From: (b) (6)
To: Moody, Dustin (Fed)
Subject: Re: more isogeny paper drama
Date: Friday, December 22, 2017 3:46:28 PM
Attachments: 0052.pdf

Hi, Dustin,

Sorry for the slow reply. and thanks for the help. Here it is.

Cheers,
Daniel

On Fri, Dec 22, 2017 at 1:18 PM, Moody, Dustin (Fed) <dustin.moody@nist.gov> wrote:

> If it's okay - I'll take a better look next week.
>
> My first thoughts on reading it: I somewhat agree with the idea that these problems are "natural". But many authors claim this for all sorts of problems, and I don't think it's that big of a deal. The most natural isogeny problems are the ones mentioned by Galbraith (it's cited in the references). The word "natural" is subjective, so doesn't matter a lot to me if somebody claims it. As to the reductions, I'd have to look at the paper again to be sure.
>
> I'm not sure I still have the .pdf. Can you send it again?
>
> ---
>
> **From:** Daniel Smith (b) (6)
> **Sent:** Friday, December 22, 2017 12:35:47 PM
> **To:** Moody, Dustin (Fed)
> **Subject:** more isogeny paper drama
>
> Hi, Dustin,
>
> I don't know if you're available at this point or not, but I wanted to update you on a comment that has just come up in discussing the isogeny paper.
>
> I don't feel as strongly negative about the paper as my subreviewer. But I do disagree with the author's claims that the problems they study are "natural" and that their "formulations give the most natural and cohesive framework for understanding these problems".
>
> In fact I find most of the problems in Section 3 not at all natural.
>
> Problem 1 is the standard problem that is studied in all papers, and problem \hat{1} is its dual. These are fine, but not new.
>
> Problem 2 and \hat{2} involve an additional isogeny, which is equivalent to giving a subgroup of the kernel of the secret isogeny. This is not at all a natural problem. It does not arise naturally in any crypto context. Now, this problem is claimed (Theorem 2) to be equivalent to Problem 1. But looking at the proof of Theorem 2 we see that phi' = ID and so X = E' (and so the "subgroup" of the kernel is the trivial group). The point is that Problem 2 has a decisional aspect ("return No otherwise") and that is all that is being used to show equivalence to KeyValidation (the decisional problem). So the proof is really just trivial change of notation between two identical problems.

> Problem 3 considers the set of all isogenies. But we already know that more than one isogeny means there is a short cycle in the isogeny graph, and that this is a collision in the Microsoft hash. Such events have been long ago proved super-rare (Charles-Goren-Lauter). So Problem 3 is essentially spurious generality. Lemmas 5 and 6 are giving an alternative way to prove a weaker result than the collision-resistance of the Microsoft hash.
>
> Do you have any opinion on this?
>
> Cheers,
> Daniel

On the Equivalence of Torsion-Point Isogeny Problems

David Urbanik[1] and David Jao[2]

[1] Department of Pure Mathematics
[2] Department of Combinatorics and Optimization
University of Waterloo, Waterloo ON, Canada
{dburbani,djao}@uwaterloo.ca

Abstract. Recently, several candidates for quantum-resistant cryptographic primitives based on computational problems involving isogenies between supersingular elliptic curves have emerged. Although computational problems involving isogenies have attracted prior interest in mathematics and cryptography, the particular problems involved in these cryptosystems are unusual in that additional information is revealed about the isogeny to the attacker. Consequently, there has been comparatively little study of the isogeny problems which underlie these quantum-resistant proposals. In this paper, we remedy this situation in two respects. First, under randomized polynomial-time reductions, we prove a six-way equivalence between certain natural candidate problems whose security underlies these cryptosystems. Secondly, using these equivalences, we give formulations of the computational and decisional versions of the Supersingular Isogeny Diffie-Hellman problem, and prove that certain oracles capable of solving these problems are also equivalent to the preceding six. We argue that our formulations give the most natural and cohesive framework for understanding these problems and their relationships, and that these results clarify the assumptions underlying the security of these cryptosystems.

Keywords: isogeny-based cryptography, SIDH, torsion points, supersingular elliptic curves, equivalence theorems

1 Introduction

In 2011, with the aim of achieving a quantum-resistant cryptosystem, Jao and De Feo proposed a key exchange protocol [11] based on the security of certain conjecturally hard problems involving isogeny computations between supersingular elliptic curves. The key exchange protocol, commonly called SIDH (Supersingular Isogeny Diffie-Hellman), functions analogously to the classical Diffie-Hellman protocol, where the difficulty of discrete log problems is replaced by the difficulty of certain "isogeny-finding" problems, and the difficulty of the computational and decisional Diffie-Hellman problems is replaced by the difficulty of the computational and decisional SIDH problems. As in the classical case, one is interested in determining whether these problems are equivalent.

To make the analogy between the two cases more explicit, we recall the case of ordinary Diffie-Hellman. In ordinary Diffie-Hellman, one is given a cyclic group generated by an element g. Alice and Bob choose private integers a and b respectively, and compute public keys g^a and g^b. They then exchange the public keys, and each computes $(g^b)^a = (g^a)^b$, which they take to be their shared secret. The difficult problems underlying such a scheme are: given (g, g^a) find a, given (g, g^b) find b, and given (g, g^a, g^b) find g^{ab}.

Intuitively, in the SIDH case, one would like a protocol which proceeds as follows. One begins with a supersingular elliptic curve E, analogous to the element g. Alice and Bob choose private subgroups A and B and compute public keys E/A and E/B, where the public keys are so-called "quotient curves" corresponding to those subgroups. They then exchange the public keys and compute $(E/A)/B = (E/B)/A$, which they take to be their shared secret. The difficult problems underlying the scheme would be: given $(E, E/A)$ find A, given $(E, E/B)$ find B, and given $(E, E/A, E/B)$ find $E/\langle A, B\rangle$.

Unfortunately, there are various technical obstructions to proceeding directly in this manner. One such obstruction is the difficulty of computing the quotient $(E/A)/B$ from knowledge of E/A and B, since B is not actually a subgroup of E/A and one needs instead the image of B under the quotient map $\phi_A \colon E \to E/A$. Another such obstruction is that the curves $(E/A)/B$ and $(E/B)/A$ are not in general equal, but only isomorphic, so one needs to take an isomorphism invariant to be the shared secret. The key insight in the Jao-De Feo paper can be viewed as a prescription for resolving these problems and making a protocol of this form computationally tractable.

When one follows the Jao-De Feo prescription, one arrives at a problem like the following. Let ℓ_1 and ℓ_2 be small distinct primes, e_1 and e_2 exponents such that $\log(\ell_1^{e_1}) \approx \log(\ell_2^{e_2})$, and $p = \ell_1^{e_1}\ell_2^{e_2} \pm 1$ a prime. Given two supersingular elliptic curves E and E' defined over $\mathbb{F}_{p^2}$, and the values of a degree $\ell_1^{e_1}$ isogeny $\phi \colon E \to E'$ on $E[\ell_2^{e_2}]$, find ϕ. If we continue with the analogy above, ϕ is one of the "quotient" maps $\phi_A \colon E \to E/A$ and $\phi_B \colon E \to E/B$. It is known that finding ϕ is equivalent to finding its kernel, which is either Alice's private subgroup A or Bob's private subgroup B (we provide a proof in Section 2). Thus, by the analogy above, we see that this problem is the analogue of the discrete logarithm problem in classical Diffie-Hellman.

Little is known about the security of this exact problem. More general isogeny problems have been studied in the literature[1, 4, 6], but the majority of such studies ignore the information provided by the action of the isogeny on $E[\ell_2^{e_2}]$. One exception is the recent work of Petit[14], but his work focuses on attacks on certain "overstreched" variants of SIDH and is orthogonal to our work. A second exception is an argument made by Thormarker[19] and Galbraith and Vercauteren[9], which shows heuristically that one can reduce a computational SIDH problem to its decisional variant. To use this reduction to find ϕ_A or ϕ_B, and hence Alice's or Bob's private key, one must assume heuristics relating to the number of isogenies of the desired form. In particular, the argument involves "backtracking" along a graph formed by all degree-ℓ_1 isogenies in hopes

of obtaining a path of length e_1 from E to E'. To ensure that the size of the graph computed is polynomially bounded, it is necessary to assume that the number of possible isogenies $\phi\colon E \to E'$ is also polynomially bounded.

In Section 3, we introduce six natural candidate problems which underlie the security of cryptosystems obtained from Jao-De Feo-like constructions, and prove that they are all equivalent under randomized polynomial-time reductions. Our arguments put the above heuristic assumptions on firm footing, and establish conditions on when the $\ell_2^{e_2}$-torsion information suffices to determine the isogeny. This, in turn, shows an unconditional equivalence between the six candidate problems and the decisional variant considered by Thormarker [19] Galbraith and Vercauteren [9].

In Section 4, we give a formulation of the SIDH protocol which is more natural in two respects. Firstly, we define the decisional problem studied by Thormarker [19] and Galbraith and Vercauteren [9] as an instance of the so-called static Key Validation Problem, which has been studied by several authors[13, 3, 7] in hopes of obtaining a static-static or non-interactive key exchange (NIKE). Using the equivalences in Section 3, we show that this problem suffices to break computational SIDH, which suggests that it should be intractable. Secondly, we show that in our formulation of SIDH, a combination of the decisional and computational SIDH problems is equivalent to the problems studied in Section 3. We argue that our results provide the first clear picture of the problem landscape underlying the SIDH cryptosystem and other cryptosystems based on similar constructions.

2 Preliminaries on Isogeny Problems

For general background on elliptic curves we refer to Silverman [16].

Let E and E' be two elliptic curves defined over a finite field in characteristic p. An isogeny $\phi\colon E \to E'$ is a surjective rational map of curves which is also a group homomorphism between the elliptic curve groups of E and E'. It can be shown that all isogenies have finite kernels. Isogenies whose kernel consists of only the identity element are called isomorphisms, and have inverse maps that are also isogenies. Isogenies have a *degree*, which is their degree as a rational map. Isogenies are called *separable* if the size of their kernels is equal to their degree. Each isogeny $\phi\colon E \to E'$ has a dual isogeny $\widehat{\phi}\colon E' \to E$ which satisfies $\phi \circ \widehat{\phi} = [\deg \phi]$ and $\widehat{\phi} \circ \phi = [\deg \phi]$, where $[\deg \phi]$ is the scalar-multiplication-by-$\deg \phi$ map on the appropriate curve.

It is known that for any finite subgroup H of an elliptic curve E, there is a unique curve up to isomorphism, denoted E/H, which is the image of a seperable isogeny $\phi_H\colon E \to E/H$ with kernel exactly H. Hence to each finite subgroup H of E we may associate an isomorphism class of curves which are the codomains of isogenies with kernel H. Vélu [20] gave formulas using which one may compute the curves E/H from H and compute the corresponding isogenies with kernel H.

In this paper, we are interested in isogenies resulting from the following construction. Pick a prime $p = \ell_1^{e_1}\ell_2^{e_2} \pm 1$, where ℓ_1 and ℓ_2 are small distinct primes, and $\log(\ell_1^{e_1}) \approx \log(\ell_2^{e_2})$. It can then be shown that every supersingular elliptic curve E in characteristic p is defined over $\mathbb{F}_{p^2}$, and that $E(\mathbb{F}_{p^2}) \cong (\mathbb{Z}/\ell_1^{e_1}\ell_2^{e_2}\mathbb{Z}) \times (\mathbb{Z}/\ell_1^{e_1}\ell_2^{e_2}\mathbb{Z})$. We are then interested in cyclic $\ell_i^{f_i}$-degree isogenies from E, which we think of as being obtained via quotients $E/\langle P\rangle$, where $P \in E[\ell_i^{e_i}]$ is a point generating a cyclic subgroup of order $\ell_i^{f_i}$. We note that $E[\ell_i^{e_i}]$ is also defined over $\mathbb{F}_{p^2}$.

The reason this particular construction is of interest is that it allows for efficient computation of the isogenies involved. Typically, p is proportional to the relevant security parameter, and so is chosen to be exponentially large. Consequently, these isogenies also have exponentially large degree for f_i sufficiently large, which could make their computation difficult. However, the fact that $\ell_i^{f_i}$ is smooth for $i = 1$ and $i = 2$ allows one to factor these isogenies as a composition of f_i isogenies of degree ℓ_i, each of which is easy to compute. To construct such a factorization, we first compute the subgroup filtration

$$\{\mathcal{O}_E\} = \langle [\ell_i^{f_i}]P\rangle \subset \langle [\ell_i^{f_i-1}]P\rangle \subset \cdots \subset \langle [\ell_i]P\rangle \subset \langle P\rangle \ .$$

We then represent ϕ as a composition of isogenies $\phi = \phi_{f_i} \circ \cdots \phi_1$, where $\phi_k \colon E_{k-1} \to E_k$, $E_0 = E$ and we set $\phi_0 = \mathrm{id}_E$. Given E_{k-1}, we compute $E_k := E_{k-1}/\langle \phi_{k-1}([\ell_i^{f_i-k}]P)\rangle$ and $\phi_k \colon E_{k-1} \to E_k$ using Vélu's formulas. Since each ℓ_i is small, this step can be done efficiently, and so the total time taken to compute ϕ is determined by f_i, which is the number of ℓ_i-degree isogenies in the factorization. Based on our construction, this is at most $\frac{1}{2}\log_{\ell_i}(p)$, and so this can be done in time logarithmic in the security parameter.

The above discussion implies that, in isogeny-based cryptography, algorithms with running time polynomial in $\log(p)$ are polynomial-time algorithms in the complexity-theoretic sense, and algorithms with running times bounded below by p^s for some positive exponent s are exponential-time algorithms. For instance, the best attacks on SIDH have classical complexity $O(p^{1/4})$ and quantum complexity $O(p^{1/6})$. When we discuss polynomial-time and exponential-time complexity throughout this paper, we will always mean in the complexity-theoretic sense, i.e. polynomial and exponential in $\log(p)$.

As discussed in Section 1, the action of isogenies on the $\ell_i^{e_i}$-torsion subgroups of elliptic curves is important in our discussion. To make working with these subgroups easier, we prove the following lemma.

Lemma 1. *For ease of notation, denote $\ell = \ell_i$, $e = e_i$, and $n = p \mp 1 = \ell_1^{e_1}\ell_2^{e_2}$, where $i \in \{1, 2\}$. Then there exists a randomized polynomial-time algorithm to compute a $\mathbb{Z}$-basis for $E[\ell^e] \cong (\mathbb{Z}/\ell^e\mathbb{Z}) \times (\mathbb{Z}/\ell^e\mathbb{Z})$.*

Proof. Consider a curve equation $y^2 = x^3 + ax + b$ for E in Weierstrass form. It is well-known that by choosing a random value of x in $\mathbb{F}_{p^2}$ and computing the square-root of the right-hand side one obtains a random point in $E(\mathbb{F}_{p^2})$ with probability asymptotically $\frac{1}{2}$. Furthermore, ignoring the case of the identity point (which is of no interest), the case when $y = 0$ (which is easily accounted

for as a special case), and making sure to choose the sign of y uniformly at random, this process will sample the elements of $E(\mathbb{F}_{p^2})$ uniformly at random. Since these steps may be computed in polynomial time, we may assume that we can efficiently sample uniformly random points of $E(\mathbb{F}_{p^2})$.

Note that ℓ^e is relatively prime to n/ℓ^e, and that there is a factorization $E(\mathbb{F}_{p^2}) = E[\ell^e] \times E[n/\ell^e]$. Thus a random point P in $E(\mathbb{F}_{p^2})$ can be thought of as corresponding in a unique way to a pair (P_1, P_2), where $P = P_1 + P_2$, $P_1 \in E[\ell^e]$ and $P_2 \in E[n/\ell^e]$. If we compute $[n/\ell^e]P$ we will get $([n/\ell^e]P_1, \mathcal{O}_E)$. Since the map $[n/\ell^e]$ restricts to an isomorphism on $E[\ell^e]$, this process can be thought of as selecting an element of $E[\ell^e]$ uniformly at random.

To complete the proof, we simply randomly choose elements of full order in $E[\ell^e]$ until we obtain two that are independent. Note that because $E[\ell^e] \cong (\mathbb{Z}/\ell^e\mathbb{Z}) \times (\mathbb{Z}/\ell^e\mathbb{Z})$, an element will have full order provided that at least one of its coefficients under such an isomorphism is not divisible by ℓ, which will happen a $1 - \frac{1}{\ell^2}$ fraction of the time. Two such full-order elements P and P' will be independent provided that $\langle P \rangle \cap \langle P' \rangle = \{\mathcal{O}_E\}$, which is equivalent to the statement that $\langle [\ell^{e-1}]P \rangle \neq \langle [\ell^{e-1}]P' \rangle$. There are $\ell + 1$ subgroups of order ℓ in $E[\ell^e]$, so this happens with probability $1 - \frac{1}{\ell+1}$. This shows that selecting random pairs of full order points will give us a basis with probability bounded below by a constant, which completes the proof. □

Remark 1. In the paper [2], the above computations are implemented for the case where $\ell_1^{e_1} = 2^{372}$ and $\ell_2^{e_2} = 3^{239}$, and the authors show that finding a basis for $E[\ell_i^{e_i}]$ requires no more than 10 milliseconds on a modern machine.

Remark 2. Given a basis for $E[\ell^e]$, one also has a basis for $E[\ell^{e-k}]$ obtained via scalar multiplication by $[\ell^k]$.

A useful fact to have when one deals with these bases is that one can efficiently represent points in terms of them. The next lemma formalizes this fact.

Lemma 2. *With the same setup as Lemma 1, and P a point in $E[\ell^e]$, there is a polynomial-time algorithm to compute the coefficients of P with respect to a $\mathbb{Z}$-basis for $E[\ell^e]$.*

Proof. Since ℓ^e is smooth, this computation is simply a 2-dimensional discrete logarithm problem in a group of smooth order, and such problems have been extensively studied; for instance, see [18]. Also see [2] for an efficient implementation in the SIDH case. □

The main purpose of representing points with respect to chosen bases is so that we can represent the action of isogenies on torsion subgroups as matrices. For instance, we have said that the situations of interest involve the action of an isogeny $\phi\colon E \to E'$ of degree dividing ℓ^e on $E[n/\ell^e]$. Since $\phi(E[n/\ell^e]) = E'[n/\ell^e]$, we can represent this portion of the map with respect to bases for $E[n/\ell^e]$ and $E'[n/\ell^e]$ as a matrix over $\mathbb{Z}/(n/\ell^e)\mathbb{Z}$. We will refer to such a representation as a *matrix representation* of $\phi\big|_{E[n/\ell^e]}$. Because ϕ is bijective between $E[n/\ell^e]$ and

$E'[n/\ell^e]$, such a matrix is invertible, and it allows us to compute the group-theoretic inverse of $\phi\big|_{E[n/\ell^e]}$. This fact is very useful in our reductions, so we catalog it for future reference in Lemma 3.

Lemma 3. *Suppose that $\phi\colon E \to E'$ is an isogeny of degree dividing ℓ^e, with the same definitions as Lemma 1. Then there is a polynomial-time randomized algorithm to compute the group isomorphism $(\phi\big|_{E[n/\ell^e]})^{-1}$.*

We mentioned earlier that one can think of separable isogenies as being in correspondence with their kernels, and also with their duals. Since we wish to use these correspondences in the context of polynomial-time reduction theorems, we will need the fact that these correspondences can be computed efficiently. Lemma 4 serves this purpose.

Lemma 4. *Suppose that $\phi\colon E \to E'$ is an isogeny with degree dividing ℓ^e, with the same definitions as Lemma 1. Then there is a randomized polynomial-time algorithm to compute any of the following following four pieces of data from knowledge of just one of them.*

(i) The kernel H of ϕ.
(ii) A sequence of prime degree rational maps $\phi_1, \ldots, \phi_s$ such that $\phi = \phi_s \circ \cdots \circ \phi_1$.
(iii) The kernel H' of $\widehat{\phi}$.
(iv) A sequence of prime degree rational maps $\phi'_1, \ldots, \phi'_s$ such that $\widehat{\phi} = \phi'_s \circ \cdots \circ \phi'_1$.

Proof. We have already seen that given (i) one may obtain (ii), and analogously given (iii) one may obtain (iv). Hence to complete the proof, it suffices to show that given (ii) we can find (iii), and analogously given (iv) we can find (i).

In the first case, we use Lemma 1 (or Remark 2) to choose a basis for $E[\deg \phi]$. We know that $\widehat{\phi} \circ \phi = [\deg \phi]$, which has kernel exactly $E[\deg \phi]$. Hence the kernel H' of $\widehat{\phi}$ is exactly $\phi(E[\deg \phi])$, which is easily computed by evaluating ϕ on the basis for $E[\deg \phi]$. The other case is analogous. □

3 Equivalence of Isogeny Problems

Throughout this section, we fix $\ell = \ell_i$ and $e = e_i$ for some $i \in \{1, 2\}$, and $n = \ell_1^{e_1}\ell_2^{e_2}$. Recalling that these primes and their exponents are chosen such that $\log(\ell_1^{e_1}) \approx \log(\ell_2^{e_2})$, we formalize this property precisely by supposing that

$$|\log(\ell_1^{e_1}) - \log(\ell_2^{e_2})| < \kappa,$$

where $\kappa = O(1)$ is constant. We note that for the most widely-used parameters, which were first suggested in [3], κ is less than 3, and one expects there to be enough primes of the right form so that asymptotically one could assume $\kappa = o(1)$.

We begin by defining three natural problems of interest in isogeny-based cryptography. We will see that the other three problems which comprise the promised six-way equivalence are in some sense "dual" to these problems. The problems we consider all involve the evaluation of isogenies on the n/ℓ^e-torsion subgroup of an elliptic curve E. Since there are exponentially many points in this subgroup, such an evaluation is represented in practice by the values of an isogeny ϕ on a basis for $E[n/\ell^e]$. For ease of terminology, we say that $P, Q \in E[n/\ell^e]$ form a *basis pair* if together they generate $E[n/\ell^e]$. We will often use the fact that if $\eta\colon E \to E'$ is any isogeny of degree relatively prime to n/ℓ^e, and $P, Q \in E[n/\ell^e]$ is a basis pair for $E[n/\ell^e]$, then $\eta(P), \eta(Q) \in E'[n/\ell^e]$ is a basis pair for $E'[n/\ell^e]$. Note that this statement applies even if $E' = E$ and η is a scalar multiplication map. Hence, for fixed E, we have the following problems of interest:

(1) Given a curve E', a basis pair $P, Q \in E[n/\ell^e]$, and a basis pair $R, S \in E'[n/\ell^e]$, either
   (i) return an isogeny $\phi\colon E \to E'$ of degree dividing ℓ^e such that $\phi(P) = R$ and $\phi(Q) = S$, or
   (ii) report that one doesn't exist.

(2) Given a curve E', a basis pair $P, Q \in E[n/\ell^e]$, a basis pair $R, S \in E'[n/\ell^e]$, and an additional map $\psi\colon E \to X$, either
   (i) return "Yes" if there exists an isogeny $\phi\colon E \to E'$ of degree dividing ℓ^e which factors through ψ, and such that $\phi(P) = R$ and $\phi(Q) = S$, or
   (ii) return "No" otherwise.

   We say that ϕ *factors through* ψ if there is a $\psi'\colon X \to E'$ such that $\phi = \psi' \circ \psi$.

(3) Given a curve E', a basis pair $P, Q \in E[n/\ell^e]$, and a basis pair $R, S \in E'[n/\ell^e]$, return the set of all isogenies $\phi\colon E \to E'$ of degree dividing ℓ^e such that $\phi(P) = R$ and $\phi(Q) = S$.

For each $i = 1, 2, 3$, let $(O_{E,i})_{\ell^e}$ denote an oracle to solve Problem (i).

Before proving reduction theorems relating these problems, we make some remarks on their naturality. In isogeny-based cryptosystems, important private information is usually represented in the form of either a secret isogeny or (equivalently by Lemma 4) a secret kernel. The attacker is then given points $R = \phi(P)$ and $S = \phi(Q)$, and is tasked with finding ϕ (equivalently, finding its kernel). The above is the spirit of Problem (1), except in principle there could possibly be more than one such ϕ, even though intuitively one suspects such an outcome to be exceedingly unlikely. Hence to find the secret isogeny[3] one may need to find the "right" ϕ, for which it suffices to solve Problem (3). Finally, Problem (2) represents a natural attack strategy on these cryptosystems, in that the fastest known attacks involve some variant of a breadth-first search on the so-called ℓ-*isogeny graph*, and a non-trivial solution to Problem (2) would allow one

[3] For technical reasons even finding the "wrong" ϕ would typically suffice to break isogeny-based schemes despite not necessarily recovering the secret isogeny.

to optimize this search. We recall that the ℓ-isogeny graph is the graph whose vertices are elliptic curves and whose edges are ℓ-degree isogenies[4].

Before we can prove any reduction theorems involving the oracle $(O_{E,3})_{\ell^e}$, we need to know that the set returned by $(O_{E,3})_{\ell^e}$ has polynomial size, since otherwise one cannot even query the oracle in a polynomial-time reduction. The next two lemmas accomplish this task.

Lemma 5. *Let $\phi, \phi' \colon E_1 \to E_2$ be isogenies of degree d from E_1 to E_2. If ϕ and ϕ' agree on N affine points, where $N > 3d^2$, then they are equal.*

Proof. The idea is to translate the problem of determining whether ϕ equals ϕ' to a problem about whether certain algebraic surfaces are equal, and then apply algebro-geometric tools from intersection theory. The proof is somewhat technical, so we give it in Appendix A. □

Lemma 6. *The set returned by $(O_{E,3})_{\ell^e}$ has size polynomial in $\log p$. Furthermore, there is a randomized polynomial-time reduction $(O_{E,3})_{\ell^e} \leq_P (O_{E,1})_{\ell^e}$.*

Proof. Suppose we are given E', a basis pair $P, Q \in E[n/\ell^e]$, a basis pair $R, S \in E'[n/\ell^e]$, and wish to answer queries as $(O_{E,3})_{\ell^e}$. If $\phi \colon E \to E'$ is an isogeny mapping P to R and Q to S, then the homomorphism property determines ϕ on the entire $E[n/\ell^e]$ torsion, and so ϕ is determined on $(n/\ell^e)^2$ points. If $(n/\ell^e)^2 > 3(\ell^e)^2$ then there can be only one such ϕ by Lemma 5. In this case, we may simply call the oracle $(O_{E,1})_{\ell^e}$.

If $(n/\ell^e)^2$ is not greater than $3(\ell^e)^2$, we simply "backtrack" along the isogeny graph until we obtain isogeny problems to which Lemma 5 applies. In particular, we recall that $|\log(\ell^e) - \log(n/\ell^e)| < \kappa$ where $\kappa = O(1)$, and so we may find $k = O(1)$ such that $(n/\ell^e)^2 > 3(\ell^{e-k})^2$. We observe that any isogeny ϕ of degree ℓ^e from E to E' must factor through some curve X as $\phi = \psi' \circ \psi$ where $\psi \colon E \to X$, $\psi' \colon X \to E'$, $\deg \psi = \ell^{e-k}$, and $\deg \psi' = \ell^k$. Moreover, the isogeny ψ is uniquely determined by X and ψ', since in order for ϕ to be a valid isogeny mapping P to R and Q to S, ψ must map P to $(\psi'\big|_{X[n/\ell^e]})^{-1}(R)$ and Q to $(\psi'\big|_{X[n/\ell^e]})^{-1}(S)$. Hence, by Lemma 5, there is at most one isogeny of the right form corresponding to the pair (X, ψ'). Up to isomorphisms of X, there are polynomially many such pairs. Isomorphic X's give the same final isogeny $\phi = \psi' \circ \psi$, since the paths on the isogeny graph, and hence the kernels of the two maps, will be the same. This completes the proof that the set returned by $(O_{E,3})_{\ell^e}$ has polynomial size.

To complete the proof of the reduction $(O_{E,3})_{\ell^e} \leq_P (O_{E,1})_{\ell^e}$, it suffices to show that one may compute the pairs (X, ψ') (up to isomorphisms of X), and also the points $(\psi'\big|_{X[n/\ell^e]})^{-1}(R)$ and $(\psi'\big|_{X[n/\ell^e]})^{-1}(S)$ so that one may query the oracle $(O_{E,1})_{\ell^e}$. To compute the curves X, it suffices to compute all isogenies of degree ℓ^k from E', which we may do since $k = O(1)$. The codomains of these isogenies are the curves we require, and the duals of these maps, which we may compute by Lemma 4, are the maps ψ'. We may then compute the points $(\psi'\big|_{X[n/\ell^e]})^{-1}(R)$ and $(\psi'\big|_{X[n/\ell^e]})^{-1}(S)$ by Lemma 3. □

[4] There is a more sophisticated definition which considers curves and isogenies up to isomorphism, but we will not need it.

The preceding two Lemmas accomplish most of the work necessary to show the equivalence between Problems (1), (2) and (3). Indeed, we have just seen that Problem (3) is no harder than Problem (1). Problem (1) is also clearly no harder than Problem (3), since (3) asks for all such isogenies, and for (1) we only require one such isogeny. Given an oracle for Problem (3), we may also solve Problem (2) simply by computing the isogeny graph corresponding to the isogenies returned by (3), and seeing if the ℓ-isogeny path corresponding to the isogeny $\psi\colon E \to X$ given as input to (2) extends to a path from E to E' of the right degree.

All that remains in order to show that the three problems are equivalent is to show that given an oracle for (2) we may solve either (1) or (3). This result follows straightforwardly from our earlier observation about the importance of Problem (2) in optimizing search algorithms for isogeny problems. Indeed, we may attempt to find isogenies from E to E' of the right form by considering a breadth-first search on the ℓ-isogeny graph starting from E. At each stage, we wish to "prune" the search tree by determining which ℓ-isogeny paths $\psi\colon E \to X$ do not extend to an isogeny $\phi\colon E \to E'$ mapping P to R and Q to S. But this question is exactly the question answered by (2), and so we may easily compute the appropriate (polynomially-sized by Lemma 6) graph which solves Problem (3) (and hence (1)). This discussion completes the proof of Theorem 1.

Theorem 1. *The oracles $(O_{E,1})_{\ell^e}, (O_{E,2})_{\ell^e}$, and $(O_{E,3})_{\ell^e}$ are equivalent under randomized polynomial-time reductions.* □

To achieve the promised six-way equivalence, we define three additional problems which are in some sense "dual" to the three we have already described. They are as follows.

$\widehat{(1)}$ Given a curve E', a basis pair $P, Q \in E[n/\ell^e]$, and a basis pair $R, S \in E'[n/\ell^e]$, either
  (i) return an isogeny $\phi'\colon E' \to E$ of degree dividing ℓ^e such that $\phi'(R) = P$ and $\phi'(S) = Q$, or
  (ii) report that one doesn't exist.

$\widehat{(2)}$ Given a curve E', a basis pair $P, Q \in E[n/\ell^e]$, a basis pair $R, S \in E'[n/\ell^e]$, and an additional map $\psi'\colon E' \to X$, either
  (i) return "Yes" if there exists an isogeny $\phi'\colon E' \to E$ of degree dividing ℓ^e which factors through ψ', and such that $\phi'(R) = P$ and $\phi'(S) = Q$, or
  (ii) return "No" otherwise.

$\widehat{(3)}$ Given a curve E', a basis pair $P, Q \in E[n/\ell^e]$, and a basis pair $R, S \in E'[n/\ell^e]$, return the set of all isogenies $\phi' : E' \to E$ of degree dividing ℓ^e such that $\phi'(R) = P$ and $\phi'(S) = Q$.

For each $i = 1, 2, 3$, let $\widehat{(O_{E,i})}_{\ell^e}$ denote an oracle to solve Problem $\widehat{(i)}$.

Interest in problems $\widehat{(1)}$ and $\widehat{(3)}$ is natural given the equivalence that exists between finding an isogeny and finding its dual by Lemma 4. The naturality of

Problem $\widehat{(2)}$ can also be justified on similar grounds as for (2), in that the best algorithm for finding isogenies between these curves at present involes performing a breadth-first search outwards from both the base curve E *and* the target curve E', and so finding non-trivial optimizations to this search is also important when searching outwards from E'. Problem $\widehat{(2)}$ is also often easier to use when proving reductions, since the strategy of working backwards from the curve E' corresponds most naturally to the backtracking strategy of Thormarker [19] and Galbraith and Vercauteren [9]. For instance, the decisional variant they study may be formulated as the following problem, which we call the *Key Validation Problem* in anticipation of its role in the next section.

Problem 1 (Key Validation). Given E', a basis pair $P, Q \in E[n/\ell^e]$, a basis pair $R, S \in E'[n/\ell^e]$, and $0 \leq k \leq e$, determine whether there exists an isogeny $\phi\colon E \to E'$ of degree dividing ℓ^{e-k} such that $\phi(P) = R$ and $\phi(Q) = S$.

Problem $\widehat{(2)}$ can then be used to prove the following equivalence:

Theorem 2. *The Key Validation problem is equivalent to Problem* $\widehat{(2)}$ *under randomized polynomial-time reductions.*

Proof. Suppose we are given E', a basis pair $P, Q \in E[n/\ell^e]$, a basis pair $R, S \in E'[n/\ell^e]$, and an additional map $\psi'\colon E' \to X$. Let $k = \deg \psi'$. We wish to answer questions as the oracle $\widehat{(O_{E,2})}_{\ell^e}$ using an oracle for the Key Validation problem. The question of whether $\phi'\colon E' \to E$ of degree dividing ℓ^e factoring through $\psi'\colon E' \to X$ exists such that $\phi'(R) = P$ and $\phi'(S) = Q$ is the same as the question of whether there is a map $\psi\colon X \to E$ of degree dividing ℓ^{e-k} which maps $\psi'(R)$ to P and $\psi'(S)$ to Q. This, in turn, is the same as asking whether the map $\widehat{\psi}\colon E \to X$, which would map P to $[\deg \psi]\psi'(R)$ and Q to $[\deg \psi]\psi'(S)$, exists. However, noting that $[\deg \psi]\psi'(R)$ and $[\deg \psi]\psi'(S)$ form a basis pair, this last question may be answered by an oracle for the Key Validation problem.

For the other direction, we suppose we are given E', a basis pair $P, Q \in E[n/\ell^e]$, a basis pair $R, S \in E'[n/\ell^e]$, an integer $0 \leq k \leq e$, and wish to solve the Key Validation problem using an oracle for $\widehat{(2)}$. The Key Validation problem asks whether $\phi\colon E \to E'$ of degree dividing ℓ^{e-k} exists mapping P to R and Q to S. The question is the same as asking whether $\widehat{\phi}$ exists mapping R to $[\deg \phi]P$ and Q to $[\deg \phi]S$. Hence we may query the oracle $\widehat{(2)}$ with the input $(E', [\deg \phi]P, [\deg \phi]Q, R, S, \psi')$, where we take $\psi' = \mathrm{id}_{E'}$, and we try all possible values of $\deg \phi$. Note again that $[\deg \phi]P$ and $[\deg \phi]Q$ form a basis pair. This completes the proof. □

Lastly, we show that the problems $\widehat{(i)}$ are all equivalent to each other, and in fact equivalent to the problems (i). This proof gives the promised six-way equivalence.

Theorem 3. *The Problems* (1), (2), (3), $\widehat{(1)}$, $\widehat{(2)}$, *and* $\widehat{(3)}$ *are all equivalent under randomized polynomial-time reductions.*

Proof. We already know (1), (2), and (3) are equivalent. We first observe that (1) is equivalent to $\widehat{(1)}$. Indeed, to find an isogeny $\phi\colon E \to E'$ mapping P to R and Q to S, it suffices by Lemma 4 to find its dual, which maps R to $[\deg\phi]P$ and S to $[\deg\phi]Q$, where again $[\deg\phi]P$ and $[\deg\phi]Q$ form a basis pair. Hence we may solve (1) using the oracle $\widehat{(O_{E,1})}_{\ell^e}$, and vice versa. The same argument shows that (3) is equivalent to $\widehat{(3)}$.

So it suffices to show that $\widehat{(O_{E,2})}_{\ell^e}$ can be used to solve $\widehat{(1)}$, and that $\widehat{(O_{E,3})}_{\ell^e}$ can be used to solve $\widehat{(2)}$. These arguments are identical to those we used to show that $(O_{E,2})_{\ell^e}$ can be used to solve (1) and that $(O_{E,3})_{\ell^e}$ can be used to solve (2), except with the roles of E and E' swapped. □

4 Equivalent Oracles for SIDH

In this section, we give a formulation of the SIDH cryptosystem and apply the results of Section 3 to show an equivalence between the problems we discussed and the computational and decisional SIDH problems.

Recall that, in our setup so far, $p = \ell_1^{e_1}\ell_2^{e_2} \pm 1$ is a prime, and E is a supersingular elliptic curve defined over $\mathbb{F}_{p^2}$. Recall also that $E[n] \cong (\mathbb{Z}/n\mathbb{Z}) \times (\mathbb{Z}/n\mathbb{Z})$, where $n = p \mp 1 = \ell_1^{e_1}\ell_2^{e_2}$, and that $E[\ell_1^{e_1}]$ and $E[\ell_2^{e_2}]$ are both defined over $\mathbb{F}_{p^2}$. We assume that there are fixed basis pairs $P_1, Q_1 \in E[\ell_1^{e_1}]$ and $P_2, Q_2 \in E[\ell_2^{e_2}]$ which are known to all parties. The SIDH protocol proceeds as follows.

1. Alice chooses a cyclic subgroup $A \subset E[\ell_1^{e_1}]$, computes $\phi_A : E \to E/A$, and sends her public key $(E/A, \phi_A(P_2), \phi_A(Q_2))$ to Bob.
2. Bob chooses a cyclic subgroup $B \subset E[\ell_2^{e_2}]$, computes $\phi_B : E \to E/B$, and sends his public key $(E/B, \phi_B(P_1), \phi_B(Q_1))$ to Alice.
3. Alice finds $\phi_B(A)$ using her knowledge of $P_1, Q_1, \phi_B(P_1), \phi_B(Q_1)$ and A.
4. Bob finds $\phi_A(B)$ using his knowledge of $P_2, Q_2, \phi_A(P_2), \phi_A(Q_2)$ and B.
5. They both compute the shared secret, namely, the common j-invariant of $(E/B)/\phi_B(A) \cong (E/A)/\phi_A(B)$.

We make a few observations. First, the only non-public piece of information needed for Alice's computations in steps 3 and 5 is Alice's secret A, so a natural candidate problem is to find A from the public information related to A. Similarly, we have a candidate problem of finding B from the public information related to B. By Lemma 4, finding A and B is equivalent to finding ϕ_A and ϕ_B, and we will prefer this formulation in terms of isogenies for consistency with our results in Section 3. We state these problems as follows.

Problem 2 (A-Isogeny Problem). Given E/A, a basis pair $P_2, Q_2 \in E[\ell_2^{e_2}]$, and a basis pair $\phi_A(P_2), \phi_A(Q_2) \in E/A[\ell_2^{e_2}]$, find $\phi_A\colon E \to E/A$.

Problem 3 (B-Isogeny Problem). Given E/B, a basis pair $P_1, Q_1 \in E[\ell_1^{e_1}]$, and a basis pair $\phi_B(P_1), \phi_B(Q_1) \in E/B[\ell_1^{e_1}]$, find $\phi_B\colon E \to E/B$.

It is not difficult to see that given the oracle $(O_{E,3})_{\ell_1^{e_1}}$ we may solve the A-Isogeny Problem, and that given an oracle to solve the A-Isogeny Problem we may solve Problem (1), so an A-Isogeny Problem oracle is equivalent to any of the six oracles with subscript $\ell_1^{e_1}$ by the results of Section 3. Similarly, a B-Isogeny Problem oracle is equivalent to any of the six oracles with subscript $\ell_2^{e_2}$.

Although the A and B Isogeny Problems are natural, this formulation is not the formulation usually given in the literature. The reason is that the typical description of the SIDH protocol requires that Alice and Bob choose cyclic subgroups of order $\ell_1^{e_1}$ and $\ell_2^{e_2}$ respectively, rather than simply any cyclic subgroup in their respective torsion groups. Consequently, the isogeny problems of interest are ones where one is also told that the degree of the isogeny is $\ell_1^{e_1}$ or $\ell_2^{e_2}$, which is a slight difference from our formulation and the formulations given in Section 3, where one is simply required to find isogenies of degree dividing ℓ^e.

However, the decision to formulate these problems as one where the isogenies have fixed degree ℓ^e is not universal. We have already mentioned the reduction of Thormarker [19] and Galbraith and Vercauteren [9], which considers isogenies of varying degrees. Another example is Petit's paper [14], which discusses attack strategies on torsion-point isogeny problems. Petit considers a more general class of problems where $\ell_1^{e_1}$ and $\ell_2^{e_2}$ are replaced by arbitrary coprime integers N_1 and N_2. The fact that Petit studies this more general class of problems allows him to consider a so-called "optimal degree variant" of the protocol, which he in turn shows is vulnerable to a certain family of attacks. We note that the theorems in Section 3 could not have been proven in this setting, since they relied crucially on the fact that there is a unique path on the isogeny graph corresponding to each isogeny kernel, which is false if the degree of the isogeny is not a prime power.

We also note that there is no harm to the security of the protocol if A and B are allowed to be arbitrary cyclic kernels, provided that Alice and Bob choose their generating point uniformly at random. Indeed, the event that a random point in $E[\ell^e] \cong (\mathbb{Z}/\ell^e\mathbb{Z}) \times (\mathbb{Z}/\ell^e\mathbb{Z})$ generates the kernel of an isogeny of small degree dividing ℓ^k is exponentially unlikely, since this outcome requires that both coefficients under such an isomorphism are divisible by ℓ^{e-k}, which happens with probability $\frac{1}{(\ell^{e-k})^2}$. This observation is analogous to how one does not typically exclude small private exponents in ordinary Diffie-Hellman, despite the fact that finding the exponent a given (g, g^a) is easy when a is sufficiently small, because the probability of Alice choosing a small private exponent a is low enough that the attacker gains no appreciable advantage if the protocol permits this possibility. Consequently, we also permit arbitrary cyclic kernels in our formulation of SIDH, as this relaxation allows us to apply the results of the previous section, and leads to a more natural and cohesive framework for studying the underlying hard problems.

Our next task is to give formulations of the decisional and computational SIDH problems. The computational SIDH problem is simply the core problem

required to break our formulation of the SIDH cryptosystem. With the notation as above, it is defined as follows.

Problem 4 (CSIDH Problem). Given

- the curves $E, E/A$ and E/B,
- basis pairs $P_1, Q_1 \in E[\ell_1^{e_1}]$ and $P_2, Q_2 \in E[\ell_2^{e_2}]$, and
- basis pairs $\phi_A(P_2), \phi_A(Q_2) \in (E/A)[\ell_2^{e_2}]$ and $\phi_B(P_1), \phi_B(Q_1) \in (E/B)[\ell_1^{e_1}]$,

find the isomorphism class of $E/\langle A, B\rangle$.

The CSIDH problem also has decisional variants. In ordinary Diffie-Hellman on a cyclic group G generated by g, the decisional Diffie-Hellman problem is to determine whether a triple $(x, y, z) \in G \times G \times G$ satisfies $\log_g(x) \log_g(y) = \log_g(z)$ modulo the order of G. To continue with the analogy, one could imagine being given supersingular curves (X, Y, Z), and being asked to determine whether the kernels of "the maps" $\psi_X \colon E \to X$, $\psi_Y \colon E \to Y$, and $\psi_Z \colon E \to Z$ satisfy $\langle \ker \psi_X, \ker \psi_Y \rangle = \ker \psi_Z$.

One issue with this formulation is that, even with the additional information specified about the images of these maps on torsion points, there is no guarantee that these maps are uniquely determined, as Lemma 5 is only strong enough to guarantee that there are at most a constant number of them. But interestingly, there is no difference between finding the "right" and the "wrong" map when it comes to computing the shared secret. The reason is that a party to the protocol only needs the other party's image curve E' and the two torsion images R and S to compute the correct shared secret, and so the shared secret cannot depend on the particular isogeny $\psi \colon E \to E'$ which takes P to R and Q to S. Thus, if there is an additional map $\psi'_X \colon E \to X$ of the right form (using the notation above), we will not have $\ker \psi'_X \ker \psi_Y = \ker \psi_Z$, but we will nevertheless have $E/(\ker \psi'_X \ker \psi_Y) \cong E/(\ker \psi_Z)$, and so an attacker with knowledge of ψ'_X could compute the secret curve by computing $\psi_Y(\ker \psi'_X)$ using the public torsion point images, and then by modding out $\psi_Y(\ker \psi'_X)$ from $E/(\ker \psi_Y)$. This suggests that to formulate the decisional SIDH problem, we need to instead interest ourselves in whether the tuple (X, Y, Z) and the associated torsion point information corresponds to a valid instance of the SIDH key exchange.

A second issue arises from the inherent asymmetry in the hard problems underlying the SIDH cryptosystem. In ordinary Diffie-Hellman, Alice and Bob's private exponent are both secured under the same discrete logarithm problem. But in SIDH, the problems securing Alice and Bob's private subgroups are in fact different, because Alice's isogenies have degree equal to a power of ℓ_1 and Bob's isogenies have degree equal to a power of ℓ_2. There does not seem to be any way to show that these two problems are equivalent; they are not equivalent generically, because as ℓ increases the number of cyclic subgroups of $E[\ell^e]$, and hence the number of possible private keys, must tend to 0 if one wishes to preserve the relationship $\log(\ell^e) \approx \log(n/\ell^e)$. Consequently, one cannot expect to prove a theorem that, say, the A-Isogeny Problem is equivalent to a "symmetric" formulation of the computational and decisional SIDH problems, since a

symmetric argument would say the same thing for the B-Isogeny Problem, and necessarily imply the equivalence of the A-Isogeny and B-Isogeny Problems. This observatoin motivates the asymmetry in our formulation of the decisional SIDH problems and the theorem that follows.

Problem 5 (A-DSIDH Problem). Suppose that E/B, and the image of the A-basis pair $\phi_B(P_1), \phi_B(Q_1) \in (E/B)[\ell_1^{e_1}]$ is known. Then given

- a curve X,
- a basis pair $R, S \in X[\ell_2^{e_2}]$,
- a curve Z, and
- an integer $0 \leq k \leq e_1$,

determine whether the tuple $(X, E/B, Z)$ is a *valid SIDH tuple*, in the sense that there is a map $\psi_X : E \to X$ of degree dividing $\ell_1^{e_1-k}$, which sends P_2 to R, Q_2 to S, and such that $Z \cong E/(\ker \psi_X B)$.

Problem 6 (B-DSIDH Problem). Suppose that E/A, and the image of the B-basis pair $\phi_A(P_2), \phi_A(Q_2) \in (E/A)[\ell_2^{e_2}]$ is known. Then given

- a curve Y,
- a basis pair $R, S \in Y[\ell_1^{e_1}]$,
- a curve Z, and
- an integer $0 \leq k \leq e_2$,

determine whether the tuple $(E/A, Y, Z)$ is a *valid SIDH tuple*, in the sense that there is a map $\psi_Y : E \to Y$ of degree dividing $\ell_2^{e_2-k}$, which sends P_1 to R, Q_1 to S, and such that $Z \cong E/(A \ker \psi_Y)$.

We now prove the main theorem of this section.

Theorem 4. *An oracle for the A-Isogeny problem is equivalent under randomized polynomial time reductions to an oracle which solves both the CSIDH Problem and the A-DSIDH Problem. Analogously, an oracle for the B-Isogeny problem is equivalent under randomized polynomial time reductions to an oracle which solves both the CSIDH Problem and the B-DSIDH Problem.*

Remark 3. The hypotheses of the A-DSIDH Problem specify that the information of Bob's public key is known. What this means in this context is that the equivalence between the A-Isogeny problem and the union of the CSIDH and A-DSIDH Problems is relative to a particular fixed public key for Bob that the A-SIDH oracle works with. The analogous fact is true for the other equivalence.

Proof. We start by assuming we have an oracle to solve the A-Isogeny Problem. We have already seen that given such an oracle one can solve the CSIDH problem, since one may find Alice's private subgroup A, and then proceed as Alice does to compute the shared secret. Hence, suppose we are given a curve X, a basis pair $R, S \in X[\ell_2^{e_2}]$, and a curve Z, and wish to determine whether $(X, E/B, Z)$ is a valid SIDH tuple. Since an oracle for the A-Isogeny Problem is equivalent to

the oracle $(O_{E,3})_{\ell_1^{e_1}}$, we may find the set of all isogenies of degree dividing $\ell_1^{e_1}$ from E to X with the correct torsion images. If this set is empty, we know that $(X, E/B, Z)$ is not a valid SIDH tuple. If it is non-empty, we may compute the kernels of these isogenies by Lemma 4, and proceed as Alice does to compute the potential shared secrets. If one (hence all) of these secret curves is isomorphic to Z, we accept, otherwise, we reject.

Next, we assume that we have an oracle which solves both the CSIDH Problem and the A-DSIDH Problem. We will show that given such an oracle we may solve the Key Validation Problem for $\ell_1^{e_1}$, which gives the desired conclusion by the equivalence between the A-Isogeny Problem and the problems in Section 3. We suppose we are given a proposed public key (X, R, S), where X purports to be a curve connected by an isogeny $\psi_X \colon E \to X$ of degree dividing $\ell_1^{e_1-k}$ such that $\psi_X(P_2) = R$ and $\psi_X(Q_2) = S$. We begin by calling the CSIDH oracle on the base curve, the base curve basis points, Bob's public key, and the proposed public key (X, R, S). One of two things may happen: the CSIDH oracle fails[5], in which case we know that (X, R, S) is invalid, or it returns some curve Z.

The curve Z could either be a correct shared secret (if the public key (X, R, S) was valid), or an arbitrary curve (if the public key (X, R, S) was invalid). It suffices to distinguish between these two cases. This is exactly the role of the A-DSIDH oracle, which we give the input $(X, E/B, Z)$ and the associated auxiliary information. If X was a valid public key, then the CSIDH oracle must have generated a valid Z, and the A-SIDH oracle will confirm this. Otherwise, the tuple must be invalid, which the A-SIDH oracle will also confirm.

This completes the proof of the first statement. The proof of the other statement involving B-type oracles proceeds in the same way. □

Remark 4. If one specializes the above proof to the case where Bob's public key is trivial (B is the identity subgroup), then there is no need for the CSIDH oracle, and the theorem proves that the A-Isogeny Problem is equivalent to its decisional variant, which in our case is the Key Validation Problem.

5 Conclusion

The torsion-point isogeny problems underlying the security of SIDH and several proposals for isogeny-based signatures [8, 12, 17, 21] have thus far undergone little

[5] Typically, one does not consider what happens when one gives an oracle invalid input. But one can easily consider what the possibilities are for a real algorithm: either the algorithm fails (produces an error, or runs longer than a worst-case bound on its running time), or gives an answer that does not solve the problem (because no answer solves the problem). Since the oracle formalism is really just a way of arguing about algorithms, we see no reason not to assume this behaviour here. Note that this sort of reasoning has appeared previously in the context of reducibility theorems in cryptology. For instance, in the reduction of the security of the Goldwasser-Micali encryption scheme [10] to the quadratic residuosity problem, one queries a cryptosystem-breaking oracle on potentially invalid public keys, which is the same situation as what is being described here.

study. One could argue this lack of study is indicative of their difficulty: the same complexity-theoretic obstructions which prevent problems from having efficient solutions can also preclude the existence of non-trivial algorithms, reductions, and security theorems. But it is nevertheless important, especially for researchers not in the isogeny-based cryptography community, that the relevant problems be formulated in a manner that emphasizes their connections and relationships. Such a formulation helps to guide both classical and quantum cryptanalysis, informs choices made when designing variants, and sheds light on which problems are likely to have tractable solutions.

Our formulations and reductions make significant progress towards these goals. In the SIDH context, our theorems provide a natural equivalence between the isogeny problems securing Alice and Bob's private keys and two problems securing the shared secret; they show that the Key Validation Problem, which has been a topic of interest in several papers [3, 7, 13], is likely to be intractable; and, by giving conditions on when torsion-point images suffice to determine isogenies (see Lemma 5), they give the first ever results concerning the uniqueness of SIDH public keys. In the more general context, our reductions show that improving the natural attack strategies, characterized by Problems (2) and $\widehat{(2)}$, is just as difficult as solving the isogeny problems themselves.

We believe that these results are important, not just because of their intrinsic value (which is itself significant), but also because they help theorists and practitioners alike understand the problem landscape. Consequently, we hope that our work will help guide and encourage further study in the field.

References

1. J.-F. Biasse, D. Jao, and A. Sankar. A quantum algorithm for computing isogenies between supersingular elliptic curves. In W. Meier and D. Mukhopadhyay, editors, *Progress in Cryptology – INDOCRYPT 2014: 15th International Conference on Cryptology in India, New Delhi, India, December 14-17, 2014, Proceedings*, pages 428–442, Cham, 2014. Springer International Publishing.
2. C. Costello, D. Jao, P. Longa, M. Naehrig, J. Renes, and D. Urbanik. Efficient Compression of SIDH Public Keys. In J.-S. Coron and J. B. Nielsen, editors, *Advances in Cryptology – EUROCRYPT 2017: 36th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Paris, France, April 30 – May 4, 2017, Proceedings, Part I*, pages 679–706, Cham, 2017. Springer International Publishing.
3. C. Costello, P. Longa, and M. Naehrig. Efficient algorithms for supersingular isogeny diffie-hellman. In M. Robshaw and J. Katz, editors, *Advances in Cryptology – CRYPTO 2016: 36th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2016, Proceedings, Part I*, pages 572–601, Berlin, Heidelberg, 2016. Springer Berlin Heidelberg.
4. C. Delfs and S. D. Galbraith. Computing isogenies between supersingular elliptic curves over $\mathbb{F}_p$. *Designs, Codes and Cryptography*, 78(2):425–440, Feb 2016.
5. D. Eisenbud and J. Harris. *3264 and all that—a second course in algebraic geometry.* Cambridge University Press, Cambridge, 2016.

6. S. D. Galbraith. Constructing isogenies between elliptic curves over finite fields. *LMS Journal of Computation and Mathematics*, 2:118–138, 1999.
7. S. D. Galbraith, C. Petit, B. Shani, and Y. B. Ti. On the security of supersingular isogeny cryptosystems. In J. H. Cheon and T. Takagi, editors, *Advances in Cryptology – ASIACRYPT 2016: 22nd International Conference on the Theory and Application of Cryptology and Information Security, Hanoi, Vietnam, December 4-8, 2016, Proceedings, Part I*, pages 63–91, Berlin, Heidelberg, 2016. Springer Berlin Heidelberg.
8. S. D. Galbraith, C. Petit, and J. Silva. Identification. protocols and signature schemes based on supersingular isogeny problems. In T. Takagi and T. Peyrin, editors, *Advances in Cryptology – ASIACRYPT 2017: 23rd International Conference on the Theory and Applications of Cryptology and Information Security, Hong Kong, China, December 3-7, 2017, Proceedings, Part I*, Cham, 2017. Springer International Publishing. To appear.
9. S. D. Galbraith and F. Vercauteren. Computational problems in supersingular elliptic curve isogenies. Cryptology ePrint Archive, Report 2017/774, 2017. https://eprint.iacr.org/2017/774.
10. S. Goldwasser and S. Micali. Probabilistic Encryption & How to Play Mental Poker Keeping Secret All Partial Information. In *Proceedings of the Fourteenth Annual ACM Symposium on Theory of Computing*, STOC '82, pages 365–377, New York, NY, USA, 1982. ACM.
11. D. Jao and L. De Feo. Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. In *Post-quantum cryptography*, volume 7071 of *Lecture Notes in Comput. Sci.*, pages 19–34. Springer, Heidelberg, 2011.
12. D. Jao and V. Soukharev. Isogeny-Based Quantum-Resistant Undeniable Signatures. In Mosca, Michele, editor, *Post-Quantum Cryptography: 6th International Workshop, PQCrypto 2014, Waterloo, ON, Canada, October 1-3, 2014. Proceedings*, pages 160–179, Cham, 2014. Springer International Publishing.
13. D. Kirkwood, B. C. Lackey, J. McVey, M. Motley, J. A. Solinas, and D. Tuller. Failure is not an option: Standardization issues for post-quantum key agreement. *Workshop on Cybersecurity in a Post-Quantum World*, 2015.
14. C. Petit. Faster algorithms for isogeny problems using torsion point images. In T. Takagi and T. Peyrin, editors, *Advances in Cryptology – ASIACRYPT 2017: 23rd International Conference on the Theory and Applications of Cryptology and Information Security, Hong Kong, China, December 3-7, 2017, Proceedings, Part II*, pages 330–353, Cham, 2017. Springer International Publishing.
15. I. R. Shafarevich. *Basic algebraic geometry. 1.* Springer, Heidelberg, third edition, 2013. Varieties in projective space.
16. J. Silverman. *The Arithmetic of Elliptic Curves*, volume 106 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 2 edition, 2009.
17. X. Sun, H. Tian, and Y. Wang. Toward quantum-resistant strong designated verifier signature from isogen ies. In *4th International Conference on Intelligent Networking and C ollaborative Systems (INCoS)*, pages 292–296, 2012.
18. E. Teske. The Pohlig-Hellman method generalized for group structure computati on. *Journal of Symbolic Computation*, 27(6):521–534, 1999.
19. E. Thormarker. Post-quantum cryptography: Supersingular isogeny diffie-hellman key exchange. Master's thesis, Stockholm University, 2017. http://kurser.math.su.se/pluginfile.php/16103/mod_folder/content/0/2017/2017_42_report.pdf.
20. J. Vélu. Isogénies entre courbes elliptiques. *C. R. Acad. Sci. Paris Sér. A-B*, 273:A238–A241, 1971.

21. Y. Yoo, R. Azarderakhsh, A. Jalali, D. Jao, and V. Soukharev. A post-quantum digital signature scheme based on supersingular isogenies. In A. Kiayias, editor, *Financial Cryptography and Data Security: 21st International Conference, FC 2017, Sliema, Malta, April 3–7, 2017, Revised Selected Papers*, Berlin, Heidelberg, 2017. Springer Berlin Heidelberg. To appear.

A Proof of Lemma 5

We will assume throughout that we work over an algebraic closure of our base field $\mathbb{F}_{p^2}$. We note that this has no effect on the content of Lemma 5, since if we manage to prove it over an algebraic closure of $\mathbb{F}_{p^2}$, then it will also be true over the base field.

To prove Lemma 5, we need some tools from intersection theory. Our main reference is Eisenbud and Harris [5]. We apply the generalized Bézout's theorem [5, Corollary 2.4], which reads:

> **Bézout's Theorem**: If $X_1, \cdots, X_k \subset \mathbb{P}^n$ are subvarieties of codimensions $c_1, \cdots, c_k$, with $\sum c_i \leq n$, and the X_i intersect generically transversely, then
>
> $$\deg(X_1 \cap \cdots \cap X_k) = \prod \deg(X_i).$$

We note that for 0-dimensional varieties (i.e., collections of points), the notion of degree simply corresponds to the number of points counted with multiplicity. Hence the theorem agrees with the classical Bézout's theorem in this case.

The term *generically transversely* needs explaining. Eisenbud and Harris first define the notion of *intersecting transversely at* p for two subvarieties A and B of X to mean A, B and X are all smooth at p and that $T_pA + T_pB = T_pX$, that is, the tangent spaces of the subvarieties sum to give the tangent space of the entire space. One then says that A and B intersect *generically transversely* if A and B meet transversely at some point of each irreducible component C of $A \cap B$. It can be shown that the set of points of $A \cap B$ where A and B are transverse is open, so because C is irreducible, the set of points in C at which $A \cap B$ meet transversely is either dense in C or empty. Hence if we find a single point in a component C where A and B meet transversely, that means that they meet transversely at a "general point" in C.

We also need the related notion of intersecting *dimensionally transversely*. If A and B are subvarieties of some variety X, then A intersects B *dimensionally transversely* if each irreducible component C of $A \cap B$ satisfies $\operatorname{codim} C = \operatorname{codim} A + \operatorname{codim} B$, where the codimension is taken with respect to X. According to [5, Prop. 1.28], subvarieties A and B are generically transverse if and only if they are dimensionally transverse and each irreducible component C of $A \cap B$ contains a point where X is smooth. In our case, we have $X = \mathbb{A}^3$, so the last condition is of no consequence.

We now apply this theory to our case of interest. We consider two affine subvarities of a 3-dimensional affine space $\mathbb{A}^3$ in the variables x, y and z given

by the equations

$$E : y^2 = x^3 + ax + b, \qquad \text{and} \qquad V : \eta(x,y)z = \psi(x,y).$$

The variety E is of course a Weierstrass-form elliptic curve (considered as a surface in $\mathbb{A}^3$). The functions $\psi(x,y)$ and $\eta(x,y)$ are the polynomials corresponding to the numerator and denominator of a component $\phi_z(x,y) = \frac{\psi(x,y)}{\eta(x,y)}$ of an isogeny from E (obtained via, say, Vélu's formulas). One can take $\psi(x,y)$ to be of degree d and $\eta(x,y)$ to be of degree strictly less than d. Hence, within $\mathbb{A}^3$, we have that E is a 2-dimensional surface of degree 3, and V is a 2-dimensional surface of degree d.

We need the fact that V is irreducible. It suffices to show that its defining polynomial is irreducible. Let $g(x,y,z) := \eta(x,y)z - \psi(x,y)$, and suppose that we have a non-trivial factorization of $g(x,y,z)$. It is clear that the factorization must be of the form

$$g(x,y,z) = (\alpha_1(x,y)z + \alpha_2(x,y))\beta(x,y).$$

Indeed, only one of the two factors may contain a term with a positive power of z, and that factor itself can be written in the form of a polynomial linear in z with the coefficients being polynomials in x and y. Expanding this expression for $g(x,y,z)$ and comparing with its definition in terms of $\eta(x,y)$ and $\psi(x,y)$, we see that both $\eta(x,y)$ and $\psi(x,y)$ contain a non-trivial factor $\beta(x,y)$. But then $\eta(x,y)$ and $\psi(x,y)$ cannot come from a degree d isogeny, since the quotient $\frac{\psi(x,y)}{\eta(x,y)}$ is a rational map of degree less than d after canceling the mutual factors of $\beta(x,y)$. Thus $g(x,y,z)$ is irreducible, and so is V.

We now show that E and V intersect generically transversely. Let $p = (p_1,p_2,p_3)$ be an arbitrary point in $E \cap V$. We compute the tangent spaces of E and V:

$$T_pE : (-3p_1^2 - a)(x - p_1) + 2p_2(y - p_2) = 0$$
$$T_pV : \frac{\partial g}{\partial x}(p_1,p_2,p_3)(x - p_1) + \frac{\partial g}{\partial y}(p_1,p_2,p_3)(y - p_2) + \eta(p_1,p_2)(z - p_3) = 0$$

Since $p \in E \cap V$, we know that (p_1,p_2) is a point on E. Since $\eta(x,y)$ is the denominator of an isogeny which is defined on all of E, we know that $\eta(p_1,p_2) \neq 0$. This tells us in particular that V is smooth at p (the tangent space is non-degenerate), and that $T_pE + T_pV = T_p\mathbb{A}^3$ at p, since the tangent space T_pE is parallel to the z-axis, and the tangent space T_pV is not. Since E is also smooth as the equation for E comes from an elliptic curve, we have thus shown that E and V intersect generically transversely.

We now return to the proof of Lemma 5.

Lemma 5. *Let $\phi, \phi' \colon E_1 \to E_2$ be isogenies of degree d from E_1 to E_2. If ϕ and ϕ' agree on N affine points, where $N > 3d^2$, then they are equal.*

Proof. We assume that E_1 and E_2 are given by the usual affine Weierstrass equations, so that the isogenies ϕ and ϕ' are represented by rational maps between

two affine algebraic curves. There is no loss of generality in doing so, since any isogeny between two elliptic curves can be represented in this way. Note that although we are working in affine space, we may identify the various subvarieties with their projective closures when applying Bézout's theorem.

Suppose the equations for E_1 and E_2 are

$$E_1 : y_1^2 = x_1^3 + a_1 x_1 + b_1, \qquad \text{and} \qquad E_2 : y_2^2 = x_2^3 + a_2 x_2 + b_2.$$

Then the maps ϕ and ϕ' each have two components, and so it suffices to show that the component functions are equal. Consider the x_2 component functions $\phi_{x_2}(x_1, y_1)$ and $\phi'_{x_2}(x_1, y_1)$. These component functions are rational functions of the form $\frac{\psi(x_1,y_1)}{\eta(x_1,y_1)}$ as described above. For ease of notation, write $\phi_{x_2}(x_1, y_1) = \frac{\psi(x_1,y_1)}{\eta(x_1,y_1)}$ and $\phi'_{x_2}(x_1, y_1) = \frac{\psi'(x_1,y_1)}{\eta'(x_1,y_1)}$, where we have suppressed the dependence on x_2.

As above, we define varieties obtained from the isogeny components:

$$V : \eta(x_1, y_1)x_2 = \psi(x_1, y_1), \qquad \text{and} \qquad V' : \eta'(x_1, y_1)x_2 = \psi'(x_1, y_1).$$

We have seen that $E_1 \cap V$ and $E_1 \cap V'$ are generically transverse intersections. Suppose also that V intersects V' generically transversely. Then by Bézout's Theorem, we see that $\deg(E_1 \cap V \cap V') \leq 3d^2$ (with equality over the projective closure). Since the intersections are generically transverse, by [5, Prop. 1.28] they are dimensionally transverse, and so all components in the intersection have codimension equal to $\operatorname{codim} E_1 + \operatorname{codim} V + \operatorname{codim} V'$ (see the paragraph under the statement of Corollary 2.4 (Bézout's Theorem) in [5]). Thus this intersection has dimension 0, and so it is a set of points of size at most $3d^2$. But we know that ϕ and ϕ' agree on at least $N > 3d^2$ points, and so $E_1 \cap V \cap V'$ must contain at least $N > 3d^2$ points, which is a contradiction. Therefore V and V' must not intersect generically transversely.

By [5, Prop. 1.28], since $\mathbb{A}^3$ is smooth everywhere, there must be an irreducible component C of $V \cap V'$ on which V and V' are not dimensionally transverse. That is, there is some irreducible component of C of $V \cap V'$ where the equality $\operatorname{codim} C = \operatorname{codim} V + \operatorname{codim} V'$ does not hold. Since V and V' are irreducible, we know that $2 \geq \dim C \geq \dim V + \dim V' - 3$, where 3 is the overall dimension of the affine space in which they are embedded [15, Theorem 1.24]. Hence the possibilities for $\operatorname{codim} C$ are 1 or 2, as $\dim V = \dim V' = 2$. Since the equality $\operatorname{codim} C = \operatorname{codim} V + \operatorname{codim} V'$ does not hold, we conclude that $\operatorname{codim} C = 1$, and hence C is 2-dimensional.

We thus know that V and V' intersect in some 2-dimensional component C. But V and V' are irreducible, so we conclude that $V = V'$. We thus see that $E_1 \cap V = E_1 \cap V'$, which is exactly the statement that $\phi_{x_2} = \phi'_{x_2}$ as functions on E_1. The argument that $\phi_{y_2} = \phi'_{y_2}$ on E is analogous. Since the component functions of ϕ and ϕ' agree on all points in E, the maps ϕ and ϕ' agree on all points in E, and so the isogenies are equal. □