On Lions and Elligators: An efficient constant-time implementation of CSIDH

Michael Meyer[1,2], Fabio Campos[1], and Steffen Reith[1]

[1] Department of Computer Science, University of Applied Sciences Wiesbaden, Germany
[2] Department of Mathematics, University of Würzburg, Germany
{Michael.Meyer, FabioFelipe.Campos, Steffen.Reith}@hs-rm.de

Abstract. The recently proposed CSIDH primitive is a promising candidate for post quantum static-static key exchanges with very small keys. However, until now there is only a variable-time proof-of-concept implementation by Castryck, Lange, Martindale, Panny, and Renes, recently optimized by Meyer and Reith, that can leak various information about the private key. Therefore, we present a constant-time implementation that samples key elements only from intervals of nonnegative numbers and uses dummy isogenies, which prevents certain kinds of side-channel attacks. We apply several optimizations, e.g. SIMBA and Elligator, in order to get a more efficient implementation.

Keywords: CSIDH · Post-Quantum Cryptography · constant-time · Supersingular Elliptic Curve Isogenies

1 Introduction

Isogeny-based cryptography is the most juvenile family of the current proposals for post-quantum cryptography. The first cryptosystem based on the hardness of finding an explicit isogeny between two given isogenous elliptic curves over a finite field was proposed in 1997 by Couveignes [9], eventually independently rediscovered by Rostovtsev and Stolbunov [18] in 2004, and therefore typically called CRS. Childs, Jao and Soukharev [6] showed in 2010, that CRS can be broken using a subexponential quantum algorithm by solving an abelian hidden shift problem. To avoid this attack, Jao and De Feo [12] described a new isogeny-based scheme SIDH (supersingular isogeny Diffie-Hellman) that works with supersingular curves over $\mathbb{F}_{p^2}$. The current state-of-the-art implementation is SIKE [11], which was submitted to the NIST post-quantum cryptography competition [16].

De Feo, Kieffer and Smith optimized CRS in 2018 [10]. Their ideas led to the development of CSIDH by Castryck, Lange, Martindale, Panny, and Renes [5], who adapted the CRS scheme to supersingular curves and isogenies defined over a prime field $\mathbb{F}_p$. They implemented the key exchange as a proof-of-concept,

This work was partially supported by Elektrobit Automotive, Erlangen, Germany.

which is efficient, but does not run in constant-time, and can therefore leak information about private keys. We note that building an efficient constant-time implementation of CSIDH is not as straightforward as in SIDH, where, speaking of running-times, only one Montgomery ladder computation depends on the private key (see [8]).

In this paper we present a constant-time implementation of CSIDH with many practical optimizations.

Organization. The rest of this paper is organized as follows. The following section gives a brief algorithmic introduction to CSIDH [5]. Two leakage scenarios based on time and power analysis are presented in Section 3. In Section 4, we suggest different methods on how to avoid these leakages and build a constant-time implementation. Section 5 contains a straightforward implementation description of our suggested methods, and various optimizations. Thereafter, we provide implementation results in Section 6 and give concluding remarks in Section 7. The appendices A and B give more details about our implementations and algorithms.

Note that there are two different notions of constant-time implementations, as explained in [2]. In our case, it suffices to work with the notion, that the running time does not depend upon the choice of the private key, but may vary due to randomness. The second notion specifies strict constant time, meaning that the running time must be the same every time, independent from private keys or randomness. When we talk about constant time throughout this paper, we refer to the first notion described here.

Related work. In [2], Bernstein, Lange, Martindale, and Panny describe constant-time implementations in the second notion from above, which is required for quantum attacks. In this paper, we follow the mentioned different approach for an efficient constant-time implementation, but reuse some of the techniques from [2].

2 CSIDH

We only cover the algorithmic aspects of CSIDH here, and refer to [5] for the mathematical background and a more detailed description.

We first choose a prime of the form $p = 4 \cdot \ell_1 \cdot ... \cdot \ell_n - 1$, where the ℓ_i are small distinct odd primes. We work with supersingular curves over $\mathbb{F}_p$, which guarantees the existence of points of the orders ℓ_i, that enable us to compute ℓ_i-isogenies from kernel generator points by Vélu-type formulas [19].

A private key consists of a tuple $(e_1, ..., e_n)$, where the e_i are sampled from an interval $[-B, B]$. The absolute value $|e_i|$ specifies how many ℓ_i-isogenies have to be computed, and the sign of e_i states, whether points on the current curve or on its twist have to be used as kernel generators. One can represent this graphically:

Algorithm 1: Evaluating the class group action.

```
Input  : a ∈ F_p and a list of integers (e_1, ..., e_n).
Output: a' such that [l_1^{e_1} ··· l_n^{e_n}]E_a = E_{a'}.
1  while some e_i ≠ 0 do
2      Sample a random x ∈ F_p.
3      if x^3 + ax^2 + x is a square in F_p then
4          s ← +1.
5      else
6          s ← −1.
7      Let S = {i | sign(e_i) = s}.
8      if S = ∅ then
9          Go to line 2.
10     P = (x : 1), k ← ∏_{i∈S} ℓ_i, P ← [(p + 1)/k]P.
11     foreach i ∈ S do
12         K ← [k/ℓ_i]P.
13         if K ≠ ∞ then
14             Compute a degree-ℓ_i isogeny φ : E_a → E_{a'} with ker(φ) = ⟨K⟩.
15             a ← a', P ← φ(P), k ← k/ℓ_i, e_i ← e_i − s.
16 return A.
```

Over $\mathbb{F}_p$, the supersingular ℓ_i-isogeny graph consists of distinct cycles. Therefore, we have to walk $|e_i|$ steps through the cycle for ℓ_i, where the sign of e_i tells us the direction.

Since this class group action is commutative, it allows a basic Diffie-Hellman-type key exchange: Starting from a supersingular curve E, Alice and Bob choose a private key as described above, and compute their public key curves E_A resp. E_B via isogenies, as described in Algorithm 1. Then Alice repeats her computations, this time starting at the curve E_B, and vice versa. Both parties then arrive at the same curve E_{AB}, which represents their shared secret. Furthermore, public keys can be verified efficiently in CSIDH (see [5]). Therefore, a static-static key-exchange is possible.

However, the quantum security is still an open problem. For our implementation, we use CSIDH-512, the parameter set from [5], that is conjectured to satisfy NIST security level 1. In the light of the subexponential attack on CRS [6], more analysis on CSIDH has been done in [3,4,2].

3 Leakage scenarios

It is clear and already mentioned in [5] that the proof-of-concept implementation of CSIDH is not side-channel resistant. In this paper we focus on two scenarios, that can leak information on the private key. Note that the second scenario

requires the attacker to have more power than the first. Further, there will naturally be many more scenarios for side-channel attacks.

Timing leakage. As the private key in CSIDH specifies how many isogenies of which degree have to be computed, it is obvious that this (up to additional effort for point multiplications due to the random choice of points) determines the running time of the algorithm. As stated in [13], the worst case running time occurs for the private key $(5, 5, ..., 5)$, and takes more than 3 times as much as in the average case. The other extreme is the private key $(0, 0, ..., 0)$, which would require no computations at all. However, in a timing-attack protected implementation, the running time should be independent from the private key.

Power analysis. Instead of focusing on the running time, we now assume that an attacker can measure the power consumption of the algorithm. We further assume that from the measurement, the attacker can determine blocks which represent the two main primitives in CSIDH, namely point multiplication and isogeny computation, and can separate these from each other. Now assume that the attacker can separate the loop iterations from each other. Then the attacker can determine which private key elements share the same sign from the isogeny blocks that are performed in the same loop, since they have variable running time based on the isogeny degree. This significantly reduces the possible key space and therefore also the complexity of finding the correct key.

4 Mitigating Leakages

In this section we give some ideas on how to fix these possible leakages in an implementation of CSIDH. We outline the most important ideas here, and give details about how to implement them efficiently in CSIDH-512 in section 5.

Dummy isogenies. First, it seems obvious that one should compute a constant number of isogenies of each degree ℓ_i, and only use the ones required by the private key, in order to obtain a constant running time. However, in this case additional multiplications are required, if normal isogenies and unused isogenies are computed in the same loop[3]. We adapt the idea of using dummy isogenies from [13] for that cause. Meyer and Reith propose to design dummy isogenies, that instead of updating the curve parameters and evaluating the curve point P, computes $[\ell_i]P$ in the degree-ℓ_i dummy isogeny. Since the isogeny algorithm computes $[\frac{\ell_i-1}{2}]K$ for the kernel generator K, one can replace K by P there, and perform two more differential additions to compute $[\ell_i]P$. The curve parameters remain unchanged.

In consequence, a dummy isogeny simply performs a scalar multiplication. Therefore, the output point $[\ell_i]P$ then has order not divided by ℓ_i, which is

[3] This is required, since otherwise, an attacker in the second leakage scenario can determine the private key easily.

important for using this point to compute correct kernel generators in following iterations. Further, one can design the isogeny and dummy isogeny algorithms for a given degree ℓ_i such that they perform the same number and sequence of operations with only minor computational overhead compared to the isogenies from [13]. This is important to make it hard for side-channel attackers to distinguish between those two cases.

Balanced vs. unbalanced private keys. Using dummy isogenies to spend a fixed time on isogeny computations in not enough for a constant-time implementation, however. Another problem lies in the point multiplications in line 10 and 12 of algorithm 1. We use an observation from [13] to illustrate this. They consider the private keys $(5, 5, 5, ...)$ and $(5, -5, 5, -5, ...)$ and observe that for the first key, the running time is 50% higher than for the second key. The reason for this is that in the first case in order to compute one isogeny of each degree, the multiplication in line 10 is only a multiplication by 4, and the multiplication in line 12 has a factor of bitlength 509 in the first iteration, 500 in the second iteration, and so on.

For the second key, we have to perform one loop through the odd i and one through the even i in order two compute one isogeny of each degree ℓ_i. Therefore, the multiplications in line 10 are by 254 resp. 259 bit factors, while the bitlengths of the factors in the multiplications in line 12 are 252, 244,..., resp. 257, 248, and so on (see Figure 1). In total, adding up the bitlenghts of all factors, we can measure the cost of all point multiplications for the computation of one isogeny per degree, where we assume that the condition in line 13 of algorithm 1 never fails, since one Montgomery ladder step is performed per bit. For the first key, we end up with 16813 bits, while for the second key we only have 9066 bits.

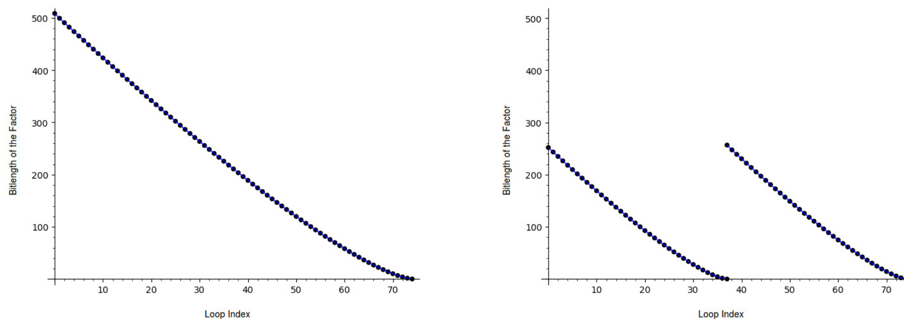

Fig. 1. Bitlengths of factors for computing one isogeny per degree.

This can be generalized to any private key: The more the key elements (or the products of the respective ℓ_i) are unbalanced, i.e. many of them share the same sign, the more the computational effort grows, compared to the perfectly

balanced case from above. This behavior depends on the private key and can therefore leak information. Hence, it is clear that we have to prevent this in order to achieve a constant-time implementation.

One way to achieve this is to use constant-time Montgomery ladders, that always run to the maximum bitlength, no matter how large the respective factor is. However, this would lead to a massive increase in running time. Another possibility for handling this is to only choose key elements of a fixed sign. Then we have to adjust the interval from which we sample the integer key elements, e.g. from $[-5, 5]$ to $[0, 10]$. This however doubles the computational effort for isogenies (combined normal and dummy isogenies). We will return to this idea later.

Determining the sign distribution. In our second leakage scenario, an attacker might determine the sign distribution of the key elements by identifying blocks of isogeny resp. dummy isogeny computations. One way of mitigating this attack would be to let each degree-ℓ_i isogeny run as long as a ℓ_{max}-isogeny, where ℓ_{max} is the largest ℓ_i. As used in [2], this is possible because of the Matryoshka-doll structure of the isogeny algorithms. This would allow an attacker in the second leakage scenario to only determine the number of positive resp. negative elements, but not their distribution, at the cost of a large increase of computational effort. We can also again restrict to the case that we only choose nonnegative (resp. only nonpositive) key elements. Then there is no risk of leaking information about the sign distribution of the elements, since in this setting the attacker knows this beforehand, at the cost of twice as many isogeny computations.

Limitation to nonnegative key elements. Since this choice eliminates both of the aforementioned possible leakages, we use the mentioned different interval to sample private key elements from. In CSIDH-512, this means using the interval $[0, 10]$ instead of $[-5, 5]$. One might ask if this affects the security properties of CSIDH. As before, there are 11^{74} different tuples to choose from in CSIDH-512. Castryck et al. argue in [5] that there are multiple vectors $(e_1, e_2, ..., e_n)$, which represent the same ideal class, meaning that the respective keys are equivalent. However, they assume by heuristic arguments that the number of short representations per ideal class is small, i.e. the 11^{74} different keys $(e_1, e_2, ..., e_n)$, where all e_i are sampled from the interval $[-5, 5]$, represent not much less than 11^{74} ideal classes. Heuristically we can conjecture that the same applies if we sample from the slightly shifted interval $[0, 10]$. Therefore, we assume that our different choice of parameters has no implications on the security of CSIDH-512.

In the following sections we focus on optimized implementations, using the mentioned countermeasures against attacks, i.e. sampling key elements from the interval $[0, 10]$ and using dummy isogenies.

5 Efficient Implementation

5.1 Straightforward Implementation

First, we describe the straightforward implementation of the evaluation of the class group action in CSIDH-512 with the choices from above, before applying various optimizations. We briefly go through the implementation aspects of the main primitives, i.e. point multiplications, isogenies and dummy isogenies, and explain why this algorithm runs in constant-time, i.e. does not leak information about the private key.

Algorithm 2: Constant-time evaluation of the class group action in CSIDH-512.

Input : $a \in \mathbb{F}_p$ and a list of integers $(e_1, ..., e_n)$.
Output: a' such that $[\mathfrak{l}_1^{e_1} \cdots \mathfrak{l}_n^{e_n}]E_a = E_{a'}$.

1. Initialize $k = 4$, $e^{iso} = (e_1, ..., e_n)$ and $e^{dum} = (e_1^{dum}, ..., e_n^{dum})$, where $e_i^{dum} = 10 - e_i$.
2. **while** some $e_i^{iso} \neq 0$ or $e_i^{dum} \neq 0$ **do**
3. Sample random values $x \in \mathbb{F}_p$ until we have some x where $x^3 + ax^2 + x$ is a square in $\mathbb{F}_p$.
4. Set $P = (x : 1)$, $P \leftarrow [k]P$, $S = \{i \mid e_i^{iso} \neq 0$ or $e_i^{dum} \neq 0\}$.
5. **foreach** $i \in S$ **do**
6. Let $m = \prod_{j \in S, j > i} \ell_i$.
7. Set $K \leftarrow [m]P$.
8. **if** $K \neq \infty$ **then**
9. **if** $e_i^{iso} \neq 0$ **then**
10. Compute a degree-ℓ_i isogeny $\varphi : E_a \rightarrow E_{a'}$ with $ker(\varphi) = \langle K \rangle$.
11. $a \leftarrow a'$, $P \leftarrow \varphi(P)$, $e_i^{iso} \leftarrow e_i^{iso} - 1$.
12. **else**
13. Compute a degree-ℓ_i dummy isogeny:
14. $a \leftarrow a$, $P \leftarrow [\ell_i]P$, $e_i^{dum} \leftarrow e_i^{dum} - 1$.
15. **if** $e_i^{iso} = 0$ *and* $e_i^{dum} = 0$ **then**
16. Set $k \leftarrow k \cdot \ell_i$.

Parameters. As described in [5], we have a prime number $p = 4 \cdot \ell_1 \cdot \ell_2 \cdot ... \cdot \ell_n - 1$, where the ℓ_i are small distinct odd primes. We further assume that we have $\ell_1 > \ell_2 > ... > \ell_n$. In CSIDH-512 we have $n = 74$, and we sample the elements of private keys $(e_1, e_2, ..., e_n)$ from $[0, 10]$.

Handling the private key. Similar to the original implementation of Castryck et al., we copy the elements of the private key in an array $e^{iso} = (e_1, e_2, ..., e_n)$,

where e_i determines how many isogenies of degree ℓ_i we have to compute. Furthermore, we set up another array $e^{dum} = (10 - e_1, 10 - e_2, ..., 10 - e_n)$, to determine how many dummy isogenies of each degree we have to compute. As we go through the algorithm, we compute all the required isogenies and dummy isogenies, reducing e_i^{iso} resp. e_i^{dum} by 1 after each degree-ℓ_i isogeny resp. dummy isogeny. We therefore end up with a total of 10 isogeny computations (counting isogenies and dummy isogenies) for each ℓ_i.

Sampling random points. In algorithm 2 line 3, we have to find curve points on the current curve, that are defined on the curve itself instead of the twist curve, i.e. their y-coordinates have to be defined over $\mathbb{F}_p$. As in [5] this can be done by sampling a random $x \in \mathbb{F}_p$, and computing y^2 by the curve equation $y^2 = x^3 + ax^2 + x$. We then check if y is defined over $\mathbb{F}_p$ by a Legendre symbol computation, i.e. by computing $(y^2)^{(p-1)/2} \pmod{p}$. If this is not the case, we simply repeat this procedure until we find a suitable point. Note that we require the curve parameter a to be in affine form. Since a will typically be in projective form after isogeny computations, we therefore have to compute the affine parameter each time before sampling a new point.

Elliptic curve point multiplications. Since we work with Montgomery curves, using only projective XZ-coordinates, and projective curve parameters $a = A/C$, we can use the standard Montgomery ladder as introduced in [14], adapted to projective curve parameters as in [8]. This means that per bit of the factor, one combined doubling and differential addition is performed.

Isogenies. For the computation of isogenies, we use the formulas presented by Meyer and Reith in [13]. They combine the Montgomery isogeny formulas by Costello and Hisil [7], and Renes [17] with the twisted Edwards formulas by Moody and Shumow [15], in order to obtain an efficient algorithm for the isogeny computations in CSIDH. For a ℓ_i-isogeny, this requires a point K of order ℓ_i as kernel generator, and the projective parameters A and C of the current curve. It outputs the image curve parameters A' and C', and the evaluation of the point P. As mentioned before, the algorithm computes all multiples of the point K up to the factor $\frac{\ell_i - 1}{2}$. See e.g. [2] for more details.

Dummy isogenies. As described before, we want the degree-ℓ_i dummy isogenies to output the scalar multiple $[\ell_i]P$ instead of an isogeny evaluation of P. Therefore, we interchange the points K and P in the original isogeny algorithm, such that it computes $[\frac{\ell_i - 1}{2}]P$. We then perform two more differential additions, i.e. compute $[\frac{\ell_i + 1}{2}]P$ from $[\frac{\ell_i - 1}{2}]P$, P, and $[\frac{\ell_i - 3}{2}]P$, and compute $[\ell_i]P$ from $[\frac{\ell_i + 1}{2}]P$, $[\frac{\ell_i - 1}{2}]P$, and P.

If one is concerned that a side-channel attacker can detect, that the curve parameters A and C are not changed for some time (meaning that a series of

dummy isogenies is performed), one could further multiply A and C by a random number[4] $0 < k < p$.

As mentioned before, we want isogenies and dummy isogenies of degree ℓ_i to have the same number and sequence of operations. Hence, we also perform the two extra differential additions in the isogeny algorithm, without using the results.

5.2 Running time

We now explain why this algorithm runs in constant time. As already explained, we perform 10 isogeny computations (counting isogenies and dummy isogenies) for each degree ℓ_i. Furthermore, isogenies and dummy isogenies have the same running time. Therefore the total computational effort for isogenies is constant, independent from the respective private key. We also set the same condition (line 8 of algorithm 2) for the kernel generator for the computation of a dummy isogeny, in order not to leak information.

Sampling random points and finding a suitable one doesn't run in constant time in algorithm 2. However, the running time only depends on randomly chosen values, and does not leak any information on the private key.

Now for simplicity assume that we always find a point of full order, i.e. a point that can be used to compute one isogeny of each degree ℓ_i. Then it is easy to see that the total computational effort for scalar multiplications in algorithm 2 is constant, independent from the respective private key. If we now allow random points, we will typically not satisfy the condition in line 8 of algorithm 2 for all i. Therefore, additional computations (sampling random points, and point multiplications) are required. However, this does not leak information about the private key, since this only depends on the random choice of curve points, but not on the private key.

Hence, we conclude, that the implementation of algorithm 2 as described here, prevents the leakage scenarios considered in section 3. It is however quite slow compared to the performance of variable-time CSIDH-512 in [13,5]. In the following section, we focus on how to optimize and speed up the implementation.

5.3 Optimizations

Sampling points with Elligator. In [2] Bernstein, Lange, Martindale, and Panny pointed out, that Elligator [1], specifically the Elligator 2 map, can be used in CSIDH to be able to choose points over the required field of definition. Since we only need points defined over $\mathbb{F}_p$, this is especially advantageous in our situation. For $a \neq 0$ the Elligator 2 map works as follows (see [2]):

- Sample a random $u \in \{2, 3, ..., (p-1)/2\}$.
- Compute $v = A/(u^2 - 1)$.

---

[4] One could actually use an intermediate non-zero number of the isogeny computation, since the factor is not required to be truly random.

- Compute e, the Legendre symbol of $v^3 + av^2 + v$.
- If $e = 1$, output v. Otherwise, output $-v - a$.

Therefore, for all $a \neq 0$, we can replace the search for a suitable point in line 3 of algorithm 2, at the cost of an extra inversion. However, as explained by Berstein et al., one can precompute $1/(u^2 - 1)$ for some values of u, e.g. for $u \in \{2, 3, 4, ...\}$. Then the cost is essentially the same as for the random choice of points, but we always find a suitable point this way, compared to the probability of 1/2 when sampling random points. This could, however, potentially lead to the case that we cannot finish the computation: Consider that we only have one isogeny of degree ℓ_i left to compute, but for all of the precomputed values of u, the order of the corresponding point is not divided by ℓ_i. Then we would have to go back to a random choice of points to finish the computation. However, our experiments suggest that it is enough to have 10 precomputed values. Note that the probability for actually finding points of suitable order is somewhat lower when using Elligator (see [2]).

For $a = 0$, Bernstein et al. also show how to adapt the Elligator 2 map to this case, but also argue that one could precompute a point of full order (divided by all ℓ_i) and simply use this point whenever $a = 0$. We follow their latter approach.

SIMBA (Splitting isogeny computations into multiple batches). In section 4, we analyzed the running time of variable-time CSIDH-512 for the keys $e_1 = (5, 5, ..., 5)$ and $e_2 = (5, -5, 5, -5, ...)$. For the latter, the algorithm is significantly faster, because of the smaller multiplications during the loop (line 12 of algorithm 1), see Figure 1. We adapt and generalize this observation here, in order to speed up our constant-time implementation.

Consider for our setting the key $(10, 10, ..., 10)$ and that we can again always choose points of full order. To split the indices as algorithm 1 does for the key e_2, we define the two sets $S_1 = \{1, 3, 5, ..., 73\}$ and $S_2 = \{2, 4, 5, ..., 74\}$. Then the loops through the ℓ_i for $i \in S_1$ resp. $i \in S_2$ require significantly smaller multiplications, while only requiring to compute $[4k]P$ with $k = \prod_{i \in S_2} \ell_i$ resp. $k = \prod_{i \in S_1} \ell_i$ beforehand. We now simply perform 10 loops for each set, and hence this gives exactly the same speedup, as algorithm 1 gives for e_2 compared to e_1.

One might ask if splitting the indices in 2 sets already gives the best speedup. We generalize the observation from above, now splitting the indices into m batches, where $S_1 = \{1, m+1, 2m+1, ...\}$, $S_2 = \{2, m+2, 2m+2, ...\}$, and so on[5]. Before starting a loop through the indices $i \in S_j$ with $1 \leq j \leq n$, one now has to compute $[4k]P$ with $k = \prod_{h \notin S_j} \ell_h$. The number and size of these multiplications grows when m grows, so we can expect, that the speedup turns into an increasing computational effort when m is too large.

[5] Note that in [2] a similar idea is described. However, in their algorithm, only two isogeny degrees are covered in each iteration. Our construction makes use of the fact that we restrict to intervals of nonnegative numbers for sampling the private key elements.

To find the best choice for m, we computed the total number of Montgomery ladder steps during the computation of one isogeny of each degree in CSIDH-512 for different m, with the assumptions from above. We did not take into account here, that when m grows, we will have to sample more points (which costs at least one Legendre symbol computation each). Table 1 shows that the optimal choice should be around $m = 5$.

Table 1. Number of Montgomery ladder steps for computing one isogeny of each degree in CSIDH-512 for different numbers of batches m.

m	1	2	3	4	5	6	7
Ladder steps	16813	9066	6821	5959	5640	5602	5721

If we now come back to the choice of points through Elligator, the assumption from above does not hold anymore, and with very high probability, we will need more than 10 loops per index set. Typically, soon after 10 loops through each set, the large isogeny degrees will be finished, while there are some small degree isogenies left to compute. In this case our optimization backfires, since in this construction, the indices of the missing ℓ_i will be distributed among the m different sets. We therefore need large multiplications in order to only check a few small degrees per set. Hence it is beneficial to define a number $\mu \geq 10$, and merge the sets after μ steps, i.e. simply going back to algorithm 2 for the computation of the remaining isogenies. We dub this construction SIMBA-m-μ.

Sampling private key elements from different intervals. Instead of sampling all private key elements from the interval $[0, 10]$, and in total computing 10 isogenies of each degree, one could also consider to choose the key elements from different intervals for each isogeny degree. For a private key $e = (e_1, e_2, ..., e_n)$, we can choose an interval $[0, B_i]$ for each e_i, in order to e.g. reduce the number of expensive large degree isogenies at the cost of computing more low degree isogenies. We require $\prod_i (B_i + 1) \approx 11^{74}$, in order to obtain the same security level as before. For the security implication of this choice, the same as in section 4 applies.
Trying to find the optimal parameters B_i leads to a large integer optimization problem, that is not likely to be solvable exactly. Therefore, we heuristically searched for parameters, that are likely to improve the performance of CSIDH-512. We present them in section 6 and Appendix A.

Note that if we choose $B = (B_1, ..., B_n)$ differently from $B_i = 10$ for all i, the benefit of our optimizations above may change accordingly. Therefore, we changed the parameters m and μ in our implementation according to the respective B.

Skip point evaluations. As described before, the isogeny algorithms compute the image curve parameters, and push a point P through the isogeny. However,

in the last isogeny per loop, this is unnecessary, since we choose a new point after the isogeny computation anyway. Therefore, it saves some computational effort, if we skip the point evaluation parts in these cases.

Application to variable-time CSIDH. Note that many of the optimizations from above are also applicable to variable-time CSIDH-512 implementations as in [13] or [5]. We could therefore also speed up the respective implementation results using the mentioned methods.

6 Implementation Results

We implemented our optimized constant-time algorithm in C, using the implementation accompanying [13], which is based on the implementation from the original CSIDH paper [5] by Castryck et al. For example the implementation of the field arithmetic in assembly is the one from [5]. Our final algorithm, containing all the optimizations from above, can be found in Appendix B.

Since we described different optimizations, that can influence one another, it is not straightforward to decide which parameters B, m and μ to use. Therefore, we implemented CSIDH-512 dynamically, and tested various choices of parameters. The parameters and implementation results can be found in Appendix A. The best parameters we found are given by

$$\begin{aligned} B = [&5, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 8, 8, 8, 8, 8, 8, 8, 11, 11, 11, \\ &11, 11, 11, 11, 11, 11, 11, 11, 11, 11, 11, 11, 11, 11, 11, 13, 13, 13, 13, 13, 13, 13 \\ &13, 13, 13, 13, 13, 13, 13, 13, 13, 13, 13, 13, 13, 13, 13, 13, 13, 13, 13, 13, 13] \end{aligned}$$

using SIMBA-5-11, where the key element e_i is chosen from $[0, B_i]$ and we assume $\ell_1 > \ell_2 > ... > \ell_n$. We do not claim that these are the best parameters; there might be better choices, that we simply did not consider in our experiments.

Table 2. Performance comparison of the class group action evaluation in CSIDH-512 with the mentioned parameters. All timings were measured on an Intel Core i7-6500 Skylake processor running Ubuntu 16.04 LTS, averaged over 1 000 runs.

Clock Cycles $\times 10^6$	wall clock time
326.5	$\sim$ 128 ms

In table 2, we give the cycle count and running time for the implementation using the parameters from above. The code is freely available at `https://zenon.cs.hs-rm.de/pqcrypto/constant-csidh-c-implementation`.[6]

[6] The provided implementation actually uses a further small speedup as explained in Appendix A, and achieved a running time of 322.6 million clock cycles in the setting from above.

To give a comparison, that mainly shows the impact of SIMBA and the different choice of B, we also ran the straightforward implementation according to algorithm 2 with $B = [10, 10, ..., 10]$, also using Elligator. In this case, we measured 621.5 million clock cycles in the same setting as above.

Compared to the performance of the variable-time implementation from [13], this means a slowdown of factor 3.10. However, as mentioned, also the variable-time implementation can benefit from the optimizations from this paper, so this comparison should not be taken too serious.

7 Conclusion

We present the first implementation of CSIDH, that prevents certain side-channel attacks, such as timing leakages. However, there might be more leakage models, depending on how powerful the attacker is. There is also more work to be done on making this implementation as efficient as possible. It may e.g. be possible to find a CSIDH-friendly prime p, that allows for faster computations in $\mathbb{F}_p$.

Also the security features of CSIDH remain an open problem. More analysis on this is required, to show if the parameters are chosen correctly for the respective security levels.

We note that our results depend on the parameters from CSIDH-512. However, it is clear that the described optimizations can be adapted to other parameter sets and security levels as well.

Acknowledgments. This work was partially supported by Elektrobit Automotive, Erlangen, Germany. We thank Joost Renes for answering some questions during the preparation of this work.

References

1. Bernstein, D.J., Hamburg, M., Krasnova, A., Lange, T.: Elligator: Elliptic-curve points indistinguishable from uniform random strings. In: Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security. pp. 967–980. ACM (2013)
2. Bernstein, D.J., Lange, T., Martindale, C., Panny, L.: Quantum circuits for the CSIDH: optimizing quantum evaluation of isogenies. Cryptology ePrint Archive, Report 2018/1059 (2018), https://eprint.iacr.org/2018/1059
3. Biasse, J.F., Jacobson Jr, M.J., Iezzi, A.: A note on the security of CSIDH. To appear at Indocrypt 2018 (2018), https://arxiv.org/abs/1806.03656
4. Bonnetain, X., Schrottenloher, A.: Quantum Security Analysis of CSIDH and Ordinary Isogeny-based Schemes. Cryptology ePrint Archive, Report 2018/537 (2018), https://eprint.iacr.org/2018/537
5. Castryck, W., Lange, T., Martindale, C., Panny, L., Renes, J.: CSIDH: An efficient post-quantum commutative group action. In: Peyrin, T., Galbraith, S. (eds.) ASIACRYPT 2018, LNCS 11274. pp. 395–427. Springer (2018)
6. Childs, A., Jao, D., Soukharev, V.: Constructing elliptic curve isogenies in quantum subexponential time. Journal of Mathematical Cryptology **8**(1), 1–29 (2014)

7. Costello, C., Hisil, H.: A simple and compact algorithm for SIDH with arbitrary degree isogenies. In: Takagi, T., Peyrin, T. (eds.) Advances in Cryptology ASIACRYPT 2017. Lecture Notes in Computer Science, vol 10624, pp. 303–329. Springer, Cham (2017)
8. Costello, C., Longa, P., Naehrig, M.: Efficient algorithms for supersingular isogeny Diffie-Hellman. In: Robshaw, M., Katz, J. (eds.) Advances in Cryptology - CRYPTO 2016, pp. 572–601. Lecture Notes in Computer Science, 9814, Springer (2016)
9. Couveignes, J.M.: Hard homogeneous spaces. Cryptology ePrint Archive, Report 2006/291 (2006), `https://eprint.iacr.org/2006/291`
10. De Feo, L., Kieffer, J., Smith, B.: Towards practical key exchange from ordinary isogeny graphs. In: Peyrin, T., Galbraith, S. (eds.) ASIACRYPT 2018, LNCS 11274. pp. 365–394. Springer (2018)
11. Jao, D., et al.: Supersingular isogeny key encapsulation, Round 1 submission, NIST Post-Quantum Cryptography Standardization (2017)
12. Jao, D., De Feo, L., Plût, J.: Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. Journal of Mathematical Cryptology **8**(3), 209–247 (2014)
13. Meyer, M., Reith, S.: A faster way to the CSIDH. To appear at Indocrypt 2018 (2018), `https://eprint.iacr.org/2018/782`
14. Montgomery, P.L.: Speeding the Pollard and Elliptic Curve Methods of Factorization. Mathematics of Computation **48**(177), 243–264 (1987)
15. Moody, D., Shumow, D.: Analogues of Vélu's Formulas for Isogenies on Alternate Models of Elliptic Curves. Math. Comput. **85**(300), 1929–1951 (2016)
16. National Institute of Standards and Technology (NIST): Submission requirements and evaluation criteria for the post-quantum cryptography standardization process (2016)
17. Renes, J.: Computing isogenies between Montgomery curves using the action of (0, 0). In: Lange, T., Steinwandt, R. (eds.) PQCrypto 2018, 9th International Conference on Post-Quantum Cryptography. pp. 229–247. Springer (2018)
18. Rostovtsev, A., Stolbunov, A.: Public-key cryptosystem based on isogenies. Cryptology ePrint Archive, Report 2006/145 (2006), `http://eprint.iacr.org/2006/145`
19. Vélu, J.: Isogénies entre courbes elliptiques. C.R. Acad. Sc. Paris, Série A. **271**, 238–241 (1971)

A Implementation Results

We tested several parameters in a dynamical implementation, as explained in the paper. The setting is the same as in section 6. For the parameters $B_0, ..., B_4$ we chose

$$B_0 = [10, 10, 10, ..., 10],$$

$$\begin{aligned} B_1 = [&1, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 8, 8, 8, 8, 8, 8, 8,\\ &12, 12, 12, 12, 12, 12, 12, 12, 12, 12, 12, 12, 12, 12, 14, 14, 14, 14, 14,\\ &14, 14, 14, 14, 14, 14, 14, 14, 14, 14, 14, 14, 14, 14, 14, 14, 14, 14, 14,\\ &14, 14, 14, 14, 14, 14, 14, 14], \end{aligned}$$

$$B_2 = [5, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 7, 8, 8, 8, 8, 8, 8, 8,$$
$$11, 11, 11, 11, 11, 11, 11, 11, 11, 11, 11, 11, 11, 11, 11, 11, 11, 11, 13,$$
$$13, 13, 13, 13, 13, 13, 13, 13, 13, 13, 13, 13, 13, 13, 13, 13, 13, 13, 13,$$
$$13, 13, 13, 13, 13, 13, 13, 13],$$

$$B_3 = [2, 4, 4, 4, 4, 4, 4, 4, 4, 4, 4, 4, 4, 4, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 10, 10,$$
$$10, 10, 10, 10, 10, 10, 10, 10, 10, 16, 16, 16, 16, 16, 16, 16, 16, 16, 16,$$
$$16, 16, 16, 16, 16, 16, 16, 16, 16, 16, 16, 16, 16, 16, 16, 16, 16, 16, 16,$$
$$16, 16, 16, 16, 16, 16, 16, 16, 16], \text{ and}$$

$$B_4 = [2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 12, 20, 20,$$
$$20, 20, 20, 20, 20, 20, 20, 20, 20, 20, 20, 20, 20, 20, 20, 20, 20, 20, 20,$$
$$20, 20, 20, 20, 20, 20, 20, 20, 20, 20, 20, 20, 20, 20, 20, 20, 20, 20, 20,$$
$$20, 20, 20, 20, 20, 20, 20, 20, 20].$$

We measured many different combinations with different m and μ, running SIMBA-m-μ as described above, averaging the running time over 100 runs per parameter set, given in 10^6 clock cycles . For each B_i, we present the three best combinations we found.

B	1st		2nd		3rd	
0	μ=10 m=5	338.1	μ=10 m=6	343.5	μ=11 m=5	343.7
1	μ=12 m=4	329.3	μ=14 m=4	330.6	μ=13 m=4	330.8
2	μ=11 m=5	326.5	μ=12 m=5	327.0	μ=11 m=4	327.6
3	μ=16 m=4	333.8	μ=17 m=4	337.6	μ=16 m=3	339.3
4	μ=20 m=3	397.5	μ=20 m=4	399.0	μ=21 m=3	399.5

For the best combinations mentioned above, we further tried to rearrange the order of the primes ℓ_i in the loops. As pointed out in [13], it is beneficial to go through the ℓ_i in descending order. However, if we suppress isogeny point evaluations in the last iteration per loop, this means that these savings refer to small ℓ_i, and therefore the impact of this is rather small. Hence, we put a few large primes at the end of the loops, therefore requiring more computational effort for point multiplications, which is however in some situations outweighed by the larger savings from not evaluating points.

In this way, the best combination we found for CSIDH-512 is $\ell_1 = 349$, $\ell_2 = 347$, $\ell_3 = 337$,..., $\ell_{69} = 3$, $\ell_{70} = 587$, $\ell_{71} = 373$, $\ell_{72} = 367$, $\ell_{73} = 359$, and $\ell_{74} = 353$,

using SIMBA-5-11 and B_2, where the B_i are swapped accordingly to the ℓ_i.

In the same setting as in section 7, we measured 322.6 million clock cycles for this combination, which saves 3.9 million clock cycles compared to the results from above.

B Algorithms

In this section we describe our constant-time algorithm, containing the optimizations from above. We split the application of SIMBA in two parts: SIMBA-I splits the isogeny computations in m batches, and SIMBA-II merges them after μ rounds. Note that in our implementation, it is actually not required to generate all the arrays from SIMBA-I.

Algorithm 5 shows the full class group action evaluation. Due to many loops and indices, it looks rather complicated. We recommend to additionally have a look at our implementation, provided in section 6.

Algorithm 3: SIMBA-I.

Input : $e = (e_1, ..., e_n)$, $B = (B_1, ..., B_n)$, m.
Output: $e^{i,iso} = (e_1^{i,iso}, ..., e_n^{i,iso})$, $e^{i,dum} = (e_1^{i,dum}, ..., e_n^{i,dum})$, k_i for $i \in \{0, ..., m-1\}$.

1 Initialize $e^{i,iso} = e^{i,dum} = (0, 0, ..., 0)$ and $k_i = 4$ for $i \in \{0, ..., m-1\}$
2 **foreach** $i \in \{1, ..., 74\}$ **do**
3 $\quad e_i^{i\%m,iso} \leftarrow e_i$
4 $\quad e_i^{i\%m,dum} \leftarrow B_i - e_i$
5 $\quad$ **foreach** $j \in \{1, ..., m\}$ **do**
6 $\quad\quad$ **if** $j \neq (i\%m)$ **then**
7 $\quad\quad\quad k_i \leftarrow k_i \cdot \ell_i$

Algorithm 4: SIMBA-II.

Input : $e^{i,iso} = (e_1^{i,iso}, ..., e_n^{i,iso})$ and $e^{i,dum} = (e_1^{i,dum}, ..., e_n^{i,dum})$ for $i \in \{0, ..., m-1\}$, m.

Output: $e^{iso} = (e_1^{iso}, ..., e_n^{iso})$, $e^{dum} = (e_1^{dum}, ..., e_n^{dum})$, and k.

1 Initialize $e^{iso} = e^{dum} = (0, 0, ..., 0)$, and $k = 4$.
2 **foreach** $i \in \{1, ..., 74\}$ **do**
3 $\quad e_i^{iso} \leftarrow e_i^{i\%m,iso}$
4 $\quad e_i^{dum} \leftarrow e_i^{i\%m,dum}$
5 $\quad$ **if** $e_i^{iso} = 0$ *and* $e_i^{dum} = 0$ **then**
6 $\quad\quad k \leftarrow k \cdot \ell_i$

Algorithm 5: Constant-time evaluation of the class group action in CSIDH-512.

```
    Input  : a ∈ F_p , e = (e_1, ..., e_n), B, m, μ.
    Output: a' such that [l_1^{e_1} ··· l_n^{e_n}]E_a = E_{a'}.
 1  Run SIMBA-I(e, B, m).
 2  foreach i ∈ {1, ..., μ} do
 3      foreach j ∈ {1, ..., m} do
 4          Run Elligator to find a point P, where y_P ∈ F_p.
 5          P ← [k_j]P
 6          S = {ι | e_ι^{m,iso} ≠ 0 or e_ι^{m,dum} ≠ 0}
 7          foreach ι ∈ S do
 8              α = ∏_{κ∈S,κ>ι} ℓ_κ
 9              K ← [α]P.
10              if K ≠ ∞ then
11                  if e_ι^{j,iso} ≠ 0 then
12                      Compute a degree-ℓ_ι isogeny φ : E_a → E_{a'} with
                          ker(φ) = ⟨K⟩.
13                      a ← a', P ← φ(P), e_ι^{j,iso} ← e_ι^{j,iso} − 1.
14                  else
15                      Compute a degree-ℓ_ι dummy isogeny:
16                      a ← a, P ← [ℓ_ι]P, e_ι^{j,dum} ← e_ι^{j,dum} − 1.
17                  if e_ι^{j,iso} = 0 and e_ι^{j,dum} = 0 then
18                      Set k_j = k_j · ℓ_ι.
19  Run SIMBA-II(e^{i,iso} and e^{i,dum} for i ∈ {0, ..., m − 1}, m).
20  while some e_i^{iso} ≠ 0 or e_i^{dum} ≠ 0 do
21      Run Elligator to find a point P, where y_P ∈ F_p.
22      Set P = (x : 1), P ← [k]P, S = {i | e_i^{iso} ≠ 0 or e_i^{dum} ≠ 0}.
23      foreach i ∈ S do
24          Let m = ∏_{j∈S,j<i} ℓ_i.
25          Set K ← [m]P.
26          if K ≠ ∞ then
27              if e_i^{iso} ≠ 0 then
28                  Compute a degree-ℓ_i isogeny φ : E_a → E_{a'} with ker(φ) = ⟨K⟩.
29                  a ← a', P ← φ(P), e_i^{iso} ← e_i^{iso} − 1.
30              else
31                  Compute a degree-ℓ_i dummy isogeny:
32                  a ← a, P ← [ℓ_i]P, e_i^{dum} ← e_i^{dum} − 1.
33              if e_i^{iso} = 0 and e_i^{dum} = 0 then
34                  Set k = k · ℓ_i.
```