Towards hyperelliptic isogeny-based cryptography

Benjamin Smith

Isogeny school, online, 2021

Inria + École polytechnique, France

Hyperelliptic Jacobians

Hyperelliptic Jacobians

Suppose $\mathcal{X} : y^2 = f(x)$ is hyperelliptic of genus $g > 1$.

From now on, we suppose f has odd degree,
so $\mathcal{X}$ has a single point ∞ at infinity.

Even degree case is (only) slightly more complicated.

Goal: to define a compact (and algebraic) representation for $\mathrm{Pic}^0(\mathcal{X})$.

Reduced representatives for classes

If $[D]$ is in $\mathrm{Pic}^0(\mathcal{X})$, then $[D]$ has a unique **reduced representative**:

$$[D] = [P_1 + \cdots + P_r - r\infty]$$

for some $P_1, \ldots, P_r \in \mathcal{X}$ depending on $[D]$ (not D) such that

- $P_i \neq \infty$ and $P_i \neq \iota(P_j)$ for $i \neq j$ (semi-reducedness)
- $r \leq g$ (reducedness)

$$[D] \in \mathrm{Pic}^0(\mathcal{X})(\Bbbk) \iff P_1 + \cdots + P_r \in \mathrm{Div}(\mathcal{X})(\Bbbk)$$

Note: the individual P_i need not be in $\mathcal{X}(\Bbbk)$!

Why?

Riemann–Roch guarantees the existence of the reduced representative.

If $[D]$ is in $\mathrm{Pic}^0(\mathcal{X})$, then applying Riemann–Roch to $D + g\infty$ yields a function f such that $D + g\infty + \mathrm{div}(f) = D'$ is effective; so $[D' - g\infty] = [D]$ with $\deg D' = g$.

$D' - g\infty$ is *almost* a reduced representative:
it remains to remove any $P + \iota(P) - 2\infty = \mathrm{div}(x - x(P))$ from D'.

The Mumford representation

Suppose we have a class $[D]$ in $\mathrm{Pic}^0(\mathcal{X})(\Bbbk)$ with reduced representative

$$D = P_1 + \cdots + P_r - r\infty \in \mathrm{Div}^0(\mathcal{X})(\Bbbk) \ .$$

The **Mumford representation** of $[D]$ is the (unique) pair of polynomials $\langle a(x), b(x)\rangle$ in $\Bbbk[x]$ such that

- $a(x) = \prod_{i=1}^{r}(x - x(P_i))$, and
- $b(x(P_i)) = y(P_i)$ for $1 \leq i \leq r$;

for each of the x-coordinates appearing as a root of a, evaluating b gives the corresponding y-coordinate.

If necessary, compute b by Lagrange interpolation.

The Mumford representation

If $\langle a(x), b(x)\rangle$ represents a class on $\mathcal{X} : y^2 = f(x)$, then

1. a is monic of degree $r \leq g$, and
2. b satisfies $\deg b < r$ and $b^2 \equiv f \pmod{a}$.

Theorem: Any pair $\langle a(x), b(x)\rangle$ in $\mathbb{k}[x]^2$ satisfying these conditions represents a divisor class in $\mathrm{Pic}^0(\mathcal{X})(\mathbb{k})$.

$\Longrightarrow$ identify divisor classes with Mumford reps of their reduced representatives: we simply write $[D] = \langle a, b\rangle$.

We associate $\langle a(x), b(x)\rangle$ *with the ideal* $(a(x), y - b(x))$.

Hyperelliptic Jacobians

We can collect the Mumford representations by degree $0 \leq d \leq g$:

$$M_d := \{\langle a, b\rangle : \deg(b) < \deg(a) = d, b^2 \equiv f \pmod{a}\} .$$

We can view the coefficients of $a(x) = x^d + a_{d-1}x^{d-1} + \cdots + a_1 x + a_0$ and $b(x) = b_{d-1}x^{d-1} + \cdots + b_0$ as coordinates on $\mathbb{A}^{2d}$.

$b^2 \pmod{a}$ and $f \pmod{a}$ are polynomials of degree $d-1$ in $\Bbbk[a_i, b_i][x]$; the vanishing of their coefficients defines d independent equations in the $2d$ coordinates, cutting out M_d as a d-dimensional subvariety in $\mathbb{A}^{2d}$.

- M_0 is a point;
- M_1 is an affine copy of $\mathcal{X}$;
- $\#M_d(\mathbb{F}_q) = O(q^d)$ for $0 \leq d \leq g$.

The Jacobian

Glueing together $M_0, \ldots, M_g$, we give $\mathrm{Pic}^0(\mathcal{X})$ the structure of a g-dimensional algebraic variety $\mathcal{J}_\mathcal{X}$, called the **Jacobian**.

Over $\mathbb{F}_q$, we have $\#\mathcal{J}_\mathcal{X} = O(q^g)$ *(more precision later).*

We want an expression of the group law on $\mathcal{J}_\mathcal{X}$ in terms of its coordinates; Cantor's algorithm does this using an explicit form of Riemann–Roch.

Cantor's algorithm: addition on $\mathcal{J}_{\mathcal{X}}$

Input: Reduced divisors $D_1 = \langle a_1, b_1 \rangle$ and $D_2 = \langle a_2, b_2 \rangle$ on $\mathcal{X}$.
Output: A reduced $D_3 = \langle a_3, b_3 \rangle$ s.t. $[D_3] = [D_1 + D_2]$ in $\mathrm{Pic}^0(\mathcal{X})$.

1. $(d, u_1, u_2, u_3) := \mathrm{XGCD}(a_1, a_2, b_1 + b_2)$
   // so $d = \mathbf{gcd}(a_1, a_2, b_1 + b_2) = u_1 a_1 + u_2 a_2 + u_3(b_1 + b_2)$.
2. Set $a_3 := a_1 a_2 / d^2$;
3. Set $b_3 := b_1 + (u_1 a_1 (b_2 - b_1) + u_3 (f - b_1^2))/d \pmod{a_3}$;
4. If $\mathbf{deg}\, a_3 \leq g$ then go to Step 9;
5. Set $\tilde{a}_3 := a_3$ and $\tilde{b}_3 := b_3$;
6. Set $a_3 := (f - b_3^2)/a_3$;
7. Let $(Q, b_3) := \mathrm{Quotrem}(-b_3, a_3)$;
8. While $\mathbf{deg}\, a_3 > g$
   - 8a Set $t := \tilde{a}_3 + Q(b_3 - \tilde{b}_3)$;
   - 8b Set $\tilde{b}_3 := b_3$, $\tilde{a}_3 = a_3$, and $a_3 := t$;
   - 8c Let $(Q, b_3) := \mathrm{Quotrem}(-b_3, a_3)$;
9. Return $\langle a_3, b_3 \rangle$.

- Step 1: $d(x(P_i)) = 0$ iff $P_i = \iota(Q_j)$ for some j
- Steps 2, 3: sum D_1 and D_2, remove contribution of d
  $\longrightarrow$ pre-reduced D_3 such that $[D_3] = [D_1 + D_2]$
- Loop: reduces degree of the representative until reduced.
- Exercise: *how many steps until the result is reduced?*

Jacobian arithmetic

Cantor's algorithm is useful for general hyperelliptic curves, but there are generally better options for cryptographers, especially in genus 2.

- Optimized straight-line programs: `https://ia.cr/2012/670`
- Fast Kummer arithmetic (see below)

Constructive pre-quantum cryptographic work in $g > 2$ has largely stopped in the last 15 years: index calculus makes DLP hardness sub-optimal in higher genus.

Embeddings of Jacobians

The Mumford representation lets us compute with a hyperelliptic Jacobian by dividing it up into affine pieces:

$$\mathcal{J}_{\mathcal{X}} = M_0 \cup M_1 \cup \cdots \cup M_g \ .$$

In fact, $\mathcal{J}_{\mathcal{X}}$ is projective (it's an *abelian variety*)
—so what are its projective embeddings?

This is a nontrivial question

$$\mathcal{J}_{\mathcal{X}} = M_0 \cup M_1 \cup \cdots \cup M_g \quad \text{with each } M_i \subset \mathbb{A}^{2i}$$

recalls the usual decomposition $\mathbb{P}^n = \mathbb{A}^0 \cup \mathbb{A}^1 \cup \cdots \cup \mathbb{A}^n$ —but it's not the same!

As cryptographers, we are used to thinking of projective coordinates as nothing more than convenient denominator elimination. To go from affine to projective, we just homogenize with respect to a new variable.

But if you just homogenize Mumford representations, then you get something totally wrong.

The Jacobi intersection model

To create projective embeddings of curves, we used divisors and Riemann–Roch.

For example: given a point $\mathcal{O}$ on an elliptic $\mathcal{E}$, we embed $\mathcal{E}$ in $\mathbb{P}^2 = \mathbb{P}(L(3\mathcal{O})) = \mathbb{P}(\langle x, y, 1\rangle)$.

Alternative embeddings: for example, use $D = 4\mathcal{O}$.

- $L(4\mathcal{O}) = \langle x, y, u, v\rangle$ (because $\dim L(4\mathcal{O}) = \deg(4\mathcal{O}) = 4$);
- $L(8\mathcal{O}) \supseteq L(4\mathcal{O})^2 = \left\langle x^2, xy, xu, xv, y^2, yu, yv, u^2, uv, v^2\right\rangle$
- but $\dim L(8\mathcal{O}) = 8 \implies$ 2 quadratic relations in x, y, u, v.
- $\implies$ the *Jacobi intersection model* of $\mathcal{E}$:

$$\mathcal{E} : F_2(x, y, u, v) = G_2(x, y, u, v) = 0 \quad \subset \mathbb{P}^3 = \mathbb{P}(L(4\mathcal{O})) \,.$$

Theta

So, if $\mathcal{E}$ is an elliptic curve and $\mathcal{O}$ is a point on $\mathcal{E}$, then:

- $L(2\mathcal{O})$ gives a double cover $\mathcal{E} \to \mathbb{P}^1$ (the x-line)
- $L(3\mathcal{O})$ embeds $\mathcal{E}$ in $\mathbb{P}^2$ with one cubic equation;
- $L(4\mathcal{O})$ embeds $\mathcal{E}$ in $\mathbb{P}^3$ with two quadratic equations.

What are the hyperelliptic analogues?

We need a divisor on $\mathcal{J}_\mathcal{X}$ to take the place of $\mathcal{O}$ on $\mathcal{E}$:

$$\Theta := \{[P_1 + \cdots + P_{g-1} - (g-1)\infty] : P_1, \ldots, P_{g-1} \in \mathcal{X}(\overline{\Bbbk})\}$$

(Note: $\Theta = M_0 \cup \cdots \cup M_{g-1}$).

Projective embeddings of $\mathcal{J}_{\mathcal{X}}$

The theta divisor is

$$\Theta := \{[P_1 + \cdots + P_{g-1} - (g-1)\infty] : P_i \in \mathcal{X}(\overline{\mathbb{k}})\} \subset \mathcal{J}_{\mathcal{X}}\,.$$

We have $\dim L(n\Theta) = n^g$, so

- $L(2\Theta)$ gives a double cover of a Kummer variety in $\mathbb{P}^{2^g-1}$
- $L(3\Theta)$ embeds $\mathcal{J}_{\mathcal{X}}$ in $\mathbb{P}^{3^g-1}$
- $L(4\Theta)$ embeds $\mathcal{J}_{\mathcal{X}}$ in $\mathbb{P}^{4^g-1}$.

The dimension of the space is exponential in g
(and so is the number of equations!)

Generally, $\mathcal{J}_{\mathcal{X}}$ does not embed in a smaller projective space than $\mathbb{P}^{3^g-1}$.

Projective embeddings of $\mathcal{J}_\mathcal{X}$ for $g = 2$

For $g = 2$: $\mathcal{J}_\mathcal{X}$ is a surface, and Θ is a copy of $\mathcal{X}$ inside $\mathcal{J}_\mathcal{X}$.

- $L(2\Theta)$ gives a double cover of the quartic **Kummer surface** in $\mathbb{P}^3$.
- $L(3\Theta)$ gives the "Grant" embedding in $\mathbb{P}^8$
  with 10 quadratic and 3 cubic equations.
- $L(4\Theta)$ gives the "Flynn" embedding in $\mathbb{P}^{15}$
  with 72 quadratic equations.
- $\mathcal{J}_\mathcal{X}$ *never* embeds in $\mathbb{P}^3$.
- $\mathcal{J}_\mathcal{X}$ embeds in $\mathbb{P}^4$ if and only if $\mathrm{End}(\mathcal{J}_\mathcal{X})$ contains $\mathbb{Z}[(1+\sqrt{5})/2]$
  (...Horrocks–Mumford bundle, etc.)

Kummer surfaces

Kummer varieties

If $\mathcal{A}$ is an abelian variety (e.g. an elliptic curve, or the Jacobian of a hyperelliptic curve), then the negation map $\ominus$ is an involution of $\mathcal{A}$.

We can take the quotient under the action of this involution. The quotient map $\mathcal{A} \to \mathcal{A}/\langle\ominus\rangle = \mathcal{A}/\langle\pm 1\rangle$ is a double cover of the **Kummer variety** of $\mathcal{A}$.

Example: Consider an elliptic curve $\mathcal{E} : y^2 = x^3 + ax + b$.
The negation map is $\ominus(x, y) = (x, -y)$, so $P \mapsto x(P)$ is the quotient by ± 1.
$\implies \mathbb{P}^1$ is the Kummer variety (in this case, *Kummer line*) of $\mathcal{E}$.

Kummer lines give the mathematical foundation for **Montgomery** (x-only) arithmetic on elliptic curves.

Kummer surfaces

Genus 2 analogue: The **Kummer surface** $\mathcal{K}_{\mathcal{X}} := \mathcal{J}_{\mathcal{X}}/\langle \pm 1 \rangle$ is a quartic surface in $\mathbb{P}^3$.

It has 16 point singularities, which are the images of the 16 points in $\mathcal{J}_{\mathcal{X}}[2]$.

From elliptic to genus-2 Jacobians:

- $\mathbb{P}^1$ is the image of $\mathcal{E}$ in $\mathbb{P}(L(2\mathcal{O}))$,
- $\mathcal{K}_{\mathcal{X}}$ is the image of $\mathcal{J}_{\mathcal{X}}$ in $\mathbb{P}(L(2\Theta))$.

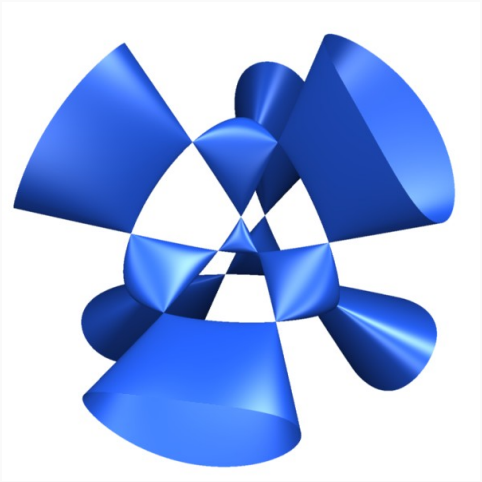

What are Kummers useful for?

We can use Kummer surfaces for a genus-2 analogue of Montgomery arithmetic.

$\mathcal{K}_{\mathcal{X}}$ is not a group, but it inherits scalar multiplication from $\mathcal{J}_{\mathcal{X}}$ *(since* $[m](\pm D) = \pm([m]D)$*).*

$\mathcal{K}_{\mathcal{X}}$ is therefore suitable for

- cryptosystems that only use scalar multiplication (like Diffie–Hellman),
- or where “sign” is not important (like much of isogeny-based crypto).

In any case: working with $\mathcal{K}_{\mathcal{X}}$, defined by one equation in $\mathbb{P}^3$, is generally much more convenient than working with $\mathcal{J}_{\mathcal{X}}$, defined by many more equations in many more variables.

Kummer surfaces

Kummer surface arithmetic: Chudnovsky & Chudnovsky, Gaudry, Robert, Cosset, ...

The "squared" Kummer surface gives highly optimized formulae:

$$4E^2 \cdot X_1X_2X_3X_4 = \left(\begin{array}{c} X_1^2 + X_2^2 + X_3^2 + X_4^2 - F(X_1X_4 + X_2X_3) \\ - G(X_1X_3 + X_2X_4) - H(X_1X_2 + X_3X_4) \end{array} \right)^2,$$

with E, F, G, H in $\mathbb{k}$ derived from **theta constants** a, b, c, d in $\mathbb{k}$.

For DLP-based systems, Kummer arithmetic is *faster* than elliptic x-line arithmetic at the same security level.

Eg. 128-bit DLP security: $\mathcal{K}_{\mathcal{X}}$ over 128-bit field beats $\mathcal{E}/\langle \pm 1 \rangle$ over 256-bit field.

Details, and an example of this in practice: `https://ia.cr/2017/518`

Doubling on the Kummer surface

The pseudo-doubling operation on $\mathcal{K}_{\mathcal{X}}$ is

$$\pm P = \left(X_1^P : X_2^P : X_3^P : X_4^P\right) \longmapsto \left(X_1^{[2]P} : X_2^{[2]P} : X_3^{[2]P} : X_4^{[2]P}\right) = \pm[2]P$$

where

$$\begin{aligned}
X_1^{[2]P} &= \epsilon_1(U_1 + U_2 + U_3 + U_4)^2\,, & U_1 &= \hat{\epsilon}_1(X_1^P + X_2^P + X_3^P + X_4^P)^2\,,\\
X_2^{[2]P} &= \epsilon_2(U_1 + U_2 - U_3 - U_4)^2\,, & U_2 &= \hat{\epsilon}_2(X_1^P + X_2^P - X_3^P - X_4^P)^2\,,\\
X_3^{[2]P} &= \epsilon_3(U_1 - U_2 + U_3 - U_4)^2\,, & U_3 &= \hat{\epsilon}_3(X_1^P - X_2^P + X_3^P - X_4^P)^2\,,\\
X_4^{[2]P} &= \epsilon_4(U_1 - U_2 - U_3 + U_4)^2\,, & U_4 &= \hat{\epsilon}_4(X_1^P - X_2^P - X_3^P + X_4^P)^2
\end{aligned}$$

for $\pm P$ with all $X_i^P \neq 0$. The ϵ_i and $\hat{\epsilon}_i$ are derived from the same theta constants.

Abelian varieties and isogenies

Abelian varieties

An **abelian variety** is an algebraic group that can be embedded in some projective space.

When an abelian variety is properly embedded in an actual projective space—that is, once we have put coordinates on it—we say that it is a **principally polarized abelian variety (PPAV)**.

A polarization on an abelian variety is analogous to the distinguished point ("at infinity") on an elliptic curve: as we saw yesterday, that choice of point gives you embeddings in projective spaces via Riemann–Roch spaces.

We need to polarize abelian varieties to compute with them, so from our point of view, PPAVs are the "correct" higher-dimensional generalization of elliptic curves.

Moduli dimensions

Recall: The j-invariant classifies elliptic curves up to $\overline{\mathbb{k}}$-isomorphism, so there are $O(q)$ non-isomorphic elliptic curves over $\mathbb{F}_q$.

For higher-genus curves, and higher-dimensional abelian varieties, there are more complicated **moduli spaces** (systems of invariants).

Objects	Moduli	isom. classes $/\mathbb{F}_q$
Hyperelliptic curves/Jacobians of genus g	$2g-1$	$O(q^{2g-1})$
General curves/Jacobians of genus $g>1$	$3g-3$	$O(q^{3g-3})$
PPAVs of dimension g	$g(g+1)/2$	$O(q^{\frac{1}{2}g(g+1)})$

Note: Torelli's theorem tells us $\mathcal{J}_{\mathcal{X}_1} \cong \mathcal{J}_{\mathcal{X}_2} \iff \mathcal{X}_1 \cong \mathcal{X}_2$.

Kinds of abelian varieties

In low dimensions, we can avoid heavy PPAV machinery by working with curves:

- $g = 1$ Every 1-dimensional PPAV is an elliptic curve.
- $g = 2$ Every 2-dimensional PPAV is isomorphic to:
  - A product of 2 elliptic curves, or
  - The Jacobian of a genus-2 curve (necessarily hyperelliptic).
- $g = 3$ Every 3-dimensional PPAV is isomorphic to:
  - A product of 3 elliptic curves, or
  - The product of an elliptic curve with a genus-2 Jacobian, or
  - The Jacobian of a genus-3 hyperelliptic curve, or
  - The Jacobian of a non-hyperelliptic genus-3 curve.

General g-dimensional PPAVs are not isomorphic to Jacobians for every $g > 3$.

Cardinalities

Let $\mathcal{A}$ be a g-dimensional abelian variety over $\mathbb{F}_q$.

The **Weil polynomial** (characteristic poly. of Frobenius) has the form

$$\chi_{\mathcal{A}}(T) = T^{2g} + a_1 T^{2g-1} + \cdots + a_g T^g + a_{g-1} q T^{g-1} + \cdots + a_1 q^{g-1} T + q^g \,.$$

The **cardinality** of $\mathcal{A}(\mathbb{F}_q)$ is

$$(q^{1/2} - 1)^{2g} \leq \#\mathcal{A}(\mathbb{F}_q) \leq (q^{1/2} + 1)^{2g} \,.$$

Computing the cardinality can be very hard (unless p is very small and $\mathcal{A}$ is a Jacobian; or $\mathcal{A}$ is a product of elliptic curves; or $\mathcal{A}$ is supersingular).

Even in dimension $g = 2$, large-characteristic point counting is hard in practice!

Torsion and group structures

If $\mathcal{A}$ is a g-dimensional abelian variety over $\mathbb{F}_q$, then

- $\mathcal{A}[\ell^k](\overline{\mathbb{F}}_q) \cong (\mathbb{Z}/\ell^k\mathbb{Z})^{2g}$ when $\ell \neq p$,
- and $\mathcal{A}[p^k](\overline{\mathbb{F}}_q) \cong (\mathbb{Z}/p^k\mathbb{Z})^r$ for $0 \leq r \leq g$ (independent of k)

The integer r is called the p-rank.

Hence,

$$\mathcal{A}(\mathbb{F}_q) \cong (\mathbb{Z}/d_1\mathbb{Z}) \times (\mathbb{Z}/d_2\mathbb{Z}) \times \cdots \times (\mathbb{Z}/d_{2g}\mathbb{Z})$$

with each $d_{i+1} \mid d_i$ ($d_i = 1$ is allowed).

Endomorphism rings

Endomorphism rings of PPAVs over $\mathbb{F}_q$ are more complicated than their elliptic counterparts.

General PPAVs are ordinary and not reducible: $\mathrm{End}(\mathcal{A}) \cong$ an order in a **CM-field** (a quadratic imaginary extension of a totally real field of degree g).
This is somewhat analogous to the elliptic CM case, with an order in the degree-g real field replacing $\mathbb{Z}$; but it can get complicated, especially when that real order has class number > 1.

Products of PPAVs:

$$\mathrm{End}(\mathcal{A}_1 \times \mathcal{A}_2) \cong \mathrm{End}(\mathcal{A}_1) \times \mathrm{End}(\mathcal{A}_2) \times \mathrm{Hom}(\mathcal{A}_1, \mathcal{A}_2) \times \mathrm{Hom}(\mathcal{A}_2, \mathcal{A}_1)\,.$$

Supersingular PPAVs are isogenous to powers of supersingular elliptic curves, so we can think of matrices of quaternions.

Concrete endomorphisms

What about computing with explicit, concrete endomorphisms?

In general, the situation is similar to the elliptic case: we can work efficiently with elements of $\mathbb{Z}[\pi]$, where π is the Frobenius endomorphism.

Computing the endomorphism ring of a higher-dimensional abelian variety is fundamentally harder than computing elliptic endomorphism rings: even the basic point-counting algorithms which we need to understand $\mathbb{Z}[\pi] \subset \mathrm{End}(\mathcal{A})$ have much higher complexity.

However, we *do* know some special families of hyperelliptic Jacobians equipped with efficiently computable **real multiplications**: that is, fixed low-degree non-integer elements of the real subring of $\mathrm{End}(\mathcal{A})$. These maps are realised in terms of **correspondences**, which are algebraic maps of divisor classes.

Isogenies of abelian varieties

An **isogeny of abelian varieties** is a morphism that:

- is **finite**,
- is (geometrically) **surjective**, and
- **maps 0 to 0**.

Basically: dimension-preserving morphisms mapping 0 to 0.

All isogenies are homomorphisms.

Isogenies of PPAVs

An **isogeny of PPAVs** is more subtle: generally, we are looking at kernels that are maximal isotropic subgroups with respect to the Weil pairing.

Example: the genus-2 analogue of an elliptic 2-isogeny is a $(2,2)$**-isogeny**: the kernel is a subgroup $S \subset \mathcal{A}[2]$ such that $S \cong (\mathbb{Z}/2\mathbb{Z})^2$ and $e_2(S,S) = 1$.

Why bother? We need to respect polarizations if we want to realize the isogeny as a projective map, or to have an isogeny of Jacobians...

Ordinary vs supersingular

Ordinary PPAVs:

- hard to compute their cardinalities / hard to derive parameters
- more complicated endomorphism rings and class group structures than in the elliptic case (see e.g. **Ionica–Thomé** `https://ia.cr/2014/230`)
- high-degree isogeny-computation is feasible, if not especially fast (see e.g. **Lubicz–Robert** `https://arxiv.org/abs/1001.2016`)
- *But don't let this put you off! More research is needed.*

Supersingular vs superspecial

Up to this point, higher-dimensional isogeny-based crypto has focused almost exclusively on the **supersingular** case with $(2,2)$- and $(3,3)$-isogenies.

We often (even just implicitly) work with the subgraph of **superspecial** abelian surfaces.

Supersingular: *isogenous* to a product of supersingular elliptic curves

Superspecial: supersingular and *isomorphic as an unpolarized PPAV* to a product of supersingular elliptic curves

The distinction is subtle, but if you start superspecial and use only separable isogenies (no p-isogenies), then you stay superspecial.

The neighbourhood of a general vertex in the superspecial $(2, 2)$-isogeny graph

Recall: the elliptic supersingular 2-isogeny graph is a 3-regular* expander.

The genus-2 superspecial $(2, 2)$-isogeny graph is 15-regular, so the neighbourhood of a typical vertex looks like this diagram on the right.

But don't think of the graph as being a 15-regular version of the elliptic graph!
Looking at the neighbourhood of a typical edge, we see many interconnections between neighbours...

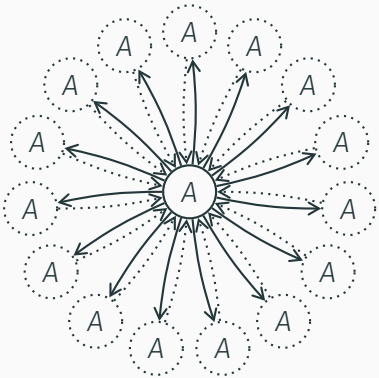

The neighbourhood of a general edge and its dual

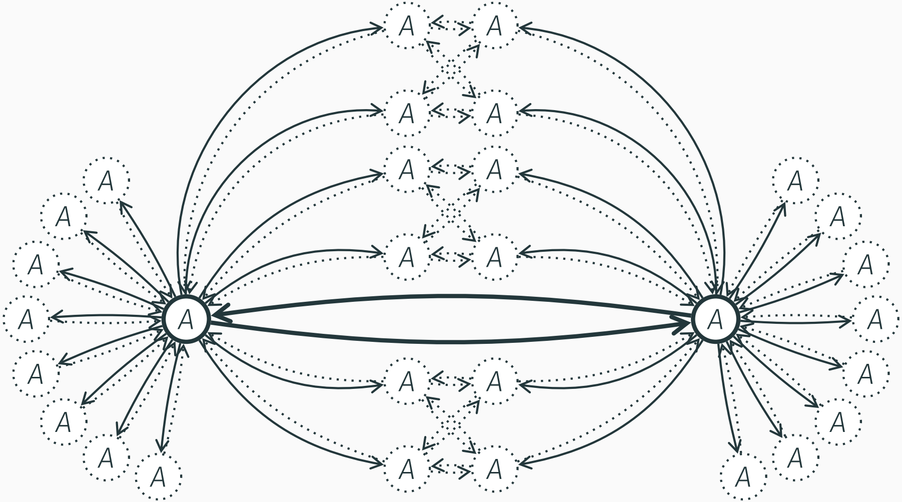

Genus-2 isogeny-based cryptography, as things stand

Problems and resources

Hash functions

Generalizations of the Charles–Goren–Lauter hash:

- **Takashima**: *Efficient Algorithms for Isogeny Sequences and Their Cryptographic Applications* (2018)
  Based on the Richelot $(2,2)$-isogeny graph of supersingular abelian surfaces
- **Castryck–Decru–Smith** `https://ia.cr/2019/296`:
  Refinement of the Takashima hash to the superspecial graph

SIDH analogues

There has been some preliminary work on genus-2 SIDH techniques:

- Costello: `https://ia.cr/2018/850`
  Using genus-2 methods to implement elliptic SIDH(!), via the Weil restriction
- Flynn–Ti: `https://ia.cr/2019/177`
  A genus-2 SIDH proposal:
  - $(2, 2)$-isogenies computed with Richelot formulae
  - $(3, 3)$-isogenies on the Kummer surface

Plenty of room for further algorithmic work here!

A closer look at the genus-2 superspecial graph

Understanding the properties of the genus-2 superspecial isogeny graph:

- Jordan–Zaytman `https://arxiv.org/abs/2005.09031`:
  The superspecial isogeny graph is connected for all g
- Florit–Smith `https://ia.cr/2021/012`:
  Statistical properties and structure for $(2,2)$-isogeny graphs
- Florit–Smith "Atlas" `https://ia.cr/2021/013`:
  Diagrams of local structures and pathological neighbourhoods/subgraphs of $(2,2)$-isogeny graphs

We still have a much less complete picture here than we do for the elliptic supersingular graph.

Attacking the higher-genus isogeny problem

The general **isogeny problem**:

Given $\mathcal{A}$ and $\mathcal{A}'$ in the same isogeny class, compute an isogeny $\mathcal{A} \to \mathcal{A}'$.

The isogeny problem is supposed to be hard for supersingular elliptic curves.

Superspecial g-dimensional PPAVs over $\mathbb{F}_{p^2}$: $O(p^{g(g+1)/2})$ vertices.

- **Costello–Smith** `https://ia.cr/2019/1387`
  Evidence that the isogeny problem might be substantially less hard for higher-dimensional superspecial PPAVs.
  Classical $\widetilde{O}(p^{g-1})$ operations
  Quantum $\widetilde{O}(p^{(g-1)/2})$ Grover oracle calls

Much more research required here!