One-way functions and malleability oracles: Hidden shift attacks on isogeny-based protocols

Péter Kutas[1], Simon-Philipp Merz[2], Christophe Petit[3,1], and Charlotte Weitkämper[1]

[1] University of Birmingham, UK
[2] Royal Holloway, University of London, UK
[3] Université libre de Bruxelles, Belgium

Abstract. Supersingular isogeny Diffie-Hellman key exchange (SIDH) is a post-quantum protocol based on the presumed hardness of computing an isogeny between two supersingular elliptic curves given some additional torsion point information. Unlike other isogeny-based protocols, SIDH has been widely believed to be immune to subexponential quantum attacks because of the non-commutative structure of the endomorphism rings of supersingular curves.

We contradict this commonly believed misconception in this paper. More precisely, we highlight the existence of an abelian group action on the SIDH key space, and we show that for sufficiently *unbalanced* and *overstretched* SIDH parameters, this action can be efficiently computed (heuristically) using the torsion point information revealed in the protocol. This reduces the underlying hardness assumption to a hidden shift problem instance which can be solved in quantum subexponential time.

We formulate our attack in a new framework allowing the inversion of one-way functions in quantum subexponential time provided a malleability oracle with respect to some commutative group action. This framework unifies our new attack with earlier subexponential quantum attacks on isogeny-based protocols, and it may be of further interest for cryptanalysis.

1 Introduction

The hardness of solving mathematical problems such as integer factorization or the computation of discrete logarithms in finite fields and elliptic curve groups guarantees the security of most currently deployed cryptographic protocols. However, these classical problems can be solved efficiently using quantum algorithms. Quantum computers with sufficient processing power to threaten cryptographic primitives currently in use do presumably not yet exist, but progress towards their realization is being made. The possibility of large scale quantum computers and the need for long-term security in some applications necessitate the development of quantum-secure cryptographic algorithms.

Different approaches to attain quantum resistance are based on problems in lattices, codes, multivariate polynomials over finite fields, and elliptic curve isogenies. Within the field of post-quantum cryptography, isogeny-based cryptography is a relatively new area which is particularly interesting due to the small key sizes required. The main problem underlying this branch of post-quantum cryptography is to find an isogeny $\varphi : E_1 \to E_2$ between two given isogenous elliptic curves E_1 and E_2 over some finite field $\mathbb{F}_q$.

An early isogeny-based cryptographic system utilizing *ordinary* elliptic curves was proposed by Couveignes but at first only circulated privately [7]. Meanwhile, the first construction using *supersingular* curves was a hash function developed by Charles, Lauter and Goren [4]. Later, Rostovtsev and Stolbunov independently rediscovered and further developed Couveignes' construction [26]. In 2010, Childs, Jao and Soukharev [5] showed how to break this scheme in quantum subexponential time using a reduction to an instance of abelian hidden shift problem. While this attack is tolerable for sufficiently large parameters, the main drawback of the Couveignes-Rostovtsev-Stolbunov (CRS) construction is its unacceptable lack of speed. Adapting the CRS scheme to supersingular elliptic curves, Castryck et al. managed to eliminate most of the performance issues allowing for larger practical parameters when introducing CSIDH [3]. While it is known that CSIDH can be attacked in quantum subexponential time, there have been several works on establishing its concrete security levels [2, 21].

The attack due to Childs, Jao and Soukharev crucially relies on the commutativity of the ideal class groups acting on the endomorphism rings of the relevant elliptic curves over $\mathbb{F}_q$. This motivated Jao and De Feo [14] to consider the full isogeny graph of supersingular elliptic curves, whose endomorphism rings are maximal orders in a quaternion algebra (in particular, the endomorphism rings are non-commutative). The result of their work, the *Supersingular Isogeny Diffie-Hellman* (SIDH) key agreement scheme, underlies the SIKE submission to NIST's post-quantum standardization process [1, 13].

The hard problem SIDH is based on is to find an isogeny between two isogenous curves, further given the images of certain torsion points under this isogeny. The best known way to break SIDH with balanced parameters on both classical and a quantum computers is a claw-finding approach on the isogeny graph [15] which does not use any torsion point information. Yet, the supply of this additional public information has fueled cryptanalytic research. It has been shown that the torsion point information can be used in *active* attacks [10] or when parameters are sufficiently overstretched [19, 22]. However, a widespread misconception amongst cryptographers assumes that due to SIDH's non-commutative nature there is no quantum attack reducing the SIDH problem to an abelian hidden shift problem. In particular, many believe that no reasonable variant of Childs-Jao-Soukharev's attack applies in the supersingular case [14, p. 18, Section 5].

Our contributions. We provide a new quantum attack on overstretched SIDH which uses a reduction of the underlying computational problem to an injective abelian hidden shift problem. This can be solved in quantum subexponential time and thus disproves the common misbelief mentioned above.

Let $\varphi : E_0 \to E_0/K$ be a secret isogeny that an attacker wishes to recover. As in SIDH, let E_0, E_0/K, $\deg(\varphi)$, and some torsion point images under the secret isogeny be known publicly. The idea underlying our cryptanalysis is to construct an abelian group G of E_0-endomorphisms acting freely and transitively on certain cyclic subgroups of E_0. These subgroups are kernels of $\deg(\varphi)$-isogenies, and therefore they can be mapped to supersingular elliptic curves $\deg(\varphi)$-isogenous to E_0. The group action of G can then be understood as an action on the curves. Forcing the endomorphisms in G to be of a certain degree, the public torsion point information allows an adversary to compute the action on E_0/K efficiently under some heuristics. Finally, solving an abelian hidden shift problem of two functions mapping G to a set of curves $\deg(\varphi)$-isogenous to E_0 containing E_0/K enables an attacker to recover K and therefore φ. We stress that this is a novel way of exploiting torsion point information.

While this attack does not threaten SIDH with balanced parameter sets as originally proposed by Jao and De Feo [14] and used in SIKE [13], it shows that an attack using a hidden shift algorithm is possible despite SIDH's non-commutative nature.

We describe our new attack as a special instance in a more general setting. This allows us to unify our new cryptanalysis with other quantum attacks on isogeny-based schemes such as the one due to Childs, Jao and Soukharev [5] constructing isogenies between ordinary curves, or a similar application of quantum hidden shift algorithms to CSIDH [2, 3, 21].

This framework might be of interest beyond isogeny-based cryptography. To define one of the key properties required, we introduce the notion of a *malleability oracle* for a function with respect to some group action. Under some additional assumptions, access to this oracle is sufficient to compute preimages of the function via solving a hidden shift problem.

Outline. In Section 2, we provide an overview of the notations we use, we recall mathematical background for isogeny-based cryptography, and we review quantum algorithms used in our attack. In Section 3, we present our general framework, namely sufficient conditions for computing preimages of one-way functions via reduction to a hidden shift problem, and then present our new attack on overstretched SIDH in Section 4. In Section 5, we additionally instantiate our general framework with the attack of Childs, Jao and Soukharev and its generalization to CSIDH. We conclude the paper in Section 6 with a discussion of potential improvements and future work.

2 Preliminaries

In this section, we introduce terminology and notation, and we recall relevant background on isogeny-based protocols and quantum algorithms.

2.1 Terminology

We call a function $\mu : \mathbb{N} \to \mathbb{R}$ *negligible* if for every positive integer c there exists an integer N_c such that $|\mu(x)| < \frac{1}{x^c}$ for every $x > N_c$. We call an algorithm *effi-*

cient if the execution time is bounded by a polynomial in the security parameter of the underlying cryptographic scheme. Given any function, by having *oracle access* to this function we mean that it is feasible to evaluate the function at any possible element in an efficient way. In particular, we assume that the oracle acts like a black box such that one query with an element from the domain outputs the corresponding value of the function.

Further, we call a function $f : \{0,1\}^* \to \{0,1\}^*$ *one-way*, if f can be computed by a polynomial-time algorithm, but for all polynomial-time randomized algorithms F, all positive integers c and all sufficiently large $n = \text{length}(x)$, $\Pr[f(F(f(x))) = f(x)] < n^{-c}$, where the probability is taken over the choice of x from the discrete uniform distribution on $\{0,1\}^n$, and the randomness of F.

2.2 Mathematical background on isogenies

For more complete introductions to elliptic curves and to isogeny-based cryptography we refer to Silverman [28] and De Feo [8], respectively.

Let $\mathbb{F}_q$ be a finite field of characteristic p. In the following we assume $p \geq 3$ and therefore an elliptic curve E over $\mathbb{F}_q$ can be defined by its short Weierstrass form

$$E(\mathbb{F}_q) = \{(x,y) \in \mathbb{F}_q^2 \mid y^2 = x^3 + Ax + B\} \cup \{\mathcal{O}_E\}$$

where $A, B \in \mathbb{F}_q$ and $\mathcal{O}_E$ is the point $(X : Y : Z) = (0 : 1 : 0)$ on the associated projective curve $Y^2Z = X^3 + AXZ^2 + BZ^3$. The set of points on an elliptic curve is an abelian group under the "chord and tangent rule" with $\mathcal{O}_E$ being the identity element. The *j-invariant* of an elliptic curve is $j(E) = 1728\frac{4A^3}{4A^3+27B^2}$ and there is an isomorphism of curves $f : E_0 \to E_1$ if and only if $j(E_0) = j(E_1)$.

Given two elliptic curves E_0 and E_1 over a finite field $\mathbb{F}_q$, an *isogeny* is a non-constant rational map $\phi : E_0 \to E_1$ defined over $\mathbb{F}_q$ which is also a group homomorphism from $E_0(\mathbb{F}_q)$ to $E_1(\mathbb{F}_q)$. Two curves are called *isogenous* if there exists an isogeny between them. The *degree* of an isogeny ϕ is its degree as a rational map. For separable isogenies, the degree is also equal to the number of elements in the kernel of ϕ. Note that we will always consider the separable case in the following.

Since an isogeny defines a group homomorphism $E_0 \to E_1$, its kernel is a subgroup of E_0. Conversely, any subgroup $S \subset E_0$ determines a (separable) isogeny $\phi : E_0 \to E_1$ with $\ker\phi = S$ and $E_1 = E_0/S$.

An *endomorphism* of an elliptic curve E defined over $\mathbb{F}_q$ is an isogeny defined over an extension of $\mathbb{F}_q$ mapping E onto itself. The set of endomorphisms of E together with the zero map forms a ring under pointwise addition and function composition. This ring is the *endomorphism ring* of E, denoted $\text{End}(E)$, and it is isomorphic either to an order in a quaternion algebra and E is called *supersingular*, or to an order in an imaginary quadratic field and E is referred to as an *ordinary* curve [28].

Let d be a positive integer. Throughout the paper, we say a supersingular elliptic curve E is *at distance* d from E_0 if there exists a separable isogeny ϕ with cyclic kernel of degree d from E_0 to E.

For any isogeny $\phi : E_0 \to E_1$, there exists another isogeny $\hat{\phi}$, called the *dual isogeny*, satisfying $\phi \circ \hat{\phi} = \hat{\phi} \circ \phi = [\deg(\phi)]$, where $[\cdot]$ denotes scalar multiplication. Therefore, the property of being isogenous is an equivalence relation on the set of isomorphism classes of elliptic curves defined over $\mathbb{F}_q$.

2.3 Hard homogeneous spaces and CSIDH

Recall the notion of Couveignes' *hard homogeneous spaces* (HHS) [7], a finite commutative group action for which some operations are easy to compute and others are hard.

Instances of Couveignes' hard homogeneous spaces can be constructed using elliptic curve isogenies and have been the basis of one branch of isogeny-based cryptography which uses the group action we will describe in the following.

Denote the set of all isomorphism classes over $\overline{\mathbb{F}_q}$ of isogenous curves with n points and endomorphism ring $\mathcal{O}$ by $\mathrm{Ell}_{q,n}(\mathcal{O})$, and represent the isomorphism class of a curve E in $\mathrm{Ell}_{q,n}(\mathcal{O})$ by the j-invariant $j(E)$. Any isogeny $\varphi : E \to E_{\mathfrak{b}}$ between curves having the same endomorphism ring in $\mathrm{Ell}_{q,n}(\mathcal{O})$ is determined by E and $\ker \varphi$ up to isomorphism. This kernel corresponds to an ideal $[\mathfrak{b}]$ in $\mathcal{O}$. Recall that the ideal class group of $\mathcal{O}$, $\mathrm{Cl}(\mathcal{O})$, is the quotient group of the abelian group of fractional $\mathcal{O}$-ideals under ideal multiplication and all principal fractional $\mathcal{O}$-ideals. Since principal ideals in $\mathcal{O}$ correspond to isomorphisms, ideals that are equivalent in $\mathrm{Cl}(\mathcal{O})$ induce the same isogeny up to isomorphism. Hence, we have a well-defined group action

$$\begin{aligned} \cdot : \mathrm{Cl}(\mathcal{O}) \times \mathrm{Ell}_{q,n}(\mathcal{O}) &\to \mathrm{Ell}_{q,n}(\mathcal{O}), \\ ([\mathfrak{b}], j(E)) &\mapsto j(E_{\mathfrak{b}}), \end{aligned}$$

which is free and transitive ([31], Thm. 4.5, and erratum Thm. 4.5 of [27]).

Given two elliptic curves E_0, E_1 in $\mathrm{Ell}_{q,n}(\mathcal{O})$ up to isomorphism, it is in general assumed to be hard to find an isogeny $\varphi : E_0 \to E_1$.

A similar construction can be performed with endomorphism rings of supersingular curves. This occurrence of hard homogenous spaces is used for the *Commutative SIDH* (CSIDH) protocol [3] proposed for post-quantum non-interactive key exchange. Since the endomorphism rings of such curves are orders in a quaternion algebra, they are non-commutative and hence yield a group action with less desirable properties than in the construction for ordinary curves. Therefore, Castryck et al. suggest restricting the endomorphism ring to the subring of $\mathbb{F}_p$-rational endomorphisms which is an order in an imaginary quadratic field, and as such commutative. Again, the ideal class group of this order $\mathcal{O}$ acts on $\mathrm{Ell}_p(\mathcal{O})$, the set of all isomorphism classes of supersingular isogenous curves over $\mathbb{F}_p$ with $\mathbb{F}_p$-rational endomorphism ring (isomorphic to) $\mathcal{O}$.

Given that the set $\mathrm{Ell}_p(\mathcal{O})$ is non-empty, the group action is free and transitive (see [3], Thm. 7, summarizing results from [31], [27]), and can be used to perform a Diffie-Hellman-type key exchange. Note that CSIDH is strictly speaking not an instance of a HHS as it is not possible to compute the group action efficiently for *all* group elements.

There have been multiple proposals to attack concrete parameter suggestions for CSIDH with quantum algorithms. Peikert [21] uses Kuperberg's collimation sieve algorithm to solve the hidden shift instance with quantum accessible classical memory and subexponential quantum time, a strategy independently also explored by Bonnetain-Schrottenloher [2].

2.4 SIDH

We recall the *Supersingular Isogeny Diffie-Hellman* (SIDH) protocol which was introduced by Jao and De Feo in [14] and forms the basis of *Supersingular Isogeny Key Encapsulation* (SIKE) [13] which has been submitted to NIST's post-quantum competition.

Fix some supersingular elliptic curve E_0 over a field $\mathbb{F}_{p^2}$, where p is a prime, and let N_1 and N_2 be two smooth integers coprime to p with $(N_1, N_2) = 1$. Further choose some points $P_A, Q_A, P_B, Q_B \in E_0$ such that P_A and Q_A generate the N_1-torsion of E_0, $E_0[N_1]$, and similarly, $\langle P_B, Q_B \rangle = E_0[N_2]$. Then the protocol is as follows:

1. Alice chooses a random cyclic subgroup of $E_0[N_1]$ generated by a point of the form $A = P_A + [x_A]Q_A$ and Bob chooses some random cyclic subgroup of $E_0[N_2]$ generated by $B = P_B + [x_B]Q_B$.
2. Alice then computes her secret isogeny $\varphi_A : E_0 \to E_0/\langle A \rangle$ and Bob computes his secret isogeny $\varphi_B : E_0 \to E_0/\langle B \rangle$.
3. Alice sends the curve $E_A := E_0/\langle A \rangle$ and the two points $\varphi_A(P_B), \varphi_A(Q_B)$ to Bob while Bob sends $\big(E_B := E_0/\langle B \rangle, \varphi_B(P_A), \varphi_B(Q_A)\big)$ to Alice.
4. Alice and Bob both compute the shared secret curve $E_{AB} := E_0/\langle A, B \rangle$ using the given torsion information, $E_{AB} = E_B/\langle \varphi_B(A) \rangle = E_A/\langle \varphi_A(B) \rangle$.

For SIDH, one chooses the prime p of the form $p = N_1 N_2 f - 1$ with N_1 and N_2 being powers of 2 and 3, respectively. As the above protocol is vulnerable to adaptive attacks (see e.g., [10]), SIKE applies a variant of the Fujisaki-Okamoto transformation due to Hofheinz, Hövelmanns and Kiltz [12] to standard SIDH. To ensure that both Alice and Bob enjoy the same level of security, the recommended parameter sets for SIDH and SIKE suggest balanced parameters, i.e., $N_1 \approx N_2$.

The active attack on standard SIDH presented by Galbraith et al. [10] utilizes the additional information on torsion points to recover a secret key through multiple executions of the protocol with malformed messages. Further, the given torsion point information is exploited in Petit's passive attack [22] on a non-standard variant of SIDH with unbalanced and comparatively large torsion parameters. The requirements on unbalancedness and size of parameters have recently been improved upon by Kutas et al. [19] who additionally show that, even with balanced parameters, there exist certain primes which facilitate an effective-torsion point attack on SIDH.

For our quantum attack to work, we will need to relax the balancedness condition of standard SIDH and require one of N_1 and N_2 to be larger than the other by a certain factor. In particular, we need $N_1 N_2 \gg p$ which prohibits choosing p

as suggested by Jao-De Feo. We call this variant of SIDH *overstretched*. Note that this variant of SIDH is still polynomial time as long as N_1 and N_2 are smooth numbers, albeit much slower in practice than with the suggested parameters.

SIDH is believed to be immune to subexponential quantum attacks [1, 13, 14]. In particular, it has been claimed and widely accepted that no reasonable variant of Childs et al.'s attack [5] exists for SIDH [14, p.18, Section 5]. Yet, we will show in this paper how to reduce SIDH with overstretched parameters to an abelian hidden shift problem.

2.5 Quantum algorithms to solve hidden shift problems

First, we recall what is meant when two functions are said to be shifts of each other, or equivalently that these two functions *hide a shift*.

Definition 2.1. *Let $F_0, F_1 \colon G \to X$ be two functions defined on some group G, such that there exists some $s \in G$ satisfying $F_0(g) = F_1(g \cdot s)$ for all $g \in G$. The* hidden shift problem *is to find s given oracle access to the functions F_0 and F_1.*

Multiple approaches utilizing quantum computation have been proposed to solve the hidden shift problem. Some of these works have considered different group structures as well as variations on the promise. We summarize some quantum algorithms solving the injective abelian hidden shift problem, i.e., where the functions F_i are injective functions and G is abelian.

The first quantum subexponential algorithm is due to Kuperberg [17] and reduces the hidden shift problem to the hidden subgroup problem in the dihedral group $D_G \simeq \mathcal{C}_2 \ltimes G$, i.e., to finding a subgroup of D_G such that a function obtained from combining the input functions of the hidden shift problem is constant exactly on its cosets. It requires quantum subexponential time, namely $2^{\mathcal{O}(\sqrt{\log |G|})}$ quantum queries, for a finite abelian group G. A modification of this method proposed by Regev [25] reduces the memory required by Kuperberg's approach (from super-polynomial to polynomial) while keeping the running time quantum subexponential. Another, slightly faster algorithm, the collimation sieve, using polynomial quantum space was proposed later by Kuperberg [18]. In this variant, parameter trade-offs between classical and quantum running time and quantumly accessible memory are possible.

These algorithms to solve the hidden shift problem when G is abelian generally begin by producing some random quantum states, each with an associated classical "label", by evaluating the group action on a uniform superposition over the group G. For this generation of states, oracle access to the two functions F_0 and F_1 is needed. Then, the hidden shift s is extracted bitwise through performing measurements on specific quantum states (i.e., ones with desirable labels) which are generated from the random states via some sieve algorithm.

3 Malleability oracles and hidden shift attacks

In this section, we introduce the notion of a *malleability oracle* for a one-way function. Under some conditions, such an oracle allows the computation of preimages of given elements in quantum subexponential time by reduction to the hidden shift problem.

3.1 Malleability oracles

Recall the definition of a free and transitive group action.

Definition 3.1. *Let G be a group with neutral element e, and let $\mathcal{I}$ be a set. A* (left) group action $\star$ *of G on $\mathcal{I}$ is a function*

$$\star \colon G \times \mathcal{I} \to \mathcal{I}, \quad (g, x) \mapsto g \star x,$$

that satisfies $e \star x = x$, and $gh \star x = g \star (h \star x)$ for all $x \in \mathcal{I}$ and $g, h \in G$.

The group action is called transitive *if and only if $\mathcal{I}$ is non-empty and for every pair of elements $x, y \in \mathcal{I}$ there exists $g \in G$ such that $g \star x = y$. The group action is called* free *if and only if $g \star x = x$ implies $g = e$.*

Next, we define an oracle capturing the main premise required for our strategy to compute preimages of one-way functions.

Definition 3.2. *Let $f : \mathcal{I} \to \mathfrak{O}$ be an injective (one-way) function and let $\star$ be the action of a group G on $\mathcal{I}$. A* malleability oracle *for G at $o := f(i)$ provides the value of $f(g \star i)$ for any input $g \in G$, i.e., the malleability oracle evaluates the map*

$$g \mapsto f(g \star i).$$

We call the function f malleable*, if a malleability oracle is available at every $o \in f(\mathcal{I})$.*

In Section 4 we show how a polynomial-time malleability oracle can be constructed in the context of SIDH with overstretched parameters, and in Section 5 we describe other contexts where such an oracle arises naturally.

For the remainder of the paper, we will denote the action of a group element $g \in G$ on a set element $i \in \mathcal{I}$ by $g \cdot i$.

3.2 Reduction to hidden shift problem

Given a malleability oracle at $o = f(i)$, computing a preimage of o reduces to a hidden shift problem in the following case.

Theorem 3.3. *Let $f : \mathcal{I} \to \mathfrak{O}$ be an injective (one-way) function and let G be a group acting transitively on $\mathcal{I}$. Given a malleability oracle for G at $o := f(i)$, the preimage of o can be computed by solving a hidden shift problem.*

Proof. Let k be an arbitrary but fixed element in $\mathcal{I}$ and define

$$F_k \; : \; G \to \mathfrak{O} \; , \; \theta \mapsto f(\theta \cdot k).$$

Since f is an injective function, $i = f^{-1}(o)$ is unique and thus F_i is well-defined. Moreover, the malleability oracle allows us to evaluate the function F_i on any $\theta \in G$, as $F_i(\theta) = f(\theta \cdot i)$.

Fix some arbitrary $j \in \mathcal{I}$. Since we know j, we can evaluate F_j on any group element θ by evaluating $f(\theta \cdot j)$ via simply computing the group action. Due to the transitivity of the group action of G, there exists $\sigma \in G$ such that $i = \sigma \cdot j$. Since for all $\theta \in G$

$$F_i(\theta) = f(\theta \cdot i) = f(\theta\sigma \cdot j) = F_j(\theta\sigma),$$

the functions F_j and F_i are shifts of each other. Hence, solving the hidden shift problem for F_i and F_j allows us to recover σ, and thus to compute $i = \sigma \cdot j$. □

The following corollary will be used in our attack on overstretched SIDH.

Corollary 3.4. *Let $f : \mathcal{I} \to \mathfrak{O}$ be an injective (one-way) function and let G be a finitely generated abelian group acting freely and transitively on $\mathcal{I}$. Given a malleability oracle for G at $o := f(i)$, the preimage of o can be computed in quantum subexponential time.*

Proof. To obtain a hidden shift instance solvable by subexponential quantum algorithm such as Kuperberg's, we only have to show that for every $k \in \mathcal{I}$ the function $F_k(\theta) = f(\theta \cdot k)$ is injective. Then the claim follows from Theorem 3.3 and the discussion in Section 2.5.

Suppose that $F_k(g) = f(g \cdot k) = f(h \cdot k) = F_k(h)$ for some $g, h \in G$. Since f is injective and the group action is free, this implies $g = h$. □

4 Attack on overstretched SIDH instances in quantum subexponential time

Despite the non-commutative nature of SIDH, we show in this section that one can find an abelian group action on its private key space. Moreover for sufficiently *overstretched* SIDH parameters, the torsion point information revealed in the protocol allows us to build a malleability oracle for this group action. This gives rise to an attack using quantum subexponential hidden shift algorithms as outlined in Section 3.2.

This section is organized as follows: We first sketch our approach to exploit the torsion point information in Section 4.1. We then solve some technical issues in Sections 4.2 through 4.4. These issues require small tweaks to our general approach, and we summarize the resulting algorithm in Section 4.5. Finally in Section 4.6, we present a hybrid approach to combine guessing part of the secret and computing the remaining part using our new attack; this allows us to slightly extend the attack to further parameter sets.

Throughout this section, we use the following notation. Let $p \equiv 3 \pmod 4$ be prime, let E_0 be the supersingular elliptic curve with j-invariant 1728 defined over $\mathbb{F}_p$, given by the equation $y^2 = x^3 + x$, and let $\mathcal{O}_0 = \mathrm{End}(E_0)$ be its endomorphism ring. Note that $\mathcal{O}_0$ is well-known. More precisely, it is the $\mathbb{Z}$-module generated by $1, \iota, \frac{1+\pi}{2}$ and $\frac{\iota+\iota\pi}{2}$, where ι denotes the non-trivial automorphism of E_0, $(x, y) \mapsto (-x, iy)$, and π is the Frobenius endomorphism, $(x, y) \mapsto (x^p, y^p)$.

Remark 4.1. The attack we describe can be expanded to other curves that are close to E_0, such as the curve used in the updated parameters of SIKE for the second round of NIST's post-quantum standardization effort [1], by computing the isogeny to E_0 and translating the problem to there.

4.1 Overview of the attack

Let $\mathcal{I}$ be the set of cyclic N_1-order subgroups of E_0, and let $\mathfrak{O}$ be the set of j-invariants of all supersingular curves that are N_1-isogenous to E_0. Let f be the function sending any element of $\mathcal{I}$ to the j-invariant of the codomain of its corresponding isogeny, i.e.,

$$f : \mathcal{I} \to \mathfrak{O}, \quad K \mapsto j(E_0/K). \tag{1}$$

The function f can be efficiently computed on any input using Vélu's formulae [29], provided N_1 is sufficiently smooth and that the N_1-torsion is defined over a sufficiently small extension field of $\mathbb{F}_p$. In SIDH, the latter is achieved by choosing $N_1 | p - 1$, but this is true more generally for sufficiently powersmooth N_1.

On the other hand, inverting f amounts to finding an isogeny of degree N_1 from E_0 to a curve in a given isomorphism class, or equivalently to finding the subgroup of E_0 defining this isogeny. The conjectured hardness of this problem is at the heart of isogeny-based cryptography.

In the SIDH protocol, additional torsion point information is transmitted publicly as part of the exchange, and thus also given to adversaries. For the security proof it is assumed that a variant of the following problem with $N_1 \approx N_2$ is hard [14].

Problem 4.2. *Let p be a large prime, let N_1 and N_2 be two smooth coprime integers such that $E_0[N_1]$ and $E_0[N_2]$ can be represented efficiently, let $K \in \mathcal{I}$ be a cyclic subgroup of order N_1 of E_0 chosen uniformly at random, and let $\varphi : E_0 \to E_0/K$. Given the supersingular elliptic curves E_0 and E_0/K together with the restriction of φ to $E_0[N_2]$, compute K.*

Our attack exploits the information provided by the restriction of the secret isogeny to $E_0[N_2]$ to construct a malleability oracle for f at the (unknown) secret. Following the framework outlined in Section 3, this gives rise to an attack on *overstretched* SIDH.

Let G be a subgroup of $(\mathcal{O}_0/N_1\mathcal{O}_0)^*$. Then G induces a group action on $\mathcal{I}$ given by

$$G \times \mathcal{I} \to \mathcal{I}\ , \ (\theta, K) \mapsto \theta(K).$$

Indeed, the degree of any non-trivial representative θ is coprime to N_1 and thus preserves the order of any generator of K.

Note that the full group $(\mathcal{O}_0/N_1\mathcal{O}_0)^*$ is not abelian. Our attack will require an abelian subgroup G acting on $\mathcal{I}$ such that G acts freely and transitively on the orbit of an isogeny kernel of an isogeny $E_0 \to E_0/K$ under this group action, as well as one element in this orbit. This leads to the following task.

Task 4.3. *Let $K \in \mathcal{I}$ be any cyclic subgroup of E_0 of order N_1 chosen uniformly at random and let $\varphi : E_0 \to E_A := E_0/K$. Compute an element $L \in \mathcal{I}$ and an abelian subgroup G of $(\mathcal{O}_0/N_1\mathcal{O}_0)^*$ such that G acts freely and transitively on the orbit $G \cdot L$, f is injective on $G \cdot L$ and $j(E_A)$ is contained in $f(G \cdot L) \subset \mathfrak{O}$.*

We solve this task in Section 4.2. More precisely, we find three subsets of $\mathcal{I}$ restricted to which f is injective, and we give abelian groups that induce the required action on these subsets. Furthermore, the image of f restricted to one of these three subsets of $\mathcal{I}$ will always contain $j(E_0/K)$.

In order to apply our general framework from Section 3, it remains to construct a malleability oracle for f at $j(E_0/K)$ for any secret $K \in \mathcal{I}$. To construct this oracle, we use both the torsion point information provided in the SIDH protocol and a solution to the following task.

Task 4.4. *Given an endomorphism $\theta \in G$ of degree coprime to N_1 and an integer N_2 coprime to N_1, compute an endomorphism θ' of degree N_2 such that θ and θ' induce the same action on the set $\mathcal{I}$ of cyclic subgroups of $E_0[N_1]$ of order N_1.*

In Appendix C, we give a direct solution to a variation of this task when using sufficiently overstretched and unbalanced parameters, i.e. $N_2 > p^2N_1^4$. However, in Section 4.3 we show that it suffices to lift elements of πG, where π is the Frobenius map. A solution to Task 4.4 for these elements requiring only $N_2 > pN_1^4$ is described in Section 4.4.

The following lemma results from the coprimality of $\deg(\theta)$ and N_1 and is depicted in Figure 1.

$$\begin{tikzcd}
E_0 \arrow[r, "\varphi"] \arrow[d, "\theta"'] & E_A \arrow[d] \\
E_0 \arrow[r] & E_0/\theta(\ker\varphi) \cong E_A/\varphi(\ker\theta)
\end{tikzcd}$$

Fig. 1. The isogenies φ and the endomorphism θ are of coprime degrees.

Lemma 4.5. *Let $\varphi : E_0 \to E_A$ be an isogeny of degree N_1 and let $\theta \in End(E_0)$ be of degree coprime to N_1. Then $E_A/\varphi(\ker\theta)$ is isomorphic to $E_0/\theta(\ker\varphi)$.*

Let N_3 be the degree of θ. We cannot compute the curve $E_0/\theta(\ker\varphi)$ in general without the knowledge of the isogeny φ or its action on the N_3-torsion. However, we can compute the curve if we find an endomorphism θ' of degree N_3' such that θ and θ' have the same action on the N_1-torsion and $\varphi|_{E_0[N_3']}$ is known. This is the motivation behind Task 4.4, as we know the action of φ on the N_2-torsion in Problem 4.2. A solution to this task yields a malleability oracle for f with respect to the previously described group action of G on $\mathcal{I}$ in the SIDH setting.

We outline the construction of the malleability oracle in Algorithm 1. Correctness will follow from the proof of Proposition 4.26 given a suitable choice of the acting group G which we will discuss in Subsection 4.2.

```
Algorithm 1: Computation of $f(\theta(K))$, given $f(K)$ and $\theta \in G$
  Let $\varphi : E_0 \to E_A := E_0/K$ be an isogeny of degree $N_1$, let $N_2$ be coprime to
    $N_1$ and $G \subset (\mathcal{O}_0/N_1\mathcal{O}_0)^*$ one of the abelian groups as in Task 4.3 that acts
    freely and transitively on $K$.
  Input: $E_0$, $f(K) = j(E_A)$, $\varphi|_{E_0[N_2]}$ and $\theta \in G$.
  Output: $f(\theta(K)) = j(E_0/\theta(K))$.
1 Compute endomorphism $\theta'$ of degree $N_2$ having the same action as $\theta$ on cyclic
    $N_1$-order subgroups of $E_0[N_1]$ as provided by a solution to Task 4.4;
2 Determine $\varphi(\ker\theta')$, using the knowledge of $\varphi$ on $E_0[N_2]$;
3 Compute $j(E_A/\varphi(\ker\theta')) = j(E_0/\theta(K))$;
4 return $f(\theta(K)) = j(E_0/\theta(K))$
```

For parameters that allow us to construct a malleability oracle, we can then solve Problem 4.2 underlying SIDH-like protocols via a reduction to an injective abelian hidden shift problem using the framework introduced in Section 3.2.

Informal result 4.6. *Suppose the parameters allow the efficient solution of Task 4.4, then Problem 4.2 can be solved in quantum subexponential time.*

We use the remainder of this section to prove this result formally under certain assumptions. To this end, we first give solutions to Task 4.3 and, for some parameters, to a variant of Task 4.4. More precisely, we show in Section 4.3 that it is sufficient to lift elements from πG instead of G. For this case, we then give a more efficient lifting procedure requiring unbalanced and overstretched parameters. We construct a malleability oracle using the torsion point information provided in SIDH and a subroutine solving our variant of Task 4.4. Apart from some technical details that we will address in the following, the informal result follows from Corollary 3.4. An overview of the attack is depicted in Algorithm 2.

4.2 A free and transitive group action

Recall that E_0 is the supersingular curve with j-invariant 1728, given by the equation $y^2 = x^3 + x$. In this section we provide a solution to Task 4.3. For

Algorithm 2: Solving SIDH's underlying hardness assumption via an abelian hidden shift problem

```
Let $\varphi : E_0 \to E_0/K$ be an $N_1$-isogeny and $N_2 \in \mathbb{Z}$ such that $\gcd(N_1, N_2) = 1$.
Input: $E_0$, $E_0/K$, $\varphi(E_0[N_2])$.
Output: Isogeny $E_0 \to E$, where $j(E) = j(E_0/K)$.
Compute an abelian group $G \subset (\mathcal{O}_0/N_1\mathcal{O}_0)^*$ acting freely and transitively on
  the orbit $G(K)$ and some $J \in G(K) \subset \mathcal{I}$;
Define $F_K : G \to \mathfrak{O}, g \mapsto f(g(K))$ and $F_J : G \to \mathfrak{O}, g \mapsto f(g(J))$;
Compute injective abelian hidden shift $\theta \in G$ of $F_K$ and $F_J$, i.e., $\theta \in G$ such
  that $F_K(g) = F_J(\theta g)$ for all $g \in G$, using a quantum algorithm such as
  Kuperberg's. To this end, one evaluates $F_K$ using Algorithm 1 and $F_J$ using
  the knowledge of $J$;
return Isogeny $E_0 \to E_0/\theta(J)$
```

simplicity, we treat N_1 as a power of 2, but the results generalize to any power of a small prime. A generalization to powers of 3 is sketched in Appendix B.

We provide the solution by identifying three subsets of $\mathcal{I}$ that are orbits under a free and transitive action of abelian subgroups of $(\mathcal{O}_0/N_1\mathcal{O}_0)^*$. More precisely, let $P \in E_0$ such that $\langle P, \iota(P)\rangle = E_0[N_1]$, where ι denotes the automorphism $(x, y) \mapsto (-x, iy)$ of E_0. Let $Q := P + \iota(P)$ and define the following three subsets of $\mathcal{I}$.

$$\mathcal{I}_1 := \{\langle P + [\alpha]\iota(P)\rangle \mid \alpha \text{ even } \}$$
$$\mathcal{I}_2 := \left\{\langle Q + \alpha\iota(Q)\rangle \mid \alpha \text{ even and } \alpha \in \left[0, \frac{N_1}{2} - 1\right]\right\}$$
$$\mathcal{I}_3 := \left\{\langle Q + \alpha\iota(Q)\rangle \mid \alpha \text{ even and } \alpha \in \left[\frac{N_1}{2}, N_1 - 1\right]\right\}$$

Recall the function f defined in (1), mapping cyclic subgroups of $E_0[N_1]$ of order N_1 to j-invariants of curves at distance N_1 from E_0,

$$f : \mathcal{I} \to \mathfrak{O}, \quad K \mapsto j(E_0/K).$$

We will show that restricting the function f to any of the subsets $\mathcal{I}_1$, $\mathcal{I}_2$, or $\mathcal{I}_3$ yields an injective function and we will prove that $f(\cup_i \mathcal{I}_i) = f(\mathcal{I})$. Furthermore, we will see that

$$G_0 := \{a + b\iota \mid a \text{ odd}, b \text{ even } \} / N_1\mathcal{O}_0^*$$

acts transitively on $\mathcal{I}_1$. In order to ensure that the action is free, we identify two endomorphisms $a + b\iota$ and $a' + b'\iota$ in G_0 if there exists an odd $\lambda \in \mathbb{Z}/N_1\mathbb{Z}$ such that $a \equiv \lambda a' \pmod{N_1}$ and $b \equiv \lambda b' \pmod{N_1}$. We denote the resulting group by G.

In order to define free and transitive group actions on $\mathcal{I}_2$, and $\mathcal{I}_3$ we define similarly to G_0

$$H_0 := \{a + b\iota \mid a \text{ odd}, b \text{ even } \} / (N_1/2)\mathcal{O}_0^*.$$

Again, we identify two endomorphisms $a + b\iota$ and $a' + b'\iota$ in H_0 if there exists an odd $\lambda \in \mathbb{Z}/(N_1/2)\mathbb{Z}$ such that $a \equiv \lambda a' \pmod{N_1/2}$ and $b \equiv \lambda b' \pmod{N_1/2}$, we obtain a group H. The group H will act freely and transitively on $\mathcal{I}_2$ and $\mathcal{I}_3$.

Hence, one of these three options will always be a solution to Task 4.3.

The map f is based on the well-known correspondence between $\mathcal{I}$ and curves at distance N_1 from E_0. However, this correspondence is not necessarily one-to-one. In particular, if E_0 has a non-scalar endomorphism of degree N_1^2, then that endomorphism can be decomposed as $\hat{\tau}_1 \circ \tau_2$, where τ_1 and τ_2 are non-isomorphic isogenies of degree N_1 from E_0 to the same curve E. For small enough N_1, the following lemma shows that two kernels correspond to the same curve if and only if they are linked by the automorphism ι.

Lemma 4.7. *Suppose that $N_1^2 < \frac{p+1}{4}$. Then the only endomorphisms of degree N_1^2 of E_0 are $[N_1]$ and $[N_1] \cdot \iota$, where $\iota : E_0 \to E_0, (x, y) \mapsto (-x, iy)$ is the non-trivial automorphism.*

Proof. Due to the condition $N_1^2 < \frac{p+1}{4}$, an endomorphism θ of degree N_1^2 lies in $\mathbb{Z}[\iota]$. Let $\theta = a + b\iota$ for some $a, b \in \mathbb{Z}$. Then the degree of θ is $a^2 + b^2$. Now we have to prove that the only ways to decompose N_1^2 as a sum of two squares are trivial, i.e., $N_1^2 = N_1^2 + 0^2 = 0^2 + N_1^2$.

Let $N_1 = 2^k$, and we prove the statement by induction on k. For $k = 1$ the statement is trivial. Suppose that $k > 1$ and that $N_1^2 = a^2 + b^2$. Then a and b cannot both be odd as N_1^2 is divisible by four. If they were both even, then dividing by four yields a decomposition of $(N_1/2)^2 = (a/2)^2 + (b/2)^2$. By the induction hypothesis, this decomposition is trivial implying that N_1^2 can also only be decomposed in a trivial way. □

Corollary 4.8. *Suppose that $N_1^2 < \frac{p+1}{4}$. Let ϕ and ϕ' be two isogenies of degree N_1 from E_0 to a curve E. Then either $\ker \phi = \ker \phi'$ or $\ker \phi = \iota(\ker \phi')$.*

Proof. Consider the endomorphism $\tau = \hat{\phi}' \circ \phi$ of E_0. The degree of τ is N_1^2, so $\tau = [N_1]$ or $\tau = [N_1] \cdot \iota$ by Lemma 4.7. In the former case, the isogenies ϕ and ϕ' are identical by the uniqueness of the dual. In the latter case, we have $\ker \phi = \iota(\ker \phi')$. □

Thus, an element in the image of f has precisely one preimage if the kernel of the corresponding isogeny is fixed by the automorphism ι.

Identifying an abelian group with $\mathcal{I}_1$: Now, we will give the free and transitive group action on $\mathcal{I}_1$ and show that f restricted to $\mathcal{I}_1$ is injective.

Let P be a point such that $\{P, \iota(P)\}$ is a basis of $E_0[N_1]$ and recall

$$\mathcal{I}_1 := \{\langle P + [\alpha]\iota(P)\rangle \mid \alpha \text{ even }\}.$$

We show that the restriction of f to $\mathcal{I}_1$ is injective.

Proposition 4.9. *Let $j(E_0) = 1728$ and suppose that $N_1^2 < \frac{p+1}{4}$. The restriction of f to $\mathcal{I}_1$ is injective.*

Proof. We apply Corollary 4.8 to show that the codomains of isogenies with kernel in $\mathcal{I}_1$ are pairwise non-isomorphic curves. It is clear that $P+\alpha\iota(P)$ and $P+\alpha'\iota(P)$ are not scalar multiples of each other if $\alpha \neq \alpha'$ as $P, \iota(P)$ generate $E_0[N_1]$. It remains to show that for any even α, α', the points $P+\alpha\iota(P)$ and $-\alpha' P+\iota(P)$ are not scalar multiples of each other. Suppose there exists an odd λ such that

$$P+\alpha\iota(P)=\lambda(-\alpha' P+\iota(P)).$$

Note that we can restrict to odd λs as the order of both points is N_1. Since $\{P,\iota(P)\}$ is a basis of the N_1-torsion, this implies that $1 \equiv -\lambda\alpha' \pmod{N_1}$. Since α' is even this is a contradiction concluding the proof. □

Clearly, $f(\mathcal{I}_1)$ does not include all elliptic curves at distance N_1 from E_0, i.e., all curves in $f(\mathcal{I})$. Every curve at distance N_1 from E_0 is of the form $E_0/\langle P+\alpha\iota(P)\rangle$ for some $\alpha \in \mathbb{Z}/N_1\mathbb{Z}$, which follows from the observation that the curves $E_0/\langle \beta_1 P+\beta_2\iota(P)\rangle$ and $E_0/\langle -\beta_2 P+\beta_1\iota(P)\rangle$ are isomorphic since their kernels are linked by ι. We first restrict ourselves to define a free and transitive group action on $\mathcal{I}_1$ and define the free and transitive group action on the kernels corresponding to the remaining curves later.

Recall that E_0 is a curve with well-known endomorphism ring, and we are interested in the endomorphisms that are of degree coprime to N_1. While there are infinitely many such endomorphisms, we are only concerned with their action on $E_0[N_1]$, i.e., we are looking at the group $(\mathcal{O}_0/N_1\mathcal{O}_0)^*$ which is isomorphic to $GL_2(\mathbb{Z}/N_1\mathbb{Z})$ [30, p. 676]. Furthermore, we are only concerned with the action of the endomorphisms on $\mathcal{I}$, i.e., on cyclic subgroups of $E_0[N_1]$ of order N_1, and we can therefore identify even more endomorphisms with each other by the following lemma.

Lemma 4.10. *Let (a,b,c,d) and (a',b',c',d') be the coefficients of θ and θ' with respect to some $\mathbb{Z}$-basis of the endomorphism ring $\mathcal{O}_0$ of E_0, and let $\mathcal{I}$ be the set of cyclic N_1-order subgroups of $E_0[N_1]$. Then $\theta(K)=\theta'(K)$ for every $K \in \mathcal{I}$ if and only if there exists some $\lambda \in (\mathbb{Z}/N_1\mathbb{Z})^*$ such that*

$$(a,b,c,d) \equiv \lambda(a',b',c',d') \pmod{N_1}.$$

Proof. Considering the respective restrictions to $E_0[N_1]$, two endomorphisms are equal if they lie in the same class in $(\mathcal{O}_0/N_1\mathcal{O}_0)^*$. Moreover, let θ_1, θ_2 be two endomorphisms such that $\theta_1=[\lambda]\theta_2$ for some integer λ, and let P be an element of order N_1. Since scalar multiplication commutes with any endomorphism, it is easy to see that $\theta_1(P)$ and $\theta_2(P)$ generate the same subgroup in $E_0[N_1]$ if and only if λ is coprime to N_1. □

Now, we are ready to give a solution to Task 4.3 if $K \in \mathcal{I}_1$.

Proposition 4.11. *Let G be the group of equivalence classes of elements*

$$\{a+b\iota \mid a \text{ odd, } b \text{ even }\} \subset \mathbb{Z}[\iota]/N_1\mathcal{O}_0^* \subset (\mathcal{O}_0/N_1\mathcal{O}_0)^*,$$

where we identify two elements if and only if they differ by multiplication by an odd scalar modulo N_1. Then G is an abelian group, and it acts freely and transitively on $\mathcal{I}_1$.

Proof. It is easy to see that the endomorphisms in $\mathbb{Z}[\iota]$ of degree coprime to N_1 form an abelian subgroup of $\mathcal{O}_0$. Using any basis for $E_0[N_1]$ of the form $\{P, \iota(P)\}$, we can write the elements of this subgroup as matrices of the form $\begin{pmatrix} a & b \\ -b & a \end{pmatrix}$, where a is odd and b is even. By identifying two endomorphisms $a_1+b_1\iota$ and $a_2+b_2\iota$ if there exists an integer λ coprime to N_1 and an endomorphism δ such that $a_1 - \lambda a_2 + (b_1 - \lambda b_2) = N_1\delta$, which is possible by Lemma 4.10, we obtain G. As G is closed under multiplication and reduction modulo N_1, it is a subgroup of an abelian group and therefore abelian itself. Note that G contains all equivalence classes under Lemma 4.10 of endomorphisms of the form $a + b\iota$ for even b, independently of the chosen basis.

To examine the orbit of an element in $\mathcal{I}$, which is a cyclic N_1-order subgroup of $E_0[N_1]$, under the action of G, it is sufficient to look at the orbit of a generator of this cyclic group in $\mathcal{I}$. We consider the orbit of P which has coordinates $(1, 0)$ with respect to our basis under the group action of G. The image of $(1, 0)$ under an element $\begin{pmatrix} 1 & b \\ -b & 1 \end{pmatrix}$ is $(1, b)$. Inspecting the cyclic subgroups of E_0 these points generate, we get $G \cdot \langle P \rangle = \mathcal{I}_1$. □

Free and transitive group action on $\mathcal{I}_2$ and $\mathcal{I}_3$: So far we have defined a free and transitive group action on $\mathcal{I}_1$ and thus for the curves in $f(\mathcal{I}_1)$. However, when the secret kernel is generated by $P+\alpha\iota(P)$ with α odd, the curve $E_0/\langle P+\alpha\iota(P)\rangle$ is not contained in $f(\mathcal{I}_1)$. This is the case we handle next.

One can show that the action of the previously defined group G acting on curves at distance N_1 from E_0 considered via f has three orbits (see Appendix A for details). We have already seen that $f(\mathcal{I}_1)$ is one orbit, but the odd-α cases will split into two orbits. Clearly, G cannot be free and transitive on both orbits, since the size of the orbits is smaller than the cardinality of the group. We avoid this issue by choosing a different (but related) group of cardinality $N_1/4$ acting on the curves corresponding to an odd α.

Lemma 4.12. *Let $Q := P + \iota(P)$ and define*

$$\mathcal{I}_2 := \left\{ \langle Q + \alpha\iota(Q) \rangle \mid \alpha \textit{ even and } \alpha \in \left[0, \frac{N_1}{2} - 1\right] \right\}$$
$$\mathcal{I}_3 := \left\{ \langle Q + \alpha\iota(Q) \rangle \mid \alpha \textit{ even and } \alpha \in \left[\frac{N_1}{2}, N_1 - 1\right] \right\}.$$

The restrictions $f_{|\mathcal{I}_2}$ and $f_{|\mathcal{I}_3}$ of f to $\mathcal{I}_2$ and $\mathcal{I}_3$ are injective.

Proof. We show that two distinct isogenies with kernel both in $\mathcal{I}_2$ (or both in $\mathcal{I}_3$) map to two non-isomorphic curves. Let α, α' be such that $\langle Q + \alpha\iota(Q) \rangle$ and

$\langle Q + \alpha'\iota(Q)\rangle$ are both in $\mathcal{I}_2$, or $\mathcal{I}_3$, respectively. Suppose there exists an odd λ such that

$$Q + \alpha\iota(Q) = \lambda(Q + \alpha'\iota(Q)).$$

This means $1 - \lambda \equiv 0 \pmod{N_1/2}$ and $\alpha - \lambda\alpha' \equiv 0 \pmod{N_1/2}$ which implies $\alpha \equiv \alpha' \pmod{N_1/2}$. We are left to show that $Q+\alpha\iota(Q)$ is never an odd multiple of $-\alpha Q + \iota(Q)$. Suppose there exists an odd λ such that

$$Q + \alpha\iota(Q) = \lambda(-\alpha' Q + \iota(Q)).$$

This implies $1 + \alpha'\lambda \equiv \alpha - \lambda \equiv 0 \pmod{N_1/2}$, which is a contradiction, since $\alpha - \lambda \equiv 0 \pmod{N_1/2}$ implies that λ is even while $1 + \alpha'\lambda \equiv 0 \pmod{N_1/2}$ implies that λ is odd. Therefore, the curves $E_0/\langle Q+\alpha\iota(Q)\rangle$ and $E_0/\langle Q+\alpha'\iota(Q)\rangle$ are pairwise non-isomorphic. □

Finally, we give a free and transitive group action on $\mathcal{I}_2$ and $\mathcal{I}_3$. We start by defining the acting group.

We identify two endomorphisms $a + b\iota$ and $a' + b'\iota$ if there exists an odd $\lambda \in \mathbb{Z}/(N_1/2)\mathbb{Z}$ such that $a \equiv \lambda a' \pmod{N_1/2}$ and $b \equiv \lambda b' \pmod{N_1/2}$ and we call the resulting group H_0. Let H be the subgroup of H_0 containing elements with even b.

Proposition 4.13. *H acts freely and transitively on $\mathcal{I}_2$ and $\mathcal{I}_3$.*

Proof. It is enough to show that H acts transitively on $\mathcal{I}_2$ and $\mathcal{I}_3$ because $H, \mathcal{I}_2$ and $\mathcal{I}_3$ have the same cardinality. We show that the orbit $H \cdot \langle Q\rangle$ contains every element in $\mathcal{I}_2$. This follows immediately from $(1 + \alpha\iota)Q = Q + \alpha\iota(Q)$. Similarly, H acts transitively on $\mathcal{I}_3$ as

$$(1+\alpha\iota)(Q+N_1\iota(Q)/2) = (1-\alpha N_1/2)Q+(\alpha+N_1/2)\iota(Q) = Q+(\alpha+N_1/2)\iota(Q),$$

where $(\alpha N_1/2)Q = 0$ as α is even. □

What remains to be shown is that every curve $E_0/\langle P + \alpha\iota(P)\rangle$ with odd α has a j-invariant contained in $f(\mathcal{I}_2)$ or $f(\mathcal{I}_3)$.

Proposition 4.14. *Let α be an odd integer. Then $f(\langle P + \alpha\iota(P)\rangle)$ is contained in $f(\mathcal{I}_2)$ or $f(\mathcal{I}_3)$.*

Proof. Observe that

$$P + \alpha\iota(P) = \frac{1+\alpha}{2}(P + \iota(P)) + \frac{\alpha-1}{2}(-P + \iota(P)) = \frac{1+\alpha}{2}Q + \frac{\alpha-1}{2}\iota(Q).$$

The sum of $\frac{1+\alpha}{2}$ and $\frac{\alpha-1}{2}$ is odd and therefore one of the fractions is even while the other one is odd. If $\frac{\alpha-1}{2}$ is even, then it is clear that the curve is contained in $f(\mathcal{I}_2)$ or $f(\mathcal{I}_3)$. In the case where $\frac{1+\alpha}{2}$ is even, $E_0/\langle \frac{1+\alpha}{2}Q + \frac{\alpha-1}{2}\iota(Q)\rangle$ is isomorphic to $E_0/\langle \frac{1-\alpha}{2}Q + \frac{\alpha+1}{2}\iota(Q)\rangle$ (because their kernels are related by ι) and thus the curve is contained in $f(\mathcal{I}_2)$ or $f(\mathcal{I}_3)$. □

In this subsection, we have identified three subsets of $\mathcal{I}$, restricted to which f is injective. Moreover, we have seen that the union $\cup_{i=1}^{3} f(\mathcal{I}_i)$ contains the j-invariants of all curves at distance N_1 from E_0. Finally, we gave an abelian subgroup of $(\mathcal{O}_0/N_1\mathcal{O}_0)^*$ for each of these subsets of $\mathcal{I}$ that acts freely and transitively on it. Thus, we solve Task 4.3 as long as one determines or guesses which of the three $f(\mathcal{I}_i)$ contains $j(E_0/K)$.

4.3 Using the Frobenius map

In the previous subsection, we described how to choose suitable abelian subgroups of $(\mathcal{O}_0/N_1\mathcal{O}_0)^*$ in order to solve Task 4.3 after guessing whether $j(E_0/K)$ is a j-invariant in $f(\mathcal{I}_1)$, $f(\mathcal{I}_2)$, or $f(\mathcal{I}_3)$.

The elements of the acting groups chosen as described in the previous section can be trivially lifted to $\mathbb{Z}[\iota] := \mathbb{Q}[\iota] \cap \mathcal{O}_0$. In Appendix C we show how these representatives can be lifted directly to elements of norm N_2 or eN_2, where e is a small positive integer, whenever the SIDH parameters N_1 and N_2 are sufficiently overstretched and unbalanced with $N_2 > p^2 N_1^4$. For these parameters, this solves a variation of Task 4.4.

In this section we reduce the required unbalancedness partially by proving that we can lift elements from $\pi\mathbb{Z}[\iota]$ instead. Assuming that $N_2 > pN_1^4$, we will show in Subsection 4.4 how an endomorphisms from $\pi\mathbb{Z}[\iota]$ can be lifted efficiently to another endomorphism of norm N_2 or eN_2, for some small integer e, inducing the same action on $\mathcal{I}$. Note that it is not possible to choose a group generated by an element in $\pi\mathbb{Z}[\iota]$ to solve Task 4.3 directly, acting freely and transitively on a large number of N_1-isogeny kernels, as such an element has multiplicative order at most 4.

As before, let $\varphi : E_0 \to E_0/K$ denote the secret N_1-isogeny we want to compute. Recall that to run our attack we need to be able to compute $E_0/\theta(K)$ for every θ in the groups G acting on $\mathcal{I}_1$, and H acting on $\mathcal{I}_2$ and $\mathcal{I}_2$. We have seen that we can represent θ as an element in $\mathbb{Z}[\iota]$.

Let π denote the Frobenius map. Assuming that we can lift $\pi\theta$ to an endomorphism of degree N_2 inducing the same action on $\mathcal{I}$, we can compute $E_0/\pi\theta(K)$ using knowledge of $\varphi(E_0[N_2])$ as described in Section 4.1. Now let $B := \theta(K)$. Given $E_0/\pi(B)$, we can compute E_0/B using the Frobenius map as follows.

Lemma 4.15. *Let E be an elliptic curve defined over $\mathbb{F}_p$, π the Frobenius map and let $B \subset E$ be a cyclic subgroup. $E/\pi(B)$ is isomorphic to the image of the Frobenius map of E/B.*

Proof. Let ϕ_1 be the isogeny with kernel B and ϕ_2 the isogeny with kernel $\pi(B)$. The isogeny ϕ_1 is separable and its kernel is contained in the kernel of $\phi_2 \circ \pi$. Then, there exists a unique isogeny $\psi : E/B \to E/\pi(B)$ satisfying $\phi_2 \circ \pi = \psi \circ \phi_1$ (see [28, Corollary III. 4.11.]), i.e., the following diagram commutes.

$$\begin{tikzcd}
E \arrow[r, "\phi_1"] \arrow[d, "\pi"'] & E/B \arrow[d, "\psi"] \\
E \arrow[r, "\phi_2"'] & E/\pi(B)
\end{tikzcd}$$

The degree of a composition of isogenies is the product of its factors which implies $\deg(\psi) = p$. Furthermore, ψ is not separable as the Frobenius map is not. As ψ can be decomposed as a composition of the Frobenius map and a separable isogeny (see [28, Corollary II.2.12.]), $\deg(\psi) = p$ implies that ψ must be a composition of Frobenius and an automorphism. Hence, E_0/B and $E_0/\pi(B)$ are linked by the Frobenius map. □

Lemma 4.15 implies that we can compute $E_0/\theta(K)$ by first computing $E_0/\pi\theta(K)$ and then applying the Frobenius map. This gives rise to the following strategy when constructing the malleability oracle.

Assume we want to compute $E_0/\theta(K)$ for some $\theta \in \mathbb{Z}[\iota]$ and unknown K, given the image of the N_2-torsion of the isogeny $\varphi : E_0 \to E_0/K$. Using the lifting algorithm of Subsection 4.4, we compute an endomorphism θ' of degree N_2 or eN_2 for a small e that induces the same action on $\mathcal{I}$ as $\pi\theta$. As described previously, the torsion point information allows us to compute $E_0/\theta'(K) = E_0/\pi\theta(K)$. By Lemma 4.15, applying the Frobenius map yields $E_0/\pi\theta'(K) = E_0/\theta(K)$.

4.4 Lifting $\boldsymbol{\theta \in \pi\mathbb{Z}[\iota]}$ to an element of norm $\boldsymbol{eN_2}$

In this subsection we give an efficient algorithm to lift endomorphisms from $\pi\mathbb{Z}[\iota] = \pi(\mathbb{Q}[\iota] \cap \mathrm{End}(E_0))$ to another endomorphism of $E_0/\mathbb{F}_p$ of degree N_2 or eN_2 that induces the same action on $\mathcal{I}$, whenever $N_2 > pN_1^4$. Here, e is the smallest positive integer such that $eN_2/p(c_0^2 + d_0^2)$ is a quadratic residue modulo $2N_1$, where $\pi(c_0 + d_0\iota) \in \pi\mathbb{Z}[\iota]$ is the endomorphism we want to lift.

This will solve the following task, which is a variant of Task 4.4, efficiently.

Task 4.16. *Let N_1, N_2 be coprime integers such that $N_2 > pN_1^4$, let $\theta := \pi(c_0 + d_0\iota) \in \pi\mathbb{Z}[\iota]$ be an E_0-endomorphism of degree coprime to N_1 and let e denote the smallest positive integer such that $eN_2/p(c_0^2 + d_0^2) \pmod{2N_1}$ is a quadratic residue. Compute an endomorphism θ' of degree N_2 or eN_2 such that $\theta(K) = \theta'(K)$ for all $K \in \mathcal{I}$.*

We have discussed in Section 4.3 that we can lift $\pi(c_0+d_0\iota)$ instead of $c_0+d_0\iota$. Therefore, this task solves Task 4.4 up to the following two relaxations. First, we require N_2 to be sufficiently large and unbalanced compared to N_1. Second, we allow θ' to be either of degree N_2 or eN_2 for some small positive integer e.

We have implemented the lifting algorithm of this section in MAGMA and made it publicly available[4].

[4] https://github.com/SimonMerz/lifting-for-malleability-oracles.

Remark 4.17. If N_1 were a prime, e could be chosen as the smallest quadratic non-residue modulo N_1. However, in our case N_1 is a composite number. Thus, the product of two quadratic non-residues might not be a quadratic residue if there are multiple cosets of the subgroup of quadratic residues in the group of units modulo $2N_1$.

We are primarily interested in the case where N_1 is a prime power ℓ^n. By Hensel's lemma, being a quadratic residue modulo ℓ^n is equivalent to being a quadratic residue modulo ℓ, if ℓ is odd, and equivalent to being a quadratic residue modulo 8, if $\ell = 2$.

Consequently, there is one coset of the quadratic residues in the group of units of $2N_1$ if ℓ is an odd prime. Therefore, e can be chosen to be the smallest quadratic non-residue modulo ℓ. For example, if N_1 is a power of 3 one can choose $e = 2$.

If $\ell = 2$, then there are three cosets of the quadratic residues in the group of units, i.e., the ones that contain $3, 5$, and 7 respectively. Consequently, e can always be chosen to be one of $3, 5$, or 7 in this case.

In case N_1 has distinct prime factors, for $eN_2/p(c_0^2 + d_0^2)$ to be a quadratic residue it has to be a quadratic residue modulo the largest prime power dividing $2N_1$ for each distinct prime factor. If the number of cosets grows, so do the possibilities for e und thus the size of the smallest e that is guaranteed to work.

We now describe an algorithm to solve Task 4.16. By Lemma 4.10, it suffices to solve the following task, which is similar to the problem solved at the core of the KLPT algorithm [16].

Task 4.18. *Given $\theta = a_0 + b_0\iota + (c_0 + d_0\iota)\pi$, find $\theta' = a_1 + b_1\iota + (c_1 + d_1\iota)\pi$ of degree N_2 or eN_2 with coefficients $(a_1, b_1, c_1, d_1) \equiv \lambda(a_0, b_0, c_0, d_0) \pmod{N_1}$ for some scalar $\lambda \in (\mathbb{Z}/N_1\mathbb{Z})^*$.*

In the following, we provide a solution to this task. Let

$$\theta' = \lambda a_0 + N_1 a_1 + \iota(\lambda b_0 + N_1 b_1) + (\lambda c_0 + N_1 c_1 + \iota(\lambda d_0 + N_1 d_1))\pi.$$

As $\mathrm{Norm}(x + y\iota) = x^2 + y^2$, its norm equals

$$\mathrm{Norm}(\theta') = (\lambda a_0 + N_1 a_1)^2 + (\lambda b_0 + N_1 b_1)^2 + p\big((\lambda c_0 + N_1 c_1)^2 + (\lambda d_0 + N_1 d_1)^2\big). \quad (2)$$

Since $\theta \in \pi\mathbb{Z}[\iota]$ implies $a_0 = b_0 = 0$, Equation (2) simplifies to

$$\mathrm{Norm}(\theta') = N_1^2(a_1^2 + b_1^2) + p\big((\lambda c_0 + N_1 c_1)^2 + (\lambda d_0 + N_1 d_1)^2\big). \quad (3)$$

Set e to be the smallest positive integer such that $eN_2/(p(c_0^2 + d_0^2))$ is a quadratic residue modulo $2N_1$.

The goal is to compute θ' such that $\mathrm{Norm}(\theta') = eN_2$. Considering Equation (3) modulo N_1, we obtain

$$eN_2 \equiv \lambda^2 p(c_0^2 + d_0^2) \pmod{N_1}. \quad (4)$$

Since $eN_2/p(c_0^2 + d_0^2)$ is a quadratic residue modulo $2N_1$ by the choice of e, there exists a solution for λ in Equation (4) modulo $2N_1$. Compute any such solution, and lift it to the integers in $[1, 2N_1 - 1]$. Note that we do not lose generality by the lift as any other lift of λ corresponds to a change in c_1, d_1 instead.

For fixed c_0, d_0 and λ, this gives an affine relation between c_1 and d_1 modulo N_1, i.e.,

$$c_0 c_1 + d_0 d_1 \equiv \frac{\mathrm{Norm}(\theta') - \lambda^2 p(c_0^2 + d_0^2)}{2\lambda p N_1} \pmod{N_1}. \tag{5}$$

Finally, one is left with the problem of representing an integer r as the sum of two squares, namely to find a solution (a_1, b_1) for

$$a_1^2 + b_1^2 = r := \frac{\mathrm{Norm}(\theta') - p\big((\lambda c_0 + N_1 c_1)^2 + (\lambda d_0 + N_1 d_1)^2\big)}{N_1^2} \tag{6}$$

where λ, c_0 and d_0 are fixed, and c_1, d_1 satisfy an affine equation modulo N_1.

As Petit and Smith pointed out at Mathcrypt 2018, the solution space to Equation (5) is a translated lattice modulo N_1 [23]. More precisely, we know that c_0 or d_0 is coprime to N_1. Without loss of generality, let d_0 be coprime to N_1. Furthermore, let C denote the right hand side of Equation (5). Then, (c_1, d_1) lies in the lattice

$$\langle (c_0/d_0, -1), (N_1, 0) \rangle + (C/d_0, 0). \tag{7}$$

Clearly, r from Equation (6) can only be represented as a sum of two squares, if it is positive. This happens when the parameters N_1 and N_2 are sufficiently overstretched and unbalanced. To find a solution, one computes close vectors (c_1, d_1) to the target vector $(-\lambda c_0/N_1, -\lambda d_0/N_1)$ in the translated lattice.

Given the factorisation of r as defined in Equation (6), Cornacchia's algorithm [6] can then efficiently solve for a_1, b_1 or determine that no such solution exists. If no solution exists, a different vector (c_1, d_1) is chosen.

Remark 4.19. Cornacchia's algorithm requires the factorization of r. This can be done in classical subexponential time or in quantum polynomial time. To avoid such computations, we apply Cornacchia's algorithm only when r is a prime and otherwise sample another close vector from the lattice.

Assuming the values of r behave like random values around pN_1^3 for the close vectors, one expects to choose $\log(pN_1^3)$ different vectors (c_1, d_1) before finding a solution for a_1, b_1 with Cornacchia's algorithm. If we do not apply Cornacchia's algorithm unless r is prime, we expect furthermore to sample roughly $\log(pN_1^3)$ values for (c_1, d_1) until r is prime.

The volume of the translated lattice is N_1. Thus, for a generic lattice for which the Gaussian heuristic holds we expect to find a lattice point at distance N_1 from $(\lambda c_0/N_1, \lambda d_0/N_1)$. Furthermore, we can use the Hermite constant for 2-dimensional lattices to trivially bound the distance between this lattice point and the next $2\log(pN_1^3)$ closest lattice points by $\frac{8}{3}\log(pN_1^3)\sqrt{N_1}$. Thus, heuristically r is positive for the expected number of vectors (c_1, d_1) that we need to sample, whenever $eN_2 > pN_1^3 + 8/3\log(pN_1^3)\sqrt{N_1^3}$.

Remark 4.20. Note that for specific lattices, the Gaussian heuristic might be violated. In the worst case, we can only expect to find a lattice point at distance N_1^2 from $(\lambda c_0/N_1, \lambda d_0/N_1)$ and overall solutions require roughly $eN_2 > pN_1^4$.

It is easy to see that a solution for (a_1, b_1, c_1, d_1) as computed with the routine described above satisfies Equation (9). The full algorithm is summarized in Algorithm 3 and an implementation in MAGMA is available[4].

Algorithm 3: Lift element from $\pi\mathbb{Z}[\iota]$ to quaternion of norm N_2 or eN_2

```
   Input: $\theta = \pi(c_0 + d_0\iota) \in \mathrm{End}(E_0)$, and parameters $p$, $\varepsilon$, $N_1$, $N_2$
   Output: $\theta' = N_1a_1 + N_1b_1\iota + (\lambda c_0 + N_1c_1)\pi + (\lambda d_0 + N_1d_1)\iota\pi$ satisfying
           $\mathrm{Norm}(\theta') = N_2$ or $eN_2$ with probability $1 - \varepsilon$ and $\perp$ otherwise
 1 $e \leftarrow$ least positive integer s.t. $eN_2/p(c_0^2 + d_0^2) \pmod{2N_1}$ is a quadratic
     residue;
 2 Compute $\lambda$ in $eN_2 \equiv \lambda^2 p(c_0^2 + d_0^2) \pmod{2N_1}$;
 3 Compute affine relation $c_0c_1 + d_0d_1 \equiv C \pmod{N_1}$;
 4 Define translated lattice $L$ containing all $(c_1, d_1)$ satisfying the affine relation;
 5 $B \leftarrow \log(\varepsilon)\log(pN_1^3)/\log(1 - \log^{-1}(pN_1^3))$;
 6 for $m = 1, \ldots, B$ do
 7     Compute next closest vector $(c_1, d_1)$ to $(-\lambda c_0/N_1, -\lambda d_0/N_1)$ in $L$;
 8     $r \leftarrow \frac{\mathrm{Norm}(\theta') - p((\lambda c_0 + N_1c_1)^2 + (\lambda d_0 + N_1d_1)^2)}{N_1^2}$;
 9     if $r$ prime then
10         Use Cornacchia's algorithm to find $a_1, b_1$ such that $a_1^2 + b_1^2 = r$ or
             determine that no solution exists;
11     if solution found then
12         return $\theta' = N_1a_1 + N_1b_1\iota + (\lambda c_0 + N_1c_1)\pi + (\lambda d_0 + N_1d_1)\iota\pi$;
13 return $\perp$
```

An examination of Algorithm 3 shows that it aborts after a fixed number of trials for pairs (c_1, d_1), which leads to the following result.

Lemma 4.21. *Algorithm 3 always terminates and is correct if it returns a solution.*

We conclude this section by investigating the heuristic probability of the lifting algorithm returning a solution or aborting unsuccessfully, as well as its complexity.

Lemma 4.22. *Let $0 < \varepsilon < 1$. Assume r in Line 8 of Algorithm 3 behaves like a random value around pN_1^3. Then we expect Algorithm 3 heuristically to return a correct lift with probability $1 - \varepsilon$ and an error $\perp$ otherwise.*

Proof. If r in Line 8 of Algorithm 3 behaves like a random value around pN_1^3, we expect it to be prime with probability roughly $1/\log(pN_1^3)$ and Cornacchia's algorithm to provide a solution with probability approximately $1/(\log(pN_1^3))$ due to Landau [20] and Ramanujan [24]. Iterating over B short vectors (c_1, d_1) of the lattice as defined in Step 6 of Algorithm 3, we therefore expect our algorithm to return $\perp$ with probability

$$\left(1 - \frac{1}{\log(pN_1^3)}\right)^{B/\log(pN_1^3)}.$$

Hence, iterating over $B \geq \log(\varepsilon)\log(pN_1^3)/\log(1 - \log^{-1}(pN_1^3))$ as in Algorithm 3, we fail to find a solution with probability less than ε heuristically. □

Remark 4.23. In Algorithm 2 the lifting of endomorphisms is used for every element of the acting group G or H with cardinality $N_1/2$ and $N_1/4$, respectively. Since we expect the lifting algorithm to fail heuristically with probability ε for every single group element and the functions in Algorithm 2 are only exact shifts of each other when it does not fail a single time, we need to choose ε sufficiently small. Assuming independence between the different executions of the lifting algorithm, we expect to find two functions satisfying the promise of a hidden shift with probability $(1-\varepsilon)^{N_1/2} \approx 1-\varepsilon N_1/2$ by first order Taylor approximation. Thus, choosing $\varepsilon < \frac{1}{N_1}$ we expect our lifting to work with probability roughly $\frac{1}{2}$ on all endomorphisms of G and similarly $\varepsilon < \frac{2}{N_1}$ for the elements in H. By the previous lemma, the lifting remains polynomial in $\log(N_1)$ and $\log(p)$ for any such ε. Choosing ε smaller allows us to heuristically achieve a larger success probability of the algorithm. The worst-case complexity of the lifting increases linearly in $|\log(\varepsilon)|$.

Lemma 4.24. *Let* $0 < \varepsilon < 1$. *Algorithm 3 runs in time polynomial in* $\log p$, $\log N_1$, *and* $|\log(\varepsilon)|$.

Proof. The worst-case runtime of the algorithm stems from sampling B (as defined in Algorithm 3, Line 5) potential values of (c_1, d_1) from a lattice of dimension 2. In each iteration one needs to run a primality test, and apply Cornacchia's algorithm to a prime of size polynomial in p and N_1. □

The main drawback of our lifting algorithm is the requirement of approximately $N_2 > pN_1^3$ in case the Gaussian heuristic is satisfied for the lattice defined in Equation (7), and roughly $N_2 > pN_1^4$ otherwise (see Remark 4.20). This bound might be partially caused by inefficiencies in the lifting algorithm. However, the following remark discusses why we can a priori not expect to find a lifting algorithm for balanced parameters.

Remark 4.25. A randomly chosen non-homogeneous quadratic equation in two variables has in general no solution. Similarly, for arbitrary endomorphisms and any N_1, N_2, we would not expect to find an endomorphism $a_1 + b_1\iota \in \mathbb{Z}[\iota]$ (in the variables a_1, b_1) inducing the same action on $\mathcal{I}$ of degree N_2. Yet, as soon as we lift an endomorphism θ to an endomorphism $\theta' = N_1(a_1 + b_1\iota + c_1\pi) + \lambda\theta$, the degree of the lift will be of degree larger than pN_1^2.

4.5 Algorithm summary

We begin the summary of our attack by proving that a solution to Task 4.4 allows us to construct a malleability oracle for f.

Proposition 4.26. *Let $f_{|\mathcal{I}'} : \mathcal{I}' \to \mathfrak{O}$ be the function defined in (1) restricted to a domain $\mathcal{I}'$ so it is injective, let G be an abelian subgroup of $(\mathcal{O}_0/N_1\mathcal{O}_0)^*$ acting freely and transitively on $\mathcal{I}'$ and let $\varphi : E_0 \to E_0/K$, where $K \in \mathcal{I}'$ is chosen uniformly at random and unknown. Suppose the public parameters allow us to solve Task 4.4 for endomorphisms in G efficiently. Given $\varphi|_{E_0[N_2]}$, we then have a polynomial-time malleability oracle for G at $f_{|\mathcal{I}'}(K)$.*

Proof. We need to show that there exists an efficient algorithm that, on input $f(K) = f_{|\mathcal{I}'}(K) = j(E_0/K)$ and $\theta \in G$, computes $f(\theta(K))$. Let φ be the isogeny corresponding to the cyclic subgroup $K \subset E_0$ of order N_1.

The endomorphism θ has degree N_2 coprime to N_1 and using the efficient solution to Task 4.4, we can compute some θ' of degree N_2 such that it has the same action on the N_1-torsion as θ. Therefore, $f(\theta(K)) = E_0/\theta(K) = E_0/\theta'(K)$ up to isomorphism. By Lemma 4.5, this equals $(E_0/K)/\varphi(\ker\theta')$. Since $\ker\theta'$ lies in $E_0[N_2]$, we can compute its image under φ and therefore we can calculate $f(\theta(K)) = (E_0/K)/\varphi(\ker\theta')$ efficiently. □

Proposition 4.26 calls for solutions to the Tasks 4.3 and 4.4. In Sections 4.2 and 4.4 we presented solutions to *variants* of these tasks. We use the remainder of this section to summarize the impact of these variations on the success of our approach.

Restricting the function $f : \mathcal{I} \to \mathfrak{O}$ to a subset $\mathcal{I}'$ such that $f_{|\mathcal{I}'}$ is injective and its image contains $j(E_0/K)$ for the K one aspires to recover requires information on the secret we do not posses. However, we gave three subsets $\mathcal{I}_1$, $\mathcal{I}_2$, $\mathcal{I}_3$ of $\mathcal{I}$ in Section 4.2 such that f restricted to any of these subsets is injective. The images of these sets under f partition all curves at distance N_1 from E_0 up to isomorphism, i.e., one of the three subsets will yield the desired result. Moreover, we provided abelian subgroups of $\mathbb{Q}[\iota] \cap (\mathcal{O}_0/N_1\mathcal{O}_0)^*$ acting freely and transitively on $\mathcal{I}_1$, $\mathcal{I}_2$, and $\mathcal{I}_3$.

We then supply an algorithm to solve Task 4.16, a variant of Task 4.4 when N_1 and N_2 are sufficiently unbalanced, lifting endomorphisms from $\pi\mathbb{Z}[\iota]$ to ones with the same action on $\mathcal{I}$ of degree N_2 or eN_2. Here, e is a small integer depending on the parameters p, N_1, N_2 and the endomorphism. As a consequence, to use the torsion point information of $E_0[eN_2]$ under the secret isogeny given the image of $E_0[N_2]$, we need to guess the action on $E_0[e]$. Furthermore, we lift all endomorphisms in the acting group and thus we need to guess the action on $E_0[E]$, where E is the least common multiple of all e appearing for the different lifts. In Remark 4.17 we discuss which e might appear depending on the factorisation of N_1. For example, E is 2 if N_1 is a power of 3, or $\operatorname{lcm}(3, 5, 7)$ if N_1 is a power of 2. Guessing the action of the secret isogeny on $E_0[E]$ takes $O(E^3)$ trials. Finally, for efficiency reasons we lift endomorphisms from $\pi\mathbb{Z}[\iota]$, whereas the elements in the abelian groups acting on $\mathcal{I}_1$, $\mathcal{I}_2$, and $\mathcal{I}_3$ have representatives in $\mathbb{Z}[\iota]$. In Section 4.3 we showed that this is no restriction via the computation of an action of the Frobenius map.

For each combination of guesses of $E_0[E]$ under the secret isogeny and whether f maps the secret K into $f(\mathcal{I}_1)$, $f(\mathcal{I}_2)$ or $f(\mathcal{I}_3)$, we can use a subexponential quantum algorithm such as Kuperberg's [18] to compute the hidden shift for the functions F_K and F_J as defined in Algorithm 2 and verify the output of the algorithm. Both functions are injective and therefore the verification can be achieved by computing both functions on a single element and its shift respectively. Once the premise of a hidden shift is satisfied, Kuperberg's algorithm [18] recovers the (correct) solution to the injective abelian hidden shift problem. Thus, we recover the secret isogeny as described in Section 4. We can summarize our result as follows.

Theorem 4.27. *Let $N_2 > pN_1^4$. Under the heuristics used for the lifting of endomorphisms in Section 4.4, the SIDH problem can be solved in quantum subexponential time via a reduction to the injective abelian hidden shift problem.*

During this section, we have made some restrictions to simplify the presentation of our cryptanalysis. We assumed the starting curve E_0 to be a supersingular curve with j-invariant 1728. However, the attack also applies to other curves with known endomorphism rings that are close to E_0. In Section 4.2, we described the required group action on $\mathcal{I}$ under the further assumption that N_1 is a power of 2, which can be generalized to powers of small primes. A sketch for powers of 3 can be found in Appendix B. Finally, we assumed that $N_1^2 < \frac{p+1}{4}$ in Lemma 4.7. However, to run our attack we can slightly ease this restriction. Namely, if $N_1^2 > \frac{p+1}{4}$, then we choose a divisor N_1' of N_1 such that $N_1'^2 < \frac{p+1}{4}$ and run the attack with N_1' instead. This will reveal the N_1'-part of the isogeny and then we can guess the remaining part. For sufficiently small $\frac{N_1}{N_1'}$, this is only a minor inefficiency.

4.6 Hybrid attacks on overstretched SIDH

In this section, we examine to what extent partial knowledge of the secret, i.e., knowledge of the most or least significant bits, renders the attack more efficient. Moreover, we describe how the attack can be adapted to some further parameters that are not quite sufficiently unbalanced. The idea is to apply exhaustive search to recover parts of the secret isogeny until the remaining part of the isogeny is of such small degree that the attack outlined in this paper can be used to recover the remaining part.

We start with the case where the most significant bits of the secret are leaked or correctly guessed. These bits correspond to the last steps of the secret isogeny in the isogeny graph. Assume N_1 is a power of a prime ℓ. If the most significant k digits of the secret with respect to their representation in base ℓ are leaked or guessed correctly, the partial isogeny which remains to be recovered is of degree

N_1/ℓ^k and we can run our attack as soon as N_1/ℓ^k fulfills the unbalancedness criterion $N_2 > p(N_1/\ell^k)^4$.

The case where the least significant digits are known or guessed requires a little more work. For simplicity of our exposition we assume again that N_1 is a power of 2 as in Section 4.2, but the results generalize to powers of small primes.

Lemma 4.28. *Let G be the group of Proposition 4.11, and let $G' \subset G$ be the subset of the form $\{a + b\iota \mid a \text{ odd, } b \text{ divisible by } 2^k\}$ where we identify two endomorphisms with each other if they differ by multiplication by an odd scalar modulo N_1. Then G' is an abelian subgroup of G.*

Proof. Since G is abelian, it suffices to show that G' is a subgroup. Consider $(a + b\iota)(a' + b'\iota) = (aa' - bb') + (ab' + a'b)\iota$. It is easy to see that $aa' - bb'$ is odd and $ab' + a'b$ is divisible by 2^k if $a + b\iota$ and $a' + b'\iota$ are in G'. □

Assume the least significant k bits of the secret, or equivalently the first k steps of the secret isogeny, are known. Kernels of isogenies of degree $N_1 > 2^k$ that share the same first k steps lie in the same 2^k-torsion subgroup and are therefore congruent modulo 2^k.

Recall the subsets of $\mathcal{I}$ introduced in Section 4.2.

Proposition 4.29. *Let $\mathcal{I}'$ be any subsets of $\mathcal{I}_1 := \{\langle P + [\alpha]Q\rangle \text{ with } 2|\alpha\}$ containing all those cyclic subgroups where the αs are congruent modulo 2^k. The group G' of Lemma 4.28 acts freely and transitively on any $\mathcal{I}'$.*

Proof. First, we need to show that $G' \times \mathcal{I}' \to \mathcal{I}'$ is well-defined. Let $(1 + b\iota)$ be a representative of some element in G' and let $P + k\iota(P)$, for some $k \in \mathbb{Z}$, be the kernel of an isogeny leading to a curve in $\mathcal{I}'$. We have

$$(1 + b\iota)(P + k\iota(P)) = P + k\iota(P) + b(\iota(P) - kP) \equiv P + k\iota(P) \pmod{b}$$

and as b is divisible by 2^k, $P + k\iota(P) \in \mathcal{I}'$ implies $(1 + b\iota)(P + k\iota(P)) \in \mathcal{I}'$. That the action is free and transitive follows from Proposition 4.11 and a counting argument as $|G|/|G'| = 2^{k-1} = |\mathcal{I}_1|/|\mathcal{I}'|$. □

Similarly, we can take subsets of $\mathcal{I}_2$ and $\mathcal{I}_3$ and restrict the acting group.

This gives rise to an attack strategy when N_2 is not large enough. Guessing k bits of the secret before applying the attack on the remaining part allows an attacker to reduce the requirements on the parameters to $N_2 > p(N_1/2^k)^4$. This is the same as when guessing the last bits of the secret.

Given such a partial isogeny, one computes the correct equivalence class $\mathcal{I}'$ from the kernel of the known part of the isogeny. Moreover, one needs to compute the lifts of elements of G' to endomorphisms of norm N_2 or eN_2. Computing the action of G' on the set $\mathcal{I}'$ allows one to test for the hidden shift property. Once it is satisfied, the secret can be recovered by solving an injective abelian hidden shift problem. Otherwise, one can make another guess on the k bits of the secret.

Apart from reducing the requirements on the unbalancedness, guessing part of the isogeny reduces the number of elements one needs to lift and the size of the hidden shift instance. Depending on the concrete parameter sets provided, one may combine exhaustive search and the attack presented in this paper to recover secrets more efficiently.

5 Childs-Jao-Soukharev's attack on HHS

We begin by providing more detail on how the algorithm proposed by Childs, Jao and Soukharev [5] succeeds to construct an isogeny between two given ordinary elliptic curves in quantum subexponential time. The provided strategy can further be applied to CSIDH [3].

Recall the free and transitive group action from Section 2.3 of the class group on the set of isogenous ordinary curves with the same endomorphism ring. The hard problem is to find an isogeny between two isogenous ordinary elliptic curves with the same endomorphism ring, i.e., reversing this group action. Childs-Jao-Soukharev provide an algorithm that constructs such an isogeny in quantum subexponential time [5] using a reduction to the hidden shift problem.

We summarize the core idea as another instance of our framework using malleability oracles. Let $\mathcal{I} := \mathrm{Cl}(\mathcal{O})$ and $\mathfrak{O} := \mathrm{Ell}_{q,n}(\mathcal{O})$. We can look at the group action defined in Section 2.3 as a one-way function

$$f : \mathcal{I} \to \mathfrak{O} \ , \ [x] \mapsto [x] \cdot j(E_0).$$

Note that the class group $\mathrm{Cl}(\mathcal{O})$ acts on itself and therefore f has a malleability oracle with respect to the class group readily available everywhere on the image, i.e., f is malleable with respect to this group action.

Finding an isogeny φ is now equivalent to finding the ideal class $[\mathfrak{b}]$ in $\mathrm{Cl}(\mathcal{O})$ containing the ideal corresponding to the kernel of φ, i.e., we would like to compute the preimage of f at $j(E_1) = [\mathfrak{b}] \cdot j(E_0)$.

Childs-Jao-Soukharev observed that the functions $F_i : \mathrm{Cl}(\mathcal{O}) \to \mathrm{Ell}_{q,n}(\mathcal{O})$, $[x] \mapsto [x] \cdot j(E_i)$ for $i = 0, 1$ are shifts of each other. Moreover, they are injective functions since the action of the class group on $\mathrm{Ell}_{q,n}(\mathcal{O})$ is free and transitive. The injective abelian hidden shift problem can be solved in quantum subexponential time, which allows one to recover $[\mathfrak{b}]$ and therefore an isogeny $\varphi : E_0 \to E_1$.

Analogously to the case for ordinary curves, the group action in CSIDH utilizing supersingular curves can be attacked this way. Recall that CSIDH uses the $\mathbb{F}_p$-rational endomorphism ring of the fixed starting curve E_0, $\mathcal{O}$. In the Diffie-Hellman-type key exchange, recovering a party's secret key constitutes of computing their secret ideal class $[\mathfrak{b}] \in \mathrm{Cl}(\mathcal{O})$ which satisfies $[\mathfrak{b}] \cdot E_0 = E_B$ for the party's public curve E_B. Through defining functions $F_0, F_1 : \mathrm{Cl}(\mathcal{O}) \to \mathrm{Ell}_p(\mathcal{O})$ by $F_0([x]) = [x] \cdot E_0$ and $F_1([x]) = [x] \cdot E_B$, it is possible to reduce finding Bob's secret key $[\mathfrak{b}]$ to an instance of the injective hidden shift problem: We have $F_1([x]) = F_0([x] \cdot [\mathfrak{b}])$ for any ideal class $[x] \in \mathrm{Cl}(\mathcal{O})$, where the functions are both injective due to the group action being free and transitive.

6 Conclusion and further work

In this paper, we constructed an abelian group action on the key space of the inherently non-commutative SIDH. Having this group action in place allows us to construct a heuristic malleability oracle using the torsion point information provided in SIDH when overstretched and sufficiently unbalanced parameters are being used. This contradicts the commonly believed misconception that no such group action exists in the branch of isogeny-based cryptography where one considers the full isogeny graph of supersingular elliptic curves. We embedded our attack in a more general framework that also captures other quantum attacks on schemes in isogeny-based cryptography.

The attack does *not* apply to balanced parameters as specified in the original SIDH proposal [14] or the NIST post-quantum candidate SIKE [13]. Furthermore, the unbalancedness condition between N_1 and N_2 is stronger than required by the attack from [19]. Interestingly, the obstruction to attack SIDH with balanced parameters in our case does not seem to be directly related to the hindrances in other attacks on unbalanced SIDH exploiting torsion point information [19, 22] but to limitations of the KLPT algorithm [16] and the ones described in Remark 4.25 instead. Improvements to the lifting subroutine included in the KLPT algorithm would not only partially decrease the required unbalancedness of SIDH parameters in this work, but also improve various isogeny-based schemes such as Galbraith-Petit-Silva's signatures [11] and SQISign [9].

Future work will extend the given quantum algorithm to more general group actions of quadratic orders that embed optimally into the (known) endomorphism ring of the starting curve. Hereby, the starting curve does not necessarily need to be of j-invariant 1728. Furthermore, we will generalize the approach to higher genus generalizations of SIDH. Finally, providing applications of this work to areas beyond isogeny-based cryptography is left for future investigation.

It remains an open problem to improve the framework further and to give conditions on the malleability oracle that are sufficient to invert one-way functions in quantum polynomial time.

Acknowledgements. We thank Lorenz Panny for helpful comments on a previous version of this paper and the anonymous reviewers of Eurocrypt2021 for their work and useful feedback. The work of Péter Kutas and Christophe Petit was supported by EPSRC grant EP/S01361X/1. Simon-Philipp Merz was supported by EPSRC grant EP/P009301/1.

Bibliography

[1] Reza Azarderakhsh, Matthew Campagna, Craig Costello, Luca De Feo, Basil Hess, Amir Jalali, David Jao, Brian Koziel, Brian LaMacchia, Patrick Longa, et al. Supersingular isogeny key encapsulation. *Updated parameters for round 2 of NIST Post-Quantum Standardization project*, 2019.

[2] Xavier Bonnetain and André Schrottenloher. Quantum security analysis of CSIDH. In *Advances in Cryptology - EUROCRYPT 2020*, pages 493–522, 2020.

[3] Wouter Castryck, Tanja Lange, Chloe Martindale, Lorenz Panny, and Joost Renes. CSIDH: an efficient post-quantum commutative group action. In *Advances in Cryptology - ASIACRYPT 2018*, pages 395–427, 2018.

[4] Denis X Charles, Kristin E Lauter, and Eyal Z Goren. Cryptographic hash functions from expander graphs. *Journal of Cryptology*, 22(1):93–113, 2009.

[5] Andrew Childs, David Jao, and Vladimir Soukharev. Constructing elliptic curve isogenies in quantum subexponential time. *Journal of Mathematical Cryptology*, 8(1):1–29, 2014.

[6] Giuseppe Cornacchia. Su di un metodo per la risoluzione in numeri interi dell'equazione $\sum_{h=0}^{n} c_h x^{n-h} y^h = p$. *Giornale di Matematiche di Battaglini*, 46:33–90, 1908.

[7] Jean-Marc Couveignes. Hard homogeneous spaces. *IACR Cryptology ePrint Archive*, 2006:291, 1999.

[8] Luca De Feo. Mathematics of isogeny based cryptography. *arXiv preprint:1711.04062*, 2017.

[9] Luca De Feo, David Kohel, Antonin Leroux, Christophe Petit, and Benjamin Wesolowski. SQISign: compact post-quantum signatures from quaternions and isogenies. *IACR Cryptology ePrint Archive*, 2020:1240, 2020.

[10] Steven D. Galbraith, Christophe Petit, Barak Shani, and Yan Bo Ti. On the security of supersingular isogeny cryptosystems. In *Advances in Cryptology - ASIACRYPT 2016*, pages 63–91, 2016.

[11] 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.

[12] Dennis Hofheinz, Kathrin Hövelmanns, and Eike Kiltz. A modular analysis of the Fujisaki-Okamoto transformation. *IACR Cryptology ePrint Archive*, 2017:604, 2017.

[13] David Jao, Reza Azarderakhsh, Matthew Campagna, Craig Costello, Luca De Feo, Basil Hess, Amir Jalali, Brian Koziel, Brian LaMacchia, Patrick Longa, Michael Naehrig, Joost Renes, Vladimir Soukharev, and David Urbanik. SIKE: Supersingular isogeny key encapsulation. `http://sike.org/`, 2017.

[14] David Jao and Luca De Feo. Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. In *International Workshop on Post-Quantum Cryptography*, pages 19–34. Springer, 2011.

[15] Samuel Jaques and John M. Schanck. Quantum cryptanalysis in the RAM model: Claw-finding attacks on SIKE. In *Advances in Cryptology - CRYPTO 2019*, pages 32–61, 2019.

[16] 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.

[17] Greg Kuperberg. A subexponential-time quantum algorithm for the dihedral hidden subgroup problem. *SIAM Journal on Computing*, 35(1):170–188, 2005.

[18] Greg Kuperberg. Another subexponential-time quantum algorithm for the dihedral hidden subgroup problem. *arXiv preprint:1112.3333*, 2011.

[19] Péter Kutas, Chloe Martindale, Lorenz Panny, Christophe Petit, and Katherine E. Stange. Weak instances of SIDH variants under improved torsion-point attacks. *IACR Cryptology ePrint Archive*, 2020:633, 2020.

[20] Edmund Landau. *Über die Einteilung der positiven ganzen Zahlen in vier Klassen nach der Mindestzahl der zu ihrer additiven Zusammensetzung erforderlichen Quadrate.* 1909.

[21] Chris Peikert. He gives C-sieves on the CSIDH. In *Advances in Cryptology - EUROCRYPT 2020*, pages 463–492, 2020.

[22] Christophe Petit. Faster algorithms for isogeny problems using torsion point images. In *Advances in Cryptology - ASIACRYPT 2017*, pages 330–353, 2017.

[23] Christophe Petit and Spike Smith. An improvement to the quaternion analogue of the l-isogeny problem. *Presentation at MathCrypt*, 2018.

[24] Srinivasa Ramanujan. First letter to G.H. Hardy. 1913.

[25] Oded Regev. A subexponential time algorithm for the dihedral hidden subgroup problem with polynomial space. *arXiv preprint:0406151*, 2004.

[26] Alexander Rostovtsev and Anton Stolbunov. Public-key cryptosystem based on isogenies. *IACR Cryptology ePrint Archive*, 2006:145, 2006.

[27] René Schoof. Nonsingular plane cubic curves over finite fields. *J. Comb. Theory, Ser. A*, 46(2):183–211, 1987.

[28] Joseph H Silverman. *The arithmetic of elliptic curves*, volume 106. Springer Science & Business Media, 2009.

[29] Jacques Vélu. Isogénies entre courbes elliptiques. *CR Acad. Sci. Paris, Séries A*, 273:305–347, 1971.

[30] John Voight. *Quaternion algebras.* Preprint, 2018.

[31] William C Waterhouse. Abelian varieties over finite fields. In *Annales scientifiques de l'École Normale Supérieure*, volume 2, pages 521–560, 1969.

A The orbits of the group action

Recall that in Section 4.2, we defined a group acting on the set $\mathcal{I}_1$ which differs from the group acting on the sets $\mathcal{I}_2$ and $\mathcal{I}_3$. The reason for having multiple group actions is that we require them to be free and transitive. Let G be the group defined at the beginning of Section 4.2 and recall that N_1 is a power of 2. Clearly, G acts on all the kernels generated by points of the form $P + \alpha\iota(P)$. Let us study the orbits of this group action in more detail. As we have already seen in Proposition 4.11, the kernels where α is even form a single orbit. Now we show that there are two more orbits occurring when α is odd. For simplicity we will refer to a kernel generated by $P + \alpha\iota(P)$ by $(1, \alpha)$.

Lemma A.1. *Let α be odd. Then $(1, \alpha)$ is either in the orbit of $(1, 1)$ or $(1, 3)$.*

Proof. First we suppose that $\alpha \equiv 1 \pmod 4$ and show that $(1, \alpha)$ is in the orbit of $(1, 1)$ in this case. For this, we must prove the existence of an odd λ and an even b such that the following system is satisfied: $\lambda(1 + b) = 1$ and $\lambda(1 - b) = \alpha$.

Solving the system, we find that $\lambda = \frac{1+\alpha}{2}$ and $b = \frac{1-\lambda}{\lambda}$. These satisfy the required criteria since $1 + \alpha \equiv 2 \pmod 4$, hence λ is odd and $1 - \lambda$ is even.

Now suppose that $\alpha \equiv 3 \pmod 4$. In this case, we show that $(1, \alpha)$ is in the orbit of $(1, 3)$. Again there must exist an odd λ and an even b such that both $\lambda(1 + 3b) = 1$ and $\lambda(3 - b) = \alpha$.

This implies that $\lambda = \frac{1+3\alpha}{10}$, which is an odd integer because α is congruent to 3 modulo 4 and so $1 + 3\alpha$ is congruent to 2 modulo 4. Now one can calculate that b equals $\frac{1-\lambda}{3\lambda}$ which is even since λ is odd, proving the second case. □

By Lemma A.1 the group action defined above has three orbits on $\mathcal{I}$. However, the action of all of G is no longer free on the (smaller) orbits corresponding to an odd α.

B Generalizing Section 4.2 to $N_1 = 3^k$

In this section we sketch a generalization of Section 4.2 to the case where N_1 is a power of 3.

Lemma 4.7 carries over to this case as 9^k can only be written as a sum of two squares in a trivial fashion. Let P be a point such that $\{P, \iota(P)\}$ is a basis of $E_0[N_1]$. We show that every curve at distance N_1 from E_0 can be reached by an isogeny with a kernel of the form $\langle P + \alpha\iota(P)\rangle$. Let $Q = \beta_1 P + \beta_2\iota(P)$ be a point of order N_1. If β_1 is coprime to 3, then we may multiply Q by an appropriate scalar such that the coordinate of P becomes 1. Suppose that β_1 is divisible by 3. Since Q has order N_1, β_2 is not divisible by 3. Observe that the points Q and $\iota(Q)$ generate kernels leading to isomorphic curves which implies that $\beta_1 P + \beta_2\iota(P)$ and $\iota(Q) = -\beta_2 P + \beta_1\iota(P)$ correspond to isogenies leading to isomorphic curves. Multiplying $\iota(Q)$ with an appropriate scalar, we obtain a kernel generator of the form $P + \alpha\iota(P)$.

However, some curves of the form $E_0/\langle P + \alpha\iota(P)\rangle$ may be pairwise isomorphic. Namely let α be coprime to 3. Then the kernels generated by $P + \alpha\iota(P)$ and $P - \alpha^{-1}\iota(P)$ correspond to isomorphic curves. On the other hand, it is easy to see that α and $-\alpha^{-1}$ are not congruent modulo 3. In particular, all curves at distance N_1 from E_0 can be reached by isogenies with kernels of the form $P + \alpha\iota(P)$ where α is congruent to 0 or 1 modulo 3. With a calculuation similar to the one in Section 4.2, it can be shown that these curves are pairwise non-isomorphic.

The acting group can be defined in a similar fashion, namely as the endomorphisms of the form $a + b\iota$ where b is divisible by 3 and two endomorphisms are identified whenever they are the same modulo N_1 up to multiplication by a scalar coprime to N_1. For simplicity we refer to the point $P + \alpha\iota(P)$ as $(1, a)$. Similarly to Appendix A one can check that the action has two orbits:

1. The orbit of $(1, 0)$ consisting of points of the form $(1, x)$, where 3 divides x.
2. The orbit of $(1, 1)$ consisting of points of the form $(1, x)$, where $x \equiv 1 \pmod 3$.

The orbit of $(1,2)$ contains points of the form $(1,x)$ where x is congruent to 2 modulo 3, but in terms of j-invariants it consists of exactly the same curves as the second orbit. Since all these orbits have the same cardinality as the acting group, the group action is free and transitive, as required.

C Lifting $\boldsymbol{\theta} \in \mathbb{Z}[\boldsymbol{\iota}]$ to an element of norm $\boldsymbol{N_2}$ or $\boldsymbol{eN_2}$

Similar to Section 4.4, let e denote the smallest positive integer such that $eN_2/p(a_0^2+b_0^2)$ is a quadratic residue modulo $2N_1$, where $a_0+b_0\iota \in \mathbb{Z}[\iota]$ is the endomorphism we want to lift. Remark 4.17 regarding the size of e still applies in this case.

In this section we describe how to lift endomorphisms in $\mathbb{Z}[\iota] = \mathbb{Q}[\iota] \cap \mathrm{End}(E_0)$ directly to another endomorphism of E_0 of degree N_2 or eN_2 which has the same action on $\mathcal{I}$. This gives an efficient solution to the following variant of Task 4.4.

Task C.1. *Let N_1, N_2 be integers such that $N_2 > p^2N_1^4$. Given an endomorphism $\theta \in G$ of degree coprime to N_1 and an integer N_2 corpime to N_1, compute an endomorphism θ' of degree N_2 or eN_2 such that $\theta(K) = \theta'(K)$ for all $K \in I$.*

As before, this is a relaxation of Task 4.4 in two ways. First, we require N_2 to be sufficiently large and unbalanced compared to N_1. Second, we allow θ' to be either of degree N_2 or eN_2 for some small integer e.

We now give the algorithm to solve Task C.1. By Lemma 4.10 it suffices to solve Task 4.18 for endomorphisms in $\mathbb{Z}[\iota]$.

Let $\theta' = \lambda a_0 + N_1a_1 + \iota(\lambda b_0 + N_1 b_1) + (\lambda c_0 + N_1 c_1 + \iota(\lambda d_0 + N_1 d_1))\pi$. Then its norm equals

$$\mathrm{Norm}(\theta') = (\lambda a_0+N_1a_1)^2+(\lambda b_0+N_1b_1)^2+p((\lambda c_0+N_1c_1)^2+(\lambda d_0+N_1d_1)^2), \quad (8)$$

as $\mathrm{Norm}(x+y\iota) = x^2+y^2$. Since $\theta \in \mathbb{Z}[\iota]$ implies $c_0 = d_0 = 0$, Equation (8) simplifies to

$$\mathrm{Norm}(\theta') = (\lambda a_0 + N_1a_1)^2 + (\lambda b_0 + N_1 b_1)^2 + pN_1^2(c_1^2+d_1^2). \quad (9)$$

We want to compute θ' such that $\mathrm{Norm}(\theta') = eN_2$. Considering Equation (9) modulo $2N_1$, we obtain

$$eN_2 \equiv \lambda^2(a_0^2+b_0^2) \pmod{2N_1}. \quad (10)$$

The choice of e implies that there exists a solution for λ. Compute any such solution, and lift it to the integers in $[1, 2N_1-1]$. This is without loss of generality as any other lift of λ corresponds to a change in a_1, b_1 instead.

Considering the equation modulo N_1^2 yields an affine relation between a_1 and b_1 modulo N_1, i.e.,

$$\lambda(a_0a_1+b_0b_1) \equiv \frac{\mathrm{Norm}(\theta') - \lambda^2(a_0^2+b_0^2)}{2N_1} \pmod{N_1}.$$

Take the affine relation between a_1 and b_1 modulo N_1, say $e_b b_1 = e_a a_1 + e_c + mN_1$ for some fixed integers e_a, e_b, e_c and a variable integer m. Assume $e_b \not\equiv 0 \pmod{p}$ as lifting would be trivial otherwise, and substitute b_1 in Equation (9) modulo the prime p, i.e.,

$$eN_2 \equiv (\lambda a_0 + N_1 a_1)^2 + (\lambda b_0 + N_1 e_b^{-1}(e_a a_1 + e_c + mN_1))^2 \pmod{p}.$$

Note that fixing any value for m leaves a quadratic equation in a_1 modulo p. Fix $m = 0$ and complete the square in the equation to solve it if there exists a solution. Otherwise, increase m by one and repeat. Heuristically, one expects this degree-2 polynomial modulo p to be split with probability 1/2 and hence we expect to iterate twice before finding a solution.

Once a solution for a_1 is obtained modulo p, lift it to the integers. One is left with the problem of representing an integer as the norm of an element in $\mathbb{Z}[\iota]$, i.e., finding c_1 and d_1 such that

$$c_1^2 + d_1^2 = r := \frac{\mathrm{Norm}(\theta') - (\lambda a_0 + N_1 a_1)^2 + (\lambda b_0 + N_1 b_1)^2}{pN_1^2}$$

if they exist. Clearly, r can only be a norm if it is positive. This happens when the parameters N_1 and N_2 are overstretched, and more precisely if $\mathrm{Norm}(\theta') > p^2 N_1^4$.

If the prime decomposition of r is known, Cornacchia's algorithm [6] can efficiently answer the question whether r can be decomposed that way and compute a solution if one exists. Assuming the value of $(\lambda a_0 + N_1 a_1)^2 + (\lambda b_0 + N_1 b_1)^2$ behaves like a random value around $p^2 N_1^4$, one expects to choose $\log(p^2 N_1^4)$ different values for m with a solution to the quadratic equation modulo p before finding a solution with Cornacchia's algorithm.

Remark C.2. Cornacchia's algorithm requires the factorization of r. This could be done in subexponential time on a classical computer or in quantum polynomial time. To avoid such computations, we apply Cornacchia's algorithm only when r is a prime and keep sampling further values for m otherwise.

Since we do not apply Cornacchia's algorithm until r is prime, we expect to sample roughly $\log(p^2 N_1^4)$ values for m until r is prime.

It is easy to see that a solution for (a_1, b_1, c_1, d_1) as computed with the routine described above satisfies Equation (3). The full algorithm is summarized in Algorithm 4.

We conclude this section by investigating the heuristic probability of the lifting algorithm returning a solution or aborting unsuccessfully, as well as its complexity.

The success probability is based on the following heuristic assumptions:

1. The discriminant of $(\lambda a_0 + N_1 a_1)^2 + (\lambda b_0 + N_1 b_1)^2$ in Line 6 of Algorithm 4 is uniformly distributed modulo p.
2. r in Line 9 of Algorithm 4 behaves like a random value around $p^2 N_1^4$.

Lemma C.3. *Let $\varepsilon > 0$. Under the heuristic assumptions mentioned in the preceding paragraph, the algorithm returns a lift with probability $1 - \varepsilon$ and an error $\perp$ otherwise.*

Algorithm 4: Lift element from $\mathbb{Z}[\iota]$ to quaternion of norm N_2 or eN_2

```
    Input: $\theta = a_0 + b_0\iota \in \mathrm{End}(E_0)$, and parameters $p$, $\varepsilon$, $N_1$, $N_2 > p^2N_1^4$
    Output: $\theta' = \lambda a_0 + N_1a_1 + (\lambda b_0 + N_1b_1)\iota + N_1c_1\pi + N_1d_1\iota\pi$ satisfying
            $\mathrm{Norm}(\theta') = N_2$ or $eN_2$ with probability $1-\varepsilon$ and $\perp$ otherwise
 1  $e \leftarrow$ least positive integer s.t. $eN_2/p(a_0^2+b_0^2) \pmod{2N_1}$ is quadratic residue;
 2  Compute $\lambda$ in $eN_2 \equiv \lambda^2(a_0^2+b_0^2) \pmod{2N_1}$;
 3  Compute linear relation between $a_1$ and $b_1$ modulo $N_1$, say $e_bb_1 \equiv e_aa_1 + e_c$
      $\pmod{N_1}$ for some integers $e_a, e_b, e_c$, using
        $$\lambda(a_0a_1 + b_0b_1) \equiv \frac{eN_2 - ((\lambda a_0)^2 + (\lambda b_0)^2)}{2N_1} \pmod{N_1};$$
 4  $B \leftarrow 2\log(\varepsilon)\log(p^2N_1^4)/\log(1-\log^{-1}(p^2N_1^4))$;
 5  for $m = 0, 1, \dots, B$ do
 6      Substitute $b_1$ using expression $e_bb_1 = e_aa_1 + e_c + mN_1$ in
          $$eN_2 \equiv (\lambda a_0 + N_1a_1)^2 + (\lambda b_0 + N_1b_1)^2 \pmod{p};$$
 7      if solution for $a_1 \pmod{p}$ exists then
 8          Compute $a_1$ and $b_1$ modulo $p$ and lift them to integers in $[-p/2, p/2]$;
 9          $r \leftarrow \frac{eN_2-((\lambda a_0+N_1a_1)^2+(\lambda b_0+N_1b_1)^2)}{pN_1^2}$;
10          if $r$ is prime then
11              Use Cornacchia's algorithm to decompose $r$ as sum of two squares
                  $c_1^2 + d_1^2$ or determine that no solution exists;
12          if solution is found then
13              return $\theta' = \lambda a_0 + N_1a_1 + (\lambda b_0 + N_1b_1)\iota + N_1c_1\pi + N_1d_1\iota\pi$;
14  return $\perp$
```

Proof. Based on the heuristic that the discriminant of $(\lambda a_0 + N_1a_1)^2 + (\lambda b_0 + N_1b_1)^2$ in Step 6 of Algorithm 4 is uniformly distributed modulo p, we expect to find a solution for $a_1 \pmod{p}$ for half of the chosen m. Moreover, if r (Line 13, Algorithm 4) behaves like a random value around $p^2N_1^4$, we expect it to be prime with probability roughly $1/\log(p^2N_1^4)$ and Cornacchia's algorithm to provide a solution with probability roughly $1/(\log(p^2N_1^4))$ due to Landau [20] and Ramanujan [24]. Iterating over B values of m, we therefore expect our algorithm to return $\perp$ with probability

$$\left(1 - \frac{1}{\log(p^2N_1^4)}\right)^{B/2(\log(p^2N_1^4)}.$$

In particular, iterating over $B \geq 2\log(\varepsilon)\log(p^2N_1^4)/\log(1-\log^{-1}(p^2N_1^4))$ as in Algorithm 4, we expect to fail to find a solution with probability less than ε heuristically. □

Finally, it is easy to observe the following result.

Lemma C.4. *Algorithm 4 always terminates and it runs in time polynomial in* $\log p$, $\log N_1$ *and* $|\log(\varepsilon)|$ *for every* $\varepsilon > 0$.

Proof. For any $\varepsilon > 0$, the worst-case runtime of the algorithm stems from the iteration over up to $2\log(\varepsilon)\log(p^2N_1^4)/\log(1-\log^{-1}(p^2N_1^4))$ values of m. In each iteration one needs to solve at most one quadratic equation modulo p, and apply Cornacchia's algorithm to a prime of size polynomial in p and N_1. □

The main drawback of this lifting algorithm is the requirement for the unbalancedness $N_2 > p^2N_1^4$. In Section 4.3, we argued why we can lift endomorphisms from $\pi\mathbb{Z}[\iota]$ instead, which is possible with an unbalancedness of $N_2 > pN_1^4$ as described in Section 4.4.