Rational Isogenies from Irrational Endomorphisms

Wouter Castryck[1(✉)], Lorenz Panny[2(✉)], and Frederik Vercauteren[1(✉)]

[1] imec-COSIC, KU Leuven, Leuven, Belgium
{wouter.castryck,frederik.vercauteren}@esat.kuleuven.be
[2] Department of Mathematics and Computer Science,
Technische Universiteit Eindhoven, Eindhoven, The Netherlands
lorenz@yx7.cc

Abstract. In this paper, we introduce a polynomial-time algorithm to compute a connecting $\mathcal{O}$-ideal between two supersingular elliptic curves over $\mathbb{F}_p$ with common $\mathbb{F}_p$-endomorphism ring $\mathcal{O}$, given a description of their full endomorphism rings. This algorithm provides a reduction of the security of the CSIDH cryptosystem to the problem of computing endomorphism rings of supersingular elliptic curves. A similar reduction for SIDH appeared at Asiacrypt 2016, but relies on totally different techniques. Furthermore, we also show that any supersingular elliptic curve constructed using the complex-multiplication method can be located precisely in the supersingular isogeny graph by explicitly deriving a path to a known base curve. This result prohibits the use of such curves as a building block for a hash function into the supersingular isogeny graph.

Keywords: Isogeny-based cryptography · Endomorphism rings · CSIDH

1 Introduction

Isogeny-based cryptography is founded on the hardness of computing an isogeny between two isogenous elliptic curves over a finite field $\mathbb{F}_q$. Since this problem appears to remain hard even for quantum computers, it is one of the main candidates for building post-quantum cryptography [26]. Although the origins of isogeny-based cryptography go back to work by Couveignes from 1997 using ordinary elliptic curves [10], the currently most efficient instantiations rely on

Author list in alphabetical order; see https://www.ams.org/profession/leaders/culture/CultureStatement04.pdf. This work was supported in part by the Commission of the European Communities through the Horizon 2020 program under project number 643161 (ECRYPT-NET) and by the Research Council KU Leuven grants C14/18/067 and STG/17/019, and by CyberSecurity Research Flanders with reference number VR20192203. The first listed author was affiliated with the Department of Mathematics at KU Leuven during part of the preparation of this paper.
Date of this document: 2020-02-20.

A. Canteaut and Y. Ishai (Eds.): EUROCRYPT 2020, LNCS 12106, pp. 523–548, 2020.
https://doi.org/10.1007/978-3-030-45724-2_18

supersingular curves. These instantiations can be broadly classified into two families, known as SIDH [19] and CSIDH [7], depending on which supersingular elliptic curves and connecting isogenies are being used.

The acronym SIDH is shorthand for "Supersingular-Isogeny Diffie–Hellman", a key-exchange protocol introduced by Jao and De Feo in 2011 [19]. SIDH works in the full supersingular ℓ-isogeny graph, i.e., one considers the graph consisting of all (isomorphism classes of) supersingular elliptic curves defined over $\overline{\mathbb{F}}_p$ for a specifically chosen prime p and connecting isogenies of small prime degree ℓ. The vertices of this graph are the j-invariants of the isomorphism classes and are all contained in $\mathbb{F}_{p^2}$. Finding a path between two given vertices $j(E_1)$ and $j(E_2)$ is equivalent to constructing an isogeny between E_1 and E_2 whose degree is a power of ℓ.

The full endomorphism ring of a supersingular elliptic curve is a maximal order in a quaternion algebra. Kohel, Lauter, Petit and Tignol [22] showed that the above path-finding problem can be solved in (heuristically) expected polynomial time when given the endomorphism rings of E_1 and E_2; we will refer to this algorithm as "KLPT". Galbraith, Petit, Shani and Ti [16] later extended the KLPT algorithm specifically for the SIDH setting and showed that knowledge of the endomorphism rings of E_1 and E_2 suffices to break SIDH. Results by Eisenträger, Hallgren, Lauter, Morrison and Petit [13] show that finding a path in the isogeny graph is essentially equivalent to computing endomorphism rings.

CSIDH stands for "Commutative SIDH" and was introduced by Castryck, Lange, Martindale, Panny, and Renes [7] in 2018. CSIDH restricts the isogeny graph under consideration to supersingular elliptic curves and isogenies *defined over* $\mathbb{F}_p$ and mimics Couveignes' construction of a "hard homogeneous space". In particular, if E is a supersingular elliptic curve over $\mathbb{F}_p$, then its ring of $\mathbb{F}_p$-rational endomorphisms is an imaginary quadratic order $\mathcal{O} \subseteq \mathbb{Q}(\sqrt{-p})$. The C in "CSIDH" refers to the commutativity of $\mathcal{O}$, which (much like the situation on ordinary curves used by Couveignes) gives rise to an action of the (commutative) ideal-class group $\mathrm{cl}(\mathcal{O})$ on the set of supersingular elliptic curves over $\mathbb{F}_p$ with $\mathcal{O}$ as their $\mathbb{F}_p$-rational endomorphisms. This class-group action immediately leads to several cryptographic primitives such as identification, non-interactive key agreement, and even signature schemes.

1.1 Contributions

Our first contribution reduces the key recovery problem in CSIDH to computing the full endomorphism ring of the target curve, where in many cases even one non-$\mathbb{F}_p$-rational endomorphism suffices. More precisely, given two supersingular elliptic curves E, E' over $\mathbb{F}_p$ with $\mathbb{F}_p$-rational endomorphism ring $\mathcal{O}$, assuming sufficient knowledge of their full endomorphism rings $\mathrm{End}(E)$ and $\mathrm{End}(E')$, we show how to compute in polynomial time an ideal $\mathfrak{a} \subseteq \mathcal{O}$ such that $E' = [\mathfrak{a}]E$. This result can be seen as an analogon of [16] for SIDH, but uses different techniques, and in particular it does not rely on the KLPT algorithm [22].

Several remarks on this result are in order:

- In CSIDH all curves have *the same known* $\mathbb{F}_p$-rational endomorphism ring $\mathcal{O}$, which therefore does not contain any information specific to E, nor to $[\mathfrak{a}]$. This explains why we require knowledge of at least one endomorphism of E that is not $\mathbb{F}_p$-rational.
- Since both $\mathrm{End}(E_0)$ and $\mathrm{End}(E)$ are assumed to be known, one can run the KLPT algorithm to obtain an isogeny $\alpha\colon E_0 \to E$. However, this isogeny is most likely not $\mathbb{F}_p$-rational and as such does not correspond to the CSIDH private key. It is easy to verify that the isogeny $\beta = \alpha \circ \pi_{E_0} + \pi_E \circ \alpha$, with π the p-power Frobenius endomorphism on the respective curves, is an $\mathbb{F}_p$-rational isogeny[1] from E_0 to E. Note that β can be evaluated efficiently on points of E_0, but it is unclear how to efficiently derive an invertible ideal $\mathfrak{b} \subseteq \mathcal{O}$ whose action on E_0 corresponds to β. Such an ideal $\mathfrak{b}$ is required to break the CSIDH Diffie–Hellman key agreement and other derived protocols, since it is essentially a curve-independent way of specifying an $\mathbb{F}_p$-rational isogeny.
- Our polynomial-time algorithm returns an ideal $\mathfrak{a}$ whose norm is not necessarily smooth. To efficiently compute the action of $[\mathfrak{a}]$ therefore requires an extra smoothing step, which obtains an ideal of smooth norm in the ideal class $[\mathfrak{a}]$. This smoothing step is standard and consists of a combination of a class-group computation and lattice reduction to solve an instance of the approximate closest-vector problem (CVP). The class-group computation requires subexponential time using classical computers [18], but runs in polynomial time on a quantum computer [21]. Using the BKZ algorithm [28], one can solve the CVP problem up to a subexponential approximation factor in subexponential time. This last step therefore implies that asymptotically, the smoothing step requires subexponential time. However, we note that for any *practical* instantiation of CSIDH, solving the approximate CVP problem can be done fairly efficiently [4].

Our second contribution is motivated by an important open problem in isogeny-based cryptography, namely how to hash into a supersingular isogeny graph without revealing a path to a known base curve. This problem remains open both in the SIDH (full isogeny graph) and the CSIDH ($\mathbb{F}_p$-rational isogeny graph) setting. The hash function introduced by Charles, Goren and Lauter [8] can be used to hash any string into the supersingular isogeny graph, but by construction, the hash function itself leaks an isogeny path from a base curve. To illustrate the issue, we can compare with the standard elliptic-curve discrete-logarithm setting: The equivalent of the CGL construction would start from the public base point $P \in E(\mathbb{F}_q)$ and construct a point Q by multiplying P with a scalar computed deterministically from the message. As such, anyone would know the discrete logarithm of Q with respect to P, which voids cryptographic applications relying on the assumption that the relationship between Q and P cannot be discovered. To remedy this, elliptic-curve cryptosystems instead hash to curve points using maps like Elligator [3], which computes a point directly without passing through a scalar first, but an equivalent of these constructions in isogeny-based cryptography is not known.

[1] Unless $\beta = 0$.

Besides the random-walk approach à la CGL, it is also possible to generate supersingular elliptic curves using the complex-multiplication (CM) method [6]. It is therefore natural to wonder whether CM can be useful to hash into the supersingular isogeny graph, and in particular, whether finding paths between the resulting curves could be computationally hard. Our second result squashes this hope by locating these curves (and therefore also a path to a base curve) in the supersingular isogeny graph, in a surprisingly explicit manner (see Theorem 26(iii) for the exact statement).

The remainder of the paper is organized as follows. In Sect. 2 we recall the necessary mathematical background. In Sect. 3 we introduce the notion of twisting endomorphisms and explain their relation to $\mathbb{F}_p$-rational isogenies. Section 4 describes our new algorithm to compute a connecting ideal between two supersingular elliptic curves over $\mathbb{F}_p$ given their endomorphism rings and argues that (at least classically) our approach appears to be optimal. Finally, Sect. 5 shows how to locate supersingular elliptic curves constructed via CM in the isogeny graph, by explicitly deriving a path to a known starting curve.

2 Preliminaries

In this section we recall the required mathematical background and fix notation. Our focus lies on supersingular elliptic curves over finite prime fields $\mathbb{F}_p$, although much of what follows readily generalizes to arbitrary elliptic curves over arbitrary finite fields. Some of the observations below seem new.

For ease of exposition, we shall assume $p > 3$ throughout, noting that this is not necessarily a requirement for all of the statements.

2.1 Quadratic Twisting

For each odd prime number p we fix a non-square element $\xi \in \mathbb{F}_p$ along with a square root $\sqrt{\xi} \in \mathbb{F}_{p^2} \setminus \mathbb{F}_p$; if $p \equiv 3 \pmod 4$ then our default choice is $\xi = -1$ and we write $\mathbf{i} = \sqrt{-1}$. For an elliptic curve $E\colon y^2 = f(x)$ over $\mathbb{F}_p$ defined by some squarefree cubic polynomial $f(x) \in \mathbb{F}_p[x]$, we call the curve $E^t\colon \xi^{-1}y^2 = f(x)$ the *quadratic twist* of E over $\mathbb{F}_p$. The map $\tau_E\colon E \to E^t$, $(x, y) \mapsto (x, \sqrt{\xi} \cdot y)$ is a non-$\mathbb{F}_p$-rational isomorphism. From $\sqrt{\xi}^p = -\sqrt{\xi}$ one easily sees that

$$\tau_E \circ \pi_E = -\pi_{E^t} \circ \tau_E, \tag{1}$$

with π_E and π_{E^t} the respective Frobenius endomorphisms of E and E^t.

It can exceptionally happen that our definition of the quadratic twist is a trivial twist in the sense of [30, § X.2]:

Lemma 1. *An elliptic curve $E/\mathbb{F}_p$ is $\mathbb{F}_p$-isomorphic to its quadratic twist E^t if and only if $p \equiv 3 \pmod 4$ and $j(E) = 1728$.*

Proof. After an $\mathbb{F}_p$-isomorphism, we can assume $E\colon y^2 = x^3 + Ax + B$ with $A, B \in \mathbb{F}_p$ satisfying $4A^3 + 27B^2 \neq 0$. Then its quadratic twist is $\mathbb{F}_p$-isomorphic

to $y^2 = x^3 + A\xi^2 x + B\xi^3$ for some non-square ξ. According to [30, Prop. III.3.1] this curve is $\mathbb{F}_p$-isomorphic to E if and only if $A\xi^2 = Au^4$ and $B\xi^3 = Bu^6$ for some $u \in \mathbb{F}_p \setminus \{0\}$. This holds if and only if $B = 0$ and ξ^2 is a fourth power, from which the lemma follows. □

2.2 Hard Homogeneous Spaces from Supersingular Curves

Fix a prime number $p > 3$ and consider the imaginary quadratic number field $K = \mathbb{Q}(\sqrt{-p})$ along with its maximal order $\mathcal{O}_K$. If E is a supersingular elliptic curve defined over $\mathbb{F}_p$, then its ring $\mathrm{End}_p(E)$ of $\mathbb{F}_p$-rational endomorphisms admits an isomorphism to an order $\mathcal{O} \subseteq K$, under which π_E is mapped to $\sqrt{-p}$. In particular, $\mathcal{O}$ always contains the subring $\mathbb{Z}[\sqrt{-p}]$, hence if $p \equiv 1 \pmod 4$ then $\mathcal{O} = \mathcal{O}_K = \mathbb{Z}[\sqrt{-p}]$, while if $p \equiv 3 \pmod 4$ then either $\mathcal{O} = \mathbb{Z}[\sqrt{-p}]$ or $\mathcal{O} = \mathcal{O}_K = \mathbb{Z}[(1+\sqrt{-p})/2]$. We write $\mathcal{E}\ell\ell_p(\mathcal{O})$ to denote the set of $\mathbb{F}_p$-isomorphism classes of supersingular elliptic curves having endomorphism $\mathcal{O}$.

Remark 2. If $p \equiv 3 \pmod 4$, then the $\mathbb{F}_p$-endomorphism ring of a supersingular elliptic curve $E/\mathbb{F}_p$ is determined by its 2-torsion; see [12]: either we have $\#E(\mathbb{F}_p)[2] = 2$, in which case $E \in \mathcal{E}\ell\ell_p(\mathbb{Z}[\sqrt{-p}])$, or $\#E(\mathbb{F}_p)[2] = 4$, in which case $E \in \mathcal{E}\ell\ell_p(\mathbb{Z}[(1+\sqrt{-p})/2])$.

Every such order $\mathcal{O}$ comes equipped with its (ideal-)class group $\mathrm{cl}(\mathcal{O})$, which consists of invertible ideals modulo non-zero principal ideals; the class of an invertible ideal $\mathfrak{a} \subseteq \mathcal{O}$ is denoted by $[\mathfrak{a}]$. The number of elements of $\mathrm{cl}(\mathcal{O})$ is called the class number and denoted by $h(\mathcal{O})$.

Lemma 3. *If $p \equiv 3 \pmod 4$ then $h(\mathcal{O})$ is odd, while if $p \equiv 1 \pmod 4$ then $\mathrm{cl}(\mathcal{O})$ has a unique element of order 2, in particular $h(\mathcal{O})$ is even.*

Proof. This follows from genus theory [11]. □

Through

$$\mathrm{cl}(\mathcal{O}) \times \mathcal{E}\ell\ell_p(\mathcal{O}) \longrightarrow \mathcal{E}\ell\ell_p(\mathcal{O}): \quad ([\mathfrak{a}], E) \longmapsto [\mathfrak{a}]E := E/E[\mathfrak{a}]$$

the class group acts in a free and transitive manner on the set $\mathcal{E}\ell\ell_p(\mathcal{O})$ of ($\mathbb{F}_p$-isomorphism classes of) supersingular elliptic curves defined over $\mathbb{F}_p$ whose ring of $\mathbb{F}_p$-endomorphisms $\mathrm{End}_p(E)$ is isomorphic to $\mathcal{O}$ [31]. Here $E[\mathfrak{a}]$ denotes the intersection of the kernels of all elements of $\mathfrak{a}$ interpreted as endomorphisms of E; to compute this intersection it suffices to consider a set of generators of $\mathfrak{a}$.

Ignoring constructive issues, this group action (for large enough p) is conjectured to turn $\mathcal{E}\ell\ell_p(\mathcal{O})$ into a "hard homogeneous space", in which the following problems are assumed to be computationally infeasible:

Definition 4. (Vectorization problem.) *Given $E, E' \in \mathcal{E}\ell\ell_p(\mathcal{O})$, find the ideal class $[\mathfrak{a}] \in \mathrm{cl}(\mathcal{O})$ for which $E' = [\mathfrak{a}]E$.*

(Parallelization problem.) *Given $E, E', E'' \in \mathcal{E}\ell\ell_p(\mathcal{O})$, find the elliptic curve $[\mathfrak{a}][\mathfrak{b}]E$ where $[\mathfrak{a}], [\mathfrak{b}] \in \mathrm{cl}(\mathcal{O})$ are such that $E' = [\mathfrak{a}]E$ and $E'' = [\mathfrak{b}]E$.*

The hardness of the parallelization problem naturally relates to the security of the Diffie–Hellman-style key exchange protocol built from the above group action: starting from a publicly known base curve $E \in \mathcal{E}\ell\ell_p(\mathcal{O})$, the two parties Alice and Bob secretly sample $[\mathfrak{a}]$ resp. $[\mathfrak{b}]$ from $\mathrm{cl}(\mathcal{O})$, compute $[\mathfrak{a}]E$ resp. $[\mathfrak{b}]E$, and publish the result. The shared secret is then $[\mathfrak{a}][\mathfrak{b}]E$, which Alice computes as $[\mathfrak{a}]([\mathfrak{b}]E)$ and which Bob computes as $[\mathfrak{b}]([\mathfrak{a}]E)$. Clearly, in order to solve the parallelization problem, it suffices to solve the vectorization problem. On a quantum computer, the converse holds as well [14].

For later use we recall the following rule, which was pointed out in [7, Rem. 5], albeit very briefly and without proof (see also [1, Prop. 3.31]).

Lemma 5. *For all* $[\mathfrak{a}] \in \mathrm{cl}(\mathcal{O})$ *and all* $E \in \mathcal{E}\ell\ell_p(\mathcal{O})$ *we have* $[\mathfrak{a}]^{-1}E = ([\mathfrak{a}]E^t)^t$.

Proof. It is convenient to assume that $\mathfrak{a}$ is generated by elements of $\mathbb{Z}[\sqrt{-p}]$, which can be done without loss of generality by scaling with an appropriate principal ideal if needed. We claim that the composition

$$E \xrightarrow{\tau_E} E^t \longrightarrow\!\!\!\!\!\rightarrow E^t/E^t[\mathfrak{a}] = [\mathfrak{a}]E^t \xrightarrow{\tau_{[\mathfrak{a}]E^t}} ([\mathfrak{a}]E^t)^t$$

is an $\mathbb{F}_p$-rational isogeny whose kernel equals the ideal $\bar{\mathfrak{a}}$ obtained from $\mathfrak{a}$ by complex conjugation. This claim implies the lemma because $\mathfrak{a}\bar{\mathfrak{a}}$ is the principal ideal generated by $\mathrm{N}(\mathfrak{a})$.

Let φ be the middle isogeny $E^t \twoheadrightarrow E^t/E^t[\mathfrak{a}]$. Two applications of (1) yield

$$\pi_{([\mathfrak{a}]E^t)^t} \circ (\tau_{[\mathfrak{a}]E^t} \circ \varphi \circ \tau_E) = (\tau_{[\mathfrak{a}]E^t} \circ \varphi \circ \tau_E) \circ \pi_E,$$

implying the $\mathbb{F}_p$-rationality. One verifies that $a+b\sqrt{-p} \in \mathfrak{a}$ if and only if $a+b\pi_{E^t}$ vanishes on $\ker\varphi$, which holds if and only if $a - b\pi_E$ vanishes on $\ker(\varphi \circ \tau_E)$, from which it follows that $\ker(\tau_{[\mathfrak{a}]E^t} \circ \varphi \circ \tau_E) = \ker(\varphi \circ \tau_E) = E[\bar{\mathfrak{a}}]$. □

2.3 CSIDH

CSIDH (pronounced "seaside") is an efficient instantiation of the more general supersingular hard-homogeneous-spaces construction described in the previous section. We let $r \in \mathbb{Z}_{\geq 1}$ and consider a prime p of the form $p = 4\ell_1\ell_2\cdots\ell_r - 1$, where the ℓ_i's are distinct odd prime numbers. This implies $p \equiv 3 \pmod 8$, so a priori there are two options for $\mathcal{O}$, namely $\mathbb{Z}[\sqrt{-p}]$ and the maximal order $\mathcal{O}_K = \mathbb{Z}[(1+\sqrt{-p})/2]$. CSIDH chooses the former option. Recall from Remark 2 that this corresponds to supersingular elliptic curves over $\mathbb{F}_p$ having a unique $\mathbb{F}_p$-rational point of order 2.

Remark 6. The set $\mathcal{E}\ell\ell_p(\mathbb{Z}[\sqrt{-p}])$ is sometimes referred to as the "floor", as opposed to $\mathcal{E}\ell\ell_p(\mathbb{Z}[(1+\sqrt{-p})/2])$ which is called the "surface". This terminology comes from the volcano structure of the 2-isogeny graph of supersingular elliptic curves over $\mathbb{F}_p$; see [12]. We stress that CSIDH can be set up equally well on the surface, although a convenient feature of the floor is that each $E \in \mathcal{E}\ell\ell_p(\mathbb{Z}[\sqrt{-p}])$ is $\mathbb{F}_p$-isomorphic to a Montgomery curve $E_A\colon y^2 = x^3 + Ax^2 + x$ for a unique coefficient $A \in \mathbb{F}_p$; furthermore, the coefficient defining E^t is then given by $-A$.

The prime p was chosen such that the primes $\ell_1, \ell_2, \ldots, \ell_r$ exhibit particularly easy splitting behaviour in $\mathbb{Z}[\sqrt{-p}]$, namely

$$(\ell_i) = (\ell_i, \sqrt{-p} - 1)(\ell_i, \sqrt{-p} + 1). \tag{2}$$

We refer to the respective factors, which are complex conjugates of each other, by $\mathfrak{l}_i$ and $\overline{\mathfrak{l}_i}$. If we define $\ell_0 := 4$ then (2) also applies to $i = 0$, so we can similarly define $\mathfrak{l}_0$ and $\overline{\mathfrak{l}_0}$. All these ideals are clearly invertible, so we can consider their classes $[\mathfrak{l}_i]$ and $[\overline{\mathfrak{l}_i}] = [\mathfrak{l}_i]^{-1}$ inside $\mathrm{cl}(\mathcal{O})$. Although this is not known in general, it seems likely that the $[\mathfrak{l}_i]$'s together generate the entire class group.

Example 7. The concrete instantiation CSIDH-512 from [7] has $r = 74$, where ℓ_1, ℓ_2, …, ℓ_{73} are the odd primes up to 373 and where $\ell_{74} = 587$. This results in a 511-bit prime p. The structure of $\mathrm{cl}(\mathbb{Z}[\sqrt{-p}])$ was computed by Beullens, Kleinjung and Vercauteren [4], who verified that $[\mathfrak{l}_1] = [(3, \sqrt{-p} - 1)]$ is in fact a generator.

The basic idea is then to let Alice and Bob choose their secrets as

$$[\mathfrak{a}] = [\mathfrak{l}_1]^{a_1}[\mathfrak{l}_2]^{a_2} \cdots [\mathfrak{l}_r]^{a_r} \qquad \text{resp.} \qquad [\mathfrak{b}] = [\mathfrak{l}_1]^{b_1}[\mathfrak{l}_2]^{b_2} \cdots [\mathfrak{l}_r]^{b_r},$$

for exponent vectors $(a_1, a_2, \ldots, a_r)$ and $(b_1, b_2, \ldots, b_r)$ sampled at random from some bounded subset of $\mathbb{Z}^r$, for instance uniformly from a hypercube $[-B; B]^r$ of size $(2B+1)^r \approx h(\mathbb{Z}[\sqrt{-p}]) \approx \sqrt{p}$. The resulting public keys and shared secret are then computed using $|a_1| + \ldots + |a_r|$ resp. $|b_1| + \ldots + |b_r|$ repeated actions of $[\mathfrak{l}_i]$ or $[\mathfrak{l}_i]^{-1} = [\overline{\mathfrak{l}_i}]$. If $E \in \mathcal{E}\ell\ell_p(\mathbb{Z}[\sqrt{-p}])$ then the subgroups

$$E[\mathfrak{l}_i] = \{\, P \in E[\ell_i] \mid \pi_E(P) = P \,\} = E(\mathbb{F}_p)[\ell_i]$$
$$E[\overline{\mathfrak{l}_i}] = \{\, P \in E[\ell_i] \mid \pi_E(P) = -P \,\}$$

consist of points having $\mathbb{F}_p$-rational x-coordinates; therefore, these actions are easy to evaluate using low-degree Vélu-type formulas and involving only arithmetic in $\mathbb{F}_p$.

As far as we know, the following class group relations have not appeared in the literature before:[2]

Lemma 8. *In* $\mathrm{cl}(\mathbb{Z}[\sqrt{-p}])$*, we have*

$$[\mathfrak{l}_1][\mathfrak{l}_2] \cdots [\mathfrak{l}_r] = [\overline{\mathfrak{l}_0}] \neq [1] \qquad \textit{and} \qquad [\mathfrak{l}_1]^3[\mathfrak{l}_2]^3 \cdots [\mathfrak{l}_r]^3 = [1].$$

Proof. One easily verifies that

$$\mathfrak{l}_1\mathfrak{l}_2 \cdots \mathfrak{l}_r = \left(\frac{p+1}{4}, \sqrt{-p} - 1\right) \qquad \text{and} \qquad \mathfrak{l}_0\mathfrak{l}_1\mathfrak{l}_2 \cdots \mathfrak{l}_r = \left(\sqrt{-p} - 1\right).$$

The latter identity implies $[\mathfrak{l}_1][\mathfrak{l}_2] \cdots [\mathfrak{l}_r] = [\mathfrak{l}_0]^{-1} = [\overline{\mathfrak{l}_0}]$, while the former shows that $[\mathfrak{l}_1][\mathfrak{l}_2] \cdots [\mathfrak{l}_r]$ is an element of order 3. Indeed, it represents a non-trivial

[2] After we posted a version of this paper online, we learned that this was observed independently and quasi-simultaneously in [27], with a more elaborate discussion.

ideal class because $\mathbb{Z}[\sqrt{-p}]$ contains no elements of norm $(p+1)/4$, while its order divides 3 since

$$\left(\frac{p+1}{4}, \sqrt{-p} - 1\right)\mathcal{O}_K = \frac{1+\sqrt{-p}}{2}\mathcal{O}_K,$$

i.e., it belongs to the kernel of the group homomorphism

$$\mathrm{cl}(\mathcal{O}) \longrightarrow \mathrm{cl}(\mathcal{O}_K),\ \mathfrak{a} \longmapsto \mathfrak{a}\mathcal{O}_K$$

which is 3-to-1 by [9, Thm. 5.2]. □

Note that this allows for reduction of the secret exponent vectors of Alice and Bob modulo $(3, 3, \ldots, 3)$. It also shows that the action of $[\mathfrak{l}_1][\mathfrak{l}_2]\cdots[\mathfrak{l}_r]$ can be evaluated using a single application of $[\overline{\mathfrak{l}_0}] = [(4, \sqrt{-p}+1)]$. The latter step can be taken using an isogeny of degree 4, or using a composition of two isogenies of degree 2, which necessarily makes us pass through the surface.

2.4 The Full Endomorphism Ring

The "full" endomorphism ring of a supersingular elliptic curve, as opposed to merely the $\mathbb{F}_p$-rational endomorphisms, plays a fundamental role in the theory of supersingular isogeny graphs.

An elliptic curve E is supersingular if and only if $\mathrm{End}(E)$ is non-commutative. In that case, $\mathrm{End}(E)$ embeds as a maximal order into a certain quaternion algebra $B_{p,\infty}$ ramified at p and infinity, which is unique up to isomorphism. Concretely, $B_{p,\infty}$ can be constructed as a four-dimensional $\mathbb{Q}$-algebra of the form $\mathbb{Q} \oplus \mathbb{Q}\mathbf{i} \oplus \mathbb{Q}\mathbf{j} \oplus \mathbb{Q}\mathbf{ij}$, subject to the multiplication rules $\mathbf{i}^2 = -q$, $\mathbf{j}^2 = -p$, and $\mathbf{ji} = -\mathbf{ij}$, for some positive integer q that depends on p. In the common case that $p \equiv 3 \pmod 4$, we can and will use $q = 1$. (Thus $B_{p,\infty}$ may be viewed as two imaginary quadratic fields "glued together" non-commutatively.) We certainly cannot stress enough that the embedding $\mathrm{End}(E) \hookrightarrow B_{p,\infty}$ is *extremely non-unique*; in fact, there are always infinitely many choices, and usually none of them sticks out as being particularly natural.

The notions of dual, degree, and trace of endomorphisms carry over to $B_{p,\infty}$: Taking the dual corresponds to conjugation, which maps $\alpha = a + b\mathbf{i} + c\mathbf{j} + d\mathbf{ij}$ to $\overline{\alpha} = a - b\mathbf{i} - c\mathbf{j} - d\mathbf{ij}$. The degree turns into $\mathrm{N}(\alpha) = \alpha\overline{\alpha} = a^2 + b^2q + c^2p + d^2qp$, and the trace is simply $\mathrm{tr}(\alpha) = \alpha + \overline{\alpha} = 2a$. Moreover, the trace yields a symmetric bilinear map $\langle \alpha, \beta \rangle = \mathrm{tr}(\overline{\alpha}\beta)$ on $B_{p,\infty}$, with respect to which the basis $1, \mathbf{i}, \mathbf{j}, \mathbf{ij}$ is orthogonal. With this, finding *an* embedding $\mathrm{End}(E) \hookrightarrow B_{p,\infty}$ when being given rational maps that span $\mathrm{End}(E)$ in some computationally effective way is easy: A variant of Schoof's point counting algorithm [29] can be used to compute traces of endomorphisms, and thereby the map $\langle \cdot, \cdot \rangle$, which can then be used in the Gram–Schmidt process to compute an orthogonal basis of the given endomorphism ring. Once the basis is orthogonal, some norm computations are necessary to align the given maps with the algebraic properties of the abstract quaternion representation. See [13, § 5.4] for details. We will commonly use the

$\mathbb{Q}$-basis $(1, \mathbf{i}, \mathbf{j}, \mathbf{ij})$ in the forthcoming algorithms to compute with $\mathrm{End}(E)$; the isomorphism to the corresponding rational maps of curves will be made explicit whenever it is realized computationally.

One reason why the endomorphism rings are interesting for cryptographic applications is because they contain all the information necessary to construct an isogeny between two curves: Given $\mathrm{End}(E)$ and $\mathrm{End}(E')$, it is easy to find a *connecting ideal* $\mathcal{I}$ between them; that is, a lattice in $B_{p,\infty}$ that is a left ideal of $\mathrm{End}(E)$ and a right ideal of $\mathrm{End}(E')$. For example, the following choice works:

Lemma 9. *Between any two maximal orders $\mathcal{Q}$ and $\mathcal{Q}'$ in $B_{p,\infty}$, the lattice $\mathcal{I} = \mathcal{Q}\mathcal{Q}' = \mathrm{span}\,\{ab \mid a \in \mathcal{Q}, b \in \mathcal{Q}'\}$ is a connecting ideal.*

Proof. This is an easy special case of [20, Algorithm 3.5]: Clearly $\mathcal{Q}\mathcal{I} \subseteq \mathcal{I}$, hence $\mathcal{O}_{\mathsf{L}}(\mathcal{I}) \supseteq \mathcal{Q}$, and equality follows since $\mathcal{Q}$ is maximal. Similarly, $\mathcal{O}_{\mathsf{R}}(\mathcal{I}) = \mathcal{Q}'$. □

The intersection of all kernels of endomorphisms contained in this ideal is a finite subgroup determining a separable isogeny $E \longrightarrow E'$. One can prove that the codomain curve of the isogeny given by such a left ideal of $\mathrm{End}(E)$ only depends on the left-ideal *class* of $\mathcal{I}$: This is what the Kohel–Lauter–Petit–Tignol algorithm [22] exploits to find a *smooth-degree*, hence efficiently computable, isogeny between E and E' given their endomorphism rings.

Since we are concerned with supersingular elliptic curves defined over $\mathbb{F}_p$, our endomorphism rings—maximal orders in $B_{p,\infty}$—will always contain a copy of the Frobenius order $\mathbb{Z}[\sqrt{-p}] \cong \mathbb{Z}[\pi_E] \subseteq \mathrm{End}_p(E)$. It thus makes sense to fix the image of the Frobenius endomorphism π_E when embedding $\mathrm{End}(E)$ into $B_{p,\infty}$ once and for all: We will always assume that π_E is mapped to $\mathbf{j}$.

3 Twisting Endomorphisms

As before, we focus on the case of finite fields $\mathbb{F}_p$ with $p > 3$ prime.

Definition 10. *Let E be an elliptic curve defined over $\mathbb{F}_p$. An endomorphism $\alpha \in \mathrm{End}(E)$ is called a* twisting endomorphism *of E if*

$$\alpha \circ \pi_E = -\pi_E \circ \alpha.$$

(Note that E must necessarily be supersingular for this to be possible.)

Lemma 11. *Let E be an elliptic curve defined over $\mathbb{F}_p$. The non-zero twisting endomorphisms of E are precisely the elements of $\mathrm{End}(E)$ that are purely imaginary over $\mathrm{End}_p(E)$.*

Proof. Write $\alpha = a + b\mathbf{i} + c\mathbf{j} + d\mathbf{ij}$ with $a, b, c, d \in \mathbb{Q}$; then using the fact that π_E is mapped to $\mathbf{j}$, the equality $\alpha \circ \pi_E = -\pi_E \circ \alpha$ implies $a = c = 0$. □

Lemma 12. *Twisting endomorphisms have kernels defined over $\mathbb{F}_p$. (Thus they always equal either the zero map or an $\mathbb{F}_p$-isogeny followed by an isomorphism.)*

Proof. Since $\pi_E^{-1}(\ker \alpha) = \ker(\alpha \circ \pi_E) = \ker(-\pi_E \circ \alpha) = \ker \alpha$, the subgroup $\ker \alpha$ is stable under the action of $\mathrm{Gal}(\overline{\mathbb{F}}_p/\mathbb{F}_p)$, hence $\mathbb{F}_p$-rational. □

Lemma 13. *Let E be an elliptic curve as above and let α be a non-zero twisting endomorphism of E. Then $\tau_E \circ \alpha \colon E \to E^t$ is an $\mathbb{F}_p$-rational isogeny of degree $\mathrm{N}(\alpha)$.*

Proof. Since τ_E is an isomorphism we have $\deg(\tau_E \circ \alpha) = \deg \alpha = N(\alpha)$, so it remains to prove the $\mathbb{F}_p$-rationality, which follows from

$$\tau_E \circ \alpha \circ \pi_E = -\tau_E \circ \pi_E \circ \alpha = \pi_{E^t} \circ \tau_E \circ \alpha$$

where the last equality uses that $\sqrt{\xi} \in \mathbb{F}_{p^2} \setminus \mathbb{F}_p$ and therefore $\sqrt{\xi}^p = -\sqrt{\xi}$.

4 Isogenies from Known Endomorphisms

In this section, we describe how to find a connecting ideal between two supersingular elliptic curves over $\mathbb{F}_p$ given their full endomorphism rings.

The basic idea behind our approach is to exploit the symmetry of the isogeny graph over $\mathbb{F}_p$ with respect to quadratic twisting; cf. Lemma 5: Intuitively, the distance between a curve and its quadratic twist tells us where in the graph it is located, and combining this information for two curves allows finding the distance between them. See Fig. 1 below for an illustration.

In more mathematical terms, the "distance" between E and its quadratic twist corresponds to an invertible ideal $\mathfrak{a} \subseteq \mathcal{O}$ that connects E to E^t, i.e., satisfies $[\mathfrak{a}]E = E^t$. We will show in Algorithm 1 how to find such an ideal, given the full endomorphism ring of E. Subsequently, given two arbitrary supersingular elliptic curves E, E' with the same $\mathbb{F}_p$-endomorphism ring $\mathcal{O}$ together with such a "twisting ideal" for each of them, Algorithm 2 can be used to find a connecting ideal from E to E', i.e., an invertible ideal $\mathfrak{c} \subseteq \mathcal{O}$ such that $[\mathfrak{c}]E = E'$.

The following lemma shows the relationship between ideals in $\mathrm{End}_p(E)$ and $\mathrm{End}(E)$ that determine the same subgroup; it is of crucial significance for the forthcoming algorithms.

Lemma 14. *Let E be a supersingular elliptic curve defined over $\mathbb{F}_p$. Consider a non-zero ideal $\mathfrak{c} \subseteq \mathrm{End}_p(E)$ and a non-zero left ideal $\mathcal{I} \subseteq \mathrm{End}(E)$ such that the corresponding subgroups $E[\mathcal{I}]$ and $E[\mathfrak{c}]$ are equal. Then $\mathcal{I} \cap \mathrm{End}_p(E) = \pi_E^k \mathfrak{c}$ for some $k \in \mathbb{Z}$.*[3]

Proof. Following [31, Thm. 4.5], we know that for every order $\mathcal{O}$ which can arise as an endomorphism ring, every ideal of $\mathcal{O}$ is a kernel ideal, and thus

$$\mathcal{I} = \{\gamma \in \mathrm{End}(E) \mid \ker \gamma \supseteq E[\mathcal{I}]\} \cdot \pi_E^r$$
$$\mathfrak{c} = \{\gamma \in \mathrm{End}_p(E) \mid \ker \gamma \supseteq E[\mathfrak{c}]\} \cdot \pi_E^s$$

with non-negative integers $r, s \in \mathbb{Z}$. Now $E[\mathcal{I}] = E[\mathfrak{c}]$ by assumption, hence it follows that $\mathcal{I} \cap \mathrm{End}_p(E) = \pi_E^{r-s}\mathfrak{c}$, which shows the claim. □

[3] One could handle the purely inseparable part—powers of π_E—in a unified way by working with scheme-theoretic kernels. Since this issue is only tangential to our work, we will for simplicity avoid this technical complication and deal with π_E explicitly.

4.1 The Algorithm

Throughout this section, we write $\mathcal{O}_E := \mathrm{End}_p(E)$ for brevity.

Recall from Sect. 2.4 that we assume $\mathrm{End}(E)$ is represented as a maximal order in $B_{p,\infty}$ with respect to the $1, \mathbf{i}, \mathbf{j}, \mathbf{ij}$ basis, and such that the Frobenius endomorphism π_E is mapped to $\mathbf{j} \in B_{p,\infty}$ under the embedding.

We start off with an algorithm to find an ideal that connects a curve to its quadratic twist, which will be used as a building block for the main algorithm to connect two arbitrary curves with the same $\mathbb{F}_p$-endomorphism ring in the $\mathbb{F}_p$-isogeny graph.

Algorithm 1: Connecting ideal of a curve and its twist.

Input: a supersingular $E/\mathbb{F}_p$ and the full endomorphism ring $\mathrm{End}(E)$.

Output: an invertible ideal $\mathfrak{a} \subseteq \mathcal{O}_E$ such that $[\mathfrak{a}]E = E^t$.

Find a non-zero element $\alpha \in \mathrm{End}(E)$ of the form $x\mathbf{i} + y\mathbf{ij}$.

Compute the ideal $\mathfrak{a} := \big(\mathrm{End}(E) \cdot \alpha\big) \cap \mathcal{O}_E$.

Return $\mathfrak{a}$.

Lemma 15. *Algorithm 1 is correct and runs in polynomial time.*

Proof. Note that $\alpha \in \mathbf{i}\mathcal{O}_E$ is a twisting endomorphism of E due to Lemma 11. Hence, $E[\mathrm{End}(E) \cdot \alpha] = \ker \alpha$ is an $\mathbb{F}_p$-rational subgroup of E giving rise to an $\mathbb{F}_p$-rational isogeny $E \longrightarrow E^t$, which is necessarily horizontal since $\mathcal{O}_E = \mathcal{O}_{E^t}$. Therefore, there exists an invertible ideal $\mathfrak{c}$ of $\mathcal{O}_E$ such that $E[\mathfrak{c}] = \ker \alpha$, and we may apply Lemma 14 to conclude that $\mathfrak{a} = \big(\mathrm{End}(E) \cdot \alpha\big) \cap \mathcal{O}_E$ in fact equals the desired ideal $\mathfrak{c}$—up to powers of π_E, which is an endomorphism.

Regarding the runtime, everything consists of basic arithmetic in $B_{p,\infty}$ and some linear algebra over $\mathbb{Q}$ and $\mathbb{Z}$. □

As mentioned before, the inherent symmetry of the $\mathbb{F}_p$-isogeny graph with respect to quadratic twisting implies that the "location" of a curve E in the graph is somehow related to the properties of ideals that connect E to its quadratic twist E^t. The following lemma makes this intuition precise, in the sense that it determines a connecting ideal between two curves almost uniquely when given a twisting ideal for each of them. This correspondence is then used in an explicit manner to compute a connecting ideal in Algorithm 2.

Lemma 16. *Let E_0 and E_1 be supersingular elliptic curves defined over $\mathbb{F}_p$ with $\mathrm{End}_p(E_0) \cong \mathrm{End}_p(E_1)$, such that we may simply write $\mathcal{O}$ for both. If $\mathfrak{b}, \mathfrak{c} \subseteq \mathcal{O}$ are invertible ideals such that $[\mathfrak{b}]E_0 = E_0^t$ and $[\mathfrak{c}]E_1 = E_1^t$, then the unique ideal class $[\mathfrak{a}]$ such that $[\mathfrak{a}]E_0 = E_1$ satisfies the equation $[\mathfrak{a}]^2 = [\mathfrak{b}][\mathfrak{c}]^{-1}$.*

Proof. By Lemma 5, applying the action of an ideal class $[\mathfrak{u}]$ to E^t gives the same result as first applying $[\overline{\mathfrak{u}}] = [\mathfrak{u}]^{-1}$ and then twisting. Hence, if $[\mathfrak{a}]E_0 = E_1$, then $[\mathfrak{a}]^{-1}E_0^t = E_1^t$. However, by the assumptions, we have $[\mathfrak{a}]^{-1}E_0^t = [\mathfrak{a}]^{-1}[\mathfrak{b}]E_0$ on the left-hand side and $E_1^t = [\mathfrak{c}]E_1 = [\mathfrak{c}][\mathfrak{a}]E_0$ on the right-hand side, which implies the claimed equality of ideal classes as the class-group action is free. See Fig. 1 for a visualization of the situation on an isogeny cycle. □

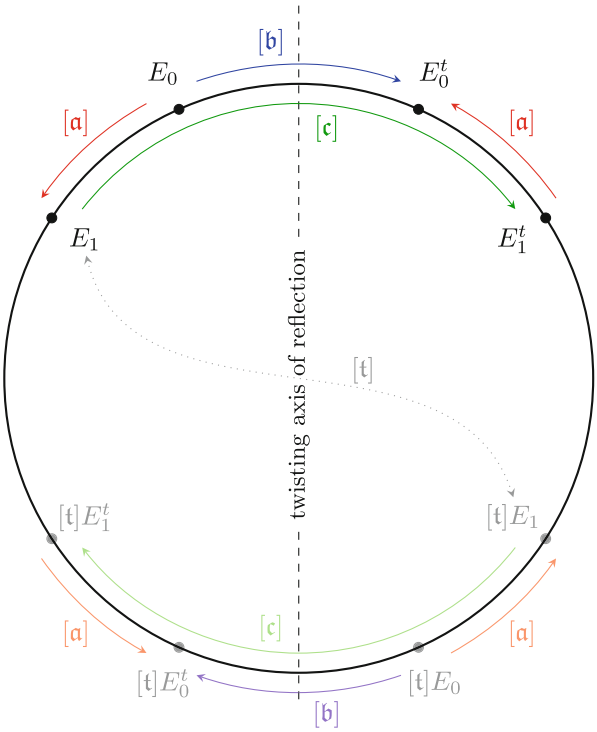

Fig. 1. Illustration of Lemma 16 and the square-root issue in Lemma 17. If the ideal $\mathfrak{t} = (2, \sqrt{-p})$ is non-principal and invertible in $\mathcal{O}$, it corresponds to a point symmetry with respect to the "center" of the isogeny cycle, and the entire relationship between $E_{0,1}$ and their twists is replicated on the "opposite" side with the "dual" curves $[\mathfrak{t}]E_{0,1}$ and their twists. This explains why the output of Algorithm 2 is a priori only correct up to multiplication by $\mathfrak{t}$; the quadratic equation determining $[\mathfrak{a}]$ simply cannot distinguish whether $[\mathfrak{a}]$ jumps between the two worlds or not.

Algorithm 2: Connecting ideal of two curves.

Input: supersingular elliptic curves $E_0, E_1/\mathbb{F}_p$ with $\mathcal{O}_{E_0} = \mathcal{O}_{E_1} = \mathcal{O}$, together with their full endomorphism rings $\mathrm{End}(E_0)$ and $\mathrm{End}(E_1)$.

Output: an invertible ideal $\mathfrak{a} \subseteq \mathcal{O}$ such that $[\mathfrak{a}]E_0 = E_1$.

Using Algorithm 1, find an invertible ideal $\mathfrak{b} \subseteq \mathcal{O}$ with $[\mathfrak{b}]E_0 = E_0^t$.
Likewise, find an invertible ideal $\mathfrak{c} \subseteq \mathcal{O}$ such that $[\mathfrak{c}]E_1 = E_1^t$.
Compute an ideal $\mathfrak{a} \subseteq \mathcal{O}$ such that $[\mathfrak{a}]^2 = [\mathfrak{b}][\mathfrak{c}]^{-1}$ in $\mathrm{cl}(\mathcal{O})$ using [5, § 6].
If $p \equiv 1 \pmod 4$ and the right order of $\mathrm{End}(E_0) \cdot \mathfrak{a}$ in $B_{p,\infty}$ is not isomorphic to $\mathrm{End}(E_1)$, then replace $\mathfrak{a}$ by $\mathfrak{a} \cdot (2, 1+\sqrt{-p})$.

Return $\mathfrak{a}$.

Lemma 17. *Algorithm 2 is correct and runs in polynomial time.*

Proof. Most of this follows from Lemmas 16 and 15. The square root in $\mathrm{cl}(\mathcal{O})$ to determine the ideal $\mathfrak{a}$ can be computed in polynomial time using the algorithm in [5, § 6].

Regarding the correctness of the output, recall from Lemma 3 that the class number of $\mathcal{O}$ is odd if $p \equiv 3 \pmod 4$, hence the square root $[\mathfrak{a}]$ is unique. On the other hand, if $p \equiv 1 \pmod 4$, then Lemma 3 implies that there are exactly two square roots. Now the order $\mathcal{O}$ has discriminant $-4p$, hence $(p) = (\sqrt{-p})^2$ and $(2) = (2, 1+\sqrt{-p})^2$ are the only ramified primes. The principal ideal $(\sqrt{-p})$ becomes trivial in $\mathrm{cl}(\mathcal{O})$. However, $\mathfrak{t} := (2, 1+\sqrt{-p})$ is non-principal as there is no element of norm 2 in $\mathcal{O}$, hence $[\mathfrak{t}]$ is an element of order 2 in $\mathrm{cl}(\mathcal{O})$. Thus the two square roots of $[\mathfrak{b}][\mathfrak{c}]^{-1}$ are $[\mathfrak{a}]$ and $[\mathfrak{at}]$. The final check in the algorithm identifies the correct choice by lifting $\mathfrak{a}$ to a left $\mathrm{End}(E_0)$-ideal and comparing its right order to the endomorphism ring of E_1; they must be isomorphic if $\mathfrak{a}$ determines an isogeny $E_0 \to E_1$ as intended. □

An Example. To illustrate the algorithms in this section, we will show their workings on a concrete, rather special example.

Lemma 18. *Assume $p \equiv 3 \pmod 4$ and let E_1 be a supersingular elliptic curve over $\mathbb{F}_p$ with $\mathbb{F}_p$-endomorphism ring $\mathcal{O}$. Let E_0 be the elliptic curve in $\mathcal{Ell}_p(\mathcal{O})$ having j-invariant 1728. If $\mathfrak{b} \subseteq \mathcal{O}$ is an invertible ideal such that $[\mathfrak{b}]E_1 = E_1^t$, then the unique ideal class $[\mathfrak{a}]$ such that $[\mathfrak{a}]E_0 = E_1$ is given by $[\mathfrak{b}]^{(h(\mathcal{O})-1)/2}$.*

Proof. This follows from Lemmas 1 and 16, together with the fact that the class number of $\mathcal{O}$ is odd. □

Example 19. Assume that $p \equiv 11 \pmod{12}$. We illustrate Algorithm 2 by computing a connecting ideal $\mathfrak{a}$ between $E_0 \colon y^2 = x^3 + x$ and $E_1 \colon y^2 = x^3 + 1$. Note that both curves are contained in $\mathcal{Ell}_p(\mathbb{Z}[\sqrt{-p}])$, as can be seen by considering $E(\mathbb{F}_p)[2]$. If $\omega \in \mathbb{F}_{p^2} \setminus \mathbb{F}_p$ denotes a primitive 3^{rd} root of unity, then E_1 admits the automorphism $(x, y) \mapsto (\omega x, y)$, which will, by abuse of notation, be denoted by ω as well. According to [25, Prop. 3.2],[4] the endomorphism ring of E_1 is isomorphic to the $B_{p,\infty}$-order

$$\mathcal{Q} = \mathbb{Z} + \mathbb{Z}\frac{-1+\mathbf{i}}{2} + \mathbb{Z}\mathbf{j} + \mathbb{Z}\frac{3+\mathbf{i}+3\mathbf{j}+\mathbf{ij}}{6},$$

where $\mathbf{i}$ corresponds to $2\omega + 1$ and satisfies[5] $\mathbf{i}^2 = -3$, and as usual $\mathbf{j}$ corresponds to the Frobenius endomorphism π_{E_1}. If we choose the twisting endomorphism $\alpha = \mathbf{i}$ in Algorithm 1, then we find $\mathcal{Q}\mathbf{i} \cap \mathbb{Z}[\mathbf{j}] = (3, \mathbf{j} - 1)$. (Of course, this also

[4] Unfortunately, the statement of [25, Prop. 3.2] wrongly attributes this description to the quadratic twist of E_1.

[5] Here we deviate from our convention that $\mathbf{i}^2 = -1$ as soon as $p \equiv 3 \pmod 4$.

follows from the fact that $2\omega+1$ is a degree-3 isogeny whose kernel $\{(0,\pm 1),\infty\}$ is $\mathbb{F}_p$-rational.) So $E_1^t = [(3,\sqrt{-p}-1)]E_1$, and we can take

$$\mathfrak{a} = (3,\sqrt{-p}-1)^{(h(\mathbb{Z}[\sqrt{-p}])-1)/2} \tag{3}$$

by Lemma 18. Thus, in the 3-isogeny graph associated with $\mathcal{E}\ell\ell_p(\mathbb{Z}[\sqrt{-p}])$, which is a union of cycles, the curve E_1 and its twist $E_1^t\colon y^2 = x^3 - 1$ can be found "opposite" of our starting curve E_0, on the same cycle. We will generalize this example in Sect. 5.

Example 20. In particular, the findings of Example 19 hold for a CSIDH prime $p = 4\ell_1\ell_2\cdots\ell_r - 1$ with $\ell_1 = 3$, so that $(3,\sqrt{-p}-1) = \mathfrak{l}_1$. Note that $E\colon y^2 = x^3+1$ is isomorphic to the Montgomery curve $E_{-\sqrt{3}}\colon y^2 = x^3 - \sqrt{3}\cdot x^2 + x$ through

$$E_{-\sqrt{3}} \longrightarrow E,\ (x,y) \longmapsto (\delta^2 x - 1, \delta^3 y),$$

where $\sqrt{3} \in \mathbb{F}_p$ denotes the square root of 3 which is a square itself, and $\delta^2 = \sqrt{3}$. In view of the class-group computation carried out in [4] for the CSIDH-512 parameter set, the previous example shows that the ideal

$$\mathfrak{l}_1^{12732622111474213758851509300531960108081025715274321179628543048779880586309 5}$$

takes the starting Montgomery coefficient 0 to the coefficient $-\sqrt{3}$, and one further application of $\mathfrak{l}_1$ takes it to $\sqrt{3}$. Smoothing this ideal using the class-group relations of $\mathrm{cl}(\mathbb{Z}[\sqrt{-p}])$ from [4] yields (for instance) the CSIDH-512 exponent vector

$$\begin{array}{cccccccccccccccccc}
(5, & -7, & -1, & 1, & -4, & -5, & -8, & 4, & -1, & 5, & 1, & 0, & -2, & -4, & -2, & 2, & -9, & 4, & 2,\\
5, & 1, & 1, & 1, & 5, & -4, & 2, & 6, & 5, & -1, & 0, & 0, & -4, & -1, & -3, & -1, & -4, & 1, & 7,\\
1, & 4, & 1, & 4, & -7, & 0, & -3, & -1, & 0, & 1, & 2, & 3, & 1, & 2, & -4, & -5, & 9, & -1, & 4,\\
0, & 5, & 1, & 0, & 1, & 1, & 3, & 0, & 2, & 2, & 2, & -1, & 2, & 1, & -1, & 11, & 3), & &
\end{array}$$

which can indeed be readily verified to connect E_0 to $E_{-\sqrt{3}}$ by plugging it into a CSIDH-512 implementation, such as that of [7], as a private key.

Example 21. If in Example 19, we instead choose the twisting endomorphism

$$\alpha = \frac{\mathbf{i}+\mathbf{ij}}{3} = -1 - \mathbf{j} + 2\frac{3+\mathbf{i}+3\mathbf{j}+\mathbf{ij}}{6} \in \mathcal{Q},$$

then we obtain a twisting ideal of norm $(p+1)/3$. In the CSIDH setting of Example 20 above, one can deduce that this ideal is nothing but $\bar{\mathfrak{l}}_0\bar{\mathfrak{l}}_2\bar{\mathfrak{l}}_3\cdots\bar{\mathfrak{l}}_r$, so this confirms the first class-group relation stated in Lemma 8.

4.2 Incomplete Knowledge of Endomorphism Rings

At first sight, there appears to be no strong reason why one requires the full endomorphism rings to be known exactly in Algorithm 1, rather than for instance a full-rank proper subring $\mathcal{Q} \subsetneq \mathrm{End}(E)$ containing $\mathcal{O}$: Twisting endomorphisms

α can clearly be found in every full-rank subring, and one can still compute the left ideal $\mathcal{Q} \cdot \alpha$, which can then be intersected with $\mathcal{O}$. The result is indeed an ideal $\mathfrak{a}$ of $\mathcal{O}$, as desired, but unfortunately it turns out that $\mathfrak{a}$ usually falls short of connecting E to its quadratic twist unless in fact $\mathcal{Q} = \mathrm{End}(E)$. This is not surprising: If $\mathcal{Q}$ is contained in multiple non-isomorphic maximal orders, then the algorithm would need to work for all those maximal orders—and therefore elliptic curves—simultaneously, which is absurd. However, luckily, one can prove that $\mathfrak{a}$ is only *locally* "wrong" at the conductor, i.e., the index $f := \big[\mathrm{End}(E) : \mathcal{Q}\big]$.

Lemma 22. *Let $\mathcal{Q} \subseteq \mathrm{End}(E)$ a full-rank subring containing $\mathcal{O}$ and $\alpha \in \mathcal{Q} \backslash \{0\}$ a twisting endomorphism. Defining $\mathfrak{a} := (\mathcal{Q} \cdot \alpha) \cap \mathcal{O}$ and $\mathfrak{b}_c := \big(\mathrm{End}(E) \cdot c\alpha\big) \cap \mathcal{O}$, we have inclusions of $\mathcal{O}$-ideals*

$$\mathfrak{b}_f \subseteq \mathfrak{a} \subseteq \mathfrak{b}_1,$$

where the norm of the quotient $(\mathfrak{b}_1 : \mathfrak{b}_f)$ divides the squared conductor f^2.

Proof. The inclusions are obvious from $\mathrm{End}(E) \cdot f \subseteq \mathcal{Q} \subseteq \mathrm{End}(E)$. Moreover,

$$f\mathfrak{b}_1 = \big(f \cdot \mathrm{End}(E) \cdot \alpha\big) \cap (f \cdot \mathcal{O}) \subseteq \big(\mathrm{End}(E) \cdot f\alpha\big) \cap \mathcal{O} = \mathfrak{b}_f,$$

and the inclusions we have just established imply a chain of surjections

$$\mathfrak{b}_1/f\mathfrak{b}_1 \twoheadrightarrow \mathfrak{b}_1/\mathfrak{b}_f \twoheadrightarrow \mathfrak{b}_1/\mathfrak{a}$$

on the quotients of $\mathfrak{b}_1$. The first module in this sequence is clearly isomorphic to $\mathbb{Z}^2/f\mathbb{Z}^2$, therefore the index $[\mathfrak{b}_1 : \mathfrak{b}_f]$ must be a divisor of $|\mathbb{Z}^2/f\mathbb{Z}^2| = f^2$. □

Note that both ideals $\mathfrak{b}_1$ and $\mathfrak{b}_f$ from Lemma 22 would be correct outputs for a generalization of Algorithm 1 to proper subrings of $\mathrm{End}(E)$, but $\mathfrak{a}$ typically is not. However, the lemma still suggests an easy strategy for guessing $\mathfrak{b}_1$ after having obtained $\mathfrak{a}$ from the subring variant of Algorithm 1, at least when factoring f is feasible and there are not too many prime factors: In that case, one may simply brute-force through all ideals $\mathfrak{c} \subseteq \mathcal{O}$ of norm dividing f^2 and output $\mathfrak{ac}$ for each of them. The lemma guarantees that a correct such $\mathfrak{c}$ exists, since the ideal $(\mathfrak{b}_1 : \mathfrak{a})$ is a good choice. This procedure is summarized in Algorithm 3.

Algorithm 3: Twisting a curve using an endomorphism *sub*ring.

Input: a supersingular $E/\mathbb{F}_p$ and a rank-4 subring $\mathcal{Q} \subseteq \mathrm{End}(E)$ with $\mathcal{Q} \supseteq \mathcal{O}_E$.
Output: a set A of invertible ideals $\mathfrak{a} \subseteq \mathcal{O}_E$ such that $\exists\, \mathfrak{a} \in A$ with $[\mathfrak{a}]E = E^t$.

Find a non-zero element $\alpha \in \mathcal{Q}$ of the form $x\mathbf{i} + y\mathbf{ij}$.
Compute the ideal $\mathfrak{a} := \big(\mathcal{Q} \cdot \alpha\big) \cap \mathcal{O}_E$.
Determine $f = [\mathrm{End}(E) : \mathcal{Q}]$ as the (reduced) discriminant of $\mathcal{Q}$ divided by p.
Factor f and iterate through all ideals $\mathfrak{c} \subseteq \mathcal{O}$ of norm dividing f^2 to compute the set $A := \{\mathfrak{ac} \mid \mathfrak{c} \subseteq \mathcal{O} \text{ ideal}, \mathrm{N}(\mathfrak{c}) \mid f^2\}$.
Return A.

We can bound the size of the set A returned by the algorithm as follows: If the conductor f factors into primes as $f = \prod_{i=1}^{r} p_i^{e_i}$, then there are at most

$$\prod_{i=1}^{r} \binom{2e_i + 2}{2} \in O\big((\log f)^{2r}\big)$$

distinct $\mathcal{O}$-ideals of norm dividing f^2. Hence, if f is factorable in polynomial time and the number of distinct prime factors r is bounded by a constant, then Algorithm 3 takes polynomial time to output polynomially many ideals, and at least one of them is guaranteed to be correct.

4.3 Can We Do Better?

It is a natural question to ask whether one can tweak the KLPT quaternion-ideal algorithm [22] to simply output an ideal corresponding to an isogeny defined over $\mathbb{F}_p$, while preserving the main characteristics of the algorithm, namely the smoothness of the ideal that is returned and the (heuristic) polynomial runtime.

In this section, we argue that the answer is likely "no", at least for classical algorithms: More concretely, we show that such an algorithm can be used as a black-box oracle to construct, under a few mild assumptions, a polynomial-time algorithm for the discrete-logarithm problem in those imaginary-quadratic class groups where the prospective KLPT variant would apply. In contrast, the currently best known algorithm is only subexponential-time [18]. Thus, the basic conclusion appears to be that either our result is essentially optimal, or there exists an improved classical algorithm to compute class-group discrete logarithms in (at least) some cases.

In a sense, this is not surprising: The requirement that the output be generated by an ideal of the two-dimensional subring $\mathrm{End}_p(E)$ removes about the same amount of freedom as was "adjoined" when moving from $\mathbb{Q}(\sqrt{-p})$ to $B_{p,\infty}$ in the first place. In fact, the KLPT algorithm makes explicit constructive use of a quadratic subring of $B_{p,\infty}$ to achieve its functionality; an advantage that can be expected to cease when imposing strong restrictions on the output.

We formalize the situation as follows. Suppose we are given an algorithm $\mathcal{A}$ with the same interface as Algorithm 2, i.e., it takes as input two supersingular elliptic curves $E, E'/\mathbb{F}_p$ with the same $\mathbb{F}_p$-endomorphism ring $\mathcal{O}$, together with their full endomorphism rings, and outputs an ideal $\mathfrak{a} \subseteq \mathcal{O}$ such that $[\mathfrak{a}]E = E'$. In addition, our hypothetical algorithm $\mathcal{A}$ now guarantees that all prime factors of the returned ideal $\mathfrak{a}$ are elements of some polynomially-sized set $S_{\mathcal{O}}$, which may depend on the prime p or the ring $\mathcal{O}$ but not on the concrete input curves E and E'. For example, $S_{\mathcal{O}}$ might consist of the prime ideals of $\mathcal{O}$ with norm bounded by a polynomial in $\log p$.[6]

[6] Under GRH, Bach [2] proved that $\mathrm{cl}(\mathcal{O})$ is generated by prime ideals of norm less than $C(\log p)^2$ for an explicitly computable small constant C. It is not known unconditionally whether a polynomial bound on the norms suffices.

Then, Algorithm 5 can use such an oracle $\mathcal{A}$ to compute discrete logarithms in the subgroup of $\mathrm{cl}(\mathcal{O})$ generated by the subset $S_{\mathcal{O}}$ in expected polynomial time, assuming that querying $\mathcal{A}$ takes polynomial time. Note that the core of the reduction is Algorithm 4, which employs $\mathcal{A}$ to decompose class-group elements as a relation over the factor base $S_{\mathcal{O}}$, and those relations are subsequently used by Algorithm 5 in a generic and fairly standard index-calculus procedure.

A remark on notation: we make use of vectors and matrices indexed by finite sets I such as $S_{\mathcal{O}}$—in real implementations this would correspond to fixing an ordering of I and simply storing normal vectors or matrices of length $|I|$. We use the notation $|_{I'}$ to restrict a vector or matrix to the columns indexed by a subset $I' \subseteq I$.

Algorithm 4: Finding a class-group relation using $\mathcal{A}$.

Input: an oracle $\mathcal{A}$ as above, and an ideal $\mathfrak{a} \subseteq \mathcal{O}$ such that $[\mathfrak{a}] \in \langle [\mathfrak{s}] \mid \mathfrak{s} \in S_{\mathcal{O}} \rangle$.
Output: a vector $(e_{\mathfrak{s}} \mid \mathfrak{s} \in S_{\mathcal{O}}) \in \mathbb{Z}^{S_{\mathcal{O}}}$ such that $[\mathfrak{a}] = \left[\prod_{\mathfrak{s} \in S_{\mathcal{O}}} \mathfrak{s}^{e_{\mathfrak{s}}}\right]$.
Find a supersingular $E/\mathbb{F}_p$ with $\mathrm{End}_p(E) = \mathcal{O}$ and known $\mathrm{End}(E)$.
Apply KLPT to $\mathrm{End}(E) \cdot \mathfrak{a}$ to get an equivalent powersmooth left ideal $\mathcal{I}$.
Find the codomain $E' = [\mathfrak{a}]E$ by computing the isogeny defined by $\mathcal{I}$.
Compute $\mathrm{End}(E')$ as the right order of $\mathcal{I}$ in $B_{p,\infty}$.
Now query $\mathcal{A}$ to find an ideal $\mathfrak{c} \in \langle S_{\mathcal{O}} \rangle$ such that $[\mathfrak{c}]E = E' = [\mathfrak{a}]E$.
By assumption, $\mathfrak{c}$ is of the form $\prod_{\mathfrak{s} \in S_{\mathcal{O}}} \mathfrak{s}^{e_{\mathfrak{s}}}$.
Return that exponent vector $\underline{e} = (e_{\mathfrak{s}} \mid \mathfrak{s} \in S_{\mathcal{O}})$.

Lemma 23. *Algorithm 4 is correct. It takes polynomial time under the heuristic that the KLPT algorithm runs in polynomial time.*

Proof. Note that finding a curve E as desired is easy: first construct an arbitrary supersingular elliptic curve over $\mathbb{F}_p$ using [6], then potentially walk to the surface or floor of a 2-volcano. Next, note that the curve E' in fact equals $[\mathfrak{a}]E$, since $\mathrm{End}(E) \cdot \mathfrak{a}$ and $\mathfrak{a}$ define the same subgroup of E and $\mathcal{I}$ is equivalent as a left ideal to $\mathrm{End}(E) \cdot \mathfrak{a}$. Computing $\mathrm{End}(E')$ given $\mathcal{I}$ is easy linear algebra. Now, $\mathfrak{c}$ is a product of ideals in $S_{\mathcal{O}}$ by assumption on $\mathcal{A}$, and it must be equivalent to $\mathfrak{a}$ in $\mathrm{cl}(\mathcal{O})$ since the latter acts freely on $\mathcal{E}\ell\ell_p(\mathcal{O})$. In conclusion, Algorithm 4 indeed returns a correct relation vector for $\mathfrak{a}$ and takes polynomial time to do so. □

Using Algorithm 4, we can then follow the generic index-calculus procedure shown in Algorithm 5 to compute discrete logarithms in $\mathrm{cl}(\mathcal{O})$:

Lemma 24. *Algorithm 5 is correct and runs in expected polynomial time.*[7]

Proof sketch. It is not hard to check that the output of the algorithm is correct if it terminates; we thus only have to bound the expected runtime.

[7] Note that this does not require *any* assumptions on the output distribution of $\Delta(\mathsf{a})$, other than that the returned vectors are correct. (The algorithm still takes polynomial time if the oracle Δ only succeeds on an inverse polynomial fraction of inputs.).

Algorithm 5: Solving DLP using index calculus (generic).

Input:
- a generating set S of a finite abelian group G.
- an upper bound B on the cardinality $|G|$.
- elements $\mathsf{g}, \mathsf{h} \in G$ such that $\mathsf{h} \in \langle \mathsf{g} \rangle$.
- a probabilistic algorithm $\Delta\colon G \to \mathbb{Z}^S$, such that for all inputs $\mathsf{a} \in G$, we have $\|\Delta(\mathsf{a})\|_\infty < B$ and $\mathsf{a} = \prod_{\mathsf{b}\in S} \mathsf{b}^{\Delta(\mathsf{a})_\mathsf{b}}$.

Output: an integer x such that $\mathsf{g}^x = \mathsf{h}$.

Fix a large integer $H \gg B^{2|S|+1}$. (In practice, use much smaller H.)
Initialize empty matrices $M \in \mathbb{Z}^{0\times 2}$ and $L \in \mathbb{Z}^{0\times S}$.
For $n = 1, 2, \ldots$ **do**
  Pick integers u, v uniformly random in $\{-H, \ldots, H\}$.
  Invoke Δ to obtain a vector $\underline{e} \in \mathbb{Z}^S$ such that $\mathsf{g}^u\mathsf{h}^v = \prod_{\mathsf{b}\in S} \mathsf{b}^{e_\mathsf{b}}$.
  Append (u, v) to M as a new row. Append $\underline{e}$ to L as a new row.
  Compute a basis matrix $K \in \mathbb{Z}^{r\times n}$ of the left kernel of L, which is a lattice in $\mathbb{Z}^n$ of rank r.
  If the row span of $K \cdot M$ contains a vector of the form $(x, -1)$ **then**
    Return x.

Since the proof is rather technical, we will merely show the overall strategy. Note that it suffices to lower bound the success probability of the algorithm when $r = 2$ by a constant: Since $r \geq n - |S|$ throughout, it is evident that running $|S| + \alpha$ iterations of Algorithm 5 has success probability at least as big as $\lfloor \alpha/2 \rfloor$ independent executions of the modified algorithm. We thus want to lower bound the probability that two entries λ_1, λ_2 in the second column of $K \cdot M$ are coprime.

First, since Δ cannot distinguish from which scalars (u, v) the element $\mathsf{g}^u\mathsf{h}^v$ was obtained, the conditional distribution of each coefficient of M after fixing a certain oracle output $\underline{e}$ is close to uniform on $\{-H, \ldots, H\}$. As the lattice spanned by the rows of $K \cdot M$ is clearly independent of a basis choice, we may without loss of generality assume that the rows of K form a shortest basis of $\mathbb{Z}^r K$; using lattice techniques, one can then show that the norms of vectors in a shortest basis of $\mathbb{Z}^r K$ are upper bounded by $B^{2|S|}$. (This uses the bound on the size of integers returned by Δ.) Hence λ_i is a "small" coprime linear combination of random variables essentially uniform on $\{-H, \ldots, H\}$, which in turn implies that λ_i is close to uniform modulo all potential prime divisors. Thus the probability that $\gcd(\lambda_1, \lambda_2) = 1$ is lower bounded by a constant, similar to the well-known fact that the density of coprime pairs in $\mathbb{Z}^2$ is $\zeta(2)^{-1} = 6/\pi^2$. □

For concreteness, we briefly spell out how to instantiate Algorithm 5 for our particular application to $\mathrm{cl}(\mathcal{O})$. Clearly, Algorithm 4 will serve as the oracle Δ, so the factor base S equals the set $S_\mathcal{O}$ from Algorithm 4. In order to keep the representation sizes limited and to obtain unique representatives of ideal classes, the required products $\mathsf{g}^u\mathsf{h}^v$ should be computed using the square-and-multiply algorithm combined with reduction of binary quadratic forms; see [11] for more

context on the correspondence between quadratic forms and ideals (§ 7.B) and the notion of reduction (§ 2.A). To select the estimate B on the group order, recall the upper bound $h(\mathcal{O}) \in O(\sqrt{p} \log p)$ from the class number formula.

5 Vectorizing CM Curves

To the best of our knowledge, there exist two practical methods for constructing supersingular elliptic curves over a large finite field $\mathbb{F}_p$: either one reduces curves having CM by some order $\mathcal{R}$ in an imaginary quadratic field F modulo (appropriately chosen) primes that do not split in F, or one performs isogeny walks starting from known supersingular curves. As pointed out earlier, outside of trusted setup, the latter method is not suitable for most cryptographic applications. In this section we focus on the former method; additional details can be found in Bröker's paper [6] and the references therein. As we will see, from a security point of view, the situation is even more problematic in this case: we show that the vectorization problem associated with a CM-constructed supersingular elliptic curve over $\mathbb{F}_p$ admits a surprisingly easy and explicit solution.

In practice, when constructing supersingular elliptic curves over $\mathbb{F}_p$ one does not explicitly write down CM curves. Rather, one computes the Hilbert class polynomial $H_{\mathcal{R}}(T) \in \mathbb{Z}[T]$ for $\mathcal{R}$, which is a monic irreducible polynomial whose roots are the j-invariants of the curves having CM by $\mathcal{R}$. This polynomial can be computed effectively, although the existing methods are practical for orders having small discriminants only, one reason being that the degree of $H_{\mathcal{R}}(T)$ equals $h(\mathcal{R})$. The roots of $H_{\mathcal{R}}(T) \bmod p \in \mathbb{F}_p[T]$ are precisely those $j \in \overline{\mathbb{F}}_p$ which arise as the j-invariant of a supersingular elliptic curve obtained by reducing an elliptic curve having CM by $\mathcal{R}$. It is well-known that all these j-invariants are in fact elements of $\mathbb{F}_{p^2}$, i.e., the irreducible factors of $H_{\mathcal{R}}(T) \bmod p$ are at most quadratic. The linear factors then correspond to elliptic curves over $\mathbb{F}_p$.

Example 25. The Hilbert class polynomial for $\mathbb{Z}[\sqrt{-17}]$ is given by

$$\begin{aligned} H_{\mathbb{Z}[\sqrt{-17}]}(T) = \; & T^4 - 1782110400000T^3 - 7584369216000000000T^2 \\ & -3185070387200000000000T - 208929750630400000000000000, \end{aligned}$$

whose reduction modulo 83 factors as $(T-28)(T-50)(T^2+7T+73)$. This gives rise to two pairs of quadratic twists of elliptic curves over $\mathbb{F}_{83}$ that appear as the reduction modulo 83 of a curve with CM by $\mathbb{Z}[\sqrt{-17}]$.

The main result of this section is the following theorem; for conciseness, our focus lies on the setting where $p \equiv 3 \pmod 4$ and where

$$\mathbb{Z}[\sqrt{-\ell}] \subseteq \mathcal{R} \subseteq \mathbb{Q}(\sqrt{-\ell})$$

for some odd prime number ℓ, i.e., we want our CM curves to come equipped with an endomorphism Ψ satisfying $\Psi \circ \Psi = [-\ell]$. This leaves us with two options for $\mathcal{R}$, namely $\mathbb{Z}[\sqrt{-\ell}]$ and $\mathbb{Z}[(1+\sqrt{-\ell})/2]$. In Remark 32 we will briefly comment on how to locate curves having CM by more general imaginary quadratic orders.

Theorem 26. *Let $p \equiv 3 \pmod 4$ and $\ell < (p+1)/4$ be primes with $\left(\frac{-p}{\ell}\right) = 1$.*

(i) If $\ell \equiv 1 \pmod 4$ then

$$H_{\mathbb{Z}[\sqrt{-\ell}]}(T) \bmod p$$

has precisely two $\mathbb{F}_p$-rational roots, both corresponding to a pair of quadratic twists of supersingular elliptic curves. One pair is contained in $\mathcal{Ell}_p(\mathbb{Z}[\sqrt{-p}])$ while the other pair is contained in $\mathcal{Ell}_p(\mathbb{Z}[(1+\sqrt{-p})/2])$.

(ii) If $\ell \equiv 3 \pmod 4$ then both

$$H_{\mathbb{Z}[(1+\sqrt{-\ell})/2]}(T) \bmod p \qquad \text{and} \qquad H_{\mathbb{Z}[\sqrt{-\ell}]}(T) \bmod p$$

have exactly one $\mathbb{F}_p$-rational root each, in both cases corresponding to a pair of quadratic twists of elliptic curves. The first such pair is contained in $\mathcal{Ell}_p(\mathbb{Z}[\sqrt{-p}])$, while the other pair is contained in $\mathcal{Ell}_p(\mathbb{Z}[(1+\sqrt{-p})/2])$.

(iii) Let $\mathcal{O} \in \{\mathbb{Z}[\sqrt{-p}], \mathbb{Z}[(1+\sqrt{-p})/2]\}$ and let $E, E^t \in \mathcal{Ell}_p(\mathcal{O})$ be a pair of supersingular elliptic curves over $\mathbb{F}_p$ arising as above. Up to order, this pair is given by the curves

$$[\mathfrak{l}]^{(h(\mathcal{O})-1)/2} E_0 \qquad \text{and} \qquad [\mathfrak{l}]^{(h(\mathcal{O})+1)/2} E_0 \tag{4}$$

for any prime ideal $\mathfrak{l} \subseteq \mathcal{O}$ lying above ℓ. Here $E_0 \colon y^2 = x^3 \pm x$ is the unique curve with j-invariant 1728 in $\mathcal{Ell}_p(\mathcal{O})$.

This theorem can be seen as a vast generalization of (3) from Example 19, where we dealt with the reduction modulo p of the curve $E \colon y^2 = x^3 + 1$ over $\mathbb{Q}$ having CM by the ring of Eisenstein integers $\mathbb{Z}[e^{2\pi \mathbf{i}/3}] = \mathbb{Z}[(1+\sqrt{-3})/2]$. Up to twisting it is the only such curve: the Hilbert class polynomial for $\mathbb{Z}[(1+\sqrt{-3})/2]$ is just T. An endomorphism Ψ satisfying $\Psi^2 = -3$ can be constructed as $2\omega + 1$, where ω is the automorphism $E \to E,\ (x, y) \mapsto (e^{2\pi \mathbf{i}/3} x, y)$.

One particularly interesting range of parameters satisfying the stated assumptions is where

- $p = 4\ell_1 \ell_2 \cdots \ell_r - 1$ is a CSIDH prime with $r \geq 2$, and
- $\ell = \ell_i$ for some $i \in \{1, 2, \ldots, r\}$.

If $r = 1$ then $\ell_1 = (p+1)/4$, so Theorem 26 can no longer be applied. However, the reasons for excluding the boundary case $\ell = (p+1)/4$ are rather superficial and the statement remains largely valid in this case (the exclusion is related to the possible occurrence of $j = 1728$ as a root of $H_{\mathcal{R}}(T) \bmod p$, which comes with some subtleties in terms of quadratic twisting; see the proof).

5.1 Twisting Endomorphisms from Deuring Reduction

Before proceeding to the proof of Theorem 26, we discuss Deuring lifting and reduction, with a focus on how the endomorphism Ψ behaves under reduction.

Theorem 27 (Deuring's reduction theorem). *Let p be a prime number and let E be an elliptic curve over a number field K which has CM by some order $\mathcal{R}$ in an imaginary quadratic number field F. Let $\mathfrak{p}$ be a prime of K above p at which E has good reduction. Then $E \bmod \mathfrak{p}$ is supersingular if and only if p ramifies or is inert in F.*

Proof. This is part of [23, Thm. 12 of Ch. 13]. □

When applying this to an elliptic curve E/K having CM by our order $\mathcal{R} \subseteq \mathbb{Q}(\sqrt{-\ell})$ from above, the endomorphism Ψ satisfying $\Psi \circ \Psi = [-\ell]$ reduces modulo $\mathfrak{p}$ to an endomorphism ψ which also satisfies $\psi \circ \psi = [-\ell]$. This is because reduction modulo $\mathfrak{p}$ induces an (injective) homomorphism of endomorphism rings; see for instance [23, § 2 of Ch. 9]. The following proposition gives sufficient conditions for ψ to be a twisting endomorphism.

Proposition 28. *Assume $K = \mathbb{Q}(j(E))$, $p > 2$ and $\ell \leq (p+1)/4$. If $E \bmod \mathfrak{p}$ is supersingular and $j(E \bmod \mathfrak{p}) \in \mathbb{F}_p$ then $\deg \mathfrak{p} = 1$ and*

$$\pi_{E \bmod \mathfrak{p}} \circ \psi = -\psi \circ \pi_{E \bmod \mathfrak{p}}, \tag{5}$$

i.e., ψ anticommutes with the p-power Frobenius endomorphism of $E \bmod \mathfrak{p}$.

The proof of this proposition relies on the following observation:

Lemma 29. *Let α be an algebraic integer and $K = \mathbb{Q}(\alpha)$. Consider a prime number p and a prime ideal $\mathfrak{p} \subseteq \mathcal{O}_K$ above p. If $\mathbb{F}_p(\alpha \bmod \mathfrak{p}) \subsetneq \mathcal{O}_K/\mathfrak{p}$, then p divides the discriminant of the minimal polynomial $f(x) \in \mathbb{Z}[x]$ of α over $\mathbb{Q}$.*

Proof. If p does not divide the discriminant of $f(x)$, then

$$\mathfrak{p} = \big(p, g(\alpha)\big),$$

where $g(x) \in \mathbb{Z}[x]$ is a monic polynomial of degree $\deg \mathfrak{p}$ whose reduction modulo p is an irreducible factor in $\mathbb{F}_p[x]$ of $f(x) \bmod p$; this is a well-known fact, see e.g. [24, Thm. 27]. But this implies that $\alpha \bmod \mathfrak{p}$ is a generator of $\mathcal{O}_K/\mathfrak{p}$ over $\mathbb{F}_p$, so the lemma follows by contradiction. □

Proof (of Proposition 28). The minimal polynomial of $j(E)$ over $\mathbb{Q}$ is precisely the Hilbert class polynomial $H_{\mathcal{R}}(T)$ for $\mathcal{R}$. The field $H = \mathbb{Q}(\sqrt{-\ell}, j(E))$ is a quadratic extension of K known as the ring class field for $\mathcal{R}$, see [11, proof of Prop. 1.32]. If $\mathcal{R}$ is a maximal order, then this is better known as the Hilbert class field.

Using that $\ell \leq (p+1)/4$, we see that p does not ramify in $\mathbb{Q}(\sqrt{-\ell})$, hence it must be inert by our assumption that $E \bmod \mathfrak{p}$ is supersingular. This implies that $p\mathcal{O}_H$ splits as a product of prime ideals $\mathfrak{P}$ of degree 2, see [11, Cor. 5.25] for a proof in case $\mathcal{R}$ is a maximal order and [11, proof of Prop. 9.4] for the general case (this is where we use the assumption $p > 2$). Our prime ideal $\mathfrak{p}$ is necessarily dominated by such a $\mathfrak{P}$, so it follows that

– either $\deg \mathfrak{p} = 1$, in which case $\mathfrak{p}$ must be inert in H, i.e., $\mathfrak{p}\mathcal{O}_H = \mathfrak{P}$,
– or $\deg \mathfrak{p} = 2$, in which case $\mathfrak{p}$ must split in H.

But the latter option would imply that

$$\mathbb{F}_p(j(E) \bmod \mathfrak{p}) = \mathbb{F}_p(j(E \bmod \mathfrak{p})) = \mathbb{F}_p \subsetneq \mathcal{O}_K/\mathfrak{p}$$

and therefore, in view of Lemma 29, it would follow that p divides the discriminant of $H_{\mathcal{R}}(T)$. This is impossible: by Gross–Zagier [17, p. 195] the primes p dividing the discriminant of $H_{\mathcal{R}}(T)$ cannot be larger than the absolute value of the discriminant of $\mathcal{R}$, which is at most 4ℓ.

We have thus established that $\deg \mathfrak{p} = 1$. Now let Σ be the non-trivial automorphism of H over K. From [23, § 4 of Ch. 10] we see that Ψ is not defined over K and therefore $\Psi^\Sigma = -\Psi$. But Σ necessarily descends to the Frobenius automorphism σ of $\mathcal{O}_H/\mathfrak{P} \cong \mathbb{F}_{p^2}$ over $\mathcal{O}_K/\mathfrak{p} \cong \mathbb{F}_p$, from which it follows that $\psi^\sigma = -\psi$. This implies (5) and thereby concludes the proof. □

We remark that the last part of the preceding proof mimics the proof of [15, Prop. 6.1]. However, the statement of [15, Prop. 6.1] is lacking an assumption on $\deg \mathfrak{p}$. For instance, in our case, if $\deg \mathfrak{p} = 2$ and therefore $\mathfrak{p}$ splits in H, the reasoning fails because the extension $\mathcal{O}_H/\mathfrak{P}$ over $\mathcal{O}_K/\mathfrak{p}$ becomes trivial. And indeed, in these cases it may happen that the reduction of Ψ mod $\mathfrak{p}$ does *not* anticommute with Frobenius:

Example 30. The discriminant of the Hilbert class polynomial for $\mathbb{Z}[\sqrt{-29}]$ is divisible by 83. More precisely, its reduction modulo 83 factors as $T(T-50)(T-67)^2(T^2+7T+73)$. One can verify that inside $K = \mathbb{Q}[T]/(H_{\mathbb{Z}[\sqrt{-29}]}(T))$, we have

$$83\mathcal{O}_K = (83, T)(83, T-50)(83, T^2-7)(83, T^2+7T+73),$$

where the third factor is a degree-2 prime ideal $\mathfrak{p}$ modulo which T reduces to 67; note that $67^2 \equiv 7 \pmod{83}$. So in this case we have $\mathbb{F}_p(T \bmod \mathfrak{p}) \subsetneq \mathcal{O}_K/\mathfrak{p}$.

Let E be any of the two elliptic curves over $\mathbb{F}_{83}$ having j-invariant 67. By exhaustive search through the possible kernels of order 29, one can check that E admits four distinct automorphisms squaring to $[-29]$. These appear in the form $\pm\psi, \pm\psi^\sigma$, where as in the proof of Proposition 28 we use σ to denote the action of the p-power Frobenius. In particular ψ does not anticommute with π_E. Nevertheless, by Deuring's lifting theorem (recalled below), the pair (E, ψ) must arise as the reduction of some CM curve along with an endomorphism Ψ satisfying $\Psi \circ \Psi = [-29]$. (Note: this also applies to the pair (E, ψ^σ), which is reflected in the fact that 67 appears as a double root of $H_{\mathbb{Z}[\sqrt{-\ell}]}(T)$ mod 83.)

Theorem 31 (Deuring's lifting theorem). *Let $E/\overline{\mathbb{F}}_p$ be an elliptic curve and let $\alpha \in \mathrm{End}(E)$. There exists an elliptic curve E' over a number field K along with an endomorphism $\alpha' \in \mathrm{End}(E')$ and a prime $\mathfrak{p}$ of K above p at which E' has good reduction, such that $E' \bmod \mathfrak{p}$ is isomorphic to E and such that α' reduces to α modulo $\mathfrak{p}$.*

Proof. See [23, Thm. 14 of Ch. 13]. □

5.2 Proof of Theorem 26

Proof (of Theorem 26). Using quadratic reciprocity one checks that

$$\left(\frac{-p}{\ell}\right) = 1 \iff \left(\frac{-\ell}{p}\right) = -1,$$

from which we see that p is inert in $\mathbb{Q}(\sqrt{-\ell})$. Hence a curve with CM by $\mathbb{Z}[\sqrt{-\ell}]$ has supersingular reduction modulo p and therefore the $\mathbb{F}_p$-rational roots of the Hilbert class polynomial

$$H_{\mathbb{Z}[\sqrt{-\ell}]}(T) \bmod p$$

should correspond to pairs of quadratic twists in either the floor $\mathcal{E}\ell\ell_p(\mathbb{Z}[\sqrt{-p}])$ or the surface $\mathcal{E}\ell\ell_p(\mathbb{Z}[(1+\sqrt{-p})/2])$. If $\ell \equiv 3 \pmod 4$, then the same conclusions apply to $\mathbb{Z}[(1+\sqrt{-\ell})/2]$.

As a side note, we remark that $\ell < (p+1)/4$ implies that $y^2 = x^3 \pm x$ does not admit any twisting endomorphisms of norm ℓ, which is easy to elaborate from [25, Prop. 3.1]. In view of Proposition 28, we therefore see that the $\mathbb{F}_p$-rational roots of the Hilbert class polynomial never include 1728. Hence by Lemma 1 there is no ambiguity in what is meant by "pairs of quadratic twists". (Apart from this ambiguity, the theorem remains true under the weaker assumption $\ell \leq (p+1)/4$.)

We first claim that $\mathcal{E}\ell\ell_p(\mathbb{Z}[\sqrt{-p}])$ and $\mathcal{E}\ell\ell_p(\mathbb{Z}[(1+\sqrt{-p})/2])$ both contain *at most* one such pair E, E^t. Indeed, using Proposition 28 we see that E comes equipped with a twisting endomorphism ψ of degree ℓ, which by Lemma 13 corresponds to an $\mathbb{F}_p$-rational degree-ℓ isogeny $E \to E^t$. Its kernel is necessarily of the form $E[\mathfrak{l}]$ for some prime ideal $\mathfrak{l}$ above ℓ, i.e., we must have $E^t = [\mathfrak{l}]E$. But then we can solve the vectorization problem $E = [\mathfrak{a}]E_0$: from Lemma 18 we get that $[\mathfrak{a}] = [\mathfrak{l}]^{(h(\mathcal{O})-1)/2}$. Since the pair

$$\left\{[\mathfrak{l}]^{(h(\mathcal{O})-1)/2}, \quad [\mathfrak{l}]^{(h(\mathcal{O})+1)/2} = [\bar{\mathfrak{l}}]^{(h(\mathcal{O})-1)/2}\right\}$$

does not depend on the choice of $\mathfrak{l}$, this shows that the pair $\{E, E^t\}$ is fully characterized by ℓ, implying the claim. At the same time this proves (iii).

Next, let us explain why $\mathcal{E}\ell\ell_p(\mathbb{Z}[\sqrt{-p}])$ and $\mathcal{E}\ell\ell_p(\mathbb{Z}[(1+\sqrt{-p})/2])$ contain *at least* one such pair E, E^t. We remark that this comes for free if $\ell \equiv 3 \pmod 4$, since in this case the Hilbert class polynomials for $\mathbb{Z}[\sqrt{-\ell}]$ and $\mathbb{Z}[(1+\sqrt{-\ell})/2]$ have odd degree and split over $\mathbb{F}_{p^2}$, their roots being supersingular j-invariants: hence they must admit at least one $\mathbb{F}_p$-rational root. In general, we can reverse the reasoning from the previous paragraph and *define* E, E^t using (4), for some choice of prime ideal $\mathfrak{l}$ above ℓ. In particular $E^t = [\mathfrak{l}]E$, which provides us with an $\mathbb{F}_p$-rational degree-ℓ isogeny $\varphi\colon E \to E^t$, which we use to construct an endomorphism $\psi = \tau_{E^t} \circ \varphi$ of E that is not $\mathbb{F}_p$-rational. In contrast, it is easily verified that $\psi \circ \psi$ is $\mathbb{F}_p$-rational. Therefore the minimal polynomial of ψ cannot admit a non-zero linear term, i.e., $\psi \circ \psi$ must be a scalar-multiplication map, necessarily of the form $[\pm\ell]$. By Deuring's lifting theorem E can be lifted to an elliptic curve over a number field carrying an endomorphism Ψ whose reduction modulo a suitable prime above p yields ψ. Since Ψ must belong to an imaginary quadratic ring we see that $\Psi \circ \Psi = [-\ell]$ as wanted.

Altogether this proves (i), while for (ii) it leaves us with the task of showing that if $\ell \equiv 3 \pmod 4$, then the unique $\mathbb{F}_p$-rational root of

$$H_{\mathbb{Z}[(1+\sqrt{-\ell})/2]}(T) \bmod p$$

corresponds to a pair of elliptic curves $\{E, E^t\}$ with endomorphism ring $\mathbb{Z}[\sqrt{-p}]$. Equivalently, we need to show that such curves admit a unique $\mathbb{F}_p$-rational point of order 2, rather than three such points. To this end, let $P \in E$ be an $\mathbb{F}_p$-rational point of order 2 and let φ be the endomorphism of E corresponding to $(1+\sqrt{-\ell})/2$. Proposition 28 implies that $\pi_E \circ \varphi = \overline{\varphi} \circ \pi_E$, where $\overline{\varphi}$ corresponds to $(1-\sqrt{-\ell})/2$. But then clearly $(\varphi + \overline{\varphi})(P) = P \neq \infty$, which implies that $\overline{\varphi}(P) \neq \varphi(P)$ and therefore that $\pi_E(\varphi(P)) \neq \varphi(P)$, i.e., $\varphi(P)$ is a non-rational point of order 2. This concludes the proof. □

Remark 32. The above ideas can be generalized to locate reductions mod p of CM curves carrying an endomorphism Ψ such that $\Psi \circ \Psi = [-\ell_1 \ell_2 \cdots \ell_s]$, where the $\ell_i \leq (p+1)/4$ are distinct odd primes for which

$$\left(\frac{-\ell_1 \ell_2 \cdots \ell_s}{p} \right) = -1. \tag{6}$$

We did not elaborate this in detail, but assume for instance that each ℓ_i splits in $\mathbb{Q}(\sqrt{-p})$; note that this implies (6). Letting $\mathcal{O} \in \{\mathbb{Z}[\sqrt{-p}], \mathbb{Z}[(1+\sqrt{-p})/2]\}$, one expects that 2^{s-1} pairs E, E^t in $\mathcal{E}\ell\ell_p(\mathcal{O})$ can be obtained as the reduction mod p of an elliptic curve carrying such an endomorphism Ψ. Fixing for each $i = 1, 2, \ldots, s$ a prime ideal $\mathfrak{l}_i \subseteq \mathcal{O}$ of norm ℓ_i, these pairs are characterized by

$$E^t = [\mathfrak{l}_1][\mathfrak{l}_2]^{e_2}[\mathfrak{l}_2]^{e_3} \cdots [\mathfrak{l}_s]^{e_s} E$$

with $(e_2, e_3, \ldots, e_s) \in \{\pm 1\}^{s-1}$. As before, an application of Lemma 18 then solves the corresponding vectorization problems.

Acknowledgements. The authors would like to thank Benjamin Wesolowski, Robert Granger, Christophe Petit, and Ben Smith for interesting discussions regarding this work, and Lixia Luo for pointing out an error in an earlier version of Lemma 22, as well as a few smaller mistakes. Thanks to Daniel J. Bernstein for providing key insights regarding the proof of Lemma 24.

References

1. Arpin, S., et al.: Adventures in Supersingularland. Cryptology ePrint Archive 2019/1056 (2018). https://ia.cr/2019/1056
2. Bach, E.: Explicit bounds for primality testing and related problems. Math. Comput. **55**(191), 355–380 (1990)
3. Bernstein, D.J., Hamburg, M., Krasnova, A., Lange, T.: Elligator: elliptic-curve points indistinguishable from uniform random strings. In: ACM Conference on Computer and Communications Security, pp. 967–980. ACM (2013). https://ia.cr/2013/325

4. Beullens, W., Kleinjung, T., Vercauteren, F.: CSI-FiSh: efficient isogeny based signatures through class group computations. In: Galbraith, S.D., Moriai, S. (eds.) ASIACRYPT 2019. LNCS, vol. 11921, pp. 227–247. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-34578-5_9
5. Bosma, W., Stevenhagen, P.: On the computation of quadratic 2-class groups. J. de Théorie des Nombres de Bordeaux **8**(2), 283–313 (1996)
6. Bröker, R.: Constructing supersingular elliptic curves. J. Comb. Number Theory **1**(3), 273–469 (2009)
7. 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. LNCS, vol. 11274, pp. 395–427. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-03332-3_15
8. Charles, D.X., Lauter, K.E., Goren, E.Z.: Cryptographic hash functions from expander graphs. J. Cryptol. **22**(1), 93–113 (2009). https://ia.cr/2006/021
9. Conrad, K.: The conductor ideal. Expository paper. https://kconrad.math.uconn.edu/blurbs/gradnumthy/conductor.pdf
10. Couveignes, J.-M.: Hard homogeneous spaces. IACR Cryptology ePrint Archive 2006/291 (1997). https://ia.cr/2006/291
11. Cox, D.A.: Primes of the Form x^2+ny^2: Fermat, Class Field Theory, and Complex Multiplication. Pure Applied Mathematics, 2nd edn. Wiley, Hoboken (2013)
12. Delfs, C., Galbraith, S.D.: Computing isogenies between supersingular elliptic curves over $\mathbb{F}_p$. Des. Codes Cryptogr. **78**(2), 425–440 (2016). https://arxiv.org/abs/1310.7789
13. Eisenträger, K., Hallgren, S., Lauter, K., Morrison, T., Petit, C.: Supersingular isogeny graphs and endomorphism rings: reductions and solutions. In: Nielsen, J.B., Rijmen, V. (eds.) EUROCRYPT 2018. LNCS, vol. 10822, pp. 329–368. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-78372-7_11
14. Galbraith, S., Panny, L., Smith, B., Vercauteren, F.: Quantum equivalence of the DLP and CDHP for group actions. Cryptology ePrint Archive 2018/1199 (2018). https://ia.cr/2018/1199
15. Galbraith, S., Rotger, V.: Easy decision Diffie-Hellman groups. LMS J. Comput. Math. **7**, 201–218 (2004). https://ia.cr/2004/070
16. Galbraith, S.D., Petit, C., Shani, B., Ti, Y.B.: On the security of supersingular isogeny cryptosystems. In: Cheon, J.H., Takagi, T. (eds.) ASIACRYPT 2016. LNCS, vol. 10031, pp. 63–91. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53887-6_3
17. Gross, B.H., Zagier, D.B.: On singular moduli. J. für die Reine und Angewandte Mathematik **355**, 191–220 (1985)
18. Hafner, J.L., McCurley, K.S.: A rigorous subexponential algorithm for computation of class groups. J. Am. Math. Soc. **2**, 837–850 (1989)
19. Jao, D., De Feo, L.: Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. In: Yang, B.-Y. (ed.) PQCrypto 2011. LNCS, vol. 7071, pp. 19–34. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-25405-5_2
20. Kirschmer, M., Voight, J.: Algorithmic enumeration of ideal classes for quaternion orders. SIAM J. Comput. **39**(5), 1714–1747 (2010). https://arxiv.org/abs/0808.3833
21. Kitaev, A.Y.: Quantum measurements and the Abelian stabilizer problem. Electron. Colloquium Comput. Complex. (ECCC) **3**(3) (1996). https://eccc.hpi-web.de/eccc-reports/1996/TR96-003

22. Kohel, D., Lauter, K., Petit, C., Tignol, J.-P.: On the quaternion ℓ-isogeny path problem. LMS J. Comput. Math. **17**(Suppl. A), 418–432 (2014). https://ia.cr/2014/505
23. Lang, S.: Elliptic Functions. Graduate Texts in Mathematics, vol. 112. Springer, Heidelberg (1987). https://doi.org/10.1007/978-1-4612-4752-4. With an appendix by John Tate
24. Marcus, D.A.: Number Fields. Universitext, 2nd edn. Springer, Heidelberg (2018). https://doi.org/10.1007/978-1-4684-9356-6. With a foreword by Barry Mazur
25. McMurdy, K.: Explicit representation of the endomorphism rings of supersingular elliptic curves (2014). Preprint. https://phobos.ramapo.edu/~kmcmurdy/research/McMurdy-ssEndoRings.pdf
26. National Institute of Standards and Technology: Post-Quantum Cryptography Standardization, December 2016. https://csrc.nist.gov/Projects/Post-Quantum-Cryptography/Post-Quantum-Cryptography-Standardization
27. Onuki, H., Takagi, T.: On collisions related to an ideal class of order 3 in CSIDH. Cryptology ePrint Archive 2019/1202 (2019). https://ia.cr/2019/1202
28. Schnorr, C.-P., Euchner, M.: Lattice basis reduction: improved practical algorithms and solving subset sum problems. Math. Program. **66**, 181–199 (1994)
29. Schoof, R.: Elliptic curves over finite fields and the computation of square roots mod p. Math. Comput. **44**(170), 483–494 (1985)
30. Silverman, J.H.: The Arithmetic of Elliptic Curves. Graduate Texts in Mathematics, vol. 106, 2nd edn. Springer, Heidelberg (2009). https://doi.org/10.1007/978-0-387-09494-6
31. Waterhouse, W.C.: Abelian varieties over finite fields. Annales scientifiques de l'École Normale Supérieure **2**, 521–560 (1969)