Getting precise about precision

Kiran S. Kedlaya (kedlaya@mit.edu)

Department of Mathematics, Massachusetts Institute of Technology

Effective methods in p-adic cohomology
Oxford, March 19, 2010

These slides available at http://math.mit.edu/~kedlaya/papers/.

Supported by NSF, DARPA, MIT, IAS.

Contents

Contents

Computing with real numbers

It is wildly impractical (if not outright impossible) to compute with exact real numbers. Instead, one typically uses *floating-point approximations*, in which only a limited number of digits are carried.

These are sufficient for many practical computations where answers need only be correct with some reasonable probability. For extra reliability, one can increase the number of digits carried.

However, floating-point calculations do give reproducible results, so one can use them in establishing proofs. One approach is to attach error bounds to floating-point numbers, yielding *interval arithmetic*.

Does precision matter in p-adic cohomology?

When working in the ring $\mathbb{Z}_p/p^n\mathbb{Z}_p$, all computations are exact. But when working in $\mathbb{Z}_p$ or $\mathbb{Q}_p$, one again makes only approximate calculations.

For numerical experiments, approximate answers are often sufficient. For provable calculations, one must add error estimates, but the difference between weak and strong error bounds often appears in asymptotics only as a constant factor.

So does precision matter?

Precision matters!

For provable computations *in practice*, bad precision estimates often lead to excessive time and memory consumption. In many cases, these can push a computation over the feasibility boundary. (This is particularly true in dimension greater than 1.)

Even for experimental computations, a proper understanding of precision allows one to optimize parameters while still retaining a high probability of reasonable results.

But my favorite reason to study precision estimates in p-adic cohomology is...

Precision matters!

For provable computations *in practice*, bad precision estimates often lead to excessive time and memory consumption. In many cases, these can push a computation over the feasibility boundary. (This is particularly true in dimension greater than 1.)

Even for experimental computations, a proper understanding of precision allows one to optimize parameters while still retaining a high probability of reasonable results.

But my favorite reason to study precision estimates in p-adic cohomology is...

... it involves some very interesting mathematics!

Plan of the talk

A typical application of p-adic cohomology to compute zeta functions would involve computation of the Frobenius action on a basis of a cohomology group, extraction of a p-adic approximation of the characteristic polynomial of the Frobenius matrix, and reconstruction of a Weil polynomial.

We will work through these steps in reverse order.

Contents

Weil polynomials and zeta functions

Let X be a variety over $\mathbb{F}_q$. The *zeta function* is the Dirichlet series

$$\zeta_X(s) = \prod_{x \in X \text{closed}} (1 - \#\kappa(x)^{-s})^{-1},$$

for $\kappa(x)$ the residue field of x. It can be represented as

$$\frac{P_1(T)P_3(T)\cdots}{P_0(T)P_2(T)\cdots} \qquad (T = q^{-s}, P_i(T) \in 1 + T\mathbb{Z}[T]).$$

If X is smooth proper, the roots of $P_i(T)$ in $\mathbb{C}$ have absolute value $q^{-i/2}$ (i.e., the reverse of P_i is a *Weil polynomial*), and

$$P_i(T) = \det(1 - FT, H^i(X))$$

for F the Frobenius action on the i-th rigid cohomology $H^i(X)$. By computing in $H^i(X)$, we can obtain a p-adic approximation of P_i.

Recovering a Weil polynomial from an approximation

How good an approximation is needed to determine $P_i(T)$ uniquely?

For instance, say X is a curve of genus g, so

$$P_1(T) = 1 + a_1 T + \cdots + a_g T^g + \cdots + a_{2g} T^{2g}$$

and the higher coefficients satisfy $a_{g+i} = q^i a_{g-i}$. For $i = 1, \ldots, g$, we have

$$|a_i| \leq \binom{2g}{i} q^{i/2}$$

so $P_1(T)$ is uniquely determined by its residue modulo q^n as long as

$$q^n > 2\binom{2g}{g} q^{g/2}.$$

But this is not optimal!

Recovering a Weil polynomial from an approximation

Write $P_1(T) = (1 - \alpha_1 T) \cdots (1 - \alpha_{2g} T)$, and define the power sums

$$s_i = \alpha_1^i + \cdots + \alpha_{2g}^i.$$

The s_i are integers of norm at most $2gq^{i/2}$, and satisfy the Newton-Viète identities

$$s_i + a_1 s_{i-1} + \cdots + a_{i-1} s_1 + i a_i = 0.$$

Once $a_1, \ldots, a_{i-1}$ are known, so are $s_1, \ldots, s_{i-1}$, so we can determine a_i by determining s_i. Consequently, $P_1(T)$ is uniquely determined by its residue modulo q^n as long as

$$q^n > \max\left\{ \frac{4g}{i} q^{i/2} : i = 1, \ldots, g \right\}.$$

Refinements

I have Sage code to find all Weil polynomials obeying a congruence. Using such code, one can determine experimentally how much precision in the congruence is needed to uniquely determine the Weil polynomial; it is typically slightly less than the best known bounds (by one or two digits).

This gap grows when one adds extra constraints on the Weil polynomial (e.g., if X is a curve whose Jacobian has extra endomorphisms).

However, fixing some initial coefficients of $P_i(T)$ (by explicit point counting) apparently *does not reduce* precision requirements in most cases.

Contents

Setup

Let A be a square matrix over $\mathbb{Z}_p$ (or more generally, any complete discrete valuation ring). How sensitive is $\det(1 - TA)$ to perturbations of A? In other words, if B is another square matrix of the same size, how do bounds on B translate into bounds on $\det(1 - T(A + B))$?

Example: If B is divisible by p^n, then so is each coefficient of $\det(1 - TA) - \det(1 - T(A + B))$.

In practice, this is often not optimal!

Enter the Hodge polygon

Suppose X is smooth proper over $\mathbb{Z}_p$ (for example), and suppose $p > i$ (for simplicity). Then

$$H^i_{\mathrm{crys}}(X_{\mathbb{F}_p}, \mathbb{Z}_p) \cong H^i_{\mathrm{dR}}(X)$$

carries a *Hodge filtration*

$$0 = \mathrm{Fil}^{-1} \subseteq \cdots \subseteq \mathrm{Fil}^i = H^i_{\mathrm{dR}}(X)$$

with $\mathrm{Fil}^j / \mathrm{Fil}^{j-1} \cong H^j(X, \Omega^{i-j}_{X/\mathbb{Z}_p})$. Frobenius does not preserve this filtration, but the image of Fil^j is divisible by p^{i-j}.

By computing with a basis of $H^i_{\mathrm{dR}}(X)$ compatible with the Hodge filtration, we pick up some p-adic divisibilities that help reduce the perturbation in the characteristic polynomial.

Exploiting the Hodge polygon

Let's go back to our square matrices A and B and impose some p-adic divisibility conditions on some columns of A. For instance, suppose

$$A = \begin{pmatrix} p^2* & p* & p* & p* & p* & * \\ p^2* & p* & p* & p* & p* & * \\ p^2* & p* & p* & p* & p* & * \\ p^2* & p* & p* & p* & p* & * \\ p^2* & p* & p* & p* & p* & * \\ p^2* & p* & p* & p* & p* & * \end{pmatrix}$$

and B is divisible by p^m. Then

$$\det(1-TA)-\det(1-T(A+B)) = p^m*T+p^m*T^2+p^{m+1}*T^3+p^{m+2}*T^4+\dots,$$

e.g., by writing the coefficient of T^k in $\det(1 - TA)$ as a signed sum of principal $k \times k$ minors of A.

Example: K3 surfaces

Let X be a K3 surface over $\mathbb{F}_q$, e.g., a smooth quartic surface in $\mathbb{P}^3$. Then $P_2(T) = (1 - qT)P_2^{\text{prim}}(T)$ where $P_2^{\text{prim}}(T)$ has degree 21. By symmetry, $P_2^{\text{prim}}(T)$ is determined by its coefficients up to T^{10}, so one expects to need about q^{10} precision in the Frobenius matrix.

However, the Hodge numbers in primitive middle cohomology are 1, 19, 1. So precision q^m in the Frobenius matrix gives precision q^{m+k-2} in the coefficient of T^k for $k = 2, \ldots, 10$. Hence one needs only about q^3 precision in the Frobenius matrix!

Contents

The Monsky-Washnitzer method

Let $\overline{X}$ be a smooth proper variety over $\mathbb{Z}_p$ (say), and let $\overline{Z}$ be a divisor of simple normal crossings. We can use Monsky-Washnitzer cohomology to find the zeta function of $\overline{U} = \overline{X} \setminus \overline{Z}$, by lifting to a smooth pair (X, Z) over $\mathbb{Q}_p$ and computing in the weak completion of the coordinate ring of $U = X \setminus Z$. In cases of interest, the de Rham cohomology of U will be described by a reduction rule for differential forms.

The action of a Frobenius lift is given by some p-adically convergent power series, which we truncate for the computation. However, the power of p dividing the unseen remainder is typically reduced by the process of reducing it into basis form.

Example: hyperelliptic curves

Consider the hyperelliptic curve $y^2 = P(x)$, where P is is a polynomial of degree $2g+1$ over $\mathbb{Z}_p$ with no repeated roots modulo p. Use the Frobenius lift $F : x \mapsto x^p$. For $i = 0, \dots, 2g-1$, write

$$F\left(\frac{x^i\,dx}{y}\right) = Q(x)y + \sum_{s=1}^{\infty} \frac{p^i R_i(x)\,dx}{y^{2s-1}},$$

with $\deg(R_i) \leq 2g$. One reduces the pole orders using the relation

$$0 \equiv d\left(\frac{A(x)}{y^{2s-1}}\right) = \frac{A'(x)\,dx}{y^{2s-1}} - \frac{(2s-1)A(x)P'(x)\,dx}{2y^{2s+1}}.$$

In doing so, $R_i(x)\,dx/y^{2s-1}$ acquires a denominator no worse than p^m for $m = \lfloor \log_p(2s+1) \rfloor$. (Proof: expand formally at each Weierstrass point.)

One may treat Q similarly by expanding at ∞. It helps to pass to a *crystalline basis*, e.g., $x^i\,dx/y^3$.

Smooth projective hypersurfaces

Assume for simplicity that $p \geq n-2$. Consider the complement in $X = \mathbb{P}^n$ of a smooth hypersurface Z of degree d defined by a polynomial $P(x_0, \dots, x_n)$. Put

$$\Omega = \sum_{i=0}^{n} (-1)^i x_i dx_0 \wedge \cdots \wedge \widehat{dx_i} \wedge \cdots \wedge dx_n.$$

The top cohomology is a quotient of the space of homogeneous degree 0 forms $A\Omega/P^d$ by all relations of the form

$$\frac{\partial A}{\partial x_i}\frac{\Omega}{P^d} - \frac{\partial P}{\partial x_i}\frac{mA\Omega}{P^{d+1}} \qquad (i = 0, \dots, n).$$

These can be used to reduce the pole order d.

We use the Frobenius lift $F : x_i \mapsto x_i^p$.

Precision loss

We can find a basis of cohomology in which each element is $A\Omega/P^d$ for some monomial A and some $d \in \{1, \dots, n\}$. As in the hyperelliptic case, when we apply Frobenius to a basis vector, we get an p-adically convergent infinite series, which we truncate and then reduce in cohomology.

We need to estimate the contribution to the Frobenius matrix from the omitted terms. This amounts to asking: when one reduces $A\Omega/P^d$ to a linear combination of basis vectors, how much of a denominator is introduced?

First answer: the denominator is at most p^m with

$$\sum_{i=1}^{n} \lfloor \log_p \max\{1, d-1\} \rfloor.$$

Local expansions revisited

To get the previous bound, we would like to "expand formally", but this requires thinking in terms of *sheaves*.

For $d \geq 0$, consider the map

$$\Omega^{\cdot}_{(X,Z)} \to \Omega^{\cdot}_{(X,Z)}(dZ)$$

of complexes of sheaves, in which the left side is the *logarithmic* de Rham complex and the right side allows poles of order d. Then pass to the homology sheaves

$$\mathcal{H}^i \to \mathcal{H}^i_d.$$

We calculate formally in local coordinates that the cokernel of each map is killed by $\mathrm{lcm}(1, \dots, d)$. We then get the previous bound by computing the de Rham cohomology of U as the hypercohomology of $\Omega^{\cdot}_{(X,Z)}$.

Refining the bound

The preceding analysis gave a denominator bound of p^m with $m \sim n \log_p d$. We can refine this to $m \sim (n-1) \log_p d$.

Sketch of proof: at each step with d divisible by p, approximate each monomial of A with all exponents congruent to -1 mod p by an element of the image of Frobenius. Using the fact that the Frobenius matrix is divisible by p (from the comparison between the de Rham cohomology of U and Z), we get savings in reducing these monomials. Each other monomial can be raised once (picking up a factor of p) and then reduced.

Can we do better?

It appears that the bounds we obtained are still not optimal. For instance, consider the smooth quartic (K3) surface in $\mathbb{P}^3$ over $\mathbb{F}_3$ defined by

$$x^4 - xy^3 + xy^2w + xyzw + xyw^2 - xzw^2 + y^4 + y^3w - y^2zw + z^4 + w^4.$$

To get final precision $3^3, 3^4, 3^5, 3^6$ in the Frobenius matrix, our refined bounds suggest that we need to truncate the Frobenius action modulo $3^7, 3^{10}, 3^{11}, 3^{12}$. However, experimentally we only need to truncate modulo $3^6, 3^7, 3^9, 3^{10}$.

This matters in runtimes. For instance, going from 3^6 to 3^7 added a factor of 3.

References

For recovering Weil polynomials: KSK, Search techniques for root-unitary polynomials.

For perturbations of characteristic polynomials: KSK, p-adic differential equations, chapter 4.

For precision loss analysis using sheaves: Abbott, KSK, Roe, Bounding Picard numbers of surfaces using p-adic cohomology.