Some notes on isogenies for MasterMath course "Selected Areas in Cryptography" 2019

Tanja Lange, Eindhoven University of Technology
tanja@hyperelliptic.org

07 May 2019

1 Introduction

Isogeny-based cryptography uses maps between elliptic curves to build public-key cryptography. The first such system dates back to 1997, but the first publicly accessible proposals by Couveignes [Cou06] and Rostovtsev–Stolbunov [RS06] are from 2006, after they were used as an attack tool [GHS02, JMV05] in 2002 and 2005 and an isogeny-based hash function [CLG09] was published in 2006. These systems use isogenies between ordinary elliptic curves over finite fields to create a key-exchange system; the key-exchange system is denoted CRS in the following. Stolbunov [Sto11] also shows how to use this construction for building identification schemes. Shor's attack [Sho97], which breaks elliptic-curve cryptography based on the discrete-logarithm problem, does not affect these constructions, but in 2010 Childs, Jao, and Soukharev [CJS14] showed that CRS can be broken with a sub-exponential quantum attack due to Kuperberg [Kup05]. This means that parameters of the CRS scheme need to be scaled up asymptotically, making this already slow system even slower, but it does not mean that the system is fundamentally broken.

In 2011, Jao and De Feo [JF11] designed a different isogeny-based system that uses isogenies between supersingular curves over extension fields and does not have the same weakness as the above-described CRS key-exchange and which according to current knowledge offers exponential security, even against quantum attacks. A minor downside compared to CRS is the more complicated data flow and that the security assumption has changed from the pure isogeny-finding problem to one where additional information is available.

In the years since, research on isogeny-based systems focused mostly on this supersingular-isogeny Diffie-Hellman (SIDH) protocol. Many speedups were found and the security is better understood now; one important attack [GPST16] showed that reusing keys requires extra care such as using the Fujisaki-Okamoto transform [FO99], but for ephemeral use no issues are known. The shared benefit of CRS and SIDH systems is that they require very little bandwidth compared to other systems that are expected to resist attacks using quantum computers.

This text explains the basics of isogeny-based cryptography and the blueprint of all attacks; it also covers state-of-the-art proposals of isogeny-based cryptography by including pointers to a submission to the NIST project on post-quantum cryptography and more recent developments.

The next section gives some mathematics background to have proper definitions; it is possible to follow the other sections by knowing that an isogeny is a map between elliptic curves that satisfies some properties. For more background on elliptic curves and isogenies as well as proofs of the statements in the next section see, e.g., Silverman [Sil09].

2 Mathematics background: elliptic curves and isogenies

Let p be a prime larger than 3 and $n > 0$ an integer. Let $\mathbb{F}_{p^n}$ denote the finite field with p^n elements. An elliptic curve E over $\mathbb{F}_{p^n}$ can be written in short Weierstrass form

$$E : y^2 = x^3 + ax + b, \quad a, b \in \mathbb{F}_{p^n} \text{ and } 4a^3 + 27b^2 \neq 0.$$

Points on the elliptic curves over $\mathbb{F}_{p^n}$ are all pairs $(x, y) \in \mathbb{F}_{p^n}^2$ which satisfy the curve equation along with an extra point ∞. These points form a group with ∞ as neutral element. This same group operation works for points over any extension field of $\mathbb{F}_{p^n}$. We use $E/\mathbb{F}_{p^n}$ to denote that E is defined over $\mathbb{F}_{p^n}$; we use $E(\mathbb{F}_{p^n})$ to denote the group of points on E over $\mathbb{F}_{p^n}$, and we use E when speaking about properties that hold independently of the extension field. The number of points on an elliptic curve over $\mathbb{F}_{p^n}$ is roughly p^n: for $\#E(\mathbb{F}_{p^n}) = p^n + 1 - t$ the integer t lies in the interval $[-2p^{n/2}, 2p^{n/2}]$. An elliptic curve over $\mathbb{F}_{p^n}$ is *supersingular* if $t \equiv 0 \bmod p$, else the curve is *ordinary*.

The *order* of a point $P \in E(\mathbb{F}_{p^n})$ is the smallest integer $k > 0$ such that $kP = \infty$, where kP means the addition of k times P. Since there are only finitely many points in $E(\mathbb{F}_{p^n})$ the order of every point is finite; furthermore, the order of any point $P \in E(\mathbb{F}_{p^n})$ divides the group order $\#E(\mathbb{F}_{p^n})$. For every prime ℓ not dividing p there are either $1, \ell$, or ℓ^2 points P with $\ell P = \infty$. The first case corresponds to only $P = \infty$ satisfying $kP = \infty$. The second case corresponds to $\ell - 1$ points of order ℓ and ∞ satisfying the equations; these points form a cyclic group that behaves like $\mathbb{Z}/\ell$. The third case corresponds to $\ell^2 - 1$ points of order ℓ and ∞ satisfying the equations; these points form a product of two cyclic groups that behaves like $\mathbb{Z}/\ell \times \mathbb{Z}/\ell$. In particular, in this case all points of order ℓ are given as a linear combination $aP + bQ, a, b \in [0, \ell - 1]$, where P and Q are points of order ℓ and Q is not a multiple of P.

There are other representations of elliptic curves, most notably Montgomery form [Mon87]

$$BY^2 = X^3 + AX^2 + X, \quad A, B \in \mathbb{F}_{p^n} \text{ and } B(A^2 - 4) \neq 0$$

and Edwards form [BL07]

$$au^2 + v^2 = 1 + du^2v^2, \quad a, d \in \mathbb{F}_{p^n} \text{ and } ad(d - a) \neq 0$$

which have advantages for implementations and, sometimes, for exposition.

Two elliptic curves $E_1/\mathbb{F}_{p^n}$ and $E_2/\mathbb{F}_{p^n}$ are *isomorphic over* $\mathbb{F}_{p^n}$ if there exists a polynomial map over $\mathbb{F}_{p^n}$ that maps points (x, y) on E_1 to points on E_2 in a one-to-one way which is compatible with the group operation. If E_1 and E_2 are given in short Weierstrass form with $ab \neq 0$ any isomorphism has the form $(x, y) \mapsto (u^2x, u^3y)$ for some $u \neq 0$.

The systems we present later will consider elliptic curves up to isomorphism, i.e., work with *isomorphism classes*. They thus require a unique representative for each class. The typical choice of invariant for isomorphism classes is the j-invariant which for curves in Weierstrass form is $j = 1728 \cdot 4a^3/(4a^3 + 27b^2)$. For the example $(x, y) \mapsto (u^2x, u^3y)$ map given above the curve coefficients satisfy $a_2 = u^4a_1$ and $b_2 = u^6b_1$ which gives the same j for E_1 and E_2.

An *isogeny* between two elliptic curves $E_1/\mathbb{F}_{p^n}$ and $E_2/\mathbb{F}_{p^n}$ is a non-constant rational function that maps points from E_1 to points on E_2 and is compatible with the group law. In particular, it maps the neutral element ∞_1 on E_1 to the neutral element ∞_2 on E_2. Unlike isomorphisms, isogenies need not be 1-to-1 but can have several points map to ∞_2. For isogenies of interest in cryptography[1], the *degree* ℓ of an isogeny is the number of points on E_1, taken over any extension field of $\mathbb{F}_{p^n}$, mapping to ∞_2. Note, this does not mean that E_2 has fewer points than E_1 over $\mathbb{F}_{p^n}$, just that some points on E_2 are not in the image of points on E_1 over $\mathbb{F}_{p^n}$. In fact $E_1/\mathbb{F}_{p^n}$ and $E_2/\mathbb{F}_{p^n}$ are isogenous if and only if they have the same number of points over $\mathbb{F}_{p^n}$, i.e., $\#E_1(\mathbb{F}_{p^n}) = \#E_2(\mathbb{F}_{p^n})$. The set of curves that are isogenous to E is called the *isogeny class of* E. Note that if E is supersingular then all curves in its isogeny class are supersingular; similarly, ordinary curves are isogenous to ordinary curves.

The requirement that an isogeny be compatible with the group operation means that points of some order m coprime to ℓ are mapped to points of order m. An isogeny ϕ of degree ℓ has as kernel (i.e., points mapped to ∞_2) a cyclic subgroup of order ℓ, i.e., containing ℓ points, and each kernel uniquely defines an isogeny. These points may be defined over $\mathbb{F}_{p^n}$ or over some extension field. One way of computing an isogeny is to start with such a subgroup and then to use Vélu's formulas [Vél71], which give an explicit equation of the image curve and ϕ in terms of the coordinates of the points in the kernel. The computational complexity of these formulas grows linearly in ℓ and requires computations in the extension field over which the points in the kernel are defined.

For each isogeny $\phi : E_1 \to E_2$ there exists a *dual isogeny* $\hat{\phi} : E_2 \to E_1$ which has the same degree ℓ and for which it holds that the composition $\hat{\phi} \circ \phi = [\ell]_{E_1}$ is the multiplication-by-ℓ map on E_1 and likewise $\phi \circ \hat{\phi} = [\ell]_{E_2}$.

The ℓ-*isogeny graph* over $\mathbb{F}_{p^n}$ is an undirected graph that has as nodes the isomorphism classes of elliptic curves over $\mathbb{F}_{p^n}$ and two such classes are connected if there exists an ℓ isogeny from one curve in the class to one in the other class. (By combining the isogeny with an isomorphism each curve in the class can be reached; note that different schemes use different extensions of $\mathbb{F}_{p^n}$ for defining the isomorphisms). The graph is undirected because for each isogeny $\phi : E_1 \to E_2$, the dual isogeny $\hat{\phi}$ provides a map back. In an

[1] technically, for separable isogenies

ℓ-isogeny graph each node has zero, one, two, or $\ell + 1$ edges. This number depends on the structure of the set of points of order ℓ on the elliptic curve, which is a subgroup of $\mathbb{Z}/\ell \times \mathbb{Z}/\ell$. Each order-$\ell$ subgroup defines an isogeny over some extension field and zero, one, two, or $\ell + 1$ curves that are ℓ-isogenous to this curve may be defined over the same field.

The set of all curves over a field $\mathbb{F}_{p^n}$ splits into multiple disjoint components consisting of curves ℓ-isogenous to one another. In the following, we only consider curves having the same number of points, i.e. curves that are isogenous under an isogeny of some degree. Note that this does not imply that all of these curves are connected under an ℓ isogeny for some fixed ℓ.

Components of the ℓ-isogeny graph consisting of ordinary curves form (parts of) *volcanoes* (see Figure 1 for an illustration). A full volcano consists of a circular rim (the top of the volcano) and each of the nodes of the rim additionally has $\ell - 1$ edges pointing downwards, each node on lower levels has ℓ edges pointing downwards. Cryptosystems use only the top of the volcano and thus pick isogeny classes where that part is large. Each node in the top has exactly two neighbors and repeated application of ℓ isogenies makes multiple steps in the same direction. For different isogeny degrees ℓ_i the rim might split into multiple rims or multiple rims get combined into a larger one. The maximal size of the rim is the class number of the endomorphism ring; for details on what this means see [Sut12]. An important property for the CRS system is that the action of isogenies on ordinary elliptic curves is commutative, i.e., the order of applying ℓ_1 and ℓ_2 isogenies does not change the isomorphism class of the image curve.

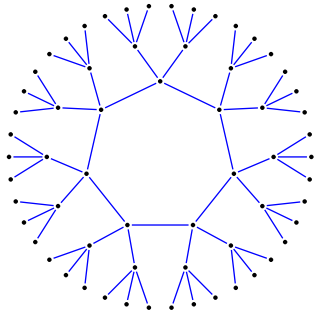

Figure 1: Isogeny graph for 2-isogenies forming a volcano with rim and two more levels. Image credit: Lorenz Panny.

All isomorphism classes of supersingular elliptic curves over extensions of $\mathbb{F}_p$ have j-invariant defined over $\mathbb{F}_{p^2}$, so considering isomorphism classes of supersingular elliptic curves over $\mathbb{F}_{p^2}$ covers all classes of supersingular curves, where the isomorphisms are taken over extension fields of $\mathbb{F}_{p^2}$. This graph consists of roughly $p/12$ isomorphism classes and is (almost) $\ell + 1$ regular and Ramanujan. This means that the graph is very well connected and any node in the graph can be reached in few steps from any other node (rapid mixing). It also means that there is no sense of direction – in computing multiple ℓ isogenies one can avoid going back but each step offers a choice of ℓ other edges forward.

3 The CRS system

The CRS system [Cou06, RS06] resembles closely the regular Diffie-Hellman key exchange. It uses the overlay of multiple isogeny graphs for the same set of ordinary curves isogenous

to one another (when allowing the isomorphism to be over a larger field). The system parameters fix a finite field $\mathbb{F}_p$, a starting curve E of known order N, and a set of primes $\ell_i > 2$, $1 \leq i \leq r$, so that for each of these primes an ℓ_i isogeny can be computed with not too much effort and in a unique manner. For a properly-chosen curve the number of curves isogenous to it are on the order of $\sqrt{p}$.

For each of the primes ℓ_i Alice picks an exponent a_i and computes the curve E_A which is $\prod \ell_i^{a_i}$ isogenous to E. This is typically computed as a sequence of a_1 isogenies of degree ℓ_1, a_2 isogenies of degree ℓ_2, etc. Computing another ℓ_1 isogeny on $\phi_{\ell_1}(E)$ works similar to the computation on E and continues on the rim of the ℓ_1 isogeny volcano. Alice's public key is E_A and her secret key is the exponent vector $(a_1, a_2, \ldots, a_r)$.

Similarly, Bob picks $(b_1, b_2, \ldots, b_r)$ and computes and publishes E_B.

To check that Bob's key is valid Alice verifies that the number of points on E_B is N. Starting from Bob's curve E_B, Alice computes the curve E_{BA} which is $\prod \ell_i^{a_i}$ isogenous to E_B; similarly Bob computes the curve E_{AB} which is $\prod \ell_i^{b_i}$ isogenous to E_A. The resulting curves are isomorphic because it does not matter whether Alice's or Bob's isogenies are applied first, thus $j(E_{BA}) = j(E_{AB})$ and Alice and Bob use a key derived from $j(E_{BA})$ as their shared secret.

The cost of computing the isogenies depends on the degrees ℓ_i and, when using Vélu's formulas, the extension field of $\mathbb{F}_p$ over which points of order ℓ_i are defined. If all a_i are chosen from an interval of size m, at most m^r different curves can be reached. In order to be able to reach all $\approx \sqrt{p}$ isogenous curves efficiently it is important to overlay multiple isogeny graphs, i.e. choose r sufficiently large, else a lot more steps are needed. This means that the system cannot stick to primes ℓ_i for which points of order ℓ_i are defined over $\mathbb{F}_p$ but needs to move to more primes and extension fields. The timing in [Sto10] makes the system too slow to be practical, and the security of the chosen parameters was later revised downwards.

4 The SIDH system

The SIDH system [JF11, DFJP14] uses the isogeny graph of supersingular curves over $\mathbb{F}_{p^2}$. Choosing $p = 2^{n_A}3^{n_B} - 1$ and $E : Y^2 = X^3 + X$ as starting curve means that the curves have $(p+1)^2 = 2^{2n_A}3^{2n_B}$ points and that all $2^{2n_A} - 1$ points of order 2^{n_A} and all $3^{2n_B} - 1$ of order 3^{n_B} are defined over $\mathbb{F}_{p^2}$. Computing degree 2^{n_A} or 3^{n_B} isogenies is very efficient using Vélu's formulas as the kernels are subgroups defined over $\mathbb{F}_{p^2}$. The same holds for all curves in the graph. The parameters are chosen so that $2^{n_A} \approx 3^{n_B} \approx \sqrt{p}$.

In SIDH, Alice works with 2^{n_A} isogenies and Bob works with 3^{n_B} isogenies. Note that unlike in CRS the degrees are known publicly, but there are about $\sqrt{p}$ choices left; to see this we need to look at the number of choices for a kernel of these maps. The points of order 2^{n_A} form a space of dimension two, we can find points P_A and Q_A of order 2^{n_A} with $Q_A \neq \langle P_A \rangle$. Then the subgroups of order 2^{n_A} are given by $\langle P_A \rangle, \langle P_A + Q_A \rangle, \langle P_A + 2Q_A \rangle, \ldots, \langle P_A + (2^{n_A} - 1)Q_A \rangle$; each of these subgroups determines a unique 2^{n_A} isogeny.

The system parameters for SIDH are p and E as above, a basis P_A, Q_A of the points of

order 2^{n_A} on E, and similarly a basis P_B, Q_B of the points of order 3^{n_B} on E.

Alice picks a secret $0 \leq a < 2^{n_A}$, computes $T_A = P_A + aQ_A$ and the isogeny ϕ_A with kernel $\langle T_A \rangle$, landing at E_A, which is (part of) her public key. Similarly, Bob picks a secret $0 \leq b < 3^{n_B}$, computes $T_B = P_B + bQ_B$ and the isogeny ϕ_B with kernel $\langle T_B \rangle$ to E_B.

One difficulty in defining this system is that Alice cannot compute an isogeny ϕ'_A on E_B that matches ϕ_A translated by ϕ_B without having more information on ϕ_B, but ϕ_B is Bob's secret, so cannot be given to Alice. The way out found by Jao and De Feo is to include additional points in the public keys of Alice and Bob, namely Bob also computes and publishes $\phi_B(P_A)$ and $\phi_B(Q_A)$. With that information, Alice can compute $T'_A = \phi_B(P_A) + a\phi_B(Q_A)$ and the isogeny ϕ'_A with kernel $\langle T'_A \rangle$, landing at E_{BA}. The use of the image points means that E_{BA} and Bob's E_{AB} are isomorphic. SIDH uses the j-invariant of the resulting curve to compute a shared key.

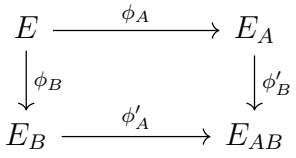

Figure 2: SIDH key exchange for Alice and Bob.

In summary, Alice's secret key is a and her public key is $(E_A, \phi_A(P_B), \phi_A(Q_B))$. Bob's secret key is b and his public key is $(E_B, \phi_B(P_A), \phi_B(Q_A))$.

One problem, recognized by [GPST16], is that Alice cannot validate the key she receives from Bob. An evil Bob can perform a reaction attack on Alice to learn a by sending her malformed public keys, e.g. sending $(E_B, \phi_B(P_A), \phi_B(Q_A) + 2^{n_A-1}\phi_B(P_A))$ has Alice compute the same j-invariant as Bob if and only if a is even (because then $T'_A = \phi_B(P_A) + a\phi_B(Q_A) = \phi_B(P_A) + a(\phi_B(Q_A) + 2^{n_A-1}\phi_B(P_A))$ as $2^{n_A}\phi_B(P_A) = \infty_B$), learning one bit of Alice's secret. See Section 6 for how to deal with this issue.

Computing ϕ_A in one go would be very inefficient because the cost grows linearly with the degree and ϕ_A has degree 2^{n_A}. Hence, SIDH decomposes this 2^{n_A} isogeny into n_A computations of 2 isogenies. These start with a 2 isogeny ϕ_2 with kernel $\langle 2^{n_A-1}T_A \rangle$ and compute the image $\phi_2(T_A)$, which is a point of order 2^{n_A-1} on a 2-isogenous curve, so that computing ϕ_A is the same as computing ϕ_2 followed by a 2^{n_A-1} isogeny with kernel $\langle \phi_2(T_A) \rangle$. Likewise, the images of P_B and Q_B are pushed through the 2 isogenies. A fast sequence of steps is proposed in [DFJP14] to reduce the cost of this computation.

5 Security analysis

For CRS, an attacker is confronted with the problem of finding E_{AB} given E, E_A, and E_B. Analogous to the situation of discrete logarithms and Diffie-Hellman, there is no known attack faster than computing Alice's or Bob's isogeny to get E_{AB}.

Obviously, the number of keys has to be large enough to protect against brute force searches or their more intelligent meet-in-the-middle variants. These search for the key in time square-root of the search space, so in roughly $\sqrt[4]{p}$ if the key space is large enough so that all isogenous curves can be reached.

A quantum attacker can use the attack by Childs, Jao, and Soukharev [CJS14] which

requires a subexponential number of calls to an oracle which compute isogenies in quantum superposition. A very recent analysis [BLMP19] for the case of CSIDH (see below) shows that the cost of each such oracle call contributes significantly to the cost of the attack so that for low security levels the main concern is to defend against the above non-quantum attacks.

For SIDH an attacker is confronted with the problem of finding E_{AB} given E, E_A, $\phi_A(P_B)$, $\phi_A(Q_B)$, E_B, $\phi_B(P_A)$, and $\phi_B(Q_A)$. The additional points have raised some concern but no attack for balanced $2^{n_A} \approx 3^{n_B}$, such as the parameters proposed in SIKE (see below), is known.

Alice's and Bob's key spaces have size $\approx \sqrt{p}$, so meet-in-the-middle attacks run in time roughly $\sqrt[4]{p}$. A recent analysis [ACC+18] has shown that the cost of these attacks is typically underestimated, meaning that smaller parameters would offer sufficient security.

On the quantum side, similarly [JS19] showed that the attack costs of the so-called claw finding attack with $\sqrt[6]{p}$ underestimate security when taking into account the full cost of quantum computation (RAM model) so that choosing parameters to protect against non-quantum attacks suffices to remain secure against quantum attackers.

6 Complete instantiations of isogeny-based encryption

The only isogeny-based submission to the NIST competition is SIKE [JAC+] by Jao, Azarderakhsh, Campagna, Costello, De Feo, Hess, Jalali, Koziel, LaMacchia, Longa, Naehrig, Renes, Soukharev, and Urbanik. SIKE is based on SIDH and uses a transformation to achieve CCA security and prove that ciphertexts are properly generated. Of the candidates that advanced to the second round, SIKE has the smallest public keys and the smallest combined size of message and public key.

SIDH for use in TLS 1.3 has been tested by Cloudflare https://blog.cloudflare.com/sidh-go/ in a hybrid construct with regular elliptic-curve cryptography.

After the submission deadline for the NIST competition, De Feo, Kieffer, and Smith published [DKS18] a new way to construct curves for CRS so that the resulting system becomes less slow. A much bigger speedup was obtained by Castryck, Lange, Martindale, Panny, and Renes in their recent CSIDH [CLM+18] system. CSIDH works with isogenies between $\mathbb{F}_p$-isomorphism classes of supersingular curves defined over $\mathbb{F}_p$. It uses the fact that isogenies when restricted to this subset of supersingular curves commute, so that a protocol with the same data flow as CRS becomes feasible, i.e. no additional points are needed. The number of points on such a curve is $p+1$ making it very easy to control which isogenies are defined over $\mathbb{F}_p$: CSIDH puts $p = 4\ell_1 \cdot \ell_2 \cdots \ell_r - 1$ for some r and distinct odd primes ℓ_i. This means that the resulting curve has points of order ℓ_i and Vélu's formulas can be used. A downside is that subexponential attacks similar to those on CRS apply but the upshot is that speeds of CSIDH are comparable to those of SIDH while the keys are smaller, at least at lower security levels, because no additional points are needed.

References

[ACC+18] Gora Adj, Daniel Cervantes-Vázquez, Jesús-Javier Chi-Domínguez, Alfred Menezes, and Francisco Rodríguez-Henríquez. On the cost of computing isogenies between supersingular elliptic curves. In *SAC*, volume 11349 of *Lecture Notes in Computer Science*, pages 322–343. Springer, 2018.

[BL07] Daniel J. Bernstein and Tanja Lange. Faster addition and doubling on elliptic curves. In *ASIACRYPT*, volume 4833 of *Lecture Notes in Computer Science*, pages 29–50. Springer, 2007.

[BLMP19] Daniel J. Bernstein, Tanja Lange, Chloe Martindale, and Lorenz Panny. Quantum circuits for the CSIDH: optimizing quantum evaluation of isogenies. In *EUROCRYPT*, Lecture Notes in Computer Science, page to appear. Springer, 2019. https://ia.cr/2018/1059.

[CJS14] Andrew M. Childs, David Jao, and Vladimir Soukharev. Constructing elliptic curve isogenies in quantum subexponential time. *J. Mathematical Cryptology*, 8(1):1–29, 2014. https://arxiv.org/abs/1012.4019.

[CLG09] Denis Xavier Charles, Kristin E. Lauter, and Eyal Z. Goren. Cryptographic hash functions from expander graphs. *J. Cryptology*, 22(1):93–113, 2009. https://ia.cr/2006/021.

[CLM+18] Wouter Castryck, Tanja Lange, Chloe Martindale, Lorenz Panny, and Joost Renes. CSIDH: An efficient post-quantum commutative group action. In *ASIACRYPT*, volume 11274 of *Lecture Notes in Computer Science*, pages 395–427. Springer, 2018. https://ia.cr/2018/383.

[Cou06] Jean-Marc Couveignes. Hard Homogeneous Spaces, 2006. IACR Cryptology ePrint Archive 2006/291. https://ia.cr/2006/291.

[DFJP14] Luca De Feo, David Jao, and Jérôme Plût. Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. *Journal of Mathematical Cryptology*, 8(3):209–247, 2014. IACR Cryptology ePrint Archive 2011/506. https://ia.cr/2011/506.

[DKS18] Luca De Feo, Jean Kieffer, and Benjamin Smith. Towards practical key exchange from ordinary isogeny graphs. In *ASIACRYPT*, volume 11274 of *Lecture Notes in Computer Science*, pages 365–394. Springer, 2018. https://ia.cr/2018/485.

[FO99] Eiichiro Fujisaki and Tatsuaki Okamoto. Secure integration of asymmetric and symmetric encryption schemes. In *CRYPTO*, volume 1666 of *Lecture Notes in Computer Science*, pages 537–554. Springer, 1999.

[GHS02] Steven D. Galbraith, Florian Hess, and Nigel P. Smart. Extending the GHS Weil descent attack. In *EUROCRYPT*, volume 2332 of *Lecture Notes in Computer Science*, pages 29–44. Springer, 2002.

[GPST16] Steven D. Galbraith, Christophe Petit, Barak Shani, and Yan Bo Ti. On the security of supersingular isogeny cryptosystems. In *ASIACRYPT (1)*, volume 10031 of *Lecture Notes in Computer Science*, pages 63–91. Springer, 2016. IACR Cryptology ePrint Archive 2016/859. https://ia.cr/2016/859.

[JAC+] David Jao, Reza Azarderakhsh, Matthew Campagna, Craig Costello, Luca De Feo, Basil Hess, Amir Jalali, Brian Koziel, Brian LaMacchia, Patrick Longa, Michael Naehrig, Joost Renes, Vladimir Soukharev, and David Urbanik. SIKE. Submission to [Nat16]. http://sike.org.

[JF11] David Jao and Luca De Feo. Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. In *PQCrypto*, volume 7071 of *Lecture Notes in Computer Science*, pages 19–34. Springer, 2011. https://eprint.iacr.org/2011/506/20110918:024142.

[JMV05] David Jao, Stephen D. Miller, and Ramarathnam Venkatesan. Do all elliptic curves of the same order have the same difficulty of discrete log? In *ASIACRYPT*, volume 3788 of *Lecture Notes in Computer Science*, pages 21–40. Springer, 2005. https://ia.cr/2004/312.

[JS19] Samuel Jaques and John M. Schanck. Quantum cryptanalysis in the RAM model: Claw-finding attacks on SIKE. Cryptology ePrint Archive, Report 2019/103, 2019. https://ia.cr/2019/103.

[Kup05] Greg Kuperberg. A subexponential-time quantum algorithm for the dihedral hidden subgroup problem. *SIAM J. Comput.*, 35(1):170–188, 2005. https://arxiv.org/abs/quant-ph/0302112.

[Mon87] Peter L. Montgomery. Speeding the Pollard and elliptic curve methods of factorization. *Mathematics of Computation*, 48(177):243–264, 1987.

[Nat16] National Institute of Standards and Technology. Post-quantum cryptography standardization, December 2016. https://csrc.nist.gov/Projects/Post-Quantum-Cryptography/Post-Quantum-Cryptography-Standardization.

[RS06] Alexander Rostovtsev and Anton Stolbunov. Public-key cryptosystem based on isogenies, 2006. IACR Cryptology ePrint Archive 2006/145. https://ia.cr/2006/145.

[Sho97] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.*, 26(5):1484–1509, 1997. https://arxiv.org/abs/quant-ph/9508027.

[Sil09] Joseph H. Silverman. *The arithmetic of elliptic curves*. Number 106 in Graduate Texts in Mathematics. Springer, 2nd edition, 2009.

[Sto10] Anton Stolbunov. Constructing public-key cryptographic schemes based on class group action on a set of isogenous elliptic curves. *Adv. in Math. of Comm.*, 4(2):215–235, 2010.

[Sto11] Anton Stolbunov. *Cryptographic Schemes Based on Isogenies*. PhD thesis, Norwegian University of Science and Technology, 2011.

[Sut12] Andrew V. Sutherland. Isogeny volcanoes. In *ANTS X*, volume 1 of *The Open Book Series*, pages 507–530. Mathematical Sciences Publishers, 2012. https://arxiv.org/abs/1208.5370.

[Vél71] Jacques Vélu. Isognies entre courbes elliptiques. *Comptes Rendus de l'Académie des Sciences de Paris*, 273:238–241, 1971.

Acknowledgement

Please let me know of any typos you find. I would like to acknowledge feedback by Chris Swart and Matt Campagna on earlier versions of this file.