Modular polynomials for abelian surfaces and related algorithms

Jean Kieffer (CNRS)

KULB seminar, Leuven, May 17, 2024

Classical modular polynomials

Fix $\ell \geq 1$ prime. The classical modular polynomial of level ℓ

$$\Phi_\ell \in \mathbb{Z}[X, Y]$$

satisfies: if k is a field of char. $\neq \ell$, and E, E' are elliptic curves over k, then

$$\Phi_\ell(j(E), j(E')) = 0 \iff E \text{ and } E' \text{ are } \ell\text{-isogenous over } \overline{k}.$$

Example

$$\begin{aligned}\Phi_2(X, Y) = X^3 + Y^3 - X^2Y^2 + 1488X^2Y + 1488XY^2 - 162000X^2 - 162000Y^2 \\ + 40773375XY + 8748000000X + 8748000000Y - 157464000000000.\end{aligned}$$

Used to navigate isogeny graphs and compute isogenies.

Navigating isogeny graphs

2-isogeny graph of supersingular elliptic curves over $\mathbb{F}_{p^2}$:

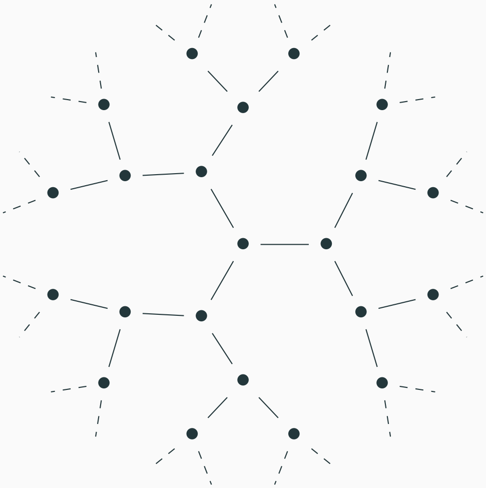

2-isogeny graph of supersingular elliptic curves over $\mathbb{F}_{p^2}$:

- Starting point: E

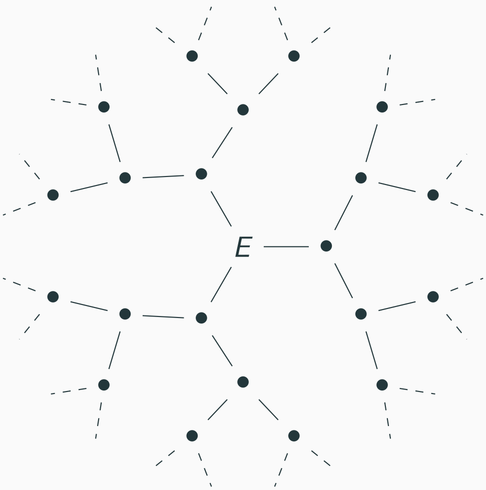

Navigating isogeny graphs

2-isogeny graph of supersingular elliptic curves over $\mathbb{F}_{p^2}$:

- Starting point: E
- Solve $\Phi_2(j(E), Y) = 0$ in $\mathbb{F}_{p^2}$: find 3 roots

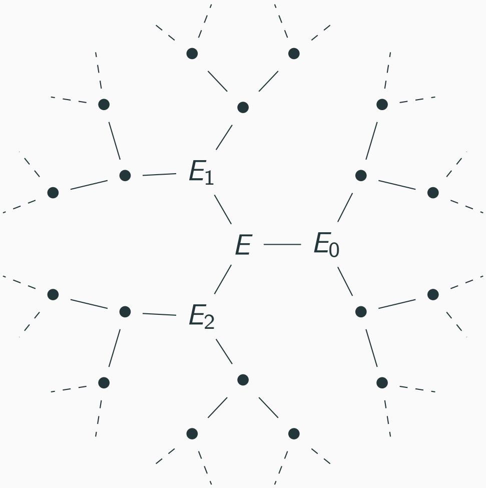

Navigating isogeny graphs

2-isogeny graph of supersingular elliptic curves over $\mathbb{F}_{p^2}$:

- Starting point: E
- Solve $\Phi_2(j(E), Y) = 0$ in $\mathbb{F}_{p^2}$: find 3 roots
- Pick path to E_0, say

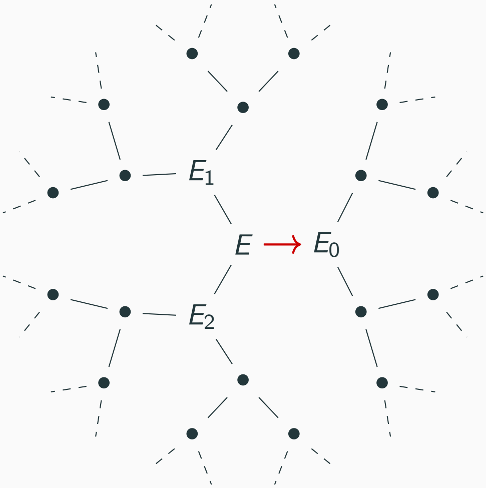

Navigating isogeny graphs

2-isogeny graph of supersingular elliptic curves over $\mathbb{F}_{p^2}$:

- Starting point: E
- Solve $\Phi_2(j(E), Y) = 0$ in $\mathbb{F}_{p^2}$: find 3 roots
- Pick path to E_0, say

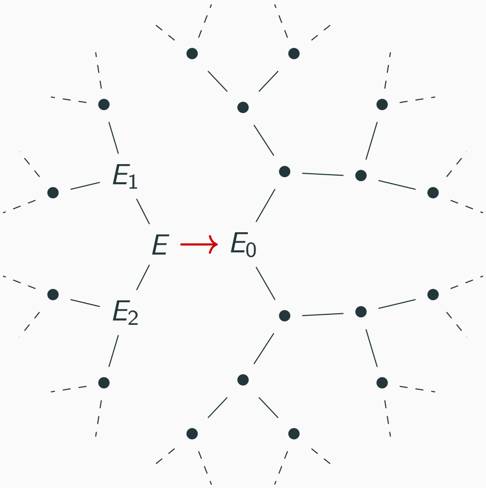

Navigating isogeny graphs

2-isogeny graph of supersingular elliptic curves over $\mathbb{F}_{p^2}$:

- Starting point: E
- Solve $\Phi_2(j(E), Y) = 0$ in $\mathbb{F}_{p^2}$: find 3 roots
- Pick path to E_0, say
- Solve $\Phi_2(j(E_0), Y)/(Y - j(E)) = 0$: find 2 roots $j(E_{00}), j(E_{01})$

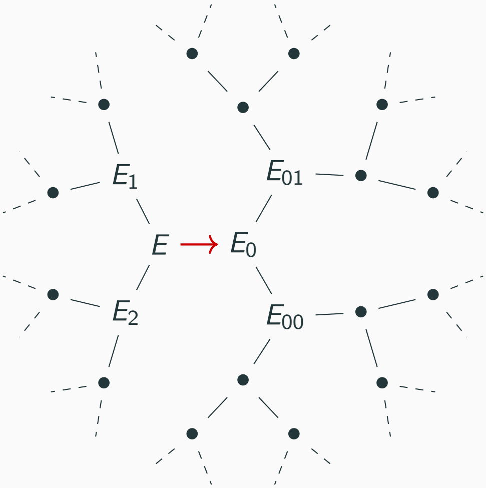

Navigating isogeny graphs

2-isogeny graph of supersingular elliptic curves over $\mathbb{F}_{p^2}$:

- Starting point: E
- Solve $\Phi_2(j(E), Y) = 0$ in $\mathbb{F}_{p^2}$: find 3 roots
- Pick path to E_0, say
- Solve $\Phi_2(j(E_0), Y)/(Y - j(E)) = 0$: find 2 roots $j(E_{00}), j(E_{01})$
- Pick path to E_{01}, say

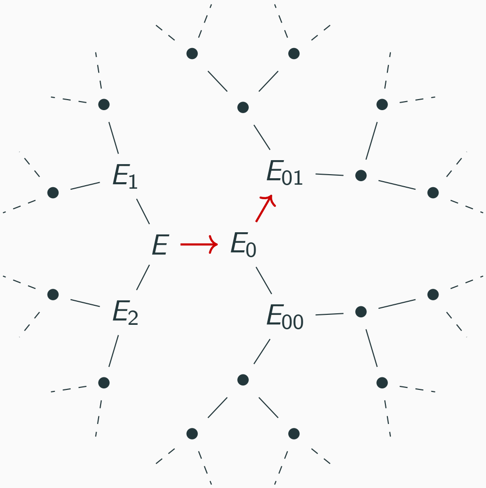

Navigating isogeny graphs

2-isogeny graph of supersingular elliptic curves over $\mathbb{F}_{p^2}$:

- Starting point: E
- Solve $\Phi_2(j(E), Y) = 0$ in $\mathbb{F}_{p^2}$: find 3 roots
- Pick path to E_0, say
- Solve $\Phi_2(j(E_0), Y)/(Y - j(E)) = 0$: find 2 roots $j(E_{00})$, $j(E_{01})$
- Pick path to E_{01}, say
- Continue!

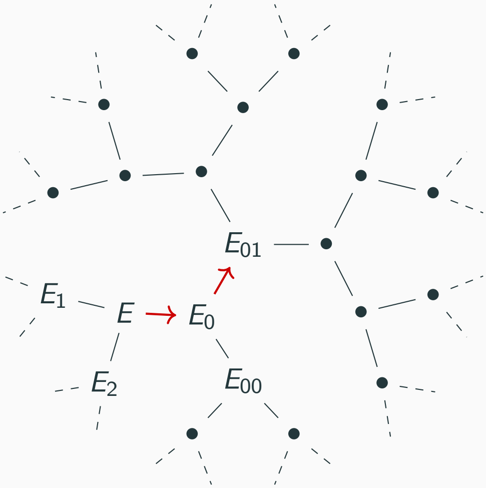

Computing isogenies

Theorem (Elkies '95, Bostan–Morain–Salvy–Schost '08)

- ℓ prime, k a field of char. 0 or $> 4\ell + 1$.
- E, E' elliptic curves over k that are ℓ-isogenous.
- Assume $\partial_X \Phi_\ell(j(E), j(E')) \neq 0$, i.e. $j(E')$ is a simple root of $\Phi_\ell(j(E), Y)$.
  This is true generically.

Then, given E, E' and $\partial_X \Phi_\ell(j(E), j(E'))$, one can compute polynomial formulas for the ℓ-isogeny

$$\varphi : E \to E',$$

in particular an equation of $\ker \varphi$, in $\widetilde{O}(\ell)$ operations in k (quasi-linear time.)

Complexity bounds

The height of $F \in \mathbb{Q}(X_1, \ldots, X_n)$ is

$$h(F) = \log(\max |c|), \quad \text{where } c \text{ runs through the coefficients of } F.$$

Complexity bounds for Φ_ℓ

- Φ_ℓ has degree $\ell + 1$ in both variables X and Y.
- $h(\Phi_\ell) \sim 6\ell \log \ell$ [Cohen '84]. Storing Φ_ℓ costs $O(\ell^3 \log \ell)$ space.
- Φ_ℓ can be computed in quasi-linear time $\widetilde{O}(\ell^3)$ [Enge '09, Bröker–Lauter–Sutherland '12, Sutherland '13].

In summary:

- Φ_ℓ allow us to manipulate isogenies without torsion input.
- Cheaper than computing (subgroups of) $E[\ell]$ from scratch: e.g. the SEA algorithm (Schoof–Elkies–Atkin '90s) computes $\#E(\mathbb{F}_q)$ in time $\widetilde{O}(\log^4 q)$.

State of the art

	dim 1	dim 2	dim g
Definition of Φ_ℓ	✓		
Complexity bounds	✓		
Evaluating $\Phi_\ell(j(E), Y)$	✓		
Isogenies without torsion input	✓		
Point counting	✓		
More compact variants of Φ_ℓ	✓ Atkin, ...		

Higher dimensions

State of the art

	dim 1	dim 2	dim g
Definition of Φ_ℓ	✓	✓ Bröker–Lauter '09, ...	
Complexity bounds	✓		
Evaluating $\Phi_\ell(j(E), Y)$	✓		
Isogenies without torsion input	✓		
Point counting	✓		
More compact variants of Φ_ℓ	✓ Atkin, ...		

State of the art

	dim 1	dim 2	dim g
Definition of Φ_ℓ	✓	✓ Bröker–Lauter '09, ...	✓K. '22
Complexity bounds	✓	✓K. '22	✓K. '22
Evaluating $\Phi_\ell(j(E), Y)$	✓	✓K. '2?	? partial
Isogenies without torsion input	✓	✓K., Page, Robert '2?	? dim 3?
Point counting	✓	✓K. '2?	? RM?
More compact variants of Φ_ℓ	✓Atkin, ...	? Theta functions?	??

Goals of this talk

- Generalize Φ_ℓ using the geometry of moduli spaces.
- Briefly talk about complexity bounds and computing isogenies.
- Present the evaluation algorithm, its performance and applications.

Geometric interpretation of Φ_ℓ (1)

It is easier to work over $\mathbb{C}$.

- $\mathcal{H}_1$ is the upper half plane $\{\mathrm{Im}(\tau) > 0\}$. Action of $\mathrm{SL}_2(\mathbb{Z})$ on $\mathcal{H}_1$:

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \cdot \tau = \frac{a\tau + b}{c\tau + d}.$$

  $\mathcal{A}_1 = \mathrm{SL}_2(\mathbb{Z})\backslash\mathcal{H}_1$ is the moduli space of elliptic curves. It is an algebraic variety defined over $\mathbb{Q}$. We view the j-invariant as a coordinate on $\mathcal{A}_1$.

- Let $\Gamma^0(\ell) \subset \mathrm{SL}_2(\mathbb{Z})$ be the subgroup of matrices such that $b = 0 \bmod \ell$. It has index $\ell + 1$ in $\mathrm{SL}_2(\mathbb{Z})$.
  The quotient $\mathcal{A}_1(\ell) = \Gamma^0(\ell)\backslash\mathcal{H}_1$ is the moduli space of pairs (E, K) where $K \subset E$ is the kernel of an ℓ-isogeny. It is a more complicated curve than $\mathcal{A}_1$.

Geometric interpretation of Φ_ℓ (2)

We have two maps $\mathcal{A}_1(\ell) \to \mathcal{A}_1$, both $(\ell+1)$-to-one:

	$A_1(\ell) \to \mathcal{A}_1$	$\Gamma^0(\ell)\backslash\mathcal{H}_1 \to \mathrm{SL}_2(\mathbb{Z})\backslash\mathcal{H}_1$
"Domain" map	$(E,K) \mapsto E$	$\tau \mapsto \tau$
"Codomain" map	$(E,K) \mapsto E/K$	$\tau \mapsto \tau/\ell.$

Geometric interpretation

Φ_ℓ is an equation for the image in $\mathcal{A}_1 \times \mathcal{A}_1$ of the joint map

$$\begin{aligned} \mathcal{A}_1(\ell) &\to \mathcal{A}_1 \times \mathcal{A}_1 \\ (E,K) &\mapsto (E, E/K), \end{aligned}$$

using the j-invariant as a coordinate on $\mathcal{A}_1$.

For every $\tau \in \mathcal{H}_1$,

$$\Phi_\ell(j(\tau), Y) = \prod_{\gamma \in \Gamma^0(\ell)\backslash\Gamma(1)} \Big(Y - j(\tfrac{1}{\ell}\gamma\tau)\Big).$$

Modular polynomials for abelian surfaces (1)

- $\mathcal{A}_2$ is the moduli space of principally polarized abelian surfaces.
  $\mathcal{A}_2$ is an algebraic variety defined over $\mathbb{Q}$ of dimension 3, consisting of Jacobians of genus 2 curves (dense open) and products $E_1 \times E_2$ (dimension 2 subvariety).
- The Igusa invariants j_1, j_2, j_3 are convenient coordinates on $\mathcal{A}_2$.
- $A_2(\ell)$ is the moduli space of pairs (A, K) where K is the kernel of an (ℓ, ℓ)-isogeny, i.e. $K \subset A[\ell]$ is isomorphic to $(\mathbb{Z}/\ell\mathbb{Z})^2$ and isotropic for the Weil pairing.

Modular polynomials for abelian surfaces

The Siegel modular polynomials $\Psi_{\ell,1}, \Psi_{\ell,2}, \Psi_{\ell,3}$ are equations for the image of

$$\begin{array}{rcl} \mathcal{A}_2(\ell) & \to & \mathcal{A}_2 \times \mathcal{A}_2 \\ (A, K) & \mapsto & (A, A/K) \end{array}$$

using the Igusa invariants as coordinates on $\mathcal{A}_2$.

Modular polynomials for abelian surfaces (2)

The image of $\mathcal{A}_2(\ell)$ is a dimension 3 subvariety in a dimension 6 ambient space, so has several possible sets of equations.

Choose the polynomials $\Psi_{\ell,k}$ such that $\Psi_{\ell,k} \in \mathbb{Q}(X_1, X_2, X_3)[Y]$ and

$$\Psi_{\ell,1}\Big(j_1(A), j_2(A), j_3(A), j_1(A/K)\Big) = 0,$$

$$j_2(A/K) = \frac{\Psi_{\ell,2}}{\partial_Y \Psi_{\ell,1}}\Big(j_1(A), j_2(A), j_3(A), j_1(A/K)\Big), \quad \text{and same for } j_3.$$

Note

- Computing the isogenous abelian surfaces is easy (no Gröbner bases!)
- Convenient analytic formulas as in the case of Φ_ℓ.
- Can play the same game with any moduli space of abelian varieties: any dimension g, real multiplication, level structures, etc. PEL Shimura varieties.

State of the talk

	dim 1	dim 2	dim g
Definition of Φ_ℓ	✓	✓ Bröker–Lauter '09, …	✓K. '22
Complexity bounds	✓		
Evaluating $\Phi_\ell(j(E), Y)$	✓		
Isogenies without torsion input	✓		
Point counting	✓		
More compact variants of Φ_ℓ	✓Atkin, …	? Theta functions?	??

Complexity bounds

Recall: $\Psi_{\ell,k} \in \mathbb{Q}(X_1, X_2, X_3)[Y]$ for $1 \leq k \leq 3$.

Theorem (K. '22)

- The degree of $\Psi_{\ell,k}$ in each variable is $O(\ell^3)$. Tight explicit bounds.
- $h(\Psi_{\ell,k}) = O(\ell^3 \log \ell)$. Explicit bounds (huge, not tight).

A general theorem applies to modular polynomials on any PEL Shimura variety.

Corollary

- The size of $\Psi_{\ell,k}$ as a 4-variable fraction is $O(\ell^{15} \log \ell)$. [Note: 410 MB for $\ell = 3$]
- If $j_1, j_2, j_3 \in \mathbb{Q}$, the size of $\Psi_{\ell,k}(j_1, j_2, j_3, Y) \in \mathbb{Q}[Y]$ is $O(\ell^6(H + \log \ell))$ where $H = \max\{h(j_1), h(j_2), h(j_3)\}$.

We need an algorithm to evaluate the modular polynomials at (j_1, j_2, j_3) directly!

Computing isogenies

Theorem (K., Page, Robert)

- ℓ prime, k a field of char. 0 or $> 8\ell + 7$.
- A, A' Jacobians of genus 2 curves over k that are (ℓ, ℓ)-isogenous.
- Assume that the 3×3 matrix $\left(\partial_{X_i} \Psi_{\ell,k}\right)_{i,k}$ evaluated at the Igusa invariants of A, A' is invertible. This is true generically.

Then, given A, A' and the above matrix, one can compute polynomial formulas for the (ℓ, ℓ)-isogeny

$$\varphi : A \to A'$$

in $\widetilde{O}(\ell)$ operations in k (quasi-linear time).

One can then compute $\ker(\varphi)$ using polynomial arithmetic (resultants...)

The evaluation algorithm should also evaluate the derivatives of $\Psi_{\ell,k}$.

State of the talk

	dim 1	dim 2	dim g
Definition of Φ_ℓ	✓	✓ Bröker–Lauter '09, ...	✓K. '22
Complexity bounds	✓	✓K. '22	✓K. '22
Evaluating $\Phi_\ell(j(E), Y)$	✓		
Isogenies without torsion input	✓	✓K., Page, Robert '2?	? dim 3?
Point counting	✓		
More compact variants of Φ_ℓ	✓Atkin, ...	? Theta functions?	??

The evaluation algorithm

Analytic formula

Recall: for $\tau \in \mathcal{H}_1$,

$$\Phi_\ell(j(\tau), Y) = \prod_{\gamma \in \Gamma^0(\ell)\backslash \mathrm{SL}_2(\mathbb{Z})} \Big(Y - j(\tfrac{1}{\ell}\gamma\tau)\Big).$$

Similar formula in dimension 2:

- $\mathcal{H}_2 = \{\tau \in \mathrm{Mat}_{2\times 2}(\mathbb{C}) : \tau \text{ symmetric}, \operatorname{Im}\tau \text{ pos. def.}\}$: Siegel upper half space.
- The symplectic group $\mathrm{Sp}_4(\mathbb{Z})$ acts on $\mathcal{H}_2$: in block notation,

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \cdot \tau = (a\tau + b)(c\tau + d)^{-1}.$$

- Subgroup $\Gamma^0(\ell) \subset \mathrm{Sp}_4(\mathbb{Z})$ defined by $b = 0 \bmod \ell$, with index $\ell^3 + \ell^2 + \ell + 1$.

For instance: for $\tau \in \mathcal{H}_2$,

$$\Psi_{\ell,1}(j_1(\tau), j_2(\tau), j_3(\tau), Y) = \prod_{\gamma \in \Gamma^0(\ell)\backslash \mathrm{Sp}_4(\mathbb{Z})} \Big(Y - j_1(\tfrac{1}{\ell}\gamma\tau)\Big).$$

Outline

The evaluation algorithm

Let $j_1, j_2, j_3 \in \mathbb{Q}$ of height H be given.

1. Find $\tau \in \mathcal{H}_2$ with these Igusa invariants (a period matrix) at high precision.
2. Enumerate the matrices $\frac{1}{\ell}\gamma\tau$ and compute their Igusa invariants.
3. Compute the modular polynomials in $\mathbb{C}[Y]$ using the analytic formula.
4. Recognize each coefficient as a rational number.

- This algorithm has been implemented in C using the libraries FLINT/Arb.
- We use interval arithmetic throughout to ensure correctness. In step 4, we can actually get integers instead of rational numbers.
- In step 1, we use the AGM method (Dupont '06) with some improvements.
- Step 2 dominates the algorithm and relies on theta functions: stay tuned.

Main result

Theorem (K.)

We can evaluate the Siegel modular polynomials of level ℓ and their derivatives at

1. *a generic point $(j_1, j_2, j_3) \in \mathbb{Q}^3$ of height at most H in time $\widetilde{O}(\ell^3 H^2 + \ell^6 H)$,*
2. *a generic point $(j_1, j_2, j_3) \in \mathbb{F}_p^3$ for p prime in time $\widetilde{O}(\ell^3 \log^2 p + \ell^6 \log p)$.*

This is almost quasi-linear time.

"Generic" means that the algorithm will fail on a closed dimension 2 subvariety of $\mathcal{A}_2$ (e.g. Igusa invariants not defined...)

Proof of 2.: lift to $\mathbb{Q}$ and apply 1.! To handle $\mathbb{F}_q$, we extend 1. to number fields.

Practical timings

Time to evaluate $\Psi_{\ell,k}(j_1, j_2, j_3, Y)$ at $(j_1, j_2, j_3) =$ random 3-digit rational numbers:

ℓ	2	3	5	7	11	13	17
Time (s)	1.3	5.1	97	1200	40000	$1.6 \cdot 10^5$	$1.1 \cdot 10^6$
$0.002\,\ell^6 \log^3(\ell) \log\log(\ell)$	-	-	62	1200	43000	$1.5 \cdot 10^5$	$1.1 \cdot 10^6$

Using related methods, we computed a Jacobians of genus 2 curves over $\mathbb{Q}$ linked by isogenies of large degree, e.g. $(19^2, 19, 19)$ or $(31, 31)$, in roughly 1h (van Bommel, Costa, Chidambaram, K. '24).

Consequences on point counting

Results

- If A is a p.p. abelian surface over $\mathbb{F}_p$ with small Igusa invariants, then we compute $\#A(F_p)$ in heuristic time $\widetilde{O}(\log^7 p)$. Improves on Schoof's method in $\widetilde{O}(\log^8 p)$ (Gaudry–Schost '12)
- If $A/\mathbb{Q}$ is fixed, then we can compute $\#A(\mathbb{F}_p)$ for several primes p (in fact $\Omega(H \log p)$ of them) in average time $\widetilde{O}(\log^6 p)$.
- If $A/\mathbb{F}_p$ has real multiplication by $\mathbb{Q}(\sqrt{5})$ or another small real quadratic field, then we compute $\#A(\mathbb{F}_p)$ in time $\widetilde{O}(\log^4 p)$ as in the dimension 1 case.

I still need an implementation to (hopefully) establish a new point-counting record.

Theta functions

Theta functions

Recall: in the evaluation algorithm, we get matrices $\tau_1, \ldots, \tau_n \in \mathcal{H}_2$. We need to evaluate their Igusa invariants in $\mathbb{C}$ at high precision N, i.e. up to an error of $\leq 2^{-N}$.

We do this in quasi-linear time $O(\mathcal{M}(N) \log N)$ using theta functions.

Definition

Fix theta characteristics $a, b \in \{0,1\}^g$. Then

$$\theta_{a,b}(\tau) = \sum_{n \in \mathbb{Z}^g + \frac{a}{2}} \exp(i\pi(n^T \tau n + n^T b)).$$

- They are 2^{2g} analytic functions on $\mathcal{H}_g$ (16 for $g = 2$.)
- Coordinates on $\mathcal{A}_g$, e.g. the Igusa invariants, can be expressed as rational fractions in terms of theta functions.

Main theorem on theta functions

Theorem (Elkies, K., in preparation)

Given $g \geq 1$, $N \geq 0$, and given $\tau \in \mathcal{H}_g$ and $z \in \mathbb{C}^g$ that are suitably reduced, one can evaluate $\theta_{a,b}(z,\tau)$ for all characteristics (a, b) to precision N in quasi-linear time $O(2^{O(g \log g)} \mathcal{M}(N) \log N)$, uniformly in τ and z.

- Implemented in FLINT 3.1: `https://flintlib.org/doc/acb_theta.html`
- The "naive" algorithm (sum the exponential series) is not quasi-linear.
- Earlier works (Dupont '06, Labrande–Thomé '14) are specific to small g and tricky to run in interval arithmetic. This new algorithm is $\sim 10\times$ faster for $g = 2$. The timings above were with Dupont's algorithm.
- When evaluating modular polynomials, we add a (negligible) reduction step.
- For general g, how do we compute τ in the first place?

Questions

Thank you for listening!
Any questions?

	dim 1	dim 2	dim g
Definition of Φ_ℓ	✓	✓ Bröker–Lauter '09, ...	✓K. '22
Complexity bounds	✓	✓K. '22	✓K. '22
Evaluating $\Phi_\ell(j(E), Y)$	✓	✓K. '2?	? partial
Isogenies without torsion input	✓	✓K., Page, Robert '2?	? dim 3?
Point counting	✓	✓K. '2?	? RM?
More compact variants of Φ_ℓ	✓Atkin, ...	? Theta functions?	??