Rational points on Jacobians of hyperelliptic curves II

Steffen Müller
Carl von Ossietzky Universität Oldenburg

September 4, 2014

Recap

Let

- k be a perfect field of characteristic $\neq 2$,
- C/k be a hyperelliptic curve of genus $g \geq 1$, given as the smooth projective model of the affine curve given by

$$y^2 = f(x), \quad \text{where}$$

  - $f \in k[x]$ has degree $2g+1$ or $2g+2$,
  - $\operatorname{disc}(f) \neq 0$.

The Jacobian J of C is an abelian variety of dimension g over k such that

$$J(k) \cong \operatorname{Pic}^0(C/k).$$

Mordell-Weil

Theorem. (Mordell-Weil) Let $k = \mathbb{Q}$. Then

$$J(\mathbb{Q}) \cong \mathbb{Z}^r \times J(\mathbb{Q})_{\mathrm{tors}},$$

where r is a nonnegative integer and the torsion subgroup $J(\mathbb{Q})_{\mathrm{tors}} \subset J(\mathbb{Q})$ is finite.

We call

- $J(\mathbb{Q})$ the Mordell-Weil group of $J/\mathbb{Q}$;
- r the rank of $J/\mathbb{Q}$.

Proof of Mordell-Weil: reminder

Recall the steps of the proof of the Mordell-Weil theorem:

(i) $J(\mathbb{Q})/2J(\mathbb{Q})$ is a finite group.

(ii) There is a quadratic form $\hat{h} : J(\mathbb{Q}) \to \mathbb{R}$ such that for all $B \in \mathbb{R}$ the set $\{P \in J(\mathbb{Q}) : \hat{h}(P) \leq B\}$ is finite.

(iii) (i) and (ii) imply the theorem.

We've already shown (i) and (iii) and discussed the computation of r.

Today we discuss

- how to prove (ii);
- how to compute generators of $J(\mathbb{Q})$, assuming that we know r.

We first discuss how to compute $J(\mathbb{Q})_{\mathrm{tors}}$ (this is the easy part).

Two-torsion

Let k be an extension field of $\mathbb{Q}$ and let $f = \prod_{i=1}^{m_k} f_i$ be the factorisation of f into irreducibles over k.

Let $J(k)[2] = \{P \in J(k) : 2P = 0\}$. Then we have

$$J(k)[2] \cong (\mathbb{Z}/2\mathbb{Z})^{m_k - 1}$$

if d is odd or if all irreducible factors of f have even degree and we have

$$J(k)[2] \cong (\mathbb{Z}/2\mathbb{Z})^{m_k - 2}$$

otherwise.

Moreover $J(k)[2]$ is generated by the points with Mumford representation $(f_i, 0)$.

Reduction I

In order to compute $J(\mathbb{Q})_{\mathrm{tors}}$, we'll use the concept of reduction.

Suppose

- p is an odd prime number,
- $C : y^2 = f(x)$ is defined over the p-adic numbers $\mathbb{Q}_p$,
- $f \in \mathbb{Z}_p[x]$ (w.l.o.g).

Let $\tilde{f}(x) \in \mathbb{F}_p[x]$ denote the polynomial obtained by reducing the coefficients of f modulo p.

Reduction II

The reduction of C modulo p is the curve over $\mathbb{F}_p$ given by

$$\tilde{C} : y^2 = \tilde{f}(x).$$

The reduction need not be hyperelliptic; it might be singular (and, in addition, reducible or non-reduced).

We say that C has good reduction if $\tilde{C}$ is nonsingular.

As $\operatorname{disc}(\tilde{f}) = \operatorname{disc}(f) \bmod p$, C has good reduction if and only if

$$p \nmid \operatorname{disc}(f).$$

Reduction III

Suppose that C has good reduction. We can define a reduction map

$$\mathrm{red} : C(\mathbb{Q}_p) \to \tilde{C}(\mathbb{F}_p)$$

in the obvious way and we extend this to a group homomorphism $\mathrm{red} : \mathrm{Div}(C/\mathbb{Q}_p) \to \mathrm{Div}(\tilde{C}/\mathbb{F}_p)$ by linearity. Then red

- commutes with the degree morphism and
- sends a principal divisor to a principal divisor.

Hence red induces a group homomorphism

$$\mathrm{red} : J(\mathbb{Q}_p) \to \tilde{J}(\mathbb{F}_p),$$

where $\tilde{J}$ is the Jacobian of $\tilde{C}$. Let $J(\mathbb{Q}_p)_1$ denote its kernel.

Torsion in the kernel of reduction

We have an exact sequence

$$0 \longrightarrow J(\mathbb{Q}_p)_1 \longrightarrow J(\mathbb{Q}_p) \xrightarrow{\mathrm{red}} \tilde{J}(\mathbb{F}_p) \longrightarrow 0,$$

so $J(\mathbb{Q}_p)_1$ has finite index in $J(\mathbb{Q}_p)$.

Why is this important for the study of torsion points?

Proposition. There is a group isomorphism $J(\mathbb{Q}_p)_1 \cong \mathbb{Z}_p^g$. In particular, the kernel of reduction is torsion-free.

Corollary. The reduction morphism $\mathrm{red} : J(\mathbb{Q}_p) \to \tilde{J}(\mathbb{F}_p)$ is injective on torsion.

Reduction of global torsion

Now suppose that $C : y^2 = f(x)$ is defined over $\mathbb{Q}$ such that

- $f \in \mathbb{Z}[x]$ (w.l.o.g.),
- p is a prime of good reduction, i.e. $p \nmid 2\,\mathrm{disc}(f)$.

We get a group homomorphism $\mathrm{red} : J(\mathbb{Q}) \hookrightarrow J(\mathbb{Q}_p) \to \tilde{J}(\mathbb{F}_p)$.

Corollary. The restriction of red to $J(\mathbb{Q})_{\mathrm{tors}}$ is injective. In particular, $J(\mathbb{Q})_{\mathrm{tors}}$ is finite and is isomorphic to a subgroup of $\tilde{J}(\mathbb{F}_p)$.

This gives us a method for computing the torsion subgroup of $J(\mathbb{Q})$, or at least narrowing down the number of possibilities, if we can compute $\tilde{J}(\mathbb{F}_p)$. Lots of algorithms exist for the latter problem.

Computing torsion I

Let $f(x) = x^7 + x^3 + 1$. Then

- $C : y^2 = f(x)$ is hyperelliptic of genus 3;
- $\mathrm{disc}(f) = 5 \cdot 41 \cdot 4051$, so $p = 3$ and $p = 7$ are primes of good reduction for C.

We compute

- $\#\tilde{J}(\mathbb{F}_3) = 36$
- $\#\tilde{J}(\mathbb{F}_7) = 229$.

Therefore $J(\mathbb{Q})_{\mathrm{tors}}$ is trivial.

Computing torsion II

Let $f(x) = x(x-2)(x+2)(x+3)(x+7)$.

- $C : y^2 = f(x)$ is hyperelliptic of genus 2;
- $\mathrm{disc}(f) = 2^{12} \cdot 3^6 \cdot 5^4 \cdot 7^2$, so $p = 11$ and $p = 13$ are primes of good reduction for C.

We compute

- $\#\tilde{J}(\mathbb{F}_{11}) = 2^4 \cdot 11$,
- $\#\tilde{J}(\mathbb{F}_{13}) = 2^4 \cdot 3 \cdot 5$.

Therefore we find $\#J(\mathbb{Q})_{\mathrm{tors}} \mid 16$.

But f factors completely over $\mathbb{Q}$, so $J(\mathbb{Q})[2] \cong (\mathbb{Z}/2\mathbb{Z})^4$.

Hence $J(\mathbb{Q})_{\mathrm{tors}} \cong (\mathbb{Z}/2\mathbb{Z})^4$.

Proof of Mordell-Weil: What's left?

To finish the proof of the Mordell-Weil theorem, we need to construct a quadratic form $\hat{h} : J(\mathbb{Q}) \to \mathbb{R}$ such that for all $B \in \mathbb{R}$ the set

$$\{P \in J(\mathbb{Q}) : \hat{h}(P) \leq B\}$$

is finite.

We're also going to look at computational aspects of $\hat{h}$. This will allow us to find generators of $J(\mathbb{Q})$, assuming we have generators of a subgroup of finite index (in particular, we know r).

For now, suppose C is defined over a perfect field of characteristic $\neq 2$.

The Kummer variety

Theorem. The quotient $\mathcal{K} = J/\{\pm 1\}$ is a projective variety over k.
It can be embedded into $\mathbb{P}^{2^g-1}$ over k.

We call $\mathcal{K}$ the Kummer variety of J.

Fix a morphism $\kappa : J \to \mathbb{P}^{2^g-1}$ over k such that

- κ factors through the canonical surjection $J \twoheadrightarrow \mathcal{K}$ and
- $\kappa(0) = (0 : \ldots : 0 : 1)$.

We then identify $\mathcal{K}$ with $\kappa(J)$.

Examples of Kummer varieties

Example. If $C = J$ is an elliptic curve, then $\mathcal{K} = \mathbb{P}^1$ and we can take $\kappa(x, y) = x$.

Example. Suppose that $g = 2$. Then the Kummer surface $\mathcal{K}$ is a classical object of 19th century geometry.

- $\mathcal{K}$ can be embedded as a singular quartic hypersurface into $\mathbb{P}^3$.
- Explicit formulas for κ and an equation for $\mathcal{K}$ are due to Flynn.
- "Pseudo-addition" on $\mathcal{K}$ (or an alternative embedding of $\mathcal{K}$ due to Gaudry) can be used for addition algorithms on J; useful in cryptography (Duquesne, Gaudry, Cosset).

Kummer variety in genus 3

Example. If $g = 3$, then the Kummer variety can be embedded into $\mathbb{P}^7$ as an intersection of a quadric and 34 quartics.

- Stubbs: Explicit embedding κ, some defining equations for $\mathcal{K}$
- M.: Complete set of equations for $\mathcal{K}$
- Stoll: Alternative formulas for κ and defining equations for $\mathcal{K}$

So far nobody has been brave enough to tackle the case $g = 4$.

Heights on projective space

Intuitively, a height function is supposed to measure the arithmetic complexity of a geometric object over $\mathbb{Q}$.

First define the height function $h : \mathbb{P}^N(\mathbb{Q}) \to \mathbb{R}_{\geq 0}$ by

$$h\left((X_0 : \ldots : X_N)\right) = \log \max\{|X_i| : i \in \{0, \ldots, N\}\},$$

where $X_0, \ldots, X_N$ are relatively prime integers.

So the height of a point $P \in \mathbb{P}^N(\mathbb{Q})$ tells us how much "space" is needed to write down coordinates of P.

Note that for every $B \in \mathbb{R}$ there are only finitely many points $P \in \mathbb{P}^N(\mathbb{Q})$ such that $h(P) \leq B$.

From now on, suppose that $C/\mathbb{Q} : y^2 = f(x)$, where $f \in \mathbb{Z}[x]$.

The naive height

The naive height $h : J(\mathbb{Q}) \to \mathbb{R}_{\geq 0}$ is defined by $h(P) := h(\kappa(P))$. It has the following properties.

- For every $B \in \mathbb{R}$ there are only finitely many points $P \in J(\mathbb{Q})$ such that $h(P) \leq B$.
- The function $P \mapsto h(2P) - 4h(P)$ is bounded on $J(\mathbb{Q})$.

So h is quadratic up to a bounded function.

Question (Néron). Is there a quadratic form on $J(\mathbb{Q})$ with bounded difference from h?

A quadratic form on $J(\mathbb{Q})$ is a function $q : J(\mathbb{Q}) \to \mathbb{R}$ such that

- $(P, Q) \mapsto q(P + Q) - q(P) - q(Q)$ is bilinear on $J(\mathbb{Q}) \times J(\mathbb{Q})$ and
- $q(-P) = q(P)$ for all $P \in J(\mathbb{Q})$.

Néron-Tate

Answer (Néron, Tate). Yes, there is: the canonical height $\hat{h}$.

Definition. (Tate) The canonical height of a point $P \in J(\mathbb{Q})$ is defined by

$$\hat{h}(P) := \lim_{n\to\infty} 4^{-n} h(2^n P).$$

We'll come back to (a variant of) Néron's construction later.

Properties of the canonical height

Theorem. (Néron, Tate) The canonical height has the following properties:

(i) The difference $h - \hat{h}$ is bounded.

(ii) $\hat{h}$ is a quadratic form on $J(\mathbb{Q})$.

(iii) For every $B \in \mathbb{R}$ there are only finitely many points $P \in J(\mathbb{Q})$ such that $\hat{h}(P) \leq B$.

(iv) $\hat{h}(P) \geq 0$, with equality if and only if P is torsion.

By (ii) and (iii), this finishes the proof of the Mordell-Weil theorem for $J(\mathbb{Q})$!

Properties of the canonical height

Theorem. (Néron, Tate) The canonical height has the following properties:

(i) The difference $h - \hat{h}$ is bounded.

Proof of (i). We have

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

but the function $P \mapsto h(2P) - 4h(P)$ is bounded.

Properties of the canonical height

Theorem. (Néron, Tate) The canonical height has the following properties:

(ii) $\hat{h}$ is a quadratic form on $J(\mathbb{Q})$.

Proof of (ii). One proves that the parallelogram law

$$\hat{h}(P+Q)+\hat{h}(P-Q)=2\hat{h}(P)+2\hat{h}(Q)$$

holds for all $P, Q \in J(\mathbb{Q})$ by first showing that it holds for h up to a bounded function and then applying the limit formula.

This suffices to prove that $\hat{h}$ is a quadratic form (exercise).

Properties of the canonical height

Theorem. (Néron, Tate) The canonical height has the following properties:

(iii) For every $B \in \mathbb{R}$ there are only finitely many points $P \in J(\mathbb{Q})$ such that $\hat{h}(P) \leq B$.

Proof of (iii). Follows from the corresponding statement for h, since $\hat{h} - h$ is bounded.

Properties of the canonical height

Theorem. (Néron, Tate) The canonical height has the following properties:

(iv) $\hat{h}(P) \geq 0$ with equality if and only if P is torsion.

Proof of (iv). Since h is nonnegative, the limit formula shows $\hat{h}$ is as well.

If $\hat{h}(P) = 0$ then $\hat{h}(nP) = 0$ for all $n \in \mathbb{Z}$. By (iii), $\{nP : n \in \mathbb{Z}\}$ is finite, so P has finite order.

Conversely, if $nP = 0$ for some $n \geq 1$, then

$$\hat{h}(P) = \frac{1}{n^2}\hat{h}(0) = 0.$$

Checking independence

Assume that we know

- generators of $J(\mathbb{Q})_{\mathrm{tors}}$,
- the rank r,
- nontorsion points $Q_1, \ldots, Q_r \in J(\mathbb{Q})$.

To check that $Q_1, \ldots, Q_r$ are independent, it suffices to show one of the following:

- Nonvanishing of their regulator $\mathrm{Reg}(Q_1, \ldots, Q_r) = \det \left(\frac{1}{2}(\hat{h}(Q_i + Q_j) - \hat{h}(Q_i) - \hat{h}(Q_j)) \right)_{1 \le i,j \le r}$
- Independence of their reductions $\tilde{Q}_1, \ldots, \tilde{Q}_r$ modulo a prime of good reduction
- Independence of their images under the two-descent map $\delta : J(\mathbb{Q}) \to H'$

Saturation

Assume that we know

- generators of $J(\mathbb{Q})_{\mathrm{tors}}$,
- the rank r,
- independent nontorsion points $Q_1, \ldots, Q_r \in J(\mathbb{Q})$.

Hence the group $\langle Q_1, \ldots, Q_r \rangle$ is a subgroup of finite index in $J(\mathbb{Q})/J(\mathbb{Q})_{\mathrm{tors}}$.

Problem. Find generators $P_1, \ldots, P_r$ of $J(\mathbb{Q})/J(\mathbb{Q})_{\mathrm{tors}}$.

Saturation II

Proposition. The canonical height can be extended $\mathbb{R}$-linearly to a positive definite quadratic form on $J(\mathbb{Q}) \otimes \mathbb{R}$.

- The group $J(\mathbb{Q})/J(\mathbb{Q})_{\text{tors}}$ sits as a lattice Λ in the r-dimensional Euclidean vector space $V = (J(\mathbb{Q}) \otimes \mathbb{R}, \hat{h})$.
- $Q_1, \ldots, Q_r$ generate a sublattice Λ' of Λ of finite index.

Goal. Find the saturation Λ of Λ' in V.

Saturation: Approach 1

One possible approach is due to Siksek (for $g = 1$, adapted by Flynn-Smart for $g = 2$):

(i) Compute an upper bound N for the index n of Λ' in Λ.
(ii) Check for every prime $p \leq N$ whether p divides n. If it does, enlarge Λ'.

For (i), we need to compute

- the regulator $\mathrm{Reg}(Q_1, \ldots, Q_r)$ and, for $i = 1, \ldots, r$:
- positive lower bounds M_i for the ith-successive minimum of Λ'.

Then the index n of Λ' in Λ satisfies

$$n \leq N := \sqrt{\frac{\mathrm{Reg}(Q_1, \ldots, Q_r)\gamma_r^r}{M_1 \cdots M_r}},$$

where γ_i is the ith Hermite constant.

Saturation: Approach 1

For step (ii), we check for every prime $p \leq N$ whether p divides the index n.

So we have to check whether there is a point $Q \in J(\mathbb{Q})$ such that $pQ \in G = \langle Q_1, \ldots, Q_r \rangle$, but $Q \notin G$.

- Find a set of primes q of good reduction such that p divides $\#\tilde{J}(\mathbb{F}_q)$.
- Try to show that the kernel of the natural map

$$G/pG \to \prod_q \tilde{J}(\mathbb{F}_q)/p\tilde{J}(\mathbb{F}_q)$$

  is trivial.
- If it is, then no such Q can exist and hence $p \nmid n$.
- If this doesn't work, try to find such a Q and start over with $G = \langle Q_1, \ldots, Q_r, Q \rangle$.

Saturation: Approach 2

The following alternative approach is due to Stoll: Let ρ denote the covering radius of Λ':

$$\rho = \max\{\sqrt{\|P - \Lambda'\|_{\hat{h}}} : P \in V\}$$

Then the ball in V of radius ρ^2 around the origin contains a fundamental domain for Λ and hence a set of generators. So it suffices to

- compute an upper bound $B \geq \rho^2$
- enumerate all points $P \in J(\mathbb{Q})$ such that $\hat{h}(P) \leq B$.

The covering radius can be computed from the Voronoi cell of Λ'. This becomes very slow when r is large.

Instead, one can split Λ' into orthogonal parts Λ'_1 and Λ'_2 of smaller rank; then

$$\rho^2 \leq \rho(\Lambda'_1)^2 + \rho(\Lambda'_2)^2.$$

Height algorithms

Note that both algorithms for saturation crucially rely on algorithms to

(i) compute $\hat{h}(P)$ for given $P \in J(\mathbb{Q})$;
(ii) enumerate $\{P \in J(\mathbb{Q}) : \hat{h}(P) \leq B\}$ for given $B \in \mathbb{R}$.

Assume that $g \leq 3$, so that we have explicit formulas for

- $\kappa : J \twoheadrightarrow \mathcal{K} \hookrightarrow \mathbb{P}^{2^g-1}$;
- defining equations for $\mathcal{K} \subset \mathbb{P}^{2^g-1}$.

Then we can

- compute $h(P) = h(\kappa(P))$ for $P \in J(\mathbb{Q})$;
- enumerate $\{P \in J(\mathbb{Q}) : h(P) \leq B'\}$ for given $B' \in \mathbb{R}$ by
  - enumerating $\{P \in \mathcal{K}(\mathbb{Q}) : h(P) \leq B'\}$ and
  - checking which of these lift to $J(\mathbb{Q})$.

The height difference

Since $\mu := h - \hat{h}$ is bounded, it suffices to have algorithms to

- compute $\mu(P)$ for a given $P \in J(\mathbb{Q})$,
- find an upper bound $\beta \geq \sup\{|\mu(P)| : P \in J(\mathbb{Q})\}$.

Then we can

- compute $\hat{h}(P)$ as $h(P) - \mu(P)$,
- enumerate the set

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

Idea. Measure locally how far away the naive height is from being a quadratic form, i.e. decompose the function $P \mapsto 4h(P) - h(2P)$ into a sum of local terms.

Duplication on the Kummer variety

Since h lives on $\mathcal{K}(\mathbb{Q})$, need to analyze duplication on $\mathcal{K}$.

Lemma. (Flynn, Stoll) There are homogeneous quartic polynomials $\delta_1, \ldots, \delta_{2^g} \in \mathbb{Z}[x_1, \ldots, x_{2^g}]$ without common nontrivial zero such that

- the map $\delta := (\delta_1, \ldots, \delta_{2^g}) : \mathbb{P}^{2^g-1} \to \mathbb{P}^{2^g-1}$ makes the following diagram commute:

$$\begin{tikzcd} J \arrow[r, "{[2]}"] \arrow[d, "\kappa"] & J \arrow[d, "\kappa"] \\ \mathcal{K} \arrow[r, "\delta"'] & \mathcal{K} \end{tikzcd}$$

- $\delta(0, \ldots, 0, 1) = (0, \ldots, 0, 1)$ (as a map on $\mathbb{A}^{2^g}$).

Local error functions I

Let v be a place of $\mathbb{Q}$. For a point $P \in J(\mathbb{Q}_v)$ such that $\kappa(P) = (\xi_1 : \ldots : \xi_{2^g})$, we define

$$\begin{aligned} \varepsilon_v(P) = &-\log \max\{|\delta_j(\xi_1, \ldots, \xi_{2^g})|_v : 1 \le j \le 2^g\} \\ &+ 4 \log \max\{|\xi_j|_v : 1 \le j \le 2^g\}. \end{aligned}$$

Then

- ε_v is bounded.
- If p is a prime, then ε_p is nonnegative.
- If p is a prime such that C has good reduction at p, then ε_p vanishes identically on $J(\mathbb{Q}_p)$.
- If $P \in J(\mathbb{Q})$, then

$$4h(P) - h(2P) = \sum_v \varepsilon_v(P).$$

Local error functions II

Hence we can define

$$\mu_v(P) = \sum_{n=0}^{\infty} 4^{-(n+1)} \varepsilon_v(2^n P) \text{ for } P \in J(\mathbb{Q}_v).$$

The properties of ε_v imply the following properties of μ_v:

- The function μ_v is bounded.
- If p is a prime, then μ_p is nonnegative.
- We have $4\mu_v(P) - \mu_v(2P) = \varepsilon_v(P)$ for all $P \in J(\mathbb{Q}_v)$.
- If p is a prime such that C has good reduction at p, then μ_p vanishes identically on $J(\mathbb{Q}_p)$.

Decomposing the height difference I

Proposition. If $P \in J(\mathbb{Q})$, then

$$\hat{h}(P) = h(P) - \sum_v \mu_v(P).$$

Proof.

$$\begin{aligned}
\hat{h}(P) &= \lim_{n\to\infty} 4^{-n} h(2^n P) \\
&= h(P) + \sum_{n=0}^{\infty} 4^{-(n+1)} \left(h(2^{n+1}P) - 4h(2^n P)\right) \\
&= h(P) - \sum_{n=0}^{\infty} 4^{-(n+1)} \sum_v \varepsilon_v(2^n P) \\
&= h(P) - \sum_v \mu_v(P)
\end{aligned}$$

Decomposing the height difference II

Recall that we need algorithms to

- compute $\mu(P)$ for a given $P \in J(\mathbb{Q})$,
- find an upper bound $\beta \geq \sup\{|\mu(P)| : P \in J(\mathbb{Q})\}$,

where $\mu = h - \hat{h} = \sum_v \mu_v$.

Hence it suffices to have algorithms to

- compute $\mu_v(P)$ for a given $P \in J(\mathbb{Q}_v)$,
- find an upper bound $\beta_v \geq \sup\{|\mu_v(P)| : P \in J(\mathbb{Q}_v)\}$

for all places v of $\mathbb{Q}$.

We first discuss how to bound $\sup\{|\mu_v(P)| : P \in J(\mathbb{Q}_v)\}$.

Bounding the height difference I

Theorem. (Stoll) If $g \leq 3$, and p is a prime number, then

$$0 \leq \varepsilon_p(P) \leq -\log|2^{2g}\,\mathrm{disc}(f)|_p \text{ for } P \in J(\mathbb{Q}_p).$$

The proof uses representation-theoretic methods.
There are several ways to improve on this; with enough effort, can get an optimal bound for $\varepsilon_v(P)$ on $\mathcal{K}(\mathbb{Q}_v)$ if $v \neq \infty$.
Using similar techniques, Stoll also gives a method for bounding $|\varepsilon_\infty|$.

Corollary. If $g \leq 3$, and p is a prime number, then we have

$$0 \leq \mu_v(P) \leq -\frac{\log|2^{2g}\,\mathrm{disc}(f))|_v}{3} \text{ for } P \in J(\mathbb{Q}_v).$$

For $g = 1$ and p prime, there are optimal bounds for $\mu_p(P)$ (and very good bounds for $\mu_\infty(P)$) due to Cremona-Prickett-Siksek.

Bounding the height difference II

We generally expect that for most "large" primes dividing $\operatorname{disc}(f)$, their multiplicity is 1.

Proposition. (Stoll) Suppose that $g \leq 3$ and that p is an odd prime such that $\operatorname{ord}_p(\operatorname{disc}(f)) \leq 1$. Then μ_p is identically zero.

Recently, Stoll and I have found optimal bounds for the most frequent reduction types for $g = 2$. We also prove

Theorem. (M.-Stoll) If $g = 2$, then we have

$$0 \leq \mu_p(P) \leq -\frac{\log |2^8 \operatorname{disc}(f))|_p}{4} \quad \text{for} \quad P \in J(\mathbb{Q}_p).$$

Stoll has also found a method that leads to major improvements over the known methods for bounding $|\mu_\infty|$.

Computing $\mu_\infty(P)$

Assuming we have a bound for $\varepsilon_\infty(P)$, we can compute $\mu_\infty(P)$ from the definition

$$\mu_\infty(P) = \sum_{n=0}^{\infty} 4^{-(n+1)} \varepsilon_\infty(2^n P)$$

to any desired accuracy using floating-point arithmetic.

If C is an elliptic curve, better algorithms exist (due to Tate, Silverman and Bost-Mestre).

The “kernel” of μ_p

We also need to compute $\mu_p(P)$ for $P \in J(\mathbb{Q}_p)$ and p a prime such that $\mathrm{ord}_p(\mathrm{disc}(f)) \geq 2$.

Theorem. (Néron, Stoll) Suppose that $g \leq 3$ and let $U = \{P \in J(\mathbb{Q}_p) : \mu_p(P) = 0\}$. Then

- U is a subgroup of $J(\mathbb{Q}_p)$ of finite index;
- μ_p factors through the quotient $J(\mathbb{Q}_p)/U$.

For elliptic curves, U is the connected component of the identity of the Néron model.

This leads to explicit formulas for $\mu_p(P)$, due to Néron and Silverman.

For $g > 1$, the relation between U and the connected component is more complicated.

Computing μ_p I

The following algorithm is due to Stoll, building on earlier work of Flynn-Smart. Let

- $m = \min\{n \geq 1 : \varepsilon_p(nP) = 0\}$,
- $m = 2^r s$, s odd,
- t be the order of 2 in $(\mathbb{Z}/s\mathbb{Z})^\times$.

Then we have $\varepsilon_p(2^{n+t}P) = \varepsilon_p(2^n P)$ for $n \geq r$ and hence (exercise)

$$\mu_p(P) = \sum_{n=0}^{\infty} \frac{\varepsilon_p(2^n P)}{4^{n+1}} = \sum_{n=0}^{r-1} \frac{\varepsilon_p(2^n P)}{4^{n+1}} + \frac{4^{-r-1}}{1-4^{-t}} \sum_{n=0}^{t-1} \frac{\varepsilon_p(2^{r+n} P)}{4^n}.$$

- We need to compute nP for $n \leq m$ and $\varepsilon_p(2^n P)$ for $n \leq r+t-1$.
- All computations can be done p-adically.

Computing μ_p II

For $g = 2$, there is the following algorithm (Stoll-M.):

- Set $B = \mathrm{ord}_p(2^4\,\mathrm{disc}(f))$ ($\geq \varepsilon_p(P)/\log p$).
- Set $M = 2\max\left\{16, \left\lfloor(\mathrm{ord}_p(2^8) + B)/3\right\rfloor\right\}$; then $\mu_p(P)/\log p \in \mathbb{Q}$ has denominator $\leq M$.
- Set $m = \lfloor\log(BM^2/3)/\log(4)\rfloor$.
- Compute
$$\mu_0 = 4^{-m-1}\,\mathrm{ord}_p(\delta^{\circ(m+1)}(\xi)), \quad \text{where}$$
  - $\xi = (\xi_1, \ldots, \xi_4) \in \mathbb{Z}_p^4$ represents $\kappa(P) \in \mathcal{K}(\mathbb{Q}_p)$ and
  - at least one ξ_i is a p-adic unit.
- Return $\mu_1 \cdot \log p$, where μ_1 is the unique fraction with denominator at most M in the interval $[\mu_0, \mu_0 + 1/M^2]$.

This only needs $\mathcal{O}(\log(\mathrm{ord}_p(\mathrm{disc}(f)^3)))$ steps.

Example

Let $C : y^2 = f(x)$, where $f = x(x-2)(x+2)(x+3)(x+7)$.
We already know

- $J(\mathbb{Q})_{\mathrm{tors}} = J(\mathbb{Q})[2] \cong (\mathbb{Z}/2\mathbb{Z})^4$.
- The points $[(e_i, 0) - \infty]$ generate $J(\mathbb{Q})_{\mathrm{tors}}$, where the e_i are the zeroes of f.
- $J(\mathbb{Q})$ has rank 2.
- The points $Q_1 = [(-1, 6) - \infty]$ and $Q_2 = [(-4, 12) - \infty]$ generate a subgroup of $J(\mathbb{Q})/J(\mathbb{Q})_{\mathrm{tors}}$ of finite index.

Goal: Compute generators of $J(\mathbb{Q})$.

- It remains to saturate $\langle Q_1, Q_2 \rangle$.
- We use the second approach discussed above (due to Stoll).

Example: height difference bound

- First set up the Euclidean vector space $V = (J(\mathbb{Q}) \otimes \mathbb{R}, \hat{h})$.
- Then saturate the lattice $\Lambda' = \langle Q_1, Q_2 \rangle$ inside V:

We go through the following steps:

- Compute the square of the covering radius

$$\rho^2 = \frac{\hat{h}(Q_1)\hat{h}(Q_2)\hat{h}(Q_1 - Q_2)}{4\,\mathrm{Reg}(Q_1, Q_2)} \approx 0.31402537597.$$

- Compute an upper bound $\beta \geq \sup\{|h(P) - \hat{h}(P)| : P \in J(\mathbb{Q})\}$:
  - The results of Stoll give $\beta \approx 9.991786718$;
  - our improved methods give $\beta \approx 8.177347056$ (modulo a not-yet-completely-proved lemma).

Example: Saturation

We enumerate $\{P \in J(\mathbb{Q}) : h(P) \le \rho^2 + \beta\}$. This took

- about 200 seconds using the first value of β,
- less than 0.1 seconds using the second value.

So small improvements in β can lead to major savings in the enumeration step (because the running time of the latter is exponential in the search bound).

It turns out that Λ' is already saturated, i.e.

$$J(\mathbb{Q})/J(\mathbb{Q})_{\mathrm{tors}} = \langle Q_1, Q_2 \rangle,$$

so we have computed a full set of generators for the Mordell-Weil group $J(\mathbb{Q})$!

Higher genus

What if $g \geq 4$?

There is an algorithm to compute $\hat{h}(P)$ for arbitrary g using arithmetic intersection theory that is practical for $g \leq 10$ (Holmes, M.).
One needs to compute

- a regular model of $C/\mathbb{Z}$,
- Gröbner bases over $\mathbb{Z}_p$,
- theta functions on $J(\mathbb{C})$.

Enumerating points of bounded height is much harder; we can't even compute the naive height!

A first step in this direction, also based on arithmetic intersection theory, is due to Holmes.