Canonical heights on Jacobians of curves of genus two

Steffen Müller (Universiteit Groningen)
joint with
Michael Stoll (Universität Bayreuth)

Arithmetic of Hyperelliptic Curves, ICTP, Trieste
7 September 2017

Mordell-Weil

Let $A/\mathbb{Q}$ be an abelian variety (everything works more generally over number fields).

Theorem. (Mordell-Weil)

$$A(\mathbb{Q}) \cong \mathbb{Z}^r \oplus T\,,$$

where

- $r \geq 0$ is the rank of $A(\mathbb{Q})$,
- $T \cong A(\mathbb{Q})_{\mathrm{torsion}}$ is finite.

Goal. Compute generators for $A(\mathbb{Q})$.

- Generators for $A(\mathbb{Q})_{\mathrm{torsion}}$ are often easy to compute.
- Suppose we have computed r and also $Q_1, \ldots, Q_r \in A(\mathbb{Q})$ which generate a subgroup of $A(\mathbb{Q})/A(\mathbb{Q})_{\mathrm{torsion}}$ of finite index.

Properties

Theorem. (Néron, Tate). There is a positive semidefinite quadratic form $\hat{h} : A(\mathbb{Q}) \to \mathbb{R}$ with the following properties:

(a) $\hat{h}(Q) = 0$ if and only if Q is torsion.

(b) $\hat{h}$ is a positive definite quadratic form on $V := A(\mathbb{Q}) \otimes_{\mathbb{Z}} \mathbb{R}$.

(c) $\{Q \in A(\mathbb{Q}) : \hat{h}(Q) \leq B\}$ is finite for every $B \in \mathbb{R}$.

We call $\hat{h}$ the canonical height (or Néron-Tate height) on $A(\mathbb{Q})$.

Recall that the regulator of $A/\mathbb{Q}$, appearing in the conjecture of Birch and Swinnerton-Dyer, is defined using $\hat{h}$.

Back to our problem:

- $\Lambda := A(\mathbb{Q})/A(\mathbb{Q})_{\text{torsion}}$ is a lattice inside the Euclidean vector space $(V, \hat{h})$.
- $Q_1, \ldots, Q_r$ generate a finite-index sublattice $\Lambda' \leq \Lambda \subset V$.

Hence saturating Λ' inside V gives Λ.

Saturation

Method 1. (Siksek, Flynn – Smart)

- Compute an upper bound u on $(\Lambda : \Lambda')$.
- For every prime $p \leq u$ check if $p \mid (\Lambda : \Lambda')$;
  - if yes, find $Q \in \Lambda \setminus \Lambda'$ such that $pQ \in \Lambda'$ set $\Lambda' := \langle \Lambda', Q \rangle$ and repeat;
  - if no, continue with the next p.

Method 2. (Stoll) The lattice Λ is generated by

$$\{Q_1, \ldots, Q_r\} \cup \{Q \in \Lambda : \hat{h}(Q) \leq \rho^2\},$$

where ρ is the covering radius of Λ' (i.e. the maximal distance a point of V can have from Λ').

Basic computational problems

For both methods, we need to

(1) construct $\hat{h}$;

(2) compute $\hat{h}(Q)$ for given points $Q \in A(\mathbb{Q})$;

(3) enumerate $\{Q \in A(\mathbb{Q}) : \hat{h}(Q) \leq B\}$ for given $B \in \mathbb{R}$.

For (1), start with the standard height function on $\mathbb{P}^N(\mathbb{Q})$:

$$h(x_0 : \ldots : x_N) := \log \max\{|x_0|, \ldots, |x_N|\}\,,$$

where $x_0, \ldots, x_N \in \mathbb{Z}$ and $\gcd(x_0, \ldots, x_N) = 1$.

Elliptic curves

Let $A/\mathbb{Q}$ be an elliptic curve, given by a Weierstrass equation with integral coefficients.

Define $\kappa : A(\mathbb{Q}) \to \mathbb{P}^1(\mathbb{Q})$ by

$$\kappa(x : y : 1) := (x : 1)\,, \qquad \kappa(0 : 1 : 0) := (1 : 0)\,.$$

The naive height of $Q \in A(\mathbb{Q})$ is given by

$$h(Q) := h(\kappa(Q)) \in \mathbb{R}_{\geq 0}\,.$$

Tate constructed the canonical height of Q by setting

$$\hat{h}(Q) := \lim_{n\to\infty} 4^{-n} h(2^n Q) \in \mathbb{R}_{\geq 0}\,.$$

Jacobians of genus 2 curves

Let $A/\mathbb{Q}$ be the Jacobian of a curve $X/\mathbb{Q}$ of genus 2 and let

$$W : y^2 = f(x) = f_0 + f_1 x + \ldots + f_6 x^6$$

be an integral Weierstrass equation for X.

Flynn: Explicit map $\kappa : A \to \mathbb{P}^3$ such that $\kappa(A)$ is a model for the Kummer surface $K := {}^{A}\!/_{\langle -1 \rangle}$ of A and $\kappa(0) = (0:0:0:1)$.

The naive height of $Q \in A(\mathbb{Q})$ is

$$h(Q) := h(\kappa(Q))$$

Again, we get the canonical height of Q by setting

$$\hat{h}(Q) := \lim_{n \to \infty} 4^{-n} h(2^n Q) \in \mathbb{R}_{\geq 0}\,.$$

Computational goals

Recall that we want to

(1) construct $\hat{h}$;

(2) compute $\hat{h}(Q)$ for given points $Q \in A(\mathbb{Q})$;

(3) enumerate $\{Q \in A(\mathbb{Q}) : \hat{h}(Q) \leq B\}$ for given $B \in \mathbb{R}$.

For (2) and (3), we use that

$$\Psi := h - \hat{h} \quad \text{is bounded.}$$

Computational goals.

(I) Compute $\Psi(Q)$ for given $Q \in A(\mathbb{Q})$.

(II) Compute an upper bound β for Ψ.

(III) Given $B \in \mathbb{R}$, enumerate

$$\{Q \in A(\mathbb{Q}) : h(Q) \leq B + \beta\} \supset \{Q \in A(\mathbb{Q}) : \hat{h}(Q) \leq B\}.$$

Néron vs. Tate

For explicit computations, Tate's simple limit construction is not suitable, as the size of the coefficients of $2^n Q$ grows exponentially.

Instead, one uses Néron's construction of Ψ as a sum of local terms Ψ_v. However, this is rather more complicated...

"Il faudrait que tu m'expliques une fois ce que sont ces symboles locaux de Néron. Je n'ai rien compris à ce que Lang en disait - et je n'avais pas compris davantage le papier de Néron que j'ai eu une fois entre les mains. Mais quel animal ce Néron! Sous ses air patauds, il ne démontre jamais que des choses fondamentales! (Letter from Serre to Grothendieck, 1964)

We construct Ψ_v explicitly when $A = \mathrm{Jac}(X)$ and X is a curve of genus 2 given by an integral Weierstrass equation $W : y^2 = f(x)$.

For this we first decompose $4h(Q) - h(2Q)$ into local terms.

Duplication on the Kummer

Recall the map $\kappa : A \to \mathbb{P}^3$ such that

- $\kappa(A)$ is a model for the Kummer surface $K = A/\langle -1 \rangle$ of A,
- $\kappa(0) = (0 : 0 : 0 : 1)$.

Flynn. There are homogeneous quartic polynomials $\delta_1, \ldots, \delta_4 \in \mathbb{Z}[x_1, \ldots, x_4]$ such that for $\delta = (\delta_1, \ldots, \delta_4)$

- the following diagram commutes:

$$\begin{array}{ccc} A & \xrightarrow{[2]} & A \\ {\scriptstyle\kappa}\downarrow & & \downarrow{\scriptstyle\kappa} \\ K & \xrightarrow[\delta]{} & K \end{array}$$

- $\delta(0, 0, 0, 1) = (0, 0, 0, 1)$.

Explicit local height correction functions

Define, for

- v a place of $\mathbb{Q}$,
- $Q \in A(\mathbb{Q}_v)$,
- $\kappa(Q) = (x_1 : \ldots : x_4)$,

$$\tilde{\varepsilon}_v(Q) := -\log \max\{|\delta_j(x_1, x_2, x_3, x_4)|_v : 1 \le j \le 4\} + 4 \log \max\{|x_j|_v : 1 \le j \le 4\}\,.$$

Then we have

$$\sum_v \tilde{\varepsilon}_v(Q) \;=\; -h(2Q) + 4h(Q) \quad \text{for } Q \in A(\mathbb{Q})\,,$$

and $\tilde{\varepsilon}_v : A(\mathbb{Q}_v) \to \mathbb{R}$ is v-adically continuous and hence bounded.

If p is a prime number, then

- $\varepsilon_p(Q) := \tilde{\varepsilon}_p(Q)/\log p \in \mathbb{Z}_{\ge 0}$,
- $\tilde{\varepsilon}_p(Q) = 0$ if A has good reduction at p.

Decomposing Ψ

Tate's telescoping trick shows for $Q \in A(\mathbb{Q})$:

$$h(Q) - \hat{h}(Q) = \sum_{n=0}^{\infty} 4^{-(n+1)} \left(4h(2^n P) - h(2^{n+1} P)\right)$$

$$= \sum_{v} \sum_{n=0}^{\infty} 4^{-(n+1)} \tilde{\varepsilon}_v(2^n Q),$$

so we define

$$\Psi_v(Q) := \sum_{n=0}^{\infty} 4^{-(n+1)} \tilde{\varepsilon}_v(2^n Q)$$

to get

$$\Psi = h - \hat{h} = \sum_{v} \Psi_v \,.$$

Local computational goals

Our computational goals

(I) compute $\Psi(Q)$ for given $Q \in A(\mathbb{Q})$;

(II) compute an upper bound for Ψ.

now reduce to

(i) compute $\Psi_p(Q)$ for given $Q \in A(\mathbb{Q}_p)$ if p is a bad prime;

(ii) compute an upper bound for Ψ_p if p is a bad prime;

(iii) compute $\Psi_\infty(Q)$ for given $Q \in A(\mathbb{R})$;

(iv) compute an upper bound for Ψ_∞

Previous algorithms are due to Flynn-Smart, Stoll, Uchida.

The "kernel" of μ_p

Let p be a prime of bad reduction and set

$$\mu_p(Q) := \frac{\Psi_p(Q)}{\log p} = \sum_{n=0}^{\infty} 4^{-n-1}\varepsilon_p(2^n Q) \in \mathbb{Q}_{\geq 0}\,.$$

Theorem. (Stoll) For

$$U := \{Q \in A(\mathbb{Q}_p) : \mu_p(Q) = 0\}$$

we have

(a) U is a finite-index subgroup of $A(\mathbb{Q}_p)$ containing the kernel of reduction (with respect to the given model).

(b) Both μ_p and ε_p factor through $A(\mathbb{Q}_p)/U$.

Can we say more about U?

More structure

Let

- $\mathcal{A}$ be the Néron model of $A_{\mathbb{Z}_p}$, with component group Φ;
- $\mathcal{A}^0$ be the identity component of $\mathcal{A}$,
- $A_0(\mathbb{Q}_p)$ denote the points in $A(\mathbb{Q}_p)$ reducing to $\mathcal{A}^0_p$.

Theorem A. Suppose that $W_{\mathbb{Z}_p}$ has rational singularities.

Then μ_p factors through $A(\mathbb{Q}_p)/A_0(\mathbb{Q}_p) \cong \Phi(\mathbb{F}_p)$.

- $W_{\mathbb{Z}_p}$ has rational singularities if $R^i\xi_*\mathcal{O}_{\mathcal{W}}$ vanishes for all $i > 0$, where $\xi : \mathcal{W} \to W_{\mathbb{Z}_p}$ is a desingularization.
- For the proof we suppose that $W_{\mathbb{Z}_p}$ is normal and reduced.
  - First show that $\varepsilon_p(Q) = 0$ if the reduction of Q is in the image of the canonical morphism $\alpha : \mathrm{Pic}^0_{W/\mathbb{Z}_p} \to \mathcal{A}^0$.
  - Then use that α is an isomorphism if and only if $W_{\mathbb{Z}_p}$ has rational singularities.

The reduction graph

We can sometimes use Theorem A to give a formula for μ_p.

If the minimal regular model $\mathcal{X}^{\min}$ of $X_{\mathbb{Z}_p}$ is semistable, then the reduction graph $\mathcal{R}$ is defined as follows:

- The vertices are the irreducible components of the special fiber $\mathcal{X}_p^{\min}$.
- Two vertices Γ_1 and Γ_2 are connected by n edges, where n is the number of
  - intersection points of Γ_1 and Γ_2 if $\Gamma_1 \neq \Gamma_2$,
  - nodes of Γ_1 if $\Gamma_1 = \Gamma_2$.
- We put a metric on $\mathcal{R}$ by giving each edge length 1

We can also interpret $\mathcal{R}$ as an electric network with unit resistance along every edge.

A resistance formula for μ_p

Theorem B. Suppose that the minimal regular model $\mathcal{X}^{\min}$ of $X_{\mathbb{Z}_p}$ is semistable and $W_{\mathbb{Z}_p}$ has rational singularities.

Let $Q \in A(\mathbb{Q}_p)$ be such that its image in $\Phi(\mathbb{F}_p)$ is represented by $\Gamma_1 - \Gamma_2$, where Γ_1 and Γ_2 are components of $\mathcal{X}_p^{\min}$. Then

$$\mu_p(Q) = r(\Gamma_1, \Gamma_2)$$

is the resistance between Γ_1 and Γ_2 on $\mathcal{R}$.

Sketch of proof. Express μ_p in terms of Zhang's admissible intersection pairing on X. The latter can then be related to the resistance, using that we can vary Q on $[\Gamma_1 - \Gamma_2]$ by Theorem A.

Every $\Theta \in \Phi$ can be represented as $\Gamma_1 - \Gamma_2$, where the Γ_i are irreducible components of $\mathcal{X}_p^{\min}$.

General bounds

From our theorems we get formulas and sharp bounds for μ_p for the most frequent reduction types. What about the general case?

Let Δ be the discriminant of W.

Proposition. If $Q \in A(\mathbb{Q}_p)$, then

$$\mu_p(Q) \leq \frac{\mathrm{ord}_p(\Delta)}{4}\,.$$

Sketch of proof. Using Theorems A and B, show that the statement holds when $W_{\mathbb{Z}_p}$ is minimal and $\mathcal{X}^{\min}$ is semistable.

Then reduce to this case over an extension by studying how μ_p changes when we change the model $W_{\mathbb{Z}_p}$.

Computing non-archimedean corrections

To compute μ_p, recall that

$$\mu_p(Q) = \sum_{n=0}^{\infty} 4^{-n-1} \varepsilon_p(2^n Q) \in \mathbb{Q}_{\geq 0}\,. \tag{1}$$

Lemma. We have

(a) $0 \leq \varepsilon_p(Q) \leq \operatorname{ord}_p(\Delta)$,

(b) $\operatorname{denom}(\mu_p(Q)) \leq \max\left\{2, \lfloor \operatorname{ord}_p(\Delta)^2/3 \rfloor\right\}$.

- By (1) and (a), we can approximate $\mu_p(Q)$ to any desired accuracy by repeatedly applying the duplication map δ.
- For sufficiently small error, (b) lets us pin down $\mu_p(Q)$ exactly using continued fractions.
- This algorithm is quasi-linear in $p \cdot \operatorname{ord}_p(\Delta)$.

Avoiding integer factorisation

Even better, we can globalize the local algorithm to compute

$$\Psi^{\mathrm{f}}(Q) := \sum_p \mu_p(Q) \log p$$

for $Q \in A(\mathbb{Q})$ efficiently without integer factorisation.

Note that

$$\Psi^{\mathrm{f}}(Q) = \sum_{n=0}^{\infty} 4^{-n-1} \log g_n,$$

where $g_n \in \mathbb{Z}$ is such that

$$\log g_n = \sum_p \varepsilon_p(2^n Q) \log p\,.$$

We can compute the numbers g_n by repeatedly applying the duplication map δ and taking gcds.

The algorithm

Theorem C. Let $Q \in A(\mathbb{Q})$. The following algorithm computes $\Psi^{\mathrm{f}}(Q)$ exactly (as a rational combination of logs) in time quasi-linear in the size of the input data.

(1) Compute bounds B, M and m, using the bounds on $\varepsilon_p(Q)$ and on $\mathrm{denom}(\mu_p(Q))$ for all bad primes p.

(2) Compute $g_0, \ldots, g_m$ by repeatedly applying δ, but mod Δ^{m+1}.

(3) Compute pairwise coprime integers $q_1, \ldots, q_s$ and $e_{i,n} \in \mathbb{Z}_{\geq 0}$ such that $g_n = \prod_{i=1}^{s} q_i^{e_{i,n}}$ for all n.

(4) For all $i \in \{1, \ldots, s\}$:
   (a) compute $a_i := \sum_{n=0}^{m} 4^{-n-1} e_{i,n}$,
   (b) let μ_i be the simplest fraction between a_i and $a_i + \frac{1}{B^2 M^2}$.

(5) Return $\sum_{i=1}^{s} \mu_i \log q_i$.

Computing archimedean correction functions

It remains to bound and compute

$$\Psi_\infty(Q) := \sum_{n=0}^{\infty} 4^{-(n+1)} \tilde{\varepsilon}_\infty(2^n Q)\,, \tag{2}$$

where

$$\begin{aligned}\tilde{\varepsilon}_\infty(Q) = &-\log \max\{|\delta_j(x_1,x_2,x_3,x_4)| : 1 \le j \le 4\} \\ &+ 4 \log \max\{|x_j| : 1 \le j \le 4\}\,,\end{aligned}$$

and $\kappa(Q) = (x_1 : x_2 : x_3 : x_4) \in K(\mathbb{R})$.

Once we have an upper bound γ_∞ for $\tilde{\varepsilon}_\infty$, we can use (2) to approximate Ψ_∞. This turns out to be quasi-quadratic in the number of correct bits of precision in the output.

Note that $\Psi_\infty \le \gamma_\infty/3$.

Bounding archimedean correction functions

Using representation theory, Stoll has found an upper bound

$$\frac{\max_j\{|x_j|\}^4}{\max_j\{|\delta_j(x_1,\ldots,x_4)|\}},$$

which gives an upper bound γ_∞ for $\tilde{\varepsilon}_\infty$.

For this, one computes quadratic forms $y_i = y_i(x_1, x_2, x_3, x_4)$ and real numbers a_{ji} and b_{ij} such that if $(x_1 : x_2 : x_3 : x_4) \in K(\mathbb{R})$, then

- $x_j^2 = \sum_i a_{ji} y_i(x_1, \ldots, x_4)$
- $y_i(x_1, \ldots, x_4)^2 = \sum_j b_{ij} \delta_j(x_1, \ldots, x_4)$.

We iterate this process to get a sequence $(b_n)_n$ in $\mathbb{R}^4_{\geq 0}$ such that

$$\Psi_\infty \leq \frac{4^n}{4^n - 1} \log \|b_n\|_\infty \quad \text{for all } n \geq 1,$$

leading to a tight upper bound on Ψ_∞ after a few iterations.

Enumeration

Recall that we also need to enumerate

$$\{P \in A(\mathbb{Q}) : h(P) \leq B + \beta\} \supset \{P \in A(\mathbb{Q}) : \hat{h}(P) \leq B\}$$

given $B \in \mathbb{R}$, where $h(P) = h(\kappa(P))$ is the naive height of P, and β is an upper bound for $h - \hat{h}$.

Idea. Use a different function h' with bounded difference from $\hat{h}$ such that

- the bound for $h' - \hat{h}$ is smaller than the bound for $h - \hat{h}$;
- the enumeration of all points of bounded h' is no more difficult than for h.

Optimizing the naive height

For a place v, set

$$|f|_v := \max\{|f_0|_v, \ldots, |f_6|_v\}.$$

For $Q \in A(\mathbb{Q})$ with $\kappa(Q) = (x_1 : x_2 : x_3 : x_4)$, we set

$$h'(Q) := \sum_v \log \max\{|x_1|_v, |x_2|_v, |x_3|_v, |x_4|_v/|f|_v\}$$

to give all Kummer coordinates roughly the same weight.

Slightly adapting the methods discussed above for bounding $h - \hat{h}$, we usually get a much smaller bound for $h' - \hat{h}$ than for $h - \hat{h}$.

For the enumeration, we use that

$$h(x_1 : x_2 : x_3) \leq h'(Q)\,.$$

Example: The record curve

Consider the curve X given by

$$y^2 = 82342800x^6 - 470135160x^5 + 52485681x^4 + 2396040466x^3 + 567207969x^2 - 985905640x + 247747600\,.$$

- $\#X(\mathbb{Q}) \geq 642$ (current record for genus 2, found by Stoll),
- $A = \mathrm{Jac}(X)$ has rank 22 (assuming GRH) and trivial torsion over $\mathbb{Q}$.
- Previous results due to Stoll give

$$h - \hat{h} < 40.1 + 7.7 = 47.8\,.$$

- We use a modified naive height h' and show

$$h' - \hat{h} < 20.43 + (-19.25) = 1.18\,.$$

- Using this smaller bound, we show that the differences of the rational points on X generate $A(\mathbb{Q})$.