Isogenies between abelian varieties – an algorithmic survey

2022/09/21 — Isogeny days, Leuven

Damien Robert

Équipe LFANT, Inria Bordeaux Sud-Ouest

Outline

Postdoc

- The ANR CIAO is looking for a one year postdoc in Bordeaux
  https://anr.fr/Projet-ANR-19-CE48-0008
- Topics: anything related to isogeny based cryptography
- Position available until 2024-04 (should be extendable by 6 months)
- Email: http://www.normalesup.org/~robert/pro/infos.html

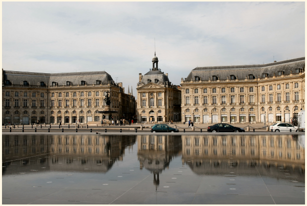

Outline

Usage of isogenies

- Speed up the arithmetic (eg split the multiplication by $[2]$ or $[3]$);
- Determine $\mathrm{End}(A)$ (volcano...);
- Point counting algorithms (ℓ-adic or p-adic: SEA, Satoh ...)
  Publicity: [Kieffer 2021] SEA like algorithm in $\widetilde{O}_K(\log^4 q)$ for abelian surfaces with RM by O_K.

- Compute class polynomials (CM-method)
- Compute modular polynomials

- Arithmetic for $\mathbb{F}_q$: construct normal basis of a finite field, irreducible polynomials, automorphism invariant smoothness basis [Couveignes-Lercier]...
- Find curves with many points
- Explore isogeny graphs (eg find a component with no Jacobians in dimension 4)
- Evaluate modular forms

Isogenies in classical cryptography

- Discrete Logarithm Problem, Pairings
- Transfer the DLP (Weil descent…)
- Reduce the impact of side channel attacks
- Random self reducibility, worst case to average case reductions.

Isogeny based cryptography

- Hash functions
- Key exchange (SIDH, CSIDH)
- Signatures (SQISign)

Higher dimensional isogenies?

- Classical cryptography: dimension 1 and 2. A bit in dimension 3 (class polynomials).
- Isogeny based cryptography: dimension 1 (hash functions in dimension 2 too).
- So mainly for algorithmic number theory (descent…)
- Certainly no use for elliptic curve based cryptosystems.

Higher dimensional isogenies?

- Classical cryptography: dimension 1 and 2. A bit in dimension 3 (class polynomials).
- Isogeny based cryptography: dimension 1 (hash functions in dimension 2 too).
- So mainly for algorithmic number theory (descent…)
- ~~Certainly no use for elliptic curve based cryptosystems.~~

The embedding lemma

- A N-isogeny $f : A \to B$ in dimension g can always be efficiently embedded into a N' isogeny $F : A' \to B'$ in dimension $8g$ (and sometimes $4g, 2g$) for any $N' \geq N$.

$$\begin{array}{ccc} A & \xrightarrow{f} & B \\ \big\downarrow & & \big\uparrow \\ A' & \xrightarrow{F} & B' \end{array}$$

- Considerable flexibility (at the cost of going up in dimension).

The embedding lemma

- A N-isogeny $f : A \to B$ in dimension g can always be efficiently embedded into a N' isogeny $F : A' \to B'$ in dimension $8g$ (and sometimes $4g, 2g$) for any $N' \geq N$.

$$\begin{array}{ccc} A & \xrightarrow{f} & B \\ \downarrow & & \uparrow \\ A' & \xrightarrow{F} & B' \end{array}$$

- Considerable flexibility (at the cost of going up in dimension).
- Write $N' - N = a_1^2 + a_2^2 + a_3^2 + a_4^2$.
- $F = \begin{pmatrix} a_1 & -a_2 & -a_3 & -a_4 & \hat{f} & 0 & 0 & 0 \\ a_2 & a_1 & a_4 & -a_3 & 0 & \hat{f} & 0 & 0 \\ a_3 & -a_4 & a_1 & a_2 & 0 & 0 & \hat{f} & 0 \\ a_4 & a_3 & -a_2 & a_1 & 0 & 0 & 0 & \hat{f} \\ -f & 0 & 0 & 0 & a_1 & a_2 & a_3 & a_4 \\ 0 & -f & 0 & 0 & -a_2 & a_1 & -a_4 & a_3 \\ 0 & 0 & -f & 0 & -a_3 & a_4 & a_1 & a_2 \\ 0 & 0 & 0 & -f & -a_4 & -a_3 & a_2 & a_1 \end{pmatrix}$

The embedding lemma

- A N-isogeny $f : A \to B$ in dimension g can always be efficiently embedded into a N' isogeny $F : A' \to B'$ in dimension $8g$ (and sometimes $4g, 2g$) for any $N' \geq N$.

$$\begin{array}{ccc} A & \xrightarrow{f} & B \\ \big\downarrow & & \big\uparrow \\ A' & \xrightarrow{F} & B' \end{array}$$

- Considerable flexibility (at the cost of going up in dimension).

- Breaks SIDH ([Castryck-Decru], [Maino-Martindale] in dimension 2, [R.] in dimension 4 or 8) $\Rightarrow$ if $N_A > N_B$, take $N' = N_A, N = N_B$
  The dimension 8 attack is in proven quasi-linear time, see http://www.normalesup.org/~robert/pro/publications/slides/2022-09-Bordeaux-SIDH.pdf for details.

- An isogeny always have a representation allowing evaluation in polylogarithmic time $\log^{O(1)} N$ [R.] $\Rightarrow$ take $N' \geq N$ powersmooth.
  (Finding this representation takes quasi-linear time.)

The embedding lemma

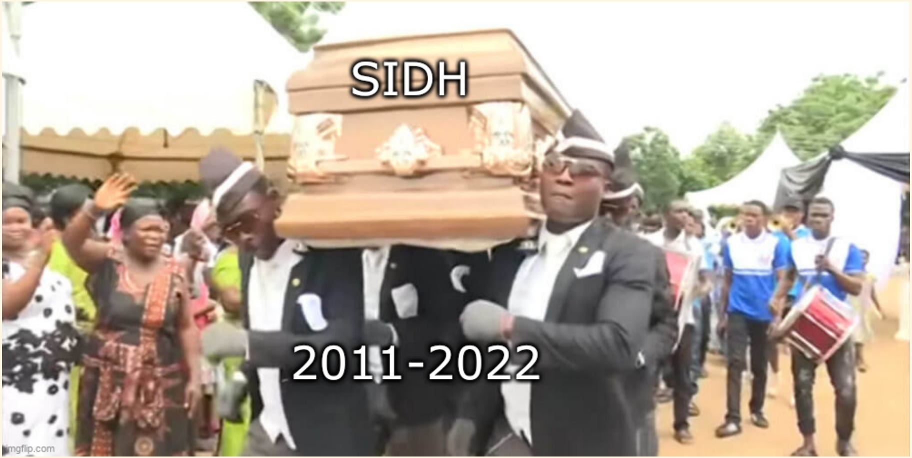

Isogeny diamonds

- $f_1 : A \to A_1$ n_1-isogeny, $f_1' : A_1 \to B$ n_1'-isogeny, $f_2 : A \to A_2$ n_2-isogeny, $f_2' : A_2 \to B$ n_2'-isogeny, $f_2' \circ f_2 = f_1' \circ f_1$.

$$\begin{array}{ccc} A & \xrightarrow{f_1} & A_1 \\ \downarrow{\scriptstyle f_2} & & \downarrow{\scriptstyle f_1'} \\ A_2 & \xrightarrow{f_2'} & B \end{array}$$

- $F = \begin{pmatrix} f_1 & \widetilde{f_1'} \\ -f_2 & \widetilde{f_2'} \end{pmatrix}$ is an $\begin{pmatrix} n_1 + n_2 & 0 \\ 0 & n_1' + n_2' \end{pmatrix}$-isogeny.

- Isogeny diamonds: If $n_1' = n_2$ (so $n_2' = n_1$), F is an N-isogeny where $N = n_1 + n_2$ ([Kani] for $g = 1$, [R.] for $g > 1$.)

Algorithms for N-isogenies

Jacobian model:

- Vélu's formula for elliptic curves [Vélu 1971]
- [Kohel, 1999]: Vélu's formula from equations of K;

- [Richelot, 1836,1837] 2-isogenies between Jacobians of genus 2 hyperelliptic curves, [Mestre 2013] for general g;
- Various explicit formula for small degree isogenies in dimension 2;
- [Smith 2008]: 2-isogenies for quartic genus 3 curves;
- [R. 2007]: the analog of Vélu's formula for genus 2 does not seem to work?

- [Couveignes-Ezome (2015)]: Algorithm in $\widetilde{O}(N^g)$ in the Jacobian model (complete algorithm for $g = 2$, [Milio 2019] for $g = 3$).
- Restricted to $g \leq 3$.

Algorithms for N-isogenies

Jacobian model:

- Vélu's formula for elliptic curves [Vélu 1971]
- [Couveignes-Ezome (2015)]: Algorithm in $\widetilde{O}(N^g)$ in the Jacobian model (complete algorithm for $g = 2$, [Milio 2019] for $g = 3$).

Theta model:

- 2-isogenies: duplication formula for theta functions [Riemann ?]
- [Mumford, 1966] isogeny formula, [Koizumi 1976, Kempf 1989] product formula (requires theta constants of higher level)
- [Lubicz-R. 2012]: ℓ^2-isogenies between abelian varieties in $O(\ell^g)$ and $\ell^{g(g+1)/2}$ ℓ-th roots.
  This corresponds to taking an ℓ-isogeny, and then each choice of roots prolongs this ℓ-isogeny into a different ℓ^2-isogeny (we get all ℓ^2-isogenies whose kernel stays of rank g), see also [Castryck, Decru, Vercauteren] work on radical isogenies.
- [Cosset-R. (2014)]: ℓ-isogenies in $O(\ell^g)$ if $\ell \equiv 1 \pmod 4$, $O(\ell^{2g})$ if $\ell \equiv 3 \pmod 4$;
- [Lubicz-R. (2022)]: An N-isogeny in dimension g can be evaluated in linear time $O(N^g)$ arithmetic operations in the theta model given generators of its kernel.
- Warning: exponential dependency 2^g or 4^g in the dimension g.
- [Lubicz-R. (2015)]: isogenies from equations of the kernel
- [Dudeanu, Jetchev, R., Vuille (2022)]: cyclic isogenies for abelian varieties with RM.

Outline

Polarised abelian varieties over $\mathbb{C}$

Definition

A complex abelian variety A of dimension g is isomorphic to a compact Lie group V/Λ with

- A complex vector space V of dimension g (linear data);
- A $\mathbb{Z}$-lattice Λ in V (of rank $2g$) (arithmetic data);
- A polarisation (quadratic data)

Example

- A vector space $V \simeq \mathbb{C}^g$ is described by a basis;
- A lattice $\Lambda = \Omega\mathbb{Z}^g \oplus \mathbb{Z}^g$ is described by a period matrix Ω;
- The quotient $\mathbb{C}^g/\Lambda$ is a torus. It is not an abelian variety in general!
- The moduli space of torus is of dimension g^2.
- If $\Omega \in \mathfrak{H}_g$, $H = \operatorname{Im} \Omega^{-1}$ is a principal polarisation.
- The moduli space of abelian varieties is of dimension $g(g+1)/2$.
- NB: when $g = 1$ both spaces have dimension 1.

Polarisations

$A = V/\Lambda$. A polarisation on A is:

- An Hermitian form H on V with $\operatorname{Im} H(\Lambda, \Lambda) \subset \mathbb{Z}$;
- A symplectic form E on H with $E(\Lambda, \Lambda) \subset \mathbb{Z}$: $E = \operatorname{Im} H$
- A (symmetric) morphism $\Phi : A \to \widehat{A}$: $\Phi = \Phi_H : z \mapsto H(z, \cdot) \in \widehat{A} = \operatorname{Hom}_{\overline{\mathbb{C}}}(V, \mathbb{C})$
- (The algebraic equivalence class of) a divisor $\mathscr{D}$ [Apell-Humbert].

Divisors and theNéron-Severi group

- To work algorithmically with an abelian variety, we need (projective) coordinates $u_1, \dots, u_m$;
- A point $P \in A$ is represented by its coordinates $(u_1(P) : \cdots, u_m(P))$.
- Coordinates are given by sections of (very ample) divisors;

- Linearly equivalent divisors $\mathcal{D} \simeq \mathcal{D}'$ give isomorphic coordinates;
- $\mathrm{Pic}(A)$: divisors modulo linear equivalence.
- $\mathcal{D} \sim \mathcal{D}'$ are algebraically equivalent $\Leftrightarrow$ $\mathcal{D}'$ is linearly equivalent to a translate of $\mathcal{D}$, ie $\mathcal{D}' \simeq t_x\mathcal{D}$ (if $\mathcal{D}$ is ample);

  $\mathcal{D}' \simeq t_x\mathcal{D} \Rightarrow \mathcal{D}' \sim \mathcal{D}$ and the converse is true if $\Phi_{\mathcal{D}}$ is surjective, ie the polarisation is non degenerate.

- Algebraically equivalent divisors = same coordinates up to translation;
- Néron-Severi group $NS(A) = \mathrm{Pic}(A) / \mathrm{Pic}^0(A)$: divisors modulo algebraic equivalence.

  More precisely: $NS(A)$ is the fppf sheaf associated to the functor $\mathrm{Pic}(A) / \mathrm{Pic}^0(A)$. Here $\mathrm{Pic}^0(A)$ is the connected component of the Picard group, it corresponds to divisors algebraically equivalent to 0, or equivalently to divisors D_0 such that $\Phi_{D_0} = 0$, ie $t_P^* D_0 \simeq D_0$ for all $P \in A$.

  So an algebraic class $\lambda = [\mathcal{D}]$ may be rational with no representative $\mathcal{D}$ defined over k. This does not happens when $k = \mathbb{F}_q$, representatives form a torsor under $\widehat{A} = \mathrm{Pic}^0(A)$, and this torsor is trivial, ie has a section, since $H^1(\mathbb{F}_q, \widehat{A}) = 0$.

  In general, the pullback $\mathcal{D}' = (1 \times \lambda)^* P$ of the Poincarre sheaf satisfy $\Phi_{\mathcal{D}'} = 2\lambda$, so 2λ is always represented by a rational divisor.

Facets of polarisations

Polarisation $\lambda =$

- a divisor Θ up to algebraic equivalence;
- a (symmetric) morphism $\lambda : A \to \widehat{A}$.
  $\lambda = \Phi_\Theta : A \to \widehat{A}, P \mapsto t_P^* \Theta - \Theta$.
  Ker $\lambda \simeq (\mathbb{Z}^g / D\mathbb{Z}^g)^2$ with $D = (d_1, \dots, d_g), d_i \mid d_{i+1}$: λ is of type $(d_1, \dots, d_g)$.
  $\deg \Theta := \prod d_i$.
- a pairing $T_\ell A \times T_\ell A \to Z(\bar{\ell})(1), (P, Q) \mapsto e_\lambda(P, Q) = e_A(P, \lambda Q)$;

Facets of polarisations

Polarisation λ =

- a divisor Θ up to algebraic equivalence;
- a (symmetric) morphism $\lambda : A \to \widehat{A}$.
  $\lambda = \Phi_\Theta : A \to \widehat{A}, P \mapsto t_P^* \Theta - \Theta$.
  Ker $\lambda \simeq (\mathbb{Z}^g / D\mathbb{Z}^g)^2$ with $D = (d_1, \dots, d_g), d_i \mid d_{i+1}$: λ is of type $(d_1, \dots, d_g)$.
  deg $\Theta := \prod d_i$.
- a pairing $T_\ell A \times T_\ell A \to Z(\bar{\ell})(1), (P,Q) \mapsto e_\lambda(P,Q) = e_A(P, \lambda Q)$;

The polarisation λ is

- Non degenerate if $\lambda : A \to \widehat{A}$ is an isogeny;
- Positive if $\lambda = \Phi_\Theta$ and Θ is ample ($\Rightarrow$ non degenerate).
- Principal if λ is (positive and) an isomorphism.

Facets of polarisations

Polarisation $\lambda =$

- a divisor Θ up to algebraic equivalence;
- a (symmetric) morphism $\lambda : A \to \widehat{A}$.
  $\lambda = \Phi_\Theta : A \to \widehat{A}, P \mapsto t_P^* \Theta - \Theta$.
  $\operatorname{Ker} \lambda \simeq (\mathbb{Z}^g / D\mathbb{Z}^g)^2$ with $D = (d_1, \dots, d_g), d_i \mid d_{i+1}$: λ is of type $(d_1, \dots, d_g)$.
  $\deg \Theta := \prod d_i$.
- a pairing $T_\ell A \times T_\ell A \to Z(\bar{\ell})(1), (P, Q) \mapsto e_\lambda(P, Q) = e_A(P, \lambda Q)$;

The polarisation λ is

- Non degenerate if $\lambda : A \to \widehat{A}$ is an isogeny;
- Positive if $\lambda = \Phi_\Theta$ and Θ is ample ($\Rightarrow$ non degenerate).
- Principal if λ is (positive and) an isomorphism.

Example

If H polarisation on $A = V/\Lambda$: $H \simeq \begin{pmatrix} \lambda_1 & & 0 \\ & \ddots & \\ 0 & & \lambda_g \end{pmatrix}$, $\lambda_i \in \mathbb{R}$, $E = \operatorname{Im} H \simeq \begin{pmatrix} 0 & D \\ -D & 0 \end{pmatrix}$ with

$D = \begin{pmatrix} d_1 & & 0 \\ & \ddots & \\ 0 & & d_g \end{pmatrix}$ on Λ, $d_1 \mid d_2 \cdots \mid d_g$, $\operatorname{Ker} \Phi_H \simeq \Lambda^\perp / \Lambda \simeq (\mathbb{Z}^g / D\mathbb{Z}^g)^2$.

- H non degenerate $\Leftrightarrow \lambda_i \neq 0$;
- H positive $\Leftrightarrow \lambda_i > 0$;

Facets of polarisations

Polarisation $\lambda =$

- a divisor Θ up to algebraic equivalence;
- a (symmetric) morphism $\lambda : A \to \widehat{A}$.
  $\lambda = \Phi_\Theta : A \to \widehat{A}, P \mapsto t_P^* \Theta - \Theta$.
  $\operatorname{Ker} \lambda \simeq (\mathbb{Z}^g / D\mathbb{Z}^g)^2$ with $D = (d_1, \dots, d_g), d_i \mid d_{i+1}$: λ is of type $(d_1, \dots, d_g)$.
  $\deg \Theta := \prod d_i$.
- a pairing $T_\ell A \times T_\ell A \to Z(\bar{\ell})(1), (P,Q) \mapsto e_\lambda(P,Q) = e_A(P, \lambda Q)$;

The polarisation λ is

- Non degenerate if $\lambda : A \to \widehat{A}$ is an isogeny;
- Positive if $\lambda = \Phi_\Theta$ and Θ is ample ($\Rightarrow$ non degenerate).
- Principal if λ is (positive and) an isomorphism.

Coordinates: if Θ is an ample divisor:

- $\dim H^0(\Theta) = \Theta^g / g! = \deg \Theta$, "degree" of the polarisation (Riemann-Roch).
  So if Θ is a principal polarisation, $\dim H^0(N\Theta) = N^g$.
  More generally, if $\mathscr{D}$ is ample, $\dim H^0(\mathscr{D}) = \prod_{i=1}^g d_i = \deg \mathscr{D} = \deg \Phi_{\mathscr{D}}^{1/2}$: the degree of the isogeny $\Phi_{\mathscr{D}}$ associated to $\mathscr{D}$ is the square of the "degree" of $\mathscr{D}$.
- 3Θ is very ample (Lefschetz).
- 2Θ descends to $K_A = A / \pm 1$ if Θ is a principal polarisation, and is very ample there if Θ is indecomposable.
- 2Θ is very ample if it is base point free;

Jacobians

- C curve of genus g.
- $\mathrm{Jac}(C) \simeq \mathrm{Pic}^0(C)$ its Jacobian.
- $\mathrm{Jac}(C) \sim C^{\langle g \rangle}$
- $\Theta_C = \{$ degenerate divisors on $C\}$ (the Theta divisor) is a principal polarisation on $\mathrm{Jac}(C)$.
  Ex: when $g = 2, C \simeq \Theta_C C \subset \mathrm{Jac}(C)$.
- C is determined by $(\mathrm{Jac}(C), \Theta_C)$ (Torelli)

  They have the same field of moduli, but if C is not hyperelliptic the field of definition of $(\mathrm{Jac}(C), \Theta_C)$ can be smaller than the field of definition of C.

Jacobians

Example

- $C/\mathbb{C}$ curve of genus g;
- V the dual of the space $V^\vee = H^0(C, \Omega^1_C)$ of holomorphic differentials of the first kind on C;
- $\Lambda \simeq H_1(C, \mathbb{Z}) \subset V$ the set of periods.

  The Abel-Jacobi map Φ is the integration of differentials on loops: $H^0(C, \Omega^1_C) \times H_1(C, \mathbb{Z}) \mapsto \mathbb{C}, (\omega, \gamma) \mapsto \int_\gamma \omega$; it induces $\Phi : H_1(C, \mathbb{Z}) \to \mathrm{Hom}(H^0(C, \Omega^1_C), \mathbb{C})$ and Λ is the image of Φ.

  By Poincare-Serre's duality: $\mathrm{Alb}(C) \simeq H^0(C, \Omega^1_C)^\vee / H_1(C, \mathbb{Z}) \simeq H^0(C, O_C)/H^1(\mathbb{C}, \mathbb{Z}) \simeq H^1(X, {O_C}^*) \simeq \mathrm{Pic}^0(C) = \mathrm{Jac}(C)$.
- The intersection pairing $H_1(C, \mathbb{Z}) \times H_1(C, \mathbb{Z}) \to \mathbb{Z}$ gives a symplectic form E on Λ;
- H the associated Hermitian form on V (via the integration pairing):

$$H^*(w_1, w_2) = \int_C w_1 \wedge w_2;$$

- $(V/\Lambda, H)$ is a principally polarised abelian variety: the Jacobian of C.

Elliptic curves vs abelian varieties

E elliptic curve

- $D \mapsto \deg D$ induces an isomorphism $NS(E) \simeq \mathbb{Z}$;
- $[(0_E)]$: unique principal polarisation
- $E \simeq \hat{E}$ via $P \mapsto (P) - (0_E)$
- $\Gamma(0_E) = \langle 1 \rangle, \Gamma(2(0_E)) = \langle 1, x \rangle$: embedding of $E/\pm 1$,
  $\Gamma(3(0_E)) = \langle 1, x, y \rangle$: Weierstrass model $y^2 + a_1xy + a_3y = x^3 + a_2x^2 + a_4x + a_6$.

The same principally polarised abelian variety A (ppav) could be, depending on its polarisation Θ_A:

- A product of elliptic curves;
- Non decomposable;
- The Jacobian of an hyperelliptic curve;
- The Jacobian of a non hyperelliptic curve ($g \geq 3$);
- Not a Jacobian ($g \geq 4$)

Outline

Isogenies and dual isogenies

- $f : A \to B$ morphism $\Leftrightarrow$ algebraic map + group morphism (it suffices to check $f(0_A) = 0_B$ by rigidity);
- f isogeny $\Leftrightarrow$ Ker f finite + surjective $\Leftrightarrow \dim A = \dim B$ and surjective $\quad \Leftrightarrow \dim A = \dim B$ and Ker f finite;
- Divisibility: $g_1 \circ f = g_2 \circ f \Rightarrow g_1 = g_2$, $f \circ g_1 = f \circ g_2 \Rightarrow g_1 = g_2$.

- Dual isogeny $\hat{f} : \hat{B} = \mathrm{Pic}^0(B) \to \widehat{A} = \mathrm{Pic}^0(A), \hat{f}(Q) := f^* D_Q$.
- $(\widehat{g \circ f}) = \hat{f} \circ \hat{g}$;

- Pairings:

  $0 \to K \to A \xrightarrow{f} B \to 0$ induces $0 \to \hat{K} \to \hat{B} \xrightarrow{\hat{f}} \widehat{A} \to 0$ with $\hat{K} \simeq \mathrm{Hom}(K, \mathbb{G}_m)$.

  Apply $\mathrm{Hom}(\cdot, \mathbb{G}_m)$ and use $\widehat{A} \simeq \mathrm{Ext}^1(A, \mathbb{G}_m)$
- $e_f : K \times \hat{K} \to \mathbb{G}_m$ Weil-Cartier pairing
- $f = [\ell]$: $e_{W,\ell} : A[\ell] \times \widehat{A}[\ell] \to \mu_\ell$ Weil pairing;
- Compatibility of pairings and isogenies: on $T_\ell A \times T_\ell \hat{B}$,

$$e_f(x,y) = e_B(f(x), y) = e_A(x, \hat{f}(y)).$$

- Biduality: $\widehat{\widehat{A}} \simeq A, \hat{\hat{f}} \simeq f$ (canonically).

  By the universal property of $\widehat{A} = \mathrm{Pic}^0(A)$, $\mathrm{id} : \widehat{A} \to \widehat{A}$ corresponds to the Poincaré sheaf P on $A \times \widehat{A}$, and P is "symmetric", $e_P((x,x'),(y,y')) = e(x,y')e(x',y)^{-1}$.

Isogenies and polarisations

- $f : A \to B$ isogeny.
- $v_1, \dots, v_m$ coordinates on B given by sections of $\mathscr{D}_B$.
- Then $u_i := v_i \circ f$ are coordinates on A given by sections of $\mathscr{D}_A := f^* \mathscr{D}_B$.
- $\deg \mathscr{D}_A = \deg f \cdot \deg \mathscr{D}_B$.

- $f : (A, \lambda_A) \to (B, \lambda_B)$ isogeny of ppavs.
- If λ_A is induced by Θ_A (resp. λ_B by Θ_B), a model of A (resp. B) will be given by coordinates of $m\Theta_A$ (resp. $m\Theta_B$), where $m = 2, 3, 4 \dots$ is small.
- We want to relate Θ_A with $f^* \Theta_B$ (or relate $m\Theta_A$ with $f^* m\Theta_B$).

N-isogenies

Definition

An isogeny $f : (A, \lambda_A) \to (B, \lambda_B)$ between ppav is an N-isogeny if $f^* \Theta_B \sim N \Theta_A$.

- $\Phi_{f^*\Theta_B}(P) = t_P^* f^* \Theta_B - f^* \Theta_B = f^*(t_{f(P)}^* \Theta_B - \Theta_B) = f^* \Phi_{\Theta_B}(f(P)) = (\hat{f} \circ \Phi_{\Theta_B} \circ f)(P)$;
- $f^* \lambda_B := \hat{f} \circ \lambda_B \circ f$;
- f is an N-isogeny $\Leftrightarrow f^* \lambda_B = N \lambda_A$;

$$\begin{array}{ccc} A & \xrightarrow{f} & B \\ \downarrow{\lambda_A} & & \downarrow{\lambda_B} \\ \widehat{A} & \xleftarrow[\hat{f}]{} & \hat{B} \end{array}$$

N-isogenies

Definition

An isogeny $f : (A, \lambda_A) \to (B, \lambda_B)$ between ppav is an N-isogeny if $f^* \Theta_B \sim N \Theta_A$.

- $\Phi_{f^*\Theta_B}(P) = t_P^* f^* \Theta_B - f^* \Theta_B = f^*(t_{f(P)}^* \Theta_B - \Theta_B) = f^* \Phi_{\Theta_B}(f(P)) = (\hat{f} \circ \Phi_{\Theta_B} \circ f)(P)$;
- $f^* \lambda_B := \hat{f} \circ \lambda_B \circ f$;
- f is an N-isogeny $\Leftrightarrow f^* \lambda_B = N \lambda_A$;
- Contragredient isogeny: $\tilde{f} = \lambda_A^{-1} \hat{f} \lambda_B : B \to A$;

$$\begin{array}{ccc} A & \xrightarrow{f} & B \\ \lambda_A^{-1} \uparrow & & \downarrow \lambda_B \\ \hat{A} & \xleftarrow[\hat{f}]{} & \hat{B} \end{array}$$

- f is an N-isogeny $\Leftrightarrow \tilde{f} f = N \Leftrightarrow f \tilde{f} = N$.

Example

An isogeny $f : E_1 \to E_2$ between elliptic curves is automatically an N-isogeny where $N = \deg f$.

N-isogenies and isotropic kernels

- Compatibility with pairings: on $T_\ell A \times T_\ell B$, $e_{\lambda_B}(f(x), y) = e_{\lambda_A}(x, \tilde{f}(y))$.
- $f : (A, \lambda_A) \to (B, \lambda_B)$ N-isogeny $\Rightarrow$ Ker f is maximal isotropic in $A[N]$ for the Weil pairing
- Ker $f = \text{Im}\, \tilde{f} \mid B[N]$, Ker f is dual to Ker $\tilde{f}$

- Conversely, if $K \subset A[N]$ maximal isotropic, $N\lambda_A$ descends to a principal polarisation on $B = A/K$.

  The pairing $e_{\lambda_A, N} = e_{\Phi_{N\lambda_A}}$ on $A[N] \times A[N]$ is also the commutator pairing of Mumford's theta group $G(N\Theta_A)$. If K is isotropic, it admits a lift $\widetilde{K}$ in $G(N\Theta_A)$, so $N\Theta_A$ descends to a divisor Θ_B on $B = A/K$. The degree relation shows that $\deg \Theta_B = 1$ if K is maximal.

- If $f : (A, \lambda_A) \to (B, \lambda_B)$ has maximal isotropic kernel in $A[N]$, $N\lambda_A$ descends to a principal polarisation λ'_B on B.
- But we may have $\lambda'_B \neq \lambda_B$.
- $\tilde{f} \circ f = N$ is a stronger condition that ensures compatibility of f with λ_B.
- f is an N-isogeny $\Leftrightarrow e_{\lambda_B}(f(x), f(y)) = e_{\lambda_A}(x, y)^N$ on $T_\ell A \times T_\ell A$.

Properties of contragredient isogenies

Biduality: $\tilde{\tilde{f}} = f$.

Composition: $f : A \to B$ a N-isogeny, $g : B \to C$ a M-isogeny, $g \circ f : A \to C$.

- $\widetilde{g \circ f} = \tilde{f} \circ \tilde{g} : C \to A$;
- $(\widetilde{g \circ f}) \circ (g \circ f) = \tilde{f} \circ \tilde{g} \circ g \circ f = NM$.
- The composition $g \circ f$ is an NM-isogeny.
- Conversely, if $g \circ f$ is an N-isogeny and f (resp. g) is an M-isogeny, then g (resp. f) is an N/M-isogeny.
- An N-isogeny is always the composition of ℓ_i-isogenies for $\ell_i \mid N$.

Product polarisation:

- $(A, \lambda_A) \times (B, \lambda_B) = (A \times B, \lambda_A \times \lambda_B)$ where $\lambda_A \times \lambda_B : A \times B \to \widehat{A} \times \hat{B}$ is the product.
- $F = \begin{pmatrix} a & c \\ b & d \end{pmatrix} : (A \times B, \lambda_A \times \lambda_B) \to (C \times D, \lambda_C \times \lambda_D)$.
- $\hat{F} = \begin{pmatrix} \hat{a} & \hat{b} \\ \hat{c} & \hat{d} \end{pmatrix} : \hat{C} \times \hat{D} \to \hat{A} \times \hat{B}$.
- $\tilde{F} = \begin{pmatrix} \tilde{a} & \tilde{b} \\ \tilde{c} & \tilde{d} \end{pmatrix} : C \times D \to A \times B$.
- Exercice: check that the 8×8-matrix at the beginning of the talk is a N'-isogeny.

Polarisations and symmetric endomorphisms

- (A, λ_A) ppav
- $\phi \in \mathrm{End}^{\lambda}(A) \mapsto \lambda_A \circ \phi$ induces a bijection between endomorphisms ϕ invariant under the Rosatti involution ($\widetilde{\phi} = \phi$) and polarisations: $NS(A) \simeq \mathrm{End}^{\lambda}(A)$.
- Let $\beta \in \mathrm{End}^{\lambda}(A)$, f is a β-isogeny if $\tilde{f}f = \beta$.
- If $f : A \to B$ is any isogeny, λ_A, λ_B principal polarisations, then f is a β-isogeny where $\beta = \tilde{f}f$. In particular $\mathrm{Ker}\, f$ is maximal isotropic for the e_β pairing on $A[\beta]$.

Example

- Via the product principal polarisation $(A \times B, \lambda_A \times \lambda_B)$, $F = \begin{pmatrix} a & c \\ b & d \end{pmatrix}$ is symmetric ($\tilde{F} = F$) iff $\tilde{a} = a, \tilde{d} = d, \tilde{b} = c$.
- $NS(A \times B) = NS(A) \times NS(B) \times \mathrm{Hom}(A, B)$.

- An ℓ-isogeny of abelian varieties has kernel of type $(\mathbb{Z}/\ell\mathbb{Z})^g$.
- An ℓ^2-isogeny of elliptic curves can have kernel of type $\mathbb{Z}/\ell^2\mathbb{Z}$ or $\mathbb{Z}/\ell\mathbb{Z} \times \mathbb{Z}/\ell\mathbb{Z}$.
- An ℓ^2-isogeny of abelian surfaces can have kernel of type $(\mathbb{Z}/\ell^2\mathbb{Z})^2$ or $\mathbb{Z}/\ell\mathbb{Z} \times \mathbb{Z}/\ell\mathbb{Z} \times \mathbb{Z}/\ell^2\mathbb{Z}$ or $(\mathbb{Z}/\ell\mathbb{Z})^4$.
- If an abelian surface (A, λ_A) has RM $\mathrm{End}^{\lambda_A}(A) = O_K$ a real quadratic order and $\ell = \beta\beta^c$, a β-isogeny will have cyclic kernel $\mathbb{Z}/\ell\mathbb{Z}$.

Outline

Algorithms for N-isogenies (overview)

- Input: generators $P_1, \dots, P_g$ of K, a maximal isotropic kernel for $A[N]$, a point $P \in A$ given by coordinates u_i, where u_i are sections of $m\Theta_A$.
- Output: A description of $B = A/K$, and the coordinates $v_i(Q)$ where $Q = f(P)$, where v_i are sections of $m\Theta_B$ (Θ_B a descent of $N\Theta_A$ by $f : A \to B$).

1. Construct $\mathcal{D} = f^* m\Theta_B$ on A.
   This is a divisor invariant by translation by K and algebraically equivalent to $Nm\Theta_A$. The converse is true by descent theory.
2. Construct the coordinates $v_i \circ f$ on A.
   These are sections of $\mathcal{D}$ invariant by translation on K, and the converse is true:
   $$\Gamma(B, m\Theta_B) \simeq \Gamma(A, f^* m\Theta_B)^K.$$
3. Evaluate these coordinates on P: $v_i(Q) = v_i \circ f(P)$.

Vélu's formula

- Weierstrass coordinates x, y on E = sections of $3(0_E)$. (x is a section of $2(0_E)$, y of $3(0_E)$.)
- K maximal isotropic in $E[N]$.
- $\mathcal{D} = \sum_{P \in K} t_P^*(3(0_E)) = \sum_{P \in K} 3(P)$ is certainly invariant by K;
- So $\mathcal{D}$ descends to $3(0_{E'})$ on $E' = E/K$;
- x, y are sections of $\mathcal{D}$ but are not invariant by translation;
- $X(P) = \sum_{T \in K} X(P+T)$ and $Y(P) = \sum_{T \in K} Y(P+T)$ are sections of $\mathcal{D}$ invariant by translation;
- They descend to Weierstrass coordinates on E';
- This is Vélu's formula (up to a constant).
- Cost: $O(N)$.
- Recover equations for E' via the formal group law.

Revisiting Vélu's formula

- Recall: $\mathscr{D} = \sum_{P \in K} t_P^* 3(0_E)$;
- We want to construct sections U of $\mathscr{D}$ that are of the form $U = v \circ f$, v a coordinate on E'.
- Equivalently: U is invariant by translation by K.
- In particular: $\operatorname{div} U$ is a divisor invariant by translation by K such that $\operatorname{div} U + \mathscr{D} \geq 0$.
- If $\mathscr{E} = \operatorname{div} f_{\mathscr{E}}$ is a principal divisor invariant by translation, $f_{\mathscr{E}}$ may not be invariant by translation!

Lemma

Let $\mathscr{E} = \sum_i a_i \sum_{T \in K}(P_i + T) = \operatorname{div} f_{\mathscr{E}}$ a principal divisor and $P_0 := \sum a_i P_i$. Then $f_{\mathscr{E}}$ is invariant by translation iff $P_0 \in K$.

Proof.

If $T \in K$, $f_{\mathscr{E}}(x+T)/f_{\mathscr{E}}(x) = e_f(T, f(P_0)) = e_N(T, P_0)$. So $f_{\mathscr{E}}$ is invariant by $K \Leftrightarrow P_0 \in E[\ell]$ is orthogonal to $K \Leftrightarrow P_0 \in K \Leftrightarrow f(P_0) = 0$. □

Revisiting Vélu's formula

- Recall: $\mathcal{D} = \sum_{P \in K} t_P^* 3(0_E)$;
- We want to construct sections U of $\mathcal{D}$ that are of the form $U = v \circ f$, v a coordinate on E'.
- Equivalently: U is invariant by translation by K.
- In particular: $\operatorname{div} U$ is a divisor invariant by translation by K such that $\operatorname{div} U + \mathcal{D} \geq 0$.
- If $\mathcal{E} = \operatorname{div} f_{\mathcal{E}}$ is a principal divisor invariant by translation, $f_{\mathcal{E}}$ may not be invariant by translation!

Lemma

Let $\mathcal{E} = \sum_i a_i \sum_{T \in K} (P_i + T) = \operatorname{div} f_{\mathcal{E}}$ a principal divisor and $P_0 := \sum a_i P_i$. Then $f_{\mathcal{E}}$ is invariant by translation iff $P_0 \in K$.

Example

- Take $Q_1, Q_2 \in E(k)$, $\mathcal{E} = \sum_{T \in K} ((Q_1 + T) + (-Q_1 + T) - (Q_2 + T) - (-Q_2 + T))$,
- $f_{\mathcal{E}} = \prod_{T \in K} \frac{x - x(Q_1 + T)}{x - x(Q_2 + T)}$ (convention: $x - 0_E := 1$).
- $f_{\mathcal{E}}$ is invariant by translation and descends to $\frac{X - f(Q_1)}{X - f(Q_2)}$ on E/K, X a Weierstrass coordinate.
- When $Q_2 = 0_E$, we recover formula from [Costello-Hisil, 2017], [Renes, 2017].
- Used by the `sqrtVelu` algorithm!

Vélu's formula in higher dimension?

- (A, Θ_A) ppav, K maximal isotropic in $A[N]$
- $\mathscr{D} = \sum_{P \in K} t_P^*(m\Theta_A)$ is certainly invariant by K;
- If u is a section of $m\Theta_A$, $U(P) = \sum_{T \in K} u(P+T)$ is certainly a section of $\mathscr{D}$ invariant by K.

- But $\mathscr{D} \sim N^g m\Theta_A$;
- So it descends to a divisor $\sim N^{g-1} m\Theta_B$!
- Our coordinates have degree too big (unless $g = 1$).

The theta group

- $Nm\Theta_A$ is not invariant by K
- So it does not descend to $m\Theta_B$
- But it is linearly equivalent to $\mathscr{D}$, a divisor invariant by K: $\mathscr{D} = Nm\Theta_A + \operatorname{div} g$.
- So $\operatorname{div}(g/t_T^* g) = t_T^* Nm\Theta_A - Nm\Theta_A$.
- Goal: construct $\mathscr{D}$. Equivalently construct g.

- Find functions g_T such that $\operatorname{div} g_T = t_T^* Nm\Theta_A - Nm\Theta_A$
- Try to glue these functions into a global function g (cocycle condition): $g_T(P) = g(P)/g(P+T)$.

- Theta group: $G(Nm\Theta_A) = \{(T, g_T) \mid \operatorname{div} g_T = t_T^* Nm\Theta_A - Nm\Theta_A\}$
- Gluing condition $\Leftrightarrow K \to G(Nm\Theta_A), T \mapsto (T, g_T)$ is a group section;

- Twisted trace: if U is a section of $Nm\Theta_A$, $U'(P) = \sum_{T \in K} g_T(P) U(P+T)$ is a section of $\mathscr{D}$ invariant by K, hence descends to $B = A/K$.

General framework for an N-isogeny algorithm

1. Find functions g_T, $\operatorname{div} g_T = t_T^* Nm\Theta_A - Nm\Theta_A$ for each $T \in K$, that glue together.
   1. Use symmetry: Θ_A symmetric divisor, g_T symmetric.
   2. Unique choice if N is odd, two choices for each T when N is even $\Rightarrow$ annoying!

Twisted Vélu's formula: if $K = \langle T \rangle$, $X(P) = \sum_{i \in \mathbb{Z}/N\mathbb{Z}} \zeta_N^i X(P+T)$, $Y(P) = \sum_{i \in \mathbb{Z}/N\mathbb{Z}} \zeta_N^i Y(P+T)$.

Eg: if N is even, $X(P) = \sum_{i \in \mathbb{Z}/N\mathbb{Z}} (-1)^i X(P+T)$ descends to a section on the symmetric divisor $2f(W)$, $W \in E[2] - K$.

General framework for an N-isogeny algorithm

1. Find functions g_T, $\operatorname{div} g_T = t_T^* Nm\Theta_A - Nm\Theta_A$ for each $T \in K$, that glue together.
2. Generate sections U of $Nm\Theta_A$.
   1. The multiplication map $\Gamma(m_1\Theta_A) \otimes \Gamma(m_2\Theta_A) \to \Gamma((m_1+m_2)\Theta_A), u \otimes v \mapsto uv$ is surjective if $m_1 \geq 3, m_2 \geq 2$ [Mumford, Koizumi, Kempf].
   2. $\Sigma_{\alpha \in \widehat{A}} \Gamma(A, m_1\Theta_A \otimes P_\alpha)\Gamma(A, m_2\Theta_A \otimes P_{-\alpha}) = \Gamma(A,(m_1+m_2)\Theta_A)$ [Mumford] for $m_1, m_2 > 0$.

   So we can always generate all sections of $\Gamma(Nm\Theta_A)$ using multiplications of sections of $\Gamma(m\Theta_A)$, eventually using also translations if $m \leq 2$.

General framework for an N-isogeny algorithm

1. Find functions g_T, $\operatorname{div} g_T = t_T^* Nm\Theta_A - Nm\Theta_A$ for each $T \in K$, that glue together.
2. Generate sections U of $Nm\Theta_A$.
3. Take the twisted traces of the sections U.
4. This gives coordinates (section of $m\Theta_B$) on B

- More work required to recover a suitable model of B (depends on the model).

General framework for an N-isogeny algorithm

1. Find functions g_T, $\operatorname{div} g_T = t_T^* Nm\Theta_A - Nm\Theta_A$ for each $T \in K$, that glue together.
2. Generate sections U of $Nm\Theta_A$.
3. Take the twisted traces of the sections U.
4. This gives coordinates (section of $m\Theta_B$) on B

- More work required to recover a suitable model of B (depends on the model).

- Summary [R. 2021]: from an effective version of the Theorem of the square:

$$t_{P+Q}^*\Theta_A + \Theta_A - t_P^*\Theta_A - t_Q^*\Theta_A = \operatorname{div} \mu_{P,Q},$$

  there is a general framework to

  1. Compute the addition law;
  2. Compute the Weil and Tate pairings;
  3. Compute isogenies.

Isogenies in the theta model

- Analytic theta functions:

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

- Universal
- Work with theta functions of level $m=2$ or $m=4$: m^g coordinates.
- Rationality: rational $\Gamma(m,2m)$-symplectic structure.
- N-isogenies in $O(N^g)$.
- Implementations in Magma (AVIsogenies) and Sage (ThetAV)

- General framework for β-isogenies but requires bootstrapping (still more work needed!).
- Theta functions $\theta_{A\times B}$ for the product theta structure on $A \times B$ are simply product of theta functions $\theta_A \cdot \theta_B$.
- $\begin{pmatrix} N_1 & 0 \\ 0 & N_2 \end{pmatrix}$-isogenies in $O(N_1^g N_2^g)$.

- Moduli: $\chi(\tau) = \prod \theta\left[\begin{smallmatrix} a/2 \\ b/2 \end{smallmatrix}\right](\tau)$ describes interesting modular locus: the locus of product of elliptic curves when $g=2$ (χ_{10}), the locus of products and Jacobians of hyperelliptic curves when $g=3$ (χ_{18}).
  The modular form $g(A,w_A) = \prod_{(B,w_B)} \chi_{10}(B,w_B)$ of weight $10(\ell^3+\ell^2+\ell+1)$ (whose product is across all normalised ℓ-isogenies) describes the locus H_{ℓ^2} of ℓ-split abelian surfaces (the Humbert surface of discriminant ℓ^2). Expressed as a polynomial P in terms of $\psi_4, \psi_6, \chi_{10}, \chi_{12}$, P is of size $\widetilde{O}(\ell^{12})$ and can be computed in quasi-linear time by evaluation-interpolation. Checking if $(A,\Theta_A)/\mathbb{F}_q$ is ℓ-split can then be done by evaluating $P(A,\Theta_A)$ in time $O(\ell^9 \log q)$, or directly via the analytic method in $\widetilde{O}(\ell^3(\log q + d^2))$.

Isogenies in the Jacobian model

- $\iota : C \to \mathrm{Jac}(C)$;
- If g is a function on C, it induces a function $\iota_* g$ on $\mathrm{Jac}(C)$ via $(\iota_* g)(\sum n_i (P_i)) = \prod g(P_i)^{n_i}$.
- All functions on $\mathrm{Jac}(C)$ can be built from $\iota_* g$ and determinants;
- NB: for pairings computations, the functions $\iota_* g$ are enough!
- N-isogenies between Jacobians in $\widetilde{O}(N^g)$ when $g = 2$ [Couveignes-Ezome 2015] and $g = 3$ [Milio 2019]
- Implementations in Magma.
- The extension to product of Jacobians should not be too hard.

Algorithms for isogenies

- Better algorithms for β-isogenies;
- $\widetilde{O}(N^{g/2})$-algorithms?
- Batch isogeny evaluation?
- More compact models of abelian varieties?

- Evaluating an isogeny on a point is only a small topic of algorithms related to isogenies: modular polynomials, explicit Kodaira-Spencer isomorphism, differential equations, …