

Counting points on curves: the general case

Jan Tuitman, KU Leuven

October 14, 2015

Algebraic curves

Let X be a smooth projective algebraic curve of genus g over some finite field \mathbf{F}_q with $q = p^n$.

Example (Projective line, $g = 0$)

$$X = \mathbb{P}_{\mathbf{F}_q}^1.$$

Example (Elliptic curve, $g = 1$)

$$X = \{(x : y : z) \in \mathbb{P}_{\mathbf{F}_q}^2 : y^2z = x^3 + axz^2 + bz^3\}$$

where $p \neq 2$ (and $4a^3 + 27b^2 \neq 0$).

Example (Non-hyperelliptic curve, $g = 4$)

$$X = \{(x : y : z : w) \in \mathbb{P}_{\mathbf{F}_q}^3 : S_2(x, y, z, w) = S_3(x, y, z, w) = 0\}$$

where $S_2, S_3 \in \mathbf{F}_q[x, y, z, w]$ are a quadric and a cubic, respectively (and some smoothness condition is satisfied).

Zeta functions

Let $|X(\mathbf{F}_{q^i})|$ denote the number of points of X with values in \mathbf{F}_{q^i} (the number of solutions of the equations for X in this field).

Recall that the zeta function of X is defined as

$$Z(X, T) = \exp \left(\sum_{i=1}^{\infty} |X(\mathbf{F}_{q^i})| \frac{T^i}{i} \right).$$

It follows from the Weil conjectures that $Z(X, T)$ is of the form

$$\frac{\chi(T)}{(1-T)(1-qT)},$$

where $\chi(T) \in \mathbf{Z}[T]$ of degree $2g$, with inverse roots that

- have complex absolute value $q^{\frac{1}{2}}$
- are permuted by the map $x \rightarrow q/x$.

Example: the projective line

Let us do an easy example.

We have

$$|\mathbb{P}^1(\mathbf{F}_{q^i})| = q^i + 1$$

so that

$$\begin{aligned} Z(\mathbb{P}_{\mathbf{F}_q}^1, T) &= \exp\left(\sum_{i=1}^{\infty} (q^i + 1) \frac{T^i}{i}\right) \\ &= \exp\left(\sum_{i=1}^{\infty} \frac{T^i}{i}\right) \exp\left(\sum_{i=1}^{\infty} \frac{(qT)^i}{i}\right) \\ &= \frac{1}{(1-T)(1-qT)} \end{aligned}$$

The problem

Problem

Compute $Z(X, T)$, or equivalently $\chi(T)$, in an efficient way.

Remark

This problem is often referred to as 'counting points'.

Remark

Let J_X denote the Jacobian variety of X . Then

$$|J_X(\mathbf{F}_q)| = \chi(1).$$

Computing $|J_X(\mathbf{F}_q)|$ is important for the Discrete Logarithm Problem on $J_X(\mathbf{F}_q)$. If this order only has small prime factors then the DLP is easy. However, in cryptography only curves of genus ≤ 2 are used, and for those curves good algorithms for counting points already exist.

Constructing p -adic cohomology

To compute zeta functions, we will use so called p -adic cohomology.

We are going to explain the construction of construction of p -adic cohomology only in the case of a smooth affine curve:

$$U = \{(x_1, \dots, x_m) \in \mathbb{A}_{\mathbf{F}_q}^m : f_1(x_1, \dots, x_m) = \dots = f_\ell(x_1, \dots, x_m) = 0\}$$

where the $f_i(x_1, \dots, x_m)$ are all elements of $\mathbf{F}_q[x_1, \dots, x_m]$ (and some smoothness condition is satisfied).

We denote

$$R = \mathbf{F}_q[x_1, \dots, x_m]/(f_1, \dots, f_\ell),$$

so that $U = \text{Spec}(R)$. First we need to lift to characteristic 0.

Lifting to characteristic 0

Let \mathbf{Q}_q denote the unique unramified extension of \mathbf{Q}_p of degree n and \mathbf{Z}_q the ring of integers of \mathbf{Q}_q .

Let $f_1, \dots, f_\ell \in \mathbf{Z}_q[x_1, \dots, x_m]$ denote lifts of f_1, \dots, f_ℓ (for which the smoothness condition is still satisfied).

We denote

$$\mathcal{U} = \{(x_1, \dots, x_m) \in \mathbb{A}_{\mathbf{Z}_q}^m : f_1(x_1, \dots, x_m) = \dots = f_\ell(x_1, \dots, x_m) = 0\}$$

and again

$$\mathcal{R} = \mathbf{Z}_q[x_1, \dots, x_m]/(f_1, \dots, f_\ell),$$

so that $\mathcal{U} = \text{Spec}(\mathcal{R})$.

Weak completion

Consider the ring of power series over \mathbf{Z}_q in m variables that converge p -adically on a disk of radius strictly greater than 1:

$$\mathbf{Z}_q\langle x_1, \dots, x_m \rangle^\dagger = \left\{ \sum a_I x^I : a_I \in \mathbf{Z}_q \text{ and } \exists \rho > 1 \text{ s.t. } \lim_{|I| \rightarrow \infty} |a_I| \rho^{|I|} = 0 \right\}$$

where $I = (i_1, \dots, i_m)$ and $|I| = i_1 + \dots + i_m$.

We then define the weak completion of \mathcal{R} as

$$\mathcal{R}^\dagger = \mathbf{Z}_q\langle x_1, \dots, x_m \rangle^\dagger / (f_1, \dots, f_\ell).$$

This is also called an overconvergent or dagger algebra.

p -adic cohomology

Now we define the overconvergent 1-forms

$$\Omega_{\mathcal{R}^\dagger}^1 = (\mathcal{R}^\dagger dx_1 \oplus \dots \oplus \mathcal{R}^\dagger dx_m) / (df_1, \dots, df_\ell)$$

and the overconvergent De Rham complex:

$$0 \longrightarrow \mathcal{R}^\dagger \xrightarrow{d} \Omega_{\mathcal{R}^\dagger}^1 \longrightarrow 0$$

where d is defined by $dg = \frac{\partial g}{\partial x_1} dx_1 + \dots + \frac{\partial g}{\partial x_m} dx_m$. The p -adic (or rigid) cohomology spaces of U are then defined as

$$H_{\text{rig}}^0(U) = \ker d \otimes \mathbf{Q}_q \quad H_{\text{rig}}^1(U) = \text{coker } d \otimes \mathbf{Q}_q.$$

It can be shown that these are finite dimensional vector spaces over \mathbf{Q}_q that do not depend on any of the choices made in their construction.

Lefschetz formula

The map F_q that sends each x_i to x_i^q defines a map from U to itself, or equivalently a homomorphism from R to itself.

One can show that F_q can be lifted to the weak completion \mathcal{R}^\dagger , i.e. that there exists a homomorphism \mathcal{F}_q from \mathcal{R}^\dagger to itself, such that \mathcal{F}_q reduces to F_q modulo p .

This homomorphism is called a Frobenius lift. It acts naturally on the p -adic cohomology spaces and the following formula holds:

$$Z(U, T) = \frac{\det(1 - (q\mathcal{F}_q^{-1})T | H_{\text{rig}}^1(U))}{(1 - qT)}$$

assuming that U is connected.

Example: the affine line minus zero

In this simple case no weak completion is needed

$$R = \mathbf{F}_q[x, 1/x]$$

$$U = \text{Spec}(R)$$

$$\mathcal{R} = \mathbf{Z}_q[x, 1/x]$$

$$\Omega_{\mathcal{R}}^1 = \mathbf{Z}_q[x, 1/x]dx$$

$$H_{\text{rig}}^0(U) = \mathbf{Q}_q$$

$$H_{\text{rig}}^1(U) = \mathbf{Q}_q \frac{dx}{x}$$

$$\mathcal{F}_q(x) = x^q$$

$$\mathcal{F}_q\left(\frac{dx}{x}\right) = \frac{d(x^q)}{x^q} = q \frac{dx}{x}$$

and we check that the Lefschetz formula gives the correct zeta function

$$Z(U, T) = \frac{(1 - T)}{(1 - qT)} = \exp\left(\sum_{i=1}^{\infty} (q^i - 1) \frac{T^i}{i}\right)$$

Some remarks

For X smooth projective (so not affine), the Lefschetz formula becomes

$$Z(X, T) = \frac{\det(1 - (q\mathcal{F}_q^{-1})T | H_{\text{rig}}^1(X))}{(1 - T)(1 - qT)}.$$

Here one may also replace $q\mathcal{F}_q^{-1}$ by \mathcal{F}_q (by Poincaré duality).

Actually, one never computes directly with \mathcal{F}_q , but instead with \mathcal{F}_p . However, \mathcal{F}_p is only σ -semilinear, where σ is the unique lift of the p -th power map from \mathbf{F}_q to \mathbf{Z}_q .

All of this is not very important for the rest of this talk.

Hyperelliptic curves

Kedlaya (2001) applied p -adic cohomology to the computation of zeta functions of hyperelliptic curves in odd characteristic.

Let \mathbf{F}_q be a finite field with $q = p^n$ and p an odd prime. Moreover, let X be the projective nonsingular curve of genus g with affine equation

$$y^2 = Q(x)$$

with $Q(x) \in \mathbf{F}_q[x]$ monic and separable of degree $2g + 1$.

Take out all of the ramification points of the map $x : X \rightarrow \mathbb{P}_{\mathbf{F}_q}^1$ from the curve and consider the open affine subset

$$U = \{(x, y) \in \mathbb{A}_{\mathbf{F}_q}^2 : y^2 = Q(x) \text{ and } y \neq 0\}$$

of X with coordinate ring

$$R = \mathbf{F}_q[x, y, 1/y]/(y^2 - Q(x)).$$

Frobenius lift and cohomology

Let $Q \in \mathbf{Z}_q[x]$ be any monic lift of Q and define

$$\mathcal{R} = \mathbf{Z}_q[x, y, 1/y]/(y^2 - Q(x)) \quad \mathcal{R}^\dagger = \mathbf{Z}_q\langle x, y, 1/y \rangle^\dagger/(y^2 - Q(x)).$$

We construct a Frobenius lift \mathcal{F}_p on \mathcal{R}^\dagger by setting

$$\mathcal{F}_p(x) = x^p$$

$$\mathcal{F}_p(y) = Q^\sigma(x^p)^{\frac{1}{2}} = y^p \left(1 + \frac{Q^\sigma(x^p) - Q(x)^p}{y^{2p}} \right)^{\frac{1}{2}}.$$

Theorem (Kedlaya)

A basis for $H_{rig}^1(U)$ is given by

$$\left[x^0 \frac{dx}{y}, \dots, x^{2g-1} \frac{dx}{y}, x^0 \frac{dx}{y^2}, \dots, x^{2g} \frac{dx}{y^2} \right]$$

and the first $2g$ vectors form a basis for the subspace $H_{rig}^1(X)$.

Kedlaya's algorithm

Algorithm

- Apply \mathcal{F}_p to the basis $[x^0 \frac{dx}{y}, \dots, x^{2g-1} \frac{dx}{y}]$ of $H_{rig}^1(X)$.
- Reduce resulting elements of $\Omega_{\mathcal{R}^\dagger}^1$ back to this basis by subtracting df with $f \in \mathcal{R}^\dagger$ and read off the matrix Φ_p of \mathcal{F}_p on $H_{rig}^1(X)$.
- Compute the matrix $\Phi_q = \Phi_p^{\sigma^{n-1}} \dots \Phi_p^\sigma \Phi_p$ of \mathcal{F}_q on $H_{rig}^1(X)$.
- Determine $\chi(T) = \det(1 - \Phi_q T)$ numerator of $Z(X, T)$.

Theorem (Kedlaya)

This algorithm runs in

$$\text{time: } \tilde{O}(pg^4 n^3) \quad \text{space: } \tilde{O}(pg^3 n^3)$$

Remark

Implemented in MAGMA by M. Harrison, quite practical.

Extensions of Kedlaya's algorithm

Kedlaya's algorithm was extended in various ways by various people, here are a few of them:

- Gaudry and Gurel (2001), superelliptic curves
- Vercauteren (2002), hyperelliptic curves in characteristic 2
- Denef and Vercauteren (2006), C_{ab} curves
- Castryck, Denef and Vercauteren (2006), nondegenerate curves

The first two algorithms in this list are small adaptations of Kedlaya's algorithm and equally practical. The third and especially the fourth are much more general, but partial implementations have shown them to be unpractical. Therefore, complete implementations do not exist, as far as we know.

General curves

We let X/\mathbf{F}_q denote the smooth projective curve birational to

$$Q(x, y) = y^{d_x} + Q_{d_x-1}(x)y^{d_x-1} + \dots + Q_0 = 0,$$

where $Q(x, y)$ is irreducible separable and $Q_i(x) \in \mathbf{F}_q[x]$ for all $0 \leq i \leq d_x - 1$.

We let $\mathcal{Q} \in \mathbf{Z}_q[x]$ denote a lift of Q that is monic of degree d_x in y .

$\Delta(x) \in \mathbf{Z}_q[x]$ denotes the resultant of \mathcal{Q} and $\frac{\partial \mathcal{Q}}{\partial y}$ with respect to the variable y and $r(x) \in \mathbf{Z}_q[x]$ the squarefree polynomial

$$r(x) = \Delta / \left(\gcd \left(\Delta, \frac{d\Delta}{dx} \right) \right).$$

Lift to characteristic 0

We take out $r(x) = 0$ from X and define

$$\mathcal{U} = \{(x, y) \in \mathbb{A}_{\mathbf{Z}_q}^2 : Q(x, y) = 0 \text{ and } r(x) \neq 0\}$$

with coordinate ring

$$\mathcal{R} = \mathbf{Z}_q[x, 1/r(x), y]/(Q).$$

For our algorithm to work we need the following condition.

Assumption

The polynomial $r(x)$ is separable (no multiple roots) over \mathbf{F}_q (so mod p).

If this is the case, we say that we have found a ‘good lift’ to characteristic 0.

p -adic cohomology

We define

$$\mathcal{R}^\dagger = \mathbf{Z}_q \langle x, 1/r(x), y \rangle^\dagger / (\mathcal{Q}).$$

Recall that

$$\Omega_{\mathcal{R}^\dagger}^1 = \frac{R^\dagger dx \oplus R^\dagger dy}{d\mathcal{Q}}$$

and that if we denote $d : \mathcal{R}^\dagger \rightarrow \Omega_{\mathcal{R}^\dagger}^1$, we have

$$H_{\text{rig}}^1(U) = \text{coker}(d) \otimes \mathbf{Q}_q.$$

Moreover, $H_{\text{rig}}^1(X)$ is the subspace of $H_{\text{rig}}^1(U)$ defined by the vanishing of a so called cohomological residue map.

Frobenius lift

To construct a Frobenius lift \mathcal{F}_p from \mathcal{R}^\dagger to itself, we set

$$\mathcal{F}_p(x) = x^p$$

and compute $\mathcal{F}_p(y)$ (to any desired precision) by Hensel lifting using the equation

$$Q^\sigma(x^p, \mathcal{F}_p(y)) = 0.$$

Note that this is possible because we have removed the zeros of $\frac{\partial Q}{\partial y}$ from the curve X by removing the zeros of $r(x)$.

After precomputing $\mathcal{F}_p(y), \dots, \mathcal{F}_p(y^{d_x-1})$ and $\mathcal{F}_p(1/r)$ it is quite easy to evaluate \mathcal{F}_p on elements of \mathcal{R}^\dagger and $\Omega_{\mathcal{R}^\dagger}^1$.

Integral bases

Let

$$W^0 \in Gl_{d_x}(\mathbf{Z}_q[x, 1/r]) \quad W^\infty \in Gl_{d_x}(\mathbf{Z}_q[x, 1/x, 1/r])$$

be matrices such that, if we denote

$$b_j^0 = \sum_{i=0}^{d_x-1} W_{i+1, j+1}^0 y^i \quad b_j^\infty = \sum_{i=0}^{d_x-1} W_{i+1, j+1}^\infty y^i$$

for all $0 \leq j \leq d_x - 1$, then:

- $[b_0^0, \dots, b_{d_x-1}^0]$ is an integral basis for $\mathbf{Q}_q(x, y)$ over $\mathbf{Q}_q[x]$,
- $[b_0^\infty, \dots, b_{d_x-1}^\infty]$ is an integral basis for $\mathbf{Q}_q(x, y)$ over $\mathbf{Q}_q[1/x]$.

Remark

MAGMA can compute such matrices already!

Finite pole order reduction

Proposition

For all $\ell \in \mathbf{Z}_{\geq 1}$ and every vector $w \in \mathbf{Q}_q[x]^{\oplus d_x}$, there exist vectors $u, v \in \mathbf{Q}_q[x]^{\oplus d_x}$ with $\deg(v) < \deg(r)$, such that

$$\frac{\sum_{i=0}^{d_x-1} w_i b_i^0}{r^\ell} \frac{dx}{r} = d \left(\frac{\sum_{i=0}^{d_x-1} v_i b_i^0}{r^\ell} \right) + \frac{\sum_{i=0}^{d_x-1} u_i b_i^0}{r^{\ell-1}} \frac{dx}{r}.$$

Remark

By repeatedly applying this proposition, we can represent any cohomology class $\in H_{rig}^1(U)$ by a 1-form that is logarithmic at all points $P \in \mathcal{X} \setminus \mathcal{U}$ with $x(P) \neq \infty$. After a precomputation, each reduction step corresponds to a matrix multiplication. One can play the same game at the points $P \in \mathcal{X} \setminus \mathcal{U}$ with $x(P) = \infty$.

p -adic precision

We can only compute to finite p -adic precision (i.e. modulo some p^N). It follows from the Weil conjectures that if we know $Z(X, T)$ to high enough precision, then we know it exactly.

Every time we divide by p , we lose a digit of p -adic precision.

We need to bound this loss of p -adic precision at every step in the algorithm. For example in the cohomological reductions.

p -adic precision: finite pole order reduction

Proposition

Let $\omega \in \Omega^1(\mathcal{U})$ be of the form

$$\omega = \frac{\sum_{i=0}^{d_x-1} w_i y^i}{r^\ell} \frac{dx}{r},$$

with $\ell \in \mathbf{Z}_{\geq 1}$, $w \in \mathbf{Z}_q[x]^{\oplus d_x}$ and $\deg(w) < \deg(r)$. We define

$$e = \max\{e_P \mid P \in \mathcal{X} \setminus \mathcal{U}, x(P) \neq \infty\},$$

where e_P denotes the ramification index of x at P .

If we represent the class of ω in $H_{rig}^1(U)$ by $\left(\sum_{i=0}^{d_x-1} u_i y^i\right) \frac{dx}{r}$, with $u \in \mathbf{Q}_q[x]^{\oplus d_x}$, then

$$p^{\lfloor \log_p(\ell e) \rfloor} u \in \mathbf{Z}_q[x]^{\oplus d_x}.$$

Our algorithm

We can now follow the same steps as in Kedlaya's algorithm. Let d_x be the degree of $Q(x, y)$ in y and d_y the degree in x .

Theorem

Our algorithm runs in

$$\text{time: } \tilde{O}(pd_x^6 d_y^4 n^3) \quad \text{space: } \tilde{O}(pd_x^4 d_y^3 n^3)$$

Remark

We have implemented this algorithm completely. MAGMA code (packages `pcc_p` and `pcc_q`) can be found at:

https://perswww.kuleuven.be/jan_tuitman/

Projects

Short term (months):

- With Wouter Castryck: construct models and lifts of curves of genus at most 5 with d_x as small as possible. This leads to faster point counting (and is interesting in itself).
- With Jennifer Balakrishnan: adapt the algorithm to apply it to the problem of Coleman integration and the Chabauty method (finding points on curves over number fields/proving they do not exist).

Long term (years):

- Developing $\tilde{O}(p^{1/2})$ and average polynomial time versions of the algorithm, following the ideas of David Harvey (who has obtained such improvements for hyperelliptic curves).