CSIDH: An Efficient Post-Quantum Commutative Group Action

Wouter Castryck[1], Tanja Lange[2], Chloe Martindale[2],
Lorenz Panny[2], and Joost Renes[3]

wouter.castryck@esat.kuleuven.be, tanja@hyperelliptic.org,
chloemartindale@gmail.com, lorenz@yx7.cc, j.renes@cs.ru.nl

[1] Department of Mathematics, KU Leuven, Belgium
[2] Department of Mathematics and Computer Science,
Technische Universiteit Eindhoven, The Netherlands
[3] Digital Security Group, Radboud Universiteit, The Netherlands

Abstract We propose an efficient commutative group action suitable for non-interactive key exchange in a post-quantum setting. Our construction follows the layout of the Couveignes–Rostovtsev–Stolbunov cryptosystem, but we apply it to supersingular elliptic curves defined over a large prime field $\mathbb{F}_p$, rather than to ordinary elliptic curves. The Diffie–Hellman scheme resulting from the group action allows for public-key validation at very little cost, runs reasonably fast in practice, and has public keys of only 64 bytes at the AES-128 security level, matching NIST's post-quantum security category I.

Keywords: Post-quantum cryptography, isogeny-based cryptography, class group action, non-interactive key exchange, key confirmation.

1 Introduction

During the past five to ten years, elliptic-curve cryptography (ECC) has taken over public-key cryptography on the internet and in security applications. Many protocols such as Signal or TLS 1.3 rely on the small key sizes and efficient computations to achieve forward secrecy, often meaning that keys are used only once. However, it is also important to notice that security does not break down if keys are reused. Indeed, some implementations of TLS, such as Microsoft's

[*] Author list in alphabetical order; see https://www.ams.org/profession/leaders/culture/CultureStatement04.pdf. This work was supported in part by the Commission of the European Communities through the Horizon 2020 program under project number 643161 (ECRYPT-NET), 645622 (PQCRYPTO), 645421 (ECRYPT-CSA), and CHIST-ERA USEIT (NWO project 651.002.004); the Technology Foundation STW (project 13499 – TYPHOON) from the Dutch government; and the Research Foundation - Flanders (FWO) through the WOG Coding Theory and Cryptography. The first listed author is affiliated on a free basis with the Department of Electrical Engineering, imec-COSIC, KU Leuven and with the Department of Mathematics, Ghent University. Date of this document: TODO.

SChannel, reuse keys for some fixed amount of time rather than for one connection [2]. Google's QUIC relies on servers keeping their keys fixed for a while to achieve quick session resumption. Finding a system that permits reusing keys for multiple key exchanges without compromising security in a post-quantum setting, while still offering decent performance, is considered an open problem. Our paper presents a solution to this problem.

Isogeny-based cryptography is a relatively new kind of elliptic-curve cryptography, whose security relies on (various disguises of) the problem of finding an explicit isogeny between two given isogenous elliptic curves over a finite field $\mathbb{F}_q$. One of the main selling points is that quantum computers do not seem to make the isogeny-finding problem substantially easier. This contrasts with regular elliptic-curve cryptography, which is based on the discrete logarithm problem in a group and therefore falls prey to a polynomial-time quantum algorithm designed by Shor in 1994 [41].

The first proposal of an isogeny-based cryptosystem was made by Couveignes in 1997 [12]. It described a non-interactive key exchange protocol where the public key space equals the set of $\mathbb{F}_q$-isomorphism classes of ordinary elliptic curves over $\mathbb{F}_q$ whose endomorphism ring is a given order $\mathcal{O}$ in an imaginary quadratic field and whose trace of Frobenius has prescribed sign. It is well-known that the ideal class group $\mathrm{cl}(\mathcal{O})$ acts freely and transitively on this set through the application of isogenies. Couveignes' central observation was that the commutativity of $\mathrm{cl}(\mathcal{O})$ naturally allows for a key exchange protocol in the style of Diffie and Hellman [17]. His work was only circulated privately and thus not picked up by the community, and the corresponding paper [12] was never formally published. The method was eventually independently rediscovered by Rostovtsev and Stolbunov in 2004 (in Stolbunov's master's thesis [44] and published on ePrint as [38] in 2006). In 2010, Childs, Jao and Soukharev [7] showed that breaking the Couveignes–Rostovtsev–Stolbunov scheme amounts to solving an instance of the abelian hidden shift problem, for which quantum algorithms with a time complexity of $L_q[1/2]$ are known to exist; see [30, 36]. While this may be tolerable (e. g., classical subexponential factorization methods have not ended the widespread use of RSA), a much bigger concern is that the scheme is unacceptably slow: despite recent clever speed-ups due to De Feo, Kieffer and Smith [15, 28], several minutes are needed for a single key exchange at a presumed classical security level of 128 bits. Nevertheless, in view of its conceptual simplicity, compactness and flexibility, it seems a shame to discard the Couveignes–Rostovtsev–Stolbunov scheme.

The attack due to Childs–Jao–Soukharev strongly relies on the fact that $\mathrm{cl}(\mathcal{O})$ is commutative, hence indirectly on the fact that $\mathcal{O}$ is commutative. This led Jao and De Feo [26] to consider the use of supersingular elliptic curves, whose full ring of endomorphisms is an order in a quaternion algebra; in particular it is non-commutative. Their resulting (interactive) key agreement scheme, which nowadays goes under the name "Supersingular Isogeny Diffie–Hellman" (SIDH), has attracted almost the entire focus of isogeny-based cryptography over the past six years. The current state-of-the-art implementation is SIKE [25], which was recently submitted to the NIST competition on post-quantum cryptography [32].

It should be stressed that SIDH is *not* the Couveignes–Rostovtsev–Stolbunov scheme in which one substitutes supersingular elliptic curves for ordinary elliptic curves; in fact SIDH is much more reminiscent of a cryptographic hash function from 2006 due to Charles, Goren and Lauter [6].

In this paper we show that such a straightforward adaptation is also possible, provided that one restricts to supersingular elliptic curves defined over a prime field $\mathbb{F}_p$. Instead of the full ring of endomorphisms, which is non-commutative, one should consider the subring of $\mathbb{F}_p$-rational endomorphisms, which is again an order $\mathcal{O}$ in an imaginary quadratic field. As before $\mathrm{cl}(\mathcal{O})$ acts via isogenies on the set of $\mathbb{F}_p$-isomorphism classes of elliptic curves whose $\mathbb{F}_p$-rational endomorphism ring equals $\mathcal{O}$, the only difference being that we now have a single orbit, rather than two orbits corresponding to opposite signs of Frobenius; see e.g. [51, Theorem 4.5], with further details to be found in [3, 16] and in Section 3 of this paper. Starting from these observations, the desired adaptation of the Couveignes–Rostovtsev–Stolbunov scheme almost unrolls itself; the details can be found in Section 4. We call the resulting scheme CSIDH, where the C stands for "commutative".[4]

While this fails to address Jao and De Feo's initial motivation for using supersingular elliptic curves, which was to avoid the $L_q[1/2]$ quantum attack due to Childs–Jao–Soukharev, we show that CSIDH remediates the main issue of the Couveignes–Rostovtsev–Stolbunov scheme, namely its inefficiency. Indeed, in Section 8 we will report on a proof-of-concept implementation which carries out a non-interactive key exchange at a presumed classical security level of 128 bits and a post-quantum security level of 64 bits in less than 100 milliseconds, while using key sizes of only 64 bytes. This is over 2000 times faster[5] than the current state-of-the-art instantiation of the Couveignes–Rostovtsev–Stolbunov scheme by De Feo, Kieffer and Smith [15, 28], which itself presents many new ideas and speedups to even achieve that speed.

For comparison, we remark that SIDH, which is already being acclaimed for its short key lengths in the post-quantum community, uses public keys of over 300 bytes. More precisely SIKE's version `p503` uses uncompressed keys of 378 bytes long [25] for achieving CCA security. The optimized SIKE implementation is about ten times faster than our proof-of-concept C implementation, but even at 100 ms, CSIDH is practical.

Another major advantage of CSIDH is that we can efficiently validate public keys, making it possible to reuse the key without the need for transformations to confirm the key was honestly generated.

Finally, CSIDH relies purely on the isogeny finding problem; no extra points are sent that could potentially harm security, as argued in [34].

[4] Since this work was started while being very close to a well-known large body of salt water, we pronounce CSIDH as [ˈsiːˌsaɪd] rather than spelling out all the letters.

[5] This speed-up is explained in part by comparing our own C implementation to the `sage` implementation of De Feo–Kieffer–Smith.

To summarize, CSIDH is a new cryptographic primitive that can serve as a drop-in replacement for ECC while maintaining security against quantum computers. It provides a *non-interactive* key exchange with full public key validation. The speed is practical while the public-key size is the smallest for key exchange or KEM in the portfolio of post-quantum cryptography. This makes CSIDH particularly attractive in the common scenario of prioritizing bandwidth over computational effort. In addition, CSIDH is compatible with 0-RTT protocols such as QUIC.

Why supersingular? To understand where the main speed-up comes from, it suffices to record that De Feo–Kieffer–Smith had the idea of choosing a field of characteristic p, where p is congruent to -1 modulo all small odd primes ℓ up to a given bound. They then look for an ordinary elliptic curve $E/\mathbb{F}_p$ such that $\#E(\mathbb{F}_p)$ is congruent to 0 modulo as many of these ℓ's as possible, i.e., such that points of order ℓ exist over $\mathbb{F}_p$. These properties ensure that $\ell\mathcal{O}$ decomposes as a product of two prime ideals $\mathfrak{l} = (\ell, \pi - 1)$ and $\bar{\mathfrak{l}} = (\ell, \pi + 1)$, where π denotes the Frobenius endomorphism. For such primes the action of the corresponding ideal classes $[\mathfrak{l}]$ and $[\bar{\mathfrak{l}}] = [\mathfrak{l}]^{-1}$ can be computed efficiently through an application of Vélu-type formulae to E (resp. its quadratic twist E^t), the reason being that only $\mathbb{F}_p$-rational points are involved. If this works for enough primes ℓ, we can expect that a generic element of $\mathrm{cl}(\mathcal{O})$ can be written as a product of small integral powers of such $[\mathfrak{l}]$, so that the class group action can be computed efficiently. However, finding an ordinary elliptic curve $E/\mathbb{F}_p$ such that $\#E(\mathbb{F}_p)$ is congruent to 0 modulo many small primes ℓ is hard, and the main focus of De Feo–Kieffer–Smith is on speeding up this search. In the end it is only practical to enforce this for 7 primes, thus they cannot take full advantage of the idea.

However, in the supersingular case the property $\#E(\mathbb{F}_p) = p + 1$ implies that $\#E(\mathbb{F}_p)$ is congruent to 0 modulo *all* primes $\ell \mid p + 1$ that we started from in building p! Concretely, our proof-of-concept implementation uses 74 small odd primes, corresponding to prime ideals $\mathfrak{l}_1, \mathfrak{l}_2, \ldots, \mathfrak{l}_{74}$ for which we heuristically expect that almost all elements of our 256-bit size class group can be written as $[\mathfrak{l}_1]^{e_1}[\mathfrak{l}_2]^{e_2} \cdots [\mathfrak{l}_{74}]^{e_{74}}$, where the exponents e_i are taken from the range $\{-5, \ldots, 5\}$; indeed, one verifies that $\log{(2 \cdot 5 + 1)^{74}} \approx 255.9979$. The action of such an element can be computed as the composition of at most $5 \cdot 74 = 370$ easy isogeny evaluations. This should be compared to using 7 small primes, where the same approach would require exponents in a range of length about $2^{256/7} \approx 2^{36}$, in view of which De Feo–Kieffer–Smith also resort to other primes with less beneficial properties, requiring to work in extensions of $\mathbb{F}_p$.

The use of supersingular elliptic curves over $\mathbb{F}_p$ has various other advantages. For instance, their trace of Frobenius t is 0, so that the absolute value of the discriminant $|t^2 - 4p| = 4p$ is as large as possible. As a consequence, generically the size of the class group $\mathrm{cl}(\mathcal{O})$ is close to its maximal possible value for a fixed choice of p. Conversely, this implies that for a fixed security level we can make a close to minimal choice for p, which directly affects the key size. Note that this contrasts with the CM construction from [4], which could in principle be used to construct ordinary elliptic curves having many points of small order, but whose endomorphism rings have very small class groups, ruling them out for the Couveignes–Rostovtsev–Stolbunov key exchange.

To explain why key validation works, note that we work over $\mathbb{F}_p$ with $p \equiv 3 \pmod 8$ and start from the curve $E_0\colon y^2 = x^3 + x$ with endomorphism ring $\mathcal{O} = \mathbb{Z}[\pi]$. As it turns out, all Montgomery curves $E_A\colon y^2 = x^3 + Ax^2 + x$ over $\mathbb{F}_p$ that are supersingular appear in the $\mathrm{cl}(\mathcal{O})$-orbit of E_0. Moreover their $\mathbb{F}_p$-isomorphism class is uniquely determined by A. So all one needs to do upon receiving a candidate public key $y^2 = x^3 + Ax^2 + x$ is check for supersingularity, which is an easy task; see Section 5. The combination of large size of $\mathrm{cl}(\mathcal{O})$ and representation by a single $\mathbb{F}_p$-element A explains the small key size of 64 bytes.

1.1 One-way group actions

Although non-interactive key exchange is the main application of our primitive, it is actually more general: It is (conjecturally) an instance of Couveignes' *hard homogeneous spaces* [12], ultimately nothing but a finite commutative group action for which some operations are easy to compute while others are hard. We summarize Couveignes' definition:

Definition 1. *A* hard homogeneous space *consists of a finite commutative group G acting freely and transitively on some set X.*
The following tasks are required to be easy (e. g., polynomial-time):

- *Compute the group operations in G.*
- *Sample randomly from G with (close to) uniform distribution.*
- *Decide validity and equality of a representation of elements of X.*
- *Compute the action of a group element $g \in G$ on some $x \in X$.*

The following problems are required to be hard (e. g., not polynomial-time):

- *Given $x, x' \in X$, find $g \in G$ such that $g * x = x'$.*
- *Given $x, x', y \in X$ such that $x' = g * x$, find $y' = g * y$.*

Any such primitive immediately implies a natural Diffie–Hellman protocol: Alice and Bob's private keys are random elements a, b of G, their public keys are $a * x_0$ resp. $b * x_0$, where $x_0 \in X$ is a public fixed element, and the shared secret is $b * (a * x_0) = a * (b * x_0)$. The private keys are protected by the difficulty of the first hard problem above, while the shared secret is protected by the second problem. Note that traditional Diffie–Hellman on a cyclic group is an instance of this, where X is the set of generators of the underlying group and G is the multiplicative group $(\mathbb{Z}/\#X)^*$ acting by exponentiation.

1.2 Notation and terminology

We stress that throughout this paper, we consider two elliptic curves defined over the same field identical whenever they are isomorphic *over that field.* Note that we do *not* identify curves that are only isomorphic over some extension field, as opposed to what is done in SIDH, for instance. In the same vein, for an elliptic curve E defined over a finite field $\mathbb{F}_p$, we let $\mathrm{End}_p(E)$ be the subring of the endomorphism ring $\mathrm{End}(E)$ consisting of endomorphisms defined over $\mathbb{F}_p$.[6] This subring is always isomorphic to an order in an imaginary quadratic number field. Conversely, for a given order $\mathcal{O}$ in an imaginary quadratic field, we let $\mathcal{E}\ell\ell_p(\mathcal{O})$ denote the set of elliptic curves E defined over $\mathbb{F}_p$ with $\mathrm{End}_p(E) \cong \mathcal{O}$. The isomorphism $\mathrm{End}_p(E) \xrightarrow{\sim} \mathcal{O}$ is not canonical; indeed, the p-power Frobenius endomorphism π of trace t can be mapped to either of the (up to) two elements $\alpha \in \mathcal{O}$ satisfying $\alpha^2 - t\alpha + p = 0$. In what follows we assume this choice to be consistent between the Frobenius elements of given trace within the set $\mathcal{E}\ell\ell_p(\mathcal{O})$; in particular, this implies that $\varphi \circ \beta = \beta \circ \varphi$ for all $\mathbb{F}_p$-isogenies φ between two curves in $\mathcal{E}\ell\ell_p(\mathcal{O})$ and all $\beta \in \mathcal{O}$.

Ideals are always assumed to be non-zero.

The notation “log” refers to the base-2 logarithm.

Acknowledgements. This project started during a research retreat on post-quantum cryptography, organized by the European PQCRYPTO and ECRYPT-CSA projects in Tenerife from 29 January until 1 February 2018. We would like to thank Jeffrey Burdges, whose quest for a flexible post-quantum key exchange protocol made us look for speed-ups of the Couveignes–Rostovtsev–Stolbunov scheme. We are grateful to Luca De Feo, Jean Kieffer and Ben Smith for sharing a draft of their paper in preparation, and to Jeroen Demeyer and Fré Vercauteren for helpful feedback.

2 Isogeny graphs

The security of an isogeny-based cryptosystem is usually assessed in terms of rapid mixing properties of the underlying isogeny graph. Just as in the original Couveignes–Rostovtsev–Stolbunov cryptosystem, in our case this graph is obtained by gluing together several large cycles, one for each prime ℓ under consideration. The precise terminology is the *Schreier graph* associated with our class group action and the chosen generators. We refer to the lecture notes of De Feo [14, §14.1] for more background and for a discussion of its rapid mixing properties [27]. One point of view on this is that switching between different cycles allows one to quickly process large group elements and thereby replaces the square-and-multiply algorithm in exponentiation-based cryptosystems (such as classical Diffie–Hellman).

The goal of this section is to analyze the structure of the individual cycles.

Definition 2. *For a field k and a prime $\ell \nmid \mathrm{char}\, k$, the* k-rational ℓ-isogeny graph *$G_{k,\ell}$ is defined as having all the elliptic curves defined over k as its vertices, and there is a directed edge (E_1, E_2) for each k-rational ℓ-isogeny from E_1 to E_2.*[7]

[6] This constraint only makes a difference for supersingular curves: in the ordinary case, all endomorphisms are defined over the base field.

[7] Due to our convention of identifying k-isomorphic curves, this means the nodes are actually k-isomorphism classes, and two isogenies give rise to the same edge if they are equal up to post-composition with an isomorphism of E_2 defined over k.

Remark 3. A priori $G_{k,\ell}$ is a directed graph, but given two elliptic curves E_1 and E_2 whose j-invariants are not in $\{0, 1728\}$, there are exactly as many edges (E_2, E_1) as (E_1, E_2), obtained by taking dual isogenies. Annoyingly, the nodes with j-invariants 0 and 1728 are more complicated, since these are exactly the curves with extra automorphisms: an elliptic curve E has fewer incoming than outgoing edges if and only if either $j(E) = 0$ and $\sqrt{-3} \in k$, or if $j(E) = 1728$ and $\sqrt{-1} \in k$. Throughout this paper, we will assume for simplicity that $\sqrt{-3}, \sqrt{-1} \notin k$, so that neither of these automorphisms is defined over k and we may view $G_{k,\ell}$ as an undirected graph. In the case of a finite prime field $k = \mathbb{F}_p$, it suffices to restrict to $p \equiv 11 \pmod{12}$, which will be satisfied in the class of instantiations we suggest.

If $k = \mathbb{F}_q$ is a finite field such that ℓ does not divide the field characteristic, then $G_{k,\ell}$ is a finite graph that is the union of ordinary connected components and supersingular connected components. The ordinary components were studied in Kohel's PhD thesis [29]. Due to their regular structure, these components later became known as *isogeny volcanoes.*

Over non-prime fields, the supersingular components on the other hand may bear no similarity at all to the volcanoes known from the ordinary case. Traditionally, following Pizer [35], one studies the unique supersingular component of $G_{k,\ell}$ where $k = \overline{\mathbb{F}}_q$, which turns out to be a finite $(\ell + 1)$-regular Ramanujan graph and forms the basis for the SIDH protocol.

However, Delfs and Galbraith [16] showed that if $k = \mathbb{F}_p$ is a finite prime field, then all connected components are volcanoes, even in the supersingular case. We present a special case of a unified statement, restricting our attention to the cases in which $G_{\mathbb{F}_p,\ell}$ is a cycle. Recall that $\mathrm{End}_p(E)$ is an order $\mathcal{O}$ in the imaginary quadratic field

$$\mathrm{End}_p(E) \otimes_{\mathbb{Z}} \mathbb{Q} \cong \mathbb{Q}(\sqrt{t^2 - 4p}) = K,$$

where $|t| \leq 2\sqrt{p}$ denotes the (absolute value of the) trace of the Frobenius endomorphism, and that two curves are isogenous over $\mathbb{F}_p$ if and only if their t is equal [47, Theorem 1].

Theorem 4 (Kohel, Delfs–Galbraith). *Let $p \geq 5$ be a prime number and let V be a connected component of $G_{\mathbb{F}_p,\ell}$. Assume that $p \equiv 11 \pmod{12}$ or that V contains no curve with j-invariant 0 or 1728 (so that it can be viewed as an undirected graph). Let t be the trace of Frobenius common to all vertices in V, and let K be as above. Assume that $\ell \nmid t^2 - 4p$.*

Then all elliptic curves in V have the same endomorphism ring $\mathcal{O} \subset K$, and $\mathcal{O}$ is locally maximal at ℓ. Moreover if $t^2 - 4p$ is a (non-zero) square modulo ℓ, then V is a cycle whose length equals the order of $[\mathfrak{l}]$ in $\mathrm{cl}(\mathcal{O})$, where $\mathfrak{l}$ is a prime ideal dividing $\ell\mathcal{O}$. If not, then V consists of a single vertex.

Proof. In the case of an ordinary component this is just a special case of [46, Theorem 7]. In the case of a supersingular component this follows from the proof of [16, Theorem 2.7]. (In both cases, we could alternatively (re)prove this theorem by proving that an ℓ-isogeny can only change the conductor of the endomorphism ring of an elliptic curve locally at ℓ and applying Theorem 7.) □

In the ordinary case a curve and its quadratic twist can never appear in the same component because they have a different trace of Frobenius. This is the main difference with the supersingular case, where this possibility is not excluded. To avoid confusion, we recall that by the quadratic twist of a given elliptic curve $E\colon y^2 = f(x)$ over $\mathbb{F}_p$ we mean the curve $E^t\colon dy^2 = f(x)$, where $d \in \mathbb{F}_p^*$ is any non-square. If (and only if) $p \equiv 3 \pmod 4$ and $j(E) = 1728$ then this may deviate from what some readers are used to, because in this case E^t and E are $\mathbb{F}_p$-isomorphic. Note that such a curve is necessarily supersingular.

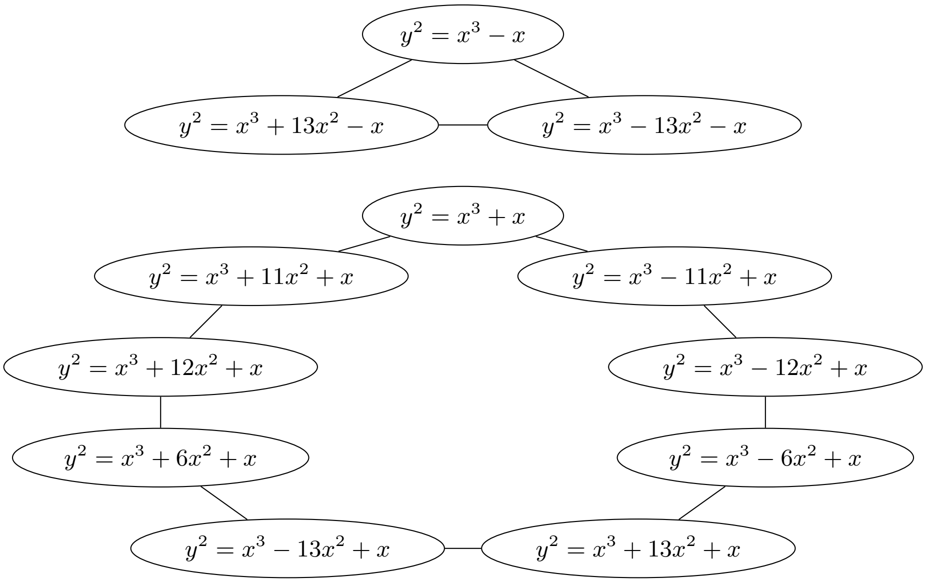

Figure 1. The two supersingular components of $G_{\mathbb{F}_{83},3}$. The curves in the top component have $\mathbb{F}_p$-rational endomorphism ring $\mathbb{Z}[(1+\sqrt{-83})/2]$, while those in the lower component correspond to $\mathbb{Z}[\sqrt{-83}]$. Running clockwise through these components corresponds to the repeated action of $[(3, \pi - 1)]$.

Remark 5. In fact, if $p \equiv 3 \bmod 4$ then there are two non-isomorphic curves over $\mathbb{F}_p$ with j-invariant 1728, namely $y^2 = x^3 - x$ and $y^2 = x^3 + x$, whose endomorphism rings are the full ring of integers $\mathbb{Z}[(1 + \sqrt{-p})/2]$ and the order $\mathbb{Z}[\sqrt{-p}]$ of conductor 2 respectively. The connected component of each curve is "symmetric": if E is n steps along $G_{\mathbb{F}_p,\ell}$ in one direction from the curve of j-invariant 1728 then the curve that is n steps in the other direction is the quadratic twist of E. In the case of $G_{\mathbb{F}_{83},3}$ we can see this in Figure 1, which is taken from [16, Figure 8].

It is also interesting to observe that the symmetry around $j = 1728$ confirms the well-known fact that the class numbers of $\mathbb{Z}[(1+\sqrt{-p})/2]$ and $\mathbb{Z}[\sqrt{-p}]$ are odd, at least in the case that $p \equiv 3 \pmod 4$.

3 The class group action

It is well-known that the ideal class group of an imaginary quadratic order $\mathcal{O}$ acts freely on the set of elliptic curves with $\mathbb{F}_p$-rational endomorphism ring $\mathcal{O}$ via isogenies. Using this group action on a set of ordinary elliptic curves for cryptographic purposes was first suggested by Couveignes [12] and independently rediscovered later by Rostovtsev and Stolbunov [44, 38]. Our proposed cryptographic group action is the equivalent of that construction in the supersingular setting, thus the following discussion covers both cases at once. For concreteness, we focus on prime fields with $p \geq 5$ and point out that the ordinary (but not the supersingular) case generalizes to all finite fields. We recall the following standard lemma:

Lemma 6. *Let $E/\mathbb{F}_q$ be an elliptic curve and G a finite $\mathbb{F}_q$-rational (i. e., stable under the action of the $\mathbb{F}_q$-Frobenius) subgroup of E. Then there exists an elliptic curve $E'/\mathbb{F}_q$ and a separable isogeny $\varphi\colon E \to E'$ defined over $\mathbb{F}_q$ with kernel G. The codomain E' and isogeny φ are unique up to $\mathbb{F}_q$-isomorphism.*[8]

Proof. [43, Proposition III.4.12, Remark III.4.13.2, and Exercise III.3.13e]. □

The ideal class group. We recall the definitions and basic properties of class groups of quadratic orders that will be needed in the following. This section is based on [13, §7]. Let K be a quadratic number field and $\mathcal{O} \subseteq K$ an order (that is, a subring which is a free $\mathbb{Z}$-module of rank 2). The *norm* of an $\mathcal{O}$-ideal $\mathfrak{a} \subseteq \mathcal{O}$ is defined as $\mathrm{N}(\mathfrak{a}) = |\mathcal{O}/\mathfrak{a}|$; it equals $\gcd(\{\mathrm{N}(\alpha) \mid \alpha \in \mathfrak{a}\})$. Norms are multiplicative: $\mathrm{N}(\mathfrak{a}\mathfrak{b}) = \mathrm{N}(\mathfrak{a})\mathrm{N}(\mathfrak{b})$.

A *fractional ideal* of $\mathcal{O}$ is an $\mathcal{O}$-submodule of K of the form $\alpha\mathfrak{a}$, where $\alpha \in K^*$ and $\mathfrak{a}$ is an $\mathcal{O}$-ideal.[9] Fractional ideals can be multiplied and conjugated in the evident way, and the norm extends to fractional ideals by requiring multiplicativity. A fractional $\mathcal{O}$-ideal $\mathfrak{a}$ is *invertible* if there exists a fractional $\mathcal{O}$-ideal $\mathfrak{b}$ such that $\mathfrak{a}\mathfrak{b} = \mathcal{O}$. If such a $\mathfrak{b}$ exists, we define $\mathfrak{a}^{-1} = \mathfrak{b}$. Clearly all *principal* fractional ideals $\alpha\mathcal{O}$, where $\alpha \in K^*$, are invertible.

By construction, the set of invertible fractional ideals $I(\mathcal{O})$ forms an abelian group under ideal multiplication. This group contains the principal fractional ideals $P(\mathcal{O})$ as a (clearly normal) subgroup, hence we may define the *ideal class group* of $\mathcal{O}$ as

$$\mathrm{cl}(\mathcal{O}) = I(\mathcal{O})/P(\mathcal{O})\,.$$

[8] This statement remains true in vast generality, but we only need it for finite fields.

[9] Note that the use of the word “ideal” is inconsistent in the literature. We make the convention that “ideal” without qualification refers to an *integral* $\mathcal{O}$-ideal (i. e., an ideal in the sense of ring theory), while fractional ideals are clearly named as such.

Every ideal class $[\mathfrak{a}] \in \mathrm{cl}(\mathcal{O})$ has an integral representative, and for any non-zero $M \in \mathbb{Z}$ there even exists an integral representative of norm coprime to M.

There is a unique *maximal order* of K with respect to inclusion, the ring of integers $\mathcal{O}_K$. The *conductor* of $\mathcal{O}$ (in $\mathcal{O}_K$) is the index $f = [\mathcal{O}_K : \mathcal{O}]$. Away from the conductor, ideals are well-behaved; every $\mathcal{O}$-ideal of norm coprime to the conductor is invertible and factors uniquely into prime ideals.

The class group action. Fix a prime $p \geq 5$ and an (ordinary or supersingular) elliptic curve E defined over $\mathbb{F}_p$. The Frobenius endomorphism π of E satisfies a characteristic equation

$$\pi^2 - t\pi + p = 0\,,$$

in $\mathrm{End}_p(E)$, where $t \in \mathbb{Z}$. (E is supersingular if and only if $t = 0$.) The $\mathbb{F}_p$-rational endomorphism ring $\mathrm{End}_p(E)$ is an order $\mathcal{O}$ in the imaginary quadratic field $K = \mathcal{O} \otimes_{\mathbb{Z}} \mathbb{Q} \cong \mathbb{Q}(\sqrt{\Delta})$, where $\Delta = t^2 - 4p$. We note that $\mathcal{O}$ always contains the Frobenius endomorphism π, and hence the order $\mathbb{Z}[\pi]$.

Any (integral) ideal $\mathfrak{a}$ of $\mathcal{O}$ splits as a product $(\pi\mathcal{O})^r\mathfrak{a}_s$, where $\mathfrak{a}_s \not\subseteq \pi\mathcal{O}$. This defines an elliptic curve $E/\mathfrak{a}$ and an isogeny

$$\varphi_{\mathfrak{a}} \colon E \to E/\mathfrak{a}$$

of degree $\mathrm{N}(\mathfrak{a})$ as follows [51]: the separable part of $\varphi_{\mathfrak{a}}$ has kernel $\bigcap_{\alpha\in\mathfrak{a}_s} \ker\alpha$, and the purely inseparable part consists of r iterations of Frobenius. The isogeny $\varphi_{\mathfrak{a}}$ and codomain $E/\mathfrak{a}$ are both defined over $\mathbb{F}_p$ and unique up to $\mathbb{F}_p$-isomorphism (by Lemma 6), justifying the notation $E/\mathfrak{a}$. Whenever the codomain $E/\mathfrak{a}$ also lies in $\mathcal{E}\ell\ell_p(\mathcal{O})$, multiplication of ideals corresponds to the composition of isogenies. Since principal ideals correspond to endomorphisms, two ideals lead to the same codomain if and only if they are equal up to multiplication by a principal fractional ideal. Moreover, every $\mathbb{F}_p$-isogeny ψ between curves in $\mathcal{E}\ell\ell_p(\mathcal{O})$ comes from an invertible $\mathcal{O}$-ideal in this way, and the ideal $\mathfrak{a}_s$ can be recovered from ψ as $\mathfrak{a}_s = \{\alpha \in \mathcal{O} \mid \ker\alpha \supseteq \ker\psi\}$. In other words:

Theorem 7. *Let $\mathcal{O}$ be an order in an imaginary quadratic field. If $\mathcal{E}\ell\ell_p(\mathcal{O})$ is non-empty, then the ideal class group $\mathrm{cl}(\mathcal{O})$ acts on $\mathcal{E}\ell\ell_p(\mathcal{O})$ via*

$$\begin{array}{ccc} \mathrm{cl}(\mathcal{O}) \times \mathcal{E}\ell\ell_p(\mathcal{O}) & \longrightarrow & \mathcal{E}\ell\ell_p(\mathcal{O}) \\ ([\mathfrak{a}], E) & \longmapsto & E/\mathfrak{a}, \end{array}$$

where $\mathfrak{a}$ is chosen as an integral representative, and this action is free. Furthermore, if $\mathcal{E}\ell\ell_p(\mathcal{O})$ contains a supersingular curve, the action is transitive, else the action has exactly two orbits.

Proof. See [51, Theorem 4.5]. Erratum: [39, Theorem 4.5]. □

The structure of the class group. The class group $\mathrm{cl}(\mathcal{O})$ is a finite abelian group whose cardinality is asymptotically [42]

$$\#\mathrm{cl}(\mathcal{O}) \approx \sqrt{|\Delta|}.$$

More precise heuristics actually predict that $\#\mathrm{cl}(\mathcal{O})$ grows a little bit faster than $\sqrt{|\Delta|}$, but the ratio is logarithmically bounded so we content ourselves with the above estimate. The exact structure of the class group can be computed in subexponential time $L_{|\Delta|}[1/2; \sqrt{2} + o(1)]$ using an algorithm of Hafner and McCurley [22]. Unfortunately, this requires too much computation for the sizes of Δ we are working with, but there are convincing heuristics concerning the properties of the class group we need. See Section 7.1 for these arguments. If the absolute value $|t|$ of the trace of Frobenius is "not too big", the discriminant Δ is about the size of p, hence by the above approximation we may assume $\#\mathrm{cl}(\mathcal{O}) \approx \sqrt{p}$. In particular, if E is supersingular, then $t = 0$, hence $|\Delta| = 4p$.

We are interested in primes ℓ that split in $\mathcal{O}$, i. e., such that there exist (necessarily conjugate) distinct prime ideals $\mathfrak{l}, \bar{\mathfrak{l}}$ of $\mathcal{O}$ with $\ell\mathcal{O} = \mathfrak{l}\bar{\mathfrak{l}}$. Such ℓ are known as *Elkies primes* in the point-counting literature. The ideal $\mathfrak{l}$ is generated as $\mathfrak{l} = (\ell, \pi - \lambda)$, where $\lambda \in \mathbb{Z}/\ell$ is an eigenvalue of the Frobenius endomorphism π on the ℓ-torsion, and its conjugate is $\bar{\mathfrak{l}} = (\ell, \pi - p/\lambda)$. Note that ℓ splits in $\mathcal{O}$ if and only if Δ is a non-zero square modulo ℓ.

Computing the group action. Any element of the class group can be represented as a product of small prime ideals [5, Propositions 9.5.2 and 9.5.3], hence we describe how to compute $E/\mathfrak{l}$ for $\mathfrak{l} = (\ell, \pi - \lambda)$. There are (at least) the following ways to proceed, which vary in efficiency depending on the circumstances [15, 28]:

- Find $\mathbb{F}_p$-rational roots of the modular polynomial $\Phi_\ell(j(E))$ to determine the two possible codomains; compute the kernel polynomial [29] $\chi \in \mathbb{F}_p[x]$ for one of them; if $(x^p, y^p) = [\lambda](x, y)$ modulo χ and the curve equation, then the codomain was correct, else the other choice is correct.
- Factor the ℓ^{th} division polynomial $\psi_\ell(E)$ over $\mathbb{F}_p$; collect irreducible factors with the right Frobenius eigenvalues (as above); use Kohel's algorithm [29] to compute the codomain.
- Find a basis of the ℓ-torsion — possibly over an extension field — and compute the eigenspaces of Frobenius; apply Vélu's formulas [50] to a basis point of the right eigenspace to compute the codomain.

As observed in [28, 15], the last method is the fastest if the necessary extension fields are small. The optimal case is $\lambda = 1$; in that case, the curve has a rational point defined over the base field $\mathbb{F}_p$. If in addition $p/\lambda = -1$, the other eigenspace of Frobenius modulo ℓ is defined over $\mathbb{F}_{p^2}$, so both codomains can easily be computed using Vélu's formulas over an at most quadratic extension (but in fact, a good choice of curve model allows for pure prime field computations, see Section 8; alternatively one could switch to the quadratic twist). Note that if $p \equiv -1 \pmod{\ell}$, then $\lambda = 1$ automatically implies $p/\lambda = -1$.

Much of De Feo–Kieffer–Smith's work [15, 28] is devoted to finding an ordinary elliptic curve E_0 with many small Elkies primes ℓ such that both E_0 and its quadratic twist have an $\mathbb{F}_p$-rational ℓ-torsion point. Despite considerable effort leading to various improvements, the results are discouraging. With the best parameters found within 17 000 hours of CPU time, evaluating one class group action still requires several minutes of computation to complete. This suggests that the original Couveignes–Rostovtsev–Stolbunov scheme will not become anything close to practical in the foreseeable future.

4 Construction and design choices

In this section, we discuss the construction of our proposed group action and justify our design decisions. For algorithmic details, see Section 8. Notice that the main obstacle to performance in the Couveignes–Rostovtsev–Stolbunov scheme — constructing a curve with highly composite order — becomes trivial when using supersingular curves instead of ordinary curves, since for $p \geq 5$ any supersingular elliptic curve over $\mathbb{F}_p$ has exactly $p + 1$ rational points.

The cryptographic group action described in the following is a straightforward implementation of this construction. Note that we require $p \equiv 3 \pmod 4$ such that an implementation may use curves in Montgomery form, and it turns out that this is also beneficial for other reasons. In principle, this constraint is not necessary for the theory to work, although the structure of the isogeny graph changes slightly (see [16] for details).

Parameters. Fix a large prime p of the form $4 \cdot \ell_1 \cdots \ell_n - 1$, where the ℓ_i are small distinct odd primes. Fix the elliptic curve $E_0\colon y^2 = x^3 + x$ over $\mathbb{F}_p$; it is supersingular since $p \equiv 3 \pmod 4$. The Frobenius endomorphism π satisfies $\pi^2 = -p$, so its $\mathbb{F}_p$-rational endomorphism ring is an order in the imaginary quadratic field $\mathbb{Q}(\sqrt{-p})$. More precisely, Proposition 8 shows $\operatorname{End}_p(E_0) = \mathbb{Z}[\pi]$.

Rational Elkies primes. By Theorem 4, the choices made above imply that the ℓ_i-isogeny graph is a disjoint union of cycles. Moreover, since $\pi^2 - 1 \equiv 0 \pmod{\ell_i}$ the ideals $\ell_i\mathcal{O}$ split as $\ell_i\mathcal{O} = \mathfrak{l}_i\overline{\mathfrak{l}_i}$, where $\mathfrak{l}_i = (\ell_i, \pi - 1)$ and $\overline{\mathfrak{l}_i} = (\ell_i, \pi + 1)$. In other words, *all* ℓ_i are Elkies primes. In particular, we can use any one of the three algorithms described above to walk along the cycles.

Furthermore, the kernel of $\varphi_{\mathfrak{l}_i}$ is the intersection of the kernels of the scalar multiplication $[\ell_i]$ and the endomorphism $\pi - 1$. That is, it is the subgroup generated by a point P of order ℓ_i which lies in the kernel of $\pi - 1$ or, in other words, is defined over $\mathbb{F}_p$. Similarly, the point generating the kernel of $\varphi_{\overline{\mathfrak{l}_i}}$ is of order ℓ_i and defined over $\mathbb{F}_{p^2}$ but not $\mathbb{F}_p$. This greatly simplifies and accelerates the implementation, since it allows performing all computations over the base field (see Section 8 for details).

Sampling from the class group. Ideally,[10] we would like to know the exact structure of the ideal class group $\operatorname{cl}(\mathcal{O})$ to be able to sample elements uniformly at random. However, such a computation is currently not feasible for the size of

[10] No pun intended.

discriminant we need, hence we resort to heuristic arguments. Assuming that the $\mathfrak{l}_i$ do not have very small order and are "evenly distributed" in the class group, we can expect ideals of the form $\mathfrak{l}_1^{e_1}\mathfrak{l}_2^{e_2}\cdots\mathfrak{l}_n^{e_n}$ for small e_i to lie in the same class only very occasionally. For efficiency reasons, it is desirable to sample the exponents e_i from a short interval centered around zero, say $\{-m,\dots,m\}$ for some integer m. We will argue in Section 7.1 that choosing m such that $2m+1 \geq \sqrt[n]{\#\mathrm{cl}(\mathcal{O})}$ is sufficient. Since the prime ideals $\mathfrak{l}_i$ are fixed global parameters, the ideal $\prod_i \mathfrak{l}_i^{e_i}$ may simply be represented as a vector $(e_1,\dots,e_n)$.

Evaluating the class group action. Computing the action of an ideal class represented by $\prod_i \mathfrak{l}_i^{e_i}$ on an elliptic curve E proceeds as outlined in Section 3. Since $\pi^2 = -p \equiv 1 \pmod{\ell_i}$, we are now in the favourable situation that the eigenvalues of Frobenius on *all* ℓ_i-torsion subgroups are $+1$ and -1. Hence we can efficiently compute the action of $\mathfrak{l}_i$ (resp. $\overline{\mathfrak{l}_i}$) by finding an $\mathbb{F}_p$-rational (resp. $\mathbb{F}_{p^2}\backslash\mathbb{F}_p$-rational) point of order ℓ_i and applying Vélu-type formulas. This step could simply be repeated for each ideal $\mathfrak{l}_i^{\pm 1}$ whose action is to be evaluated, but see Section 8 for a more efficient method.

5 Representing and validating $\mathbb{F}_p$-isomorphism classes

A major unsolved problem of SIDH is its lack of public-key validation, i. e., the inability to verify that a public key was honestly generated. This shortcoming leads to polynomial-time active attacks [20] on static variants for which countermeasures are expensive. For example, the actively secure variant SIKE [25] applies a transformation proposed by Hofheinz, Hövelmanns and Kiltz [24] which is similar to the Fujisaki–Okamoto transform [19], essentially doubling the running time on the recipient's side compared to an ephemeral key exchange.

Moreover, there is no canonical choice for the representative of the public key, leading to potential difficulties in communicating between different implementations (e. g., using different curve models) of the SIKE protocol. The following proposition tackles both these problems simultaneously for our family of CSIDH instantiations; it shows that the Montgomery coefficient forms a unique representative for the $\mathbb{F}_p$-isomorphism class resulting from the group action (hence may be hashed to obtain a shared secret), and it simplifies public-key validation.

Proposition 8. *Let $p \equiv 3 \pmod 8$ and let $E/\mathbb{F}_p$ be a supersingular elliptic curve. Then $\mathrm{End}_p(E) = \mathbb{Z}[\pi]$ if and only if there exists $A \in \mathbb{F}_p$ such that E is $\mathbb{F}_p$-isomorphic to the curve $E_A\colon y^2 = x^3 + Ax^2 + x$. Moreover, if such an A exists then it is unique.*

Proof. First suppose that E is isomorphic over $\mathbb{F}_p$ to E_A for some $A \in \mathbb{F}_p$. If E_A has full $\mathbb{F}_p$-rational 2-torsion, then Table 1 of [11] shows that either E_A or its quadratic twist must have order divisible by 8. However, both have cardinality $p+1 \equiv 4 \pmod 8$. Hence E_A can only have one $\mathbb{F}_p$-rational point of order 2. With Theorem 2.7 of [16], we can conclude $\mathrm{End}_p(E) = \mathrm{End}_p(E_A) = \mathbb{Z}[\pi]$.

Now assume that $\mathrm{End}_p(E) = \mathbb{Z}[\pi]$. By Theorem 7, the class group $\mathrm{cl}(\mathbb{Z}[\pi])$ acts transitively on $\mathcal{E\!\ell\!\ell}_p(\mathbb{Z}[\pi])$, so in particular there exists $[\mathfrak{a}] \in \mathrm{cl}(\mathbb{Z}[\pi])$ such that $[\mathfrak{a}]E_0 = E$, where $E_0 \colon y^2 = x^3 + x$. Choosing a representative $\mathfrak{a}$ that has norm coprime to $2p$ yields a separable $\mathbb{F}_p$-isogeny $\varphi_{\mathfrak{a}} \colon E_0 \to E$ of odd degree. Thus, by [37, Proposition 1] there exists an $A \in \mathbb{F}_p$ and a separable isogeny $\psi \colon E_0 \to E_A \colon y^2 = x^3 + Ax^2 + x$ defined over $\mathbb{F}_p$ such that $\ker \psi = \ker \varphi_{\mathfrak{a}}$. As isogenies defined over $\mathbb{F}_p$ with given kernel are unique up to post-composition with $\mathbb{F}_p$-isomorphisms (Lemma 6), we conclude that E is $\mathbb{F}_p$-isomorphic to E_A.

Finally, let $B \in \mathbb{F}_p$ such that $E_A \cong E_B \colon Y^2 = X^3 + BX^2 + X$. Then by [43, Proposition III.3.1(b)] there exist $u \in \mathbb{F}_p^*$ and $r, s, t \in \mathbb{F}_p$ such that

$$x = u^2X + r\,, \quad y = u^3Y + su^2X + t\,.$$

Substituting this into the curve equation for E_A and subtracting the equation of E_B (scaled by u^6) equals zero in the function field and thus leads to a linear relation over $\mathbb{F}_p$ between 1, X, X^2, Y, and XY. Writing ∞ for the base point of E_B, it follows from Riemann–Roch [43, Theorem 5.4] that $\mathcal{L}(5(\infty))$ is a 5-dimensional $\mathbb{F}_p$-vector space with basis $\{1, X, Y, X^2, XY\}$. Hence the obtained linear relation must be trivial, and a straightforward computation yields the relations

$$\begin{aligned} s = t &= 0\,, & 3r^2 + 2Ar + 1 &= u^4\,, \\ 3r + A &= Bu^2\,, & r^3 + Ar^2 + r &= 0\,. \end{aligned}$$

But since E_A only has a single $\mathbb{F}_p$-rational point of order 2, the only $r \in \mathbb{F}_p$ such that $r^3 + Ar^2 + r = 0$ is simply $r = 0$. In that case $u^4 = 1$, and hence $u = \pm 1$ since $p \equiv 3 \pmod 8$. In particular, $u^2 = 1$ and thus $A = B$. □

Therefore, by choosing public keys to consist of a Montgomery coefficient $A \in \mathbb{F}_p$, Proposition 8 guarantees that the represented curve has the right endomorphism ring under the assumption that it is supersingular.

Verifying supersingularity. As $p \geq 5$, an elliptic curve E defined over $\mathbb{F}_p$ is supersingular if and only if $\#E(\mathbb{F}_p) = p + 1$ [43, Exercise 5.10]. In general, proving that an elliptic curve has a given order N is easy if the factorization of N is known; exhibiting a subgroup (or in particular, a single point) whose order d is a divisor of N greater than $4\sqrt{p}$ implies the order must be correct. Indeed, the condition $d > 4\sqrt{p}$ implies that there exists only one multiple of d in the Hasse interval $[p + 1 - 2\sqrt{p};\, p + 1 + 2\sqrt{p}]$ [23]. This must be the group order by Lagrange's theorem.

Now note that a random point generally has very large order d. In our case $E(\mathbb{F}_p) \cong \mathbb{Z}/4 \times \prod_{i=1}^n \mathbb{Z}/\ell_i$, so that $\ell_i \mid d$ with probability $(\ell_i - 1)/\ell_i$. Ignoring the even part, this shows that the expected order is lower bounded by

$$\prod_{i=1}^{n} \left(\ell_i - 1 + \frac{1}{\ell_i}\right).$$

This expression is about the same size as $p+1$, and it is easily shown that a random point will with overwhelming probability have order (much) greater than $4\sqrt{p}$. This observation leads to straightforward verification, see Algorithm 1.

Algorithm 1: Verifying supersingularity.

```
Input: An elliptic curve $E/\mathbb{F}_p$ (where $p = 4 \cdot \ell_1 \cdots \ell_n - 1$).
Output: supersingular or ordinary.
Randomly pick a point $P \in E(\mathbb{F}_p)$ and set $d \leftarrow 1$.
for each $\ell_i$ do
    Set $Q_i \leftarrow [(p+1)/\ell_i]P$.
    If $[\ell_i]Q_i \neq \infty$ then return ordinary.        // since $\#E(\mathbb{F}_p) \nmid p+1$
    If $Q_i \neq \infty$ then set $d \leftarrow \ell_i \cdot d$.        // since $\ell_i \mid \operatorname{ord} P$
    If $d > 4\sqrt{p}$ then return supersingular.
```

If the condition $d > 4\sqrt{p}$ does not hold at the end of Algorithm 1, the point P had too small order to prove $\#E(\mathbb{F}_p) = p+1$. In this case one may retry with a new random point P (although this outcome has negligible probability and could just be ignored). There is no possibility of wrongly classifying an ordinary curve as supersingular.

Note moreover that if x-only Montgomery arithmetic is used (as we suggest) and the point P is obtained by choosing a random x-coordinate in $\mathbb{F}_p$, there is no need to differentiate between points defined over $\mathbb{F}_p$ and $\mathbb{F}_{p^2}$. Any point that has an x-coordinate in $\mathbb{F}_p$ but is only defined over $\mathbb{F}_{p^2}$ corresponds to an $\mathbb{F}_p$-rational point on the quadratic twist, which is supersingular if and only if the original curve is supersingular. Therefore one does not need to distinguish between a point on the curve E and its twist; any x-coordinate in $\mathbb{F}_p$ works.

There are more optimized variants of this algorithm; the bulk of the work are the scalar multiplications required to compute the points $Q_i = [(p+1)/\ell_i]P$. Since they are all multiples of P with shared factors, one may more efficiently compute all Q_i at the same time using a divide-and-conquer strategy (at the expense of higher memory usage). See Section 10 (TODO unnumbered ref...), and in particular Algorithm 3, for details.

6 Non-interactive key exchange

Starting from the class group action on supersingular curves and the parameter choices outlined in Sections 3 and 4, one obtains the following non-interactive key exchange protocol.

Setup. Global parameters of the scheme are a large prime $p = 4 \cdot \ell_1 \cdots \ell_n - 1$, where the ℓ_i are small distinct odd primes, and the supersingular elliptic curve $E_0 \colon y^2 = x^3 + x$ over $\mathbb{F}_p$ with endomorphism ring $\mathcal{O} = \mathbb{Z}[\pi]$.

Key generation. The private key is a sequence $(e_1, \ldots, e_n)$ of integers sampled randomly from an interval $\{-m, \ldots, m\}$. These integers represent the ideal class

$[\mathfrak{a}] = [\mathfrak{l}_1^{e_1} \cdots \mathfrak{l}_n^{e_n}] \in \mathrm{cl}(\mathcal{O})$, where $\mathfrak{l}_i = (\ell_i, \pi - 1)$. The public key is the Montgomery coefficient $A \in \mathbb{F}_p$ of the elliptic curve $[\mathfrak{a}]E_0 \colon y^2 = x^3 + Ax^2 + x$ obtained by applying the action of $[\mathfrak{a}]$ to the curve E_0.

Key exchange. Suppose Alice and Bob have keypairs $([\mathfrak{a}], A)$ and $([\mathfrak{b}], B)$. Upon receiving Bob's public key $B \in \mathbb{F}_p$, Alice verifies that the elliptic curve $E_B \colon y^2 = x^3 + Bx^2 + x$ is indeed in $\mathcal{E}\!\ell\!\ell_p(\mathcal{O})$ using Algorithm 1. She then applies the action of her own secret $[\mathfrak{a}]$ to E_B to compute the curve $[\mathfrak{a}]E_B = [\mathfrak{a}][\mathfrak{b}]E_0$. Bob proceeds analogously with his secret $[\mathfrak{b}]$ and Alice's public key A to compute the curve $[\mathfrak{b}]E_A = [\mathfrak{b}][\mathfrak{a}]E_0$. The shared secret is the Montgomery coefficient S of the common secret curve $[\mathfrak{a}][\mathfrak{b}]E_0 = [\mathfrak{b}][\mathfrak{a}]E_0 = [\mathfrak{a}\mathfrak{b}]E_0$ in the form $y^2 = x^3 + Sx^2 + x$, which is the same for Alice and Bob due to the commutativity of $\mathrm{cl}(\mathcal{O})$ and Proposition 8.

Remark 9. Key exchange is not the only application of a cryptographic group action. We briefly discuss other possible protocols.

One can define a 1-bit identification scheme based on the group action, using a key pair $([\mathfrak{a}], A)$ as above. One randomly samples[11] an element $[\mathfrak{b}] \in \mathrm{cl}(\mathcal{O})$ and commits to a curve $E' = [\mathfrak{b}]E_0$. Depending on a challenge bit b, one then releases either $[\mathfrak{b}]$ or $[\mathfrak{c}]$ (see Figure 2). We can turn this into a signature scheme by repeated application of the 1-bit protocol and by applying the Fiat–Shamir [18] or Unruh [49] transformation. Although this appears much simpler than current signature schemes based on isogenies, it is still quite inefficient, therefore we do not pursue it further here.

On the other hand, given how much the group action resembles traditional Diffie–Hellman, we expect DH-based authenticated key exchange protocols to carry over relatively effortlessly (especially due to the ease of verifying the correctness of public keys).

$$\begin{tikzcd}
E_0 \arrow[rr, "{[\mathfrak{a}]}"] \arrow[dr, dashed, "{[\mathfrak{b}]}"'] & & E_A \arrow[dl, dashed, "{[\mathfrak{c}]}"] \\
& E' &
\end{tikzcd}$$

Figure 2. A 1-bit identification protocol.

7 Security

The central problem of our new primitive is the following analogue to the classical discrete logarithm problem.

[11] Note it is not immediately clear how to do this in such a way that both $[\mathfrak{b}]$ and $[\mathfrak{c}]$ are efficiently computable and $[\mathfrak{c}]$ leaks no information about the secret key $[\mathfrak{a}]$.

Problem 10 (Key recovery). Given two supersingular elliptic curves E, E' defined over $\mathbb{F}_p$ with the same $\mathbb{F}_p$-rational endomorphism ring $\mathcal{O}$, find an ideal $\mathfrak{a}$ of $\mathcal{O}$ such that $[\mathfrak{a}]E = E'$. This ideal must be represented in such a way that the action of $[\mathfrak{a}]$ on a curve can be evaluated efficiently, for instance $\mathfrak{a}$ could be a smooth ideal.

Note that just like in the classical group-based scenario, security notions of Diffie–Hellman schemes built from our primitive rely on slightly different hardness assumptions (cf. Section 1.1) that are straightforward translations of the computational and decisional Diffie–Hellman problem. However, continuing the analogy with the classical case, and since we are not aware of any ideas to attack the key exchange without recovering one of the keys, we will assume in the following analysis that the best approach to breaking the key exchange protocol is to solve Problem 10.

No torsion point images. One of the most worrying properties of SIDH seems to be that Alice and Bob publish the images of known points under their secret isogenies along with the codomain curve, i. e., a public key is of the form $(E', \varphi(P), \varphi(Q))$ where $\varphi\colon E \to E'$ is a secret isogeny and $P, Q \in E$ are publicly known points. Although thus far nobody has succeeded in making use of this extra information to break the original scheme, Petit presented an attack using these points when overstretched, highly asymmetric parameters are used [34]. The Couveignes–Rostovtsev–Stolbunov scheme, and consequently our new scheme CSIDH, does not transmit such additional points — a public key consists of *only* an elliptic curve. Thus we are confident that a potential future attack against SIDH based on these torsion points would not apply to CSIDH.

Chosen-ciphertext attacks. As explained in Section 5, the CSIDH group action features efficient public-key validation. This implies it can be used without applying a CCA transform such as the Fujisaki–Okamoto transform [19], thus enabling efficient static–static key exchange and other applications in a post-quantum world.

7.1 Classical security

We begin by considering classical attacks.

Exhaustive key search. The most obvious approach to attack any cryptosystem is to simply search through all possible keys. In the following, we will argue that our construction provides sufficient protection against key search attacks, including dumb brute force and (less naïvely) a meet-in-the-middle approach.

As explained in Section 4, a private key of our scheme consists of a coefficient vector $(e_1, \ldots, e_n)$ where each e_i is in the range $\{-m, \ldots, m\}$, representing the ideal class $[\mathfrak{l}_1^{e_1}\mathfrak{l}_2^{e_2}\cdots\mathfrak{l}_n^{e_n}] \in \mathrm{cl}(\mathcal{O})$. There may (and typically will) be multiple such vectors that represent the same ideal class and thus form equivalent private keys.

However, we argue (heuristically) that the number of *short* representations per ideal class is small. Here and in the following, "short" means that all e_i are in the range $\{-m, \ldots, m\}$. The maximum number of such short representations immediately yields the min-entropy[12] of our sampling method, which measures the amount of work a brute-force attacker has to do while conducting an exhaustive search for the key.

We assume in the following discussion that $\mathrm{cl}(\mathcal{O})$ is "almost cyclic" in the sense that it has a very large cyclic subgroup, say of order N not much smaller than $\#\mathrm{cl}(\mathcal{O})$. According to a heuristic of Cohen and Lenstra, this is true with high probability for a "random" imaginary quadratic field [8, §9.I], and this conjecture is in line with our own experimental evidence. So suppose

$$\varphi\colon \mathrm{cl}(\mathcal{O}) \twoheadrightarrow (\mathbb{Z}/N, +)$$

is a surjective group homomorphism (which may be thought of as a projection to the large cyclic subgroup followed by an isomorphism) and define $\alpha_i = \varphi([\mathfrak{l}_i])$. We may assume that $\alpha_1 = 1$; this can be done without loss of generality whenever at least one of the $[\mathfrak{l}_i]$ has order $\geq N$ in the class group. For some fixed $[\mathfrak{a}] \in \mathrm{cl}(\mathcal{O})$, any short representation $[\mathfrak{l}_1^{e_1}\mathfrak{l}_2^{e_2}\cdots\mathfrak{l}_n^{e_n}] = [\mathfrak{a}]$ yields a short solution to the linear congruence

$$e_1 + e_2\alpha_2 + \cdots + e_n\alpha_n \equiv \varphi([\mathfrak{a}]) \pmod{N},$$

so counting solutions to this congruence gives an upper bound on the number of short representations of $[\mathfrak{a}]$. These solutions are exactly the points in some shifted version (i. e., a coset) of the integer lattice spanned by the rows of the matrix

$$L = \begin{pmatrix} N & 0 & 0 & \cdots & 0 \\ -\alpha_2 & 1 & 0 & \cdots & 0 \\ -\alpha_3 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & & \ddots & \vdots \\ -\alpha_n & 0 & 0 & \cdots & 1 \end{pmatrix},$$

so by applying the Gaussian heuristic [33, Chapter 2, Definition 8] one expects

$$\mathrm{vol}\,[-m; m]^n \,/\, \det L = (2m+1)^n/N$$

short solutions. Since we assumed $\mathrm{cl}(\mathcal{O})$ to be almost cyclic, this ratio is not much bigger than $(2m+1)^n/\#\mathrm{cl}(\mathcal{O})$, which is upper bounded by a small constant for our choice of $2m+1 \geq \sqrt[n]{\#\mathrm{cl}(\mathcal{O})}$.

As a result, we expect the complexity of a brute-force search to be around $2^{\log\sqrt{p}-\varepsilon}$ for some positive ε that is small relative to $\log\sqrt{p}$. To verify our claims, we performed computer experiments with many choices of p of up to 40 bits (essentially brute-forcing the number of representations for all elements) and found no counterexamples to the heuristic result that our sampling method loses only a few bits of brute-force security compared to uniform sampling from the

[12] The min-entropy of a random variable is the negative logarithm of the probability of the most likely outcome.

class group. For our sizes of p, the min-entropy was no more than 4 bits less than that of a perfectly uniform distribution on the class group (i. e. $\varepsilon \leq 4$). Of course this loss factor may grow in some way with bigger choices of p (a plot of the data points for small sizes suggests an entropy loss proportional to $\log\log p$), but we see no indication for it to explode beyond a few handfuls of bits.

Meet-in-the-middle key search. Since a private key trivially decomposes into a product of two smooth ideals drawn from smaller sets (e. g. splitting $[\mathfrak{l}_1^{e_1}\mathfrak{l}_2^{e_2}\cdots\mathfrak{l}_n^{e_n}]$ as $[\mathfrak{l}_1^{e_1}\cdots\mathfrak{l}_\nu^{e_\nu}]\cdot[\mathfrak{l}_{\nu+1}^{e_{\nu+1}}\cdots\mathfrak{l}_n^{e_n}]$ for some $\nu \in \{1,\dots,n\}$), the usual time-memory trade-offs à la baby-step giant-step [40] with an optimal time complexity of $O(\sqrt{\#\mathrm{cl}(\mathcal{O})}) \approx O(\sqrt[4]{p})$ apply.[13] Details, including a memoryless variation of this concept, can be found in Delfs–Galbraith's paper [16].

Remark 11. The algorithms mentioned thus far scale worse than the quantum subexponential attack outlined below, hence one could possibly take a key space smaller than $\#\mathrm{cl}(\mathcal{O})$ without any loss of security (unless the key space is chosen particularly badly, e. g., as a subgroup), which leads to improved performance. We leave a more thorough analysis of this idea for future work.

Pohlig–Hellman-style attacks. Notice that the set $\mathcal{E}\ell\ell_p(\mathcal{O})$ we are acting on does not form a group with efficiently computable operations (that are compatible with the action of $\mathrm{cl}(\mathcal{O})$). Thus there seems to be no way to apply Pohlig–Hellman style algorithms making use of the decomposition of finite abelian groups. In fact, the Pohlig–Hellman algorithm relies on efficiently computable homomorphisms to proper subgroups, which in the setting at hand would correspond to an efficient algorithm that "projects" a given curve to the orbit of E_0 under a *sub*group action. Therefore, we believe the structure of the class group to be largely irrelevant (assuming it is big enough); in particular, we do not require it to have a large prime-order subgroup.

7.2 Quantum security

We also discuss the state of quantum algorithms to solve Problem 10.

The abelian hidden shift problem. A crucial result by Kuperberg [30] is that any abelian hidden shift problem reduces to a dihedral hidden subgroup problem (on a different but closely related oracle). Given an abelian group H of order N, he then presents an algorithm to solve the abelian hidden shift problem with time, query and space complexity $2^{O(\sqrt{\log N})}$. Another subsequent alternative algorithm by Regev [36] achieves polynomial space complexity with a (asymptotically slightly worse) time and query complexity of $2^{O(\sqrt{\log N \log\log N})}$.

[13] Strictly speaking, the complexity depends on the size of the subset one samples private keys from, rather than the size of the class group, but as was argued before, these are approximately equal for our choice of m and n.

This, in turn, has been generalized by Kuperberg [31] to an algorithm using $2^{O(\sqrt{\log N})}$ time, queries and classical space, but only $O(\log N)$ quantum space. All these algorithms have subexponential time and space complexity.

Attacking the isogeny problem. The relevance of these quantum algorithms to Problem 10 has been observed by Childs, Jao and Soukharev [7]. By defining functions $f_0, f_1 \colon \mathrm{cl}(\mathcal{O}) \to \mathcal{E}\!\ell\!\ell_p(\mathcal{O})$ as $f_0 \colon [\mathfrak{b}] \mapsto [\mathfrak{b}]E$ and $f_1 \colon [\mathfrak{b}] \mapsto [\mathfrak{b}]E'$, the problem can be viewed as an abelian hidden shift problem with respect to f_0 and f_1. This requires evaluating the functions f_i on arbitrary ideal classes (i. e. without being given a representative that is a product of ideals of small prime norm) which is non-trivial. However, Childs–Jao–Soukharev show this can be done in subexponential time [7, §4].

From subexponential to practical. An important remark about all those quantum attacks is that they do not immediately lead to estimates for runtime and memory requirements of concrete instantiations on $H = \mathrm{cl}(\mathcal{O})$. Although the algorithms by Kuperberg and Regev are shown to have subexponential complexity, this asymptotic behaviour is not enough to understand the space and time complexity on actual (small) instances. Similarly, although Childs–Jao–Soukharev provide a worst-case bound for evaluating arbitary isogenies, it remains unclear what the (optimal) runtime and memory usage is. We believe further research in this direction is necessary and important, since it will directly impact the cost of an attack. However, we consider this to be out of scope for this work.[14] For the sake of argument we present the security level given in [7, Theorem 5.2], which has a query complexity of

$$L_N\left[1/2, \sqrt{2} + o(1)\right] = \exp\left[\left(\sqrt{2} + o(1)\right)\sqrt{\ln N \ln\ln N}\right], \quad \text{where } N = \#\mathrm{cl}(\mathcal{O})\,.$$

This ignores the (subexponential) cost of evaluating the quantum oracle. Hence it provides a first estimate, whose applicability we consider an open problem. Note that adjusting parameters only involves changing the prime p (and a few numbers derived from it) and is therefore very simple, should it turn out that our initial estimates are insufficient.

Grover's algorithm and claw finding. Applying Grover search [21] via claw finding as described in [26] is fully applicable to CSIDH as well, leading to an attack on Problem 10 in $O(\sqrt[6]{p})$ calls to the quantum oracle f_0. The idea is to split the search space for collisions into a classical $O(\sqrt[6]{p})$ target part and a $O(\sqrt[3]{p})$ search part on which a quantum search is applied. Our choices of p that lead to classical security are also immediately large enough to imply quantum security against this attack (cf. [32, §4.A.5 in Call for Proposals]).

[14] The page margins are certainly too narrow to contain such an analysis.

7.3 Instantiations

Finally we present estimates for some choices of p.

Security levels. We approximate security levels as proposed by NIST for the post-quantum standardization effort [32, §4.A.5]. That is, we define to have a k-bit security level if the required effort for the best attacks is larger than those needed for a key retrieval attack on an ideal block cipher with a k-bit key (e. g. AES-k for $k \in \{128, 192, 256\}$). In other words, under the assumption that the attacks query an oracle on a circuit at least as costly as AES, we should have a query complexity of at least 2^{k-1} resp. $\sqrt{2^k}$ to a classical resp. quantum oracle.

Security estimates. As explained in §7.1, the best classical attack has complexity $O(\sqrt[4]{p})$, while the best currently known quantum attack has complexity

$$L_N\left[1/2, \sqrt{2} + o(1)\right] = \exp\left[\left(\sqrt{2} + o(1)\right)\sqrt{\ln N \ln\ln N}\right], \quad \text{where } N = \#\mathrm{cl}(\mathcal{O})\,.$$

We present this for completeness (ignoring the $o(1)$ factor), but we note again that we expect this complexity to be subject to more careful analysis, taking into account the (in-)feasibility of long sequential quantum operations and the large memory requirement. A recent analysis shows that the classical attack on SIDH (which is the same for CSIDH) is likely slower in practice than current parameter estimates assumed, which is due to the huge memory requirements of the searches [1]. We summarize the resulting attack complexities (before [1]) for some sizes of p in Table 1.

Table 1. Estimated classical and quantum attack complexities rounded to whole bits, where $N = \#\mathrm{cl}(\mathcal{O})$ is approximated as $\sqrt{p}$.

NIST	k	$\log p$	$\log 2^{k-1}$	$\log \sqrt[4]{p}$	$\log \sqrt{2^k}$	$\log L_N[1/2, \sqrt{2}]$
1	128	512	127	128	64	62
3	192	1024	191	256	96	94
5	256	1792	255	448	128	129

Recall that public keys consist of a single element $A \in \mathbb{F}_p$, which may be represented using $\lceil \log p \rceil$ bits. A private key is represented as a list of n integers in $\{-m, \ldots, m\}$, where m was chosen such that $n \log(2m + 1) \approx \log \sqrt{p}$, thus it may be stored using roughly $\log p/2$ bits. Hence the rows in the table correspond to public key sizes of 64, 128 and 224 bytes, and private keys are approximately half that size when encoded optimally.

8 Implementation

In this section, we outline the most important tricks to make the system easier to implement or faster. As pointed out earlier, the crucial step is to use a field of size $4 \cdot \ell_1 \cdots \ell_n - 1$, where the ℓ_i are small distinct odd primes; this implies that *all* ℓ_i are Elkies primes for a supersingular elliptic curve over $\mathbb{F}_p$ and that the action of ideals $(\ell_i, \pi \pm 1)$ can efficiently be computed using $\mathbb{F}_p$-rational points. See Section 4 for these design decisions. The following section focuses on lower-level implementation details.

Montgomery curves. The condition $p + 1 \equiv 4 \pmod 8$ implies that Montgomery curves of the form $y^2 = x^3 + Ax^2 + x$ can be used for the implementation (cf. Proposition 8). As was shown by Costello–Longa–Naehrig [10, §3] we can do all arithmetic on the projective line, and we simply use exactly the same formulas for group operations on curves. For isogeny computations on Montgomery curves we use a projectivized variant (to avoid almost all inversions) of the formulas from Costello and Hisil [9] and Renes [37]. This can be done as follows.

For a fixed prime $\ell \geq 3$, a kernel $G = \langle P \rangle$ of order ℓ and an index $k \in \{1, \ldots, \ell - 1\}$, we write $[k]P = (X_k : Z_k)$. Then by defining $c_i \in \mathbb{F}_p$ such that

$$\prod_{k=1}^{\ell-1} (xZ_i + X_i) = \sum_{i=0}^{\ell-1} c_i x^i$$

we observe that

$$(\pi(a - 3\sigma) : 1) = \big(ac_0c_\ell - 3(c_0c_{\ell-1} - c_1c_\ell) : c_\ell^2\big),$$

following the notation from [37, Proposition 1]. By noticing that $x([k]P) = x([\ell - k]P)$ for all $k \in \{1, \ldots, (\ell - 1)/2\}$ we can reduce the computation needed by about half, and it is easy to see that we can compute $(\pi(a - 3\sigma) : 1)$ iteratively in about $5\ell\mathbf{M} + \ell\mathbf{S}$ operations.[15] We refer to the implementation for more details.

If necessary, a single division at the end of the computation suffices to obtain an affine curve constant. Note that for a given prime ℓ, we could reduce the number of field operations by finding an appropriate representative of the isogeny formulas modulo (a factor of) the ℓ-division polynomial ψ_ℓ (as done in [10] for 3- and 4-isogenies). Although this would allow for a more efficient implementation, we do not pursue this now for the sake of simplicity.

Rational points. Recall that the goal is to evaluate the action of (the class of) an ideal $\mathfrak{l}_1^{e_1} \cdots \mathfrak{l}_n^{e_n}$ on a curve $E \in \mathcal{E}\ell\ell_p(\mathbb{Z}[\sqrt{-p}])$, where each $\mathfrak{l}_i = (\ell_i, \pi - 1)$ is a prime ideal of small odd norm ℓ_i and the e_i are integers in a short interval $\{-m, \ldots, m\}$. We suppose E is given in the form $E_A \colon y^2 = x^3 + Ax^2 + x$.

The obvious way to do this is to consider each factor $\mathfrak{l}_i^{\pm 1}$ in this product and to find the abscissa of a point P of order ℓ_i on E, which (depending on the sign)

[15] Here $\mathbf{M}$ and $\mathbf{S}$ denote a multiplication resp. squaring in $\mathbb{F}_p$.

is defined over $\mathbb{F}_p$ or $\mathbb{F}_{p^2} \setminus \mathbb{F}_p$. This exists by choice of p and ℓ_i (cf. Section 4). Finding such an abscissa amounts to sampling a random x-coordinate, checking whether x^3+Ax^2+x is a square or not (for $\mathfrak{l}_i^{+1}$ resp. $\mathfrak{l}_i^{-1}$) in $\mathbb{F}_p$ (and resampling if it was wrong), followed by a multiplication by $(p+1)/\ell_i$ and repeating from the start if the result is ∞. The kernel of the isogeny given by $\mathfrak{l}_i^{\pm 1}$ is then $\langle P \rangle$, so the isogeny may be computed using Vélu-type formulas. Repeating this procedure for all $\mathfrak{l}_i^{\pm 1}$ gives the result.

However, fixing a sign before sampling a random point effectively means wasting about half of all random points, including an ultimately useless square test. Moreover, deciding on a prime ℓ_i before sampling a point and doing the cofactor multiplication wastes another proportion of the points, including both an ultimately useless square test and a scalar multiplication. Both of these issues can be remedied by not fixing an ℓ_i before sampling a point, but instead taking *any* x-coordinate, determining the smallest field of definition (i. e. $\mathbb{F}_p$ or $\mathbb{F}_{p^2}$) of the corresponding point, and then performing whatever isogeny computations are possible using that point (based on its field of definition and order). The steps are detailed in Algorithm 2.

Algorithm 2: Evaluating the class group action.

```
Input: $A \in \mathbb{F}_p$ and a list of integers $(e_1, \ldots, e_n)$.
Output: $B$ such that $[\mathfrak{l}_1^{e_1} \cdots \mathfrak{l}_n^{e_n}]E_A = E_B$ (where $E_B\colon y^2 = x^3 + Bx^2 + x$).
While some $e_i \neq 0$ do
    Sample a random $x \in \mathbb{F}_p$.
    Set $s \leftarrow +1$ if $x^3 + Ax^2 + x$ is a square in $\mathbb{F}_p$, else $s \leftarrow -1$.
    Let $S = \{i \mid \mathrm{sign}(e_i) = s\}$. If $S = \emptyset$ then start over with a new $x$.
    Let $k \leftarrow \prod_{i \in S} \ell_i$ and compute $Q \leftarrow [(p+1)/k]P$.
    For each $i \in S$ do
        Compute $R \leftarrow [k/\ell_i]Q$. If $R = \infty$ then skip this $i$.
        Compute an isogeny $\varphi\colon E_A \to E_B\colon y^2 = x^3 + Bx^2 + x$ with $\ker \varphi = R$.
        Set $A \leftarrow B$, $Q \leftarrow \varphi(Q)$, $k \leftarrow \ell_i k$, and finally $e_i \leftarrow e_i - s$.
Return $A$.
```

Due to the commutativity of $\mathrm{cl}(\mathcal{O})$, and since we only decrease (the absolute value of) each e_i once we successfully applied the action of $\mathfrak{l}_i^{\pm 1}$ to the current curve, this algorithm indeed computes the action of $[\mathfrak{l}_1^{e_1} \mathfrak{l}_2^{e_2} \cdots \mathfrak{l}_n^{e_n}]$.

Remark 12. Since the probability that a random point has order divisible by ℓ_i (and hence leads to an isogeny step in Algorithm 2) grows with ℓ_i, the isogeny steps for big ℓ_i are typically completed before those for small ℓ_i. Hence it may make sense to sample the exponents e_i for ideals $\mathfrak{l}_i$ from intervals of different lengths depending on the size of ℓ_i, or to not include any very small ℓ_i in the factorization of $p+1$ at all to reduce the expected number of repetitions of the loop above. Note moreover that doing so may also improve the performance of straightforward constant-time adaptions of our algorithms, since it yields stronger upper bounds on the maximum number of required loop iterations (at the expense of slightly higher cost per isogeny computation). Varying the choice of the ℓ_i can also lead to performance immprovements if the resulting prime p has lower Hamming weight. Finding a suitable p is a significant computational effort but needs to be done only once; all users can use the same finite field.

Remark 13. Algorithm 2 is obviously strongly variable-time when implemented naïvely. Indeed, the number of points computed in the isogeny formulas is linear in the degree, hence the iteration counts of certain loops in our implementation are very directly related to the private key. We note that it would not be very hard to create a constant-time implementation based on this algorithm by always performing the maximal required number of iterations in each loop and only storing the results that were actually needed (using constant-time conditional instructions), although this incurs quite a bit of useless computation. We leave the design of optimized constant-time algorithms for future work.

Public-key validation. Recall that the public-key validation method outlined in Section 5 essentially consists of computing $[(p+1)/\ell_i]P$ for each i, where P is a random point on E. Performing this computation in the straightforward way is simple and effective. On the other hand, a divide-and-conquer approach, such as the following recursive algorithm, yields better speeds at the expense of higher memory usage. Note that Algorithm 3 only operates on public data, hence need not be constant-time in a side-channel resistant implementation.

Algorithm 3: Batch cofactor multiplication. [45, Algorithm 7.3]

```
Input: An elliptic curve point $P$ and positive integers $(k_1, \dots, k_n)$.
Output: The points $(Q_1, \dots, Q_n)$, where $Q_i = [\prod_{j \neq i} k_j]P$.
If $n = 1$ then return $(P)$.                                        // base case
Set $m \leftarrow \lceil n/2 \rceil$ and let $u \leftarrow \prod_{i=1}^{m} k_i,\ v \leftarrow \prod_{i=m+1}^{n} \ell_i$.
Compute $R \leftarrow [u]P$ and $L \leftarrow [v]P$.
Recurse with input $L, (k_1, \dots, k_m)$ giving $(Q_1, \dots, Q_m)$.              // left half
Recurse with input $R, (k_{m+1}, \dots, k_n)$ giving $(Q_{m+1}, \dots, Q_n)$.      // right half
Return $(Q_1, \dots, Q_n)$.
```

This routine can be used for verifying that an elliptic curve $E/\mathbb{F}_p$ is supersingular as follows; pick a random point $P \in E(\mathbb{F}_p)$ and run Algorithm 3 on input $[4]P$ and $(\ell_1, \dots, \ell_n)$ to obtain the points $Q_i = [(p+1)/\ell_i]P$. Then continue like in Algorithm 1 to verify that E is supersingular using these precomputed points.

We note that the improved performance of this algorithm essentially comes from a space-time tradeoff, so the memory usage is quite high (cf. Section 8.1). On memory-constrained devices one may instead opt for the naïve algorithm, which requires less space but is slower.

8.1 Performance results

On top of a minimal implementation in the `sage` computer algebra system [48] for demonstrative purposes, we created a somewhat optimized proof-of-concept implementation of the CSIDH group action for a 512-bit prime p.[16] While this implementation features field arithmetic written in assembly (for the Intel Skylake architecture), the rest of the code is in generic C and can easily be ported to other architectures or parameter sets by plugging in a different base field implementation.

The prime p is chosen as $p = 4 \cdot \ell_1 \cdots \ell_{74} - 1$ where ℓ_1 through ℓ_{73} are the smallest 73 odd primes and $\ell_{74} = 587$ is the smallest choice distinct from the other ℓ_i that renders p prime. This parameter choice implies that public keys have a size of 64 bytes, whereas private keys are stored in 37 bytes for simplicity (but an optimal encoding would reduce this to only 32 bytes). Table 2 summarizes performance numbers for our proof-of-concept implementation. Note that private key generation is not listed as it only consists of sampling n random integers in a small interval $\{-m, \ldots, m\}$, which has negligible cost.

Table 2. Performance numbers of our proof-of-concept implementation, averaged over 10 000 runs on an Intel Skylake i5 processor running at 3.5 GHz.

	Clock cycles	**Wall-clock time**	**Stack memory**
Key validation	$8.7 \cdot 10^6$ cc	3.3 ms	12 336 bytes
Group action	$118 \cdot 10^6$ cc	45.4 ms	2 432 bytes

We emphasize that both our implementations are intended as a proof of concept and unfit for production use; in particular, they are explicitly *not side-channel resistant* and may contain any number of bugs. We leave the design of hardened and more optimized implementations for future work.

References

[1] 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, 2018. IACR Cryptology ePrint Archive 2018/313 https://ia.cr/2018/313.

[2] Daniel J. Bernstein, Bernard van Gastel, Wesley Janssen, Tanja Lange, Peter Schwabe, and Sjaak Smetsers. TweetNaCl: A crypto library in 100 tweets. In *LATINCRYPT*, volume 8895 of *Lecture Notes in Computer Science*, pages 64–83. Springer, 2014.

[3] Reinier Bröker. A p-adic algorithm to compute the Hilbert class polynomial. *Math. Comput.*, 77(264):2417–2435, 2008.

[16] All our code is published in the public domain and is available for download at https://yx7.cc/code/csidh/csidh-20180427.tar.xz.

[4] Reinier Bröker and Peter Stevenhagen. Efficient CM-constructions of elliptic curves over finite fields. *Math. Comput.*, 76(260):2161–2179, 2007.

[5] Johannes Buchmann and Ulrich Vollmer. *Binary quadratic forms: an algorithmic approach*, volume 20 of *Algorithms and Computation in Mathematics*. Springer, 2007.

[6] Denis X. 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.

[7] 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.

[8] Henri Cohen and Hendrik W. Lenstra, jr. Heuristics on class groups of number fields. In Hendrik Jager, editor, *Number Theory Noordwijkerhout 1983*, pages 33–62. Springer, 1984.

[9] Craig Costello and Hüseyin Hisil. A simple and compact algorithm for SIDH with arbitrary degree isogenies. In *ASIACRYPT (2)*, volume 10625 of *Lecture Notes in Computer Science*, pages 303–329. Springer, 2017. https://ia.cr/2017/504.

[10] Craig Costello, Patrick Longa, and Michael Naehrig. Efficient algorithms for Supersingular Isogeny Diffie–Hellman. In *CRYPTO (1)*, volume 9814 of *Lecture Notes in Computer Science*, pages 572–601. Springer, 2016. https://ia.cr/2016/413.

[11] Craig Costello and Benjamin Smith. Montgomery curves and their arithmetic: The case of large characteristic fields, 2017. IACR Cryptology ePrint Archive 2017/212 https://ia.cr/2017/212.

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

[13] David A. Cox. *Primes of the form $x^2 + ny^2$: Fermat, class field theory, and complex multiplication.* Pure and applied mathematics. Wiley, 2nd edition, 2013.

[14] Luca De Feo. Mathematics of isogeny based cryptography. *CoRR*, abs/1711.04062, 2017. https://arxiv.org/abs/1711.04062.

[15] Luca De Feo, Jean Kieffer, and Benjamin Smith. Towards practical key exchange from ordinary isogeny graphs, 2018. Preprint.

[16] Christina Delfs and Steven D. Galbraith. Computing isogenies between supersingular elliptic curves over $\mathbb{F}_p$. *Des. Codes Cryptography*, 78(2):425–440, 2016. https://arxiv.org/abs/1310.7789.

[17] Whitfield Diffie and Martin E. Hellman. New directions in cryptography. *IEEE Trans. Information Theory*, 22(6):644–654, 1976.

[18] Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *CRYPTO*, volume 263 of *Lecture Notes in Computer Science*, pages 186–194. Springer, 1986.

[19] 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.

[20] 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, 2016. https://ia.cr/2016/859.

[21] Lov K. Grover. A fast quantum mechanical algorithm for database search. In *STOC*, pages 212–219. ACM, 1996. https://arxiv.org/abs/quant-ph/9605043.

[22] James L. Hafner and Kevin S. McCurley. A rigorous subexponential algorithm for computation of class groups. *J. Amer. Math. Soc.*, 2(4):837–850, 1989.

[23] Helmut Hasse. Zur Theorie der abstrakten elliptischen Funktionenkörper III. Die Struktur des Meromorphismenrings. Die Riemannsche Vermutung. *Journal für die reine und angewandte Mathematik*, 175:193–208, 1936.

[24] Dennis Hofheinz, Kathrin Hövelmanns, and Eike Kiltz. A modular analysis of the Fujisaki–Okamoto transformation. In *TCC (1)*, volume 10677 of *Lecture Notes in Computer Science*, pages 341–371. Springer, 2017. https://ia.cr/2017/604.

[25] 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 [32]. http://sike.org.

[26] 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://ia.cr/2011/506.

[27] David Jao, Stephen D. Miller, and Ramarathnam Venkatesan. Expander graphs based on GRH with an application to elliptic curve cryptography. *J. Number Theory*, 129(6):1491–1504, 2009. https://arxiv.org/abs/0811.0647.

[28] Jean Kieffer. *Étude et accélération du protocole d'échange de clés de Couveignes–Rostovtsev–Stolbunov*. Mémoire du Master 2, Université Paris VI, 2017. https://arxiv.org/abs/1804.10128.

[29] David Kohel. *Endomorphism rings of elliptic curves over finite fields*. Ph.D. thesis, University of California at Berkeley, 1996. http://iml.univ-mrs.fr/~kohel/pub/thesis.pdf.

[30] 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.

[31] Greg Kuperberg. Another subexponential-time quantum algorithm for the Dihedral Hidden Subgroup Problem. In *TQC*, volume 22 of *LIPIcs*, pages 20–34. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2013. https://arxiv.org/abs/1112.3333.

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

[33] Phong Q. Nguyen and Brigitte Vallée, editors. *The LLL Algorithm*. Springer, 2010.

[34] Christophe Petit. Faster algorithms for isogeny problems using torsion point images. In *ASIACRYPT (2)*, volume 10625 of *Lecture Notes in Computer Science*, pages 330–353. Springer, 2017. https://ia.cr/2017/571.

[35] Arnold K. Pizer. Ramanujan graphs and Hecke operators. *Bull. Amer. Math. Soc. (N.S.)*, 23(1):127–137, 1990. https://projecteuclid.org:443/euclid.bams/1183555725.

[36] Oded Regev. A subexponential time algorithm for the dihedral hidden subgroup problem with polynomial space, 2004. https://arxiv.org/abs/quant-ph/0406151.

[37] Joost Renes. Computing isogenies between Montgomery curves using the action of $(0,0)$. In *PQCrypto*, volume 10786 of *Lecture Notes in Computer Science*, pages 229–247. Springer, 2018. https://ia.cr/2017/1198.

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

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

[40] Daniel Shanks. Class number, a theory of factorization, and genera. In *Proc. Symp. Pure Math*, volume 20, pages 415–440, 1971.

[41] 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.

[42] Carl Siegel. Über die Classenzahl quadratischer Zahlkörper. *Acta Arithmetica*, 1(1):83–86, 1935.

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

[44] Anton Stolbunov. Public-key encryption based on cycles of isogenous elliptic curves. Master's thesis, Saint-Petersburg State Polytechnical University, 2004. In Russian.

[45] Andrew V. Sutherland. *Order computations in generic groups.* PhD thesis, Massachusetts Institute of Technology, 2007. https://groups.csail.mit.edu/cis/theses/sutherland-phd.pdf.

[46] Andrew V. Sutherland. Isogeny volcanoes. In *ANTS X*, pages 507–530, 2012. https://arxiv.org/abs/1208.5370.

[47] John Tate. Endomorphisms of abelian varieties over finite fields. *Inventiones mathematicae*, 2(2):134–144, 1966.

[48] The Sage Developers. *SageMath, the Sage Mathematics Software System (version 8.1)*, 2018. https://sagemath.org.

[49] Dominique Unruh. Quantum proofs of knowledge. In *EUROCRYPT*, volume 7237 of *Lecture Notes in Computer Science*, pages 135–152. Springer, 2012. https://ia.cr/2010/212.

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

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