GPS and SQISign signature schemes. Notes for isogeny summer school, week 4.

Antonin Leroux

August 19, 2021

The goal of this document is to introduce the two isogeny-based signature schemes GPS and SQISign. Most of the content of these notes is extracted from the articles that first introduced these two constructions ([4] and [1] respectively). Another useful reference is [3] for a good exposition on the various problems involved in the effective Deuring correspondence together with several algorithms and reductions. For more on the subject of quaternion algebra and the `KLPT` algorithm, we direct the reader to the notes written by Laia Amorós and David Kohel for the 2nd week of the isogeny school. Some background and preliminary definitions on identification and signature schemes can be found in the notes of Ward Beullens on SIDH and CSIFiSh (first half of week 4). The rest of these notes is structured in two parts. First, we review all the necessary technical content and algorithms on the effective Deuring correspondence. Then, we introduce the two constructions and study their security and efficiency.

1 The Effective Deuring Correspondence

Throughout this section, we are going to look at supersingular elliptic curves over $\mathbb{F}_{p^2}$. We denote $E(j)$ the curve of j-invariant equal to j and write $\mathcal{B}_{p,\infty}$ for the quaternion algebra ramified at p and infinity.

1.1 Endomorphism Rings and Kernel Ideals

The initial Deuring correspondence can be summarized with Theorem 1, which was first proven by Deuring in [2].

Theorem 1. *The set of maximal orders in $\mathcal{B}_{p,\infty}$ under isomorphisms is in bijection with the set of supersingular j-invariants over $\mathbb{F}_{p^2}$ under galois conjugacy.*

This bijection is explicitly obtained by sending j to the endomorphism ring $\mathrm{End}(E(j))$. In Theorem 1, we are considering j-invariants over $\mathbb{F}_{p^2}$. The only non-trivial element in the galois group of $\mathbb{F}_{p^2}/\mathbb{F}_p$ is the Frobenius morphism. Thus, given a isomorphism class of maximal orders $\mathcal{O}$, there is either one or two j-invariants such that $\mathrm{End}(E(j)) \cong \mathcal{O}$. It is unique when $j \in \mathbb{F}_p$ and the Frobenius $\pi : (x, y) \mapsto (x^p, y^p)$ is an endomorphism of $E(j)$. Otherwise, the two j-invariants are equal to j and j^p.

Remark 1. In these notes, we sometimes abuse notations and assimilate an endomorphism α of a curve E with the corresponding element in $\mathcal{B}_{p,\infty}$ contained in some maximal order $\mathcal{O} \cong \mathrm{End}(E)$. To be technically correct we should have used an explicit bijection between $\mathcal{O}$ and $\mathrm{End}(E)$ every time.

Kernel ideals were first introduced by Waterhouse in [6]. They can be used to extend the original result from Deuring to isogenies and integral ideals of maximal orders. Given an isogeny $\varphi : E \to \star$, we define the associated kernel ideal as

$$I_\varphi = \{\alpha \in \mathrm{End}(E) \ : \ \alpha(P) = 0 \text{ for all } P \in \ker(\varphi)\}.$$

With this notations we can state Theorem 2.

Theorem 2. *Given a supersingular curve E, a maximal order $\mathcal{O}$ such that $\mathrm{End}(E) \cong \mathcal{O}$ and a separable isogeny $\varphi : E \to E'$ of degree D, I_φ is a left integral $\mathcal{O}$-ideal of norm D. The right order $O_R(I_\varphi)$ is isomorphic to $\mathrm{End}(E')$ and I_φ is isomorphic to $\mathrm{Hom}(E', E) \circ \varphi$. Additionally, every integral left $\mathcal{O}$-ideal arises in this way.*

Remark 2. Over $\mathbb{F}_{p^2}$, the only inseparable isogeny is the frobenius. For supersingular curves, it makes no sense to talk about kernel ideals for the frobenius as its kernel is trivial (this is one of the equivalent ways to define a supersingular elliptic curve). Yet, we can still formulate a result similar to Theorem 2 by looking at the prime ideals $\mathfrak{p}$ over p contained in maximal orders (in fact, there is exactly one such ideal by isomorphism class of maximal orders and it corresponds to the frobenius isogeny).

So far, we have explained how to define ideal from isogenies but, as mentioned in Theorem 2, for every integral ideal I there is also a corresponding isogeny φ_I. This isogeny can be defined from its kernel that we write $E[I]$ and sometimes call the kernel of I. It is defined as

$$E[I] = \{P \in E(\overline{\mathbb{F}}_{p^2}) : \alpha(P) = 0 \text{ for all } \alpha \in I\}.$$

Then, we simply take $\varphi_I : E \to {}^{E}\!/_{E[I]}$

The most important part of the Deuring Correspondence is covered with Theorems 1 and 2. There are several smaller notions that we can interpret through this approach. For instance, we saw in Theorem 2 that norm and degree were associated. Similarly, we can link the dual isogeny with the conjugate ideal and isogeny composition with ideal multiplication. Theorem 2 can also be seen through the prism of ideal classes. Since equivalent $\mathcal{O}$-ideals have isomorphic right orders, it is easy to see that they correspond to isogenies between the same pair of curves. Thus, we can put the set of supersingular j-invariants in bijection with the class set $\mathrm{Cl}(\mathcal{O})$ of any maximal order $\mathcal{O}$. This approach underlies the usual formulation of the quaternion ℓ-isogeny path problem introduced in [5] which is different from the one we use for Problem 4. We summarize the main properties of this correspondence in Table 1 from [1].

1.2 Problems, algorithms and reductions

The goal of this section is to look at the algorithmic problems that arises in the context of the Deuring correspondence. There are four problems of importance.

Supersingular j-invariants over $\mathbb{F}_{p^2}$	Maximal orders in $\mathcal{B}_{p,\infty}$
$j(E)$ (up to galois conjugacy)	$\mathcal{O} \cong \mathrm{End}(E)$ (up to isomorpshim)
(E_1, φ) with $\varphi : E \to E_1$	I_φ integral left $\mathcal{O}$-ideal and right $\mathcal{O}_1$-ideal
$\theta \in \mathrm{End}(E_0)$	Principal ideal $\mathcal{O}\theta$
$\deg(\varphi)$	$n(I_\varphi)$
$\hat{\varphi}$	$\overline{I_\varphi}$
$\varphi : E \to E_1, \psi : E \to E_1$	Equivalent Ideals $I_\varphi \sim I_\psi$
Supersingular j-invariants over $\mathbb{F}_{p^2}$	$\mathrm{Cl}(\mathcal{O})$
$\tau \circ \rho : E \to E_1 \to E_2$	$I_{\tau\circ\rho} = I_\rho \cdot I_\tau$

Table 1. The Deuring correspondence, a summary.

The first one to appear in isogeny-based cryptography was the supersingular isogeny path problem (Problem 2) (in fact the assumed-hardness of this problem is the main motivation behind isogeny-based cryptography). However, given Proposition 2, the endomorphism ring problem (Problem 1) is now considered as the fundamental hard problem of isogeny-based cryptography.

Problem 1. Given a curve E, find a compact representation of $\mathcal{O}$.

Remark 3. It is not completely clear what is meant by compact representation and there are several possible definitions. Here, we are looking for a basis of $\mathcal{O}$ over the basis $1, i, j, k$ of $\mathcal{B}_{p,\infty}$ with coefficients of polynomial size in $\log(p)$. It was shown in [3] that such a representation always exists.

Problem 2. Given two curves E_1, E_2, find an isogeny $\varphi : E_1 \to E_2$ of given degree D.

Problems 3 and 4 are simply the translation of Problems 1 and 2 through the Deuring correspondence.

Problem 3. Given a maximal order $\mathcal{O} \subset \mathcal{B}_{p,\infty}$, find a curve E with $\mathrm{End}(E) \cong \mathcal{O}$.

Problem 4. Given two maximal orders $\mathcal{O}_1, \mathcal{O}_2$, find an ideal I of norm D with $\mathcal{O}_L(I) \cong \mathcal{O}_1$ and $\mathcal{O}_R(I) \cong \mathcal{O}_2$.

When $D = \ell^n$, Problem 4 is called the Quaternion ℓ-isogeny path problem, it was the focus of [5] where the `KLPT` algorithm was introduced to solve it in polynomial time. More generally, it was shown in [4] that `KLPT` could be adapted to find a solution of power-smooth norm. The `KLPT` algorithm is the most important building block for the algorithms and reductions introduced in [3] and for the signature constructions that interests us. The second crucial algorithm is the one arising from Problem 5 below.

Problem 5. Given an integral left $\mathcal{O}$-ideal I of norm D, a curve E such that $\mathrm{End}(E) \cong \mathcal{O}$ and the action of $\mathrm{End}(E)$ on $E[D]$, find the corresponding isogeny $\varphi_I : E \to E/E[I]$.

In Section 1.3, we are going to introduce an algorithm `IdealToIsogeny` solving Problem 5. For now, let us assume that it exists and works in polynomial time for power-smooth norm D. For Proposition 1 we also assume that there exists one curve E_0 of known endomorphism ring $\mathcal{O}_0$ (in practice this is always the case, for instance when $p = 3 \mod 4$ we can take the curve $y^2 = x^3 + x$ with $\mathrm{End}(E) \cong \langle 1, i, \frac{1+j}{2}, \frac{i+k}{2} \rangle$).

Proposition 1. *Problem 3 can be solved in polynomial time.*

Proof. (sketch) We use `KLPT` between $\mathcal{O}_0$ and $\mathcal{O}$ to solve Problem 4 and find an ideal of power-smooth norm connecting $\mathcal{O}_0$ and $\mathcal{O}$. Then, we apply `IdealToIsogeny` to compute the corresponding isogeny between E_0 and a curve E. The codomain E of this isogeny is the answer.

For our next result we need to solve the inverse of Problem 5 which is Problem 6. We will write `IsogenyToIdeal` for the algorithm solving this problem.

Problem 6. Given an isogeny $\varphi : E \to E'$ of degree D, a maximal order $\mathcal{O}$ such that $\mathrm{End}(E) \cong \mathcal{O}$ and the action of $\mathrm{End}(E)$ on $E[D]$, find the corresponding $\mathcal{O}$-ideal I_φ.

Proposition 2. *Problem 1 and Problem 2 are equivalent.*

Proof. (sketch) To find an isogeny between E and E' it suffices to apply the solver for Problem 1 to find the endomorphism rings $\mathcal{O}$ and $\mathcal{O}'$ of E and E'. Then, with `KLPT`, we can find an ideal of norm D connecting these two maximal orders. The corresponding isogeny can be computed using `IdealToIsogeny` and is a suitable solution to Problem 2.

Conversely, we can find the endomorphism ring of a curve E from a solution to Problem 2 when E' is equal to E_0. Then, we can translate the dual of this isogeny to the corresponding ideal with `IsogenyToIdeal`. The desired maximal order is simply the right order of this ideal.

Remark 4. In the proof of Proposition 2, we omitted to explain how we can compute the action of $\mathrm{End}(E)$ on $E[D]$ for the first reduction. A brief discussion on the problem of computing the action of endomorphisms on torsion subgroups can be found in Section 1.3.

All the algorithms and reductions mentioned above were introduced in [5,4,3] but some of the proofs are only heuristic. Very recently, Wesolowski in [7] filled this gap by introducing a provable version of `KLPT` and used it to prove formally Proposition 2.

1.3 Ideal to isogeny translation: an algorithm

In this section, we study the algorithm `IdealToIsogeny` that we mentioned in Section 1.2. As formulated in Problem 5, one of the prerequisite for this algorithm is the action of the endomorphism ring on the D-torsion. In the end of this

section, we try to outline how one can compute this action for any endomorphism ring. This is also the purpose of Exercise 3.4. For now, let us just assume that it is given.

We focus on the case of cyclic isogenies as it is the relevant one for isogeny-based cryptography. As usual when the degree is power-smooth, we can compute the desired isogeny from a generator of its kernel (which is equal to $E[I]$). If we recall the definition of $E[I]$, it is easy to see that a generator is simply a point of the correct order sent to 0_E by every element of I. We obtain the following algorithm:

1. Select $\alpha \in I$ such that $\gcd(n(\alpha), D^2) = D$.
2. Select a basis P, Q of $E[D]$.
3. Compute $\alpha(P), \alpha(Q)$.
4. If $\alpha(P)$ has order D find a such that $\alpha(Q) = [a]\alpha(P)$. If not, swap P and Q and try again.
5. Output the isogeny of kernel generated by $Q - [a]P$.

The fact that α of the correct norm can be found in Step 1 is simply a consequence of the definition of the norm of an ideal. The norm condition $\gcd(n(\alpha), D^2) = D$ implies that either $\alpha(P)$ or $\alpha(Q)$ has order D. Finally, since α is in I, we know there must be a subgroup sent to 0_E by α and so the DLP in Step 4 has a solution. This proves that the above algorithm terminates.

The complexity depends mainly on the value of D. First, it it clear that we must have a polynomial-size representation of points in the D-torsion. Then, as we assumed that we can evaluate α on the D-torsion efficiently, the remaining hard operations are the DLP in the subgroup of order D and the D-isogeny computation. When D is power-smooth everything is efficient and we obtain a polynomial-time algorithm. Overall, as soon as D is smooth and the D-torsion is defined over an extension of polynomial degree, our algorithm is polynomial-time.

We conclude this section by talking about endomorphism ring explicit representation and computation of the action of endomorphisms on torsion points. More explicitly, given E and $\mathcal{O} \cong \mathrm{End}(E)$, we need a concrete basis such that $\mathcal{O} = \langle \omega_1, \omega_2, \omega_3, \omega_4 \rangle$, where each ω_i corresponds to an endomorphism $\rho_i \in \mathrm{End}(E)$ that can be efficiently evaluated on any point. In full generality, we only how to do this in special cases such as $j = 1728$ where there exists some endomorphisms with very simple expression. We now describe a way to do that for other endomorphisms using an approach introduced in [3]. The method is based on the existence of a special curve E_0 admitting an explicit representation of $\mathrm{End}(E_0) \cong \langle \omega_1^0, \omega_2^0, \omega_3^0, \omega_4^0 \rangle$ and an isogeny $\varphi : E_0 \to E$ of degree N_φ. The ideal I_φ is a left $\mathcal{O}_0$-ideal and right $\mathcal{O}$-ideal with $\mathcal{O} \simeq \mathrm{End}(E)$. Since I_φ is integral, it is contained in both $\mathcal{O}_0$ and $\mathcal{O}$. From there, it is easy to see that $N_\varphi \mathcal{O} \subset I \subset \mathcal{O}_0$. We will use that fact to represent and compute elements of $\mathcal{O}$. An element $\alpha \in \mathcal{O}$ can be written as an element of $\frac{\mathcal{O}_0}{N_\varphi}$ with $\alpha = \frac{1}{N_\varphi} \sum_{i=1}^4 a_i \omega_i^0$ with $a_i \in \mathbb{Z}$ for $i \in \{1, 2, 3, 4\}$. Using that, it is possible to evaluate an endomorphism α at a point P of order coprime with N_φ as $\alpha(P) = \frac{1}{N_\varphi^2} \sum_{i=1}^4 [a_i] \varphi \circ \rho_i \circ \hat{\varphi}(P)$. Exercise 3.4 is focused on the problem of evaluating endomorphisms on torsion points.

2 Signatures from Endomorphism Ring Proof of Knowledge: GPS and SQISign

The conclusion of the story that we tried to tell in Section 1.2 is that everything becomes easy when the inputs are expressed in the world of quaternions (mainly because we can apply `KLPT` and `IdealToIsogeny`). Thus, as endomorphism rings (through Problem 1) are the keys to go from elliptic curves to quaternions, we can see them as trapdoors. The quite natural question from this observation is: can we use that principle to show that we know the endomorphism ring of a supersingular elliptic curve ? The first attempt to achieve this idea led to the GPS signature scheme [4]. The overall principle (which is shared by SQISign) is to use the `KLPT` algorithm to show that one knows the endomorphism ring of some public key curve.

2.1 GPS

The GPS signature is obtained by applying the fiat-shamir transform to the repetition of a two-special sound interactive identification protocol that we describe below. As we explained, this identification scheme is based on endomorphism ring proof-of-knowledge and it follows the sigma protocol standard framework. The public key is some supersingular curve E_A and the secret key is $\mathcal{O} \cong \mathrm{End}(E_A)$ (or equivalently a secret isogeny τ from the public curve E_0 to E_A).

Commitment The prover generates a random isogeny walk $\sigma_0 : E_A \to E_1$ of degree D_c, and sends E_1 to the verifier.

Challenge The verifier sends a bit $b \in \{0, 1\}$ to the prover.

Response If $b = 0$, the prover reveals σ_0 to the verifier. Otherwise he uses the secret key to compute another isogeny $\sigma_1 : E_0 \to E_1$ of degree D and reveal it to the verifier.

Verification The verifier accepts if the received isogeny σ_b is from E_A to E_1 and has degree D_c when $b = 0$ or from E_0 to E_1 and has degree D when $b = 1$. They reject otherwise.

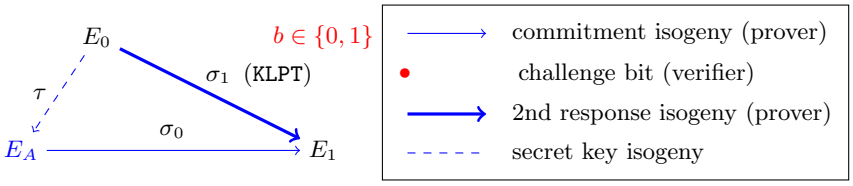

Fig. 1. A picture of GPS's identification protocol

For this protocol, the technical challenge lies in the computation of σ_1. After reading Section 1.2, the method should be clear to the reader. The isogeny σ_1 is

going to be translated with `IdealToIsogeny` from an ideal obtained by applying `KLPT` on $\mathcal{O}_0 \cong \mathrm{End}(E_0)$ and $\mathcal{O}_1 \cong \mathrm{End}(E_1)$. In order to apply the `KLPT` algorithm, we need to compute $\mathcal{O}_1$, which can be done by computing I_{σ_1} (and its right order) with `IsogenyToIdeal`. Note that if one would rather use `IdealToIsogeny` for this step as well, one could go the other way around and start by generating I_{σ_1} before translating it into an isogeny.

In terms of security, there are two important properties: soundness and honest-verifier zero-knowledge. Intuitively, the first one ensures that the verification cannot be satisfied without the prover knowing the secret while the second one ensures that the proof does not leak too much information on the secret key to the verifier. More precisely, to prove 2-special soundness, one needs to show that if two correct answers are revealed for the same commitment but different challenges, then the secret can be recovered. For GPS, two correct answers for different challenges imply that the path $\hat{\sigma}_0 \circ \sigma_1 : E_0 \to E_A$ is revealed. As we explained in the sketch of proof for Proposition 2, this is enough for anyone to compute the endomorphism ring of E_A. As such, the soundness holds under the hardness of Problem 1. For honest-verifier zero-knowledge, we need to argue that a valid transcript can be generated by a simulator with the sole knowledge of the public key. When $b = 0$, it is clear that generating a random σ_1 can be done by anyone knowing E_A. When $b = 1$, the proof is a bit more tricky. It uses the fact that the ideal obtained from `KLPT` only depends on the class of equivalence of the ideal given in input (and not the actual representative of this class). As such, a simulator can simply generate an ideal belonging to a random class and then apply `KLPT` on it.

In terms of efficiency, the GPS signature scheme suffers from one major flaw which is that the underlying identification scheme has only a challenge space of size 2. Intuitively, it means that anyone has $1/2$ chance of successfully cheating (by guessing the correct challenge). This means that to obtain λ bits of security, we need to repeat λ times this protocol. Furthermore, this repetition also implies that the signature will not be compact (in the end, the size is going to be quadratic in λ). Additionally, at the time of GPS, it was not clear how to efficiently instantiate this protocol (and the `IdealToIsogeny` part in particular). Indeed, the problem is that the degree D of solutions obtained with `KLPT` are roughly equal to p^3. This is too big to ensure that D can have a very small smoothness bound while having the D-torsion defined on a small fields extension which are the two conditions for `IdealToIsogeny` too be efficient. As a result, there hadn't been any implementation of GPS.

2.2 SQISign

SQISign is built on a principle similar to GPS but improves drastically the efficiency by addressing the issues we outlined in the previous section. The only real downside compared to GPS is security (honest-verifier zero-knowledge in particular) which is based upon a new ad hoc assumption.

Before introducing the underlying identification scheme, we try to explain and motivate the main difference between SQISign and GPS which is the soundness of the underlying identification scheme. More specifically, for GPS the low soundness is due to the fact that the challenge space has only size two. To improve the soundness, we need to have a challenge space that can be arbitrarily big. Before explaining how to do that, let us go back a little bit and try to analyze why GPS was designed in that way. In fact, it stems from a limitation of the `KLPT` algorithm from [5]. The main contribution of [5] is an algorithm that solves Problem 4 when one of the orders is a special order (that we denoted $\mathcal{O}_0$ throughout these notes). From there, one can solve the generic quaternion path problem between $\mathcal{O}_1$ and $\mathcal{O}_2$ by applying the special case algorithm twice on $\mathcal{O}_0, \mathcal{O}_1$ and $\mathcal{O}_0, \mathcal{O}_2$ before combining the solutions together. While this idea is sufficient to solve Problem 4, it is not generic enough for our application. Explaining why this method is problematic is the purpose of Exercise 3.3. As a consequence, we are stuck to use the special case `KLPT` algorithm which does not seem to allow us to devise a much better protocol than GPS. To go beyond that, we are in need of a new way to solve Problem 4 in the generic case. In reality, the reasoning that we just unfolded is exactly what lead to SQISign: use a new generalized `KLPT` algorithm in order to do the more efficient protocol presented below. This new algorithm is beyond the scope of these short notes, the full details can be found in [1].

Commitment The prover generates a random (secret) isogeny walk $\psi : E_0 \to E_1$, and sends E_1 to the verifier.

Challenge The verifier sends the description of a cyclic isogeny $\varphi : E_1 \to E_2$ of degree D_c to the prover.

Response From the isogeny $\varphi \circ \psi \circ \hat{\tau} : E_A \to E_2$, the prover constructs a new isogeny $\sigma : E_A \to E_2$ of degree D such that $\hat{\varphi} \circ \sigma$ is cyclic, and sends σ to the verifier.

Verification The verifier accepts if σ is an isogeny of degree D from E_A to E_2 and $\hat{\varphi} \circ \sigma$ is cyclic. They reject otherwise.

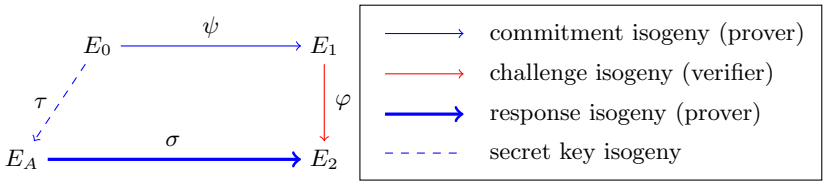

Fig. 2. A picture of SQISign's identification protocol

Similarly to GPS, the difficult operation in the above scheme is the computation of the answer isogeny σ. It can be done in the following way: the prover computes the ideals corresponding to ψ and φ with `IsogenyToIdeal`, and use

them to compute the maximal order $\mathcal{O}_2$ isomorphic to $\mathrm{End}(E_2)$. Then, the prover applies the new generalized `KLPT` on $\mathcal{O}_2$ and the endomorphism ring of E_A to obtain an ideal connecting them. The answer σ is the isogeny corresponding to this ideal and is obtained by applying `IdealToIsogeny`.

For security, we look again at soundness and honest-verifier zero-knowledge. For the former, we study answers to different challenges under the same commitment. From the diagram in Fig. 2, we see that this yields two distinct isogenies from E_1 to E_A. By combining these paths, we obtain a non-trivial endomorphism of E_A. Thus, soundness holds under the hardness of the endomorphism computation problem which is heuristically assumed to be as hard as the full endomorphism ring computation problem (see [3] for more details). Since the challenge space is now the set of isogenies of degree D_c, it can be expanded arbitrarily by increasing the size of D_c. Thus, we can obtain λ bits of soundness with just one iteration of the protocol. Unfortunately, zero-knowledge is harder to obtain in the case of SQISign as we need to argue that a simulator can produce from the public key a valid transcript that is indistinguishable from a real one. Without the knowledge of $\mathrm{End}(E_A)$, there is little more a simulator can do than generating a random isogeny of the desired degree. Unfortunately, the isogenies σ obtained from the generic KLPT are far from being random ones. Thus, zero-knowledge of SQISign's identification scheme is basically based on the assumption that an isogeny obtained by applying `IdealToIsogeny` to the output of the generic `KLPT` is indistinguishable from random. In fact, the authors obtain a slightly more generic assumption by showing that σ is in fact distributed as a random isogeny among a set of isogenies of degree D (which highly depends on the new generalized `KLPT` algorithm). Yet, in the end, the assumption remains very ad hoc and highly non-standard. For now, there are no known ways to reduce this problem to one of the classical isogeny-based assumptions.

Efficiency improvement was one of the main motivation behind SQISign. The high-soundness of the new identification protocol is already a huge leap forward as the identification scheme only needs to be executed once. The size of σ is only slightly bigger than the one of σ_1 for GPS ($15/4\log(p)$ against $3\log(p)$) so one iteration of the SQISign identification scheme is only marginally slower than one iteration of the GPS identification scheme. In both cases, the major bottleneck is the execution of `IdealToIsogeny` and so the second important contribution of the SQISign paper [1] is a new algorithm to perform that step efficiently when the degree D is smooth but the D-torsion is not necessarily defined over a small extension. The trick is to factorize D as $D_1 D_2 \dots D_n$ where each of the D_i is such that we can apply the algorithm introduced in Section 1.3. This principle bring some complications as we need to compute the action of several endomorphism rings on the D_i torsion (and the D_i are not necessarily pairwise coprime) but it can be done quite efficiently in the end using (once again) the `KLPT` algorithm (all the details can be found in [1]). The solution of Exercise 3.4 is a first step toward the final method introduced in [1].

A concrete instantiation of SQISign was also part of [1]. To obtain fast verification, the authors choose to take the degree D of σ to be a power of two (heuristically it was estimated in [1] that taking $D = 2^{1000}$ should be enough). Then, to apply the new efficient `IsogenyToIdeal`, the most important requirement is to use a special prime p such that $p^2 - 1 = 2^f Tn$ where $T \sim p^{3/2}$ is an odd smooth number (this requirement stems from the `IdealToIsogeny` improvement mentioned above) and f is as big as possible (that factorization of $p^2 - 1$ implies that the $2^f T$ torsion is defined over $\mathbb{F}_{p^2}$). To reach NIST-1 level of security, it is argued in [1] that a p of 256 bits should be enough. With that size, SQISign is by far the most compact post-quantum signature scheme targetting NIST-1 level of security. The key sizes can be found in Table 2. The SQISign paper exhibits an example of such a prime where $f = 33$, $\log T = 395$ and T is 2^{13}-smooth. This prime was used in a C implementation whose performances are reported in Table 3. The code is available at https://github.com/SQISign/sqisign.

Secret Key (bytes)	Public Key (bytes)	Signature (bytes)
16	64	204

Table 2. Size of SQISign keys and signature for the NIST-1 level of security.

		Keygen	Sign	Verify
ms	1st quartile	564	2,256	41
	median	575	2,279	42
	3rd quartile	587	2,321	43

Table 3. Performance of SQISign in millions of cycles and in milliseconds. Statistics over 100 runs for key generation and signature, and over 250 runs for verification.

3 Exercices

3.1 An example of ideal to isogeny translation.

Take $p = 163$, and let E_0 be the supersingular curve $y^2 = x^3 + x$ over $\mathbb{F}_{p^2}$. Then, we known that $\mathcal{B}_{p,\infty} = \mathbb{Q}\langle 1, i, j, k\rangle$ with $i^2 = -1$, $j^2 = -163$ and $k = ij$. The endomorphism ring of E_0 is isomorphic to $\langle 1, i, \frac{1+j}{2}, \frac{i+k}{2}\rangle$ through two endomorphisms $\iota : (x, y) \mapsto (-x, \sqrt{-1}y)$ and $\pi : (x, y) \mapsto (x^p, y^p)$ (respectively corresponding to i and j). Find a generator of the kernel of $I = \langle 2, 1 + i, i + k, \frac{3+i+j+k}{2}\rangle$ and compute the codomain of the corresponding isogeny.

3.2 Isogeny to ideal translation

Using the `IdealToIsogeny` algorithm described in Section 1.3 as inspiration, propose an algorithm `IsogenyToIdeal` to solve Problem 6 when the degree D is power-smooth.

3.3 Attack on a broken version of SQISign

This exercise uses some of the notations given in the beginning of Section 2.2. Assume that there exists a `KLPTSpecial` algorithm to solve Problem 4 in the special case where one of the orders is a special order $\mathcal{O}_0$, derive (and present it step by step) a simple algorithm `KLPTGeneric` to solve Problem 4 for any maximal orders. We want ideals of norm a power of ℓ for some small prime ℓ and no additional constraint. Describe an attack to recover the endomorphism ring of the public key E_A when SQISign is instantiated with this algorithm `KLPTGeneric`.

3.4 Endomorphism ring action on torsion points

Given $\varphi : E_0 \to E$ an isogeny of degree D (the D-torsion is assumed to be defined over $\mathbb{F}_{p^2}$) where E_0 is a special curve of known endomorphism, imagine a polynomial-time algorithm (using `KLPT`, `IsogenyToIdeal` and `IdealToIsogeny`) to compute the action of $\mathrm{End}(E)$ on $E[D]$.

References

1. De Feo, L., Kohel, D., Leroux, A., Petit, C., Wesolowski, B.: Sqisign: compact post-quantum signatures from quaternions and isogenies. In: International Conference on the Theory and Application of Cryptology and Information Security. pp. 64–93. Springer (2020)
2. Deuring, M.: Die Typen der Multiplikatorenringe elliptischer Funktionenkörper. Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg **14**(1), 197–272 (Dec 1941)
3. Eisenträger, K., Hallgren, S., Lauter, K., Morrison, T., Petit, C.: Supersingular isogeny graphs and endomorphism rings: Reductions and solutions. In: Nielsen, J.B., Rijmen, V. (eds.) Advances in Cryptology – EUROCRYPT 2018. pp. 329–368. Springer International Publishing, Cham (2018)
4. Galbraith, S.D., Petit, C., Silva, J.: Identification protocols and signature schemes based on supersingular isogeny problems. In: ASIACRYPT (2017)
5. Kohel, D., Lauter, K.E., Petit, C., Tignol, J.P.: On the quaternion ℓ-isogeny path problem. IACR Cryptology ePrint Archive **2014**, 505 (2014)
6. Waterhouse, W.C.: Abelian varieties over finite fields. Annales scientifiques de l'École Normale Supérieure **2**(4), 521–560 (1969)
7. Wesolowski, B.: The supersingular isogeny path and endomorphism ring problems are equivalent. Cryptology ePrint Archive, Report 2021/919 (2021), https://eprint.iacr.org/2021/919

A Corrections

Correction of Exercise 3.1 We apply the algorithm outlined in Section 1.3 on $\alpha = 1 + i$. With $n(\alpha) = 2$, we have a suitable candidate and it suffices to find the point of two torsion P sent to 0 by α. This point is such that $P = \iota(P)$ with $\iota : (x, y) \mapsto (\sqrt{-1}x, -y)$. It is easy to see that $P = (0, 0)$ is the point we are looking for. Then, the codomain of the isogeny can be computed using the usual Vélu formula. In fact, this curve is isomorphic to E_0 as can be verified by computing its j-invariant. This last fact, could have been proven without any isogeny computation just by realizing that since $n(\alpha) = 2 = n(I)$ we must have $I = \mathcal{O}_0\alpha$, a principal ideal that corresponds to an endomorphism through the Deuring Correspondence (see Table 1).

Correction of Exercise 3.2 Since the degree D is powersmooth, the D torsion is going to be defined over a small $\mathbb{F}_p$-extension and operations such as D-isogeny computations and DLP over the D-torsion are efficient. As for Section 1.3, we assume that we have the ability to evaluate efficiently D-torsion points through any endomorphism of the base curve E. We write $\mathcal{O} = \langle \alpha_1, \alpha_2, \alpha_3, \alpha_4 \rangle \cong \mathrm{End}(E)$ and $\varphi : E \to E'$ is the isogeny given in input. We also assume that a generator P of $\ker \varphi$ is given (if not, it could have easily been computed by solving a few DLP).

1. Compute $\alpha_1(P), \alpha_2(P), \alpha_3(P), \alpha_4(P)$
2. Find x_1, x_2, x_3, x_4 such that $\sum_{i=1}^{4} [x_i]\alpha_i(P) = 0$ and $\gcd(n(\sum_{i=1}^{4} x_i\alpha_i), D^2) = D$.
3. Output $I = \mathcal{O}D + \mathcal{O}\sum_{i=1}^{4} x_i\alpha_i$.

There are several ways of finding a suitable linear combination. A possibility is to select α_i, α_j such that $\alpha_i(P), \alpha_j(P)$ is a basis of $E[D]$ and then express $\alpha_k(P)$ (with $k \neq i, k \neq j$) in this basis by solving a bi-dimensional DLP. Since $\alpha_1, \alpha_2, \alpha_3, \alpha_4$ is a basis of $\mathrm{End}(E)$, it is easy to see that there always exists such indices i, j. The gcd condition should be satisfied with good probability.

Correction of Exercise 3.3 First, we describe the `KLPTGeneric` algorithm. The special maximal order $\mathcal{O}_0$ is such that we can apply `KLPTSpecial`. The two maximal orders in input are $\mathcal{O}_1, \mathcal{O}_2$.

1. Compute $I_1 = \texttt{KLPTSpecial}(\mathcal{O}_0, \mathcal{O}_1)$ and $I_2 = \texttt{KLPTSpecial}(\mathcal{O}_0, \mathcal{O}_2)$.
2. Output $J = \overline{I_1} \cdot I_2$.

The attack on SQISign works in the following way: given $\sigma : E_A \to E_2$ an isogeny computed from an ideal J obtained as the output of `KLPTGeneric`, decompose it as $\varphi_1 \circ \varphi_2$ where $\varphi_1 : E_A \to E_0$ and $\varphi_2 : E_0 \to E_2$. Use the dual of φ_1 to compute $\mathrm{End}(E_A)$ as described in the proof of Proposition 2.

Correction of Exercise 3.4 The idea is to use the effective representation for End(E) described in Section 1.3. The main obstacle is that this representation requires scalar division which is not well-defined in general. In fact, it only make sense to divide a point P by a scalar n when the order of P and n are coprime. Thus, if we want to apply the formula on D-torsion point we need another isogeny $\psi : E_0 \to E$ of degree coprime with D. To do that we proceed in the following way: convert φ into the corresponding ideal I_φ with `IsogenyToIdeal` and use I_φ to apply `KLPT` in order to obtain $J \sim I_\varphi$ of coprime power-smooth degree $D'0$. Then, execute `IdealToIsogeny` on J to obtain $\psi : E_0 \to E'$ of degree D'. Compute the inverse λ of $D' \mod D$. Then, use the formula to compute given in Section 1.3 to evaluate any endomorphism of End(E) on the D-torsion.