Enumerating hyperelliptic curves over finite fields in quasilinear time

Everett W. Howe

Unaffiliated mathematician based in San Diego, California (unceded Kumeyaay land)

16th Algorithmic Number Theory Symposium (ANTS XVI)

Massachusetts Institute of Technology
15–19 July 2024

Email: however@alumni.caltech.edu Web site: ewhowe.com Bluesky: @however.bsky.social Mastodon: @however@tech.lgbt

Hyperelliptic curves

Definition

A curve C over a field k is *hyperelliptic* if its genus is at least 2 and it has an involution ι such that $C/\langle\iota\rangle$ has genus 0.

- Such an involution is unique if it exists.
- Over a finite field $\mathbf{F}_q$, every genus-0 curve is isomorphic to $\mathbf{P}^1$, so. . .
- We get a double cover $\varphi\colon C \to \mathbf{P}^1$, unique up to $\operatorname{Aut} C$ and $\operatorname{Aut} \mathbf{P}^1 \cong \mathrm{PGL}_2(\mathbf{F}_q)$.
- If characteristic is odd: $y^2 = f(x, z)$ with f homogeneous, degree $2g + 2$.
- $\operatorname{div} f \subset \mathbf{P}^1$ is the ramification divisor of φ. It is effective, reduced, degree $2g + 2$.

Theorem 1

C is determined up to quadratic twist by the $\mathrm{PGL}_2(\mathbf{F}_q)$ orbit of $\operatorname{div} f$.

So: enumerating hyperelliptic curves of genus g over $\mathbf{F}_q$
$\iff$ enumerating $\mathrm{PGL}_2(\mathbf{F}_q)$ orbits of effective reduced divisors of degree $2g + 2$.

Moduli spaces of hyperelliptic curves of genus g

Classical strategy:

- Compute invariants for hyperelliptic curves of genus g. (Framework for this goes back to Gordan.)
- Curve k-rational $\Longrightarrow$ invariants k-rational
- Mestre (1990): For curves with no automorphisms other than ι, converse holds if a certain conic has a rational point. Always true over a finite field.
- Making the converse effective involves solving a system of polynomial equations.
- Curves with larger automorphism groups must be dealt with separately.

The case of genus 2

Igusa invariants (Igusa 1972)

- Elements $[J_2 : J_4 : J_6 : J_8 : J_{10}]$ of weighted projective space defined over $\mathbf{Z}$
- $J_{10} \neq 0$ and $4J_8 - J_6J_2 + J_4^2 = 0$
- Easy to enumerate all elements over a field

Curves from invariants

- Mestre (1990): Details of case where $\#\,\mathrm{Aut}\, C = 2$
- Cardona and Quer (2005): Handle the larger automorphism groups

Implemented in Magma

- Uses "G_2 invariants" of Cardona/Quer for convenience: triples $(a, b, c) \in \mathbf{F}_q^3$
- All triples $(a, b, c) \in \mathbf{F}_q^3$ are legal G_2 invariants
- `Twists(HyperellipticCurveFromG2Invariants([a,b,c]))`

The case of genus 3

Invariants

- Shioda (1967) worked out one set of invariants
- Shioda's basis does not include the discriminant
- Lercier and Ritzenthaler (2012): Invariants in weighted projective space

Curves from invariants

- Worked out by Lercier/Ritzenthaler: A tour de force!

Implemented in Magma

- Start with invariants in weighted projective space
- Compute Shioda invariants; check to see whether discriminant is nonzero
- `TwistedHyperellipticPolynomialsFromShiodaInvariants(S)`

Moduli spaces versus enumeration over finite fields

Advantages and disadvantages of moduli space approach

- Works over essentially all fields
- Requires new math to be done for every genus
- Daunting to think of generalizing even just to genus 4

Advantages and disadvantages of our approach to enumeration

- Specific to finite fields — a case of particular interest
- Can handle arbitrary genera with no additional work
- *For fixed g*: Enumerate all genus-g hyperelliptic curves $/\mathbf{F}_q$ in time $\widetilde{O}(q^{2g-1})$
- In practice, orders of magnitude faster than preceding approach
- Have not yet worked out dependence on the genus
- ANTS version uses $O(q^{2g-1})$ memory — improved to $O(\log q)$ in followup

Galois structure of Weierstrass points is relevant

Assume that $2g + 1$ is not divisible by the characteristic of $\mathbf{F}_q$.

Normal form for a hyperelliptic curve with a rational Weierstrass point

- Use $\mathrm{PGL}_2(\mathbf{F}_q)$ to move a rational ramification point of $C \to \mathbf{P}^1$ to ∞.
- Get $y^2 = f(x)$ with $\deg f = 2g + 1$.
- Use translations to eliminate coefficient of x^{2g} in f.
- Up to twists, have $y^2 = x^{2g+1} + a_{2g-1}x^{2g-1} + \cdots + a_1 x + a_0$.
- If $a_0 \neq 0$, scale x so that a_0 is in a fixed set of representatives for $\mathbf{F}_q^\times / \mathbf{F}_q^{\times(4g+2)}$.
- If $a_0 = 0$, scale x so that a_1 is in a fixed set of representatives for $\mathbf{F}_q^\times / \mathbf{F}_q^{\times(4g)}$.
- At most $(2g + 2)(4g + 2)$ ways of doing this. Choose "smallest" f we get.

To enumerate curves, loop through all degree-$(2g + 1)$ polynomials f with $a_{2g} = 0$ and with a_0 (or a_1) in given set of reps, and output those that are in normal form.

Quasilinear time algorithm for about 63.2% of all hyperelliptic curves.

The fundamental case: An irreducible Weierstrass divisor

We can enumerate curves with a rational Weierstrass point.

Furthest away from that case: $C \to \mathbf{P}^1$ ramified at a single place of degree $2g + 2$.

Main question:

How do we quickly enumerate orbit representatives for $\mathrm{PGL}_2(\mathbf{F}_q)$ acting on:

- Degree-$2m$ places of $\mathbf{P}^1$ over $\mathbf{F}_q$?
- Or, equivalently, monic irreducible polynomials of degree $2m$?

As argued in the paper, if we can do this, we can quickly enumerate all curves.

The cross polynomial

Definition

Let f be a monic irreducible polynomial over $\mathbf{F}_q$ of degree $n > 3$.
Let $\alpha \in \mathbf{F}_{q^n}$ be a root of f, and let $\chi \in \mathbf{F}_{q^n}$ be the cross ratio of α, α^q, α^{q^2}, and α^{q^3}:

$$\chi := \frac{(\alpha^{q^3} - \alpha^q)(\alpha^{q^2} - \alpha)}{(\alpha^{q^3} - \alpha)(\alpha^{q^2} - \alpha^q)}.$$

The *cross polynomial* $\operatorname{Cross} f$ of f is the characteristic polynomial of χ.

Theorem 2

Two monic irreducible polynomials over $\mathbf{F}_q$ of degree at least 4 are in the same $\mathrm{PGL}_2(\mathbf{F}_q)$ orbit if and only if their cross polynomials are equal.

Computing PGL_2 orbits of places, but not quite fast enough

Algorithm: Representatives for PGL_2 orbits of degree-n irreducibles

- Input: q and $n > 3$.
- Construct basis $\alpha_1, \ldots, \alpha_n$ of $\mathbf{F}_{q^n}$ such that $1 = a_1\alpha_1 + \cdots + a_n\alpha_n$ with $a_1 \neq 0$.
- Set L to be the empty list.
- For every $(b_2, \ldots, b_n) \in \mathbf{F}_q^{n-1}$ such that first nonzero coordinate is 1:
  - Set f to be the minimal polynomial of $b_2\alpha_2 + \cdots + b_n\alpha_n$.
  - If f has degree n then append the pair $(\mathrm{Cross}\, f, f)$ to L.
- Sort L.
- Delete $(\mathrm{Cross}\, f, f)$ from L if $\mathrm{Cross}\, f$ appears earlier on list.
- Output the second elements of each pair remaining on L.

Easy to see: Output is correct. Requires time $\widetilde{O}(q^{n-2})$ and space $O(q^{n-2})$.

Orbit reps for even-degree places in quasilinear time: The idea

Suppose f is monic irreducible polynomial over $\mathbf{F}_q$ of degree $2m$.

Then over $\mathbf{F}_{q^2}$, f factors as a product $g \cdot g^{(q)}$, where g is monic of degree m.

To enumerate degree-$2m$ irreducibles over $\mathbf{F}_q$ up to $\mathrm{PGL}_2(\mathbf{F}_q)$,
enumerate degree-m irreducibles over $\mathbf{F}_{q^2}$ up to $\mathrm{PGL}_2(\mathbf{F}_q)$.

(We will see why this is helpful.)

First take orbit reps for the degree-m irreducibles over $\mathbf{F}_{q^2}$ up to $\mathrm{PGL}_2(\mathbf{F}_{q^2})$,
then expand them using right coset represetatives for $\mathrm{PGL}_2(\mathbf{F}_q)$ in $\mathrm{PGL}_2(\mathbf{F}_{q^2})$.

It's easy to produce an explicit list of these coset representatives: See the paper.

Orbit reps for even-degree places in quasilinear time: Details

Algorithm: Representatives for PGL_2 orbits of degree-$2m$ irreducibles

- Input q, $m > 2$, and list R of right coset reps of $\mathrm{PGL}_2(\mathbf{F}_q)$ in $\mathrm{PGL}_2(\mathbf{F}_{q^2})$.
- Construct list M of orbit reps for degree-m irreducibles $/\mathbf{F}_{q^2}$ up to $\mathrm{PGL}_2(\mathbf{F}_{q^2})$.
- Set L to be the empty list.
- For every $h \in M$ and $\Gamma \in R$:
  - Set $g = \Gamma(h)$ and set $f = gg^{(q)}$.
  - Append the pair $(\mathrm{Cross}\, f, f)$ to L.
- Sort L.
- Delete $(\mathrm{Cross}\, f, f)$ from L if $\mathrm{Cross}\, f$ appears earlier on list.
- Output the second elements of each pair remaining on L.

First step takes time $\widetilde{O}((q^2)^{(m-2)}) = \widetilde{O}(q^{n-4})$ using earlier algorithm.

$O(q^{n-6})$ elements in M and $O(q^3)$ elements in R.

Output is correct. Requires time $\widetilde{O}(q^{n-3})$ and space $O(q^{n-3})$.

Timings

Sample timings (in seconds) to compute all hyperelliptic curves of genus 2 and 3 over $\mathbf{F}_q$.

"Magma" columns: timings for Mestre/Cardona/Quer and Lercier/Ritzenthaler as built into Magma.
"Divisors"columns: timings for our method of computing orbit reps for PGL_2 acting on divisors.
"Curves" columns: timings for deriving curves from divisors (i.e. checking twists).

Timings with an asterisk are estimates based on extrapolation from 10,000 random examples.

Genus 2

		Our method		
q	Magma	Divisors	Curves	Total
17	8	0.2	0.02	0.2
31	52	0.8	0.06	0.8
59	327	3.8	0.25	4.1
127	3308	36	2	38
257	27448*	290	10	300
509	211655*	2307	76	2384

Genus 3

		Our method		
q	Magma	Divisors	Curves	Total
17	5274	20	1	21
31	99463*	304	14	318
59	2408665*	5932	479	6411

What about odd-degree places?

For enumerating hyperelliptic curves in quasilinear time, we have all we need.

Yet we might still wonder:
How to enumerate PGL_2 orbits of odd-degree places in quasilinear time?

Interesting on its own as a question.
But also useful (for example) for enumerating cyclic covers of $\mathbf{P}^1$ of higher degree.

This is covered in the followup paper:
Enumerating places of $\mathbf{P}^1$ *up to automorphisms of* $\mathbf{P}^1$ *in quasilinear time*
arXiv: 2407.05534 [math.NT]

Frobenius functions and Frobenius divisors

Theorem/Definition

Let f be a monic irreducible degree-n polynomial with $n > 1$ odd. Among the rational functions of degree at most $(n-1)/2$, there is a unique F such that $\alpha^q = F(\alpha)$ for all roots α of f. We call F the *Frobenius function* for f.

The uniqueness depends on n being odd.

Definition

Let $F = g/h$ be the Frobenius function for f, viewed as a rational function on $\mathbf{P}^1$, so g and h are homogeneous polynomials in $\mathbf{F}_q[x, z]$.

The divisor of the homogeneous polynomial $xh - zg$ is the *Frobenius divisor* of f.

The Frobenius divisor is the "divisor of fixed points" of F; its degree is $\leq (n+1)/2$.

Theorem 3

The map from irreducible odd-degree polynomials to their Frobenius divisors is PGL_2*-equivariant under the natural action of* PGL_2 *on both sets.*

Frobenius functions and fixed points

Suppose F is the Frobenius function for a degree-n polynomial f.

- Let α be a root of f.
- Then $\alpha^q = F(\alpha)$, and more generally $\alpha^{q^i} = F^{(i)}(\alpha)$.
- α is a fixed point of $F^{(n)}$.
- f divides the numerator of $x - F^{(n)}$.

Finding the degree-n polynomials with Frobenius function F

For every degree-n irreducible factor f of the numerator of $x - F^{(n)}$, check whether F is the Frobenius function for f.

Warning: If $\deg F = 1$ this doesn't work (why?), and instead we do something else.

Orbit reps for odd-degree places in quasilinear time

Algorithm: PGL_2 orbit representatives for irreducibles of odd degree n

- Input q, odd $n > 1$, and list M of orbit representatives for $\mathrm{PGL}_2(\mathbf{F}_q)$ acting on effective divisors of degree up to $(n+1)/2$.
- For each divisor D on the list M, find the functions F with fixed-point divisor D.
- (Only use one F from each orbit of $\mathrm{Aut}\, D \subset \mathrm{PGL}_2(\mathbf{F}_q)$ acting by conjugation.)
- For each such F, output the degree-n polynomials f with Frobenius function F.

How to compute the list M required as input?

When $n > 5$, naïve methods are fast enough! But we can also use recursion.