Supplementary Algorithms in Supersingular Reductions

April 20, 2019

Abstract

We look over some of the algorithms used in [EHL+18]. These algorithms appear in [KV12, KLPT14, GPS17].

1 Smoothing Ideals in a Quaternion Algebra

In this section we summarize the main algorithm in [KLPT14].

Let $\mathcal{O}$ be a maximal order in a quaternion algebra B over $\mathbb{Q}$ with prime discriminant p (so $\mathbb{Q}$ ramifies at p and ∞), and let ℓ be a small prime.

Problem 1. Given an left $\mathcal{O}$-ideal I, find a representative $I\beta$ in it's left ideal class with ℓ-power norm.

We first apply a few reductions to a special case. First write $B = \mathbb{Q}\langle i, j\rangle$ with $i^2 = -q$ and $j^2 = -p$. Here q (which may be 1 or a small prime) is chosen so that B ramifies exactly at p and ∞. Let R be the ring of integers in $\mathbb{Q}[i] = \mathbb{Q}(\sqrt{-q})$. Then $R + Rj$ is a (possibly non-maximal) order in B.

1. We may assume that $\mathcal{O}$ is p-extremal[1] and contains $R + Rj$.
2. We may assume that I has (possibly large) prime reduced norm N
3. We may assume that N is coprime to ℓ, p, and $|\mathrm{disc}(R)|$.
4. We may assume that I is given by a pair of generators of the form (N, α).[2]

Our method for finding β is as follows:

1. Find any random $\gamma \in \mathcal{O}$ with $\mathrm{nrd}(\gamma) = N\ell^{e_0}$.
2. Solve for $[\mu] \in (\mathcal{O}/N\mathcal{O})^\times$ such that $\mathcal{O}[\gamma][\mu] = [I]$.[3]

[1] This means $\mathcal{O}$ contains a square root of $-p$ and also has trivial two-sided class number. Equivalently, $\mathcal{O}$ is isomorphic to the endomorphism ring of a supersingular curve with j-invariant in $\mathbb{F}_p$.

[2] See exercise 16.5 of Voight's book and [GPS17, Sec. 2.1].

[3] We use $[\mu]$ instead of the usual quotient notation $\overline{\mu}$ to distinguish residue modulo N from conjugation.

3. Lift $[\mu]$ to some $\mu \in \mathcal{O}$ with $\mathrm{nrd}(\mu) = \ell^{e_1}$ (strong approximation).

After these steps, we set $\beta = \overline{\mu\gamma}/N$. It is clear that $\mathrm{nrd}(I\beta) = \mathrm{nrd}(I)\mathrm{nrd}(\beta) = \ell^{e_0+e_1}$. To see that $I\beta \subseteq \mathcal{O}$, notice that $\mathcal{O}[\gamma][\mu] = I \Rightarrow \gamma\mu \in I \Rightarrow \overline{\mu\gamma} \in \overline{I}$.

$$I\beta = I\frac{\overline{\mu\gamma}}{N} \subseteq \frac{I\overline{I}}{N} = \mathcal{O}.$$

The last step follows from 16.6.14 of Voight's quaternion book.

1.1 Finding γ

In this section, we focus on finding a random $\gamma \in \mathcal{O}$ with specified reduced norm M.

Recall that $B = \mathbb{Q}\langle i, j\rangle/\{i^2 = -q, j^2 = -p\}$ and our maximal order $\mathcal{O}$ contains $R + Rj$, where R is the ring of integers of $\mathbb{Q}[i] = \mathbb{Q}(\sqrt{-q})$.

It follows from the relation $ij = -ji$ that $j\alpha = \overline{\alpha}j$ for any $\alpha \in \mathbb{Q}[i]$. If $\alpha, \beta \in R$, then

$$\begin{aligned}
\mathrm{nrd}(\alpha + \beta j) &= (\alpha + \beta j)\left(\overline{\alpha} + \overline{j\beta}\right) \\
&= (\alpha + \beta j)\left(\overline{\alpha} - \beta j\right) \\
&= \alpha\overline{\alpha} - \alpha\beta j + \beta j\overline{\alpha} - \beta j\beta j \\
&= \alpha\overline{\alpha} - \alpha\beta j + \beta\alpha j - \beta\overline{\beta}j^2 \\
&= \mathrm{nrd}(\alpha) + \mathrm{nrd}(\beta)p.
\end{aligned}$$

Algorithm 1 Sample integer with given reduced norm

Require: A "nice" maximal order $\mathcal{O}$, integer M.
Ensure: Element $\gamma \in \mathcal{O}$ with $\mathrm{nrd}(\gamma) = M$.

```
repeat
    β ← random element of R.
    r ← M − nrd(β)p.
until r splits into principal ideals in R.
α ← generator of a prime in R lying over r.
return γ = α + βj.
```

If R has small discriminant, then the class number will be small so most ideals will be principal. Also, the norm form on R will only include small coefficients. So the run time is determined by how often many β we have to sample. Essentially, we are choosing integers x, y hoping that some quadratic polynomial $f(x, y)$ is prime. If x, y are chosen in a box $[0, m]^2$, then we expect to iterated the inner loop $\log(|M - m^2p|)$ times. The asking for a prine that splits into principle ideals adds a factor determined by the class number of $\mathbb{Q}\otimes R$ (which will be small).

1.2 Finding $[\mu]$

Recall that we are assuming that $\mathcal{O}$ contains R, the ring of integers of $\mathbb{Q}[i] = \mathbb{Q}(\sqrt{-q})$ for a small q. Since N is large, we may assume that N is coprime to $\mathrm{disc}(R)$. Then one can show that $\mathcal{O}/N\mathcal{O} \cong (R+Rj)/N(R+Rj) \cong M_2(\mathbb{Z}/N\mathbb{Z})$. These isomorphisms can be made explicit. Because $I = (\alpha, N)$, it follows that $[I]$ is generated by $[\alpha]$. Let $A_{[\alpha]}$ be the matrix in $M_2(\mathbb{Z}/N\mathbb{Z})$ corresponding to $[\alpha]$ under the choice of isomorphism.

We are looking for $[\mu] \in (\mathcal{O}/N\mathcal{O})^\times$ such that $\mathcal{O}[\gamma][\mu] = [I]$. Since we have γ, and $[I] = [\alpha]$, we are essentially solving a linear equation over $\mathbb{Z}/N\mathbb{Z}$: $A_{[\gamma]} \cdot X = A_{[\alpha]}$. The solution can be turned back into an element of $\mathcal{O}/N\mathcal{O}$.

Remark 2. Note that $A_{[}\gamma]$ and $A_{[\alpha]}$ are rank 1 matrices (because their reduced norm lies in $N\mathcal{O}$).

Algorithm 2 Sample residue class with restriction

Require: A "nice" maximal order $\mathcal{O}$, $\gamma \in \mathcal{O}$, and left $\mathcal{O}$-ideal $I = (\alpha, N)$.
Ensure: Element $\mu \in (\mathcal{O}/N\mathcal{O})^\times$ with $\mathcal{O}[\gamma][\mu] = [I]$.
1: $\phi \leftarrow$ explicit isomorphism $\mathcal{O}/N\mathcal{O} \to M_2(\mathbb{Z}/N\mathbb{Z})$.
2: $A \leftarrow$ solution in $M_2(\mathbb{Z}/N\mathbb{Z})$ to $\phi([\gamma])X = \phi([\alpha])$.
3: $[\mu] \leftarrow \phi^{-1}(A)$.
4: **return** $[\mu]$.

With slightly more care, we can even choose $[\mu]$ in $(R/NR)^\times[j]$.

1.3 Lifting $[\mu]$

Assume now we have $\gamma \in \mathcal{O}$ with reduced norm $N\ell^{e_0}$ and $[\mu] \in (\mathcal{O}/N\mathcal{O})^\times$ with $\mathcal{O}[\gamma][\mu] = [I]$. We further assume that $[\mu] \in (R/NR)^\times[j]$. Our goal is to lift $[\mu]$ to some $\mu \in \mathcal{O}$ with ℓ power reduced norm.

Let $\mu_0 \in R+Rj$ be any lift of $[\mu]$. If ω is a generator for R, then we can write $\mu_0 = x_0+y_0\omega+(z_0+w_0\omega)j$. Our goal is to find $\lambda \in \mathbb{Z}$ and $\mu_1 \in R+Rj$ such that $\mu = \lambda\mu_0 + N\mu_1$ has ℓ-power reduced norm. Write $\mu_1 = x_1 + y_1\omega + (z_1 + w_1\omega)j$.

Recall that we assumed $[\mu] \in (R/NR)^\times[j]$. This means we can choose μ_0 such that $x_0 = y_0 = 0$. Then

$$\mathrm{nrd}(\mu) = N^2 f(x_1, y_1) + pf(\lambda z_0 + Nz_1, \lambda w_0 + Nw_1),$$

where f is the norm form on R with respect to the basis $\{1, \omega\}$. Then

$$\mathrm{nrd}(\mu) \equiv p\lambda^2 f(z_0, w_0) \mod N.$$

Setting this equal to ℓ^e, for an appropriate choice of e, allows us to find λ by taking a squareroot in $\mathbb{Z}/N\mathbb{Z}$.

Next consider the equation

$$\ell^e - pf(\lambda z_0 + Nz_1, \lambda w_0 + Nw_1) \equiv 0 \mod N^2.$$

This is linear in z_1, w_1. We choose a random solution such that $|\lambda z_0 + N z_1|$ and $|\lambda w_0 + N w_1|$ are less than N^2. Finally, it remains to find x_1, y_1 such that

$$f(x_1, y_1) = \frac{\ell^e - pf(\lambda z_0 + N z_1, \lambda w_0 + N w_1)}{N^2}.$$

Thus, we are looking for $x_1 + y_1\omega \in R$ with a particular norm. This can be done efficiently using standard methods such as Cornaccia's algorithm.

Algorithm 3 Strong Approximation

Require: $[\mu] \in (R/NR)^\times[j]$.
Ensure: A lift $\mu \in \mathcal{O}$ of $[\mu]$ such that $\mathrm{nrd}(\mu) = \ell^e$ for some e.
1: $\mu_0 \leftarrow$ an arbitrary lift of μ in Rj of the form $(z_0 + w_0\omega)j$.
2: $\lambda \leftarrow$ solution to $p\lambda^2 f(z_0, w_0) \equiv \ell^e \mod N$ (choose e such that λ exists).
3: $z_1, w_1 \leftarrow$ solution to $\ell^e - pf(\lambda z_0 + N z_1, \lambda w_0 + N w_1) \equiv 0 \mod N^2$.
4: $x_1, y_1 \leftarrow$ solution to $f(x_1, y_1) = \frac{\ell^e - pf(\lambda z_0 + N z_1, \lambda w_0 + N w_1)}{N^2}$.
5: **return** $\mu = \lambda\mu_0 + N(x_1 + y_1\omega + (w_1 + z_1\omega)j)$

1.4 Reductions

Earlier we assumed that $\mathcal{O}$ was nice. Here we will show how to solve the problem for a general order.

Suppose we have two maximal orders $\mathcal{O}_1$ and $\mathcal{O}_2$. Assume that we have an algorithm $\mathcal{A}$ which takes a left $\mathcal{O}_1$-ideal I and returns β such that $I\beta$ has ℓ-power reduced norm.

Algorithm 4 General orders

Require: Maximal orders $\mathcal{O}_1$ and $\mathcal{O}_2$, a left $\mathcal{O}_2$-ideal J, and the algorithm $\mathcal{A}$.
Ensure: β such that $J\beta$ has ℓ-power reduced norm.
1: $I \leftarrow \{\alpha \in B \colon \alpha\mathcal{O}_2\overline{\alpha} \subseteq [\mathcal{O}_1 \colon \mathcal{O}_1 \cap \mathcal{O}_2]\mathcal{O}_1\}$, this is a left $\mathcal{O}_1$-ideal and right $\mathcal{O}_2$-ideal.
2: $\beta_1 \leftarrow$ output of $\mathcal{A}$ applied to I as a left $\mathcal{O}_1$-ideal.
3: $\beta_2 \leftarrow$ output of $\mathcal{A}$ applied to IJ as a left $\mathcal{O}_1$-ideal.
4: **return** $\beta = \beta_2\overline{\beta}_1\mathrm{nrd}(I)$.

Next we will verify this algorithm produces the correct output. It is clear that $J\beta$ has ℓ-power norm, so we only check that $J\beta \subseteq \mathcal{O}_2$. It's enough to show that $\overline{\beta}\mathrm{nrd}(J) \in J$ since then $J\beta\mathrm{nrd}(J) \subseteq J\overline{J} = \mathrm{nrd}(J)\mathcal{O}_2$. So to show this, note that

$$\overline{\beta}\mathrm{nrd}(J) = \beta_1\overline{\beta}_2\mathrm{nrd}(I)\mathrm{nrd}(J) = \beta_1\mathrm{nrd}(I)\overline{\beta}_2\mathrm{nrd}(IJ)\frac{1}{\mathrm{nrd}(I)} \in (\overline{I})(IJ)\frac{1}{\mathrm{nrd}(I)} = J$$

Here we are using that if I is a left $\mathcal{O}$-ideal then $I\overline{I} = \mathrm{nrd}(I)\mathcal{O}$, and if I' is a right $\mathcal{O}$-ideal then $\overline{I}'I' = \mathrm{nrd}(I')\mathcal{O}$

Algorithm 5 Prime reduced norm

Require: A left $\mathcal{O}$-ideal I
Ensure: β such that $I\beta$ has prime reduced norm.

1: **repeat**
2: $\quad\beta' \leftarrow$ random element of I.
3: **until** $\mathrm{nrd}(\beta')/\mathrm{nrd}(I)$ is prime.
4: **return** $\beta = \overline{\beta}'/\mathrm{nrd}(I)$.

Another reduction we used was we assumed that I had prime reduced norm. We can do this as follows. It is clear that $I\beta$ will have the appropiate reduced norm. To see that $I\beta$ is an $\mathcal{O}$-ideal, note that

$$I\beta = I\frac{\overline{\beta}'}{\mathrm{nrd}(I)} \subseteq I\overline{I}\frac{1}{\mathrm{nrd}(I)} = \mathcal{O}.$$

2 ID Protocols and Signature Schemes

In this section we summarize the main results in [GPS17].

2.1 Signature Scheme

In this section, we describe the signature scheme in [GPS17, Sec. 4]. Suppose we have two parties, Sam and Victor. The goal of this protocol is to allow Sam to sign a document and Victor to validate the signature.

Before giving the protocol, we first motivate it by a 0-knowledge proof of an isogeny. Suppose Sam knows an isogeny $\phi : E_0 \to E_1$. In order to prove she knows ϕ, Sam will choose a random isogeny $E_0 \to E_2$ and publish E_0, E_1, E_2 (but none of the isogenies). Victor can challenge Sam by giving her a bit $b \in \{0, 1\}$. Sam responds by giving an isogeny $E_b \to E_2$. It would be difficult to for Sam to generate isogenies $E_0 \to E_2$ and $E_1 \to E_2$ without already knowing an isogeny $E_0 \to E_1$.

To make sure that Sam does not reveal any information about ϕ, she needs to choose a "fresh" isogeny $E_b \to E_2$, i.e. one that does not factor through E_1. This is done by a clever application of the algorithm in [KLPT14].

The protocol is as follows:

1. (Setup) Sam chooses a random isogeny $\phi : E_0 \to E_1$ of supersingular curves. She publishes E_0 and E_1 but keeps ϕ hidden.

2. (Sign) To sign an t bit message $m = m_1 \cdots m_t$, Sam does the following. For each m_i, Sam computes a random walk $\phi_i : E_0 \to E_{2,i}$. Then sets h to be the hash of m and $j(E_1), \ldots, j(E_{2,i})$. Let $h_1 \cdots h_s$ be the bits of h. For each i, Sam sets z_i to a "fresh" path from $E_{h_i} \to E_{2,i}$. The signature is $(h, z_1, \ldots, z_s)$.

3. (Verify) To verify a signature, Victor computes the j-invariant of the end of each path z_i, which is $j(E_{2,i})$. Then Victor computes the hash of the message with these j-invariants and checks it matches h.

It is important that the paths z_i are "random" compositions of low degree isogenies. In particular, paths from E_0 should not go through E_1 and similarly for paths from E_1 should avoid E_0.

Suppose we have an isogeny or path from E to E'. A "fresh" path can be constructed as follows.

1. Convert the path $E \to E'$ into an left ideal I of $\mathcal{O} = \mathrm{End}E$.
2. Apply the algorithm of [KLPT14] to find a left equivalent ideal $J = I\beta$ for some $\beta \in \mathrm{End}E \otimes \mathbb{Q}$ such that the reduced norm of J is a power of a small prime.
3. Convert the ideal J into a path $E \to E'$.

We will describe these steps in more detail below.

The path generated by the algorithm in [KLPT14] only depends on the left equivalence class of the input ideal I. Also, the distribution of outputs is actually quite small. Therefore the output should be considered "independent" of I and a "pseudo-canonical" path from $E \to E'$.

2.2 Security

The security of this scheme relies on the following problem:

Problem 3. Given two random supersingular elliptic curves E_1, E_2 over $\mathbb{F}_{p^2}$, construct an isogeny $E_1 \to E_2$.

Theorem 4. *If this problem is hard, then the signature scheme is secure under chosen message attacks in the random oracle model.*

2.3 Isogeny Path to Ideal

Note that we may evaluate $\alpha(Q)$ because $\alpha \in \mathcal{O}$ so we may assume it's given as a linear combination of some known basis for $\mathcal{O}$ for which we have corresponding rational maps in $\mathrm{End}E$.

2.4 Ideal to Isogeny Path

Let E be a supersingular elliptic curve, and let suppose we have an isomorphism from $\mathrm{End}E$ to a maximal order $\mathcal{O}$ in a quaternion algebra. Given an ideal I of ℓ-power reduced norm, we want to compute an isogeny path corresponding to I. The main difficulty is that the isogeny corresponding to I has large degree, so we need to factor it as a sequence of degree ℓ-isogenies.

The point Q can be found by picking a basis $\{P_1, P_2\}$ for $E[\ell^e]$ and computing the matrices A_j for the α_j. Then Q corresponds to a vector in the kernel

Algorithm 6 Isogeny Path to Ideal

Require: A supersingular elliptic curve E, an isomorphism from End$E$ to a maximal order $\mathcal{O}$, and a sequence of isogenies $\varphi_i : E_{i-1} \to E_i$ for $i = 1, \ldots, r$ with $E_0 = E$ and $\deg \varphi_i = \ell_i^{e_i}$.

Ensure: A sequence of ideals $I_0, \ldots, I_r$ corresponding to the path.

$I_0 \leftarrow \mathcal{O}_0 = \mathcal{O}$.
$n \leftarrow \prod \ell_i^{e_i}$.
for $i = 1, \ldots, r$ **do**
  $Q \leftarrow$ generator of $\ker \varphi_i$ in $E[\ell_i^{e_i}]$.
  $\beta_1, \ldots, \beta_4 \leftarrow \mathbb{Z}$-basis for I_{i-1}.
  $f_i(w, x, y, z) \leftarrow n(w\beta_1 + x\beta_2 + y\beta_3 + z\beta_4)$.
  repeat
    $x_1, \ldots, x_4 \leftarrow$ random choice such that $n \sum x_i\beta_i \equiv 0 \mod \ell_i^{e_i}$.
    $\alpha \leftarrow \sum x_i\beta_i$
  until $\alpha(Q) = 0$
  $I_i \leftarrow I_{i-1}\ell_i^{e_i} + \mathcal{O}_0\alpha$
end for
return $I_1, \ldots, I_r$.

Algorithm 7 Ideal to Isogeny Path

Require: A supersingular elliptic curve E, an isomorphism from End$E$ to a maximal order $\mathcal{O}$, and a left $\mathcal{O}$-ideal I.

Ensure: An isogeny path $E \to E/E[I]$ consisting of low degree isogenies.

Apply KLPT to replace I with an ideal of powersmooth norm, $\mathrm{nrd}(I) = \prod_{i=1}^{r} \ell_i^{e_i}$.
$\alpha_1, \ldots, \alpha_4 \leftarrow \mathbb{Z}$-basis for I as endomorphisms in End$E$.
$E_0, \varphi_0 \leftarrow E, id_E$.
for $i = 1, \ldots, r$ **do**
  $Q \leftarrow$ any element in $E[\ell_i^{e_i}] \bigcap \ker \alpha_j$ (see below).
  $\varphi_i : E_{i-1} \to E_i \leftarrow$ isogeny on E_{i-1} with kernel $\varphi_{i-1} \circ \cdots \circ \varphi_1(Q)$.
end for
return $\varphi_1, \ldots, \varphi_r$.

of the A_j. Note that computing A_j requires solving a DLP (since we need $\alpha_j(P_i) = a_i P_1 + b_i P_2$). This is possible because ℓ^e is small.

References

[EHL+18] Kirsten Eisenträger, Sean Hallgren, Kristin Lauter, Travis Morrison, and Christophe Petit. Supersingular isogeny graphs and endomorphism rings: reductions and solutions. In *Advances in cryptology—EUROCRYPT 2018. Part III*, volume 10822 of *Lecture Notes in Comput. Sci.*, pages 329–368. Springer, Cham, 2018.

[GPS17] Steven D. Galbraith, Christophe Petit, and Javier Silva. Identification protocols and signature schemes based on supersingular isogeny problems. In *Advances in cryptology—ASIACRYPT 2017. Part I*, volume 10624 of *Lecture Notes in Comput. Sci.*, pages 3–33. Springer, Cham, 2017.

[KLPT14] David Kohel, Kristin Lauter, Christophe Petit, and Jean-Pierre Tignol. On the quaternion ℓ-isogeny path problem. *LMS J. Comput. Math.*, 17(suppl. A):418–432, 2014.

[KV12] Markus Kirschmer and John Voight. Corrigendum: Algorithmic enumeration of ideal classes for quaternion orders [mr2592031]. *SIAM J. Comput.*, 41(3):714, 2012.