduced in [DDF⁺21], where it was shown that the security of many cryptosystems reduces to this problem, including SIDH [JD11], OSIDH [CK20] and Seta [DDF⁺21].

1.3 Contribution

The main contribution of this article is the various reductions summarised in Figure 1, under the generalised Riemann hypothesis. Some of these reductions generalise and strengthen previously known results:

- The article [CPV20] was the first to investigate the relation between ENDRING and the vectorisation problem, in the particular case of curves defined over \mathbf{F}_p (i.e., $\sqrt{-p} \in \mathfrak{D}$). They prove that knowledge of the endomorphism ring of a CSIDH public key allows one to recover the ideal class of the secret key. This surprising result, however, only implies a subexponential reduction from breaking CSIDH to computing endomorphism rings, because it is not easy to find a *good* ideal class representative of the key. In essence, they prove a reduction from the vectorisation problem, but not from its effective variant. They argue that this effectiveness seems hard to reach, because if an efficient reduction could find good class representatives, then there would be an efficient algorithm to compute discrete logarithms in class groups of large discriminant. We circumvent this issue, proving in Section 6 that the effective vectorisation problem (hence breaking CSIDH) does reduce to the endomorphism ring problem in polynomial time. Our reductions not only apply to CSIDH or close variants restricted to \mathbf{F}_p , but to arbitrary orientations, including generalisations such as [CS21]. To reach this level of generality, we introduce the notion of \mathfrak{D} -twists and prove that they enjoy similar properties to quadratic twists, and can be used to extend some of the techniques introduced in [CPV20].
- Considering vectorisation as a group-action analog of the discrete logarithm problem, there is a corresponding Diffie–Hellman analog, \mathfrak{D} -DIFFIEHELLMAN (sometimes called *parallelisation*). Properly instantiated, it corresponds to the problem of recovering the shared secret in CSIDH. In [GPSV21], it was proved that if the action of $\text{Cl}(\mathfrak{D})$ is efficiently computable, then the (non-effective) \mathfrak{D} -VECTORIZATION problem reduces to \mathfrak{D} -DIFFIEHELLMAN in quantum polynomial time. This result hit a similar wall as [CPV20]: the action is only efficiently computable for hard-to-find good ideal class representatives. Here again, our reductions bypass this obstacle, proving in Section 7 that \mathfrak{D} -DIFFIEHELLMAN is actually equivalent to \mathfrak{D} -VECTORIZATION (and to its effective variant, and to \mathfrak{D} -ENDRING) under quantum polynomial time reductions.

Finally, we focus in Section 8 on the case where \mathfrak{D} is a non-maximal order, proving reductions from our problems of interest to their *a priori* easier counterpart for superorders (with smaller discriminants and class groups).

1.4 Notation

We denote by \mathbf{Z} and \mathbf{Q} the ring of integers and the field of rational numbers. For any prime power q , we denote by \mathbf{F}_q the finite field with q elements, and $\overline{\mathbf{F}}_q$ its algebraic closure. We write $f = O(g)$ for the classic big O notation, which is equivalent to $g = \Omega(f)$. The size of a set S is denoted by $\#S$. Let a and b be two integers. We write $a \mid b$ if a divides b , and $a \parallel b$ if $a \mid b$ and $\gcd(a, b/a) = 1$. All statements containing the mention (GRH) assume the generalised Riemann hypothesis.

2 Preliminaries

In this section, we recall relevant notions related to quaternion algebras, supersingular elliptic curves, their endomorphism rings, and orientations.

2.1 Quadratic fields and orders

Let K be a quadratic number field. We write $\alpha \mapsto \bar{\alpha}$ for the conjugation, i.e., the unique non-trivial involution of K . Given any $\alpha \in K$, we write its norm $N(\alpha) = \alpha\bar{\alpha}$, and its trace $\text{Tr}(\alpha) = \alpha + \bar{\alpha}$. An *order* in K is a discrete subring $\mathfrak{D} \subset K$ such that $K = \mathbf{Q}\mathfrak{D}$. A *generator* of \mathfrak{D} is an element ω such that $\mathfrak{D} = \mathbf{Z}[\omega]$. The ring of integers \mathfrak{D}_K of K is the unique maximal order in K . Given an order \mathfrak{D} , the (ideal) class group $\text{Cl}(\mathfrak{D})$ consists of the invertible ideals of \mathfrak{D} modulo principal ideals. The class of an ideal \mathfrak{a} is denoted $[\mathfrak{a}]$.

2.2 Quaternion algebras

To any prime number p , one can associate a quaternion algebra $B_{p,\infty}$ defined as

$$B_{p,\infty} = \mathbf{Q} + \mathbf{Q}i + \mathbf{Q}j + \mathbf{Q}ij,$$

with the multiplication rules $i^2 = -q$, $j^2 = -p$ and $ji = -ij$, where q is a positive integer that depends on p . More precisely, the latter is given by

$$q = \begin{cases} 1 & \text{if } p \equiv 3 \pmod{4}, \\ 2 & \text{if } p \equiv 5 \pmod{8}, \\ q_p & \text{if } p \equiv 1 \pmod{8}, \end{cases}$$

where q_p is the smallest prime such that $q_p \equiv 3 \pmod{4}$ and $\left(\frac{p}{q_p}\right) = -1$ (see [Piz80]). Assuming GRH, it follows from [LO77] that $q_p = O((\log p)^2)$, which can thus be computed in polynomial time in $\log p$. For the general theory of quaternion algebras, we refer the reader to [Vig06] or [Voi21].

Conjugation, trace and norm. Let $\alpha = x_1 + x_2i + x_3j + x_4ij$ be a generic element in $B_{p,\infty}$. The algebra $B_{p,\infty}$ has a canonical involution $\alpha \mapsto \bar{\alpha} = x_1 - x_2i - x_3j - x_4ij$. It induces the *reduced trace* and the *reduced norm*

$$\begin{aligned}\mathrm{Trd}(\alpha) &= \alpha + \bar{\alpha} = 2x_1, \text{ and} \\ \mathrm{Nrd}(\alpha) &= \alpha\bar{\alpha} = x_1^2 + qx_2^2 + p(x_3^2 + qx_4^2).\end{aligned}$$

The latter is a positive definite quadratic map, which makes $B_{p,\infty}$ a quadratic space, and endows its finitely generated \mathbf{Z} -submodules with a lattice structure.

Maximal orders. An *order* in $B_{p,\infty}$ is a full-rank lattice that is also a subring. It is *maximal* if it is not contained in any other order. For any full-rank lattice $\Lambda \subset B_{p,\infty}$, the *left order* and the *right order of Λ* are

$$\begin{aligned}\mathcal{O}_L(\Lambda) &= \{\alpha \in B_{p,\infty} \mid \alpha\Lambda \subseteq \Lambda\}, \text{ and} \\ \mathcal{O}_R(\Lambda) &= \{\alpha \in B_{p,\infty} \mid \Lambda\alpha \subseteq \Lambda\}.\end{aligned}$$

Let \mathcal{O} be an order. A lattice $I \subseteq \mathcal{O}$ is a left \mathcal{O} -ideal if $\mathcal{O}I \subseteq I$, or a right \mathcal{O} -ideal if $I\mathcal{O} \subseteq I$. If \mathcal{O} is a maximal order, and I is a left ideal in \mathcal{O} , then $\mathcal{O}_L(I) = \mathcal{O}$ and $\mathcal{O}_R(I)$ is another maximal order. Given two maximal orders \mathcal{O}_1 and \mathcal{O}_2 , their *connecting ideal* is the ideal

$$I(\mathcal{O}_1, \mathcal{O}_2) = \{\alpha \in B_{p,\infty} \mid \alpha\mathcal{O}_2\bar{\alpha} \subseteq [\mathcal{O}_2 : \mathcal{O}_1 \cap \mathcal{O}_2]\mathcal{O}_1\},$$

which satisfies $\mathcal{O}_L(I(\mathcal{O}_1, \mathcal{O}_2)) = \mathcal{O}_1$ and $\mathcal{O}_R(I(\mathcal{O}_1, \mathcal{O}_2)) = \mathcal{O}_2$. Let \mathcal{O} be a maximal order. Two left \mathcal{O} -ideals I and J are *equivalent* if there exists $\alpha \in B_{p,\infty}$ such that $I = \alpha J$. If I and J are equivalent, then $\mathcal{O}_R(I) \cong \mathcal{O}_R(J)$.

2.3 Elliptic curves

Recall that an elliptic curve is an abelian variety of dimension 1, isogenies are non-trivial morphisms between them, and endomorphisms are isogenies from a curve to itself. For a detailed reference on elliptic curves, we refer the reader to [Sil86].

Isogenies and endomorphisms. The set of all isogenies from E to E' (over the algebraic closure of the field of definition), together with the trivial map of kernel E , is written $\mathrm{Hom}(E, E')$. It forms a \mathbf{Z} -module for point-wise addition $+$. The *endomorphism ring* $\mathrm{End}(E)$ of an elliptic curve E is the \mathbf{Z} -module $\mathrm{Hom}(E, E)$ together with the composition of maps \circ . We have an embedding $\mathbf{Z} \hookrightarrow \mathrm{End}(E) : m \mapsto [m]$, where $[m]$ is the multiplication-by- m map. *Dividing* by m means finding a preimage through $[m]$, which in general is not unique.

The *degree* $\deg(\varphi)$ of an isogeny $\varphi : E \mapsto E'$ is the smallest positive element in $\mathbf{Z} \cap (\mathrm{Hom}(E', E) \circ \varphi)$. There is a unique isogeny $\hat{\varphi}$ such that $\hat{\varphi} \circ \varphi = [\deg(\varphi)]$, called the *dual* of φ . The degree is an integral quadratic map; it thereby endows

the \mathbf{Z} -module $\text{Hom}(E, E')$ with the structure of a lattice, with associated bilinear form

$$\langle \varphi, \psi \rangle = \frac{1}{2} \left(\hat{\varphi} \circ \psi + \hat{\psi} \circ \varphi \right).$$

If $\alpha \in \text{End}(E)$, we write $E[\alpha] = \ker(\alpha)$. We also write $E[m] = \ker([m])$ for $m \in \mathbf{Z}$, and $E[S] = \bigcap_{\alpha \in S} E[\alpha]$ for $S \subseteq \text{End}(E)$.

Efficient representation of isogenies. There are many ways to computationally represent isogenies. Rather than imposing a particular encoding, let us specify the required properties. As in [Wes22], we say that an isogeny $\varphi : E \rightarrow E'$ is given in an *efficient representation* if there is an algorithm to evaluate $\varphi(P)$ for any $P \in E(\mathbf{F}_{p^k})$ in time polynomial in the length of the representation of φ and in $k \log(p)$. We also assume that an efficient representation of φ has length $\Omega(\log(\deg(\varphi)))$. With these properties, the quadratic structure of $\text{Hom}(E, E')$ is computationally available, thanks to the following lemma.

Lemma 1. *Given $\varphi, \psi \in \text{Hom}(E, E')$ in efficient representation, one can compute $\langle \varphi, \psi \rangle$ in time polynomial in the length of the representation of φ and ψ , and in $\log p$.*

Proof. This is a straightforward generalisation of [EHL⁺18, Lemma 4]. □

Supersingular curves. Fix a prime number p . If E is an elliptic curve defined over $\overline{\mathbf{F}}_p$, it is supersingular if and only if its endomorphism ring $\text{End}(E)$ is isomorphic to a maximal order in the quaternion algebra $B_{p,\infty}$ (hence $B_{p,\infty} \cong \text{End}(E) \otimes \mathbf{Q}$). Up to $\overline{\mathbf{F}}_p$ -isomorphism, all supersingular elliptic curves over $\overline{\mathbf{F}}_p$ are defined over \mathbf{F}_{p^2} , and there are $\lfloor p/12 \rfloor + \varepsilon$ of them, with $\varepsilon \in \{0, 1, 2\}$.

To any isogeny $\varphi : E_1 \rightarrow E_2$, one associates a left $\text{End}(E_1)$ -ideal $I_\varphi = \text{Hom}(E_2, E_1)\varphi$, which satisfies $\text{End}(E_2) \cong \mathcal{O}_R(I_\varphi)$. Reciprocally, given a left $\text{End}(E_1)$ -ideal I , one can construct an isogeny $\varphi_I : E_1 \rightarrow E_2$ of kernel $E[I] = \bigcap_{\alpha \in I} \ker(\alpha)$ and degree $\text{Nrd}(I)$. These two constructions are mutual inverses, meaning that for any I and φ , we have $I_{\varphi_I} = I$ and $\varphi_{I_\varphi} = \varphi$ (up to an isomorphism of the target).

Remark 1. If I and J are coprime to p , we have $E[I] \subseteq E[J]$ if and only if $J \subseteq I$.

Remark 2. An equivalent definition of supersingular is that $E[p^n]$ is trivial for any n . Then, for any isogeny φ , the size of its kernel (i.e., its separable degree) is the largest factor of $\deg(\varphi)$ coprime with p . In particular, given an ideal I , we have that $\#E[I]$ is equal to the largest factor of $N(I)$ coprime with p .

2.4 Orientations

Let K be a quadratic number field, with ring of integers \mathfrak{O}_K , and let \mathfrak{O} be an arbitrary order in K .

Definition 1 (Orientation). A K -orientation on an elliptic curve E is an embedding $\iota : K \hookrightarrow \text{End}(E) \otimes \mathbf{Q}$. It is an \mathfrak{D} -orientation if $\iota(\mathfrak{D}) = \iota(K) \cap \text{End}(E)$. Such a pair (E, ι) is called an \mathfrak{D} -oriented elliptic curve, and we say that E is \mathfrak{D} -orientable.

Note that \mathfrak{D} -orientations as defined above correspond to the *primitive* \mathfrak{D} -orientations of [CK20]. If (E, ι) is an \mathfrak{D} -oriented elliptic curve, we will often consider ι as an embedding of \mathfrak{D} into $\text{End}(E)$ (which naturally extends to an embedding of K into $\text{End}(E) \otimes \mathbf{Q}$). This relates to the notion of primitive embedding.

Definition 2. Given two lattices Λ_1 and Λ_2 , an embedding $j : \Lambda_1 \hookrightarrow \Lambda_2$ is primitive if the group $\Lambda_2/j(\Lambda_1)$ is torsion-free.

Then, one can equivalently define the notion of \mathfrak{D} -orientation as a primitive embedding $\iota : \mathfrak{D} \hookrightarrow \text{End}(E)$.

Given a K -oriented elliptic curve (E, ι) , any isogeny $\varphi : E \rightarrow E'$ induces a K -orientation $\varphi_*(\iota)$ on E' defined as

$$\varphi_*(\iota)(\alpha) = (\varphi \circ \iota(\alpha) \circ \hat{\varphi}) \otimes \frac{1}{\deg(\varphi)}.$$

From an oriented curve (E, ι) , one can naturally define two others. First, its twist $(E, \bar{\iota})$, defined by $\bar{\iota}(\alpha) = \iota(\bar{\alpha})$, further studied in Section 6.1. Second, its Frobenius $(E, \bar{\iota})^{(p)} = (E^{(p)}, (\phi_p)_*(\bar{\iota}))$, where $\phi_p : E \rightarrow E^{(p)}$ is the Frobenius isogeny.

Definition 3 (Oriented isogeny). Given two K -oriented elliptic curves (E, ι) and (E', ι') , an isogeny $\varphi : (E, \iota) \rightarrow (E', \iota')$ is K -oriented if $\iota' = \varphi_*(\iota)$. If $\deg(\varphi)$ is prime, ι is an \mathfrak{D} -orientation and ι' an \mathfrak{D}' -orientation, then the isogeny is horizontal when $\mathfrak{D} = \mathfrak{D}'$, ascending when $\mathfrak{D} \subsetneq \mathfrak{D}'$, and descending when $\mathfrak{D} \supsetneq \mathfrak{D}'$. We say that an isogeny of composite degree is horizontal, ascending or descending if it factors as prime degree isogenies all of that type.

We write $\text{SS}_{\mathfrak{D}}(p)$ the set of \mathfrak{D} -oriented supersingular elliptic curves over $\overline{\mathbf{F}}_p$ up to K -oriented isomorphism. We write $(E, \iota) \cong (E', \iota')$ if there is a K -oriented isomorphism between them. Abusing notation, we often write $(E, \iota) \in \text{SS}_{\mathfrak{D}}(p)$ to mean that (E, ι) is a representative of an isomorphism class in $\text{SS}_{\mathfrak{D}}(p)$.

Proposition 1 ([Onu21, Proposition 3.2]). The set $\text{SS}_{\mathfrak{D}}(p)$ is not empty if and only if p does not split in K and does not divide the conductor of \mathfrak{D} .

Throughout, we suppose that p does not split in K and does not divide the conductor of \mathfrak{D} .

Encoding orientations. Computationally, an orientation is encoded by a generator ω of \mathfrak{D} (i.e., $\mathfrak{D} = \mathbf{Z}[\omega]$) together with an efficient representation of the endomorphism $\iota(\omega)$. We call this an *efficient representation* of ι .

Volcanoes. The following classification of horizontal, ascending or descending isogenies is commonly referred to as the *volcano* structure of oriented ℓ -isogeny graphs. Let $(E, \iota) \in \text{SS}_{\mathfrak{D}}(p)$, let $\ell \neq p$ be a prime, and let Δ be the discriminant of \mathfrak{D} . Let $\left(\frac{\Delta}{\ell}\right)$ be the Legendre symbol. From [CK20], the K -oriented isogenies of degree ℓ from (E, ι) are distributed as follows:

- There are $\ell - \left(\frac{\Delta}{\ell}\right)$ descending isogenies,
- If \mathfrak{D} is maximal at ℓ , there are $\left(\frac{\Delta}{\ell}\right) + 1$ horizontal, and no ascending isogeny.
- If \mathfrak{D} is non-maximal at ℓ , there is no horizontal, and one ascending isogeny.

This result implies that each connected component of the K -oriented ℓ -isogeny graph (with vertices $\text{SS}_{\mathfrak{D}}(p)$ and with edges for each isogeny of degree ℓ) consists of a cycle of horizontal isogenies (called the *surface*, or the *crater*), to each vertex of which is attached an infinite tree of vertical isogenies. The *volcano* terminology refers to the shape of this graph.

3 Class groups acting on sets of elliptic curves

Fix an oriented curve $(E, \iota) \in \text{SS}_{\mathfrak{D}}(p)$. An \mathfrak{D} -ideal \mathfrak{a} induces a subgroup

$$E[\mathfrak{a}] = \bigcap_{\alpha \in \mathfrak{a}} \ker(\iota(\alpha)),$$

and an isogeny $\varphi_{\mathfrak{a}} : E \rightarrow E^{\mathfrak{a}}$ of kernel $E[\mathfrak{a}]$ and degree $N(\mathfrak{a})$ called the \mathfrak{a} -multiplication. The target $E^{\mathfrak{a}}$ is the \mathfrak{a} -transform of (E, ι) . This construction induces an action of \mathfrak{D} -ideals on the set $\text{SS}_{\mathfrak{D}}(p)$, defined by

$$\mathfrak{a} \star (E, \iota) = (E^{\mathfrak{a}}, (\varphi_{\mathfrak{a}})_*(\iota)),$$

which factors through $\text{Cl}(\mathfrak{D})$. This action, well understood for ordinary elliptic curves with complex multiplication, was first studied in the context of oriented supersingular curves in [CK20] and [Onu21].

Theorem 1 ([Onu21]). *The action*

$$\text{Cl}(\mathfrak{D}) \times \text{SS}_{\mathfrak{D}}(p) \longrightarrow \text{SS}_{\mathfrak{D}}(p) : ([\mathfrak{a}], (E, \iota)) \longmapsto \mathfrak{a} \star (E, \iota)$$

is free and has at most two orbits. For any orbit A , and any $(E, \iota) \in \text{SS}_{\mathfrak{D}}(p)$, either $(E, \iota) \in A$, or both $(E, \bar{\iota})$ and $(E^{(p)}, \iota^{(p)})$ are in A .

Proof. This theorem combines [Onu21, Proposition 3.3] and [Onu21, Theorem 3.4]. The statement about $(E, \bar{\iota})$ is from the proof of [Onu21, Proposition 3.3]. \square

Computing the action. The image $\mathfrak{a} \star (E, \iota)$ can be computed in time polynomial in the length of the encoding of (E, ι) , in $\log(N(\mathfrak{a}))$ and in the largest prime-power factor of $N(\mathfrak{a})$. This is done by evaluating the action of $\iota(\mathfrak{a})$ on $E[N(\mathfrak{a})]$ to deduce $E[\mathfrak{a}]$, as in [CK20]. Evaluating the induced orientation $(\varphi_{\mathfrak{a}})_*(\iota)$ requires

a division by $\deg(\varphi_{\mathfrak{a}})$, which can be done in time polynomial in B if $N(\mathfrak{a})$ is B -powersmooth (meaning that all its prime-power factors are smaller than B). Yet, it should be noted that the efficiency of this representation degrades after applying the action of several such ideals, say n , because their product may only be B^n -powersmooth. This will not be an issue when only a constant number of actions are applied in this way (as in our forthcoming reductions), or when the endomorphism ring of the curves are known (in which case an efficient representation of the orientation can be recomputed, that does not depend on the ideal).

3.1 Computational problems

We now define two problems that translate whether or not the action is *one-way*.

Problem 1 (\mathfrak{D} -VECTORIZATION). Given $(E, \iota), (E', \iota') \in \text{SS}_{\mathfrak{D}}(p)$, find an \mathfrak{D} -ideal \mathfrak{a} such that $E' \cong E^{\mathfrak{a}}$.

Per Theorem 1, a solution \mathfrak{a} of \mathfrak{D} -VECTORIZATION always exists since the isomorphism $E' \cong E^{\mathfrak{a}}$ does not care about the orientation. The *vectorization* terminology comes from Couveignes' work [Cou06], even though proper use of this terminology should require $(E', \iota') \cong \mathfrak{a} \star (E, \iota)$. We will see that this modification makes little difference, as \mathfrak{D} -VECTORIZATION will turn out to be equivalent to the following stronger problem.

Problem 2 (EFFECTIVE \mathfrak{D} -VECTORIZATION). Given three \mathfrak{D} -oriented supersingular curves $(E, \iota), (E', \iota'), (F, j) \in \text{SS}_{\mathfrak{D}}(p)$, find an \mathfrak{D} -ideal \mathfrak{a} (or decide that it does not exist) such that $(E', \iota') \cong \mathfrak{a} \star (E, \iota)$, and an efficient representation of $\varphi_{\mathfrak{a}} : (F, j) \rightarrow \mathfrak{a} \star (F, j)$.

Now, a solution to EFFECTIVE \mathfrak{D} -VECTORIZATION does not necessarily exist, since (E, ι) and (E', ι') could be in the two distinct orbits described in Theorem 1. At first glance, the EFFECTIVE \mathfrak{D} -VECTORIZATION problem seems harder than \mathfrak{D} -VECTORIZATION for two reasons. First, an arbitrary ideal \mathfrak{a} is unlikely to induce an efficient representation of $\varphi_{\mathfrak{a}}$. This has already proved to be a serious obstacle in the literature [CPV20, GPSV21], where given an ideal class $[\mathfrak{a}]$, heavy work goes into finding a *good* representative of $[\mathfrak{a}]$ by a *smoothing* step which we do not know how to solve in polynomial time. Second, an isomorphism $E' \cong E^{\mathfrak{a}}$ does not imply that $(E', \iota') \cong \mathfrak{a} \star (E, \iota)$, and the information lost may be substantial as $h(\mathfrak{D})$ can be arbitrarily large while there are only approximately $p/12$ supersingular curves up to isomorphism. Despite these obstacles, we will show that these two problems are equivalent, by showing they are both equivalent to an oriented version of the endomorphism ring problem.

These two vectorisation problems are closely related to the following analog of the Diffie-Hellman problem.

Problem 3 (\mathfrak{D} -DIFFIEHELLMAN). Given an oriented curve $(E, \iota) \in \text{SS}_{\mathfrak{D}}(p)$, and its images $\mathfrak{a} \star (E, \iota)$ and $\mathfrak{b} \star (E, \iota)$ for the action of two unknown ideals \mathfrak{a} and \mathfrak{b} , compute $\mathfrak{a}\mathfrak{b} \star (E, \iota)$.

In Couveignes' terminology, this is the *parallelisation* problem. It is clear that both problems \mathfrak{O} -VECTORIZATION and \mathfrak{O} -DIFFIEHELLMAN reduce to the EFFECTIVE \mathfrak{O} -VECTORIZATION problem, and the converse reductions are the object of Sections 5, 6 and 7.

Now, one may consider the same problems when no orientation ι' for E' is provided. This modification seems to make the problems much harder, and this presumed hardness (for large $\text{disc}(\mathfrak{O})$) has been introduced in [DDF⁺21] as the *Uber isogeny assumption*.

Problem 4 (\mathfrak{O} -UBER). Given $(E, \iota) \in \text{SS}_{\mathfrak{O}}(p)$ and an \mathfrak{O} -orientable elliptic curve E' , find an \mathfrak{O} -ideal \mathfrak{a} such that $E' \cong E^{\mathfrak{a}}$.

The original Uber isogeny assumption from [DDF⁺21] also asks for an effective way to apply the action of \mathfrak{a} , as in the following effective variant.

Problem 5 (EFFECTIVE \mathfrak{O} -UBER). Given two \mathfrak{O} -oriented curves $(E, \iota), (F, j) \in \text{SS}_{\mathfrak{O}}(p)$ and an \mathfrak{O} -orientable curve E' , find an \mathfrak{O} -ideal \mathfrak{a} such that $E' \cong E^{\mathfrak{a}}$, and an efficient representation of $\varphi_{\mathfrak{a}} : (F, j) \rightarrow \mathfrak{a} \star (F, j)$.

The main interest for these *Uber* problems is that the security of most isogeny based cryptosystems reduce to them [DDF⁺21, Section 5], even systems such as SIDH [JD11] for which no class group action is immediately visible. For instance, SIKEp434 [JAC⁺17] could be broken by solving an instance of EFFECTIVE \mathfrak{O} -UBER with $\mathfrak{O} = \mathbf{Z} + 2^n \mathbf{Z}[\sqrt{-1}]$, for some $n \leq 217$.

Remark 3. Complexities will often be expressed as a function of the length of the input. Whenever the input contains an \mathcal{O} -orientable elliptic curve over \mathbf{F}_{p^2} , we assume that p and \mathcal{O} are part of the input, so its length is always $\Omega(\log p + \log |\text{disc}(\mathfrak{O})|)$. For \mathcal{O} -oriented curves, the length may be longer, depending on the quality of the provided orientations.

3.2 Some known or simple algorithms

Let us briefly present algorithms to solve some of the problems introduced above.

Proposition 2. *The EFFECTIVE \mathfrak{O} -UBER problem can heuristically be solved in expected time $l^{O(1)} |\text{disc}(\mathfrak{O})|^{1/2}$, with l the length of the input.*

Proof. This is the running time of an exhaustive search restricted to powersmooth ideals, with the estimate $\#(\text{Cl}(\mathfrak{O})) = O(\log(|\text{disc}(\mathfrak{O})|) |\text{disc}(\mathfrak{O})|^{1/2})$, as discussed in [DDF⁺21, Section 5.3]. \square

Proposition 3. *The \mathfrak{O} -VECTORIZATION problem can heuristically be solved in expected time $l^{O(1)} |\text{disc}(\mathfrak{O})|^{1/4}$, with l the length of the input.*

Proof. This is the running time of the meet-in-the-middle approach, as described for instance in [DG16, CLM⁺18], but using only powersmooth walks. \square

In the following, we use the classic notation

$$L_x(\alpha) = \exp\left(O\left((\log x)^\alpha (\log \log x)^{1-\alpha}\right)\right)$$

for subexponential complexities.

Proposition 4. *The \mathfrak{D} -VECTORIZATION problem can heuristically be solved in quantum subexponential time $l^{O(1)}L_{|\text{disc}(\mathfrak{D})|}(1/2)$, with l the length of the input.*

Proof. The \mathfrak{D} -VECTORIZATION problem reduces to the hidden shift problem with respect to the functions $f, f' : \text{Cl}(\mathfrak{D}) \rightarrow \mathbf{F}_{p^2}$ defined by $f([\mathfrak{a}]) = j(\mathfrak{a} \star (E, \iota))$ and $f'([\mathfrak{a}]) = j(\mathfrak{a} \star (E', \iota'))$. It only remains to prove that the action can be evaluated in quantum subexponential time, then apply Kuperberg's algorithm [Kup05]. It is tempting to simply adapt the method of [CJS14], but they find smooth class representatives, when we need powersmooth class representatives. We take a cruder, simpler, but heuristic route. One can randomize the class representative of $[\mathfrak{a}]$, until it has a powersmooth norm. The number of $L_x(1/2)$ -powersmooth numbers at most x is $L_x(1/2)$ (see [CN00, Section 3.1]), so under the heuristic assumption that norms of random class representatives behave like random integers of the same size, we may find an $L_d(1/2)$ -powersmooth representative in time $L_d(1/2)$, with $d = \text{disc}(\mathfrak{D})$. Its action can then be evaluated in time $L_d(1/2)$. \square

4 Oriented versions of the endomorphism ring problem

4.1 The endomorphism ring problem

To define the endomorphism ring problem in its strongest form, we introduce the notion of ε -basis, thereby unifying the two variants ENDRING and MAXORDER proved to be equivalent under the generalised Riemann hypothesis in [Wes22] (and heuristically since [EHL⁺18]).

Definition 4 (ε -basis). *Let $\varepsilon : B_{p,\infty} \rightarrow \text{End}(E) \otimes \mathbf{Q}$ be an isomorphism. Given a lattice $L \subseteq B_{p,\infty}$, an ε -basis of L is a pair (α, θ) where $(\alpha_i)_{i=1}^{\text{rank}(L)}$ is a basis of L and $\theta_i = \varepsilon(\alpha_i)$. Abusing language, we also call (α, θ) an ε -basis of the image lattice $\varepsilon(L)$.*

Remark 4. We will often talk about an ε -basis without specifying a priori an isomorphism ε . The ε is then implicit, and when L has full rank, it is uniquely determined by the ε -basis.

Encoding. Computationally, we suppose that elements in $B_{p,\infty}$ are encoded as vectors of rational numbers with respect to the basis $(1, i, j, ij)$. We assume that elements $\eta \otimes n^{-1}$ of $\text{End}(E) \otimes \mathbf{Q}$ are encoded as pairs (η, n) where n is an integer and η is an endomorphism in efficient representation.

The endomorphism ring problem can be defined as either finding a basis of a maximal order \mathcal{O} of $B_{p,\infty}$ isomorphic to $\text{End}(E)$, or finding four endomorphisms that generate $\text{End}(E)$. Both have their own advantages. On one hand, knowledge of the order \mathcal{O} enables the use of powerful algorithms for orders in quaternion algebras (notably, finding ideals connecting orders is much easier than finding isogenies connecting elliptic curves [KLPT14]). On the other hand, an explicit basis of $\text{End}(E)$ provides actual endomorphisms that can be evaluated. It was proved in [Wes22] that a basis of either \mathcal{O} or $\text{End}(E)$ can be transformed into an ε -basis of $\text{End}(E)$, assuming GRH. We therefore define the endomorphism ring problem as follows.

Problem 6 (ENDRING). Given a supersingular elliptic curve E over \mathbf{F}_{p^2} , find an ε -basis of $\text{End}(E)$.

Proposition 5. *The ENDRING problem can heuristically be solved in expected time $(\log p)^{O(1)}p^{1/2}$.*

Proof. This is the running time of the best known algorithms for ENDRING, for instance [DG16] and [EHL⁺20]. □

4.2 Oriented variants of the endomorphism ring problem

The ENDRING problem can naturally be restricted to \mathfrak{D} -orientable curves, resulting in the following problem.

Problem 7 (ENDRING $_{|\mathfrak{D}}$). Given an \mathfrak{D} -orientable E , find an ε -basis of $\text{End}(E)$.

Now, if an orientation is provided, we obtain the following variant.

Problem 8 (\mathfrak{D} -ENDRING). Given $(E, \iota) \in \text{SS}_{\mathfrak{D}}(p)$, find an ε -basis of $\text{End}(E)$.

One could require solutions to \mathfrak{D} -ENDRING to be in some way compatible with the orientation. It is unnecessary: as formalised in the following lemma, it is actually easy to express a given orientation in terms of a given ε -basis.

Lemma 2. *Given $(E, \iota) \in \text{SS}_{\mathfrak{D}}(p)$ and an ε -basis of $\text{End}(E)$, one can find an embedding $j : \mathfrak{D} \hookrightarrow B_{p,\infty}$ such that $\varepsilon \circ j = \iota$ in time polynomial in the length of the input.*

Proof. We can compute scalar products between endomorphisms with Lemma 1, so we can express the orientation ι in terms of the ε -basis. □

Finally, we consider the seemingly harder problem of computing the endomorphism ring *and* an orientation.

Problem 9 (\mathfrak{D} -ENDRING^{}).* Given a supersingular \mathfrak{D} -orientable curve E , find an ε -basis of $\text{End}(E)$ and an embedding $j : \mathfrak{D} \hookrightarrow B_{p,\infty}$ such that $\varepsilon \circ j$ is an \mathfrak{D} -orientation.

Clearly, \mathfrak{D} -ENDRING reduces to ENDRING $_{|\mathfrak{D}}$, which reduces to \mathfrak{D} -ENDRING^{*}.

4.3 Computing an orientation from the endomorphism ring

While $\text{ENDRING}|_{\mathfrak{D}}$ reduces to $\mathfrak{D}\text{-ENDRING}^*$, the converse boils down to the following question.

Question 1. Given an \mathfrak{D} -orientable curve E and an ε -basis of its endomorphism ring, can one compute an \mathfrak{D} -orientation of E in probabilistic polynomial time in $\log(p)$ and $\log|\text{disc}(\mathfrak{D})|$?

We only provide a positive answer to this question when the discriminant of the order \mathfrak{D} is small enough, in Proposition 6. The general case may be more difficult.

Proposition 6. *Given an \mathfrak{D} -orientable curve E and an ε -basis of its endomorphism ring, if $|\text{disc}(\mathfrak{D})| < 2p^{1/2} - 1$, then one can compute an \mathfrak{D} -orientation of E in probabilistic polynomial time in $\log(p)$.*

Proof. Let ι be an \mathfrak{D} -orientation of E . Let $d = \text{disc}(\mathfrak{D})$. Let ω be a Minkowski reduced generator of \mathfrak{D} (of trace either 0 or 1). Then, we have $\text{Nrd}(\omega) \leq (d+1)/4$. For any $\beta \in \text{End}(E)$, we have $|\text{disc}(\mathbf{Z}[\beta])| = 4\text{Nrd}(\beta) - \text{Trd}(\beta)^2 \leq 4\text{Nrd}(\beta)$. Also, for any $\beta \in \text{End}(E) \setminus \iota(\mathfrak{D})$, it follows from [Kan89, Theorem 2'] that $\text{disc}(\mathfrak{D}) \text{disc}(\mathbf{Z}[\beta]) \geq 4p$, hence

$$\text{Nrd}(\beta) \geq \frac{p}{d} > \frac{d+1}{4} \geq \text{Nrd}(\omega).$$

This proves that the shortest vector in $\text{End}(E) \setminus \mathbf{Z}$ is a generator of $\iota(\mathfrak{D})$, which can be recovered in polynomial time. Expressing this generator as a linear combination of the ε -basis of $\text{End}(E)$ provides an efficient representation of the orientation ι . \square

Proposition 6 has the following immediate consequence.

Corollary 1. *If $|\text{disc}(\mathfrak{D})| < 2p^{1/2} - 1$, then $\mathfrak{D}\text{-ENDRING}^*$ and $\text{ENDRING}|_{\mathfrak{D}}$ are equivalent.*

5 Endomorphism rings from orientations

In this section, we prove reductions from the family of endomorphism ring problems to the family of vectorisation problems. A key ingredient is the constructive Deuring correspondence in the ‘order-to-curve’ direction, Lemma 3, a result first heuristically proved in [EHL⁺18, Proposition 13]. Observing that it is easy to produce orders with a primitive embedding of \mathfrak{D} , we deduce in Lemma 4 that we can construct \mathfrak{D} -oriented elliptic curves of known endomorphism ring, to be used as starting points for vectorisation problems.

Lemma 3 (GRH). *There is an algorithm that given p and a maximal order $\mathcal{O} \subset B_{p,\infty}$, returns an elliptic curve E such that $\text{End}(E) \cong \mathcal{O}$ together with an ε -basis of $\text{End}(E)$, and runs in time polynomial in $\log(p)$ and the length of the basis of \mathcal{O} .*

Proof. From [EHL⁺18, Proposition 13] (but using [Wes22] instead of [KLPT14] to get rid of the heuristic assumptions), we get an elliptic curve E such that $\text{End}(E) \cong \mathcal{O}$. From [Wes22, Algorithm 6], we deduce an ε -basis of $\text{End}(E)$. \square

Definition 5. Let \mathcal{O} be a maximal order in an algebra $B \cong B_{p,\infty}$, and $j : \mathfrak{D} \hookrightarrow \mathcal{O}$ a primitive embedding. Let I be a left \mathcal{O} -ideal of prime norm ℓ , and let $\mathfrak{D}' = \mathcal{O}_R(I) \cap (j(\mathfrak{D}) \otimes \mathbf{Q})$. The ideal I is j -descending if $\mathfrak{D}' \subsetneq j(\mathfrak{D})$, j -horizontal if $\mathfrak{D}' = j(\mathfrak{D})$, and j -ascending if $\mathfrak{D}' \supsetneq j(\mathfrak{D})$.

Remark 5. It immediately follows that an \mathcal{O} -oriented isogeny $\varphi : (E, \iota) \rightarrow (E, \iota')$ of prime degree is descending (respectively horizontal or ascending) if and only if the kernel ideal $I_\varphi = \{\alpha \in \text{End}(E) \mid \ker \varphi \subseteq \ker \alpha\}$ is ι -descending (respectively ι -horizontal or ι -ascending). In particular, given \mathcal{O} and j , almost all ideals of norm ℓ are j -descending, with at most two exceptions [Onu21, Proposition 4.1]. Therefore, a j -descending ideal can be found in time polynomial in $\log(\ell)$ by listing three ideals of norm ℓ , and testing them using [Rón92, Theorem 3.2].

Lemma 4 (GRH). Given \mathfrak{D} and the factorisation of its discriminant, one can find some \mathfrak{D} -oriented curve $(E, \iota) \in \text{SS}_{\mathfrak{D}}(p)$ together with an ε -basis of $\text{End}(E)$ in probabilistic polynomial time in $\log(p)$ and $\log(|\text{disc}(\mathfrak{D})|)$.

Remark 6. There is a heuristic algorithm [LB20, Algorithm 1] that solves this task in the case where \mathfrak{D} is maximal. Our approach in the proof below is different, assumes only GRH, and avoids potentially hard factorisations.

Proof. We start by computing some arbitrary maximal order \mathcal{O}_0 in $B_{p,\infty}$ (for instance, a special order as in [KLPT14, Section 2.3]). Let K be the quadratic field containing \mathfrak{D} , with ring of integers \mathfrak{D}_K . Let ω_K be a Minkowski reduced generator of \mathfrak{D}_K , with minimal polynomial $x^2 - tx + n$. To find a quaternion $\omega \in B_{p,\infty}$ with the same minimal polynomial, solve

$$(t/2)^2 + qb^2 + p(c^2 + qd^2) = n$$

for $b, c, d \in \mathbf{Q}$ with [Sim06] (using the factorisation of $\text{disc}(\mathfrak{D})$). Let $\omega = t/2 + bi + cj + dij$, which has same norm and trace as ω_K , hence same minimal polynomial. Let a be the smallest integer such that $a\omega \in \mathcal{O}_0$. Let $I = \mathcal{O}_0\omega a + \mathcal{O}_0a$. We have $I\omega \subseteq I$, so $\omega \in \mathcal{O}_R(I)$, and $\mathcal{O}_R(I)$ can be computed with [Rón92, Theorem 3.2]. The corresponding embedding $j : \mathfrak{D}_K \hookrightarrow \mathcal{O}_R(I) : \omega_K \mapsto \omega$ is primitive since \mathfrak{D}_K is maximal, and it remains to descend to \mathfrak{D} . Let c be the conductor of \mathfrak{D} . For any prime power $\ell^k \parallel c$,

- let \tilde{J}_ℓ be any j -descending $\mathcal{O}_R(I)$ -ideal of norm ℓ (see Remark 5), and
- let $J_\ell \subseteq \tilde{J}_\ell$ an ideal of norm ℓ^k such that $J_\ell \not\subseteq \ell\mathcal{O}_R(I)$.

By construction, each J_ℓ is the kernel ideal of a cyclic isogeny of norm ℓ^k whose first step (hence all steps, because of the volcano structure) is descending. The cyclicity of each J_ℓ comes from $J_\ell \not\subseteq \ell\mathcal{O}_R(I)$ and Remark 1. Then, $J = \bigcap_{\ell|c} J_\ell$ is the kernel ideal of a descending isogeny of degree c , hence $\mathfrak{D} \hookrightarrow \mathcal{O}_R(J) : c\omega_K \mapsto c\omega$ is a primitive embedding. So we define $\mathcal{O} = \mathcal{O}_R(J)$. Applying Lemma 3, we can construct E with an ε -basis of $\text{End}(E) \cong \mathcal{O}$. The orientation ι is provided by the induced efficient representation of the endomorphism $\iota(c\omega_K) = \varepsilon(c\omega)$. \square

Proposition 7 (GRH). *Given the factorisation of $\text{disc}(\mathfrak{D})$, the \mathfrak{D} -ENDRING problem reduces to \mathfrak{D} -VECTORIZATION in probabilistic polynomial time in the length of the instance.*

Proof. Suppose we are given an instance $(E, \iota) \in \text{SS}_{\mathfrak{D}}(p)$ of \mathfrak{D} -ENDRING. Find $(E', \iota') \in \text{SS}_{\mathfrak{D}}(p)$ together with an ε -basis of $\mathcal{O}' \cong \text{End}(E')$ using Lemma 4. Solving \mathfrak{D} -VECTORIZATION, find an \mathfrak{D} -ideal \mathfrak{a} such that $E = E'^{\mathfrak{a}}$. Then, the kernel ideal $I = \mathcal{O}' \cdot \iota'(\mathfrak{a})$ of $\varphi_{\mathfrak{a}} : E' \rightarrow E$ is a connecting ideal between $\mathcal{O}' \cong \text{End}(E')$ and $\mathcal{O}_R(I) \cong \text{End}(E)$. The right-order $\mathcal{O}_R(I)$ can be computed with [Rón92, Theorem 3.2], thereby solving \mathfrak{D} -ENDRING for (E, ι) . \square

Proposition 8 (GRH). *Given the factorisation of $\text{disc}(\mathfrak{D})$, \mathfrak{D} -ENDRING* reduces to \mathfrak{D} -UBER in probabilistic polynomial time in the length of the instance.*

Proof. We proceed as in Proposition 7, except that no \mathfrak{D} -orientation of E is provided, which still allows us to apply \mathfrak{D} -UBER instead of \mathfrak{D} -VECTORIZATION. \square

6 Reducing vectorisation to endomorphism ring

It is shown in [CPV20] that in the particular case of curves defined over \mathbf{F}_p , and $\sqrt{-p} \in \mathfrak{D}$, solving the endomorphism ring problem allows one to solve the \mathfrak{D} -VECTORIZATION problem. They note however that in general, the resulting ideal does not necessarily have a smooth norm, so it is hard to compute its action. They conclude that this approach necessitates an expensive smoothing step, hence only provides a sub-exponential reduction of the security of CSIDH to the endomorphism ring problem. A priori, their methods seem very specific to the CSIDH setting, exploiting the action of Frobenius and quadratic twists. Introducing an appropriate generalisation of twisting, we prove in this section that \mathfrak{D} -VECTORIZATION reduces to \mathfrak{D} -ENDRING in all generality. Pushing the results farther, we circumvent the smoothness obstruction by proving polynomial time reductions between these problems and EFFECTIVE \mathfrak{D} -VECTORIZATION, observing that the action of non-smooth ideals can be efficiently evaluated on elliptic curves of known endomorphism ring. The idea that endomorphisms allow one to evaluate non-smooth isogenies had been observed in [FKM21], and the following proposition extends it to the action of ideals on oriented curves.

Proposition 9 (GRH). *Given $(E, \iota) \in \text{SS}_{\mathfrak{D}}(p)$, an ε -basis of $\text{End}(E)$, and an \mathfrak{D} -ideal \mathfrak{a} , one can compute $\mathfrak{a} \star (E, \iota)$ and an efficient representation of $\varphi_{\mathfrak{a}}$ in probabilistic polynomial time in the length of the input (i.e., $\log p, \log(|\text{disc}(\mathfrak{D})|)$), $\log(N(\mathfrak{a}))$ and the length of ι and of the ε -basis of $\text{End}(E)$).*

Proof. We are given an isomorphism $\varepsilon : \mathcal{O} \rightarrow \text{End}(E)$. Let $j : \mathfrak{D} \rightarrow \mathcal{O}$ such that $\varepsilon \circ j = \iota$ (see Lemma 2). Let $I = \mathcal{O}j(\mathfrak{a})$, and use [Wes22, Theorem 6.4] to find $\alpha \in I$ such that $J = I\bar{\alpha}/N(\mathfrak{a})$ has powersmooth norm. Then, J is

the kernel ideal of an efficiently computable isogeny φ_J (for instance by adapting [GPS20, Lemma 5] to the provided ε -basis instead of the special \mathcal{O}_0^1). We have $\varphi_{\mathfrak{a}} \circ [\mathrm{Nrd}(J)] = \varphi_J \circ \varepsilon(\alpha)$, so $E^{\mathfrak{a}} = \varphi_J(E)$. It only remains to compute the \mathfrak{D} -orientation on $E^{\mathfrak{a}}$, i.e., an efficient representation of $((\varphi_{\mathfrak{a}})_*(\iota))(\omega)$ for some generator ω of \mathfrak{D} . We have

$$((\varphi_{\mathfrak{a}})_*(\iota))(\omega) = (\varphi_{\mathfrak{a}} \circ \iota(\omega) \circ \hat{\varphi}_{\mathfrak{a}}) / N(\mathfrak{a}) = (\varphi_J \circ \varepsilon(\alpha) \circ \bar{\alpha}) \circ \hat{\varphi}_J / (\mathrm{Nrd}(\mathfrak{a}) \mathrm{Nrd}(J)^2).$$

Now, since $\alpha \in \mathcal{O}_J(\mathfrak{a})$, we have $\alpha \circ \bar{\alpha} \in \mathcal{O}N(\mathfrak{a})$. Let $\delta = (\alpha \circ \bar{\alpha}) / N(\mathfrak{a})$, so

$$((\varphi_{\mathfrak{a}})_*(\iota))(\omega) = (\varphi_J \circ \varepsilon(\delta) \circ \hat{\varphi}_J) / \mathrm{Nrd}(J)^2.$$

One can compute the quaternion $\delta \in \mathcal{O}$, express it as a linear combination of the ε -basis, and deduce an efficient representation of the numerator $\gamma = \varphi_J \circ \varepsilon(\delta) \circ \hat{\varphi}_J$. One can evaluate $((\varphi_{\mathfrak{a}})_*(\iota))(\omega)$ at any input point with [EHL⁺18, Algorithm 5], in polynomial time because the denominator $\mathrm{Nrd}(J)^2$ is powersmooth. Therefore we have an efficient representation of the orientation $(\varphi_{\mathfrak{a}})_*(\iota)$. \square

6.1 \mathfrak{D} -twists, a generalisation of quadratic twists

We now introduce the notion of \mathfrak{D} -twists, that enjoys properties similar to quadratic twists.

Definition 6 (\mathfrak{D} -twist). *We define the \mathfrak{D} -twisting involution as the map $\tau : \mathrm{SS}_{\mathfrak{D}}(p) \rightarrow \mathrm{SS}_{\mathfrak{D}}(p)$ defined as $\tau(E, \iota) = (E, \bar{\iota})$, where $\bar{\iota}(\alpha) = \iota(\bar{\alpha})$. The oriented curve $\tau(E, \iota)$ is the \mathfrak{D} -twist of (E, ι) .*

Lemma 5. *For any $(E, \iota) \in \mathrm{SS}_{\mathfrak{D}}(p)$ and \mathcal{O} -ideal \mathfrak{a} , we have $\tau(\mathfrak{a} \star (E, \iota)) = \bar{\mathfrak{a}} \star \tau(E, \iota)$.*

Proof. It follows from the fact that $\bigcap_{\alpha \in \bar{\mathfrak{a}}} \ker(\iota(\alpha)) = \bigcap_{\alpha \in \mathfrak{a}} \ker(\bar{\iota}(\alpha))$. \square

Recall that sets of isogenies $\mathrm{Hom}(E, E')$ (and in particular $\mathrm{End}(E)$) are lattices. They have a quadratic structure, hence an associated notion of orthogonality. Given $S \subseteq \mathrm{Hom}(E, E')$, we write $S^{\perp} \subseteq \mathrm{Hom}(E, E')$ the set of isogenies orthogonal to all elements of S .

The following lemma can be seen as an analog of [CPV20, Lemma 11].

Lemma 6. *Let $(E, \iota) \in \mathrm{SS}_{\mathfrak{D}}(p)$. For any non-zero $\theta \in \mathrm{End}(E)$, we have $\theta \in \iota(\mathfrak{D})^{\perp}$ if and only if $\theta_*(\iota) = \bar{\iota}$.*

Proof. First suppose $\theta \in \iota(\mathfrak{D})^{\perp}$. Since $1 \in \mathfrak{D}$, we have $\theta \in 1^{\perp}$, i.e., $\theta = -\hat{\theta}$. Then, for any $\alpha \in \mathfrak{D}$, we have $0 = \theta \circ \bar{\iota}(\alpha) + \iota(\alpha) \circ \hat{\theta} = \theta \circ \bar{\iota}(\alpha) - \iota(\alpha) \circ \theta$. We get

$$\theta_*(\iota)(\alpha) = (\theta \circ \iota(\alpha) \circ \hat{\theta}) \otimes \frac{1}{\deg(\theta)} = (\bar{\iota}(\alpha) \circ \theta \circ \hat{\theta}) \otimes \frac{1}{\deg(\theta)} = \bar{\iota}(\alpha),$$

¹ Alternatively, one can observe that the treatment of \mathcal{O}_0 in [GPS20, Lemma 5] is sufficient. Indeed, the ideal J is constructed in [Wes22, Theorem 6.4] as a composition of two \mathcal{O}_0 -ideals, which can each be translated to an isogeny via [GPS20, Lemma 5].

which proves the first implication. For the converse, suppose that $\theta_*(\iota) = \bar{\iota}$. Then, $\theta \circ \iota(\alpha) = \iota(\bar{\alpha}) \circ \theta$ for any $\alpha \in \mathfrak{D}$. Let $\omega \in \mathfrak{D}$ be a non-zero element of trace 0, so $\bar{\omega} = -\omega$. Write $\theta = x + \theta_0$ where $x \in \mathbf{Q}$ and $\hat{\theta}_0 = -\theta_0$ (concretely, $x = \text{Trd}(\theta)/2$ and $\theta_0 = (\theta - \hat{\theta})/2$). Then,

$$\begin{aligned} [x] \circ \iota(\omega) + \theta_0 \circ \iota(\omega) &= \theta \circ \iota(\omega) = \iota(\bar{\omega}) \circ \theta = -\iota(\omega) \circ \theta \\ &= -\iota(\omega) \circ [x] - \iota(\omega) \circ \theta_0, \end{aligned}$$

hence $[2x] \circ \iota(\omega) = \text{Trd}(\hat{\theta}_0 \circ \iota(\omega)) \in \mathbf{Q}$, which implies $x = 0$ because $\iota(\omega) \notin \mathbf{Q}$. This proves $\theta = \theta_0$, which is orthogonal to 1. The above also implies that $\theta \circ \iota(\omega) = -\iota(\omega) \circ \theta$, which means that θ is orthogonal to $\iota(\omega)$, so $\theta \in \iota(\mathfrak{D})^\perp$. \square

An issue with \mathfrak{D} -twisting is that it might not preserve $\text{Cl}(\mathfrak{D})$ -orbits. We introduce the following involution to resolve this.

Definition 7. We define the involution $\tau_p : \text{SS}_{\mathfrak{D}}(p) \rightarrow \text{SS}_{\mathfrak{D}}(p)$ as $\tau_p(E, \iota) = (E, \bar{\iota})^{(p)}$, with $(E, \bar{\iota})^{(p)} = (E^{(p)}, (\phi_p)_*(\bar{\iota}))$ where $\phi_p : E \rightarrow E^{(p)}$ is the Frobenius isogeny.

Proposition 10. There exists an ideal \mathfrak{a} such that $\tau_p(E, \iota) = \mathfrak{a} \star (E, \iota)$.

Proof. The only troublesome case is if $\text{SS}_{\mathfrak{D}}(p)$ partitions into two $\text{Cl}(\mathfrak{D})$ -orbits A and B . In that case, both τ and the Frobenius involution interchange A and B (Theorem 1). It follows that τ_p stabilizes A and B . \square

It still enjoys some of the useful properties of \mathfrak{D} -twisting.

Lemma 7. We have $\tau_p(\mathfrak{a} \star (E, \iota)) = \bar{\mathfrak{a}} \star \tau_p(E, \iota)$.

Proof. It follows from Lemma 5 and the fact that $\mathfrak{a} \star (E, \iota)^{(p)} = (\mathfrak{a} \star (E, \iota))^{(p)}$. \square

Corollary 2. An isogeny $\varphi : (E, \iota) \rightarrow \tau_p(E, \iota)$ is K -oriented if and only if $\varphi \in (\phi_p \circ \iota(\mathfrak{D}))^\perp$, where $\phi_p : E \rightarrow E^{(p)}$ is the Frobenius isogeny.

Proof. By the definition of τ_p , the isogeny φ is K -oriented if and only if $\varphi_*(\iota) = (\phi_p)_*(\bar{\iota})$. The latter is equivalent to $(\hat{\phi}_p)_*(\varphi_*(\iota)) = (\hat{\phi}_p)_*((\phi_p)_*(\bar{\iota}))$. Since

$$\begin{aligned} (\hat{\phi}_p)_*(\varphi_*(\iota)) &= (\hat{\phi}_p \circ \varphi)_*(\iota), \text{ and} \\ (\hat{\phi}_p)_*((\phi_p)_*(\bar{\iota})) &= [p]_*(\bar{\iota}) = \bar{\iota}, \end{aligned}$$

we deduce that $(\hat{\phi}_p \circ \varphi)_*(\iota) = \bar{\iota}$. From Lemma 6, this is equivalent to $\hat{\phi}_p \circ \varphi \in \iota(\mathfrak{D})^\perp$, i.e., $\varphi \in \phi_p \circ (\iota(\mathfrak{D}))^\perp = (\phi_p \circ \iota(\mathfrak{D}))^\perp$. \square

Corollary 3. The integral lattice $(\phi_p \circ \iota(\mathfrak{D}))^\perp \subset \text{Hom}(E, E^{(p)})$ is primitive (i.e., the greatest common divisor of the integers represented by the associated integral quadratic form is 1).

Proof. There exist two coprime ideals \mathfrak{a} and \mathfrak{b} such that $\tau_p(E, \iota) = \mathfrak{a} \star (E, \iota) = \mathfrak{b} \star (E, \iota)$ (see for instance [Cox11, Corollary 7.17]). We deduce from Corollary 2 that there are two lattice vectors $\varphi_{\mathfrak{a}}, \varphi_{\mathfrak{b}} \in (\phi_p \circ \iota(\mathfrak{D}))^\perp \subset \text{Hom}(E, E^{(p)})$ of coprime norm. \square

6.2 Vectorisation problems from endomorphism rings

We now prove that vectorisation problems reduce to endomorphism ring problems, with a strategy similar to that of [CPV20].

Lemma 8. *Let $(E, \iota) \in \text{SS}_{\mathfrak{D}}(p)$, a non-zero \mathfrak{D} -ideal \mathfrak{a} , and a left $\text{End}(E)$ -ideal I such that $E[\mathfrak{a}] = E[I]$. Then, $I \cap \iota(\mathfrak{D}) = \iota(\mathfrak{p}^k \mathfrak{a})$ for some $k \in \mathbf{Z}$, where \mathfrak{p} is the prime ideal above p .*

Proof. This lemma is a generalisation of [CPV20, Lemma 14], from which we adapt the proof. Write $\mathfrak{a} = \mathfrak{p}^i \mathfrak{b}$ where \mathfrak{b} is coprime to p . Similarly, write $I = I_p \cap J$, where $I_p = I + p^{\text{val}_p(\text{Nrd}(I))} \mathcal{O}$ is the p -part of I (i.e., $\text{Nrd}(I_p)$ is a power of p), and $\text{Nrd}(J)$ is not divisible by p . Following the observation of Remark 2, we have $\#E[\mathfrak{a}] = \#E[\mathfrak{b}]$ and $\#E[I] = \#E[J]$, hence

$$E[\mathfrak{b}] = E[\mathfrak{a}] = E[I] = E[J].$$

From the correspondence between ideals coprime to p and kernels of separable isogenies, we have

$$\begin{aligned} \mathfrak{b} &= \{\alpha \in \mathfrak{D} \mid E[\mathfrak{b}] \subseteq \ker(\iota(\alpha))\}, \text{ and} \\ J &= \{\theta \in \text{End}(E) \mid E[J] \subseteq \ker(\theta)\}. \end{aligned}$$

Together with $E[\mathfrak{b}] = E[J]$, we deduce that

$$J \cap \iota(\mathfrak{D}) = \{\theta \in \iota(\mathfrak{D}) \mid E[\mathfrak{b}] \subseteq \ker(\theta)\} = \iota(\mathfrak{b}).$$

There exists j such that $I_p \cap \iota(\mathfrak{D}) = \iota(\mathfrak{p}^j)$, hence $I \cap \iota(\mathfrak{D}) = \iota(\mathfrak{p}^{j-i} \mathfrak{a})$, proving the lemma. \square

Lemma 9. *A separable K -oriented isogeny of prime degree is horizontal if and only if it is of the form $\psi \circ \varphi_{\mathfrak{a}}$ where ψ is a K -isomorphism and \mathfrak{a} is an invertible ideal in an order of K .*

Proof. Let $\varphi : (E, \iota) \rightarrow (E', \iota')$ be a separable K -oriented isogeny of prime degree ℓ . It can be written as $\varphi = \psi \circ \varphi_0$ where $\varphi_0 : E \rightarrow E/\ker(\varphi)$ is the canonical projection, and ψ is an isomorphism. Suppose it is horizontal. Then, its kernel is of the form $E[\mathfrak{a}]$ with $N(\mathfrak{a}) = \ell$, so $\varphi_0 = \varphi_{\mathfrak{a}}$. We have

$$\psi_*((\varphi_{\mathfrak{a}})_*(\iota)) = (\psi \varphi_{\mathfrak{a}})_*(\iota) = \varphi_*(\iota) = \iota',$$

so $\psi : \mathfrak{a} \star (E, \iota) \rightarrow (E', \iota')$ is a K -isomorphism. The converse is clear. \square

Proposition 11 (GRH). *Given $(E, \iota) \in \text{SS}_{\mathfrak{D}}(p)$ and an ε -basis of $\text{End}(E)$, one can find \mathfrak{a} such that $\tau_p(E, \iota) = \mathfrak{a} \star (E, \iota)$ in probabilistic polynomial time in the length of the input.*

Proof. From Corollary 3, the lattice $\Lambda = (\phi_p \circ \iota(\mathfrak{D}))^\perp \cap \text{Hom}(E, E^{(p)})$ is primitive, so one can find $\varphi \in \Lambda$ of prime degree coprime to p and to the conductor of \mathfrak{D} (for instance with the algorithm [Wes22, Proposition 3.6], which ensures that $\deg(\varphi)$ is a large enough prime, assuming GRH). The algorithm returns

$$\mathfrak{a} = \iota^{-1}((\text{Hom}(E^{(p)}, E) \circ \varphi) \cap \iota(\mathfrak{D})).$$

It remains to show that this output is correct. From Corollary 2, we have that $\varphi : (E, \iota) \rightarrow \tau_p(E, \iota)$ is an \mathfrak{D} -isogeny, and it is horizontal. Applying Lemma 9, the isogeny φ is induced by an invertible \mathfrak{D} -ideal \mathfrak{b} . We have $E[\mathfrak{b}] = \ker(\varphi) = E[(\text{Hom}(E^{(p)}, E) \circ \varphi)]$, so from Lemma 8, we have $\iota(\mathfrak{a}) = \iota(\mathfrak{p}^k \mathfrak{b})$ for some $k \in \mathbf{Z}$. But $N(\mathfrak{b}) = \deg(\varphi)$ is coprime with p , and $\mathfrak{b} \subseteq \mathfrak{a}$, so $k = 0$. Therefore $\mathfrak{a} = \mathfrak{b}$, hence $\tau_p(E, \iota) = \mathfrak{b} \star (E, \iota) = \mathfrak{a} \star (E, \iota)$. \square

Lemma 10. *Suppose $\tau_p(E_1, \iota_1) = \mathfrak{a} \star (E_1, \iota_1)$ and $\tau_p(E_2, \iota_2) = \mathfrak{b} \star (E_2, \iota_2)$. Any ideal \mathfrak{c} such that $\mathfrak{c} \star (E_1, \iota_1) = (E_2, \iota_2)$ satisfies $[\mathfrak{c}^2] = [\mathfrak{a}\bar{\mathfrak{b}}]$.*

Proof. We have the chain of equalities

$$\begin{aligned} \bar{\mathfrak{b}}\mathfrak{a} \star (E_1, \iota_1) &= \bar{\mathfrak{b}} \star \tau_p(E_1, \iota_1) = \bar{\mathfrak{b}} \star \tau_p(\bar{\mathfrak{c}} \star (E_2, \iota_2)) \\ &= \mathfrak{c}\bar{\mathfrak{b}} \star \tau_p(E_2, \iota_2) = \mathfrak{c} \star (E_2, \iota_2) \\ &= \mathfrak{c}^2 \star (E_1, \iota_1), \end{aligned}$$

and we conclude from the fact that the action of the class group is free. \square

Theorem 2 (GRH, EFFECTIVE \mathfrak{D} -VECTORIZATION reduces to \mathfrak{D} -ENDRING).

Given \mathfrak{D} and the factorisation of its discriminant, three \mathfrak{D} -oriented elliptic curves $(E, \iota), (E', \iota'), (F, j) \in \text{SS}_{\mathfrak{D}}(p)$, together with ε -bases of $\text{End}(E), \text{End}(E')$ and $\text{End}(F)$, one can compute (or assert that it does not exist) an \mathfrak{D} -ideal \mathfrak{c} such that $(E', \iota') = \mathfrak{c} \star (E, \iota)$ and an efficient representation of $\varphi_{\mathfrak{c}} : (F, j) \rightarrow \mathfrak{c} \star (F, j)$ in probabilistic polynomial time in the length of the input and in $\#(\text{Cl}(\mathfrak{D})[2])$.

Proof. Suppose we are given $(E, \iota), (E', \iota') \in \text{SS}_{\mathfrak{D}}(p)$, together with $\text{End}(E)$ and $\text{End}(E')$. We can compute \mathfrak{a} and \mathfrak{b} such that $\tau_p(E, \iota) = \mathfrak{a} \star (E, \iota)$ and $\tau_p(E', \iota') = \mathfrak{b} \star (E', \iota')$ with Proposition 11. From Lemma 10, the ideal class of \mathfrak{c} is one of the $\#(\text{Cl}(\mathfrak{D})[2])$ square roots of $[\mathfrak{a}\bar{\mathfrak{b}}]$. They can be enumerated following [BS96, Section 6], and each of them can efficiently be checked for correctness with Proposition 9. Once the ideal \mathfrak{c} has been found, compute an efficient representation of $\varphi_{\mathfrak{c}} : (F, j) \rightarrow \mathfrak{c} \star (F, j)$ with Proposition 9. \square

Corollary 4 (GRH). *Given the factorisation of $\text{disc}(\mathfrak{D})$, EFFECTIVE \mathfrak{D} -UBER reduces to \mathfrak{D} -ENDRING* in probabilistic polynomial time in the length of the instance and in $\#(\text{Cl}(\mathfrak{D})[2])$.*

Proof. Suppose we are given $(E, \iota), (F, j) \in \text{SS}_{\mathfrak{D}}(p)$ and an \mathfrak{D} -orientable elliptic curve E' . Solving \mathfrak{D} -ENDRING*, one can find ε -bases of $\text{End}(E), \text{End}(F)$ and $\text{End}(E')$, and an \mathfrak{D} -orientation ι' of E' . The result follows from Theorem 2. \square

7 The oriented Diffie-Hellman problem

In this section, we study the relation of \mathfrak{D} -DIFFIEHELLMAN with other \mathfrak{D} -oriented problems, proving that it is essentially quantum-equivalent to the problem \mathfrak{D} -ENDRING. First, we have the following simple reduction from the problem \mathfrak{D} -DIFFIEHELLMAN to EFFECTIVE \mathfrak{D} -VECTORIZATION.

Proposition 12. *\mathfrak{D} -DIFFIEHELLMAN reduces to EFFECTIVE \mathfrak{D} -VECTORIZATION in probabilistic polynomial time in the length of the instance.*

Proof. Suppose we are given an oriented curve $(E, \iota) \in \text{SS}_{\mathfrak{D}}(p)$, and its images $\mathfrak{a} \star (E, \iota)$ and $\mathfrak{b} \star (E, \iota)$. Solving EFFECTIVE \mathfrak{D} -VECTORIZATION, one can recover the class of \mathfrak{a} , and apply its action on $\mathfrak{b} \star (E, \iota)$, thereby obtaining $(\mathfrak{a}\mathfrak{b}) \star (E, \iota)$. \square

Now, it remains to prove that the EFFECTIVE \mathfrak{D} -VECTORIZATION problem reduces to \mathfrak{D} -DIFFIEHELLMAN. In [GPSV21], it was proved that if the action of $\text{Cl}(\mathfrak{D})$ is efficiently computable, then (non-effective) \mathfrak{D} -VECTORIZATION reduces to \mathfrak{D} -DIFFIEHELLMAN in quantum polynomial time. Unfortunately, the action of $\text{Cl}(\mathfrak{D})$ is not efficiently computable in general, since only the action of *smooth* class representatives can be computed efficiently. Therefore the reduction does not run in polynomial time, and it does not apply to the effective variant of \mathfrak{D} -VECTORIZATION. We resolve both limitations. First, we prove that the \mathfrak{D} -VECTORIZATION problem does reduce to \mathfrak{D} -DIFFIEHELLMAN in quantum polynomial time because an oracle for \mathfrak{D} -DIFFIEHELLMAN provides an efficient way to evaluate the group action (by a trick similar to what is done in [GPSV21, Lemma 1]). Second, Proposition 7 and Theorem 2 immediately enhance the reduction from \mathfrak{D} -VECTORIZATION to a reduction from EFFECTIVE \mathfrak{D} -VECTORIZATION.

Theorem 3 (GRH). *\mathfrak{D} -VECTORIZATION reduces to \mathfrak{D} -DIFFIEHELLMAN in quantum polynomial time in the length of the instance.*

Proof. This is essentially an application of a generalisation of Shor's algorithm for the discrete logarithm problem [Sho97], with the observation that an oracle for \mathfrak{D} -DIFFIEHELLMAN makes the implicit group structure of a $\text{Cl}(\mathfrak{D})$ -orbit efficiently computable. More precisely, let $(E, \iota) \in \text{SS}_{\mathfrak{D}}(p)$, and $(E', \iota') = \mathfrak{a} \star (E, \iota)$ be an instance of \mathfrak{D} -VECTORIZATION. From [Bac90], assuming GRH, there is a bound B polynomial in $\log(\text{disc}(\mathfrak{D}))$ such that $\mathfrak{B} = \{\mathfrak{p} \mid N(\mathfrak{p}) < B \text{ is prime}\}$ is a generating set of the group $\text{Cl}(\mathfrak{D})$. Let

$$f : \mathbf{Z}^{\mathfrak{B}} \times \mathbf{Z} \longrightarrow \text{SS}_{\mathfrak{D}}(p)$$

$$((e_{\mathfrak{p}})_{\mathfrak{p}}, k) \longmapsto \left(\mathfrak{a}^k \cdot \prod_{\mathfrak{p}} \mathfrak{p}^{e_{\mathfrak{p}}} \right) \star (E, \iota).$$

We prove below that one can evaluate f in quantum polynomial time. Suppose we also have an efficiently computable injective function $\text{enc} : \text{SS}_{\mathfrak{D}}(p) \rightarrow \{0, 1\}^*$. Then, solving the Abelian hidden subgroup problem (see [Mos08]) for $\text{enc} \circ f$ leads to a vector $(e_{\mathfrak{p}})_{\mathfrak{p}}$ such that $\mathfrak{a} \sim \prod_{\mathfrak{p}} \mathfrak{p}^{e_{\mathfrak{p}}}$. One can then return any ideal in the class of $\prod_{\mathfrak{p}} \mathfrak{p}^{e_{\mathfrak{p}}}$ (thereby solving the \mathfrak{D} -VECTORIZATION instance).

Unique encoding of oriented curves. We now define the function `enc`. Any input $(F, j) \in \text{SS}_{\mathfrak{D}}(p)$ is assumed to be in efficient representation; the issue is that this representation is not unique. Let ω be a Minkowski-reduced generator of \mathfrak{D} . To find a unique encoding $\text{enc}(F, j)$, one can first turn F into some canonical form, such as $F_{\text{can}} : y^2 + xy = x^3 - (36x + 1)/(j(F) - 1728)$ [Sil86, page 47]. For ease of exposition, assume that all curves under consideration are already in canonical form $F = F_{\text{can}}$. From there, one can, in a deterministic way, find a point $P \in F$ of order larger than $4N(\omega)$. Since (F, j) is in efficient representation, one can compute $Q = \iota(\omega)(P)$ in polynomial time. The triple (F, P, Q) is a unique encoding of the K -isomorphism class² of (F, j) . Indeed, suppose that (F, j) and (F, j') are both encoded as (F, P, Q) . Let $\alpha = j(\omega) - j'(\omega)$. We have $\alpha(P) = Q - Q = 0$, so if $\alpha \neq 0$, we get

$$4N(\omega) < \text{ord}(P) \leq \deg(\alpha) \leq \left(\deg(j(\omega))^{1/2} + \deg(j'(\omega))^{1/2} \right)^2 = 4N(\omega),$$

a contradiction, so $\alpha = 0$ and $j(\omega) = j'(\omega)$.

Evaluating f . It remains to prove that f can indeed be computed in polynomial time. This is only feasible thanks to the \mathfrak{D} -DIFFIEHELLMAN oracle \mathcal{O} , which makes the implicit group multiplication \odot on the orbit $\text{Cl}(\mathfrak{D}) \star (E, \iota)$ efficiently computable, as

$$(\mathfrak{b} \star (E, \iota)) \odot (\mathfrak{c} \star (E, \iota)) = (\mathfrak{bc}) \star (E, \iota) = \mathcal{O}((E, \iota), \mathfrak{b} \star (E, \iota), \mathfrak{c} \star (E, \iota)).$$

Therefore, given any $k \geq 0$ and $\mathfrak{b} \star (E, \iota)$, one can compute $\mathfrak{b}^k \star (E, \iota) = (\mathfrak{b} \star (E, \iota))^{\odot k}$ by square-and-multiply. Then, for any $(e_{\mathfrak{p}})_{\mathfrak{p}} \in \mathbf{Z}_{\geq 0}^{\mathfrak{B}}$ and $k \geq 0$, one can efficiently compute

$$f((e_{\mathfrak{p}})_{\mathfrak{p}}, k) = \left(\bigotimes_{\mathfrak{p}} (\mathfrak{p} \star (E, \iota))^{\odot e_{\mathfrak{p}}} \right) \odot (E', \iota')^{\odot k},$$

given the oriented curves $\mathfrak{p} \star (E, \iota)$. Since each \mathfrak{p} has small norm, these $\mathfrak{p} \star (E, \iota)$ can be computed in polynomial time. To deal with negative exponents, we note that the class number $h(\mathfrak{D})$ is computable in quantum polynomial time [BS16], so all exponents can be reduced modulo $h(\mathfrak{D})$. \square

Corollary 5 (GRH). *The problem CSIDH-DIFFIEHELLMAN of recovering CSIDH shared secrets reduces to the problem $\text{ENDRING}_{\mathbf{F}_p}$ of computing the full endomorphism ring of supersingular elliptic curves defined over \mathbf{F}_p , under a probabilistic polynomial time reduction in $\log p$. Conversely, $\text{ENDRING}_{\mathbf{F}_p}$ reduces to CSIDH-DIFFIEHELLMAN in quantum polynomial time in $\log p$.*

Proof. The problem CSIDH-DIFFIEHELLMAN is equal to \mathfrak{D} -DIFFIEHELLMAN for some order \mathfrak{D} containing $\sqrt{-p}$. There are at most two possibilities for \mathfrak{D} :

² In the rare case where F has non-trivial automorphisms $\{\sigma\}$, we have that $(F, \sigma_*\iota)$ are all K -isomorphic, so one should replace Q with the set $\{(\sigma_*\iota)(\omega)(P)\}$.

either $\mathbf{Z}[\sqrt{-p}]$ or $\mathbf{Z}[(1 + \sqrt{-p})/2]$. The latter is only possible when $p \equiv 3 \pmod{4}$. They differ by an index 2, and correspond to CSIDH on the *floor* or on the *surface* — see CSURF [CD20].

In either case, for such an order \mathfrak{D} , we now prove that $\text{ENDRING}_{\mathfrak{D}}$ reduces to $\text{ENDRING}_{\mathbf{F}_p}$. Let E be an $\text{ENDRING}_{\mathfrak{D}}$ -instance. Since E is \mathfrak{D} -orientable, there exists $\alpha \in \text{End}(E)$ of degree p . Since E is supersingular, α is purely inseparable, so it factors as $\alpha = \beta \circ \phi_p^E$ where $\phi_p^E : E \rightarrow E^{(p)}$ is the Frobenius and $\beta : E^{(p)} \rightarrow E$ is an isomorphism. This proves that $E \cong E^{(p)}$, hence $j(E) \in \mathbf{F}_p$ and one can compute an isomorphism $\gamma : E \rightarrow E'$ to a curve E' defined over \mathbf{F}_p . Therefore, the $\text{ENDRING}_{\mathfrak{D}}$ -instance E reduces to the $\text{ENDRING}_{\mathbf{F}_p}$ -instance E' .

Finally, we prove that $\text{ENDRING}_{\mathbf{F}_p}$ reduces to \mathfrak{D} - ENDRING . If E is defined over \mathbf{F}_p , then either $\iota : \sqrt{-p} \rightarrow \phi_p^E$ is an \mathfrak{D} -orientation on E , or there exists an isogeny $\varphi : E \rightarrow E'$ of degree 2 such that $\varphi_*(\iota)$ is an \mathfrak{D} -orientation on E' (see [DG16, Theorem 2.7]). So $\text{ENDRING}_{\mathbf{F}_p}$ for E reduces to \mathfrak{D} - ENDRING either for E or for one of its three 2-neighbours.

These new reductions at hand, the corollary follows from the other reductions summarised in Figure 1, given that the factorisation of $\text{disc}(\mathfrak{D})$ is either $-p$ or $-4p$, and from genus theory, $\#(\text{Cl}(\mathfrak{D})[2]) \leq 2$. \square

8 The case of non-maximal orders

The OSIDH cryptosystem [CK20] exploits elliptic curves oriented by an order of the form $\mathbf{Z} + \ell^e \mathfrak{D}$, where ℓ is a small prime, and \mathfrak{D} has small discriminant. It is observed however that with such parameters, the $(\mathbf{Z} + \ell^e \mathfrak{D})$ -VECTORIZATION problem is not hard, hence it would be unsafe for the protocol to provide full efficient encodings of $(\mathbf{Z} + \ell^e \mathfrak{D})$ -orientations. In this section, we generalise this fact and study its consequences for relevant variants of the endomorphism ring problem.

Lemma 11. *Let c be a positive integer. Given $(E, \iota) \in \text{SS}_{\mathbf{Z} + c\mathfrak{D}}(p)$ in efficient representation, one can compute the kernel of an isogeny $\varphi : E \rightarrow E'$ of degree c such that $\varphi_*(\iota)$ is an \mathfrak{D} -orientation of E' in probabilistic polynomial time in the length of the input, the largest prime factor of c , and, for each $\ell^e \parallel c$, the degree of the extension of \mathbf{F}_p over which $E[\ell^e]$ is defined.*

Proof. Let $\varphi : (E, \iota) \rightarrow (E', \iota')$ be the unique ascending K -isogeny of degree c , where $\iota' = \varphi_*(\iota)$ is an \mathfrak{D} -orientation. We are given a generator ω of $\mathbf{Z} + c\mathfrak{D}$ and an efficient representation of $\iota(\omega)$. The generator is of the form $\omega = a + c\omega_0$ where ω_0 is a generator of \mathfrak{D} and without loss of generality, $a = 0$. We have

$$\iota(\omega) = (\hat{\varphi}_*(\iota'))(\omega) = \frac{\hat{\varphi} \circ \iota'(\omega) \circ \varphi}{c} = \hat{\varphi} \circ \iota'(\omega_0) \circ \varphi.$$

It implies $\ker(\varphi) \subseteq \ker(\iota(\omega))$. Now $\ker(\iota(\omega))$ is cyclic (otherwise $\iota(\omega)$ would be divisible by an integer, hence ι would not be a primitive embedding), so $\ker(\varphi) = \ker(\iota(\omega)) \cap E[c] = \bigcap_{\ell^e \parallel c} (\ker(\iota(\omega)) \cap E[\ell^e])$ can be recovered in time polynomial in the largest prime factor of c , and, for each $\ell^e \parallel c$, in the degree of the extension of \mathbf{F}_p over which $E[\ell^e]$ is defined. \square

In particular, if c is smooth and $E[c]$ is defined over a small extension of \mathbf{F}_p , given $(E, \iota) \in \text{SS}_{\mathbf{Z}+c\mathfrak{D}}(p)$, it is easy to find the unique isogeny ascending to an \mathfrak{D} -orientable curve. This is the crux of so-called *torsion point attacks* on SIDH-like cryptosystems [Pet17,KMP⁺21], which can be reinterpreted as an attempt to recover a $(\mathbf{Z}+c\mathfrak{D})$ -orientation from some information on the action of isogenies on torsion points, and some carefully chosen \mathfrak{D} .

Lemma 12 (GRH). *Given the kernel of an isogeny $\varphi : E \rightarrow E'$, and an ε -basis of $\text{End}(E)$, one can compute an ε -basis of $\text{End}(E')$ in probabilistic polynomial time in the length of the input and the largest prime factor of $\deg(\varphi)$.*

Proof. This statement seems folklore; we give a proof for completeness. Let us describe a simple (and certainly not optimal) algorithm. We may assume that $\ker(\varphi)$ is cyclic. Choose any prime $\ell \mid \deg(\varphi)$, and let $\varphi_1 : E \rightarrow E_1 = E/(\ker(\varphi) \cap E[\ell])$. One can compute the left $\text{End}(E)$ -ideals I_1 of norm ℓ corresponding to φ_1 (for instance with an exhaustive search among the $\ell + 1$ possibilities, checking each guess by evaluating a basis on the kernel of φ_1). Now, $\mathcal{O}_R(I_1) \cong \text{End}(E_1)$, and one can find an ε -basis of $\text{End}(E_1)$ (for instance with [Wes22, Algorithm 6]). The isogeny φ factors as $\varphi' \circ \varphi_1$ with $\deg(\varphi') = \deg(\varphi)/\ell$, and one can iterate the procedure. \square

Theorem 4 (GRH). *$(\mathbf{Z}+c\mathfrak{D})$ -ENDRING reduces to \mathfrak{D} -ENDRING* in probabilistic polynomial time polynomial in the length of the input, the largest prime factor of c , and, for each $\ell^e \parallel c$, the degree of the extension of \mathbf{F}_p over which $E[\ell^e]$ is defined.*

Proof. Let $(E, \iota) \in \text{SS}_{\mathbf{Z}+c\mathfrak{D}}(p)$ be an instance of $(\mathbf{Z}+c\mathfrak{D})$ -ENDRING. From Lemma 11, we can compute (within the claimed running time) an isogeny $\varphi : E \rightarrow E'$ where E' is \mathfrak{D} -orientable. One can solve \mathfrak{D} -ENDRING* for E' to find an ε -basis of $\text{End}(E')$. Now, Lemma 12 allows us to find an ε -basis of $\text{End}(E)$ thanks to the ε -basis of $\text{End}(E')$ and the kernel of $\hat{\varphi}$. \square

Lemma 13. *The \mathfrak{D} -ENDRING* problem can heuristically be solved in expected time $l^{O(1)} |\text{disc}(\mathfrak{D})|^{1/2}$, where l is the length of the input.*

Proof. By Proposition 8, one can reduce \mathfrak{D} -ENDRING* to \mathfrak{D} -UBER in time polynomial in $\log p$ and $\log |\text{disc}(\mathfrak{D})|$. Then, one can solve \mathfrak{D} -UBER with Proposition 2, under the same heuristics. \square

Corollary 6. *$(\mathbf{Z}+c\mathfrak{D})$ -ENDRING can heuristically be solved in probabilistic polynomial time in the length of the input, $|\text{disc}(\mathfrak{D})|$, the largest prime factor of c , and, for each $\ell^e \parallel c$, the degree of the extension of \mathbf{F}_p over which $E[\ell^e]$ is defined.*

Proof. It immediately follows from Theorem 4 and Lemma 13. \square

This corollary implies that if \mathfrak{D} has small discriminant and c is powersmooth, then knowledge of a $(\mathbf{Z}+c\mathfrak{D})$ -orientation leaks the whole endomorphism ring.

Theorem 5 (GRH). *Suppose c is $(\log p)^{O(1)}$ -powersmooth. Then, the problem $(\mathbf{Z} + c\mathfrak{D})$ -ENDRING reduces to \mathfrak{D} -ENDRING in probabilistic polynomial time in the length of the input.*

Proof. We proceed as in the proof of Theorem 4, but reducing to \mathfrak{D} -ENDRING using the \mathfrak{D} -orientation $\varphi_*(\iota)$ on E' . We only need to prove that we have an efficient representation for $\varphi_*(\iota)$. For any $\alpha \in \mathfrak{D}$, we can write

$$\varphi_*(\iota)(\alpha) = (\varphi \circ \iota(c\alpha) \circ \hat{\varphi}) / c^2.$$

We have $c\alpha \in \mathbf{Z} + c\mathfrak{D}$, so the efficient representation of ι implies that we have an efficient representation of the numerator $\varphi \circ \iota(c\alpha) \circ \hat{\varphi}$. The powersmoothness of c allows one to divide by c^2 . \square

9 Acknowledgements

This work was supported by the Agence Nationale de la Recherche under grants ANR MELODIA (ANR-20-CE40-0013) and ANR CIAO (ANR-19-CE48-0008). The author would like to thank Katherine E. Stange, Pierrick Dartois, and Jean-François Biasse for valuable discussions, feedback and corrections, that helped improve the quality of this work.

References

- Bac90. Eric Bach. Explicit bounds for primality testing and related problems. *Mathematics of Computation*, 55(191):355–380, 1990.
- BKV19. Ward Beullens, Thorsten Kleinjung, and Frederik Vercauteren. CSI-FiSh: Efficient isogeny based signatures through class group computations. In Steven D. Galbraith and Shiho Moriai, editors, *Advances in Cryptology - ASIACRYPT 2019 - 25th International Conference on the Theory and Application of Cryptology and Information Security*, volume 11921 of *Lecture Notes in Computer Science*, pages 227–247. Springer, 2019.
- BS96. Wieb Bosma and Peter Stevenhagen. On the computation of quadratic 2-class groups. *Journal de théorie des nombres de Bordeaux*, 8(2):283–313, 1996.
- BS16. J.-F. Biasse and F. Song. Efficient quantum algorithms for computing class groups and solving the principal ideal problem in arbitrary degree number fields. In Robert Krauthgamer, editor, *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms - SODA 2016*, pages 893–902. SIAM, 2016.
- CD20. Wouter Castryck and Thomas Decru. CSIDH on the surface. In *PQCrypto 2020 - International Conference on Post-Quantum Cryptography*, volume 12100, pages 111–129. Springer, 2020.
- CJS14. Andrew Childs, David Jao, and Vladimir Soukharev. Constructing elliptic curve isogenies in quantum subexponential time. *Journal of Mathematical Cryptology*, 8(1):1–29, 2014.
- CK20. Leonardo Colò and David Kohel. Orienting supersingular isogeny graphs. *Journal of Mathematical Cryptology*, 14(1):414–437, 2020.

- CLG09. Denis X. Charles, Kristin E. Lauter, and Eyal Z. Goren. Cryptographic hash functions from expander graphs. *Journal of Cryptology*, 22(1):93–113, Jan 2009.
- CLM⁺18. Wouter Castryck, Tanja Lange, Chloe Martindale, Lorenz Panny, and Joost Renes. CSIDH: an efficient post-quantum commutative group action. In Thomas Peyrin and Steven D. Galbraith, editors, *Advances in Cryptology - ASIACRYPT 2018 - 24th International Conference on the Theory and Application of Cryptology and Information Security*, volume 11274 of *Lecture Notes in Computer Science*, pages 395–427. Springer, 2018.
- CN00. Jean-Sébastien Coron and David Naccache. Security analysis of the Gennaro-Halevi-Rabin signature scheme. In Bart Preneel, editor, *Advances in Cryptology - EUROCRYPT 2000*, volume 1807 of *Lecture Notes in Computer Science*, pages 91–101. Springer, 2000.
- Cou06. Jean Marc Couveignes. Hard homogeneous spaces. IACR Cryptology ePrint Archive, Report 2006/291, 2006. <https://eprint.iacr.org/2006/291>.
- Cox11. David A Cox. *Primes of the form $x^2 + ny^2$: Fermat, class field theory, and complex multiplication*, volume 34. John Wiley & Sons, 2011.
- CPV20. Wouter Castryck, Lorenz Panny, and Frederik Vercauteren. Rational isogenies from irrational endomorphisms. In Anne Canteaut and Yuval Ishai, editors, *Advances in Cryptology - EUROCRYPT 2020 - 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques*, volume 12106 of *Lecture Notes in Computer Science*, pages 523–548. Springer, 2020.
- CS21. Mathilde Chenu and Benjamin Smith. Higher-degree supersingular group actions. In *MathCrypt 2021 - Mathematical Cryptology*, 2021.
- DDF⁺21. Luca De Feo, Cyprien Delpech de Saint Guilhem, Tako Boris Fouotsa, Péter Kutas, Antonin Leroux, Christophe Petit, Javier Silva, and Benjamin Wesolowski. Séta: Supersingular encryption from torsion attacks. To appear in *Advances in Cryptology - ASIACRYPT 2021*, 2021.
- DG16. Christina Delfs and Steven D Galbraith. Computing isogenies between supersingular elliptic curves over \mathbb{F}_p . *Designs, Codes and Cryptography*, 78(2):425–440, 2016.
- EHL⁺18. Kirsten Eisenträger, Sean Hallgren, Kristin Lauter, Travis Morrison, and Christophe Petit. Supersingular isogeny graphs and endomorphism rings: Reductions and solutions. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology - EUROCRYPT 2018*, pages 329–368, Cham, 2018. Springer International Publishing.
- EHL⁺20. Kirsten Eisenträger, Sean Hallgren, Chris Leonardi, Travis Morrison, and Jennifer Park. Computing endomorphism rings of supersingular elliptic curves and connections to path-finding in isogeny graphs. *ANTS XIV: Proceedings of the Fourteenth Algorithmic Number Theory Symposium, Open Book Series*, 4(1):215–232, 2020.
- FKM21. Tako Boris Fouotsa, Péter Kutas, and Simon-Philipp Merz. On the isogeny problem with torsion point information. IACR Cryptology ePrint Archive, Report 2021/153, 2021. <https://eprint.iacr.org/2021/153>.
- 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, 2020.
- GPST16. Steven D. Galbraith, Christophe Petit, Barak Shani, and Yan Bo Ti. On the security of supersingular isogeny cryptosystems. In Jung Hee Cheon

- and Tsuyoshi Takagi, editors, *Advances in Cryptology - ASIACRYPT 2016 - 22nd International Conference on the Theory and Application of Cryptology and Information Security*, volume 10031 of *Lecture Notes in Computer Science*, pages 63–91, 2016.
- GPSV21. Steven Galbraith, Lorenz Panny, Benjamin Smith, and Frederik Vercauteren. Quantum equivalence of the DLP and CDHP for group actions. *Mathematical Cryptology*, 1(1):40–44, 2021.
- JAC⁺17. David Jao, Reza Azarderakhsh, Matt Campagna, Craig Costello, Luca De Feo, Basil Hess, Amir Jalili, Brian Koziel, Brian Lamacchia, Patrick Longa, et al. SIKE: Supersingular isogeny key encapsulation. 2017.
- JD11. David Jao and Luca De Feo. Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. In Bo-Yin Yang, editor, *International Workshop on Post-Quantum Cryptography – PQCrypto 2011*, pages 19–34, 2011.
- Kan89. Masanobu Kaneko. Supersingular j -invariants as singular moduli mod p . *Osaka Journal of Mathematics*, 26(4):849–855, 1989.
- KLPT14. David Kohel, Kristin Lauter, Christophe Petit, and Jean-Pierre Tignol. On the quaternion ℓ -isogeny path problem. *LMS Journal of Computation and Mathematics*, 17(A):418–432, 2014.
- KMP⁺21. Péter Kutas, Chloe Martindale, Lorenz Panny, Christophe Petit, and Katherine E. Stange. Weak instances of SIDH variants under improved torsion-point attacks. In *to appear in Advances in Cryptology - CRYPTO 2021*, Lecture Notes in Computer Science, 2021.
- Kup05. Greg Kuperberg. A subexponential-time quantum algorithm for the dihedral hidden subgroup problem. *SIAM Journal on Computing*, 35(1):170–188, 2005.
- LB20. Jonathan Love and Dan Boneh. Supersingular curves with small noninteger endomorphisms. *ANTS XIV: Proceedings of the Fourteenth Algorithmic Number Theory Symposium, Open Book Series*, 4(1):7–22, 2020.
- LO77. J.C. Lagarias and A.M. Odlyzko. Effective versions of the Chebotarev density theorem. In *Algebraic number fields: L-functions and Galois properties (Proc. Sympos., Univ. Durham, Durham, 1975)*, pages 409–464. Academic Press, London, 1977.
- Mos08. Michele Mosca. Abelian hidden subgroup problem. In *Encyclopedia of Algorithms*, pages 1–4, Boston, MA, 2008. Springer US.
- Onu21. Hiroshi Onuki. On oriented supersingular elliptic curves. *Finite Fields and Their Applications*, 69:101777, 2021.
- Pet17. Christophe Petit. Faster algorithms for isogeny problems using torsion point images. In Tsuyoshi Takagi and Thomas Peyrin, editors, *Advances in Cryptology - ASIACRYPT 2017*, volume 10625 of *Lecture Notes in Computer Science*, pages 330–353. Springer, 2017.
- Piz80. Arnold Pizer. An algorithm for computing modular forms on $\gamma_0(n)$. *Journal of algebra*, 64(2):340–390, 1980.
- Rón92. Lajos Rónyai. Algorithmic properties of maximal orders in simple algebras over \mathbf{Q} . *Computational Complexity*, 2(3):225–243, 1992.
- Sho97. Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, 1997.
- Sil86. Joseph H. Silverman. *The Arithmetic of Elliptic Curves*, volume 106 of *Graduate Texts in Mathematics*. Springer-Verlag, 1986.

- Sim06. Denis Simon. Quadratic equations in dimensions 4, 5 and more. Preprint, 2006. See [Wat13] for a published review.
- Vig06. M.-F. Vignéras. *Arithmétique des algèbres de quaternions*, volume 800. Springer, 2006.
- Voi21. John Voight. *Quaternion Algebras*. Springer International Publishing, 2021. Graduate Texts in Mathematics, No. 288.
- Wat13. Mark Watkins. Some comments about indefinite LLL. *Diophantine Methods, Lattices, and Arithmetic Theory of Quadratic Forms*, 587(233):32, 2013.
- Wes22. Benjamin Wesolowski. The supersingular isogeny path and endomorphism ring problems are equivalent. In *FOCS 2021-62nd Annual IEEE Symposium on Foundations of Computer Science*, 2022.