Computing Euler factors of genus 2 curves at odd primes of almost good reduction

Andrew V. Sutherland (with Céline Maistret)

Massachusetts Institute of Technology

Sixteenth Algorithmic Number Theory Symposium (ANTS XVI)

Key takeaways from this talk

- Computing L-functions of genus $g \geq 2$ is hard ($g = 0$ is trivial, $g = 1$ is easy).
- Primes of almost good reduction are everywhere!
- Factoring is easier when the input is not squarefree.
- **Center** and **dig**
  (to reduce computations in $\mathbb{Z}_p[x]$ to computations in $\mathbb{Z}[x]$ and $\mathbb{F}_p[x]$).

The L-function of a nice curve $C/\mathbb{Q}$ of genus $g \geq 1$

The L-function of C is defined by

$$L(X, s) := \sum_{n \geq 1} a_n n^{-s} := \prod_p L_p(p^{-s})^{-1}.$$

For good primes p the zeta function

$$Z(X_p, T) := \exp\left(\sum_{r \geq 1} \#C_p(\mathbb{F}_{p^r}) \frac{T^r}{r}\right) = \frac{L_p(T)}{(1-T)(1-pT)},$$

determines the L-polynomial $L_p \in \mathbb{Z}[T]$. It satisfies

$$L_p(T) = T^{2g}\chi_p(1/T) = 1 - a_p T + \cdots + p^g T^{2g},$$

where $\chi_p(T)$ is the charpoly of the Frobenius endomorphism of $\mathrm{Jac}\,(C_p)$.

What about the bad primes?

At bad primes for $\operatorname{Jac}(C)$ (those dividing the conductor N) we have $\deg L_p < 2g$. The worse the reduction at p is, the higher $v_p(N)$ and the lower the degree of L_p.

Information-theoretically, bad reduction makes it easier to compute $L_p(T)$, since there are fewer candidates; indeed, sufficiently bad reduction forces $L_p(T) = 1$.

Thus the primes where $\operatorname{Jac}(C)$ has good reduction are arguably the hardest. But if C also has good reduction, we can compute $L_p(T)$ very quickly. Average polynomial-time algorithms compute $L_p(T)$ for $p \leq B$ of good reduction for C using $O(\log^4 p)$ bit operations per prime. For $g \leq 3$ there are practical implementations that are very fast (see papers in ANTS XI,XII,XIV,XV).

But if p is a prime of almost good reduction (good for $\operatorname{Jac}(X)$ but bad for X) we are stuck; none of Magma, Sage, Pari/GP efficiently and correctly handle this case.

Primes of almost good reduction are plentiful and may be large

Among the roughly five million genus 2 curves we know with conductor $N \leq 10^6$, nearly 3.5 million primes of almost good reduction arise. These occur frequently, even for curves with small coefficients, including the modular curve

$$X_0(22)\colon y^2 + (x^3 + x^2 + x + 1)y = -2x^6 + 4x^5 + 2x^4 + 5x^3 + 2x^2 + x$$

which has conductor 11^2 (so 2 is a prime of almost good reduction). But most are not modular curves and do not have any extra endomorphisms, including the curve

$$y^2 = -318x^6 - 450x^5 + 108x^4 + 150x^3 + 432x^2 - 162x + 66$$

with geometric endomorphism ring $\mathbb{Z}$, conductor $43 \cdot 8599$, and $2, 3, 5$ as primes of almost good reduction. There are examples with conductor $N \leq 10^6$ that have primes of almost good reduction much larger than N, as large as $43\,858\,540\,753$.

Setup

Recall that every nice genus 2 curve $C/\mathbb{Q}$ has a model of the form $y^2 = f(x)$, where $f \in \mathbb{Z}[x]$ is a squarefree sextic. Henceforth

- $f = \sum_i f_i x^i \in \mathbb{Z}[x]$ is a squarefree sextic;
- the cluster picture of f means the cluster picture of $C\colon y^2 = f(x)$;
- p is an odd prime of almost good reduction, so $\log p \leq \log|\Delta(f)\| = O(\|f\|)$, where $\|f\| := \max_i \|\log f_i\|$ is the size of the input to our algorithms

We say that f is p-normalized if its outer cluster has depth 0 and

$$v_p(f_6) = \min_i \{v_p(f_i)\} \leq 1.$$

Given f we can efficiently compute a p-normalized g defining an isomorphic curve.

Using GCDs to quickly find repeated roots/factors

Definition

For each positive integer k and polynomial $f \in \mathbb{F}_p[x]$ we define

$$\mathrm{gcd}_k(f) := \prod_{g|f} g^{\max(v_g(f)-k+1,0)} \in \mathbb{F}_p[x],$$

where g ranges over monic irreducibles in $\mathbb{F}_p[x]$ and $v_g(f) = \max\{e \in \mathbb{Z} : g^e|f\}$.

For $p > \deg(f)$ we have

$$\mathrm{gcd}_k(f) = \gcd\left(f, f^{(1)}, \ldots, f^{(k-1)}\right),$$

and for $p \leq \deg(f)$ we compute $\mathrm{gcd}_k(f)$ by brute force (note $\deg f = 6 = O(1)$).

If $\deg(f) = O(1)$ this takes quasi-linear time (versus quasi-quadratic for factoring).

p-normalization

Let $v = v_p(f_6)$. If $v > 1$ or $v \neq \min_i\{f_p(f_i)\}$ then let

$$e := \max\left\{\left\lceil \frac{v - v_p(f_i)}{6-i} \right\rceil : 0 \le i \le 5\right\}$$

and replace f by $p^{6e-w} f(x/p^e) \in \mathbb{Z}[x]$ where $w = 2\lfloor v/2 \rfloor$.

Now $v := v_p(f_6) = \min_i\{v_p(f_i)\} \le 1$. Let $h = p^{-v} f \in \mathbb{Z}[x]$. Then $v_p(h_6) = 0$ and the outer cluster has depth zero iff $\mathrm{gcd}_6(h) = 1$ (no root of multiplicity 6 mod p).

While $\bar{u} = \mathrm{gcd}_6(\bar{h}) \neq 1$ replace h by $p^{-6} h(px + a) \in \mathbb{Z}[x]$,
where $\bar{u} = x - \bar{a} \in \mathbb{F}_p[x]$ and $a \in \mathbb{Z}$ is any lift of $\bar{a} \in \mathbb{F}_p$.

Then $g = p^v h$ is p-normalized and $y^2 = g(x)$ is isomorphic to $y^2 = f(x)$.

We henceforth further assume that f is p-normalized.

Center and dig

Let $u \in [0, p-1]$ be distinct from $a_1, \ldots, a_j \in \mathbb{Z}$ be modulo p.

- given:
  $f(x) = (x - a_1) \cdots (x - a_j)(x - pr_1 - u) \cdots (x - pr_k - u)$
  $\bar{f}(x) = (x - a_1) \cdots (x - a_j)(x - u)^k$
- center:
  $f(x + u) = (x - a_1 + u) \cdots (x - a_j + u)(x - pr_1) \cdots (x - pr_k)$
  $\bar{f}(x + u) = (x - a_1 + u) \cdots (x - a_j + u)x^k$
- dig:
  $p^{-k}f(px + u) = (px - a_1 + u) \cdots (px - a_j + u)(x - r_1) \cdots (x - r_k)$
  $\overline{p^{-k}f(px + u)} = c(x - r_1) \cdots (x - r_k)$

Reduction types

Let $f \in \mathbb{Z}[x]$ by p-normalized, let $\bar{c} := f_6 p^{-v_p(f_6)} \in \mathbb{F}_p^\times$, and let $\bar{f} = p^{-v_p(f_6)} f \in \mathbb{F}_p[x]$.
Then exactly one of the following holds, with $m \geq n$ of the same parity as $v_p(f_6)$.

type	picture	$\bar{f}$	$L_p(C,T)$
1		$\bar{c}(x-\bar{r})^3\bar{u}$	$L_p(E_1,T)L_p(E_2,T)$ over $\mathbb{F}_p$
2a	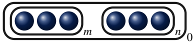	$\bar{c}(c-\bar{r})^3(x-\bar{s})^3$	$L_p(E_1,T)L_p(E_2,T)$ over $\mathbb{F}_p$
2b	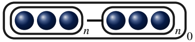	$\bar{c}\bar{q}^3$	$L_p(E_1,T^2)$ over $\mathbb{F}_{p^2}$
4	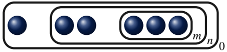	$\bar{c}(x-\bar{r})^5(x-\bar{s})$	$L_p(E_1,T)L_p(E_2,T)$ over $\mathbb{F}_p$

with $\bar{r}, \bar{s} \in \mathbb{F}_p$ distinct, $\bar{u} \in \mathbb{F}_p[x]$ a squarefree monic cubic with $\bar{u}(\bar{r}) \neq 0$,
and $\bar{q} \in \mathbb{F}_p[x]$ an irreducible monic quadratic.

Computing L-polynomials for the split types

Let $\tilde{f} = p^{-v_p(f_6)} f \in \mathbb{Z}[x]$ and let L be the splitting field of f over $\mathbb{Q}_p$.

Let $r_1 \in \mathcal{O}_L$ be a root of depth n and $r_2 \in \mathcal{O}_L$ a root of relative depth m (if any).

Let $s_1, s_2 \in \mathbb{Z}$ satisfy $r_1 \equiv s_1 \bmod p^n \mathcal{O}_L$ and $r_2 \equiv s_2 \bmod p^m \mathcal{O}_L$ and define

picture	$\bar{g}_1 \in \mathbb{F}_p[x]$	$\bar{g}_2 \in \mathbb{F}_p[x]$
	$\mathrm{sqf}(\bar{f})$	$\tilde{f}(p^n x + s_1)/p^{3n}$
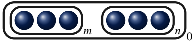	$\tilde{f}(p^n x + s_1)/p^{3n}$	$\tilde{f}(p^m x + s_2)/p^{3m}$
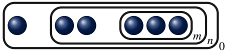	$\mathrm{sqf}(\tilde{f}(p^n x + x_1)/p^{5n} \bmod p)$	$\tilde{f}(p^m x + s_2)/p^{3m+2n}$

Then $L_p(C,T) = L_p(E_1,T)L_p(E_2,T)$ for $E_1\colon y^2 = \bar{g}_1(x)$ and $E_2\colon y^2 = \bar{g}_2(x)$.

Computing L-polynomials for the non-split type

Let $\tilde{f} = p^{-v_p(f_6)} f \in \mathbb{Z}[x]$, let L be the splitting field of f over $\mathbb{Q}_p$.

Let $\bar{f} = \tilde{f} \in \mathbb{F}_p[x]$ and $\bar{q} \in \mathbb{F}_p[x]$ the irreducible monic quadratic for which $\bar{f} = \bar{c}\bar{q}^3$.

Let $q \in \mathbb{Z}[x]$ be any lift of $\bar{q}$, let $F := \mathbb{Q}_p[z]/(q(z)) \subseteq L$ and $\mathcal{O} := \mathbb{Z}[z]/(q(z)) \subseteq \mathcal{O}_L$.

Let $\kappa := \mathbb{F}_p[z]/(\bar{q}(z)) \simeq \mathbb{F}_{p^2}$, let $r \in \mathcal{O}_L$ be a root of f, and $s \in \mathcal{O}$ with $r \equiv s \bmod p\mathcal{O}_L$.

Let $\hat{f}$ be the image of f in $\mathcal{O}[x]$ via $\mathbb{Z}[z] \subseteq \mathcal{O}[z]$ induced by $\mathbb{Z} \subseteq \mathcal{O}$, and let

$$\bar{g} = \hat{f}(p^n x + s)/p^{3n} \bmod p\mathcal{O} = \kappa[x] \simeq \mathbb{F}_{p^2}[x].$$

Then for the elliptic curve $E\colon y^2 = \bar{g}(x)$ over $\mathbb{F}_{p^2}$ we have $L_p(C, T) = L_p(E, T^2)$.

Main result

Theorem (Maistret-S)

Let $C/\mathbb{Q}$ be a genus 2 curve $y^2 = f(x) = \sum_i f_i x^i \in \mathbb{Z}[x]$ with almost good reduction at an odd prime p. There is a deterministic algorithm that, given a nonsquare element of $\mathbb{F}_p^\times$, computes the L-polynomial $L_p(C,T)$ in time

$$O(\|f\|^2 \log^2 \|f\| / \log p + \log^5 p),$$

where $\|f\| = \max_i \log \|f_i|$. There is a Las Vegas algorithm with the same expected running time that does not require a nonsquare element of $\mathbb{F}_p^\times$.

Timings

Timings for computing $3\,454\,506$ Euler factors of genus 2 curves $C/\mathbb{Q}$ of small conductor at odd primes of almost good reduction.

method	total time	average time	median time	maximum time
EulerFactor	242 days	6.1 s	0.9 s	over 8 hours
New alg (Magma)	1.23 hours	1.3 ms	1.2 ms	24 ms
New alg (C)	27.1 s	7.8 μs	3.4 μs	21 ms

https://github.com/AndrewVSutherland/Genus2Euler