TOWARDS QUANTUM-RESISTANT CRYPTOSYSTEMS FROM SUPERSINGULAR ELLIPTIC CURVE ISOGENIES

LUCA DE FEO, DAVID JAO, AND JÉRÔME PLÛT

ABSTRACT. We present new candidates for quantum-resistant public-key cryptosystems based on the conjectured difficulty of finding isogenies between supersingular elliptic curves. The main technical idea in our scheme is that we transmit the images of torsion bases under the isogeny in order to allow the parties to construct a shared commutative square despite the noncommutativity of the endomorphism ring. Our work is motivated by the recent development of a subexponential-time quantum algorithm for constructing isogenies between ordinary elliptic curves. In the supersingular case, by contrast, the fastest known quantum attack remains exponential, since the noncommutativity of the endomorphism ring means that the approach used in the ordinary case does not apply. We give a precise formulation of the necessary computational assumptions along with a discussion of their validity, and prove the security of our protocols under these assumptions. In addition, we present implementation results showing that our protocols are multiple orders of magnitude faster than previous isogeny-based cryptosystems over ordinary curves.

This paper is an extended version of [19]. We add a new zero-knowledge identification scheme, and detailed security proofs for the protocols. We also present a new, asymptotically faster, algorithm for key generation, a thorough study of its optimization, and new experimental data.

Keywords: elliptic curves, isogenies, quantum-resistant public-key cryptosystems

1. INTRODUCTION

The Diffie-Hellman scheme is a fundamental protocol for public-key exchange between two parties. Its original definition over finite fields is based on the hardness of computing the map $g, g^a, g^b \mapsto g^{ab}$ for $g \in \mathbb{F}_p^*$, while its elliptic curve analogue depends on the difficulty of computing $P, aP, bP \mapsto abP$ for points P on an elliptic curve. Recently, Stolbunov [43] proposed a Diffie-Hellman type system based on the difficulty of computing isogenies between ordinary elliptic curves, with the stated aim of obtaining quantum-resistant cryptographic protocols. The fastest known (classical) probabilistic algorithm for solving this problem is the algorithm of Galbraith and Stolbunov [17], based on the algorithm of Galbraith, Hess, and Smart [16]. This algorithm is exponential, with a worst-case running time of $O(\sqrt[4]{q})$. However, on a quantum computer, recent work of Childs et al. [10] has shown that the private keys in Stolbunov's system can be recovered in subexponential time. Moreover, even if we only consider classical attacks in assessing security levels, Stolbunov's scheme requires 229 seconds (even with precomputation) to perform a single key exchange operation at the 128-bit security level on a desktop PC [43, Table 1].

In this work we present isogeny-based key-exchange, encryption and identification schemes that address both the performance and security drawbacks of Stolbunov's system. Our primitive achieves performance on the order of 60 milliseconds (cf. Section 7) at the 128-bit security level (as measured against the fastest known quantum attacks) using desktop PCs, making our schemes far faster than Stolbunov's. In terms of security, our schemes are not vulnerable to the algorithm of Childs et al. [10], nor to any algorithm of this type, since they are not based on a group action. The fastest known attacks against our schemes, even on quantum computers, require fully exponential time. Our schemes involve new computational assumptions upon which their quantum resistance is based, and like all new computational assumptions, further study and the passage of time is needed for validation. Nevertheless, we believe our proposal represents a promising candidate for quantum-resistant isogeny-based public-key cryptography.

Our proposal, presented in Section 3, uses isogenies between *supersingular* elliptic curves rather than ordinary elliptic curves. The main technical difficulty is that, in the supersingular case, the endomorphism

ring is noncommutative, whereas Diffie-Hellman type protocols require commutativity. We show how to overcome this obstacle by providing the outputs of the isogeny on certain points as auxiliary input to the protocols. To the best of our knowledge, nothing similar to this idea has ever previously appeared in the literature. Providing this auxiliary input does not seem to make the problem of finding isogenies any easier, and the added difficulty arising from noncommutativity actually makes our protocols stronger; see Section 5.1 for a full discussion. The multiple orders of magnitude of performance gains in our scheme arise from the fact that supersingular isogeny graphs are much faster to navigate than ordinary graphs, as described in Section 4. In Sections 5 and 6 we provide formal statements of the hardness assumptions and security reductions for our system. Finally, in Section 7 we present implementation results confirming the correctness and performance of our protocol.

2. Isogenies

Let E_1 and E_2 be elliptic curves defined over a finite field $\mathbb{F}_q$. An isogeny $\phi : E_1 \to E_2$ defined over $\mathbb{F}_q$ is a non-constant rational map defined over $\mathbb{F}_q$ which is also a group homomorphism from $E_1(\mathbb{F}_q)$ to $E_2(\mathbb{F}_q)$ [37, III.4]. The degree of an isogeny is its degree as a rational map. For separable isogenies, to have degree ℓ means to have kernel of size ℓ. Every isogeny of degree greater than 1 can be factored into a composition of isogenies of prime degree over $\bar{\mathbb{F}}_q$ [11].

An endomorphism of an elliptic curve E defined over $\mathbb{F}_q$ is an isogeny $E \to E$ defined over $\mathbb{F}_{q^m}$ for some m. The set of endomorphisms of E together with the zero map forms a ring under the operations of pointwise addition and composition; this ring is called the endomorphism ring of E and denoted $\mathrm{End}(E)$. The ring $\mathrm{End}(E)$ is isomorphic either to an order in a quaternion algebra or to an order in an imaginary quadratic field [37, V.3.1]; in the first case we say E is supersingular and in the second case we say E is ordinary.

Two elliptic curves E_1 and E_2 defined over $\mathbb{F}_q$ are said to be isogenous over $\mathbb{F}_q$ if there exists an isogeny $\phi : E_1 \to E_2$ defined over $\mathbb{F}_q$. A theorem of Tate states that two curves E_1 and E_2 are isogenous over $\mathbb{F}_q$ if and only if $\#E_1(\mathbb{F}_q) = \#E_2(\mathbb{F}_q)$ [45, §3]. Since every isogeny has a dual isogeny [37, III.6.1], the property of being isogenous over $\mathbb{F}_q$ is an equivalence relation on the finite set of $\bar{\mathbb{F}}_q$-isomorphism classes of elliptic curves defined over $\mathbb{F}_q$. Accordingly, we define an isogeny class to be an equivalence class of elliptic curves, taken up to $\bar{\mathbb{F}}_q$-isomorphism, under this equivalence relation.

The ℓ-torsion group of E, denoted $E[\ell]$, is the set of all points $P \in E(\bar{\mathbb{F}}_q)$ such that ℓP is the identity. For ℓ such that $p \nmid \ell$, we have $E[\ell] \cong \mathbb{Z}/\ell\mathbb{Z} \oplus \mathbb{Z}/\ell\mathbb{Z}$.

Curves in the same isogeny class are either all supersingular or all ordinary. We assume for the remainder of this paper that we are in the supersingular case.

Supersingular curves are all defined over $\mathbb{F}_{p^2}$, and for every prime $\ell \nmid p$, there exist $\ell+1$ isogenies (counting multiplicities) of degree ℓ originating from any given such supersingular curve. Given an elliptic curve E and a finite subgroup Φ of E, there is up to isomorphism a unique isogeny $E \to E'$ having kernel Φ [37, III.4.12]. Hence we can identify an isogeny by specifying its kernel, and conversely given a kernel subgroup the corresponding isogeny can be computed using Vélu's formulas [48]. Typically, this correspondence is of little use, since the kernel, or any representation thereof, is usually as unwieldy as the isogeny itself. However, in the special case of kernels generated by $\mathbb{F}_{p^2}$-rational points of smooth order, specifying a generator of the kernel allows for compact representation and efficient computation of the corresponding isogeny, as we demonstrate below.

2.1. **Ramanujan graphs.** Let $G = (\mathcal{V}, \mathcal{E})$ be a finite graph on h vertices $\mathcal{V}$ with undirected edges $\mathcal{E}$. Suppose G is a regular graph of degree k, i.e., exactly k edges meet at each vertex. Given a labeling of the vertices $\mathcal{V} = \{v_1, \ldots, v_h\}$, the adjacency matrix of G is the symmetric $h \times h$ matrix A whose ij-th entry $A_{i,j} = 1$ if an edge exists between v_i and v_j and 0 otherwise.

It is convenient to identify functions on $\mathcal{V}$ with vectors in $\mathbb{R}^h$ via this labeling, and therefore also think of A as a self-adjoint operator on $L^2(\mathcal{V})$. All of the eigenvalues of A satisfy the bound $|\lambda| \leq k$. Constant vectors are eigenfunctions of A with eigenvalue k, which for obvious reasons is called the trivial eigenvalue λ_{triv}. A family of such graphs G with $h \to \infty$ is said to be a sequence of *expander graphs* if all other

eigenvalues of their adjacency matrices are bounded away from $\lambda_{\text{triv}} = k$ by a fixed amount.[1] In particular, no other eigenvalue is equal to k; this implies the graph is connected. A Ramanujan graph is a special type of expander which has $|\lambda| \leq 2\sqrt{k-1}$ for any nontrivial eigenvalue which is not equal to $-k$ (this last possibility happens if and only if the graph is bipartite). The Ramanujan property was first defined in [26]. It characterizes the optimal separation between the two largest eigenvalues of the graph adjacency matrix, and implies the expansion property.

A fundamental use of expanders is to prove the rapid mixing of the random walk on $\mathcal{V}$ along the edges $\mathcal{E}$. The following rapid mixing result is standard but we present it below for completeness. For the proof, see [20] or [12, 25, 35].

Proposition 2.1. Let G be a regular graph of degree k on h vertices. Suppose that the eigenvalue λ of any nonconstant eigenvector satisfies the bound $|\lambda| \leq c$ for some $c < k$. Let S be any subset of the vertices of G, and x be any vertex in G. Then a random walk of length at least $\frac{\log 2h/|S|^{1/2}}{\log k/c}$ starting from x will land in S with probability at least $\frac{|S|}{2h} = \frac{|S|}{2|G|}$.

2.2. **Isogeny graphs.** An isogeny graph is a graph whose nodes consist of all elliptic curves in $\mathbb{F}_q$ belonging to a fixed isogeny class, up to $\bar{\mathbb{F}}_q$-isomorphism (so that two elliptic curves which are isomorphic over $\bar{\mathbb{F}}_q$ represent the same node in the graph). In practice, the nodes are represented using j-invariants, which are invariant up to isomorphism. Isogeny graphs for supersingular elliptic curves were first considered by Mestre [27], and were shown by Pizer [31, 32] to have the Ramanujan property.

Every supersingular elliptic curve in characteristic p is defined over either $\mathbb{F}_p$ or $\mathbb{F}_{p^2}$ [37], so it suffices to fix $\mathbb{F}_q = \mathbb{F}_{p^2}$ as the field of definition for this discussion. Thus, in contrast to ordinary curves, there are a finite number of isomorphism classes of supersingular curves in any given isogeny class; this number is in fact $g+1$, where g is the genus of the modular curve $X_0(p)$, which is roughly $p/12$. It turns out that all supersingular curves defined over $\mathbb{F}_{p^2}$ belong to the same isogeny class [27]. For a fixed prime value of $\ell \neq p$, we define the vertices of the supersingular isogeny graph $\mathcal{G}$ to consist of these g isomorphism classes of curves, with edges given by isomorphism classes of degree-ℓ isogenies, defined as follows: two isogenies $\phi_1, \phi_2 \colon E_i \to E_j$ are isomorphic if there exists an automorphism $\alpha \in \text{Aut}(E_j)$ (i.e., an invertible endomorphism) such that $\phi_2 = \alpha \phi_1$. Pizer [31, 32] has shown that $\mathcal{G}$ is a connected $k = \ell + 1$-regular multigraph satisfying the Ramanujan bound of $|\lambda| \leq 2\sqrt{\ell} = 2\sqrt{k-1}$ for the nontrivial eigenvalues of its adjacency matrix.

3. Public-key cryptosystems based on supersingular curves

In this section we present a key-exchange protocol and a public-key cryptosystem analogous to those of [34, 43], and a zero-knowledge identification scheme, all using supersingular elliptic curves.

Our protocols require supersingular curves of smooth order. Such curves are normally unsuitable for cryptography since discrete logarithms on them are easy. However, since the discrete logarithm problem is unimportant in our setting, this issue does not affect us. In the supersingular setting, it is easy to construct curves of smooth order, and using a smooth order curve will give a large number of isogenies that are fast to compute. Specifically, we fix $\mathbb{F}_q = \mathbb{F}_{p^2}$ as the field of definition, where p is a prime of the form $\ell_A^{e_A} \ell_B^{e_B} \cdot f \pm 1$. Here ℓ_A and ℓ_B are small primes, and f is a cofactor such that p is prime. Then we construct a supersingular curve E defined over $\mathbb{F}_q$ of cardinality $(\ell_A^{e_A} \ell_B^{e_B} f)^2$. By construction, $E[\ell_A^{e_A}]$ is $\mathbb{F}_q$-rational and contains $\ell_A^{e_A-1}(\ell_A + 1)$ cyclic subgroups of order $\ell_A^{e_A}$, each defining a different isogeny; the analogous statement holds for $E[\ell_B^{e_B}]$.

---

[1]Expansion is usually phrased in terms of the number of neighbors of subsets of G, but the spectral condition here is equivalent for k-regular graphs and also more useful for our purposes.

Our protocols revolve around the following commutative diagram

(1)
$$\begin{array}{ccc} E & \xrightarrow{\ \phi\ } & E/\langle P\rangle \\ \downarrow{\psi} & & \downarrow \\ E/\langle Q\rangle & \longrightarrow & E/\langle P,Q\rangle \end{array}$$

where ϕ and ψ are random walks in the graphs of isogenies of degrees ℓ_A and ℓ_B respectively. Their security is based on the difficulty of finding a path connecting two given vertices in a graph of supersingular isogenies. We refer to Section 4 for low-level algorithmic details, and Section 5 for a full discussion of security.

3.1. **Zero-knowledge proof of identity.** We begin with the protocol which is easiest to understand. Peggy knows a cyclic degree $\ell_A^{e_A}$ isogeny $\phi : E \to E/\langle S\rangle$, with the curves E and $E/\langle S\rangle$ publicly known, and wants to prove Vic that she knows a generator for $\langle S\rangle$, without revealing it.

Our protocol is loosely inspired by the zero-knowledge proof of membership for Graph Isomorphism [18]. In that protocol, Peggy shows that she knows a graph isomorphism $G \simeq G'$ by first publishing a random H such that the following diagram commutates

(2)
$$\begin{array}{ccc} G & \longleftrightarrow & G' \\ & \nwarrow\phi \quad \psi\nearrow & \\ & H & \end{array}$$

and then revealing only one among ϕ and ψ. Intuitively, this protocol is *perfectly* zero-knowledge because the information that Peggy reveals (i.e., a random permutation of G or G') could be easily computed by anyone without her help.

In an analogous way, our protocol consists in publishing the vertices of diagram (1), and then revealing some, but not all, of its arrows. Unlike the case of Graph Isomorphism, in our protocol Peggy needs to use her secret knowledge to create the diagram, thus we cannot achieve a *perfect* zero-knowledge. Nevertheless, we will show in Section 6 that, under suitable assumptions, our protocol is *computationally* zero-knowledge.

We show below the diagram used in our protocol. $\langle S\rangle$ is the kernel of the secret isogeny ϕ of degree $\ell_A^{e_A}$, while $\langle R\rangle$ is a cyclic group of order $\ell_B^{e_B}$.

(3)
$$\begin{array}{ccc} E & \xrightarrow{\ \phi\ } & E/\langle S\rangle \\ \downarrow{\psi} & & \downarrow{\psi'} \\ E/\langle R\rangle & \xrightarrow{\ \phi'\ } & E/\langle S,R\rangle \end{array}$$

Peggy can compute the diagram as follows:

- She uses Vélu's formulas to compute the isogeny $\psi : E \to E/\langle R\rangle$;
- She computes $R' = \phi(R)$ and the isogeny $\psi' : E/\langle S\rangle \to E/\langle S,R\rangle$;
- She computes $S' = \psi(S)$ and the isogeny $\phi' : E/\langle R\rangle \to E/\langle S,R\rangle$.

Now, the natural question is: which arrows of the diagram can Peggy reveal without compromising her secret ϕ? It is not hard to see, and we will show it in Theorem 6.3, that the knowledge of (ψ, ϕ') or (ψ', ϕ') allows anyone to compute the kernel of ϕ. However, we will argue that there is no obvious way to compute ϕ from the sole knowledge of ϕ'. Revealing one of ψ or ψ' is no problem either, however revealing (ψ, ψ') altogether is more subtle. Indeed, revealing the points R and $\phi(R)$ uncovers some information on the action of ϕ on $E[\ell_B^{e_B}]$: it is to be expected that after a few iterations Peggy will reveal a basis (P, Q) of $E[\ell_B^{e_B}]$ and the respective images $\phi(P), \phi(Q)$, thus allowing anyone to evaluate ϕ on the whole $E[\ell_B^{e_B}]$. Nevertheless, we

conjecture that this leakage does not compromise Peggy's secret either, and we make these data part of the public parameters.[2]

Finally, we present our protocol.

Secret parameters: A supersingular curve E defined over $\mathbb{F}_q$ and a primitive $\ell_A^{e_A}$-torsion point S defining an isogeny $\phi : E \to E/\langle S \rangle$.

Public parameters: The curves E and $E/\langle S \rangle$. Generators P, Q of $E[\ell_B^{e_B}]$ and their images $\phi(P), \phi(Q)$.

Identification: Repeat m times:
(1) Peggy chooses a random primitive $\ell_B^{e_B}$-torsion point R and computes diagram (3).
(2) Peggy sends the curves $E_1 = E/\langle R \rangle$ and $E_2 = E/\langle S, R \rangle$ to Vic.
(3) Vic selects a random bit b and sends it to Peggy.
(4) If $b = 0$, Peggy reveals the points R and $\phi(R')$. Vic accepts if they have order $\ell_b^{e_B}$ and generate the kernels of isogenies $E \to E_1$ and $E/\langle S \rangle \to E_2$, respectively.
(5) If $b = 1$, Peggy reveals the point $\psi(S)$. Vic accepts if it has order $\ell_A^{e_A}$ and generates the kernel of an isogeny $E_1 \to E_2$.

3.2. **Key exchange.** The key exchange protocol is a variation *à la* Diffie-Hellman over diagram (1). The idea is to let Alice choose ϕ, while Bob chooses ψ. Although similar in spirit to the protocol based on the action of the class group on ordinary elliptic curves of [43], a main technical difference is that, since ideal classes no longer commute (or indeed even multiply together) in the supersingular case, extra information must be communicated as part of the protocol in order to ensure that both parties arrive at the same common value.

We fix as public parameters a supersingular curve E_0 defined over $\mathbb{F}_{p^2}$, and bases $\{P_A, Q_A\}$ and $\{P_B, Q_B\}$ which generate $E_0[\ell_A^{e_A}]$ and $E_0[\ell_B^{e_B}]$ respectively, so that $\langle P_A, Q_A \rangle = E_0[\ell_A^{e_A}]$ and $\langle P_B, Q_B \rangle = E_0[\ell_B^{e_B}]$. Alice chooses two random elements $m_A, n_A \in_R \mathbb{Z}/\ell_A^{e_A}\mathbb{Z}$, not both divisible by ℓ_A, and computes an isogeny $\phi_A : E_0 \to E_A$ with kernel $K_A := \langle [m_A]P_A + [n_A]Q_A \rangle$. Alice also computes the image $\{\phi_A(P_B), \phi_A(Q_B)\} \subset E_A$ of the basis $\{P_B, Q_B\}$ for $E_0[\ell_B^{e_B}]$ under her secret isogeny ϕ_A, and sends these points to Bob together with E_A. Similarly, Bob selects random elements $m_B, n_B \in_R \mathbb{Z}/\ell_B^{e_B}\mathbb{Z}$ and computes an isogeny $\phi_B : E_0 \to E_B$ having kernel $K_B := \langle [m_B]P_B + [n_B]Q_B \rangle$, along with the points $\{\phi_B(P_A), \phi_B(Q_A)\}$. Upon receipt of E_B and $\phi_B(P_A), \phi_B(Q_A) \in E_B$ from Bob, Alice computes an isogeny $\phi'_A : E_B \to E_{AB}$ having kernel equal to $\langle [m_A]\phi_B(P_A) + [n_A]\phi_B(Q_A) \rangle$; Bob proceeds *mutatis mutandis*. Alice and Bob can then use the common j-invariant of

$$E_{AB} = \phi'_B(\phi_A(E_0)) = \phi'_A(\phi_B(E_0)) = E_0/\langle [m_A]P_A + [n_A]Q_A, [m_B]P_B + [n_B]Q_B \rangle$$

to form a secret shared key.

The full protocol is given in Figure 1. We denote by A and B the identifiers of Alice and Bob, and use sID to denote the unique session identifier.

3.3. **Public-key encryption.** The key-exchange protocol of Section 3.2 can easily be adapted to yield a public-key cryptosystem, in much the same way that Elgamal encryption follows from Diffie-Hellman. We briefly give the details here. All notation is the same as above. Stolbunov [43] uses a similar construction, upon which ours is based.

Setup: Choose $p = \ell_A^{e_A} \ell_B^{e_B} \cdot f \pm 1$, E_0, $\{P_A, Q_A\}$, $\{P_B, Q_B\}$ as above. Let $\mathcal{H} = \{H_k : k \in K\}$ be a hash function family indexed by a finite set K, where each H_k is a function from $\mathbb{F}_{p^2}$ to the message space $\{0, 1\}^w$.

Key generation: Choose two random elements $m_A, n_A \in_R \mathbb{Z}/\ell_A^{e_A}\mathbb{Z}$, not both divisible by ℓ_A. Compute $E_A, \phi_A(P_B), \phi_A(Q_B)$ as above, and choose a random element $k \in_R K$. The public key is the tuple $(E_A, \phi_A(P_B), \phi_A(Q_B), k)$ and the private key is (m_A, n_A, k).

---

[2]An alternative solution, that intuitively leaks less information on ϕ, would be to publish random generators of $\langle R \rangle$ and $\langle \phi(R) \rangle$. However, it is not clear that this idea would considerably improve the security of the protocol, and we will not pursue it further for coherence with the protocols that will follow.

$\mathcal{A}$		$\mathcal{B}$
Input: A, B, sID		**Input:** B
$m_A, n_A \in_R \mathbb{Z}/\ell_A^{e_A}\mathbb{Z}$		$m_B, n_B \in_R \mathbb{Z}/\ell_B^{e_B}\mathbb{Z}$
$\phi_A := E_0/\langle[m_A]P_A + [n_A]Q_A\rangle$		$\phi_B := E_0/\langle[m_B]P_B + [n_B]Q_B\rangle$

$$\xrightarrow{\substack{A,\text{sID}\\ \phi_A(P_B),\\ \phi_A(Q_B),\\ E_A}}$$

$$\xleftarrow{\substack{B,\text{sID}\\ \phi_B(P_A),\\ \phi_B(Q_A),\\ E_B}}$$

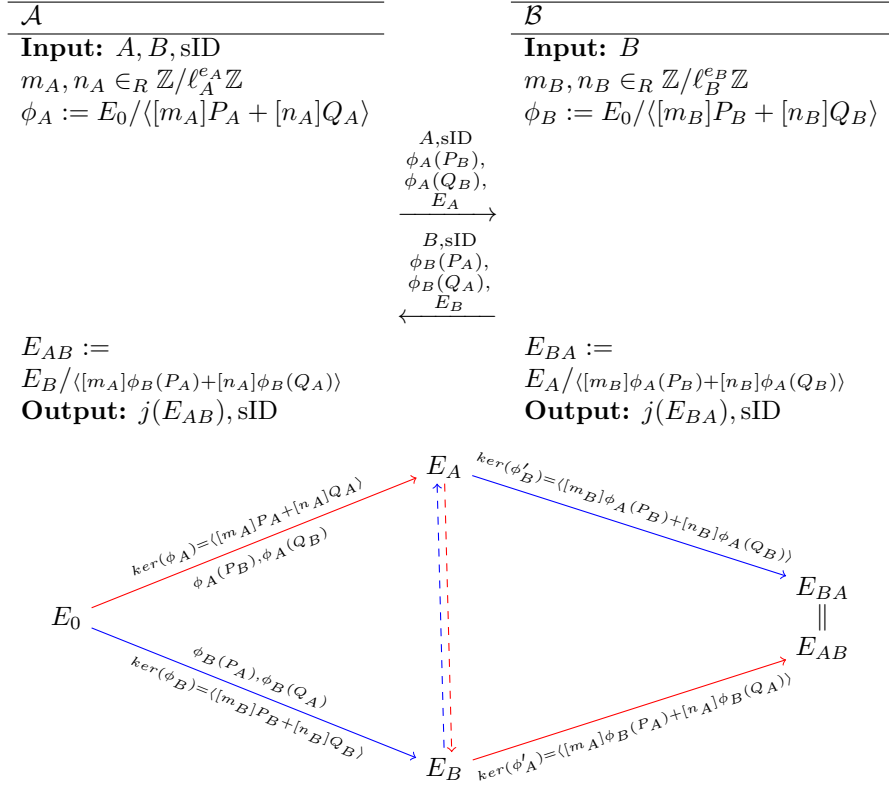| $E_{AB} :=$		$E_{BA} :=$
$E_B/\langle[m_A]\phi_B(P_A)+[n_A]\phi_B(Q_A)\rangle$		$E_A/\langle[m_B]\phi_A(P_B)+[n_B]\phi_A(Q_B)\rangle$
Output: $j(E_{AB}), \text{sID}$		**Output:** $j(E_{BA}), \text{sID}$

FIGURE 1. Key-exchange protocol using isogenies on supersingular curves.

Encryption: Given a public key $(E_A, \phi_A(P_B), \phi_A(Q_B), k)$ and a message $m \in \{0,1\}^w$, choose two random elements $m_B, n_B \in_R \mathbb{Z}/\ell_B^{e_B}\mathbb{Z}$, not both divisible by ℓ_B, and compute

$$h = H_k(j(E_{AB})),$$
$$c = h \oplus m.$$

The ciphertext is $(E_B, \phi_B(P_A), \phi_B(Q_A), c)$.

Decryption: Given a ciphertext $(E_B, \phi_B(P_A), \phi_B(Q_A), c)$ and a private key (m_A, n_A, k), compute the j-invariant $j(E_{AB})$ and set

$$h = H_k(j(E_{AB})),$$
$$m = h \oplus c.$$

The plaintext is m.

4. Algorithmic aspects

We now give specific algorithms to implement the aforementioned steps efficiently.

4.1. **Parameter generation.** For any fixed choice of $\ell_A^{e_A}$ and $\ell_B^{e_B}$, one can easily test random values of f (of any desired cryptographic size) until a value is found for which $p = \ell_A^{e_A}\ell_B^{e_B} \cdot f - 1$ or $p = \ell_A^{e_A}\ell_B^{e_B} \cdot f + 1$ is prime. The prime number theorem in arithmetic progressions (specifically, the effective version of Lagarias and Odlyzko [24]) provides a sufficient lower bound for the density of such primes.

Once the prime $p = \ell_A^{e_A}\ell_B^{e_B} \cdot f \pm 1$ is known, Bröker [6] has shown that it is easy to find a supersingular curve E over $\mathbb{F}_{p^2}$ having cardinality $(p \mp 1)^2 = (\ell_A^{e_A}\ell_B^{e_B} \cdot f)^2$. Starting from E, one can select a random

---

Algorithm 1 Three-point ladder to compute $P + [t]Q$.

Input: t, P, Q;

1: Set $A = 0, B = Q, C = P$;
2: Compute $Q - P$;
3: **for** i decreasing **from** $|t|$ **to** 1 **do**
4:     Let t_i be the i-th bit of t;
5:     **if** $t_i = 0$ **then**
6:         $A = 2A$, $B = \mathrm{dadd}(A, B, Q)$, $C = \mathrm{dadd}(A, C, P)$;
7:     **else**
8:         $A = \mathrm{dadd}(A, B, Q)$, $B = 2B$, $C = \mathrm{dadd}(B, C, Q - P)$;
9:     **end if**
10: **end for**

Output: $C = P + [t]Q$.

---

supersingular curve E_0 over $\mathbb{F}_{p^2}$ by means of random walks on the isogeny graph (cf. Proposition 2.1); alternatively, one can simply take $E_0 = E$. In either case, E_0 has group structure $(\mathbb{Z}/(p \mp 1)\mathbb{Z})^2$. To find a basis for $E_0[\ell_A^{e_A}]$, choose a random point $P \in_R E_0(\mathbb{F}_{p^2})$ and multiply it by $(\ell_B^{e_B} \cdot f)^2$ to obtain a point P' of order dividing $\ell_A^{e_A}$. With high probability, P' will have order exactly $\ell_A^{e_A}$; one can of course check this by multiplying P' by powers of ℓ_A. If the check succeeds, then set $P_A = P'$; otherwise try again with another P. A second point Q_A of order $\ell_A^{e_A}$ can be obtained in the same way. To check whether Q_A is independent of P_A, simply compute the Weil pairing $e(P_A, Q_A)$ in $E[\ell_A^{e_A}]$ and check that the result has order $\ell_A^{e_A}$; as before, this happens with high probability, and if not, just choose another point Q_A. Note that the choice of basis has no effect on the security of the scheme, since one can convert from one basis to another using extended discrete logarithms, which are easy to compute in $E_0[\ell_A^{e_A}]$ by [46].

4.2. **Key exchange and other protocols.** The key exchange is performed in two rounds. In each round Alice and Bob do the following operations on each side:

(1) Compute $\langle R \rangle = \langle [m]P + [n]Q \rangle$ for some points P, Q;
(2) Compute the isogeny $\phi : E \to E/\langle R \rangle$ for a curve E;
(3) In the first round (only), compute $\phi(R)$ and $\phi(S)$ for some points R, S;

where the curve E and the points P, Q, R, S depend on the round and the player, as shown in Figure 1. Similar operations are needed by the other protocols of Section 3. We demonstrate how to implement efficiently each of these steps.

4.2.1. *Computing $\langle [m]P + [n]Q \rangle$.* We first observe that any generator of $\langle [m]P + [n]Q \rangle$ will do. Without loss of generality we can assume that m is invertible modulo the order of the group, in which case $R' = P + [m^{-1}n]Q$ is one such generator. Computing R' by a standard double-and-add approach requires half the effort of computing $[m]P + [n]Q$ naively (see [13, 38, 1], though, for better ways of computing the latter).

However, computing $P + [m^{-1}n]Q$ by double-and-add has one major drawback: it is trivially vulnerable to simple power analysis (SPA) [22]. To avoid SPA, one can use a Montgomery ladder [28] to compute $[m^{-1}n]Q$, and then add P, but this is significantly slower.

Instead, we propose in Algorithm 1 a much more efficient ladder, computing $P + [m^{-1}n]Q$ directly. The idea behind it is simple: at each iteration the registers A, B and C contain respectively the values $[x]Q$, $[x+1]Q$ and $P + [x]Q$, for x equal to the leftmost bits of $m^{-1}n$. The function $\mathrm{dadd}(A, B, C)$ is a *differential addition*: it computes the sum $A + B$ knowing $C = A - B$. Montgomery curves have a very efficient differential addition [28], making the implementation of our ladder as efficient as the naive double and add on twisted Edwards curves (see Section 4.3).

4.2.2. *Computing smooth degree isogenies.* It remains to describe how Alice and Bob can compute and evaluate the isogenies. Let E be an elliptic curve, and let R be a point of order ℓ^e. Our goal is to compute the image curve $E/\langle R \rangle$ and to evaluate the isogeny $\phi : E \to E/\langle R \rangle$ at some points of E.

FIGURE 2. Computational structure of the construction of $\phi = \phi_5 \circ \cdots \circ \phi_0$.

Since the degree of ϕ is smooth, it is best to decompose it as a chain of ℓ-isogenies. Set $E_0 = E$, $R_0 = R$ and, for $0 \le i < e$, let

$$E_{i+1} = E_i/\langle \ell^{e-i-1} R_i \rangle, \qquad \phi_i : E_i \to E_{i+1}, \qquad R_{i+1} = \phi_i(R_i).$$

Then $E/\langle R \rangle = E_e$ and $\phi = \phi_{e-1} \circ \cdots \circ \phi_0$.

The curve E_{i+1} and the isogeny ϕ_i can be computed using Vélu's formulas [48], once the ℓ-torsion subgroup $\langle R_i \rangle$ of E_i is known. This immediately suggests two strategies having quadratic complexity in e, as described in [19].

However, we can do much better. Figure 2 summarizes the computational structure of the problem for $e = 6$. Bullets represent points, with points on the same horizontal level having the same order, and points on the same left diagonal belonging to the same curve. Dashed edges are directed and go from top to bottom; leftward edges represent multiplication by ℓ, and rightward edges represent evaluation of an ℓ-isogeny.

At the beginning of the algorithm, only the point R_0 is known. Our goal is to compute all the points on the bottom line. Indeed, from the knowledge of the point $[\ell^{e-i-1}]R_i$ we can compute the kernel of ϕ_i using $O(\ell)$ point additions. We then apply Vélu's formulas to compute ϕ_i and E_{i+1}. From this point on, we can forget about the number theoretic nature of the problem and purely focus on its combinatorial structure. Lemma 4.2 guarantees that any solution to the combinatorial problem yields a valid strategy to compute ϕ. Before stating it, we first formally rephrase the picture in Figure 2.

Definition 4.1. Let T_n be the portion of the unit triangular equilateral lattice contained between the x-axis, the line $y = \sqrt{3}x$ and the line $y = -\sqrt{3}(x - n + 1)$. We call T_n the *discrete equilateral triangle* (DET) of side n.

An *edge* is any segment of unit length directed downwards to the x-axis connecting two points of T_n. We say an edge is a *left edge* if it has positive slope, a *right edge* otherwise. This definition yields a directed acyclic graph (DAG) structure on T_n.

We equip the points of any directed graph G with the ordering $\to_G$ defined by: $x \to_G y$ if, and only if, there exists a path in G from x to y. We use $\to$ as a shorthand for $\to_{T_n}$. The *leaves* of G are the final points, and the *roots* of G are the initial points. The graph T_n has the n points on the x-axis as leaves and the top-most point as its unique root. We write $|G|$ for the number of leaves of G.

For any two points y, y' of T_n, there exists a most final point x such that $x \to y$ and $x \to y'$. We write $x = y \wedge y'$.

A *strategy* S is a sub-graph of T_n having a unique root. It is *full* if its leaves contain those of T_n; in this case, the root of S is the same as that of T_n. The *fork* of S is the final-most point x of S that is comparable (for the ordering $\to_S$) to all points of S.

Lemma 4.2. *Any full strategy yields a valid algorithm to compute the isogeny $\phi = \phi_{n-1} \circ \cdots \circ \phi_0$ by decorating the DET as shown in Figure 2.*

Proof. Travel the graph in depth-first left-first order. Upon reaching the bottom, apply Vélu's formulas before going right. □

At this point it should be clear that a strategy that passes twice through an interior point does unnecessary computations. Another waste of resources is a strategy having a leaf distinct from the leaves of T_n.

FIGURE 3. Two ill-formed strategies

We define a *well-formed* strategy to be one that has no such useless edge[3]. Figure 4 shows all the seven well-formed full strategies for $n = 4$, the leftmost and the rightmost ones corresponding respectively to the *isogeny-based* and the *multiplication-based* algorithms from [19].

FIGURE 4. The seven well-formed full strategies for $n = 4$.

It is clear that some strategies are computationally better than others. From the figures, we see that the multiplication-based and isogeny-based strategies have a number of edges quadratic in n, whereas the middle stategy in Figure 4 extends to a family of *balanced strategies* with asymptotically $\frac{1}{2\log 2} n \log n$ left and right edges.

We are thus interested in computing the "best" full strategy, according to some measure of computational effort. We first observe that any well-formed strategy has a binary tree topology, obtained by forgetting the internal nodes of out-degree less than two and by keeping the same connectivity structure. Conversely, to any binary tree A with n leaves one can canonically associate a strategy S on T_n as follows: if $n = 1$, then $S = T_1$; else, let S' be the strategy associated to the left branch of A, S'' be the translated to the right by $|S'|$ of the strategy associated to the right branches of A, and r', r'' be their roots. Then $r' \wedge r''$ is the root of T_n and we define $S = rr' \cup rr'' \cup S' \cup S''$, where rr' and rr'' are the (unique) paths from r to r' and r'' in T_n.

For example, in Figure 4 the three middle strategies share the same tree topology, and the middle one is the canonical one.

In our original problem, left edges are multiplications by ℓ, while right edges are ℓ-isogeny evaluations, and these steps have different costs. Thus it makes sense to assign different weight to left and right edges.

We fix a *measure* on T_n, i.e., a pair (p, q) of positive real numbers, where p represents the weight of a left edge (i.e., a point multiplication) and q represents the weight of a right edge (i.e., an isogeny). For any set of edges S of T_n, we write (S) for the sum of the weights of all edges of S.

If $x \to y$, then all paths going from x to y in T_n have the same measure; we write (xy) for this measure. For all x, y, y', y'' such that $x \to y$, $y \to y''$ and $y \to y''$, we have the inequality $(xy) + (yy') + (yy'') \leq (xy') + (xy'')$.

We find optimal strategies in two steps. The first step (Lemma 4.3) shows that, for any measure, only canonical strategies are interesting. Then we determine the optimal full strategy for a given measure (Proposition 4.6).

---

[3]We do not know much about well-formed full strategies. One remarkable fact is that they are in one-to-one correspondence with certain instances of Gelfand-Tsetlin patterns [29].

Lemma 4.3. *The canonical strategy is minimal, with respect to any measure, among all the strategies sharing the same tree topology.*

Lemma 4.3 follows from Lemma 4.4.

Lemma 4.4. *Let S be a strategy with root r. Let $\widehat{S}$ be the canonical strategy associated to S and $\widehat{r}$ be the root of $\widehat{S}$. Then $r \to \widehat{r}$, and $(r\widehat{r}) + (\widehat{S}) \leq (S)$.*

Proof. We prove this by induction on S. Let t be the fork of S. If t has no children in S, then it is the unique leaf of S and the lemma is obviously true. Therefore, we may assume that t has two children in S.

FIGURE 5. Proof of Lemma 4.4.

Let S', S'' be the right and left sub-strategies of S with root t. We apply the induction hypothesis on S' and S'' and use $\widehat{r} = \widehat{r}' \wedge \widehat{r}''$ to derive the following (in)equalities:

(4)
$$
\begin{aligned}
(S') + (S'') + (rt) &= (S) && \text{(definition of (S)),} \\
(t\widehat{r}') + (\widehat{S}') &\leq (S') && \text{(induction hypothesis on $\widehat{S}'$),} \\
(t\widehat{r}'') + (\widehat{S}'') &\leq (S'') && \text{(induction hypothesis on $\widehat{S}''$),} \\
(t\widehat{r}) + (\widehat{r}\widehat{r}') + (\widehat{r}\widehat{r}'') &\leq (t\widehat{r}') + (t\widehat{r}'') && \text{(definition of $\widehat{r}$),} \\
(r\widehat{r}) &= (rt) + (t\widehat{r}) && \text{(path $r \to t \to \widehat{r}$),} \\
(\widehat{S}) &= (\widehat{r}\widehat{r}') + (\widehat{r}\widehat{r}'') + (\widehat{S}') + (\widehat{S}'') && \text{(definition of $\widehat{S}$).}
\end{aligned}
$$

By summing all of these and cancelling all extra terms, we obtain the desired result. $\square$

Now that we have ruled out non-canonical strategies, we proceed to determine the best ones. From this point on, we are only interested in full strategies. Thus, we shall call *optimal* a strategy that is minimal, among all full strategies with n leaves, with respect to a measure (p, q).

In practice, we can compute optimal strategies using the fact that, for a fixed measure, optimality is a local property. For a full canonical strategy $S \neq T_1$, having i leaves to the left of its root, we define its *left branch* as $S \cap T_i$. We also define its *right branch* as $S \cap T'$, where T' is $T_{|S|-i}$ translated by i to the right. Because S is full and canonical, both branches are full canonical strategies too.

Lemma 4.5. *Let S be an optimal strategy. Let S' and S'' be, respectively, its left and right branch. Then S' and S'' are optimal strategies.*

Proof. By Lemma 4.3, we know that S is a canonical strategy. Hence, S' and S'' are well defined. Now, suppose that S' were not optimal. By substituting an optimal strategy for S' inside S, we obtain a strategy with weight lower than (S). The same argument holds for S''. $\square$

Define $C_{p,q}(n)$ to be the cost of the optimal strategies with n leaves. Lemma 4.5 says that

(5) $$ C_{p,q}(n) = \min_{i=1,\ldots,n-1} \big(C_{p,q}(i) + C_{p,q}(n-i) + (n-i)p + iq \big). $$

This equality suggests a dynamic-programming algorithm that, given p and q, computes $C(n)$ in time $O(n^2)$. A straightforward Python implementation computes the optimal strategies for $n = 1024$ in less than one

FIGURE 6. Optimal strategy for $n = 512, p = 4.6, q = 2.8$.

second, indicating that the dynamic-programming approach is very satisfying in practice. Figure 6 shows an optimal strategy for $n = 512, p = 4.6, q = 2.8$.

However, it is also possible to mathematically characterize the optimal strategies. Since the optimal strategies for (p, q) are symmetrical to those for (q, p), we may assume that $p < q$. For simplicity, we will also assume that p and q are integers: the generalization to the case of rational p and q is straightforward.

Let (u_m) be the sequence defined by $u_{p+1} = \cdots = u_{p+q} = 1$ and $u_m = u_{m-p} + u_{m-q}$, and $f(n)$ be the function defined by $f(1) = 0$, $f(n) = f(u_m) + m(n - u_m)$ for all $n \in [u_m, u_{m+1}]$. We check that f satisfies the recurrence relation

$$(6) \qquad f(u_m) = f(u_{m-p}) + f(u_{m-q}) + pu_{m-p} + qu_{m-q}.$$

Proposition 4.6. *A full strategy S is optimal if, and only if, it is canonical, both its left and right branches S' and S'' are optimal, and*

$$u_{m-q} \le |S'| \le u_{m-q+1}, \quad u_{m-p} \le |S''| \le u_{m-p+1},$$

where m is such that $u_m \le |S| < u_{m+1}$. In this case, the weight of S is $f(|S|)$.

For example, in the case where $p = 1$ and $q = 2$, the sequence (u_m) is the sequence of Fibonacci numbers, starting at $u_2 = 1$, $u_3 = 1$, $u_4 = 2...$, and a strategy S with $u_m \le |S| < u_{m+1}$ is optimal if, and only if, the left and right branches S' and S'' verify $u_{m-2} \le |S'| \le u_{m-1}$ and $u_{m-1} \le |S''| \le u_m$.

Proof. We prove this by induction on $n = |S|$. Let $n = n' + n''$ with $n', n'' > 0$ and let $S = S_{n', n''}$ be any canonical full strategy with optimal left and right branches S', S'', such that $|S'| = n'$ and $|S''| = n''$. Then any optimal strategy is one of the $S_{n', n''}$. Therefore, we find them by looking at the sign of $\delta_n = (S_{n'+1, n''-1}) - (S_{n', n''})$.

Let m', m'' be such that $u_{m'} \le n' < u_{m'+1}$ and $u_{m''} \le n'' < u_{m''+1}$. Using the induction hypothesis on S' and S'', we have $(S_{n', n''}) = f(n') + f(n'') + n'q + n''p$, and therefore

$$(7) \qquad \delta_n = \begin{cases} (m' + q) - (m'' + p), & \text{if } n'' \ge u_{m''} + 1, \\ (m' + q) - (m'' + p) + 1, & \text{if } n'' = u_{m''}. \end{cases}$$

We find that $\delta_n \le 0$ exactly when $n' < u_{m-q+1}$ and $n'' > u_{m-p}$, and $\delta_n \ge 0$ exactly when $n' \ge u_{m-q}$ and $n'' \le u_{m-p+1}$. The optimality condition follows from this. It remains only to check that $(S) = f(|S|)$: this results from relation (6). $\qquad\square$

In particular, there are in general several optimal strategies of a given size. Moreover, with z being the (unique) root in $[0, 1]$ of the equation $z^p + z^q - 1 = 0$, this gives the asymptotic equivalents

$$(8) \qquad\qquad |S'| \quad \sim \quad z^q |S|,$$

$$(9) \qquad\qquad C_{p,q}(n) = (S) \quad \sim \quad -\frac{1}{\log z}\, n \log n.$$

The relative cost of an optimal stategy is therefore asymptotically equivalent to $-\frac{2 \log 2}{\log(z^{p+q})}$ times the cost of the balanced strategy. In the case where $(p, q) = (4.6, 2.8)$, as in Table 1 with $\ell = 2$, this gives an improvement of 2.1% over the balanced strategy.

4.3. **Choice of the models.** At each phase of the key generation it is important to use models for elliptic curves that offer the fastest formulas for doubling, addition, isogeny computation and evaluation, etc.

To measure efficiency, we count the number of elementary operations in $\mathbb{F}_{p^2}$: we write I, M, S for the costs of one inversion, multiplication and squaring respectively, and we make the assumption $S \leq M \leq I$. We neglect additions, subtractions and comparisons. We refer to the Explicit Formulas Database (EFD) [4] for operation counts of elliptic point addition, doubling, etc. in various models and coordinate systems. Contrary to the convention taken in the EFD, we count multiplications by constants (other than small integers) as ordinary multiplications.

Any curve used in our cryptosystem has group structure $(\mathbb{Z}/(p \mp 1)\mathbb{Z})^2$. Hence either the curve or its twist has a point of order 4. Consequently, it is isomorphic to a twisted Edwards curve and to a Montgomery curve [3].

Twisted Edwards curves [3] have equation

$$(10) \qquad\qquad E_{a,d} \ : \ ax^2 + y^2 = 1 + dx^2 y^2.$$

They have many interesting properties, but what interests us most is their very efficient addition and doubling formulas. Using projective coordinates, one point addition costs $12M + 1S$ and one point doubling costs $4M + 4S$. When one of the points is scaled to have Z-coordinate equal to 1, these costs drop to $11M + 1S$ and $3M + 4S$ respectively.

Montgomery curves [28] have equation

$$(11) \qquad\qquad M_{B,A} \ : \ By^2 = x^3 + Ax^2 + x.$$

They have very efficient arithmetic on their Kummer line, i.e., by representing points by the coordinates $(X : Z)$ where $x = X/Z$. Using this representation, a point can be doubled using $3M + 2S$, or $2M + 2S$ when it is scaled to have Z-coordinate equal to 1.

The Kummer line identifies P with $-P$; thus it is not possible to add two distinct points. However it is still possible to perform what is called *differential addition*, i.e., adding points P and Q for which the difference $P - Q$ is known. One differential addition can be computed using $4M + 2S$, or $3M + 2S$ when the difference $P - Q$ is scaled to have Z-coordinate equal to 1.

By using doublings and differential additions, it is possible to compute any scalar multiple of a point using a Montgomery ladder [28]. Also observe that since P and $-P$ generate the same subgroup, isogenies can be defined and evaluated correctly on the Kummer line. We shall give formulas for this operation later.

It is shown in [3] that the twisted Edwards curve $E_{a,d}$ is birationally equivalent to the Montgomery curve $M_{A,B}$ where

$$(12) \qquad\qquad A = \frac{2(a + d)}{a - d}, \qquad B = \frac{4}{a - d},$$

and where the transformation is given by the following maps (in affine coordinates)

$$(13) \qquad
\begin{aligned}
\psi \ &: \ E_{a,d} \to M_{A,B}, & \psi^{-1} \ &: \ M_{A,B} \to E_{a,d}, \\
(x, y) &\mapsto \left(\frac{1+y}{1-y}, \frac{1+y}{(1-y)x} \right), & (x, y) &\mapsto \left(\frac{x}{y}, \frac{x-1}{x+1} \right).
\end{aligned}$$

Hence, after the models $E_{a,d}$ and $M_{A,B}$ are computed, points in projective coordinates can be moved from one model to the other at a cost of a few multiplications (and no inversions).

Now we detail the cost of each step of the key generation algorithm.

4.3.1. *Computing $\langle [m]P + [n]Q \rangle$.* We have already mentioned two algorithms to compute the value of $P + [m^{-1}n]Q$. One solution is to express the points P and Q in projective Edwards coordinates and perform a double and add followed by an addition. If we scale Q to have Z-coordinate equal to 1, computing $[m^{-1}n]Q$ costs $9.5M + 4.5S$ per bit on average.

Because addition by P is performed last, this approach cannot be used with points on the Montgomery curve in Kummer coordinates, but we can use the ladder given in Figure 1 instead. We first compute $P - Q$ using point addition in full projective coordinates (either by using the standard chord-and-tangent law, or by using the equivalence with twisted Edwards curves). Then we scale P, Q and $P - Q$ to have Z-coordinate equal to 1. We can now discard the Y coordinate and work on the Kummer line. At each iteration we perform one doubling and two differential additions (with one of the scaled points P, Q and $P - Q$). The total cost of one iteration is thus $9M + 6S$.

In general the cost of one squaring is close to the cost of one multiplication. Thus the ladder algorithm is slightly slower than double-and-add. However the advantages of the ladder are SPA resistance and an implementation simplified by not using Edwards coordinates at all for key generation.

4.3.2. *Isogenies of Montgomery curves.* The literature on efficient formulas for evaluating small degree isogenies is much less extensive than for point multiplication. We now give explicit formulas for isogenies of Montgomery curves, and optimize the degree 2 and 3 cases. Our goal is to obtain the most efficient formulas for isogeny evaluation, and thus we seek to avoid inversions and square root computations as much as possible.

Let E be the Montgomery curve defined by Eq. (11). It has a point of order two $P_2 = (0,0)$, and a point of order four $P_4 = (1, \sqrt{(A+2)/B})$—eventualy defined over a quadratic extension—such that $[2]P_4 = P_2$. Montgomery curves have twists of the form $\tilde{y} = \sqrt{c}y$; these are isomorphisms when c is a square. The change of coordinates $\tilde{x} = x/B, \tilde{y} = y/B$ brings the curve E to the Weierstrass form

$$(14) \qquad \tilde{y}^2 = \tilde{x}^3 + \frac{A}{B}\tilde{x}^2 + \frac{1}{B^2}\tilde{x},$$

and the point P_4 to $P'_4 = (1/B, \ldots)$. Conversely, given a Weierstrass curve with equation $\tilde{y}^2 = \tilde{x}^3 + a\tilde{x}^2 + b\tilde{x}$, together with a point $P_4 = (1/\beta, \ldots)$—with its ordinate possibly lying in a quadratic extension—such that $[2]P_4 = (0,0)$, the change of variables $\tilde{x} = x/\beta, \tilde{y} = y/\beta$ brings the curve to the Montgomery form $\beta y^2 = x^3 + a\beta x^2 + x$.

Given a point $P \neq \pm P_4$ of order 4, we will need to compute the isomorphism of Montgomery curves that brings $[2]P$ in $(0,0)$ and P in $(1,\ldots)$. Let X be the abscissa of P and X_0 the abscissa of $[2]P$; by a straightforward calculation, we find that this isomorphism is given by the map

$$(15) \qquad \begin{aligned} \iota : E &\to E', \\ (x, y) &\mapsto \left(\frac{x - X_0}{X - X_0}, \frac{y}{X - X_0} \right), \end{aligned}$$

and the new curve has equation

$$E' \; : \; \frac{B}{X - X_0}y^2 = x^3 + \frac{3X_0 + A}{X - X_0}x^2 + x.$$

By precomputing $1/(X - X_0)$, and sharing common subexpressions, the above map can be evaluated on a point in full projective coordinates using $3M$ operations, or $2M$ using Kummer coordinates.

Let G be a subgroup of the Montgomery curve E of odd cardinality ℓ and let h be the degree $(\ell - 1)/2$ polynomial vanishing on the abscissas of G. With a twist $y = \tilde{y}/\sqrt{B}$, we can put E in the form

$$E \; : \; \tilde{y}^2 = \tilde{x}^3 + A\tilde{x}^2 + \tilde{x},$$

and this does not change the abscissas of G nor the polynomial h. Now, composing the twist with Vélu's formulas gives an isogeny

$$\phi : E \to E/G,$$

$$(x, y) \mapsto \left(\frac{g(x)}{h(x)^2}, y\sqrt{B} \left(\frac{g(x)}{h(x)^2} \right)' \right).$$

We now need to express E/G in Montgomery form. Because ℓ is odd, the point $(0,0)$ of E is sent to a point of order two in E/G, and the change of variables $\tilde{x} = x - g(0)/h(0)^2$ brings this point to $(0,0)$.

Now, $\phi(P_4)$ is a point of order four lying above $(0,0)$ (possibly in a quadratic extension). Its abscissa is rational and is given by

(16)
$$\frac{1}{\beta} = \frac{g(1)}{h(1)^2} - \frac{g(0)}{h(0)^2},$$

so we further apply the change of variables $\tilde{x} = \bar{x}/\beta, \tilde{y} = \bar{x}/\beta$ to obtain a Montgomery curve. Finally, we have to twist back the model in order to obtain a curve isogenous over the base field: the twist $\bar{y} = y\sqrt{B}$ cancels with the first one and leaves us with square-root-free formulas.

Given h as input, the cost of evaluating the whole path is $O(\ell)$ operations in the base field, using the formula in [5, Proposition 4.1] to evaluate g/h.

Let now P be a 3-torsion point, let $G = \{0, P, -P\}$ and let X be the abscissa of P (and $-P$). If we specialize the formula to the case $\ell = 3$, we have $h = x - X$ and

(17)
$$\phi : E \to E/G,$$
$$(x, y) \mapsto \left(\frac{x \left(x - \frac{1}{X} \right)^2}{(x - X)^2} X^2, \ y \frac{\left(x - \frac{1}{X} \right) \left(\left(x - \frac{1}{X} \right)(x - X) + 2x \left(\frac{1}{X} - X \right) \right)}{(x - X)^3} X^2 \right),$$

and the curve E/G has equation

$$E/G \ : \ BX^2 y^2 = x^3 + \left(A + \frac{6}{X} - 6X \right) X^2 x^2 + x.$$

By precomputing X and X^2, and sharing common subexpressions, the above isogeny can be evaluated on a point in full projective coordinates using $11M + 2S$ operations, or $4M + 2S$ using Kummer coordinates.

When ℓ is even, things get more complicated. Recall that $P_2 = (0,0)$ and $P_4 = (1, \sqrt{(A+2)/B})$. The isogeny of degree 2 vanishing on P_2 and mapping P_4 to $(0,0)$ is readily seen as being

(18)
$$F \ : \ By^2 = x^3 + (A+6)x^2 + 4(2+A)x,$$
$$\phi : E \to F,$$

(19)
$$(x, y) \mapsto \left(\frac{(x-1)^2}{x}, y \left(1 - \frac{1}{x^2} \right) \right).$$

It is not immediately evident how to put F in Montgomery form without computing square roots. If P_8 is a point satisfying $[2]P_8 = P_4$, then $\phi(P_8) = (2\sqrt{2+A}, \ldots)$, and F can be put in the form

$$\frac{B}{2\sqrt{2+A}} y^2 = x^3 + \frac{A+6}{2\sqrt{2+A}} x^2 + x,$$

with the point $(1, \ldots)$ being the image of P_8. Any other isogeny of degree 2 can be treated by applying Eq. (15) to move the generator of the kernel in $(0,0)$.

By precomputing $1/\phi(P_8)$, and sharing common subexpressions, the above isogeny can be evaluated on a point in full projective coordinates using $5M + 3S$ operations, or $2M + S$ using Kummer coordinates. Unfortunately, this formula takes no square roots only if an 8-torsion point above P_2 is known.

Alternatively, ϕ and F being as before, we consider the isogeny $\psi : F \to F/\langle(0,0)\rangle$ given by

$$(20) \qquad G \ : \ \frac{B}{2-A}y^2 = x^3 - 2\frac{A+6}{2-A}x^2 + x,$$

$$\psi : F \to G,$$

$$(21) \qquad (x,y) \mapsto \left(\frac{1}{2-A}\frac{(x+4)(x+(A+2))}{x}, \frac{y}{2-A}\left(1 - \frac{4(2+A)}{x^2}\right)\right).$$

Then $\phi_4 = \psi \circ \phi$ is an isogeny of degree four $\phi_4 \colon E \to E/\langle P_4\rangle$, and the point $\left(1, \sqrt{-4(A+2)/B}\right)$ of G generates the kernel of the dual isogeny $\hat{\phi}_4$. Any other isogeny of degree 4 can be treated by applying Eq. (15) to move the kernel point to $(1,\dots)$.

By precomputing $1/(2-A)$, and sharing common subexpressions, the isogeny ϕ_4 can be evaluated on a point in full projective coordinates using $10M + 4S$ operations, or $4M + S$ using Kummer coordinates.

Unfortunately, this formula cannot be applied twice to obtain a degree 16 cyclic isogeny: indeed, a double application yields the multiplication-by-4 isogeny. If a chain of degree 4 isogenies is wanted, as for the algorithm in Subsection 4.2.2, this formula must be combined with the isomorphism in Eq. (15).

Isogenies of composite smooth degree are computed by composing small degree isogenies as discussed in Subsection 4.2.2. We focus on the fine tuning of this algorithm when the small isogenies have degrees $2, 3$ and 4.

It is important to remark that, after a generator R of the kernel of the isogeny has been computed, the algorithm does not use its ordinate at all: indeed, all the previous formulas for small degree isogenies only use the abscissas of the kernel points. Hence, we can throw away the ordinate of R altogether and only use scalar multiplication and isogeny evaluation formulas for points in Kummer coordinates.

As usual, some more care must be taken when computing cyclic isogenies of degree 2^e. In principle, one could compose either degree 2 (Eq. 18) or degree 4 (Eq. 20) isogenies; however we have already pointed out some caveats:

- Both equations require the kernel point to be moved to some specific coordinates. This can be achieved using the isomorphism in Eq. (15).
- After the first change of variables, Eq. (18) can be repeatedly chained with itself; however, from a point of order 2^e only $e - 2$ degree 2 isogenies can be computed this way, because the formula requires the knowledge of a point of order 8.
- Eq. (20) cannot be directly chained with itself to compute a cyclic isogeny of degree 4^e. It must be composed, instead, with Eq. (15) at each step.

Having this in mind, two obvious strategies to compute degree 2^e isogenies are:

(1) Use degree 2 isogenies as much as possible; use one degree 4 isogenies for the last two steps.
(2) Use one degree 2 isogeny if e is odd, then use only degree 4 isogenies composed with isomorphisms.

We argue that, given the current parameters, both approaches are likely to take roughly the same computational effort. We only have used the first one in our tests.

Finally, it should be noted that both approaches, if implemented as described above, leak two bits of security. Indeed, the point $(1,\dots)$ of the image curve is a point of order 4 in the kernel of the dual of the computed isogeny (the secret). It is easy to mask this by taking a random change of coordinates; however, we found this more costly than simply adding two bits of security and letting the information leak.

Table 1 summarizes the costs of the isogeny evaluation and scalar multiplication formulas in light of these remarks. Because squaring in $\mathbb{F}_{p^2}$ is faster than multiplication, we also report the cost obtained by taking $S = 0.8M$ (a factor that roughly approximates the fact that squaring requires 2 multiplications in $\mathbb{F}_p$ instead of 3).

Table 2 compares the total cost of isogeny evaluations and scalar multiplications in a balanced and an optimal strategy for $\ell = 2, 3, 4$, based on the costs given in Table 1, at the classical 256-bit security level (see Section 5 for our complexity assumptions). We make the following observations, backed up by the benchmarks in Section 7:

TABLE 1. Comparative costs for multiplication and isogeny evaluation in projective Kummer coordinates, in number of multiplications and squarings, and assuming $S = 0.8M$.

ℓ	2	3	4
Isogeny	$2M + S$	$4M + 2S$	$6M + S$
	2.8	5.6	6.8
Multiplication	$3M + 2S$	$7M + 4S$	$6M + 4S$
	4.6	10.2	9.2

TABLE 2. Comparative costs of the balanced and the optimal strategies for computing a degree 2^{514} ($\ell = 2, 4$) or 3^{323} ($\ell = 3$) isogeny, assuming $S = 0.8M$.

	optimal strategy			balanced strategy		
ℓ	2	3	4	2	3	4
Isogenies	2741	1610	1166	2323	1430	1033
Multiplications	1995	1151	921	2307	1288	1025
Total cost	16851.8	20756.2	16402.0	17116.6	21145.6	16454.4

- There may be a small advantage (less than 2%) in using isogenies of degree 4 instead of 2;
- The gain in using an optimal strategy instead of a balanced one is consistent with the predictions of Subsection 4.2.2;
- The difference between 2^e and 3^e isogenies is more significant (about 20%), suggesting that degree 2^e isogenies may be preferable for constrained devices.

5. Complexity assumptions

As before, let p be a prime of the form $\ell_A^{e_A}\ell_B^{e_B} \cdot f \pm 1$, and fix a supersingular curve E_0 over $\mathbb{F}_{p^2}$ together with bases $\{P_A, Q_A\}$ and $\{P_B, Q_B\}$ of $E_0[\ell_A^{e_A}]$ and $E_0[\ell_B^{e_B}]$ respectively. In analogy with the case of isogenies over ordinary elliptic curves, we define the following computational problems, adapted for the supersingular case:

Problem 5.1 (Decisional Supersingular Isogeny (DSSI) problem). Let E_A be another supersingular curve defined over $\mathbb{F}_{p^2}$. Decide whether E_A is $\ell_A^{e_A}$-isogenous to E_0.

Problem 5.2 (Computational Supersingular Isogeny (CSSI) problem). Let $\phi_A \colon E_0 \to E_A$ be an isogeny whose kernel is $\langle [m_A]P_A + [n_A]Q_A \rangle$, where m_A and n_A are chosen at random from $\mathbb{Z}/\ell_A^{e_A}\mathbb{Z}$ and not both divisible by ℓ_A. Given E_A and the values $\phi_A(P_B)$, $\phi_A(Q_B)$, find a generator R_A of $\langle [m_A]P_A + [n_A]Q_A \rangle$.

We remark that given a generator $R_A = [m_A]P_A + [n_A]Q_A$, it is easy to solve for (m_A, n_A), since E_0 has smooth order and thus extended discrete logarithms are easy in E_0 [46].

Problem 5.3 (Supersingular Computational Diffie-Hellman (SSCDH) problem). Let $\phi_A \colon E_0 \to E_A$ be an isogeny whose kernel is equal to $\langle [m_A]P_A + [n_A]Q_A \rangle$, and let $\phi_B \colon E_0 \to E_B$ be an isogeny whose kernel is $\langle [m_B]P_B + [n_B]Q_B \rangle$, where m_A, n_A (respectively m_B, n_B) are chosen at random from $\mathbb{Z}/\ell_A^{e_A}\mathbb{Z}$ (respectively $\mathbb{Z}/\ell_B^{e_B}\mathbb{Z}$) and not both divisible by ℓ_A (respectively ℓ_B). Given the curves E_A, E_B and the points $\phi_A(P_B)$, $\phi_A(Q_B)$, $\phi_B(P_A)$, $\phi_B(Q_A)$, find the j-invariant of $E_0/\langle [m_A]P_A + [n_A]Q_A, [m_B]P_B + [n_B]Q_B \rangle$.

Problem 5.4 (Supersingular Decision Diffie-Hellman (SSDDH) problem). Given a tuple sampled with probability 1/2 from one of the following two distributions:

- $(E_A, E_B, \phi_A(P_B), \phi_A(Q_B), \phi_B(P_A), \phi_B(Q_A), E_{AB})$, where E_A, E_B, $\phi_A(P_B)$, $\phi_A(Q_B)$, $\phi_B(P_A)$, and $\phi_B(Q_A)$ are as in the SSCDH problem and

$$E_{AB} \cong E_0/\langle [m_A]P_A + [n_A]Q_A, [m_B]P_B + [n_B]Q_B \rangle,$$

- $(E_A, E_B, \phi_A(P_B), \phi_A(Q_B), \phi_B(P_A), \phi_B(Q_A), E_C)$, where E_A, E_B, $\phi_A(P_B)$, $\phi_A(Q_B)$, $\phi_B(P_A)$, and $\phi_B(Q_A)$ are as in the SSCDH problem and

$$E_C \cong E_0 / \langle [m_A']P_A + [n_A']Q_A, [m_B']P_B + [n_B']Q_B \rangle,$$

where m_A', n_A' (respectively m_B', n_B') are chosen at random from $\mathbb{Z}/\ell_A^{e_A}\mathbb{Z}$ (respectively $\mathbb{Z}/\ell_B^{e_B}\mathbb{Z}$) and not both divisible by ℓ_A (respectively ℓ_B),

determine from which distribution the tuple is sampled.

The ordinary case analogue of the following problem is trivially solvable in polynomial time. Its supposed difficulty in the supersingular case is at the hearth of the security of our identification scheme.

Problem 5.5 (Decisional Supersingular Product (DSSP) problem)**.** Given an isogeny $\phi : E_0 \to E_3$ of degree $\ell_A^{e_A}$ and a tuple sampled with probability $1/2$ from one of the following two distributions:

- (E_1, E_2, ϕ'), where the product $E_1 \times E_2$ is chosen at random among those $\ell_B^{e_B}$-isogenous to $E_0 \times E_3$, and where $\phi' : E_1 \to E_2$ is an isogeny of degree $\ell_A^{e_A}$, and
- (E_1, E_2, ϕ'), where E_1 is chosen at random among the curves having the same cardinality as E_0, and $\phi' : E_1 \to E_2$ is a random isogeny of degree $\ell_A^{e_A}$,

determine from which distribution the tuple is sampled.

We conjecture that these problems are computationally infeasible, in the sense that for any polynomial-time solver algorithm, the advantage of the algorithm is a negligible function of the security parameter $\log p$. The resulting security assumptions are referred to as the DSSI assumption, CSSI assumption, etc.

5.1. **Hardness of the underlying assumptions.** Given a CSSI (respectively, SSCDH) solver, it is trivial to solve SSCDH (respectively, SSDDH). It is also trivial to solve SSDDH given a DSSI solver. There are no known reductions in the other direction, and given that the corresponding question of equivalence for discrete logarithms and Diffie-Hellman has not yet been completely resolved in all cases, it is reasonable to assume that the question of equivalence of CSSI, SSCDH, and SSDDH is at least hard to resolve. For the purposes of this discussion, we will presume that DSSI and CSSI are equivalent to SSDDH. Concerning DSSP, there is an evident reduction to DSSI. However, it is reasonable to assume that DSSP is easier than the latter.

In the context of cryptography, the problem of computing an isogeny between isogenous supersingular curves was first considered by Galbraith [15] in 1999. The first published cryptographic primitive based on supersingular isogeny graphs is the hash function proposal of Charles et al. [9], which remains unbroken to date (the cryptanalysis of [30] applies only to the LPS graph-based hash function from [9], and not to the supersingular isogeny graph-based hash functions). The fastest known algorithm for finding isogenies between supersingular curves in general takes $O(\sqrt{p}\log^2 p)$ time [9, §5.3.1]; however our problem is less general because the degree of the isogeny is known in advance and is smooth. In addition, the distribution of isogenous curves obtained from taking kernels of the form $\langle [m_A]P_A + [n_A]Q_A \rangle$ is not quite uniform: a simple calculation against Proposition 2.1 indicates that a sequence of e_A isogenies of degree ℓ_A falls short of the length needed to ensure uniform mixing, regardless of the value of p. Since we are the first to propose using isogenies of this type, there is no existing literature addressing the security of the isogenies of the special form that we propose.

There are easy exponential attacks against DSSI and CSSI that improve upon exhaustive search. To find an isogeny of degree $\ell_A^{e_A}$ between E and E_A, an attacker builds two trees of all curves isogenous to E (respectively, E_A) via isogenies of degree $\ell_A^{e_A/2}$. Once the trees are built, the attacker tries to find a curve lying in both trees. Since the degree of the isogeny ϕ_A is $\sim \sqrt{p}$ (much shorter than the size of the isogeny graph), it is unlikely that there will be more than one isogeny path—and thus more than one match—from E to E_A. Given two functions $f : A \to C$ and $g : B \to C$ with domain of equal size, finding a pair (a, b) such that $f(a) = g(b)$ is known as the *claw problem* in complexity theory. The claw problem can obviously be solved in $O(|A| + |B|)$ time and $O(|A|)$ space on a classical computer by building a hash table holding $f(a)$ for any $a \in A$ and looking for hits for $g(b)$ where $b \in B$. This gives a $O(\ell_A^{e_A/2}) = O(\sqrt[4]{p})$ classical attack

against our cryptosystems. With a quantum computer, one can do better using the algorithm in [44], which has complexity $O(\sqrt[3]{|A||B|})$, thus giving an $O(\ell_A^{e_A/3}) = O(\sqrt[6]{p})$ quantum attack against our cryptosystems. These complexities are optimal for a black-box claw attack [49].

We consider the question of whether the auxiliary data points $\phi_A(P_B)$ and $\phi_A(Q_B)$ might assist an adversary in determining ϕ_A. Since (P_B, Q_B) forms a basis for $E_0[\ell_B^{e_B}]$, the values $\phi_A(P_B)$ and $\phi_A(Q_B)$ allow the adversary to compute ϕ_A on all of $E_0[\ell_B^{e_B}]$. This is because any element of $E_0[\ell_B^{e_B}]$ is a (known) linear combination of P_B and Q_B (known since extended discrete logarithms are easy [46]). However, there does not appear to be any way to use this capability to determine ϕ_A. Even on a quantum computer, where finding abelian hidden subgroups is easy, there is no hidden subgroup to find, since ϕ_A has degree $\ell_A^{e_A}$, and thus does not annihilate any point in $E_0[\ell_B^{e_B}]$ other than the identity. Of course, if one could evaluate ϕ_A on arbitrary points of $E_0[\ell_A^{e_A}]$, then a quantum computer could easily break the scheme, and indeed in this case the scheme is also easily broken classically by using a few calls to the oracle to compute a generator of the kernel of the dual isogeny $\hat{\phi}_A$. However, it does not seem possible to translate the values of ϕ_A on $E_0[\ell_B^{e_B}]$ into values on $E_0[\ell_A^{e_A}]$.

For both ordinary and supersingular curves, there is a natural bijection between isogenies (up to isomorphism) and (left) ideals in the endomorphism ring. In the ordinary case the endomorphism ring is commutative, and ideal classes form a finite abelian group. This property has been used by Childs et al. [10] to solve the ordinary analogue of CSSI in quantum subexponential time. It is natural to ask whether their algorithm can be adapted to the supersingular setting, but, in this case, the endomorphism ring is a maximal order in a noncommutative quaternion algebra, and the left ideal classes do not form a group at all (they form a groupoid, though). Since the algorithm of Childs et al. depends crucially on the properties of abelian groups, we believe that no reasonable variant of this strategy would apply to supersingular curves.

The same correspondence between isogenies and ideals can be applied to DSSP. Indeed, deciding DSSP amounts to decide whether the ideals S, S' associated to ϕ, ϕ' are conjugated, i.e., whether there exists a left ideal $R \in \text{End}(E_0)$ such that $S = RS'R^{-1}$. Although it can be hoped that deciding conjugacy of ideal classes in the quaternion algebra $\mathbb{Q}_{p,\infty}$ is feasible, we are still faced with the problem that the best known algorithms to compute the endomorphism rings of supersingular curves are exponential in $\log p$ [23, 8, 2]. Hence, we deem DSSP secure given the current knowledge.

The fact that it is possible to obtain a zero-knowledge identification scheme from CSSI comes as no surprise, since it is well known that a zero-knowledge protocol can be obtained from any problem in NP [18]. Nevertheless, the generic construction is not very efficient, and many efforts have been made to obtain efficient *ad-hoc* schemes from NP-complete problems [36, 40, 41, 33]. While the security of most of these schemes is based on two solid assumptions, namely that P $\neq$ NP and that *secure commitment schemes* exist, our identification scheme stands on a much weaker ground: the CSSI and DSSP problems. As performances go, it is reasonable to assume that our scheme will be some orders of magnitude slower than the best zero-knowledge protocols. We can thus conclude that our scheme is of a purely theoretical and pedagogical interest. Yet it is remarkable that an efficient identification scheme based on graphs of supersingular isogenies simply exists, while the analogous construction for ordinary curves is trivially broken and no other identification scheme is currently known to work in that case [43].

6. SECURITY PROOFS

In this section we state formal security reductions relating the security of our protocols to the hardness of the appropriate underlying isogeny computation problem. The security proofs are routine, but tedious, and contain little original contribution on our part. For this reason, we only prove a representative selection of our theorems.

The statements of the theorems are as follows:

Theorem 6.1. *Under the SSDDH assumption, the key-agreement protocol of Section 3.2 is session-key secure in the authenticated-links adversarial model of Canetti and Krawczyk [7].*

Theorem 6.2. *If the SSDDH assumption holds, and the hash function family $\mathcal{H}$ is entropy-smoothing, then the public-key cryptosystem of Section 3.3 is IND-CPA.*

Theorem 6.3. *Under the CSSI and DSSP assumptions, the identification scheme of Section 3.1 is zero-knowledge.*

6.1. **Proof of Theorem 6.1.** The proofs of Theorems 6.1 and 6.2 are easily adapted from the corresponding proofs given by Stolbunov [42]. As an illustration of the proof techniques, we provide a proof of Theorem 6.1.

We recall the definition of session-key security in the authenticated-links adversarial model of Canetti and Krawczyk [7]. We consider a finite set of *parties* $P_1, P_2, \ldots, P_n$ modeled by probabilistic Turing machines. The adversary $\mathcal{I}$, also modeled by a probabilistic Turing machine, controls all communication, with the exception that the adversary cannot inject or modify messages (except for messages from corrupted parties or sessions), and any message may be delivered at most once. Parties give outgoing messages to the adversary, who has control over their delivery via the Send query. Parties are activated by Send queries, so the adversary has control over the creation of protocol sessions, which take place within each party. Two sessions s and s' are *matching* if the outgoing messages of one are the incoming messages of the other, and vice versa.

We allow the adversary access to the queries SessionStateReveal, SessionKeyReveal, and Corrupt. The SessionStateReveal($\mathfrak{s}$) query allows the adversary to obtain the contents of the session state, including any secret information. The query is noted and $\mathfrak{s}$ produces no further output. The SessionKeyReveal($\mathfrak{s}$) query enables the adversary to obtain the session key for the specified session $\mathfrak{s}$, so long as $\mathfrak{s}$ holds a session key. The Corrupt(P_i) query allows the adversary to take over the party P_i, i.e., the adversary has access to all information in P_i's memory, including long-lived keys and any session-specific information still stored. A corrupted party produces no further output. We say a session $\mathfrak{s}$ with owner P_i is *locally exposed* if the adversary has issued SessionKeyReveal($\mathfrak{s}$), SessionStateReveal($\mathfrak{s}$), or Corrupt(P_i) before $\mathfrak{s}$ is expired. We say $\mathfrak{s}$ is *exposed* if $\mathfrak{s}$ or its matching session have been locally exposed, and otherwise we say $\mathfrak{s}$ is *fresh*.

We allow the adversary $\mathcal{I}$ a single Test($\mathfrak{s}$) query, which can be issued at any stage to a completed, fresh, unexpired session $\mathfrak{s}$. A bit b is then picked at random. If $b = 0$, the test oracle reveals the session key, and if $b = 1$, it generates a random value in the key space. $\mathcal{I}$ can then continue to issue queries as desired, with the exception that it cannot expose the test session. At any point, the adversary can try to guess b. Let GoodGuess$^{\mathcal{I}}(k)$ be the event that $\mathcal{I}$ correctly guesses b, and define

$$\text{Advantage}^{\mathcal{I}}(k) = \max\left\{0, \left|\Pr[\text{GoodGuess}^{\mathcal{I}}(k)] - \frac{1}{2}\right|\right\},$$

where k is a security parameter.

The definition of security is as follows:

Definition 6.4. A key exchange protocol Π in security parameter k is said to be *session-key secure* in the authenticated-links adversarial model of Canetti and Krawczyk if for any polynomial-time adversary $\mathcal{I}$,

    (1) If two uncorrupted parties have completed matching sessions, these sessions produce the same key as output;

    (2) Advantage$^{\mathcal{I}}(k)$ is negligible.

Proof of Theorem 6.1. The proof is based on the proof given by Canetti and Krawczyk [7, §5.1] for two-party Diffie-Hellman over $\mathbb{Z}_q^*$. A similar strategy was used by Stolbunov [42] in the case of ordinary elliptic curves.

It has been shown in Section 3 that two uncorrupted parties in matching sessions output the same session key, and thus the first part of Definition 6.4 is satisfied. To show that the second part of the definition is satisfied, assume that there is a polynomial-time adversary $\mathcal{I}$ with a non-negligible advantage ε. We claim that Algorithm 2 forms a polynomial-time distinguisher for SSDDH having non-negligible advantage.

To prove the claim, we must show that Algorithm 2 has non-negligible advantage (it is clear that it runs in polynomial time). We consider separately the cases where the r-th session is (respectively, is not) chosen by $\mathcal{I}$ as the test session. If the r-th session is not the test session, then Algorithm 2 outputs a random bit, and thus its advantage in solving the SSDDH is 0. If the r-th session is the test session, then $\mathcal{I}$ will succeed with advantage ε, since the simulated protocol provided to $\mathcal{I}$ is indistinguishable from the real protocol. Since the latter case occurs with probability $1/k$, the overall advantage of the SSDDH distinguisher is ε/k, which is non-negligible. □

Algorithm 2 SSDDH distinguisher

Input: $E_A, E_B, \phi_A(P_B), \phi_A(Q_B), \phi_B(P_A), \phi_B(Q_A), E$

1: $r \xleftarrow{R} \{1, \ldots, k\}$, where k is an upper bound on the number of sessions activated by $\mathcal{I}$ in any interaction.
2: Invoke $\mathcal{I}$ and simulate the protocol to $\mathcal{I}$, except for the r-th activated protocol session.
3: For the r-th session, let Alice send A, i, E_A, $\phi_A(P_B)$, $\phi_A(Q_B)$ to Bob, and let Bob send B, i, E_B, $\phi_B(P_A)$, $\phi_B(Q_A)$ to Alice, where i is the session identifier.
4: **if** the r-th session is chosen by $\mathcal{I}$ as the test session **then**
5:     Provide $\mathcal{I}$ as the answer to the test query.
6:     $d \leftarrow \mathcal{I}$'s output.
7: **else**
8:     $d \xleftarrow{R} \{0, 1\}$.
9: **end if**
Output: d

6.2. **Proof of Theorem 6.3.** Using classical techniques from [18, 14], the theorem is proved in three steps, known by the names of *completeness*, *soundness* and *zero-knowledge*.

Proof of Theorem 6.3 (sketch). Completeness is obvious: using the algorithms of Section 4, Peggy can always compute the diagram (3) in polynomial time and make Vic accept.

To prove soundness, we let Charles be any polynomially bounded adversary capable of convincing Vic with a non-negligible probability. We use Charles as a black-box of which we can control the random coin tosses: by restarting it a polynomial number of times with the same random input, and by asking each time a different set of questions, we learn with overwhelming probability a diagram

(22)
$$
\begin{array}{ccc}
E & & E/\langle S \rangle \\
\psi \downarrow & & \downarrow \psi' \\
E/\langle R \rangle & \xrightarrow{\phi'} & E/\langle S, R \rangle
\end{array}
$$

It is then straightforward to compute the secret. Let R be a generator of the kernel of ϕ'. Using the theorem of the dual isogeny it is easy to compute $\hat{\psi}$, then $\langle \hat{\psi}(R) \rangle$ is the kernel of an isogeny $E \to E/\langle S \rangle$ of degree $\ell_A^{e_a}$. This contradicts the CSSI assumption.

To prove zero-knowledge, we use a *cheating verifier* (CV) as a black-box to construct a *simulator* (S). At each iteration, S makes a (uniformly) random guess at what the next question by V will be.

If S guesses $b = 0$, it chooses a random primitive $\ell_B^{e_B}$-torsion point $R \in E$ and computes $\phi(R)$ (recall that the action of ϕ on $E[\ell_B^{e_B}]$ is part of the public data). Then it constructs the isogenies $\psi : E \to E/\langle R \rangle$ and $\psi' : E/\langle S \rangle \to E/\langle S, R \rangle$

(23)
$$
\begin{array}{ccc}
E & \xdashrightarrow{\phi} & E/\langle S \rangle \\
\psi \downarrow & & \downarrow \psi' \\
E/\langle R \rangle & & E/\langle S, R \rangle
\end{array}
$$

and sends $E_1 = E/\langle R \rangle$ and $E_2 = E/\langle S, R \rangle$ to CV.

If S guesses $b = 1$, it chooses a random supersingular curve E' having the same cardinality as E, and a random primitive $\ell_A^{e_A}$-torsion point $R \in E'$. Then, it constructs the isogeny $\phi' : E' \to E'/\langle R \rangle$

$$E \dashrightarrow^{\phi} E/\langle S \rangle$$

(24)

$$E' \xrightarrow{\phi'} E'/\langle R \rangle$$

and sends $E_1 = E'$ and $E_2 = E'/\langle R \rangle$ to CV.

If CV does not ask the expected question, S simply discards the attempt and restarts. If CV asks the expected question, S writes (E_1, E_2, b, R) on its output. S stops whenever CV rejects, or after m successful interactions with CV.

To prove zero-knowledge, we must show that S runs in polynomial time and that its output is polynomially indistinguishable from the transcript of a conversation between CV and Peggy.

To show that S runs in polynomial time, it is enough to show that, at any iteration, for any guess b made by S, the probability that CV asks question $1 - b$ is exponentially close to $1/2$. Suppose this were not the case, then CV can be used as an oracle for DSSP.

To prove indistinguishability, using the *hybrid* technique of [18, Claim 4.2], it is enough to prove that no polynomial-time distinguisher exists for a single round of the identification scheme. It is obvious that no such distinguisher can exist for questions of type $b = 0$, because the outputs of S and Peggy are identical in this case. Suppose, now, that there exists a distinguisher D which, on input $\phi' : E_1 \to E_2$, can tell with non-negligible probability whether it comes from S or from a conversation between CV and Peggy, then D can be used as an oracle for DSSP. □

7. Implementation results and example

We implemented the key exchange protocol in the computer algebra system Sage [39] using a mixed C/Cython/Sage architecture. This allows us to access the large palette of number theoretic algorithms distributed with Sage, while still permitting very efficient code in C/Cython for the critical parts such as the algorithms of Section 4.2. The source code can be downloaded from De Feo's web page.

Arithmetic in $\mathbb{F}_{p^2}$ is written in C. We use the library GMP for arithmetic modulo p. The field $\mathbb{F}_{p^2}$ is implemented as $\mathbb{F}_{p^2}[X]/(X^2+1)$ (this requires $p = 3 \mod 4$); using this representation, one multiplication in $\mathbb{F}_{p^2}$ requires three multiplications in $\mathbb{F}_p$, one squaring requires two multiplications in $\mathbb{F}_p$, and one inversion requires one inversion, two squarings, and two multiplications in $\mathbb{F}_p$. Our experiments show that, for the sizes we are interested in, $I = 10M$ and $S = 0.8M$.

The computation of the optimal strategies as described in Section 4.2.2 is done in pure Python, using a dynamic programming algorithm.

We implemented the key exchange algorithm in C for $\ell = 2, 3$, and in Cython for any ℓ. Our experiments show that, in the C implementation, the ratio between scalar multiplications and isogeny evaluations is about 2, which is consistent with the predictions made in Table 1. We used this ratio to tune the optimal strategies for $\ell = 2, 3$: starting from 768 bits, a gain over the balanced strategy, comprised between 1% and 3%, starts getting noticeable. The performances of 3-isogenies are comparable to those of 2-isogenies, thus contradicting the prediction made by Table 2; this is explained by the operations not accounted for in Table 2, such as the computation of the isogenies and of the isogenous curves (3-isogenies gain a factor of about $\log_3 2$ on these operations). This suggests that well optimized 4-isogeny formulas may eventually outperform 2-isogenies.

Finally, the parameter generation is implemented in plain Sage. Because Sage is a collection of many open source mathematical systems, various of its subsystems are involved in this last part. Of these, Pari [47] plays an important role because it is used to compute Hilbert class polynomials and to factor polynomials over finite fields.

TABLE 3. Benchmarks for various key sizes. Alice uses $\ell = 2$, Bob uses $\ell = 3$.

	tuned 2/1			balanced	
	512 bits	768 bits	1024 bits	768 bits	1024 bits
Alice round 1	28.1 ms	65.7 ms	122 ms	66.8 ms	123 ms
Alice round 2	23.3 ms	54.3 ms	101 ms	55.5 ms	102 ms
Bob round 1	28.0 ms	65.6 ms	125 ms	67.1 ms	128 ms
Bob round 2	22.7 ms	53.7 ms	102 ms	55.1 ms	105 ms

All tests ran on a 2.4 GHz Opteron running in 64-bit mode. The results are summarized in Table 3. At the quantum 128-bit security level (768-bit p), our numbers improve upon Stolbunov's reported performance figures [43, Table 1] by over three orders of magnitude (.066 seconds vs. 229 seconds). This is the highest security level appearing in [43, Table 1], so comparisons at higher levels are difficult. Nevertheless, it seems safe to assume that the improvement is even greater at the 256-bit security level. Our results demonstrate that the proposed scheme is practical.

7.1. **Example.** As a convenience, we provide an example calculation of a key-exchange transaction. Let $\ell_A = 2$, $\ell_B = 3$, $e_A = 63$, $e_B = 41$, and $f = 11$. We use the starting curve $E_0 : y^2 = x^3 + x$. For the torsion bases, we use

$$P_A = (2374093068336250774107936421407893885897i + 2524646701852396349308425328218203569693,$$
$$1944869260414574206229153243510104781725i + 1309099413211767078055232768460483417201)$$
$$P_B = (1556716033657530876728525059284431761206i + 1747407329595165241335131647929866065215,$$
$$3456956202852028835529419995475915388483i + 1975912874247458572654720717155755005566)$$

and $Q_A = \psi(P_A), Q_B = \psi(P_B)$, where $i = \sqrt{-1}$ in $\mathbb{F}_{p^2}$ and $\psi(x, y) = (-x, iy)$ is a distortion map [21]. The secret values are

$$m_A = 2575042839726612324, \; n_A = 8801426132580632841,$$
$$m_B = 4558164392438856871, \; n_B = 20473135767366569910$$

The isogeny $\phi_A \colon E_0 \to E_A$ is specified by its kernel, and thus the curve E_A is only well defined up to isomorphism; its exact value may vary depending on the implementation. In our case, the curve is $E_A : y^2 = x^3 + ax + b$ where

$$a = 4281282453562248945620618211807181141127i + 2147708009907711790134986624604674525769$$
$$b = 3230359267202197460304331835170424053093i + 1577264336482370197045362359104894884862$$

and the values of $\phi_A(P_B)$ and $\phi_A(Q_B)$ are

$$\phi_A(P_B) = (1216243037955078292900974859441066026976i + 1666291136804738684832637187674330905572,$$
$$3132921609453998361853372941893500107923i + 2823164938573549485619800034616855236

6)$$
$$\phi_A(Q_B) = (2039728694420930519155732965018291910660i + 2422092614322988112492931615528155727388,$$
$$1688115812694355145549889238510457034272i + 1379185984608240638912948890349738467536)$$

Similarly, in our implementation $E_B : y^2 = x^3 + ax + b$ is the curve with

$$a = 2574722398094022968578313861884608943122i + 4645075571495590621841741325716474277

22$$
$$b = 2863478907513088792144998311229772886197i + 1767078036714109405796777065089868386753$$

and the values of $\phi_B(P_A)$ and $\phi_B(Q_A)$ are

$$\phi_B(P_A) = (2519086003347973214770499154162540098181i + 1459702974009609198723981125457548440872,$$
$$2072057067933292599326928766255155081380i + 891622100638258849401618552145232311395)$$
$$\phi_B(Q_A) = (5379399452280339324392143298279854366

6i + 3698741609788138685588489568343190504844,$$
$$2853868073971808398649663652161215323750i + 1869730480053624141372373282795858691139)$$

The common j-invariant of $E_{AB} \cong E_{BA}$, computed by both Alice and Bob, is equal to

$$j(E_{AB}) = 1437145494362655119168482808702111413744i + 8334980967783864529517222853105920563

51.$$

8. Conclusion

We propose a new family of conjecturally quantum-resistant public-key cryptographic protocols using isogenies between supersingular elliptic curves of smooth order. In order to compensate for the noncommutative endomorphism rings that arise in this setting, we introduce the idea of providing the images of torsion bases as part of the protocol. Against the fastest known attacks, the resulting key exchange scheme improves upon all previous isogeny-based schemes by orders of magnitude in performance at conventional security levels, making it the first practical isogeny-based public-key cryptosystem. Unlike prior such schemes, our proposal admits no known subexponential-time attacks even in the quantum setting.

References

[1] Adrian Antipa, Daniel Brown, Robert Gallant, Rob Lambert, René Struik, and Scott Vanstone. Accelerated verification of ECDSA signatures. In *Selected Areas in Cryptography 2005*, volume 3897 of *Lecture Notes in Computer Science*, pages 307–318, Berlin, Heidelberg, 2006. Springer Berlin / Heidelberg.

[2] Juliana V. Belding. *Number Theoretic Algorithms for Elliptic Curves*. PhD thesis, University of Maryland, 2008.

[3] Daniel Bernstein, Peter Birkner, Marc Joye, Tanja Lange, and Christiane Peters. Twisted Edwards Curves. In *Progress in Cryptology AFRICACRYPT 2008*, pages 389–405, 2008.

[4] Daniel J. Bernstein and Tanja Lange. Explicit-Formulas Database, 2007. http://www.hyperelliptic.org/EFD/index.html.

[5] Alin Bostan, François Morain, Bruno Salvy, and Éric Schost. Fast algorithms for computing isogenies between elliptic curves. *Math. Comp.*, 77(263):1755–1778, 2008.

[6] Reinier Bröker. Constructing supersingular elliptic curves. *J. Comb. Number Theory*, 1(3):269–273, 2009.

[7] Ran Canetti and Hugo Krawczyk. Analysis of key-exchange protocols and their use for building secure channels. In Birgit Pfitzmann, editor, *EUROCRYPT*, volume 2045 of *Lecture Notes in Computer Science*, pages 453–474. Springer, 2001.

[8] Juan M. Cerviño. On the correspondence between supersingular elliptic curves and maximal quaternionic orders, April 2004. http://arxiv.org/abs/math/0404538.

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

[10] Andrew Childs, David Jao, and Vladimir Soukharev. Constructing elliptic curve isogenies in quantum subexponential time, 2010. http://arxiv.org/abs/1012.4019/.

[11] Jean-Marc Couveignes. Hard homogeneous spaces, 2006. http://eprint.iacr.org/2006/291/.

[12] Giuliana Davidoff, Peter Sarnak, and Alain Valette. *Elementary number theory, group theory, and Ramanujan graphs*, volume 55 of *London Mathematical Society Student Texts*. Cambridge University Press, Cambridge, 2003.

[13] Taher ElGamal. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE transactions on information theory*, 31(4):469–472, 1985.

[14] Uriel Feige, Amos Fiat, and Adi Shamir. Zero-knowledge proofs of identity. *Journal of Cryptology*, 1(2):77–94, June 1988.

[15] Steven D. Galbraith. Constructing isogenies between elliptic curves over finite fields. *LMS J. Comput. Math.*, 2:118–138 (electronic), 1999.

[16] Steven D. Galbraith, Florian Hess, and Nigel P. Smart. Extending the GHS Weil descent attack. In *Advances in cryptology—EUROCRYPT 2002 (Amsterdam)*, volume 2332 of *Lecture Notes in Comput. Sci.*, pages 29–44. Springer, Berlin, 2002.

[17] Steven D. Galbraith and Anton Stolbunov. Improved algorithm for the isogeny problem for ordinary elliptic curves, 2011. http://arxiv.org/abs/1105.6331/.

[18] Oded Goldreich, Silvio Micali, and Avi Wigderson. Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems. *Journal of the Association for Computing Machinery*, 38(3):690–728, July 1991.

[19] David Jao and Luca De Feo. Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. In Bo-Yin Yang, editor, *PQCrypto*, volume 7071 of *Lecture Notes in Computer Science*, pages 19–34. Springer, 2011.

[20] 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.

[21] Antoine Joux. The Weil and Tate pairings as building blocks for public key cryptosystems. In *Algorithmic number theory (Sydney, 2002)*, volume 2369 of *Lecture Notes in Comput. Sci.*, pages 20–32. Springer, Berlin, 2002.

[22] Paul Kocher, Joshua Jaffe, and Benjamin Jun. Differential Power Analysis. In *Advances in Cryptology CRYPTO' 99*, volume 1666 of *Lecture Notes in Computer Science*, pages 388–397, Berlin, Heidelberg, December 1999. Springer Berlin Heidelberg.

[23] David Kohel. *Endomorphism rings of elliptic curves over finite fields*. PhD thesis, University of California, Berkeley, 1996.

[24] Jeffrey C. Lagarias and Andrew M. Odlyzko. Effective versions of the Chebotarev density theorem. In *Algebraic number fields: L-functions and Galois properties (Proc. Sympos., Univ. Durham, Durham, 1975)*, pages 409–464. Academic Press, London, 1977.

[25] Alexander Lubotzky. *Discrete groups, expanding graphs and invariant measures*, volume 125 of *Progress in Mathematics*. Birkhäuser Verlag, Basel, 1994. With an appendix by Jonathan D. Rogawski.

[26] Alexander Lubotzky, Ralph Phillips, and Peter Sarnak. Ramanujan graphs. *Combinatorica*, 8(3):261–277, 1988.

[27] Jean-François Mestre. La méthode des graphes. Exemples et applications. In *Proceedings of the international conference on class numbers and fundamental units of algebraic number fields (Katata, 1986)*, pages 217–242, Nagoya, 1986. Nagoya Univ.

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

[29] OEIS Foundation Inc. The on-line encyclopedia of integer sequences. `http://oeis.org/A130715`, 2012.

[30] Christophe Petit, Kristin Lauter, and Jean-Jacques Quisquater. Full cryptanalysis of LPS and Morgenstern hash functions. In *Proceedings of the 6th international conference on Security and Cryptography for Networks*, SCN '08, pages 263–277, Berlin, Heidelberg, 2008. Springer-Verlag.

[31] Arnold K. Pizer. Ramanujan graphs and Hecke operators. *Bull. Amer. Math. Soc. (N.S.)*, 23(1):127–137, 1990.

[32] Arnold K. Pizer. Ramanujan graphs. In *Computational perspectives on number theory (Chicago, IL, 1995)*, volume 7 of *AMS/IP Stud. Adv. Math.*, pages 159–178. Amer. Math. Soc., Providence, RI, 1998.

[33] David Pointcheval. A new identification scheme based on the perceptrons problem. In *Advances in Cryptology EURO-CRYPT '95*, volume 921 of *Lecture Notes in Computer Science*, pages 319–328, Berlin, Heidelberg, 1995. Springer Berlin / Heidelberg.

[34] Alexander Rostovtsev and Anton Stolbunov. Public-key cryptosystem based on isogenies, 2006. `http://eprint.iacr.org/2006/145/`.

[35] Peter Sarnak. *Some applications of modular forms*, volume 99 of *Cambridge Tracts in Mathematics*. Cambridge University Press, Cambridge, 1990.

[36] Adi Shamir. An efficient identification scheme based on permuted kernels (extended abstract). In *Proceedings on Advances in cryptology*, CRYPTO '89, pages 606–609, New York, NY, USA, 1989. Springer-Verlag New York, Inc.

[37] Joseph H. Silverman. *The arithmetic of elliptic curves*, volume 106 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1992. Corrected reprint of the 1986 original.

[38] Jerome A. Solinas. Low-weight binary representations for pairs of integers. Technical report, National Security Agency, USA, 2001.

[39] William A. Stein et al. *Sage Mathematics Software (Version 4.6.2)*. The Sage Development Team, 2011. `http://www.sagemath.org`.

[40] Jacques Stern. Designing identification schemes with keys of short size. In *Advances in Cryptology CRYPTO '94*, volume 839 of *Lecture Notes in Computer Science*, pages 164–173, Berlin, Heidelberg, 1994. Springer Berlin / Heidelberg.

[41] Jacques Stern. A new identification scheme based on syndrome decoding. In *Advances in Cryptology CRYPTO' 93*, volume 773 of *Lecture Notes in Computer Science*, pages 13–21, Berlin, Heidelberg, 1994. Springer Berlin / Heidelberg.

[42] Anton Stolbunov. Reductionist security arguments for public-key cryptographic schemes based on group action. In Stig F. Mjølsnes, editor, *Norsk informasjonssikkerhetskonferanse (NISK)*, pages 97–109, 2009.

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

[44] Seiichiro Tani. Claw Finding Algorithms Using Quantum Walk. `http://arxiv.org/abs/0708.2584`, March 2008.

[45] John Tate. Endomorphisms of abelian varieties over finite fields. *Invent. Math.*, 2:134–144, 1966.

[46] Edlyn Teske. The pohlig-hellman method generalized for group structure computation. *Journal of Symbolic Computation*, 27(6):521–534, 1999.

[47] The PARI Group, Bordeaux. *PARI/GP, version* `2.4.3`, 2008. available from `http://pari.math.u-bordeaux.fr/`.

[48] Jacques Vélu. Isogénies entre courbes elliptiques. *C. R. Acad. Sci. Paris Sér. A-B*, 273:A238–A241, 1971.

[49] Shengyu Zhang. Promised and Distributed Quantum Search Computing and Combinatorics. In *Proceedings of the Eleventh Annual International Conference on Computing and Combinatorics*, volume 3595 of *Lecture Notes in Computer Science*, pages 430–439, Berlin, Heidelberg, 2005. Springer Berlin / Heidelberg.

LABORATOIRE PRISM, UNIVERSITÉ DE VERSAILLES, 78035 VERSAILLES, FRANCE
E-mail address: `Luca.De-Feo@prism.uvsq.fr`
URL: `http://www.prism.uvsq.fr/~dfl/`

UNIVERSITY OF WATERLOO, WATERLOO, ONTARIO, N2L 3G1, CANADA
E-mail address: `djao@math.uwaterloo.ca`
URL: `http://www.math.uwaterloo.ca/~djao/`

LABORATOIRE PRISM, UNIVERSITÉ DE VERSAILLES, 78035 VERSAILLES, FRANCE
E-mail address: `Jerome.Plut@prism.uvsq.fr`
URL: `http://www.prism.uvsq.fr/~plut/`