Counting points on curves using a map to $\mathbb{P}^1$.

Jan Tuitman, KU Leuven

February 24, 2014

Zeta functions

Suppose that

- $\mathbf{F}_q$ finite field of cardinality $q = p^n$.
- $X/\mathbf{F}_q$ a smooth proper algebraic curve of genus g.

Recall that the zeta function of X is defined as

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

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

The defining equation

We let $X/\mathbf{F}_q$ denote the smooth projective curve given by the (affine) equation

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

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

We let $\mathcal{Q} \in \mathbf{Z}_q[x]$ denote a lift of Q containing the same monomials.

Proposition

The $\mathbf{Z}_q[x]$-module $\mathbf{Z}_q[x,y]/(\mathcal{Q})$ is free with basis $[1, y, \ldots, y^{d-1}]$.

The discriminant

Definition

We let $\Delta(x) \in \mathbf{Z}_q[x]$ *denote 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* $r = \Delta/(\gcd(\Delta, \frac{d\Delta}{dx}))$.

Some notation

We define

$$\mathcal{S} = \mathbf{Z}_q[x, 1/r], \qquad \mathcal{R} = \mathbf{Z}_q[x, 1/r, y]/(\mathcal{Q}),$$
$$\mathcal{S}^\dagger = \mathbf{Z}_q\langle x, 1/r\rangle^\dagger, \qquad \mathcal{R}^\dagger = \mathbf{Z}_q\langle x, 1/r, y\rangle^\dagger/(\mathcal{Q}),$$
$$\mathcal{V} = \operatorname{Spec}\mathcal{S}, \qquad \mathcal{U} = \operatorname{Spec}\mathcal{R},$$

so that x defines a finite étale morphism from $\mathcal{U}$ to $\mathcal{V}$. Moreover, we denote the special and generic fibres by

$$U = \mathcal{U} \otimes \mathbf{F}_q, \quad V = \mathcal{V} \otimes \mathbf{F}_q \quad \mathbb{U} = \mathcal{U} \otimes \mathbf{Q}_q, \quad \mathbb{V} = \mathcal{V} \otimes \mathbf{Q}_q.$$

Assumption I

Assumption

1. *There exists a smooth proper curve* $\mathcal{X}$ *over* $\mathbf{Z}_q$ *and a smooth relative divisor* $\mathcal{D}_{\mathcal{X}}$ *on* $\mathcal{X}$ *such that* $\mathcal{U} = \mathcal{X} \setminus \mathcal{D}_{\mathcal{X}}$.
2. *There exists a smooth relative divisor* $\mathcal{D}_{\mathbf{P}^1}$ *on* $\mathbf{P}^1_{\mathbf{Z}_q}$ *such that* $\mathcal{V} = \mathbf{P}^1_{\mathbf{Z}_q} \setminus \mathcal{D}_{\mathbf{P}^1}$.

Assumption II

Assumption

The zero locus of $\mathcal{Q}$ in $\mathbf{A}^2_{\mathbf{Q}_q}$ is smooth.

Proposition

The element $s = r/\frac{\partial \mathcal{Q}}{\partial y}$ of $\mathbf{Q}_q(x, y)$ is contained in $\mathbf{Z}_q[x, y]/(\mathcal{Q})$.

Sketch of the proof: $\Delta/\frac{\partial \mathcal{Q}}{\partial y}$ is contained in $\mathbf{Z}_q[x, y]/(\mathcal{Q})$ by the definition of Δ as the determinant of the Sylvester matrix. By the assumption, $[1, y, \ldots, y^{d-1}]$ is an integral basis of $\mathbf{Q}_q[x, y]/(\mathcal{Q})$ over $\mathbf{Q}_q[x]$. So for any monic irreducible polynomial $\pi \in \mathbf{Z}_q[x]$, the element $\frac{\partial \mathcal{Q}}{\partial y}/\pi$ of $\mathbf{Q}_q(x, y)$ is not integral at (π) because of the term $(d/\pi)y^{d-1}$, hence its inverse $\pi/\frac{\partial \mathcal{Q}}{\partial y}$ is integral (even zero) at (π). Since $\prod_{\pi|\Delta} \pi = r$, this proves the Proposition.

Assumption III

Assumption

We assume that a matrix $W^{\infty} \in Gl_d(\mathbf{Z}_q[x, x^{-1}])$ *is known such that if we denote*

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

for all $0 \leq j \leq d-1$, *then* $[b_0^{\infty}, \cdots, b_{d-1}^{\infty}]$ *is an integral basis for* $\mathbf{Q}_q(x, y)$ *over* $\mathbf{Q}_q[x^{-1}]$.

When Q is nondegenerate with respect to its Newton polygon, W^{∞} can be written down directly. In general there are good algorithms to compute integral bases in function fields.

Example: hyperelliptic case

p odd

$\mathcal{Q} = y^2 - f(x)$

$f(x) \in \mathbf{Z}_q[x]$ of degree $2g+1$ with $(f, f') = 1$

Assumptions are satisfied

$\Delta(x) = r(x) = f(x) = y^2$

$s(x, y) = y/2$

$$W^{\infty} = \begin{pmatrix} 1 & 0 \\ 0 & x^{-(g+1)} \end{pmatrix}$$

Frobenius lift

Let σ be the p-power Frobenius on $\mathbf{Z}_q$ Define sequences $(\alpha_i)_{i\geq 0}$, $(\beta_i)_{i\geq 0}$, with $\alpha_i \in S^\dagger$ and $\beta_i \in \mathcal{R}^\dagger$, by the following recursion:

$$\begin{aligned}
\alpha_0 &= \frac{1}{r^p}, \\
\beta_0 &= y^p, \\
\alpha_{i+1} &= \alpha_i(2 - \alpha_i r^\sigma(x^p)) && \pmod{p^{2^{i+1}}}, \\
\beta_{i+1} &= \beta_i - \mathcal{Q}^\sigma(x^p, \beta_i)s^\sigma(x^p, \beta_i)\alpha_i && \pmod{p^{2^{i+1}}}.
\end{aligned}$$

Then one easily checks that the σ-semilinear ringhomomorphism $\mathrm{F}_p : \mathcal{R}^\dagger \to \mathcal{R}^\dagger$ defined by

$$\mathrm{F}_p(x) = x^p, \qquad \mathrm{F}_p(\frac{1}{r}) = \lim_{i\to\infty} \alpha_i, \qquad \mathrm{F}_p(y) = \lim_{i\to\infty} \beta_i,$$

is a Frobenius lift.

Effective convergence bounds

Proposition

Let $N \in \mathbf{N}$. Then modulo p^N:

1. $\mathrm{F}_p(1/r)$ *is congruent to* $\sum_{i=p}^{pN} \frac{\rho_i(x)}{r^i}$, *where for all* $p \leq i \leq pN$ *the polynomial* $\rho_i \in \mathbf{Z}_q[x]$ *satisfies* $\deg(\rho_i) < \deg(r)$.
2. $\mathrm{F}_p(y^i)$ *is congruent to* $\sum_{j=0}^{d-1} \phi_{i,j}(x) y^j$, *where*

$$\phi_{i,j} = \sum_{k=0}^{p(N-1)} \frac{\phi_{i,j,k}(x)}{r^k},$$

for all $0 \leq i, j \leq d-1$ *and* $\phi_{i,j,k} \in \mathbf{Z}_q[x]$ *satisfies* $\deg(\phi_{i,j,0}) \leq -\operatorname{ord}_\infty(W^\infty) - p \operatorname{ord}_\infty((W^\infty)^{-1})$ *and* $\deg(\phi_{i,j,k}) < \deg(r)$, *for all* $0 \leq i, j \leq d-1$ *and* $1 \leq k \leq p(N-1)$.

Rigid cohomology

We define the overconvergent Kähler differentials

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

and the overconvergent De Rham complex

$$\Omega^\bullet_{\mathcal{R}^\dagger} : \quad 0 \longrightarrow \mathcal{R}^\dagger \xrightarrow{d} \Omega_{\mathcal{R}^\dagger} \longrightarrow 0.$$

We then have

$$H^1_{\text{rig}}(U) = H^1(\Omega^\bullet_{\mathcal{R}^\dagger} \otimes \mathbf{Q}_q) = \text{coker}(d) \otimes \mathbf{Q}_q.$$

Computing in the cohomology: finite points

Proposition

For all $\ell \in \mathbf{N}$ *and every vector* $w \in \mathbf{Q}_q[x]^{\oplus d}$, *there exist (unique) vectors* $u, v \in \mathbf{Q}_q[x]^{\oplus d}$ *with* $\deg(v) < \deg(r)$, *such that*

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

Sketch of the proof: r is separable, so r' is invertible in $\mathbf{Q}_q[x]/(r)$. v has to satisfy $\left(\frac{M}{r'} - \ell I\right) v \equiv \frac{u}{r'} \pmod{r}$ over $\mathbf{Q}_q[x]/(r)$. The finite exponents of $(M/r)dx$ are contained in $[0, 1)$, hence $\det(\ell I - M/r')$ is invertible in $\mathbf{Q}_q[x]/(r)$, so there is a unique solution v. We now take

$$u = \frac{w - \left(M - \ell r' I\right) v}{r} - \frac{dv}{dx}.$$

Precision loss: finite points

Proposition

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

$$\omega = \frac{\sum_{i=0}^{d-1} w_i(x) y^i}{r^\ell} \frac{dx}{r},$$

where $\ell \in \mathbf{N}$ *and* $w_i \in \mathbf{Z}_q[x]$ *satisfies* $\deg(w_i) < \deg(r)$ *for all* $0 \leq i \leq d-1$. *We define* $e = \max\{e_P | P \in \mathcal{X} \setminus \mathcal{U}, x(P) \neq \infty\}$. *If we represent the class of* ω *in* $H^1_{rig}(U)$ *by* $\left(\sum_{i=0}^{d-1} u_i y^i\right) \frac{dx}{r}$ *with* $u \in \mathbf{Q}_q[x]^{\oplus d}$ *using the previous Proposition then*

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

Computing in the cohomology: infinite points

Proposition

For every vector $w \in \mathbf{Q}_q[x, x^{-1}]^{\oplus d}$ *with*

$$\mathrm{ord}_\infty(w) \leq -\deg(r),$$

there exist vectors $u, v \in \mathbf{Q}_q[x, x^{-1}]^{\oplus d}$ *with* $\mathrm{ord}_\infty(u) > \mathrm{ord}_\infty(w)$ *such that*

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

We do not actually want to reduce this far as new finite poles can appear, we stop when $\mathrm{ord}_\infty(u) > \mathrm{ord}_0(W^\infty) - \deg(r) + 1$.

Precision loss: infinite points

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

$$\omega = (\sum_{i=0}^{d-1} w_i(x, x^{-1}) b_i^{\infty}) \frac{dx}{r},$$

where $w \in \mathbf{Z}_q[x, x^{-1}]^{\oplus d}$ satisfies
$\mathrm{ord}_{\infty}(w) \leq \mathrm{ord}_0(W^{\infty}) - \deg(r) + 1$. We define

$$m = -\mathrm{ord}_{\infty}(w) - \deg(r) + 1,$$
$$e_{\infty} = \max\{e_P | P \in \mathcal{X} \setminus \mathcal{U}, x(P) = \infty\}.$$

If we represent the class of ω in $H^1_{\mathrm{rig}}(U)$ by $\left(\sum_{i=0}^{d-1} u_i y^i\right) \frac{dx}{r}$, with $u \in \mathbf{Q}_q[x, x^{-1}]^{\oplus d}$ such that $\mathrm{ord}_{\infty}(u) > \mathrm{ord}_0(W^{\infty}) - \deg(r) + 1$ using the previous proposition, then

$$p^{\lfloor \log_p(m e_{\infty}) \rfloor} u \in \mathbf{Z}_q[x, x^{-1}]^{\oplus d}.$$

Algorithm

1. The Frobenius lifting and pole order reduction procedures allow us to compute the action of the Frobenius F_p on $H^1_{\mathrm{rig}}(U)$ to any desired p-adic precision.
2. Using a cohomological residue map, we can identify $H^1_{\mathrm{rig}}(X)$ inside $H^1_{\mathrm{rig}}(U)$ as in Kedlaya's algorithm. We then only have to compute with this subspace.
3. So we can compute $\chi(T) = \det(1 - \mathrm{F}_p\, T | H^1_{\mathrm{rig}}(X))$ to any desired p-adic precision.
4. Explicit bounds on the absolute values of the coefficients of $\chi(T)$ are known from the Weil conjectures, so if known to high enough p-adic precision then $\chi(T)$ is determined exactly.

Complexity

Recall that $d = \deg_y(Q)$, $\delta = \deg_x(Q)$, $q = p^n$.

Theorem

The runtime of the algorithm is $\tilde{\mathcal{O}}(pd^6\delta^4 n^3)$.

Note that this is exponential in $\log(p)$, so bad for gathering statistics on Frobenius distributions.

However, it is very good (polynomial, small exponents) in everything else.

It would be interesting to try and combine the average polynomial time machinery of Harvey with this algorithm to gather statistics on Frobenius distributions of (more) complicated curves.

Implementation

I have implemented a somewhat restricted version of the algorithm in MAGMA.

It seems to run 2-3 orders of magnitude faster than alternatives like Castryck-Denef-Vercauteren for nondegenerate curves, while it can be applied more generally.