

Computation of BSD invariants in higher genus

Raymond van Bommel

Massachusetts Institute of Technology

Simons Collaboration on Arithmetic Geometry, Number Theory, and Computation

11 December 2020



Computation of BSD invariants in higher genus

Raymond van Bommel

Massachusetts Institute of Technology


Simons Collaboration on Arithmetic Geometry, Number Theory, and Computation

These slides can be downloaded at
`raymondvanbommel.nl/talks/Leiden2020.pdf`



Elliptic curves

An elliptic curve over \mathbb{Q} is a curve in 2-dimensional space given by an equation of the shape $y^2 = f(x)$, where $f(x)$ is a polynomial of degree 3. The points $E(\mathbb{Q})$ are the solutions with $x, y \in \mathbb{Q}$ to this equation together with a point at infinity. There is an addition law defined on this set of points, which gives $E(\mathbb{Q})$ the structure of an abelian group.



It is known that $E(\mathbb{Q})$ is a finitely generated group, hence it is of the shape $\mathbb{Z}^r \times T$, where T is a finite group. The integer r is called the *algebraic rank* of E .

Elliptic curves (continued)

If p is a prime number, and the denominators of the coefficients of $f(x)$ do not contain a factor p , then you could count the number points N_p of the equation $y^2 = f(x)$ with $x, y \in \mathbb{F}_p$. Here we need to again add a point at infinity.

In the 60s, Peter Swinnerton-Dyer used a computer to calculate this number of points for a lot of primes p and curves E . He observed that the following asymptotic law seems to hold

$$\prod_{p \leq x} \frac{N_p}{p} \sim \log(x)^r \quad \text{as } x \rightarrow \infty.$$

This eventually led to a conjecture, which is now known as the Birch and Swinnerton-Dyer conjecture. This conjecture predicts equality between the rank of E and order of vanishing of the L -function $L(E, s)$ at $s = 1$.

Generalised Birch and Swinnerton-Dyer conjecture

Tate has generalised the Birch and Swinnerton-Dyer conjecture to abelian varieties over number fields. For this talk, we consider the case where J is the Jacobian of a curve C over \mathbb{Q} . Then the conjecture links:

- the special value of the L -function of J ,
- the real period Ω ,
- the regulator R ,
- the Tamagawa numbers c_p ,
- the size of $J(\mathbb{Q})_{\text{tors}}$,
- the (algebraic) rank r of $J(\mathbb{Q})$, and
- the size of the Tate-Shafarevich group $\text{III}(J)$,

through the formula:
$$\lim_{s \rightarrow 1} (s-1)^{-r} L(J, s) = \frac{\Omega \cdot R \cdot |\text{III}(J)| \cdot \prod_p c_p}{|J(\mathbb{Q})_{\text{tors}}|^2}$$

Weil conjectures for curves

Let C_p be a projective curve over \mathbb{F}_p . Then we can define its zeta function

$$Z_{C_p}(T) = \exp \left(\sum_{n \geq 1} \frac{|C_p(\mathbb{F}_{p^n})|}{n} T^n \right) \in \mathbb{Q}[[T]].$$

Theorem (Weil conjectures)

There is a polynomial $L_{C_p} \in \mathbb{Z}[T]$ such that $Z_{C_p}(T) = \frac{L_{C_p}(T)}{(1-T)(1-qT)}$.
Moreover,

$$L_{C_p}(T) = \prod_{i=1}^{\text{genus}(C_p)} (1 - \alpha_i T) \left(1 - \frac{p}{\alpha_i} T \right),$$

for certain $\alpha_i \in \mathbb{C}$ with $|\alpha_i| = \sqrt{p}$.

In Deligne's original proof, it turns out that $Z_{C_p}(T)$ is the characteristic polynomial of the action of $1 - \text{Frob} \cdot T$ on the ℓ -adic cohomology group $H^1(C_p, \mathbb{Q}_\ell)$.

L-function of a curve over \mathbb{Q}

Let C be a smooth projective curve over \mathbb{Q} now. Then the L -function of Jac is defined as the Euler product

$$L(s) = \prod_{p \text{ prime}} L_p(p^{-s})^{-1}.$$

If p is a prime of good reduction for C , i.e. if the reduction C_p/\mathbb{F}_p of C is smooth, then the local factor $L_p(T)$ is exactly defined as $L_{C_p}(T)$ as in the previous slide. For primes of bad reduction, we need to look at the action of Frobenius on dual of the Tate-module

$$V_\ell(\text{Jac}(C)) = \mathbb{Q}_\ell \otimes_{\mathbb{Z}_\ell} (\lim A[\ell^n](\overline{\mathbb{Q}})).$$

It is known that $L(s)$ defines a holomorphic function on the part of the complex plane with $\text{Im}(s) > \frac{3}{2}$.

Conjecture

The function $L(s)$ extends to an entire function and satisfies a functional equation relating $L(s)$ with $L(2-s)$.

Computing the zeta function

Computing the zeta function boils down to counting points on $C_p(\mathbb{F}_{p^n})$. It suffices to do this only for $n = 1, \dots, \text{genus}(C_p)$. For fixed genus, the best algorithms to do this run in time polynomial in $\log p$:

- Pila generalised Schoof's point counting algorithm for elliptic curves to the case of abelian varieties. This algorithm runs in polynomial time, but is not practical at all, due to large constants in the runtime.
- Harvey and Sutherland developed an algorithm to count points on superelliptic curves $y^m = f(x)$ in average polynomial time. This means that they can compute the point counts for all $p \leq x$ at the same time, taking $\log(x)$ time per prime on average. The idea is to compute the coefficients of the Cartier-Manin matrix, which for the case $m = 2$ are certain coefficients of the polynomial $f(x)^{\frac{p-1}{2}} \bmod p$.

In practice, if one wants to compute the point counts for just one prime p , it is better to resort to one of the older exponential time algorithms using the p -adic cohomology.

Evaluating the L -function

Common strategy: find the local factors L_p for all good primes p up to a certain bound, and then use the conjectural functional equation to guess the L_p for the bad primes.

Of course, the Euler product does not converge for $s = 1$, and the function is not even known to extend analytically to a neighbourhood of $s = 1$. However, if we assume the function extends, then one can use a trick by Tim Dokchitser involving the Mellin transform to compute the special value. The Mellin transform of a function $f(t)$ is

$$\mathcal{M}(f)(s) = \int_0^\infty t^s f(t) \frac{dt}{t}.$$

A problem that occurs is that with any numerical computation, we could never prove that $L_p(1) = 0$. In particular, the order of vanishing cannot be computed provably in many cases.

Recent work: Costa and Platt implemented a calculator for L -functions, which uses arbitrary-precision ball arithmetic to give provable error bounds, see github.com/edgarcosta/lfunctions.

Regular models

There are different ways to write down equations with coefficients in \mathbb{Z} for the curve C over \mathbb{Q} . If there is a way such that the same equations over \mathbb{F}_p define a smooth curve, then C has *good reduction* at p . For finitely many primes p this is not possible, and then C has *bad reduction* at p . For these primes, we can still find a *regular model*.

Example

Consider the affine curve E with the following equation over \mathbb{Z} :

$$y^2 = x^3 - 2.$$

The point $(0, 0) \in \mathbb{F}_2^2$ is not smooth. This point corresponds to the maximal ideal $\mathfrak{m} = (2, x, y)$ in the ring $\mathbb{Z}[x, y]$. The equation for E is contained in \mathfrak{m} , but not in \mathfrak{m}^2 . This means that the tangent space of this model at the point $(0, 0) \in \mathbb{F}_2^2$ still has the expected dimension.

Regular models (continued)

Definition (regular model)

A *regular model* of C over $\mathbb{Z}_{(p)}$ is a flat proper regular scheme $\mathcal{C}/\mathbb{Z}_{(p)}$ together with an isomorphism $\mathcal{C}_{\mathbb{Q}} \cong C$. Here regular means that the tangent spaces have the expected dimension everywhere.

Theorem (Lipman)

For a smooth curve over \mathbb{Q} and any prime p there exists a regular model over $\mathbb{Z}_{(p)}$.

In practice, we can obtain such a regular model as follows. We start with any proper flat model \mathcal{C} of C , and then we repeatedly blow-up the non-regular subschemes of \mathcal{C} . The function `RegularModel` in Magma, implemented by Donnelly, does this computation.

There are also other ways to obtain regular models, e.g. using the cluster picture machinery for hyperelliptic curves due to Dokchitser, Dokchitser, Maistret and Morgan.

Tamagawa numbers

Definition (Tamagawa number)

Let \mathcal{J} be a Néron model of $\text{Jac}(C)$ over $\mathbb{Z}_{(p)}$. Then we define the Tamagawa number c_p of $\text{Jac}(C)$ at p to be the number of \mathbb{F}_p -points of the component group scheme $\mathcal{J}_{\mathbb{F}_p}/\mathcal{J}_{\mathbb{F}_p}^0$

The Tamagawa number can only be non-trivial for primes of bad reduction. In this case, we can a regular model \mathcal{C} to compute them.

Theorem (Raynaud)

Let I be the set of components of $\mathcal{C}_{\overline{\mathbb{F}_p}}$. Then there are certain linear maps $\alpha: \mathbb{Z}^I \rightarrow \mathbb{Z}^I$ and $\beta: \mathbb{Z}^I \rightarrow \mathbb{Z}$, such that $\text{im } \alpha$ and $\ker \beta$ fit in an exact sequence of $\text{Gal}(\overline{\mathbb{F}_p}/\mathbb{F}_p)$ -modules

$$0 \rightarrow \text{im } \alpha \rightarrow \ker \beta \rightarrow (\mathcal{J}_{\mathbb{F}_p}/\mathcal{J}_{\mathbb{F}_p}^0)(\overline{\mathbb{F}_p}) \rightarrow 0.$$

The maps α and β , and the map from $\text{im } \alpha$ to $\ker \beta$ can be expressed explicitly in terms of the multiplicity and intersection numbers of the components in I .

Implementation

This method has been implemented by the speaker in Magma, on top of the existing regular model package. The computation has been done for all 66 158 curves of genus 2 in the LMFDB. This took about 2.02 seconds on average, and for the worst curve this took about 1600 seconds.

The main difficulty turned out to be the following:

Problem


Given an integer n , construct a number field K of degree n such that:

- *for every $d \mid n$, there is a subfield $K_d \subset K$ of absolute degree d ;*
- *the ring of integers of K can be computed in a practical way.*

For more details, see my preprint [arXiv:2002.04667](https://arxiv.org/abs/2002.04667).

Real period

Over \mathbb{C} , the genus g curve C is a ball with g handles. Its homology group $H_1(C, \mathbb{Z})$ is generated by $2g$ elements $\gamma_1, \dots, \gamma_{2g}$.



On the other hand, if we take a basis $\omega_1, \dots, \omega_g \in \Omega^1(C, \mathbb{Q})$ of holomorphic differentials, we can compute the period matrix

$$\begin{pmatrix} \int_{\gamma_1} \omega_1 & \dots & \int_{\gamma_1} \omega_g & \int_{\gamma_1} \overline{\omega_1} & \dots & \int_{\gamma_1} \overline{\omega_g} \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \int_{\gamma_{2g}} \omega_1 & \dots & \int_{\gamma_{2g}} \omega_g & \int_{\gamma_{2g}} \overline{\omega_1} & \dots & \int_{\gamma_{2g}} \overline{\omega_g} \end{pmatrix}.$$

Real period (continued)

Now we want to choose the basis $\omega_1, \dots, \omega_g$ such that it respects the arithmetic structure on C . If we take a regular model $\mathcal{C}/\mathbb{Z}_{(p)}$ of C , we want to be able to view the differentials as differentials on \mathcal{C} and, for example, reduce them to the special fibre. We also don't want any differential to vanish identically on the special fibre.

More precisely, for each prime p suppose the basis $\omega_1, \dots, \omega_g$ is the restriction to C of a basis of $\omega_{\mathcal{C}/\mathbb{Z}_{(p)}}(\mathcal{C})$.

Definition (real period)

If we take such a basis $\omega_1, \dots, \omega_g$, then *real period* of C is the absolute value of the determinant of the period matrix multiplied with the number of connected components of $J(\mathbb{R})$.

Example

For an elliptic curve with short Weierstraß equation $y^2 = x^3 + ax + b$, we can always take the differential $\frac{dx}{2y}$, and integrate it along a basis of homology.

Implementation

Methods for numerical integration of the differentials have been implemented by several people. For example, by Van Wamelen in Magma for hyperelliptic curves, by Neurohr in Magma for general curves, and by Zotine, N. Bruin and Sijsling in SageMath.

The speaker implemented a method to compute the right basis of differentials in Magma, on top of the existing regular model package. this has been used to compute real periods for all 66 158 curves of genus 2 in the LMFDB. On average this took 1.67 seconds per curve, and about 6668 seconds in the hardest case.

Main difficulty: the regular model that we obtain for C can be quite complicated. For example, for one curve we obtained a patch for which we needed an equation in 3 variables of degree 56 with 428 terms. We could not compute a Gröbner basis for this patch within 24 hours. The solution was to use a number of tricks to circumvent the most difficult Gröbner basis computations.

Torsors and Sha

Definition (torsor)

A torsor of $J := \text{Jac}(C)$ is a scheme T/\mathbb{Q} together with an action of J on T , such that $T_{\overline{\mathbb{Q}}} \cong J_{\overline{\mathbb{Q}}}$ as a scheme with an action of $J_{\overline{\mathbb{Q}}}$.

Definition (Tate-Shafarevich group)

The *Tate-Shafarevich group* of J is defined as

$$\text{III}_J = \{\text{torsors } T \text{ of } J \mid T(\mathbb{Q}_p) \neq \emptyset \neq T(\mathbb{R})\} / \cong$$

Example

The genus 1 curve given by $3X^3 + 4Y^3 + 5Z^3 = 0$ has no \mathbb{Q} -points, even though it has points over \mathbb{R} and \mathbb{Q}_p for any p . Hence, it is a non-trivial torsor of its Jacobian, and it gives rise to a non-trivial element of the Tate-Shafarevich group.

It is conjectured that III_J is finite, and we can only compute it in a very limited number of cases.

Descent

Example

Consider smooth projective curves given by

$$E: y^2 = (x - \alpha)(x - \beta)(x - \gamma),$$

$$E_{a,b} := \{ay_1^2 = x - \alpha, by_2^2 = x - \beta, \frac{1}{ab}y_3^2 = x - \gamma\}.$$

For each pair of squarefree integers (a, b) , there is an degree 4 cover map $\varphi_{a,b}: E_{a,b} \rightarrow E$ given by $(x, y_1, y_2, y_3) \mapsto (x, y_1y_2y_3)$. These covering maps are in fact twists of each other (i.e. they become isomorphic over $\overline{\mathbb{Q}}$). In fact, they are twists of the multiplication-by-2 map $E \rightarrow E$.

The maps have a special property: each rational points in $E(\mathbb{Q})$ lies in the image of $E_{a,b}(\mathbb{Q})$ for exactly one pair of squarefree integers (a, b) .

The curve $E_{a,b}$ only has points everywhere locally (i.e. a point over \mathbb{R} and a point over \mathbb{Q}_p for every p) for an explicitly computable finite set of pairs (a, b) .

Selmer groups

Definition (2-Selmer group)

The *2-Selmer group* of $J := \text{Jac}(C)$ is defined by

$$\text{Sel}_2(J) = \{\text{twists } T \rightarrow J \text{ of } J \xrightarrow{2} J \mid T(\mathbb{Q}_p) \neq \emptyset \neq T(\mathbb{R})\} / \cong$$

Lemma

The 2-Selmer group is (in theory) computable, and fits into an exact sequence

$$0 \rightarrow J(\mathbb{Q})/2J(\mathbb{Q}) \rightarrow \text{Sel}_2(J) \rightarrow \text{III}_J[2] \rightarrow 0.$$

The lemma tells us that Selmer groups can be used to find upper bounds for the rank of J . These upper bounds are not always sharp, depending on $\text{III}_J[2]$.

Implementation

In case the upper bounds are not sharp, one could try the following strategies.

The Weil restriction

$$\mathrm{Res}_{\mathbb{Q}}^{\mathbb{Q}(\sqrt{d})}(J_{\mathbb{Q}(\sqrt{q})})$$

is isomorphic to $J \times J_d$, where d is a quadratic twist of J . Sometimes a lower bound for $\mathrm{rk}(J_d)$ and an upper bound for $\mathrm{rk}(J_{\mathbb{Q}(\sqrt{d})})$ can give a better upper bound for $\mathrm{rk}(J)$. This method might also provide information about $\mathrm{III}_J[2]$. This idea is due to Cremona and Mazur.

Sometimes applying these methods to abelian varieties that are isogenous to J may provide a better upper bound.

These methods have been implemented in Magma by Stoll, and have been used to compute ranks for all 66 158 curves of genus 2 in the LMFDB.

Néron-Tate height

Identify each point of J with its inverse to obtain the so-called Kummer variety $K = J/\pm$ associated to J . Let Θ be a Theta divisor on J . Then 2Θ descends to a very ample divisor on K , with an associated closed embedding $\iota: K \hookrightarrow \mathbb{P}^{2g-1}$, where g is the genus of C .

Definition (Néron-Tate height)

We define a naive height $h_{\text{naive}}(x) = \log(\max(|x_1|, \dots, |x_{2g}|))$, where $(x_1 : \dots : x_{2g})$ are primitive integer coordinates for $\iota(x)$.

The *Néron-Tate height* is then defined by:

$$h_{\text{NT}}(x) = \lim_{n \rightarrow \infty} \frac{1}{n^2} h_{\text{naive}}(nx) \quad \text{for } x \in J(\mathbb{Q}).$$

Remark. *It is not practical to compute the Néron-Tate height using this definition for curves of genus greater than 2.*

There are theorems giving an upper bound for $|h_{\text{naive}}(x) - h_{\text{NT}}(x)|$. These bounds can be used to effectively search for all points up to a certain height.

Regulator

Definition (height pairing)

For $x, y \in J(\mathbb{Q})$ we define

$$\langle x, y \rangle := \frac{1}{2}(h_{\text{NT}}(x + y) - h_{\text{NT}}(x) - h_{\text{NT}}(y)),$$

where h_{NT} is the Néron-Tate height on $J(\mathbb{Q})$.

This defines a non-degenerate bilinear form

$$\langle \cdot, \cdot \rangle : J(\mathbb{Q})/J(\mathbb{Q})_{\text{tors}} \times J(\mathbb{Q})/J(\mathbb{Q})_{\text{tors}} \rightarrow \mathbb{R}.$$

Definition (regulator)

Let r be the rank of $J(\mathbb{Q})$. If $x_1, \dots, x_r \in J(\mathbb{Q})$ are generators of the free part of $J(\mathbb{Q})$, then the *regulator* of $J(\mathbb{Q})$ is defined as

$$\left| \det \begin{pmatrix} \langle x_1, x_1 \rangle & \langle x_1, x_2 \rangle & \dots & \langle x_1, x_r \rangle \\ \vdots & \vdots & \ddots & \vdots \\ \langle x_r, x_1 \rangle & \langle x_r, x_2 \rangle & \dots & \langle x_r, x_r \rangle \end{pmatrix} \right|.$$

Computation of the Néron-Tate height

For the computation, we use a decomposition into local heights.

Theorem (Faltings (1984), Hriljac (1985))

Let D and E be divisors on C of degree 0, with disjoint support. Then

$$h_{\text{NT}}([D], [E]) = - \sum_v \langle D, E \rangle_v,$$

where we sum over all places, finite and infinite, of \mathbb{Q} .

For the finite primes, the local height $\langle D, E \rangle_p$ is defined in terms of intersection theory on a regular model \mathcal{C} of C over $\mathbb{Z}_{(p)}$.

For the infinite prime, we need to compute Riemann surface \mathbb{C}^g/Λ representing J together with an Abel-Jacobi map from $C(\mathbb{C})$ to \mathbb{C}^g/Λ . To compute $\langle D, E \rangle_\infty$, one then has to evaluate a number of Jacobi theta functions on this Riemann surface.


This has been implemented in Magma by Holmes, Müller and the speaker, building on a Magma package by Neurohr.

Computation of the regulator

The main challenge now is to find generators for the free part of $J(\mathbb{Q})$.
The general strategy is:

1. Search for points until they generate a finite index subgroup of the Mordell-weil group $J(\mathbb{Q})/J(\mathbb{Q})_{\text{tors}}$ (using the rank bounds).
2. For small primes p , calculate the saturation of the subgroup at p .
I.e. we enlarge our subgroup such that we know that the index inside $J(\mathbb{Q})/J(\mathbb{Q})_{\text{tors}}$ does not contain a factor p anymore.
3. If the subgroup is still not complete, we can find an upper bound for the height of the missing points, using the fact that the index the subgroup in $J(\mathbb{Q})/J(\mathbb{Q})_{\text{tors}}$ must be big now. Using an exhaustive search on points, we can finish the computation.

The methods have been implemented by Stoll in Magma. For some of the curves of genus 2 in the LMFDB, we ended up looking for points in an isogenous Jacobian. This turned out to be easier for the hardest cases. In the worst case, curve 900617.a.900617.1, we ended up with a generator of height 65.12498.


Δ → Genus 2 curves → Q → 900617 → a → 900617 → 1
Feedback · Hide Menu

Genus 2 curve 900617.a.900617.1

Introduction

Overview Random
Universe Knowledge

L-functions

Degree 1 Degree 2
Degree 3 Degree 4
ζ zeros

Modular forms

Classical Maass
Hilbert Bianchi

Varieties

Elliptic curves over Q
Elliptic curves over Q(α)
Genus 2 curves over Q
Higher genus families
Abelian varieties over \mathbb{F}_q

Fields

Number fields
p-adic fields

Representations

Dirichlet characters
Artin representations

Groups

Galois groups
Sato-Tate groups

Minimal equation

$$y^2 + (x^2 + \alpha)y = x^5 - 65x^4 + 224x^3 + 30x^2 + x \quad (\text{homogenize, simplify})$$

Invariants

Conductor: $N = 900617 = 37 \cdot 101 \cdot 241$
 Discriminant: $\Delta = 900617 = 37 \cdot 101 \cdot 241$

Igusa-Clebsch Invariants

$I_2 = 2670244 = 2^2 \cdot 667561$
 $I_4 = 217112281 = 217112281$
 $I_6 = 192196130709037 = 7 \cdot 13 \cdot 83 \cdot 25446330029$
 $I_{10} = 115278976 = 2^7 \cdot 37 \cdot 101 \cdot 241$

(Igusa invariants, G2 invariants)

Automorphism group

$\text{Aut}(X) \simeq C_2$
 $\text{Aut}(X_{\overline{\mathbb{Q}}}) \simeq C_2$

Rational points

Known points: $(1 : 0 : 0), (0 : 0 : 1)$

Number of rational Weierstrass points: 2

This curve is [locally solvable](#) everywhere.

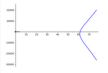
Mordell-Weil group of the Jacobian

Group structure: $\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$

Generator	D_h	Height	Order
$D_h = 2 \cdot (1 : 0 : 0)$	$47 \cdots 16x^2 - 18 \cdots 85xz + 23 \cdots 21z^2 = 0, 97 \cdots 08y = 15 \cdots 81x^2 + 79 \cdots 79z^2$	65.12498	∞
79643530293244152625559560327008189879			
1506336190309404255903375214327742890681			
971668119226160466248499255531481387008			
23857394798558414003864121			
182691020911837408458454185			
47162606537913965088399616			

Properties

Label 900617.a.900617.1



Conductor 900617
Discriminant 900617
Mordell-Weil group $\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$
Sato-Tate group $\text{USp}(4)$
 $\text{End}(A_p) \otimes \mathbb{R}$ \mathbb{R}
 $\text{End}(A_p) \otimes \mathbb{Q}$ \mathbb{Q}
 $\text{End}(A) \otimes \mathbb{Q}$ \mathbb{Q}
 $\overline{\rho}$ -simple yes
 GL_n -type no

Related objects

Isogeny class 900617.a
 Twists
 L-function

Learn more about

Completeness of the data
 Source of the data
 Reliability of the data
 Genus 2 curve labels

Strategies to find the torsion subgroup

Strategy 1. Torsion points have Néron-Tate height 0. Using an upper bound $|h_{\text{NT}}(x) - h_{\text{naive}}(x)| < B$, we look for points of naive height at most B .

Problem with strategy 1. When the genus gets greater, these height bounds may become impractical. Even if the torsion points are not so difficult to find in practice, the enumeration of all points might take a long time.

Strategy 2. Use the fact that $J_{\text{tors}}(\mathbb{Q}) \rightarrow J_p(\mathbb{F}_p)$ is injective, when p is an odd prime of good reduction, and J_p is the reduction of J at p . We get an upper bound by taking the greatest common divisor of $|J_p(\mathbb{F}_p)|$.

Problem with strategy 2. The upper bound might not be sharp, even if one is allowed to take as many primes as wanted. Even if one also looks at the group structure of $J_p(\mathbb{F}_p)$, it is possible that there is not enough information to prove that $J_{\text{tors}}(\mathbb{Q})$ has been found.

Why does strategy 2 fail?

Example


Suppose J has two-torsion points P , Q , and R such that

$$\mathbb{Q}(P) = \mathbb{Q}(\sqrt{2}), \quad \mathbb{Q}(Q) = \mathbb{Q}(\sqrt{3}), \quad \mathbb{Q}(R) = \mathbb{Q}(\sqrt{6}).$$

This happens for example on the hyperelliptic curve

$$y^2 = x(x^2 - 2)(x^2 - 3)(x^2 - 6).$$

We have the following subfield/subgroup diagram:



For each prime p , Frobenius acts as one of the four elements of V_4 , and we see that it must fix one of the torsion points P , Q , or R . This shows that $|J_p(\mathbb{F}_p)[2]| > |J(\mathbb{Q})[2]|$, for each prime p .

Improved strategy

Our new strategy is to find these torsion points over extensions of \mathbb{Q} . The following proposition gives us an idea how far one has to search.

Proposition (vB)

Suppose C has genus 3. There is a finite subset $S \subset J(\overline{\mathbb{Q}})[2]$ and a prime p of good reduction such that $[\mathbb{Q}(P) : \mathbb{Q}] \leq 12$ for all $P \in S$, and the reduction map induces a bijection

$$\langle J(\mathbb{Q})[2], S \rangle \longrightarrow J_p(\mathbb{F}_p)[2].$$

In other words: there is a prime p such that the 2-torsion modulo p can be explained by a collection of 2-torsion points over number fields of degree at most 12.

Proof.

The proof is a computation. One goes through all potential Galois groups of the 2-torsion, i.e. all subgroups of $\mathrm{Sp}(6, \mathbb{F}_2)$, and check that the conditions are satisfied. □

Inverting Abel-Jacobi

Through the rest of the talk, we will assume C has genus 3, and a divisor E of degree 1. We can use this divisor to uniquely write points in the Jacobian as $D - \deg(D) \cdot E$, where we take D to be effective and of minimal degree (which is at most 3 by Riemann-Roch).

There are multiple strategies to find torsion points over number fields. The first strategy is: inverting the Abel-Jacobi map

$$\iota: C(\mathbb{C})^3 \rightarrow \mathbb{C}^3/\Lambda: (P_1, P_2, P_3) \mapsto [P_1 + P_2 + P_3 - 3E].$$

If we have a point $P \in C(\mathbb{C})^3$ corresponding to a point in $J(\mathbb{Q})[\ell^n]$, then we can find all $Q \in C(\mathbb{C})^3$ such that $\ell \cdot \iota(Q) = \iota(P)$ using the Newton-Raphson method.¹

We can then use a lattice algorithm, e.g. LLL, to try to find algebraic relations for the coordinates of Q , and in this way produce torsion points over number fields.

¹For numerical stability reasons, it is actually better to take another base divisor $Q_1 + Q_2 + Q_3$ for the Abel-Jacobi map. We pick $Q_1, Q_2, Q_3 \in C(\mathbb{C})$ random.

Algebraic reconstruction with CRT

Another way to find torsion points is algebraic reconstruction. This time we look at all candidates for Q modulo p , i.e. we enumerate all points $\overline{Q}_p \in J_p(\mathbb{F}_p)$ such that $\ell \cdot \overline{Q}_p = \overline{P}$.

For any finite sets of primes S , we can combine information about potential reductions \overline{Q}_p with the Chinese remainder theorem to get a point modulo $\prod_{p \in S} p$. If the points \overline{Q}_p happen to be the reduction of the same algebraic point Q , we can use lattice techniques to try to reconstruct relations for the coordinates of Q .

Example

To illustrate how the method works, I will show how to reconstruct $\sqrt{2}$ from its reductions modulo p . In this case I picked $S = \{10007, 10009\}$. The roots of $x^2 - 2$ are $2641, 7366 \in \mathbb{F}_{10007}$ and $4419, 5590 \in \mathbb{F}_{10009}$.

$$\text{Let } I = (x - 2641, 10007) \cap (x - 4419, 10009) \subset \mathbb{Z}[x].$$

Then I is generated by $f_1 = x + 8893582$ and $f_2 = 100160063$, and an LLL-reduced basis for the lattice $\langle f_1x, f_1, f_2x^2, f_2x, f_2 \rangle$ contains the polynomial $x^2 - 2$, which is exactly the polynomial we were looking for.

Comparison of methods

Inverting Abel Jacobi:

- Advantage: it does not take a lot of extra computational power to obtain more precision.
- Disadvantage: one has to check all possible ℓ^{2g} possible ℓ -divisors of a torsion point P . For $g = 3$ this is only practical for $\ell = 2, 3$.

Algebraic reconstruction with the Chinese remainder theorem:

- Advantage: there are fewer possible ℓ -divisors of a torsion point P , when the primes are chosen carefully. By considering a quotient group $J_{\text{tors}}/J_{\text{known tors}}$, this number can be reduced even further.
- Disadvantage: to increase the precision of the result, an extra prime has to be added, which causes combinatorial explosion of the number of possible combinations of ℓ -divisors. This limits the degree of the number field to about 6 in practice.

Algebraic reconstruction with Hensel lifting:

- Using Hensel lifting, Reitsma computes torsion points modulo ℓ^N , for hyperelliptic curves of genus 3, and for N large enough to conclude completeness by using height bounds of Stoll. This method of producing torsion points might also work here.

Implementation

The implementation uses:

- Code of Costa to compute the L -polynomial of plane quartics modulo p for primes p of size $\approx 10^6$. This is needed to compute the order $|J_p(\mathbb{F}_p)|$ of the reduction of the Jacobian.
- Code of several people incorporated in Magma, e.g. to compute Riemann-Roch spaces for divisors on curves, LLL-reduced bases for lattices, Gröbner bases for ideals, and an analytic Abel-Jacobi map.
- The method `polredbest` in PARI/GP to compute small polynomials for the number fields involved in the computation.

The code has been used to try to compute the torsion for the 82 240 non-hyperelliptic curves of genus 3 in Sutherland's database. This took about 8 core months and succeeded for all but ~ 1200 of the curves. For ~ 600 of the failing cases, the problem was that no degree 1 divisor has been found on the curve.

For the successful cases, the torsion subgroup and completeness proofs have been stored. These can be verified without repeating the whole computation.

Statistics

We counted how often each order occurred.

1	2	3	4	5	6	7	8	9				
58662	8785	5087	2545	1105	1394	616	686	447				
10	11	12	13	14	15	16	17	18	19	20	21	
214	51	354	42	128	76	116	7	125	30	50	55	
22	23	24	25	26	27	28	29	30	31	32	33	35
17	2	83	14	14	19	29	1	23	3	19	12	5
36	38	39	40	41	42	44	45	48	49	50	51	52
31	4	7	15	1	15	2	7	14	2	2	2	2
54	56	57	60	62	64	65	66	70	72	75	80	84
2	4	4	8	1	6	1	3	3	4	2	2	1
				96	98	105	160					
				3	1	1	1					