Computing cyclic isogenies using real multiplication

Notes of a talk given for the ANR Peace project

Damien Robert

2013–04–19; Updated 2013–04–23

1 Introduction

This notes are an expanded version of a talk [Rob13] I gave the 11 April 2013 for the PEACE meeting in Paris. Since several people who could not attend have asked for informations about this talk I give here a public version. **A word of warning:** these are preliminary notes so they are bound to have mistakes. More importantly I lack a concrete implementation yet.

1.1 Computing isogenies with maximal isotropic kernel

In [CR11] we gave an algorithm to compute isogenies between abelian varieties. More precisely, let A/k be an abelian variety A/k of dimension g represented by its theta null point of level n (in particular A is polarized). Then given a basis $e_1,\dots,e_g$ of a rational kernel $K\subset A[\ell]$ maximally isotropic for the ℓ-Weil pairing (with ℓ prime to $2n$), we explain how to compute $B=A/K$ (via its theta null point of level n) and how to compute the image of a point $x\in A(\overline{k})$ via the isogeny $f:A\to B$. This can be seen as a generalisation of the well known Vélu's formulas [Vél71] to compute isogenies between abelian varieties.

This algorithm needs a polynomial (in the size of the kernel K) number of operations in the field where the geometric points of K live. Actually, the article [CR11] focus on the case of dimension $g=2$, because in this case every (generic) abelian variety is a Jacobian of an hyperelliptic curve, and we explain how to use Thomae's formulas to convert between the Mumford representation and the theta representation (see also [Wam99]).

More details on this algorithm are also given in [Cos11] (using analytic theta functions), and in [Rob10] (using algebraic theta functions). The algorithm given in [CR11] builds on result from [FLR11; LR12b] by applying a result from [Koi76] (in the analytic setting) and [Kem89] (in the algebraic setting).

The above algorithm was implemented in [BCR10a] to compute isogenies between abelian varieties of dimension 2 over finite fields. Note that when one use the theory of analytic theta functions, to extend the results to varieties over a finite field, one need to assume that they are ordinary so that a lift to characteristic zero can be taken. The advantage of algebraic theta functions is that the resulting theory will work over any algebraically closed field of characteristic prime to the level n. Since $n=2$ or $n=4$ this handle all fields of odd characteristics. For an ordinary abelian variety over $\mathbb{F}_{2^n}$, one can lift to characteristic zero, but the formulas from the isogeny algorithm have bad reduction in this case, so we need to make a change of variable. The resulting algorithm to compute isogenies in characteristic two is described in [BCR10b] (for the dimension 2 case).

The condition ℓ prime to $2n$ is purely technical, we explain in [Rob10] how to compute an isogeny when this is not the case (in this case we need more than juste the geometric points of the kernel, we will see why in Section 2).

Finally, an improvement of this algorithm so that only operations over the field of definition of the kernel K are needed (provided we have the equations of K) is given in [Rob12] (in collaboration with David Lubicz).

1.2 The case of cyclic isogenies

At the end of [CR11], we concluded that it would be worthwhile to investigate the case of isogenies with cyclic kernel; they are needed to have a full description of the isogeny graph (otherwise we don't even have a connected

subgraph), which has many applications: [LR12a]... The problem here is that the pullback of a line bundle by a cyclic kernel is not as easy to describe algebraically as when the kernel is maximally isotropic.

It is easier to explain why if we use the theory of complex multiplication [Shi98]. Let K be a (primitive) CM field of degree $2g$ (a totally imaginary quadratic extension of a totally real field K_0). Then the moduli space of abelian varieties with complex multiplication by O_K is a torsor under the Shimura class group

$$\mathfrak{C} = \{(I,\rho) \mid I \text{ a fractional } O_K\text{-ideal with } I\bar{I} = (\rho),\ \rho \in K^+ \text{ totally positive}\}/K^*$$

(In particular, it is of dimension 0.)

If $(A, \mathscr{L}_0)$ is a principally polarized abelian variety of dimension g with CM by O_K, the element $(I,\rho) \in \mathfrak{C}$ acts on A in the following way: I give the kernel K of the corresponding isogeny on A, and ρ explain the action on the polarization. I corresponds to a maximal isotropic kernel (for the Weil pairing on $\mathscr{L}_0^\ell$) iff I is of relative norm ℓ. In this case the element (I,ℓ) give an isogeny between the polarized abelian variety $(A, \mathscr{L}_0^\ell)$ and $(B, \mathscr{M}_0)$ (where $\mathscr{M}_0$ is a principal polarization), so the action of ℓ on the polarization is easy to describe. For a general element (I,ρ), one would need to understand what the polarization "$\mathscr{L}_0^\rho$" such that we have an isogeny $(A, \mathscr{L}_0^\rho)$ and $(B, \mathscr{M}_0)$ of polarized abelian varieties would mean. Note that $\mathscr{L}_0^\rho$ is not isomorphic to $\rho^*\mathscr{L}_0$ (Think about the case $\rho = \ell$ and $\mathscr{L}_0$ symmetric where $\rho^*\mathscr{L}_0 = \mathscr{L}_0^{\ell^2} \neq \mathscr{L}_0^\ell$).

When $\rho = \ell$, one could compute an isogeny (with maximal isotropic kernel for $\mathscr{L}_0^\ell$) the following way: find a matrix $F \in \mathrm{Mat}_r(\mathbb{Z})$ such that ${}^tFF = \ell\,\mathrm{Id}$. Then the Koizumi-Kempf formula applied to F give a link between the theta functions of level ℓn on $\mathscr{L}^\ell$ (where $\mathscr{L} = \mathscr{L}_0^n$) and the theta functions of level n on $\mathscr{L}$, we will call this "changing the level" or the "level formulas". (Basically we just have to apply the isogeny theorem on the isogeny $F : A^r \to A^r$ given composant by composant by the matrix F. Here A^r is given the product polarization $\mathscr{L} \star \ldots \star \mathscr{L}$, so the isogeny theorem give relations between products of r theta functions on A.) Then once we are in $(A, \mathscr{L}^\ell)$ we can just apply the isogeny theorem to get into $(B, \mathscr{M})$ ($\mathscr{M} = \mathscr{M}_0^n$). In [CR11] we do things the other way around because we get a more efficient algorithm this way, we will explain why latter.

In the case of complex multiplication, one could try to adopt a similar strategy for a cyclic isogeny coming from the action of (I,ρ): find a matrix $F \in \mathrm{Mat}_r(O_K)$ such that ${}^tFF = \rho\,\mathrm{Id}$ and apply a Koizumi like formula to get from $(A, \mathscr{L})$ to $(A, \mathscr{L}^\rho)$. We have two problem here: the Koizumi formula comes from the isogeny formula on A^r, but when F is not an integral matrix, there is no reason that F respect the underlying symplectic decomposition, so we may not apply the isogeny theorem. The second problem, is that even if it does, to compute the corresponding change of level, we need a way to compute the action of elements of O_K on affine lifts uniformly. For an action of $\gamma \in \mathbb{Z}$ we know how to do it using differential additions, but it is not clear how to do that for a more general γ. If γ itself correspond to an isogeny with maximal isotropic kernel, then one solution is to use [CR11], because the isogeny algorithm given here actually work with affine coordinates (this is clear given the way we keep track of the projective factors), so it would be doable but would need branching isogeny computations inside the level formula of our current cyclic isogeny computation. All in all this seemed like a cumbersome computation, and it only guides us in the case of fixed CM, whereas I was interested in moving vertically in the isogeny graph using cyclic isogenies.

In November 2011, Dimitar Jetchev contacted me about the possibility of computing cyclic isogenies in dimension 2, and this is basically the response I gave: that in the restricted case of known CM and horizontal isogeny, it should theoretically be feasible but rather cumbersome.

1.3 Real multiplication to the rescue!

In July 2012, while I was visiting Microsoft Research, I discussed with Sorina Ionica who showed me wonderful graphs of cyclic isogenies between abelian varieties having the same real multiplication (RM) in dimension 2. These graphs were obtained in collaboration with Emmanuel Thomé, following an idea from John Boxall to use real multiplication to compute isogenies.

While Sorina and Thomé obtained their graphs by working over $\mathbb{C}$ (and with lattices coming from the Hilbert space $\mathfrak{H}_1^g$ via the real multiplication O_{K_0}), this discussion made clear that the case of computing the action of ρ on the polarization (in order to compute a cyclic isogeny) was much better than I thought.

Indeed, it is clear from the definition of the Shimura class group that ρ is a totally positive element in K_0. It is well known that every such element is a sum of squares, and it is also well known how from such a sum of

squares one can use Clifford's algebra to compute a matrix F such that $t_F F = \rho\,\mathrm{Id}$. The important part here is that $F \in \mathrm{Mat}_r(O_{K_0})$ is composed of totally real elements. This has two important consequences, first since complex conjugation on an ideal $I \subset K$ correspond to the dual isogeny, an element $\gamma \in K_0$ commutes with the Rosatti involution. In particular, the action of the elements of K_0 on $\mathbb{C}^g$ is given by symmetric matrices for the hermitian form H associated to the principal polarization on $(A, \mathscr{L}_0)$. In particular they are all codiagonalisable, so it is immediate that the matrix F respect the symplectic decomposition and we can apply the isogeny theorem to obtain Koizumi-like formulas. Secondly, computing the action of an element in K_0 on affine points of the abelian variety is much easier than for a general element in K as we will see.

Independently, Alina Dudeanu and Dimitar Jetchev have also been working on obtaining a Koizumi-like formula in the analytic setting using real multiplication.

This notes are heavily indebted to helpful discussions with John Boxall and Sorina Ionica; and even more importantly to my on-going collaboration with David Lubicz in the use of algebraic theta functions for cryptographic applications. We will assume known the standard results of analytic theta functions [Igu72; Mum83; Mum84; Mum91; BL04] and algebraic theta functions [Mum66; Mum67a; Mum67b; Mum70]. We use the standard acronyms ppav for principally polarized abelian variety and pav for polarized abelian variety. We will also always assume that the line bundles are symmetric.

2 Symmetric theta structures and the isogeny theorem

Let A be an abelian variety of dimension g defined over an algebraically closed field $\overline{k}$. Let $\mathscr{L}_0$ be a symmetric ample line bundle of degree one on A, $\mathscr{L}_0$ defines a principal polarization: $A \to \hat{A}$. If n is even $\mathscr{L} = \mathscr{L}_0^n$ is then totally symmetric, and the kernel $K(\mathscr{L})$ of the polarization associated to $\mathscr{L}$ is $A[n]$.

From now on, we assume that n is prime to the characteristic of k, so that $\mathscr{L}$ defines a separable polarisation. Since $\mathscr{L}$ is totally symmetric, there exist a symmetric theta structure on the theta group $G(\mathscr{L})$. Fixing such a structure fix a unique projective basis of theta functions [Mum66] that we call theta functions of level n. Note: the theta structure induces an isomorphism between the symplectic spaces $Z(\overline{n}) \times \hat{Z}(\overline{n})$ and $K(\mathscr{L}) = A[n]$ where $Z(\overline{n}) = (\mathbb{Z}/n\mathbb{Z})^g$ and $\hat{Z}(\overline{n})$ is the Cartier dual of $Z(\overline{n})$. We note $K(\mathscr{L}) = K_1(\mathscr{L}) \oplus K_2(\mathscr{L})$ where $K_1(\mathscr{L})$ corresponds to $Z(\overline{n})$ and $K_2(\mathscr{L})$ to $\hat{Z}(\overline{n})$. Usually the canonical basis of the theta functions of level n are indexed by $i \in Z(\overline{n})$, but in these notes we will index them by $i \in K_1(\mathscr{L})$ which permit us to not track explicitly the isomorphism between $Z(\overline{n})$ and $K_1(\mathscr{L})$.

If $n > 2$ then the theta functions of level n give a projective embedding of A into $\mathbb{P}^{n^g-1}_{\overline{k}}$, while if $n = 2$ we only get an embedding of the Kummer variety $A/\pm 1$ (the $n = 2$ case assume that A is absolutely simple, see [BL04]). Under a generic condition (the even theta null coordinates are non zero), this embedding of the Kummer variety is actually projectively normal (see [Koi76]).

Theorem 2.1 :
The symmetric theta structure on $G(\mathscr{L})$ is uniquely determined by a choice of symplectic basis $(e_1, \dots e_g, e'_1, \dots e'_g)$ on $A[n]$ and a choice of symplectic basis $(f_1, \dots f_g, f'_1, \dots f'_g)$ on $A[2n]$ such that $e_i = 2f_i, e'_i = 2f'_i$. (Here symplectic mean for the commutator pairing $e_{\mathscr{L}}$ and $e_{\mathscr{L}^2}$ respectively).

Moreover, changing these symplectic basis do not change the resulting symmetric theta structure if and only if

- *The symplectic basis of $A[n]$ is left invariant;*
- *The f_i are replaced by points $f_i + t_i$ with $t_i \in A[2]$ such that $e_{\mathscr{L}}(e_i, t_i) = 1$.*

In particular, fixing a symplectic basis of $A[n]$ and a symplectic decomposition $A[2n] = A_1[2n] \oplus A_2[2n]$ of the $2n$-torsion into a sum of maximal isotropic subspaces is enough (and even stronger) to fix the symmetric theta structure.

Proof : This is implicit in [Mum66, Section 3]. A symmetric theta structure comes from an isomorphism between the Heisenberg group and the theta group that commutes with the action of $[-1]$. It induces an isomorphism between the symplectic spaces $Z(\overline{n}) \times \hat{Z}(\overline{n})$ and $K(\mathscr{L}) = A[n]$ and hence fix a symplectic basis of the n-torsion.

Conversely, having fixed a symplectic basis of the n-torsion, since $\mathcal{L}$ is totally symmetric, there is always a symmetric theta structure respecting this symplectic basis. Such a choice of a symmetric theta structure can be seen as a choice of a symmetric element above each of the element of the basis $(e_1,\dots e'_g)$; since there is only two symmetric elements $\pm g_i$ above each e_i a symmetric theta structure above the symplectic basis can be seen as a choice of sign for each element of the basis.

If $g_i \in G(\mathcal{L}^2)$ is a symmetric element of the theta group above a point f_i such that $e_i = 2f_i$, then $(g_i)^2$ determines a symmetric element of the theta group above e_i that uniquely depends on the choice of f_i (since the other symmetric element above f_i is $-g_i$ which gives rise to $(-g_i)^2 = (g_i)^2$ above e_i. Via the transfer map δ_2 from [Mum66], we see how the choices of the f_i above the e_i are enough to determine the symmetric theta structure on $G(\mathcal{L})$.

It is a straightforward verification to see that replacing f_i by $f_i + t_i$ where t_i is a point of 2-torsion involve replacing $(g_i)^2$ by $e_{\mathcal{L}^2}(f_i,t_i)(g_i)^2$ which concludes the proof.

(One could also replace the application δ_2 by the isogeny $[2]$ which would involve working in $G(\mathcal{L}^4)$, as in [Kem89].) ■

Of course Theorem 2.1 also work for any totally symmetric line bundle $\mathcal{L}$ on A, defining a polarization of type $\delta = (\delta_1,\dots,\delta_g)$. The idea is that if $\mathcal{L} = \mathcal{L}_0^n$ (say with $n=2$ or $n=4$), $\mathcal{L}^\ell$ is of type $(\ell n, \ell n)$ and allows to compute isogenies with maximal isotropic kernels, but for a cyclic isogeny we need a polarisation of type $(n,\ell n)$ (like the type of $\mathcal{L}^\rho$ from Section 1.3).

Theorem 2.2 :
Let $f : (A,\mathcal{L}) \to (B,\mathcal{M})$ be an isogeny between pav, with $\mathcal{L}$ totally symmetric. Then $K = \operatorname{Ker} f$ is isotropic in $K(\mathcal{L})$ for the commutator pairing $e_{\mathcal{L}}$, and $K(\mathcal{M}) \simeq K^\perp/K$.

Assume that we have a symmetric theta structure on $G(\mathcal{L})$ coming from a symplectic basis (f_i,f'_i) on $K(\mathcal{L}^2)$. Assume that K is compatible with the induced symplectic decomposition $K(\mathcal{L}) = K_1(\mathcal{L}) \oplus K_2(\mathcal{L})$ into maximal isotropic subspaces in the sense that $K = K_1 \oplus K_2$ where $K_i = K_i(\mathcal{L}) \bigcap K$. In this case $K(\mathcal{M}) \simeq K^{2,\perp}/K_1 \oplus K^{1,\perp}/K_2$ where $K^{2,\perp} = K_2^\perp \bigcap K_1(\mathcal{L})$ and $K^{1,\perp} = K_1^\perp \bigcap K_2(\mathcal{L})$

Let $\widetilde{K}$ be the level subgroup above K induced by this theta structure; the corresponding descent data give a line bundle $\mathcal{M}'$ algebraically equivalent to $\mathcal{M}$. For simplicity we assume here that $K \subset 2K(\mathcal{L})$ (or equivalently that $A[2] \subset K^\perp$), so that $\mathcal{M}'$ is the unique totally symmetric line bundle equivalent to $\mathcal{M}$. (The isogeny theorem is valid in a more general setting, but we will only need this case in the following).

We can define a symmetric theta structure on $\mathcal{M}'$ as follow: from the symplectic basis of $K(\mathcal{L}^2)$ one derives a "canonical" basis $(g_1,\dots,g'_g)$ of $[2]^{-1}K^\perp$. Pushing this basis via the isogeny f gives a symplectic basis on $K(\mathcal{M}'^2)$, which determines the symmetric theta structure on $\mathcal{M}'$. It is easy to see that by construction, it is compatible with the theta structure on $\mathcal{L}$.

We can then apply the isogeny theorem: there exist λ such that for all $i \in K_1(\mathcal{M}')$

$$\vartheta_i^{\mathcal{M}'} = \lambda \sum_{j \in K_1(\mathcal{L}) | f(j) = i} \vartheta_j^{\mathcal{L}}.$$

Proof : This is [Mum66, Section 1]. The version stated here is from [Kem89]. See also [Rob10, Chapter 3–4] for a summary. ■

3 Computing isogenies with maximal isotropic kernel

In this section we review the algorithm of [CR11]. This is because we will see that the tools used to compute cyclic isogenies are extremely similar, and also because we will need to be able to compute maximally isotropic isogenies in order to compute cyclic isogenies.

Let $(A,\mathcal{L}_0)$ be a ppav, and K a maximal isotropic kernel for $\mathcal{L}_0^\ell$. Let n be even and $\mathcal{L} = \mathcal{L}_0^n$. Assume that we have a principal polarization $\mathcal{M}_0$ on $B = A/K$, and let $\mathcal{M} = \mathcal{M}_0^n$. For simplicity we assume here that ℓ is prime to $2n$. We note $\Phi_{\mathcal{L}} : A \to \hat{A},\ x \mapsto t_x^*\mathcal{L} \otimes \mathcal{L}^{-1}$ the polarization associated to $\mathcal{L}$.

To have an algorithm for the isogeny $f : A \to B$ mean that we want to find relations between theta functions of level n on A (for $\mathscr{L}$) and theta functions of level n on B (for $\mathscr{M}$).

First we need to have some sort of compatibility between $\mathscr{L}$ and $\mathscr{M}$. More exactly, we want the following diagram to commute:

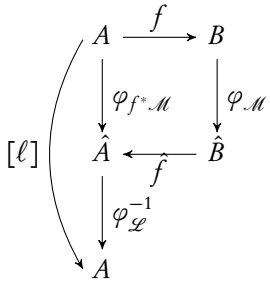

It is easy to see that since we have the following diagram:

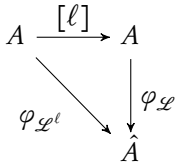

this is the case iff $\mathscr{L}^\ell = f^*\mathscr{M}$.

Now we have two tools. The Koizumi formula explain the relations between the theta functions of level ℓn for $\mathscr{L}^\ell$ and the theta functions of level n for $\mathscr{L}$.

Concretely, assume given a symmetric theta structure on $\mathscr{L}^\ell$, by Theorem 2.1 this induces a symmetric theta structure on $\mathscr{L}$. Let $F \in \mathrm{Mat}_r(\mathbb{Z})$ be a matrix such that ${}^tFF = \ell\,\mathrm{Id}$, and note also F the isogeny $A^r \to A^r$ induced by F. (In practice $r = 2$ when ℓ is a sum of two squares, and $r = 4$ otherwise). The theta structures on $\mathscr{L}$ and $\mathscr{L}^\ell$ induce product theta structures on $\mathscr{L} \star \ldots \mathscr{L}$ and $\mathscr{L}^\ell \star \ldots \mathscr{L}^\ell$. In this setting, Theorem 2.2 gives us

Proposition 3.1 :
Let $(i_1, \ldots, i_r) \in K_1(\mathscr{L})^r$. Let $x = (x_1, \ldots, x_r)$ be a geometric point of A^r and let $y = Fx$. Then (up to a constant λ)

$$\vartheta^{\mathscr{L}}_{i_1}(y_1) \cdot \ldots \cdot \vartheta^{\mathscr{L}}_{i_r}(y_r) = \lambda \sum_{\substack{(j_1,\ldots,j_r) \in K_1(\mathscr{L}^\ell) \\ F(j_1,\ldots,j_r) = (i_1,\ldots,i_r)}} \vartheta^{\mathscr{L}^\ell}_{j_1}(x_1) \cdot \ldots \cdot \vartheta^{\mathscr{L}^\ell}_{j_r}(x_r).$$

Proof : From the theorem of the square we have that $F^*(\mathscr{L} \star \mathscr{L} \star \ldots) = \mathscr{L}^\ell \star \mathscr{L}^\ell \star \ldots$. The rest is immediate from Theorem 2.2. ■

The isogeny formula explain the relations between the theta functions for $\mathscr{L}^\ell$ on A and the theta functions for $\mathscr{M}$ on B.

Proposition 3.2 :
Assume that the symmetric theta structure on $\mathscr{L}^\ell$ is such that $K \subset K_2(\mathscr{L}^\ell)$ (this is always possible). Then the symmetric theta structure on $\mathscr{L}^\ell$ induces a symmetric theta structure on $\mathscr{M}$ by Theorem 2.2 (this may require to replace $\mathscr{M}$ by an equivalent line bundle).

Let $f : A \to B$ be the isogeny of kernel K, and x a geometric point in A. Fix $i \in K_1(\mathscr{M})$, and let $j \in K_1(\mathscr{L}^\ell)$ be the unique preimage of i by f that is in $K_1(\mathscr{L}^\ell)$. We have (up to a constant λ)

$$\vartheta^{\mathscr{M}}_i(f(x)) = \lambda \vartheta^{\mathscr{L}^\ell}_j(x)$$

Proof : Immediate by Theorem 2.2. ■

Example 3.3 :
Let $A = \mathbb{C}^g/(\mathbb{Z}^g + \Omega\mathbb{Z}^g)$ be a ppav over $\mathbb{C}$ with $\Omega \in \mathfrak{H}_g$. The polarization associated to Ω correspond to an hermitian form H_0 on $\mathbb{C}^g$. More generally, a polarization comes from an hermitian form H on $\mathbb{C}^g$ such that $H(ix, iy) = H(x, y)$ and $H(\Lambda, \Lambda) \subset \mathbb{Z}$ where $\Lambda = \mathbb{Z}^g + \Omega\mathbb{Z}^g$.

An isogeny correspond to a matrix M acting on $\mathbb{C}^g$, and the dual isogeny correspond to $H_0(M\cdot, \cdot)$ acting on $\hat{A} \simeq \mathrm{Hom}_{\overline{\mathbb{C}}}(\mathbb{C}^g, \mathbb{C})$. Pulling back the dual isogeny via the principal polarization, we get that it acts on $\mathbb{C}^g$ by $M^* = {}^t\overline{M}$. (We see that we recover the action by F on $\mathscr{L} \star \mathscr{L} \star \dots$ from Proposition 3.1).

A basis of level n theta function corresponding to $H = nH_0$ (and the characteristic $c = 0$ in the sense of [BL04]) is given by $(\vartheta \left[\begin{smallmatrix} 0 \\ b \end{smallmatrix}\right](\cdot, \Omega/n)_{b \in Z(\overline{n})}$ where

$$\vartheta \left[\begin{smallmatrix} a \\ b \end{smallmatrix}\right](z, \Omega) = \sum_{n \in \mathbb{Z}^g} e^{\pi i\, {}^t(n+a)\Omega(n+a) + 2\pi i\, {}^t(n+a)(z+b)}.$$

Up to an action of the symplectic group $\mathrm{Sp}_{2g}(\mathbb{Z})$ we can assume that the kernel K corresponds to $\frac{1}{\ell}\Omega\mathbb{Z}^g/\Omega\mathbb{Z}^g$ so that the isogenous abelian variety is $B = \mathbb{C}^g/(\mathbb{Z}^g + \frac{\Omega}{\ell}\mathbb{Z}^g)$.

Comparing the basis of theta functions of level n on B

$$(\vartheta \left[\begin{smallmatrix} 0 \\ b \end{smallmatrix}\right](\cdot, \frac{\Omega}{\ell}/n))_{b \in Z(\overline{n})}$$

and the basis of theta functions of level $n\ell$ on A

$$(\vartheta \left[\begin{smallmatrix} 0 \\ b \end{smallmatrix}\right](\cdot, \Omega/\ell n))_{b \in Z(\overline{\ell n})}$$

immediately give Proposition 3.2.

Now the natural thing to compute the isogeny $A \to B$ would be to combine Propositions 3.1 and 3.2: inverse the formulas from Proposition 3.1 to go from theta coordinates on $\mathscr{L}$ to theta coordinates on $\mathscr{L}^\ell$ and then apply Proposition 3.1 on $\mathscr{L}^\ell$.

Inversing the level formula could be done as follow: first try to find a theta null point of level ℓn associated to a symmetric theta structure on $G(\mathscr{L}^\ell)$ compatible with the one on $G(\mathscr{L})$. Since we know how the moduli space of theta null point of a certain level look like (by [Mum67a] it is given by Riemann's relations + the symmetries) this can be done by a Gröbner basis algorithm. Since the fiber is finite, we are in a favorable case for Gröbner computations. Then once we have fixed a theta null point of level ℓn in the fiber, we can lift a geometric point x on A given by level n theta coordinates to level ℓn coordinates. This can also be done by a Gröbner basis algorithm since the projective equations of A embedded by theta functions is described in [Mum66] (and only need the coordinates of the theta null point).

In fact, in [FLR11] we inverse the isogeny formula from Proposition 3.2 instead. This is because it is simpler, so it allows to speed-up the Gröbner basis computation related by using the extra information we have about the system (in [FLR11] we only care about lifting the theta null point since we were only interested in describing some modular correspondances). In other words, rather than looking at the isogeny $f : (A, \mathscr{L}^\ell) \to (B, \mathscr{M})$, we look at the contragredient isogeny $\widetilde{f} : B \to A$. (This whole part and what follows is because Ben Smith complained during the talk that we think with "arrows reversed", this is to try to justify why it is a good idea in our situation!)

Still we would like to have an algorithm that does not need Gröbner basis. We note here that both Proposition 3.1 and Proposition 3.2 loose information (they go to a lower level), but only in a finite way (the associated fibers are finites). Theorem 2.1 allow us to keep track of exactly which information is lost. This suggest the following strategy: work on $(A, \mathscr{L})$ directly to recover the extra information needed to lift to level ℓn.

For instance, if we suppose that ℓ is prime to level $2n$, then it is clear from Theorem 2.1 that the choice of a compatible symmetric theta structure on $G(\mathscr{L}^\ell)$ is exactly the choice of a symplectic basis of $A[\ell]$ (we assume here that $\mu_\ell \subset k$). But since K is a maximal isotropic subgroup of the ℓ-torsion, this is the same as a choice of a basis $(e_1, \dots, e_g)$ of K and a supplementary isotropic subgroup of K in $A[\ell]$. This explain the technical condition ℓ prime to $2n$ of Section 1.1; for the general case we need to find a (compatible) symplectic basis of the full $A[2\ell n]$ torsion.

Now let's think "with arrow reversed", and let $K' = f(A[\ell])$ be the kernel of the contragredient isogeny $\widetilde{f} : B \to A$; from K' we want to compute $\widetilde{f}$ algorithmically.

Starting from theta functions of level n on B (from $\mathscr{M}$), we then want to go to theta functions of level ℓn on A (from $\mathscr{L}^\ell$). But the exact same information as before is also enough to fix a symmetric theta structure on $G(\mathscr{L}^\ell)$. Namely, fix a basis of the maximal isotropic group $K' \subset B[\ell]$ and a decomposition $B[\ell] = B_1[\ell] \oplus B_2[\ell]$ with $B_1[\ell] = K'$. This determines a full symplectic basis of the ℓ-torsion. The decomposition of $B[\ell]$ fixes a decomposition of $B[\ell^2]$ and thus a decomposition of $A[\ell]$ via $\widetilde{f}$, and the image of the basis of $B_2[\ell]$ give a basis of $K = A_2[\ell]$.

Concretely, let's look at an example with $g = 1$, $n = 2$ and $\ell = 3$. Then from Proposition 3.2 we readily see that the isogeny f is given by $(x_0, \dots x_5) \mapsto (x_0, x_3)$. Moreover by definition of a theta structure of level n, we can compute the action by translation by any point of n-torsion. In our situation, we are on level ℓn on A and we have a decomposition $A[\ell] = A_1[\ell] \oplus A_2[\ell]$ with $A_2[\ell] = K$. The isomorphism $\mathrm{Z}(\overline{\ell n}) \to A_1[\ell n]$ give us that $A_1[\ell]$ is generated by a point of 3-torsion T such that $(x+T)_i = (x)_{i+2}$ for $i \in \mathrm{Z}(\overline{\ell n})$ (2 being of 3-torsion in $\mathbb{Z}/6\mathbb{Z}$). Then the kernel K' of the contragredient isogeny $\widetilde{f}$ is generated by $f(T)$. We have $f(x+T) = (x_2, x_5)$ and $f(x+2T) = (x_4, x_1)$. We see that we could recover the coordinates of x from the knownledge of $f(x)$ and $f(T)$ if we were able to take "compatible" affine lifts of $f(x)$, $f(x)+f(T)$ and $f(x)+2f(T)$. But this is exactly what the theory of differential addition allow us to do as we explain in [LR12b].

Of course, a similar method applies to go from $(A, \mathscr{L})$ to $(A, \mathscr{L}^\ell)$ by taking uniform affine lifts of points of ℓ-torsion given by their level n theta coordinates. More details are given in [Rob10; Cos11]. So we don't really need to work with "arrow reversed", but in practice it is easier to do so; from a theta null point of level ℓn on A we readily get points of ℓ-torsion in level n on B, but it is a bit more complicated to get points of ℓ-torsion in level n on A. Once again, this come from the difference between the simplicity of the equation in Proposition 3.2 compared to Proposition 3.1.

Now we are almost finished describing the isogeny algorithm. By definition of the contragredient isogeny, the following diagram commutes:

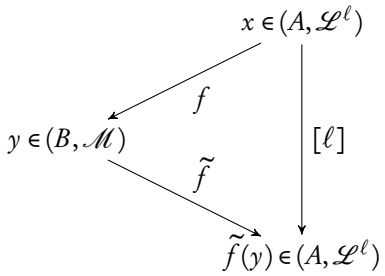

As mentioned, in [LR12b] we explain how to compute from y a point x such that $f(x) = y$. There is some ℓ-root involved, other choices of the root corresponds to different preimages (the preimage does not matter because we multiply it by $[\ell]$ afterwards).

Now $\widetilde{f}(y) = [\ell]x$. We are not quite finished because here $\widetilde{f}(y)$ is given by level ℓn theta functions. So we use the following diagram

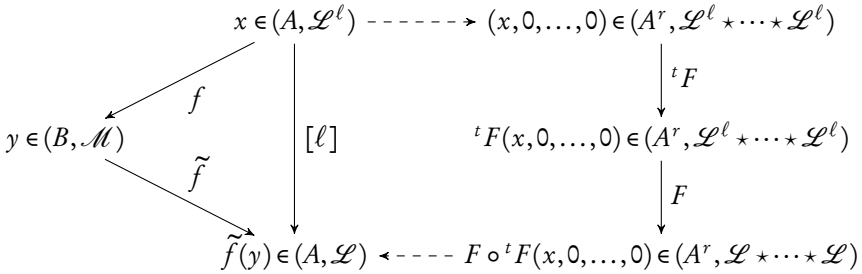

Here the computation of ${}^tF$ is done in $\mathscr{L}^\ell$ while we use Proposition 3.1 to compute the action of F in order to go back to level n.

Now fix a basis of K'. There is some ℓ-roots involved for lifting the theta null point of $(B,\mathscr{M})$ to $(A,\mathscr{L}^\ell)$ which correspond to different choices of a supplementary of K' in $B[\ell]$. Now of course these choices does not affect the end result of the computation of $\widetilde{f}(y)\in(A,\mathscr{L})$. In other words, rather than going up on $(A,\mathscr{L}^\ell)$ and then down in $(A,\mathscr{L})$ we only need to have enough informations from $(A,\mathscr{L}^\ell)$ in order to be able to go down to $(A,\mathscr{L})$. We explain how to do that in [CR11], where we get the following

Proposition 3.4 :
Let $(B,\mathscr{M}_0)$ be a ppav with a symmetric theta structure on $G(\mathscr{M})$ where $\mathscr{M}=\mathscr{M}_0^n$ is of level n even. Let $K'\subset B[\ell]$ be a maximal isotropic subgroup and $\widetilde{f}:B\to A=B/K'$ be the associated isogeny. Assume that ℓ is prime to $2n$; then the theta structure on $G(\mathscr{M})$ induces a unique polarization $\mathscr{L}$ of level n on A and a unique compatible symmetric theta structure on $G(\mathscr{L})$. Let $F\in\mathrm{Mat}_r(\mathbb{Z})$ be such that ${}^tFF=\ell\,\mathrm{Id}$.

Let $i\in K_1(\mathscr{L})$ and $(j_1,\dots,j_r)\in K_1(\mathscr{M})^r$ be the unique preimage of $(i,0,\dots,0)$ by F. Let y be a geometric point of B and let $Y={}^tF(y,0,\dots,0)\in B^r$. Then (up to a constant λ)

$$\vartheta_i^{\mathscr{L}}(\widetilde{f}(y))\cdot\dots\cdot\vartheta_0^{\mathscr{L}}(0)=\lambda\sum_{\substack{(t_1,\dots,t_r)\in K'^r\\ F(t_1,\dots,t_r)=(0,\dots,0)}}\vartheta_{j_1}^{\mathscr{M}}(Y_1+t_1)\cdot\dots\cdot\vartheta_{j_r}^{\mathscr{M}}(Y_r+t_r).\tag{1}$$

Proof : From the hypothesis ℓ prime to $2n$, Theorem 2.2 show that every compatible symmetric theta structure on $G(\mathscr{M}^\ell)$ induce the same totally symmetric line bundle $\mathscr{L}$ on A and the induced symmetric theta structure on $G(\mathscr{L})$ depends only on the choice of the symmetric theta structure on $G(\mathscr{M})$.

Now we just need to apply the diagram from above. In this diagram we apply Proposition 3.1 with $X={}^tF(x,0,\dots,0)$ where $x\in A$ is such that $f(x)=y$.

Now since ℓ is prime to n, an element $h\in K_1(\mathscr{L}^\ell)$ is of the form $h=j+T$ where $j\in K_1(\mathscr{L})$ and $T\in K_1(\mathscr{L}^\ell)[\ell]$. But by Proposition 3.2, $\vartheta_h^{\mathscr{L}^\ell}(X_i)=\vartheta_{f(j)}^{\mathscr{M}}(Y_i+f(T))$ (think about our $g=1$, $n=2$, $\ell=3$ example. Looking at the equation in Proposition 3.1 we get Equation 1.

Note that in Equation 1, the coordinates of the right hand term are not the projective coordinates of the points (it would not make sense in a sum) but of suitably normalized affine lifts. More details are given in [CR11] where we explain how to use differential additions to normalize the affine lifts.

In total, the complexity cost is given by normalizing affine lifts of the geometric points of K $O(\ell^g)$ and the changing level formula costing $O(\ell^{gr/2})$. (For an improvement, in [Rob12] we explain with David Lubicz how to adapt the formula to only need the equations of the kernel K). ■

Example 3.5 :
If $\ell=a^2+b^2$, we can take $F=\begin{pmatrix}a & b\\ -b & a\end{pmatrix}$, so that Equation 1 become

$$\vartheta_i^{\mathscr{L}}(f(y))\cdot\vartheta_0^{\mathscr{L}}(0)=\lambda\sum_{t\in K}\vartheta_{j_1}^{\mathscr{M}}(ay+at)\cdot\vartheta_{j_2}^{\mathscr{M}}(by+bt).\tag{2}$$

4 Computing isogenies with cyclic kernel

Let $f:A\to B$ be an isogeny with cyclic kernel, and assume that we have principal polarization $\mathscr{L}_0$ and $\mathscr{M}_0$ on A and B. Let $\mathscr{L}=\mathscr{L}_0^n$ and $\mathscr{M}=\mathscr{M}_0^n$.

Then there exist φ such that the following diagram commutes:

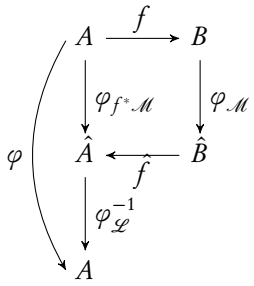

By construction, φ commutes with the Rosatti involution, so it is a (totally positive) totally real element of $\mathrm{End}^0(A)$. We note $\mathscr{L}^\varphi = f^*\mathscr{M}$ so that we have the following diagram

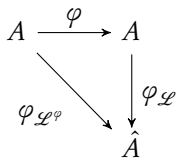

Since the commutator pairing $e_{\mathscr{L}^\varphi}$ is non degenerate (or since $\mathrm{Ker}\,\hat{f}$ is the Cartier dual of $K = \mathrm{Ker}\, f$), we see that $\mathrm{Ker}\,\varphi \subset A[\ell]$ is non isotropic for the Weil pairing. However, $K = \mathrm{Ker}\, f$ is maximally isotropic for $e_{\mathscr{L}^\varphi}$. So in Section 3 we explained how to compute an isogeny from a maximal isotropic kernel K (implicitly for Weil pairing $e_{\mathscr{L}^\ell}$), this suggest that we will be able to compute the isogeny with kernel K maximally isotropic for $e_{\mathscr{L}^\varphi}$ by replacing $[\ell]$ with φ everywhere.

Of course at one point we will need to explain how to construct $\mathscr{L}^\varphi$ without using the isogeny f, because we want to compute f from $\mathscr{L}^\varphi$.

First, the analog of Proposition 3.2 is immediate:

Proposition 4.1 :
Assume that the symmetric theta structure on $\mathscr{L}^\varphi$ is such that $K \subset K_2(\mathscr{L}^\varphi)$ (this is always possible). Then the symmetric theta structure on $\mathscr{L}^\varphi$ induces a symmetric theta structure on $\mathscr{M}$ by Theorem 2.2 (this may require to replace $\mathscr{M}$ by an equivalent line bundle).

Let $f : A \to B$ be the isogeny of kernel K, and x a geometric point in A. Fix $i \in K_1(\mathscr{M})$, and let $j \in K_1(\mathscr{L}^\varphi)$ be the unique preimage of i by f that is in $K_1(\mathscr{L}^\varphi)$. We have (up to a constant λ)

$$\vartheta_i^{\mathscr{M}}(f(x)) = \lambda \vartheta_j^{\mathscr{L}^\varphi}(x)$$

Moreover, if we introduce the φ-contragredient isogeny $\widetilde{f}$ has the isogeny $\widetilde{f} : B \to A$ such that $\widetilde{f} \circ f = \varphi$, we have the following diagram

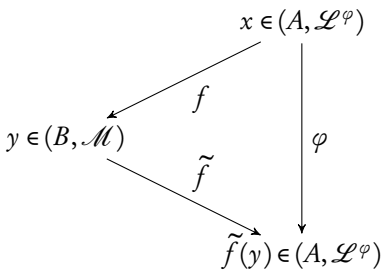

The exact same techniques as in Section 3 allow to find from $y \in B$ a preimage x, and such compute $\widetilde{f}(y)$ in coordinates from $\mathscr{L}^\varphi$. Now we just need to apply a change level formula using an equivalent of Proposition 4.

First we need to find an equivalent of the matrix F. To simplify we now assume that the division algebra $\mathrm{End}^0(A)$ is a field K, and we let K_0 be the associated totally real field. Furthermore, we assume that $O_0 = K_0 \bigcap \mathrm{End}(A)$ is the maximal order O_{K_0} of K_0 (A has maximum real multiplication).

Lemma 4.2 :
Let $\varphi \in O_{K_0}$ be a totally positive element. Then there exist $F \in \mathrm{Mat}_r(O_{K_0})$ such that ${}^tFF = \varphi\,\mathrm{Id}$.

Proof : It is well known that such a φ is a sum of m squares in O_{K_0}. We may assume that $m = 2^d$ is a power of 2. Now using the theory of Clifford's algebra for the quadratic form $Q(x_1,\dots,x_t) = -x_1^2 - x_2^2 - \cdots - x_t^2$ with $t \geqslant d$ sufficiently large, we obtain the matrix F with $r = 2^t$.

[Update 2013-04-23: as remarked by Dimitar Jetchev, a paper of Siegel show that except in $Q(\sqrt{5})$ for some elements of O_{K_0} a sum of squares can only be found using non integral elements. If we have such an element α/m, to compute its action on the ℓ-torsion, we need to compute the action of α on the ℓm-torsion, so we would like m to be as small as possible. Intuitively, for a larger r we can get a smaller m, but a large r also increase the complexity.] ■

Remark 4.3 :

- φ is a sum of two squares iff it is the norm of an element of $K_0(i)$. This is purely a local question, so it should be pretty easy to test in practice.
- In general, $\mathbb{Q}(\sqrt{5})$ is the only real quadratic field whose every integral element is a sum of 4 integral squares [TODO: check if this is correct] . So we way need to take $d > 2$.
- Also, the generic formula converting a sum of 2^d squares into a matrix of length 2^d involves denominator. That's why in the proof of the lemma we need to assume that t may be larger than d (the exact formula is given by the size of the representations of the associated Clifford's algebra).
- Still, the following will make clear that we only need to work locally on $\mathbb{Z}[\frac{1}{2\ell n}]$ so we can look for F in $\mathrm{Mat}_r(O_{K_0} \otimes \mathbb{Z}[\frac{1}{2\ell n}]$.
- All in all, I lack a clear bound on how big r could be at worse. Note that the size of r directly influence the cost of the changing level formulas (see Proposition 3.4).
- To look for smaller r, Christophe Ritzenthaler suggested looking at matrix F such that (for instance) ${}^tFF = \mathrm{diag}(\varphi, 1, \dots, 1)$. ♢

Now assume the matrix F is fixed, we have

Proposition 4.4 :
Let $(i_1,\dots,i_r) \in K_1(\mathscr{L})^r$. Let $x = (x_1,\dots,x_r)$ be a geometric point of A^r and let $y = Fx$. Then (up to a constant λ)

$$\vartheta^{\mathscr{L}}_{i_1}(y_1)\cdot\dots\cdot\vartheta^{\mathscr{L}}_{i_r}(y_r) = \lambda \sum_{\substack{(j_1,\dots,j_r)\in K_1(\mathscr{L}^\varphi)\\ F(j_1,\dots,j_r)=(i_1,\dots,i_r)}} \vartheta^{\mathscr{L}^\varphi}_{j_1}(x_1)\cdot\dots\cdot\vartheta^{\mathscr{L}^\varphi}_{j_r}(x_r).$$

Proof : It is a bit easier to look at a proof over $\mathbb{C}$. The action of F on the polarization H is given by ${}^t\overline{F}F = \varphi\,\mathrm{Id}$ (because the elements of F are real), so we have $F^*\mathscr{L} \star \mathscr{L} \cdots = \mathscr{L}^\varphi \star \mathscr{L}^\varphi \cdots$.

Note that this give the construction of $\mathscr{L}^\varphi$ we were looking for. Now the real elements of K_0 acts on $\mathbb{C}^g$ by symmetric matrixes, so they are codiagonalizable in respect to the principal polarization H_0.

In particular, the isogeny induced by F on A^r respect the symplectic decomposition given on A, so we can apply the isogeny theorem. ■

Now we just have to combine everything in the following diagram

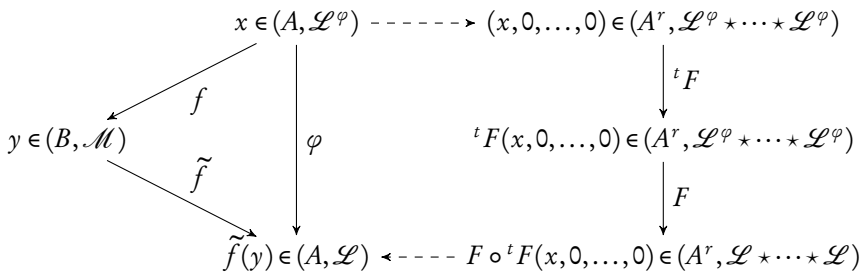

Proposition 4.5 :
Let $(B, \mathscr{M}_0)$ be a ppav with a symmetric theta structure on $G(\mathscr{M})$ where $\mathscr{M} = \mathscr{M}_0^n$ is of level n even. Let $K' \subset B[\ell]$ be a maximal isotropic subgroup for $\mathscr{M}^\varphi$ and $\widetilde{f} : B \to A = B/K'$ be the associated isogeny. Assume that ℓ is prime to $2n$; then the theta structure on $G(\mathscr{M})$ induces a unique polarization $\mathscr{L}$ of level n on A and a unique compatible symmetric theta structure on $G(\mathscr{L})$. Let $F \in \mathrm{Mat}_r(O_{K_0})$ be such that ${}^tFF = \varphi\,\mathrm{Id}$.

Let $i \in K_1(\mathscr{L})$ and $(j_1, \dots, j_r) \in K_1(\mathscr{M})^r$ be the unique preimage of $(i,0,\dots,0)$ by F. Let y be a geometric point of B and let $Y = {}^tF(y,0,\dots,0) \in B^r$. Then (up to a constant λ that may depend on y this time)

$$\vartheta_i^{\mathscr{L}}(\widetilde{f}(y)) \cdot \dots \cdot \vartheta_0^{\mathscr{L}}(0) = \lambda \sum_{\substack{(t_1,\dots,t_r) \in K'^r \\ F(t_1,\dots,t_r) = (0,\dots,0)}} \vartheta_{j_1}^{\mathscr{M}}(Y_1 + t_1) \cdot \dots \cdot \vartheta_{j_r}^{\mathscr{M}}(Y_r + t_r). \tag{3}$$

Note that the condition of having maximal real multiplication is too strong, we only need to have a matrix F corresponding to φ. In particular, we don't really need to have maximal real multiplication, nor even that A and B have the same real multiplication. Of course, we do need to have φ in $\mathrm{End}(A)$ and $\mathrm{End}(B)$, where we abuse the same notation to note $\varphi = \widetilde{f} \circ f \in \mathrm{End}(A)$ and $f \circ \widetilde{f} \in \mathrm{End}(B)$. Perhaps the following diagram is clearer:

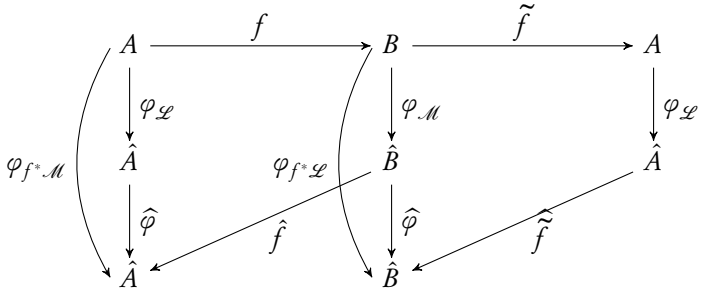

4.1 Computing the isogeny in practice

Of course in Proposition 4.5 we have hidden all the difficulties in the computation of

$$\vartheta_{j_1}^{\mathscr{M}}(Y_1 + t_1) \cdot \dots \cdot \vartheta_{j_r}^{\mathscr{M}}(Y_r + t_r),$$

where we need to have a way to compute the action of the elements of O_{K_0} giving F in a "compatible affine manner".

The easy case is if we only need the isogenous theta null point. In which case $y = 0$ and $Y = (0,\dots,0)$ so that we only need to evaluate on points of K' but we have already seen how to normalize the affine lifts [CR11]. But to compute the image of a point y we need to work harder.

We give an example on how to do that with $O_{K_0} = \mathbb{Q}(\sqrt{d})$ (d prime to ℓ) and $\varphi = a^2 + b^2$ (the generalization to a sum of more squares is immediate) so that we can take $F = \begin{pmatrix} a & b \\ -b & a \end{pmatrix}$ as in Example 3.5 We need to evaluate

$$\sum_{t \in K} \vartheta_{j_1}^{\mathscr{M}}(ay + at) \cdot \vartheta_{j_2}^{\mathscr{M}}(by + bt). \tag{4}$$

We want to compute affine coordinates of $ay+at$ and $by+bt$, where the eventual projective factor depends only on y, not on t. Let $a=\alpha+\beta\sqrt{d}$ and let's see what we can compute.

Since we have normalized all the points of K', we know αt, $\beta\sqrt{d}t$ and $\alpha+\beta\sqrt{d}$ already. We also know the "affine coordinates" of αy and $\alpha(y+t)$, this only use differential additions.

We also can compute $\sqrt{d}y$ since $\sqrt{d}$ correspond to a (d,d)-isogeny (a normal one with maximal isotropic kernel for the Weil pairing). The important point here is that Proposition 3.4 gives us the isogeny for **affine theta coordinates** (since λ is a constant). From $\sqrt{d}y$ we get $\beta\sqrt{d}y$ using differential additions. Likewise we can compute $\beta\sqrt{d}(y+t)$. If $\alpha t=\beta\sqrt{d}t'$, then $\beta\sqrt{d}y+\alpha t$ is simply $\beta\sqrt{d}(y+t')$ so we can also compute it.

Finally we can compute $\alpha+\beta\sqrt{d}y$ but only in a projective way, so we have take an arbitrary affine lift. The important point here is that we can fix it once and for all, it does not depend on t.

In the sum of four terms $\alpha y+\beta\sqrt{d}y+\alpha t+\beta\sqrt{d}t$, we have seen how to compute each of the two by two subsum. Now this is what we call a MultiWayAddition, and we claim that by using Riemann relations, this is enough to compute the whole sum. Indeed, it is easy to see that a MultiWayAddition reduces to several ThreeWayAdditions (compute $x+y+z$ from x, y, z, $x+y$, $x+z$, $y+z$) and we showed how to do that in [Rob10; LR13] (generically in level 2, for any geometric point in level $n>2$.)

Remark 4.6 :

- Finding the matrix F requires that we know what the full real endomorphism order look like, which may be expensive. Over a finite field, it should be possible by Tate's theorem to work on the ℓ-Tate module to find the action of F on the ℓ-torsion, which is what we need if we only want the isogenous theta null point (we also need the action on the 2-torsion, so we'll need to glue things).
- It would be interesting to have a purely analytic version of Proposition 4.5. Note that the analytic version of Koizumi [Koi76] is a bit stronger as stated than the algebraic version of Kempf [Kem89] (for instance to recover the usual Riemman's relation, one need to apply Kempf's version twice). ♢

References

[BL04] C. Birkenhake and H. Lange. *Complex abelian varieties*. Second. Vol. 302. Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]. Berlin: Springer-Verlag, 2004, pp. xii+635. ISBN: 3-540-20488-1 (cit. on pp. 3, 6).

[BCR10a] G. Bisson, R. Cosset, and D. Robert. "AVIsogenies (Abelian Varieties and Isogenies)". Packet magma dédié au calcul explicite d'isogénies entre variétés abéliennes. 2010. URL: http://avisogenies.gforge.inria.fr. Licence libre (LGPLv2+), enregistré à l'APP (référence IDDN.FR.001.440011.000.R.P.2010.000.10000) (cit. on p. 1).

[BCR10b] G. Bisson, R. Cosset, and D. Robert. "On the Practical Computation of Isogenies of Jacobian Surfaces". Article explaining the computations done in the AVIsogenies package, found in the source code. 2010. URL: http://avisogenies.gforge.inria.fr (cit. on p. 1).

[Cos11] R. Cosset. "Application des fonctions thêta à la cryptographie sur courbes hyperelliptiques". PhD thesis. 2011 (cit. on pp. 1, 7).

[CR11] R. Cosset and D. Robert. "An algorithm for computing (ℓ,ℓ)-isogenies in polynomial time on Jacobians of hyperelliptic curves of genus 2". Mar. 2011. URL: http://www.normalesup.org/~robert/pro/publications/articles/niveau.pdf. HAL: hal-00578991, eprint: 2011/143 (cit. on pp. 1, 2, 4, 8, 11).

[FLR11] J.-C. Faugère, D. Lubicz, and D. Robert. "Computing modular correspondences for abelian varieties". In: *Journal of Algebra* 343.1 (Oct. 2011), pp. 248–277. DOI: 10.1016/j.jalgebra.2011.06.031. arXiv: 0910.4668 [cs.SC]. URL: http://www.normalesup.org/~robert/pro/publications/articles/modular.pdf. HAL: hal-00426338 (cit. on pp. 1, 6).

[Igu72] J.-i. Igusa. *Theta functions*. Die Grundlehren der mathematischen Wissenschaften, Band 194. New York: Springer-Verlag, 1972, pp. x+232 (cit. on p. 3).

[Kem89] G. Kempf. "Linear systems on abelian varieties". In: *American Journal of Mathematics* 111.1 (1989), pp. 65–94 (cit. on pp. 1, 4, 12).

[Koi76] S. Koizumi. "Theta relations and projective normality of abelian varieties". In: *American Journal of Mathematics* (1976), pp. 865–889 (cit. on pp. 1, 3, 12).

[LR12a] K. Lauter and D. Robert. "Improved CRT Algorithm for class polynomials in genus 2". In: *ANTS* (2012). Accepted for publication at the Tenth Algorithmic Number Theory Symposium ANTS-X. University of California, San Diego, July 9 – 13, 2012 http://math.ucsd.edu/~kedlaya/ants10/. URL: http://www.normalesup.org/~robert/pro/publications/articles/classCRT.pdf. Slides http://www.normalesup.org/~robert/publications/slides/2012-07-ANTS-SanDiego.pdf, eprint: 2012/443, HAL: hal-00734450 (cit. on p. 2).

[LR12b] D. Lubicz and D. Robert. "Computing isogenies between abelian varieties". In: *Compositio Mathematica* 148.05 (Sept. 2012), pp. 1483–1515. DOI: 10.1112/S0010437X12000243. arXiv: 1001.2016 [math.AG]. URL: http://www.normalesup.org/~robert/pro/publications/articles/isogenies.pdf. HAL: hal-00446062 (cit. on pp. 1, 7).

[LR13] D. Lubicz and D. Robert. "A generalisation of Miller's algorithm and applications to pairing computations on abelian varieties". Mar. 2013. URL: http://www.normalesup.org/~robert/pro/publications/articles/optimal.pdf. HAL: hal-00806923, eprint: 2013/192 (cit. on p. 12).

[Mum66] D. Mumford. "On the equations defining abelian varieties. I". In: *Invent. Math.* 1 (1966), pp. 287–354 (cit. on pp. 3, 4, 6).

[Mum67a] D. Mumford. "On the equations defining abelian varieties. II". In: *Invent. Math.* 3 (1967), pp. 75–135 (cit. on pp. 3, 6).

[Mum67b] D. Mumford. "On the equations defining abelian varieties. III". In: *Invent. Math.* 3 (1967), pp. 215–244 (cit. on p. 3).

[Mum70] D. Mumford. *Abelian varieties*. Tata Institute of Fundamental Research Studies in Mathematics, No. 5. Published for the Tata Institute of Fundamental Research, Bombay, 1970, pp. viii+242 (cit. on p. 3).

[Mum83] D. Mumford. *Tata lectures on theta I*. Vol. 28. Progress in Mathematics. With the assistance of C. Musili, M. Nori, E. Previato and M. Stillman. Boston, MA: Birkhäuser Boston Inc., 1983, pp. xiii+235. ISBN: 3-7643-3109-7 (cit. on p. 3).

[Mum84] D. Mumford. *Tata lectures on theta II*. Vol. 43. Progress in Mathematics. Jacobian theta functions and differential equations, With the collaboration of C. Musili, M. Nori, E. Previato, M. Stillman and H. Umemura. Boston, MA: Birkhäuser Boston Inc., 1984, pp. xiv+272. ISBN: 0-8176-3110-0 (cit. on p. 3).

[Mum91] D. Mumford. *Tata lectures on theta III*. Vol. 97. Progress in Mathematics. With the collaboration of Madhav Nori and Peter Norman. Boston, MA: Birkhäuser Boston Inc., 1991, pp. viii+202. ISBN: 0-8176-3440-1 (cit. on p. 3).

[Rob10] D. Robert. "Fonctions thêta et applications à la cryptographie". PhD thesis. Université Henri-Poincarré, Nancy 1, France, July 2010. URL: http://www.normalesup.org/~robert/pro/publications/academic/phd.pdf. Slides http://www.normalesup.org/~robert/pro/publications/slides/2010-07-phd.pdf, TEL: tel-00528942. (Cit. on pp. 1, 4, 7, 12).

[Rob12] D. Robert. "Computing rational isogenies from the equations of the kernel". ANR Peace Meeting, Paris. Nov. 2012. URL: http://www.normalesup.org/~robert/pro/publications/slides/2012-11-Peace.pdf (cit. on pp. 1, 8).

[Rob13] D. Robert. "Computing cyclic isogenies using real multiplication". ANR Peace Meeting, Paris. Notes available on http://www.normalesup.org/~robert/pro/publications/notes/2013-04-cyclic-isogenies.pdf. Apr. 2013 (cit. on p. 1).

[Shi98] G. Shimura. *Abelian varieties with complex multiplication and modular functions*. Vol. 46. Princeton University Press, 1998 (cit. on p. 2).

[Vél71] J. Vélu. “Isogénies entre courbes elliptiques”. In: *Compte Rendu Académie Sciences Paris Série A-B* 273 (1971), A238–A241 (cit. on p. 1).

[Wam99] P. Wamelen. “Equations for the Jacobian of a hyperelliptic curve”. In: *AMS* 350.8 (Aug. 1999), pp. 3083–3106 (cit. on p. 1).