SoK: The Problem Landscape of SIDH

David Urbanik
University of Waterloo
Waterloo, Ontario
dburbani@uwaterloo.ca

David Jao
University of Waterloo & evolutionQ, Inc.
Waterloo, Ontario
djao@uwaterloo.ca
david.jao@evolutionq.com

ABSTRACT

The Supersingular Isogeny Diffie-Hellman protocol (SIDH) has recently been the subject of increased attention in the cryptography community. Conjecturally quantum-resistant, SIDH has the feature that it shares the same data flow as ordinary Diffie-Hellman: two parties exchange a pair of public keys, each generated from a private key, and combine them to form a shared secret. To create a potentially quantum-resistant scheme, SIDH depends on a new family of computational assumptions involving isogenies between supersingular elliptic curves which replace both the discrete logarithm problem and the computational and decisional Diffie-Hellman problems. As in the case of ordinary Diffie-Hellman, one is interested in knowing if these problems are related. In fact, more is true: there is a rich network of reductions between the isogeny problems securing the private keys of the participants in the SIDH protocol, the computational and decisional SIDH problems, and the problem of validating SIDH public keys. In this article we explain these relationships, which do not appear elsewhere in the literature, in hopes of providing a clearer picture of the SIDH problem landscape to the cryptography community at large.

CCS CONCEPTS

• **Security and privacy** → *Public key encryption; Cryptanalysis and other attacks;*

KEYWORDS

isogeny-based cryptography, post-quantum cryptography, SIDH, torsion points, supersingular elliptic curves, equivalence theorems

ACM Reference Format:

David Urbanik and David Jao. 2018. SoK: The Problem Landscape of SIDH. In *APKC'18: 5th ACM ASIA Public-Key Cryptography Workshop, June 4, 2018, Incheon, Republic of Korea*. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3197507.3197516>

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 Supersingular Isogeny

Diffie-Hellman (SIDH), 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 $E/\langle A, B \rangle$, where $\langle A, B \rangle$ denotes the subgroup generated by the set $A \cup B$.

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 Bob’s secret subgroup B under Alice’s secret quotient map $\phi_A: E \rightarrow E/A$. Another such obstruction is that quotient curves in general are only well-defined up to isomorphism, 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 cryptosystem based upon the following hard problem (see Section 2 for notation). Let ℓ_1 and ℓ_2 be small distinct primes, and let e_1 and e_2 be exponents such that $\log(\ell_1^{e_1}) \approx \log(\ell_2^{e_2})$, and such that one of $p = \ell_1^{e_1} \ell_2^{e_2} \pm 1$ is 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 $\phi: E \rightarrow E'$ on $E[\ell_2^{e_2}]$, find ϕ . If we continue with the analogy above, ϕ is one of the “quotient” maps $\phi_A: E \rightarrow E/A$ or $\phi_B: E \rightarrow E/B$. 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

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

APKC'18, June 4, 2018, Incheon, Republic of Korea

© 2018 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-5756-2/18/06.

<https://doi.org/10.1145/3197507.3197516>

that this problem is the analogue of the discrete logarithm problem in classical Diffie-Hellman. The computational and decisional SIDH problems, which also result from this construction, are discussed further in Section 4.

Little is known about the security of this exact problem. More general isogeny problems have been studied in the literature [1, 5, 6], but the majority of such studies ignore the information provided by the values of the isogeny on $E[\ell_2^{e_2}]$. One exception is the recent work of Petit [15], whose work focuses on attacks on certain “overstretched” variants of the SIDH construction. A second exception is an argument made independently by Thormarker [19] and Galbraith and Vercauteren [9], which shows heuristically that one can reduce a computational isogeny problem to a decisional variant. And while there are a handful of articles discussing these torsion-point isogeny problems, analysis of the computational and decisional SIDH problems is even more scarce: aside from a few brief comments in the original paper [11] by Jao and De Feo, it is difficult to find any reference to them at all.

Our article gives a systematic discussion of these problems, and others, which are of interest in the cryptography community. We begin in Section 2 by reviewing standard facts about supersingular elliptic curves, isogeny computation, and basis finding which occur throughout the SIDH literature. 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. Finally, in Section 4, we give a formulation of the SIDH protocol which is more natural in two respects. Firstly, 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. Secondly, in the formulation we present, the problem of validating public keys, which has been studied by several authors [4, 7, 13] in hopes of obtaining a static-static or non-interactive key exchange (NIKE), is shown to be hard, in that an efficient solution to this problem suffices to break the cryptosystem. This observation explains why efforts to validate the public keys obtained from the SIDH construction have thus far been unsuccessful. We argue that our formulations 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 [17].

Let E and E' be two elliptic curves defined over a finite field in characteristic p . An isogeny $\phi: E \rightarrow E'$ is defined to be a non-constant rational map of curves which is also a group homomorphism between the elliptic curve groups of E and E' . The kernel of an isogeny is its kernel in the sense of group theory: $\ker \phi = \{P \in E : \phi(P) = O_{E'}\}$, where O_E denotes the identity element of E . One can show that all isogenies have finite kernels, and that all isogenies are surjective over an algebraic closure. Isogenies have a *degree*, which is their degree as a rational map; this number is always a non-negative integer. Isogenies are called *separable* if the size of their kernels is equal to their degree. Separable isogenies whose kernel consists of only the identity element are called isomorphisms, and such isogenies have inverse maps that

are also isogenies (and isomorphisms). Each isogeny $\phi: E \rightarrow E'$ has a dual isogeny $\hat{\phi}: E' \rightarrow E$ which satisfies $\phi \circ \hat{\phi} = [\deg \phi]$ and $\hat{\phi} \circ \phi = [\deg \phi]$, where the notation $[m]$ for any integer m denotes the scalar multiplication by m map, that is, $P \mapsto mP$. The kernel of $[m]: E \rightarrow E$ is denoted $E[m]$, and is the set of points $P \in E$ such that $mP = O_E$. The phrase “ m -torsion subgroup” which is widely used in the literature also refers to $E[m]$.

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 separable isogeny $\phi_H: E \rightarrow E/H$ 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 . Vélú [20] gave formulas using which one may compute the curves E/H from H and compute the rational maps corresponding to isogenies with kernel H .

In isogeny-based cryptography, one is typically only concerned with *supersingular* elliptic curves. These are curves which are distinguished by the fact that their endomorphism ring, the ring formed by the collection of all isogenies from a curve to itself together with the zero map, has maximum rank 4. Up to isomorphism, all such curves can be defined over a quadratic finite field $\mathbb{F}_{p^2}$. Curves which are not supersingular are called *ordinary*. Supersingular curves have a few special properties which make them suitable for quantum-resistant cryptography. Firstly, a sub-exponential quantum attack is known for isogeny-based cryptosystems derived from ordinary curves [2], but not for those derived from supersingular curves. Secondly, if one chooses p to be of the form $p = n \pm 1$, one can show that every isomorphism class of such curves contains a representative such that $E(\mathbb{F}_{p^2}) \cong (\mathbb{Z}/n\mathbb{Z}) \times (\mathbb{Z}/n\mathbb{Z})$. This fact allows one to construct curves with highly structured groups by choosing n (and hence p) appropriately, which is very useful in designing isogeny-based cryptosystems.

In this paper, we are interested in curves and 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 $\log(\ell_1^{e_1}) \approx \log(\ell_2^{e_2})$. Find a supersingular elliptic curve E defined over $\mathbb{F}_{p^2}$ such 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 , that is, isogenies obtained via quotients of the form $E/\langle P \rangle$, where $P \in E[\ell_i^{e_i}]$ is a point generating a cyclic subgroup $\langle P \rangle$ of order $\ell_i^{f_i}$. We note that, in general on any elliptic curve, $E[m] \cong (\mathbb{Z}/m\mathbb{Z}) \times (\mathbb{Z}/m\mathbb{Z})$ over an algebraic closure; therefore $E[\ell_i^{e_i}] \cong (\mathbb{Z}/\ell_i^{e_i} \mathbb{Z}) \times (\mathbb{Z}/\ell_i^{e_i} \mathbb{Z})$ is defined over $\mathbb{F}_{p^2}$, since the points Q in $E(\mathbb{F}_{p^2})$ satisfying $\ell_i^{e_i} Q = O_E$ already account for all of $(\mathbb{Z}/\ell_i^{e_i} \mathbb{Z}) \times (\mathbb{Z}/\ell_i^{e_i} \mathbb{Z})$.

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 would normally make their computation difficult; for example, writing down rational maps of degree $\ell_i^{f_i}$ directly is infeasible. However, the fact that $\ell_i^{f_i}$ is smooth 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

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

We then represent ϕ as a composition of isogenies $\phi = \phi_{f_i} \circ \cdots \circ \phi_1$, where $\phi_k: E_{k-1} \rightarrow E_k$, $E_0 = E$ and we set $\phi_0 = \text{id}_E$. Given E_{k-1} , we compute $E_k := E_{k-1} / \langle \phi_{k-1}([\ell_i^{f_i-k}]P) \rangle$ and ϕ_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. In this construction, the value of f_i is at most $\frac{1}{2} \log_{\ell_i}(p)$, and so the isogeny can be computed in time polynomial 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 usual 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 usual 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 2.1. *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 $\mathbb{Z}/\ell^e \mathbb{Z}$ -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 equal to $\frac{1}{2}$. 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 $([n/\ell^e]P_1, O_E)$. 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 $\langle P \rangle \cap \langle P' \rangle = \langle O_E \rangle$, 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. $\square$

REMARK 2.2. *In Costello et al. [3], optimized versions of 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 2.3. *Given a basis for $E[\ell^e]$, one also has a basis for $E[\ell^{e-k}]$ obtained via scalar multiplication by $[\ell^k]$.*

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 2.4 serves this purpose.

LEMMA 2.4. *Suppose that $\phi: E \rightarrow E'$ is an isogeny with degree dividing ℓ^e , with the same definitions as Lemma 2.1. Then there is a randomized polynomial-time algorithm to compute any of the 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 $\phi_1, \dots, \phi_s$ 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 $\phi'_1, \dots, \phi'_s$ 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 2.1 (or Remark 2.3) 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. $\square$

One last useful tool will be the Weil pairing. The Weil pairing of order m is a surjective bilinear map $e_m: E[m] \times E[m] \rightarrow \mu_m$, where μ_m is the m th roots of unity in the (algebraic closure of the) underlying field. A key fact about the Weil pairing is that it satisfies a compatibility condition with isogenies: if $\phi: E \rightarrow E'$ is an isogeny and $P, Q \in E[m]$, then $e_m(\phi(P), \phi(Q)) = e_m(P, Q)^{\deg \phi}$, where the first pairing is on the curve E' . Since these pairings are efficiently computable, they can allow us to recover information about the degree of ϕ from the pairings of basis points. For instance, we have the following Lemma.

LEMMA 2.5. *Suppose that there is an isogeny $\phi: E \rightarrow E'$ and the points $P, Q \in E[m]$ form a basis. Suppose additionally that m is smooth. Then if we know the image points $\phi(P)$ and $\phi(Q)$, we can recover the degree of ϕ modulo m .*

PROOF. Using Miller's algorithm for computing the Weil pairing [14], we compute $e_m(P, Q)$ and $e_m(\phi(P), \phi(Q)) = e_m(P, Q)^{\deg \phi}$. It is a standard fact that applying e_m to a basis produces a primitive

m th root of unity, and so $e_m(P, Q)^{\deg \phi}$ is determined exactly by the value of $\deg \phi$ modulo m . Because m is smooth, we can compute the discrete logarithm of $e_m(P, Q)^{\deg \phi}$ with respect to $e_m(P, Q)$ (using for example Pohlig-Hellman [16]) and hence recover the degree of ϕ modulo m . $\square$

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 by Costello et al. [4], κ is less than 5. For technical reasons, we also assume that $\ell_1^{e_1}, \ell_2^{e_2} > 4 \exp(\kappa)$. For realistic parameters, both $\ell_1^{e_1}$ and $\ell_2^{e_2}$ are exponentially sized, so this assumption presents no issue.

We define 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 $E[n/\ell^e]$ 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 $\eta: E \rightarrow E'$ 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 three 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 $\phi: E \rightarrow E'$ 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 $\psi: E \rightarrow X$, either
 - (i) return “Yes” if there exists an isogeny $\phi: E \rightarrow E'$ 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 $\psi': X \rightarrow E'$ 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 $\phi: E \rightarrow E'$ of degree dividing ℓ^e such that $\phi(P) = R$ and $\phi(Q) = S$.

For $i = 1, 2, 3$, 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 2.4) 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 isogeny¹ one could in principle 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 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 isogenies².

Our first task will be to show that the oracles $(O_{E,1})_{\ell^e}$ and $(O_{E,3})_{\ell^e}$ are equivalent; indeed, there is at most one such isogeny. This is the content of the next two lemmas.

LEMMA 3.1. *Let $\phi, \phi': E_1 \rightarrow E_2$ be isogenies of degree at most d from E_1 to E_2 . If ϕ and ϕ' agree on N points, where $N > 4d$, then they are equal.*

PROOF. To say that ϕ and ϕ' agree on N points is to say that the isogeny $\phi - \phi'$, where subtraction is defined pointwise, sends N points to the identity element of E_2 . Combining Corollary III.6.3 and Lemma V.1.2 in Silverman's book [17], one has the bound

$$|\deg(\phi - \phi') - \deg \phi - \deg \phi'| \leq 2\sqrt{\deg \phi \deg \phi'}.$$

Simplifying, one finds that $\deg(\phi - \phi') \leq 4d$. Since the size of the kernel of a non-zero isogeny is bounded by its degree, one has either $\phi - \phi' = 0$ or $N \leq 4d$. Since we are given that $N > 4d$, we conclude that $\phi - \phi' = 0$, and so $\phi = \phi'$. $\square$

LEMMA 3.2. *The set returned by $(O_{E,3})_{\ell^e}$ has exactly one isogeny. Thus, the oracles $(O_{E,1})_{\ell^e}$ and $(O_{E,3})_{\ell^e}$ are equivalent.*

PROOF. Let $\phi, \phi': E \rightarrow E'$ be isogenies returned by $(O_{E,3})_{\ell^e}$. Then ϕ and ϕ' agree on the entirety of $E[n/\ell^e]$, and so agree on $(n/\ell^e)^2$ points. By assumption, we have that $(n/\ell^e) > 4 \exp(\kappa)$, which gives

$$\log(n/\ell^e) > \log(4) + \kappa > \log(4) + \log(\ell^e) - \log(n/\ell^e).$$

Rearranging, one finds that $\log((n/\ell^e)^2) > \log(4\ell^e)$, or that $(n/\ell^e)^2 > 4\ell^e$. Applying Lemma 3.1 with $N = (n/\ell^e)^2$ we see that $\phi = \phi'$, which completes the proof. $\square$

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 problems (1) and (3) are equivalent — in fact, there is only ever one desired isogeny in cases of interest. To solve problem (2) given an oracle for (3), it simply suffices to check if the map ψ given as input to (2) extends to the map ϕ returned by the oracle for (3). This is the same as checking if the kernel of ψ is contained in the kernel of ϕ , which can be done efficiently using a combination of Lemmas 2.1 and 2.4 (simply evaluate ϕ on a generator for the kernel of ψ and check if the result is the identity element of the target curve).

¹For technical reasons even finding the “wrong” ϕ would typically suffice to break isogeny-based schemes despite not necessarily recovering the secret isogeny.

²There is a more sophisticated definition which considers curves and isogenies up to isomorphism, but we will not need it.

So 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 the required isogeny from E to E' 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 $\psi: E \rightarrow X$ do not extend to an isogeny $\phi: E \rightarrow E'$ 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 graph. This discussion completes the proof of Theorem 3.3.

THEOREM 3.3. *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.*

Each of the above problems has a corresponding “dual” problem which is obtained from Problems (1), (2) and (3) by effectively “reversing the arrows”, which can be done efficiently using Lemma 2.4. These problems are as follows.

- (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 $\phi': E' \rightarrow E$ of degree dividing ℓ^e such that $\phi'(R) = P$ and $\phi'(S) = Q$, 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 $\psi': E' \rightarrow X$, either
 - (i) return “Yes” if there exists an isogeny $\phi': E' \rightarrow E$ of degree dividing ℓ^e which factors through ψ' , and such that $\phi'(R) = P$ and $\phi'(S) = Q$, or
 - (ii) return “No” otherwise.
- (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 $\phi': E' \rightarrow E$ of degree dividing ℓ^e such that $\phi'(R) = P$ and $\phi'(S) = Q$.

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

The equivalences between problems $(\widehat{1})$, $(\widehat{2})$ and $(\widehat{3})$ follow from the same arguments used to prove the equivalence between (1) , (2) and (3) ; indeed, this process simply amounts to a change of notation. Furthermore, by Lemma 2.4, each problem is actually equivalent to its dual. However, the dual problems can still be useful. One reason is that the fastest known algorithm for finding isogenies between supersingular elliptic curves at present involves 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]. As an example of this, we consider the following decisional isogeny problem, which we call the *Key Validation Problem* in anticipation of its role in the next section, and which we show is equivalent to the problem $(\widehat{2})$.

PROBLEM 3.4 (KEY VALIDATION). *Given E' , a basis pair $P, Q \in E[n/\ell^e]$, and a basis pair $R, S \in E'[n/\ell^e]$, determine whether there*

exists an isogeny $\phi: E \rightarrow E'$ of degree dividing ℓ^e such that $\phi(P) = R$ and $\phi(Q) = S$.

THEOREM 3.5. *The Key Validation Problem is equivalent to Problem $(\widehat{2})$ under randomized polynomial-time reductions, and the reduction from Problem $(\widehat{2})$ to the Key Validation Problem succeeds with overwhelming probability.*

PROOF. We begin by showing that an oracle for the Key Validation Problem suffices to solve Problem $(\widehat{2})$. 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 $\psi': E' \rightarrow X$. Let $\ell^k = \deg \psi'$. We are interested in determining whether there is a map $\phi': E' \rightarrow E$ of degree dividing ℓ^e which factors through ψ' , for which it suffices to decide if there is a map $\psi: X \rightarrow E$ of degree dividing ℓ^{e-k} that sends the basis pair $\psi'(R), \psi'(S) \in X[n/\ell^e]$ to the basis pair $P, Q \in E[n/\ell^e]$ (note that degree is multiplicative over composition).

To decide whether the map ψ exists, we first use Lemma 2.5 on the basis pairs $\psi'(R), \psi'(S) \in X[n/\ell^e]$ and $P, Q \in E[n/\ell^e]$ to determine the value of $\deg \psi$ (if it exists) modulo n/ℓ^e . Since we know that $\deg \psi$ must divide ℓ^{e-k} , and hence must divide ℓ^e , we may check to see if the computed value for $\deg \psi$ agrees with the value of some divisor of ℓ^e modulo n/ℓ^e .

We claim that with overwhelming probability, all divisors of ℓ^e have distinct residues modulo n/ℓ^e . Note that ℓ is a unit modulo n/ℓ^e , and so if two divisors of ℓ^e are congruent modulo n/ℓ^e then the order of ℓ modulo n/ℓ^e must be less than e . Recall that $n/\ell^e = m^f$ where m is a prime relatively prime to ℓ , and so the unit group of $(\mathbb{Z}/(n/\ell^e)\mathbb{Z})$ is either cyclic of order $m^{f-1}(m-1)$ (if $m \neq 2$) or isomorphic to $(\mathbb{Z}/2\mathbb{Z}) \times (\mathbb{Z}/2^{f-2}\mathbb{Z})$ if $m = 2$. In either case, the number of elements with order less than e is an exponentially small fraction of the total number of units, so with overwhelming probability ℓ is not one of them.

But if the divisors of ℓ^e are distinct modulo n/ℓ^e , then the value we computed for $\deg \psi$ modulo n/ℓ^e tells us that either there is no such isogeny (if $\deg \psi$ is not congruent to a divisor of ℓ^{e-k} modulo n/ℓ^e) or we have isolated a single possible value for $\deg \psi$, and hence $\deg \widehat{\psi}$. Since determining whether ψ exists is equivalent to determining whether there exists a dual map $\widehat{\psi}: E \rightarrow X$ which maps the basis pair $P, Q \in E[n/\ell^e]$ to the basis pair $[\deg \psi]\psi'(R), [\deg \psi]\psi'(S) \in X[n/\ell^e]$, it then suffices to call the Key Validation Problem oracle on this input. The Key Validation oracle will tell us whether $\widehat{\psi}$ exists with degree dividing ℓ^e , but since we already know that any such $\widehat{\psi}$ must have degree dividing ℓ^{e-k} , we may simply return the evaluation of the Key Validation oracle. This completes one direction of the reduction.

For the other direction, it suffices to note that the Key Validation problem is exactly the same as Problem (2) when the map given as input is the identity map on E , and so the proof follows from the equivalence between Problems (2) and $(\widehat{2})$. $\square$

The above argument shows that the Key Validation problem is equivalent to the preceding six isogeny problems. In the next section, we will give a formulation of SIDH in which Problem (1) is the hard problem securing Alice and Bob's private keys, and the Key Validation problem is the problem that must be solved to validate them. If the keys are not validated, it is possible to perform

an active attack [7] on static-static or non-interactive variants of the scheme, and due to the absence of good validation techniques, SIDH has thus far been limited to ephemeral exchanges. The above theorem explains why: validating the keys seems to be just as hard as breaking the scheme itself.

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 $\phi_A: E \rightarrow E/A$, 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 $\phi_B: E \rightarrow E/B$, 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 2.4, 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 4.1 (A-ISOGENY PROBLEM). *Given the curve 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 $\phi_A: E \rightarrow E/A$.*

PROBLEM 4.2 (B-ISOGENY PROBLEM). *Given the curve 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 $\phi_B: E \rightarrow E/B$.*

It is not difficult to see these are instances of the same problems we studied in Section 3. But although these 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 [15], 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 .

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 4.3 (CSIDH PROBLEM). *Given*

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

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 $\psi_X: E \rightarrow X$, $\psi_Y: E \rightarrow Y$, and $\psi_Z: E \rightarrow Z$ satisfy $\langle \ker \psi_X, \ker \psi_Y \rangle = \ker \psi_Z$.

An issue with this formulation is that it doesn't respect 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. 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 observation motivates the asymmetry in our formulation of the decisional SIDH problems and the theorem that follows.

PROBLEM 4.4 (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 $R, S \in X[\ell_2^{e_2}]$,
- and a curve Z ,

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

PROBLEM 4.5 (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 $R, S \in Y[\ell_1^{e_1}]$,
- and a curve Z ,

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

We now prove the main theorem of this section.

THEOREM 4.6. *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 4.7. *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. To do this, we use the fact that the oracle for the A -Isogeny problem is the same as the oracle $(O_{E,1})_{\ell_1^{e_1}}$ to either find an isogeny $\psi_X : E \rightarrow X$ with the correct torsion images, or find that one doesn't exist. If one doesn't exist, then we know the tuple $(X, E/B, Z)$ is invalid. If one does exist, we may compute the resulting secret curve and check that it is isomorphic to Z . Since we know by Lemma 3.1 that there is only one possibility for ψ_X , we may return failure if the secret curve is not isomorphic to Z , and success otherwise.

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 $\psi_X : E \rightarrow X$ of degree dividing $\ell_1^{e_1}$ 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 fails³, 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. $\square$

5 CONCLUSION

The torsion-point isogeny problems underlying the security of SIDH and several proposals for isogeny-based signatures [8, 12, 18, 21] have thus far undergone little 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.

In this article, we provide formulations and reductions that make significant progress towards these goals. In particular, the formulation we have given provides an explanation of why the Key Validation Problem, which has been a topic of interest in several papers [4, 7, 13], is likely to be intractable, and suggests that one should regard the computational and decisional SIDH problems as being no easier than the hard isogeny problems which are usually studied.

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

³Typically, 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.

landscape. Consequently, we hope that this article will help guide and encourage further study in the field.

6 ACKNOWLEDGMENTS

We thank an anonymous referee for the proof of Lemma 3.1. This research was supported by NSERC, Public Works and Government Services Canada, and the Royal Bank of Canada.

REFERENCES

- [1] Jean-François Biasse, David Jao, and Anirudh Sankar. 2014. A Quantum Algorithm for Computing Isogenies between Supersingular Elliptic Curves. In *Progress in Cryptology — INDOCRYPT 2014: 15th International Conference on Cryptology in India, New Delhi, India, December 14-17, 2014, Proceedings*, Willi Meier and Debdeep Mukhopadhyay (Eds.). Springer International Publishing, Cham, 428–442. https://doi.org/10.1007/978-3-319-13039-2_25
- [2] Andrew Childs, David Jao, and Vladimir Soukharev. 2014. Constructing elliptic curve isogenies in quantum subexponential time. *Journal of Mathematical Cryptology* 8, 1 (2014), 1–29. <https://doi.org/10.1515/jmc-2012-0016>
- [3] Craig Costello, David Jao, Patrick Longa, Michael Naehrig, Joost Renes, and David Urbanik. 2017. Efficient Compression of SIDH Public Keys. In *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*, Jean-Sébastien Coron and Jesper Buus Nielsen (Eds.). Springer International Publishing, Cham, 679–706. https://doi.org/10.1007/978-3-319-56620-7_24
- [4] Craig Costello, Patrick Longa, and Michael Naehrig. 2016. Efficient Algorithms for Supersingular Isogeny Diffie-Hellman. In *Advances in Cryptology — CRYPTO 2016: 36th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2016, Proceedings, Part I*, Matthew Robshaw and Jonathan Katz (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 572–601. https://doi.org/10.1007/978-3-662-53018-4_21
- [5] Christina Delfs and Steven D. Galbraith. 2016. Computing isogenies between supersingular elliptic curves over $\mathbb{F}_p$. *Designs, Codes and Cryptography* 78, 2 (01 Feb 2016), 425–440. <https://doi.org/10.1007/s10623-014-0010-1>
- [6] Steven D. Galbraith. 1999. Constructing Isogenies between Elliptic Curves Over Finite Fields. *LMS Journal of Computation and Mathematics* 2 (1999), 118–138. <https://doi.org/10.1112/S1461157000000097>
- [7] Steven D. Galbraith, Christophe Petit, Barak Shani, and Yan Bo Ti. 2016. On the Security of Supersingular Isogeny Cryptosystems. In *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*, Jung Hee Cheon and Tsuyoshi Takagi (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 63–91. https://doi.org/10.1007/978-3-662-53887-6_3
- [8] Steven D. Galbraith, Christophe Petit, and Javier Silva. 2017. Identification Protocols and Signature Schemes Based on Supersingular Isogeny Problems. In *Advances in Cryptology — ASIACRYPT 2017*, Tsuyoshi Takagi and Thomas Peyrin (Eds.). Springer International Publishing, Cham, 3–33.
- [9] Steven D. Galbraith and Frederik Vercauteren. 2017. Computational problems in supersingular elliptic curve isogenies. Cryptology ePrint Archive, Report 2017/774. (2017). <https://eprint.iacr.org/2017/774>.
- [10] Shafi Goldwasser and Silvio Micali. 1984. Probabilistic encryption. *J. Comput. System Sci.* 28, 2 (1984), 270 – 299. [https://doi.org/10.1016/0022-0000\(84\)90070-9](https://doi.org/10.1016/0022-0000(84)90070-9)
- [11] David Jao and Luca De Feo. 2011. Towards Quantum-Resistant Cryptosystems from Supersingular Elliptic Curve Isogenies. In *Post-Quantum Cryptography*, Bo-Yin Yang (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 19–34.
- [12] David Jao and Vladimir Soukharev. 2014. Isogeny-Based Quantum-Resistant Undeniable Signatures. In *Post-Quantum Cryptography: 6th International Workshop, PQCrypto 2014, Waterloo, ON, Canada, October 1-3, 2014. Proceedings*, Mosca, Michele (Ed.). Springer International Publishing, Cham, 160–179.
- [13] Daniel Kirkwood, Bradley C. Lackey, John McVey, Mark Motley, Jerome A. Solinas, and David Tuller. April, 2015. Failure is not an Option: Standardization issues for post-quantum key agreement. Talk at NIST workshop on Cybersecurity in a Post-Quantum World: <http://www.nist.gov/itl/csd/ct/post-quantum-crypto-workshop-2015.cfm>. (April, 2015).
- [14] Victor S. Miller. 2004. The Weil Pairing, and Its Efficient Calculation. *Journal of Cryptology* 17, 4 (01 Sep 2004), 235–261. <https://doi.org/10.1007/s00145-004-0315-8>
- [15] Christophe Petit. 2017. Faster Algorithms for Isogeny Problems Using Torsion Point Images. In *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*, Tsuyoshi Takagi and Thomas Peyrin (Eds.). Springer International Publishing, Cham, 330–353. https://doi.org/10.1007/978-3-319-70697-9_12
- [16] Stephen Pohlig and Martin Hellman. 1978. An improved algorithm for computing logarithms over $GF(p)$ and its cryptographic significance (Corresp.). *IEEE Transactions on Information Theory* 24, 1 (January 1978), 106–110. <https://doi.org/10.1109/TIT.1978.1055817>
- [17] Joseph Silverman. 2009. *The Arithmetic of Elliptic Curves* (2nd ed.). Graduate Texts in Mathematics, Vol. 106. Springer-Verlag, New York.
- [18] Xi Sun, Haibo Tian, and Yumin Wang. 2012. Toward Quantum-Resistant Strong Designated Verifier Signature from Isogenies. In *2012 Fourth International Conference on Intelligent Networking and Collaborative Systems*. 292–296. <https://doi.org/10.1109/iNCOS.2012.70>
- [19] Erik Thormarker. 2017. *Post-Quantum Cryptography: Supersingular Isogeny Diffie-Hellman Key Exchange*. Master's thesis. Stockholm University. http://kurser.math.su.se/pluginfile.php/16103/mod_folder/content/0/2017/2017_42_report.pdf.
- [20] Jacques Vélou. 1971. Isogénies entre courbes elliptiques. *C. R. Acad. Sci. Paris Sér. A-B* 273 (1971), A238–A241.
- [21] Youngho Yoo, Reza Azarderakhsh, Amir Jalali, David Jao, and Vladimir Soukharev. 2017. A Post-quantum Digital Signature Scheme Based on Supersingular Isogenies. In *Financial Cryptography and Data Security*, Aggelos Kiayias (Ed.). Springer International Publishing, Cham, 163–181.