A survey of 15 years of p-adic point counting

Kiran S. Kedlaya

Department of Mathematics, University of California, San Diego
kedlaya@ucsd.edu
http://kskedlaya.org/slides/

p-adic Methods in Number Theory
A Conference Inspired by the Mathematics of Robert Coleman
UC Berkeley, Saturday, May 30, 2015

Robert Coleman, computational number theorist?

In his later career, Robert's mathematical productivity was largely enabled by the existence of the personal computer (e.g., to produce overhead transparencies and lecture notes).

However, while Robert's work does not include (m)any papers which are explicitly about computational aspects of number theory...

... I nonetheless contend that Robert's inimitable *style* of mathematics is very much in the spirit of computational number theory, and number theorists with computational tendencies would be well-served by understanding some of Robert's insights.

Robert Coleman, computational number theorist?

In his later career, Robert's mathematical productivity was largely enabled by the existence of the personal computer (e.g., to produce overhead transparencies and lecture notes).

However, while Robert's work does not include (m)any papers which are explicitly about computational aspects of number theory...

... I nonetheless contend that Robert's inimitable *style* of mathematics is very much in the spirit of computational number theory, and number theorists with computational tendencies would be well-served by understanding some of Robert's insights.

Robert Coleman, computational number theorist?

In his later career, Robert's mathematical productivity was largely enabled by the existence of the personal computer (e.g., to produce overhead transparencies and lecture notes).

However, while Robert's work does not include (m)any papers which are explicitly about computational aspects of number theory...

... I nonetheless contend that Robert's inimitable *style* of mathematics is very much in the spirit of computational number theory, and number theorists with computational tendencies would be well-served by understanding some of Robert's insights.

Zeta functions of algebraic varieties

Let $\mathbb{F}_q$ be a finite field of characteristic p. After (Artin, Schmidt, and) Weil, we define the *zeta function* of a variety X over $\mathbb{F}_q$ as the formal Dirichlet series

$$\zeta(X, s) = \prod_x (1 - \#\kappa(x)^{-s})^{-1},$$

where x runs over closed points of X and $\kappa(x)$ denotes the residue field. (Equivalently, x runs over Galois orbits of $\overline{\mathbb{F}}_q$-rational points and $\kappa(x)$ denotes the minimal field of definition.)

From now on, we write ζ as a formal power series in $T = q^{-s}$. Then

$$\zeta(X, T) = \exp\left(\sum_{n=1}^{\infty} \frac{T^n}{n} \#X(\mathbb{F}_{q^n})\right).$$

Zeta functions of algebraic varieties

Let $\mathbb{F}_q$ be a finite field of characteristic p. After (Artin, Schmidt, and) Weil, we define the *zeta function* of a variety X over $\mathbb{F}_q$ as the formal Dirichlet series

$$\zeta(X, s) = \prod_x (1 - \#\kappa(x)^{-s})^{-1},$$

where x runs over closed points of X and $\kappa(x)$ denotes the residue field. (Equivalently, x runs over Galois orbits of $\overline{\mathbb{F}}_q$-rational points and $\kappa(x)$ denotes the minimal field of definition.)

From now on, we write ζ as a formal power series in $T = q^{-s}$. Then

$$\zeta(X, T) = \exp\left(\sum_{n=1}^{\infty} \frac{T^n}{n} \#X(\mathbb{F}_{q^n}) \right).$$

Examples of zeta functions

From the formula

$$\zeta(X,T) = \exp\left(\sum_{n=1}^{\infty} \frac{T^n}{n} \#X(\mathbb{F}_{q^n})\right),$$

one can compute $\zeta(X,T)$ in some explicit examples. For one:

$$\zeta(\mathbb{P}^d_{\mathbb{F}_q}, T) = \frac{1}{(1-T)(1-qT)\cdots(1-q^dT)}.$$

For another, if X is an elliptic curve over $\mathbb{F}_q$, then

$$\zeta(X,T) = \frac{1-aT+qT^2}{(1-T)(1-qT)}, \qquad a = q+1-\#X(\mathbb{F}_q).$$

Based on these (and more) examples, Weil predicted that $\zeta(X,T)$ obeys analogues of the properties of the Riemann zeta function (analytic continuation, functional equation, Riemann hypothesis).

Examples of zeta functions

From the formula

$$\zeta(X,T) = \exp\left(\sum_{n=1}^{\infty} \frac{T^n}{n} \#X(\mathbb{F}_{q^n})\right),$$

one can compute $\zeta(X,T)$ in some explicit examples. For one:

$$\zeta(\mathbb{P}^d_{\mathbb{F}_q}, T) = \frac{1}{(1-T)(1-qT)\cdots(1-q^dT)}.$$

For another, if X is an elliptic curve over $\mathbb{F}_q$, then

$$\zeta(X,T) = \frac{1-aT+qT^2}{(1-T)(1-qT)}, \qquad a = q+1-\#X(\mathbb{F}_q).$$

Based on these (and more) examples, Weil predicted that $\zeta(X,T)$ obeys analogues of the properties of the Riemann zeta function (analytic continuation, functional equation, Riemann hypothesis).

Examples of zeta functions

From the formula

$$\zeta(X, T) = \exp\left(\sum_{n=1}^{\infty} \frac{T^n}{n} \#X(\mathbb{F}_{q^n})\right),$$

one can compute $\zeta(X, T)$ in some explicit examples. For one:

$$\zeta(\mathbb{P}^d_{\mathbb{F}_q}, T) = \frac{1}{(1-T)(1-qT)\cdots(1-q^d T)}.$$

For another, if X is an elliptic curve over $\mathbb{F}_q$, then

$$\zeta(X, T) = \frac{1 - aT + qT^2}{(1-T)(1-qT)}, \qquad a = q + 1 - \#X(\mathbb{F}_q).$$

Based on these (and more) examples, Weil predicted that $\zeta(X, T)$ obeys analogues of the properties of the Riemann zeta function (analytic continuation, functional equation, Riemann hypothesis).

Rationality of the zeta function

The first of the *Weil conjectures* on zeta functions of algebraic varieties is:

Theorem

The power series $\zeta(X, T)$ represents a rational function of T.

This is widely known as a consequence of the construction of *étale cohomology* by Grothendieck et al. However, that was not the first proof!

Theorem (Dwork, 1960)

The power series $\zeta(X, T)$ is p-adic meromorphic: it is the ratio of two power series over $\mathbb{Q}_p$ with infinite radii of convergence.

Since $\zeta(X, T)$ converges for $T \in \mathbb{C}$ small (trivially), an argument of Borel (1894) then shows that $\zeta(X, T)$ is rational.

Rationality of the zeta function

The first of the *Weil conjectures* on zeta functions of algebraic varieties is:

Theorem

The power series $\zeta(X, T)$ represents a rational function of T.

This is widely known as a consequence of the construction of *étale cohomology* by Grothendieck et al. However, that was not the first proof!

Theorem (Dwork, 1960)

The power series $\zeta(X, T)$ is p-adic meromorphic: it is the ratio of two power series over $\mathbb{Q}_p$ with infinite radii of convergence.

Since $\zeta(X, T)$ converges for $T \in \mathbb{C}$ small (trivially), an argument of Borel (1894) then shows that $\zeta(X, T)$ is rational.

Rationality of the zeta function

The first of the *Weil conjectures* on zeta functions of algebraic varieties is:

Theorem

The power series $\zeta(X, T)$ represents a rational function of T.

This is widely known as a consequence of the construction of *étale cohomology* by Grothendieck et al. However, that was not the first proof!

Theorem (Dwork, 1960)

The power series $\zeta(X, T)$ is p-adic meromorphic: it is the ratio of two power series over $\mathbb{Q}_p$ with infinite radii of convergence.

Since $\zeta(X, T)$ converges for $T \in \mathbb{C}$ small (trivially), an argument of Borel (1894) then shows that $\zeta(X, T)$ is rational.

Rationality of the zeta function

The first of the *Weil conjectures* on zeta functions of algebraic varieties is:

Theorem

The power series $\zeta(X, T)$ represents a rational function of T.

This is widely known as a consequence of the construction of *étale cohomology* by Grothendieck et al. However, that was not the first proof!

Theorem (Dwork, 1960)

The power series $\zeta(X, T)$ is p-adic meromorphic: it is the ratio of two power series over $\mathbb{Q}_p$ with infinite radii of convergence.

Since $\zeta(X, T)$ converges for $T \in \mathbb{C}$ small (trivially), an argument of Borel (1894) then shows that $\zeta(X, T)$ is rational.

Rationality of the zeta function

The first of the *Weil conjectures* on zeta functions of algebraic varieties is:

Theorem

The power series $\zeta(X, T)$ represents a rational function of T.

This is widely known as a consequence of the construction of *étale cohomology* by Grothendieck et al. However, that was not the first proof!

Theorem (Dwork, 1960)

The power series $\zeta(X, T)$ is p-adic meromorphic: it is the ratio of two power series over $\mathbb{Q}_p$ with infinite radii of convergence.

Since $\zeta(X, T)$ converges for $T \in \mathbb{C}$ small (trivially), an argument of Borel (1894) then shows that $\zeta(X, T)$ is rational.

Zeta functions: the computational problem

Can one produce an algorithm that, given an explicit definition of X (i.e., defining equations), returns the rational function $\zeta(X, T)$?

Yes: one can compute a bound on the degree of $\zeta(X, T)$ and then enumerate $X(\mathbb{F}_{q^n})$ for enough values of n to determine the coefficients.

But is there an *efficient* algorithm? Enumeration is impractical unless both q and the degree bound are fairly small.

If you ask the question carelessly, probably no: the length of the answer can be exponential in the length of the input. Even computing $\#X(\mathbb{F}_q)$ is dicey: one can reduce NP-complete problems (e.g., 3-SAT) to it.

However, more modest versions of this question are of both intrinsic and extrinsic interest. E.g., interest grew rapidly circa 2000 due to applications to (hyper)elliptic curve cryptography. There is also interest in computing motivic L-functions, e.g., to test special values conjectures (BSD et al.).

Zeta functions: the computational problem

Can one produce an algorithm that, given an explicit definition of X (i.e., defining equations), returns the rational function $\zeta(X, T)$?

Yes: one can compute a bound on the degree of $\zeta(X, T)$ and then enumerate $X(\mathbb{F}_{q^n})$ for enough values of n to determine the coefficients.

But is there an *efficient* algorithm? Enumeration is impractical unless both q and the degree bound are fairly small.

If you ask the question carelessly, probably no: the length of the answer can be exponential in the length of the input. Even computing $\#X(\mathbb{F}_q)$ is dicey: one can reduce NP-complete problems (e.g., 3-SAT) to it.

However, more modest versions of this question are of both intrinsic and extrinsic interest. E.g., interest grew rapidly circa 2000 due to applications to (hyper)elliptic curve cryptography. There is also interest in computing motivic L-functions, e.g., to test special values conjectures (BSD et al.).

Zeta functions: the computational problem

Can one produce an algorithm that, given an explicit definition of X (i.e., defining equations), returns the rational function $\zeta(X, T)$?

Yes: one can compute a bound on the degree of $\zeta(X, T)$ and then enumerate $X(\mathbb{F}_{q^n})$ for enough values of n to determine the coefficients.

But is there an *efficient* algorithm? Enumeration is impractical unless both q and the degree bound are fairly small.

If you ask the question carelessly, probably no: the length of the answer can be exponential in the length of the input. Even computing $\#X(\mathbb{F}_q)$ is dicey: one can reduce NP-complete problems (e.g., 3-SAT) to it.

However, more modest versions of this question are of both intrinsic and extrinsic interest. E.g., interest grew rapidly circa 2000 due to applications to (hyper)elliptic curve cryptography. There is also interest in computing motivic L-functions, e.g., to test special values conjectures (BSD et al.).

Zeta functions: the computational problem

Can one produce an algorithm that, given an explicit definition of X (i.e., defining equations), returns the rational function $\zeta(X, T)$?

Yes: one can compute a bound on the degree of $\zeta(X, T)$ and then enumerate $X(\mathbb{F}_{q^n})$ for enough values of n to determine the coefficients.

But is there an *efficient* algorithm? Enumeration is impractical unless both q and the degree bound are fairly small.

If you ask the question carelessly, probably no: the length of the answer can be exponential in the length of the input. Even computing $\#X(\mathbb{F}_q)$ is dicey: one can reduce NP-complete problems (e.g., 3-SAT) to it.

However, more modest versions of this question are of both intrinsic and extrinsic interest. E.g., interest grew rapidly circa 2000 due to applications to (hyper)elliptic curve cryptography. There is also interest in computing motivic L-functions, e.g., to test special values conjectures (BSD et al.).

Zeta functions: the computational problem

Can one produce an algorithm that, given an explicit definition of X (i.e., defining equations), returns the rational function $\zeta(X, T)$?

Yes: one can compute a bound on the degree of $\zeta(X, T)$ and then enumerate $X(\mathbb{F}_{q^n})$ for enough values of n to determine the coefficients.

But is there an *efficient* algorithm? Enumeration is impractical unless both q and the degree bound are fairly small.

If you ask the question carelessly, probably no: the length of the answer can be exponential in the length of the input. Even computing $\#X(\mathbb{F}_q)$ is dicey: one can reduce NP-complete problems (e.g., 3-SAT) to it.

However, more modest versions of this question are of both intrinsic and extrinsic interest. E.g., interest grew rapidly circa 2000 due to applications to (hyper)elliptic curve cryptography. There is also interest in computing motivic L-functions, e.g., to test special values conjectures (BSD et al.).

Zeta functions: the computational problem (continued)

Problem (Ill-posed)

Fix a positive integer n. Is there an algorithm that, given an algebraic variety X of dimension n, returns the rational function $\zeta(X, T)$ in "polynomial time"?

To quantify "polynomial time" (in the length of the input), we must specify an input mechanism for X. Note that if $Z \subseteq X$ is a closed subscheme, then

$$\zeta(X, T) = \zeta(Z, T)\zeta(X - Z, T);$$

we can thus reduce to working with affine hypersurfaces.

Problem (Well-posed, and open!)

Fix a positive integer n. Is there an algorithm that, given $f \in \mathbb{F}_q[x_1, \ldots, x_n]$ of degree d, returns the rational function $\zeta(X, T)$ for $X = \operatorname{Spec} \mathbb{F}_q[x_1, \ldots, x_n]/(f)$ in time $\operatorname{poly}(d, \log q)$?

Zeta functions: the computational problem (continued)

Problem (Ill-posed)

Fix a positive integer n. Is there an algorithm that, given an algebraic variety X of dimension n, returns the rational function $\zeta(X, T)$ in "polynomial time"?

To quantify "polynomial time" (in the length of the input), we must specify an input mechanism for X. Note that if $Z \subseteq X$ is a closed subscheme, then

$$\zeta(X, T) = \zeta(Z, T)\zeta(X - Z, T);$$

we can thus reduce to working with affine hypersurfaces.

Problem (Well-posed, and open!)

Fix a positive integer n. Is there an algorithm that, given $f \in \mathbb{F}_q[x_1, \dots, x_n]$ of degree d, returns the rational function $\zeta(X, T)$ for $X = \operatorname{Spec} \mathbb{F}_q[x_1, \dots, x_n]/(f)$ in time $\operatorname{poly}(d, \log q)$?

Zeta functions: the computational problem (continued)

Problem (Ill-posed)

Fix a positive integer n. Is there an algorithm that, given an algebraic variety X of dimension n, returns the rational function $\zeta(X, T)$ in "polynomial time"?

To quantify "polynomial time" (in the length of the input), we must specify an input mechanism for X. Note that if $Z \subseteq X$ is a closed subscheme, then

$$\zeta(X, T) = \zeta(Z, T)\zeta(X - Z, T);$$

we can thus reduce to working with affine hypersurfaces.

Problem (Well-posed, and open!)

Fix a positive integer n. Is there an algorithm that, given $f \in \mathbb{F}_q[x_1, \dots, x_n]$ of degree d, returns the rational function $\zeta(X, T)$ for $X = \operatorname{Spec} \mathbb{F}_q[x_1, \dots, x_n]/(f)$ in time $\operatorname{poly}(d, \log q)$?

Zeta functions: the computational problem (continued)

In principle, one can compute $\zeta(X, T)$ by computing the action of Frobenius on mod-ℓ étale cohomology for a few small primes ℓ. This yields important but rather specific cases.

Theorem (Schoof, 1985; Pila, 1990)

Fix a positive integer d. There is an algorithm which, for $f \in \mathbb{F}_q[x_1, x_2]$ of degree d, returns the rational function $\zeta(X, T)$ for $X = \operatorname{Spec} \mathbb{F}_q[x_1, x_2]/(f)$ in time $\mathrm{poly}(\log q)$.

By contrast, if one agrees to concede p for $\log p$, then p-adic methods apply quite generally.

Theorem (Lauder-Wan, 2000–2008)

Fix a positive integer n. There is an algorithm which, for $f \in \mathbb{F}_q[x_1, \dots, x_n]$ of degree d, returns the rational function $\zeta(X, T)$ for $X = \operatorname{Spec} \mathbb{F}_q[x_1, \dots, x_n]/(f)$ in time $\mathrm{poly}(d, p, \log_p q)$.

Zeta functions: the computational problem (continued)

In principle, one can compute $\zeta(X, T)$ by computing the action of Frobenius on mod-ℓ étale cohomology for a few small primes ℓ. This yields important but rather specific cases.

Theorem (Schoof, 1985; Pila, 1990)

Fix a positive integer d. There is an algorithm which, for $f \in \mathbb{F}_q[x_1, x_2]$ of degree d, returns the rational function $\zeta(X, T)$ for $X = \operatorname{Spec} \mathbb{F}_q[x_1, x_2]/(f)$ in time $\operatorname{poly}(\log q)$.

By contrast, if one agrees to concede p for $\log p$, then p-adic methods apply quite generally.

Theorem (Lauder-Wan, 2000–2008)

Fix a positive integer n. There is an algorithm which, for $f \in \mathbb{F}_q[x_1, \dots, x_n]$ of degree d, returns the rational function $\zeta(X, T)$ for $X = \operatorname{Spec} \mathbb{F}_q[x_1, \dots, x_n]/(f)$ in time $\operatorname{poly}(d, p, \log_p q)$.

Zeta functions: the computational problem (continued)

In principle, one can compute $\zeta(X,T)$ by computing the action of Frobenius on mod-ℓ étale cohomology for a few small primes ℓ. This yields important but rather specific cases.

Theorem (Schoof, 1985; Pila, 1990)

Fix a positive integer d. There is an algorithm which, for $f \in \mathbb{F}_q[x_1, x_2]$ of degree d, returns the rational function $\zeta(X,T)$ for $X = \operatorname{Spec} \mathbb{F}_q[x_1,x_2]/(f)$ in time $\operatorname{poly}(\log q)$.

By contrast, if one agrees to concede p for $\log p$, then p-adic methods apply quite generally.

Theorem (Lauder-Wan, 2000–2008)

Fix a positive integer n. There is an algorithm which, for $f \in \mathbb{F}_q[x_1, \dots, x_n]$ of degree d, returns the rational function $\zeta(X,T)$ for $X = \operatorname{Spec} \mathbb{F}_q[x_1, \dots, x_n]/(f)$ in time $\operatorname{poly}(d, p, \log_p q)$.

Zeta functions: the computational problem (continued)

In principle, one can compute $\zeta(X, T)$ by computing the action of Frobenius on mod-ℓ étale cohomology for a few small primes ℓ. This yields important but rather specific cases.

Theorem (Schoof, 1985; Pila, 1990)

Fix a positive integer d. There is an algorithm which, for $f \in \mathbb{F}_q[x_1, x_2]$ of degree d, returns the rational function $\zeta(X, T)$ for $X = \operatorname{Spec} \mathbb{F}_q[x_1, x_2]/(f)$ in time $\operatorname{poly}(\log q)$.

By contrast, if one agrees to concede p for $\log p$, then p-adic methods apply quite generally.

Theorem (Lauder-Wan, 2000–2008)

Fix a positive integer n. There is an algorithm which, for $f \in \mathbb{F}_q[x_1, \dots, x_n]$ of degree d, returns the rational function $\zeta(X, T)$ for $X = \operatorname{Spec} \mathbb{F}_q[x_1, \dots, x_n]/(f)$ in time $\operatorname{poly}(d, p, \log_p q)$.

p-adic methods for zeta functions

Lauder and Wan prove their theorem by transcribing Dwork's proof of rationality. To date, the resulting algorithm has not been made practical.

However, Dwork's approach was later adapted into a *p-adic Weil cohomology theory* bearing a more formal resemblance to étale cohomology, and also yielding more effective computational methods. This includes work of myself, Lauder, Denef, Harrison, Vercauteren, Castryck, Gerkmann, Harvey, Tuitman, Xiao, Bradshaw, Balakrishnan, Hubrechts, Besser, de Jeu, Kloosterman, Escriva, Shieh, Costa, ...

We will restrict our discussion to the case where X is a curve. In this case, there are important links to Coleman's theory of p-adic abelian integrals (cf. talks of Gross, Besser), the Chabauty-Coleman method (cf. talks of Poonen, Stoll, Zureick-Brown) and nonabelian Chabauty (cf. talk of Kim).

p-adic methods for zeta functions

Lauder and Wan prove their theorem by transcribing Dwork's proof of rationality. To date, the resulting algorithm has not been made practical.

However, Dwork's approach was later adapted into a *p-adic Weil cohomology theory* bearing a more formal resemblance to étale cohomology, and also yielding more effective computational methods. This includes work of myself, Lauder, Denef, Harrison, Vercauteren, Castryck, Gerkmann, Harvey, Tuitman, Xiao, Bradshaw, Balakrishnan, Hubrechts, Besser, de Jeu, Kloosterman, Escriva, Shieh, Costa, ...

We will restrict our discussion to the case where X is a curve. In this case, there are important links to Coleman's theory of p-adic abelian integrals (cf. talks of Gross, Besser), the Chabauty-Coleman method (cf. talks of Poonen, Stoll, Zureick-Brown) and nonabelian Chabauty (cf. talk of Kim).

p-adic methods for zeta functions

Lauder and Wan prove their theorem by transcribing Dwork's proof of rationality. To date, the resulting algorithm has not been made practical.

However, Dwork's approach was later adapted into a *p-adic Weil cohomology theory* bearing a more formal resemblance to étale cohomology, and also yielding more effective computational methods. This includes work of myself, Lauder, Denef, Harrison, Vercauteren, Castryck, Gerkmann, Harvey, Tuitman, Xiao, Bradshaw, Balakrishnan, Hubrechts, Besser, de Jeu, Kloosterman, Escriva, Shieh, Costa, ...

We will restrict our discussion to the case where X is a curve. In this case, there are important links to Coleman's theory of p-adic abelian integrals (cf. talks of Gross, Besser), the Chabauty-Coleman method (cf. talks of Poonen, Stoll, Zureick-Brown) and nonabelian Chabauty (cf. talk of Kim).

p-adic cohomology for curves

Let X be a curve of genus g over $\mathbb{F}_q$. Lift X to a smooth proper curve $\tilde{X}$ over $\mathbb{Z}_q$ (the étale extension of $\mathbb{Z}_p$ with residue field $\mathbb{F}_q$). The p-adic cohomology of X "is" the algebraic de Rham cohomology of the generic fiber $\tilde{X}_{\mathbb{Q}}$. The important part is H^1, which sits in an exact sequence

$$0 \to H^0(\tilde{X}_{\mathbb{Q}}, \Omega) \to H^1_{\mathrm{dR}}(\tilde{X}_{\mathbb{Q}}, \mathbb{Q}_q) \to H^1(\tilde{X}_{\mathbb{Q}}, \mathcal{O}) \to 0.$$

As in the classical theory of Riemann surfaces, elements of H^1_{dR} can be represented by certain meromorphic differential forms.

There is a canonical endomorphism Frob_q of H^1_{dR} such that

$$\zeta(X, T) = \frac{\det(1 - T\,\mathrm{Frob}_q, H^1_{\mathrm{dR}})}{(1-T)(1-qT)}.$$

One way to produce Frob_q is using *crystalline cohomology* (Grothendieck, Berthelot, Ogus, etc.). However, from this interpretation, it is not straightforward to compute the matrix of action of Frob_q on a basis.

p-adic cohomology for curves

Let X be a curve of genus g over $\mathbb{F}_q$. Lift X to a smooth proper curve $\tilde{X}$ over $\mathbb{Z}_q$ (the étale extension of $\mathbb{Z}_p$ with residue field $\mathbb{F}_q$). The p-adic cohomology of X "is" the algebraic de Rham cohomology of the generic fiber $\tilde{X}_{\mathbb{Q}}$. The important part is H^1, which sits in an exact sequence

$$0 \to H^0(\tilde{X}_{\mathbb{Q}}, \Omega) \to H^1_{\mathrm{dR}}(\tilde{X}_{\mathbb{Q}}, \mathbb{Q}_q) \to H^1(\tilde{X}_{\mathbb{Q}}, \mathcal{O}) \to 0.$$

As in the classical theory of Riemann surfaces, elements of H^1_{dR} can be represented by certain meromorphic differential forms.

There is a canonical endomorphism Frob_q of H^1_{dR} such that

$$\zeta(X, T) = \frac{\det(1 - T\,\mathrm{Frob}_q, H^1_{\mathrm{dR}})}{(1-T)(1-qT)}.$$

One way to produce Frob_q is using *crystalline cohomology* (Grothendieck, Berthelot, Ogus, etc.). However, from this interpretation, it is not straightforward to compute the matrix of action of Frob_q on a basis.

p-adic cohomology for curves

Let X be a curve of genus g over $\mathbb{F}_q$. Lift X to a smooth proper curve $\tilde{X}$ over $\mathbb{Z}_q$ (the étale extension of $\mathbb{Z}_p$ with residue field $\mathbb{F}_q$). The p-adic cohomology of X "is" the algebraic de Rham cohomology of the generic fiber $\tilde{X}_{\mathbb{Q}}$. The important part is H^1, which sits in an exact sequence

$$0 \to H^0(\tilde{X}_{\mathbb{Q}}, \Omega) \to H^1_{\mathrm{dR}}(\tilde{X}_{\mathbb{Q}}, \mathbb{Q}_q) \to H^1(\tilde{X}_{\mathbb{Q}}, \mathcal{O}) \to 0.$$

As in the classical theory of Riemann surfaces, elements of H^1_{dR} can be represented by certain meromorphic differential forms.

There is a canonical endomorphism Frob_q of H^1_{dR} such that

$$\zeta(X, T) = \frac{\det(1 - T\,\mathrm{Frob}_q, H^1_{\mathrm{dR}})}{(1-T)(1-qT)}.$$

One way to produce Frob_q is using *crystalline cohomology* (Grothendieck, Berthelot, Ogus, etc.). However, from this interpretation, it is not straightforward to compute the matrix of action of Frob_q on a basis.

Structure of the rigid analytification

By rigid GAGA, $H^1_{\mathrm{dR}}(\tilde{X}_{\mathbb{Q}}, \mathbb{Q}_q)$ is also the de Rham cohomology of the rigid analytification Y of $\tilde{X}_{\mathbb{Q}}$. There is a *reduction* map $\mathrm{red} : Y \to X$ whose inverse images are open *residue discs*.

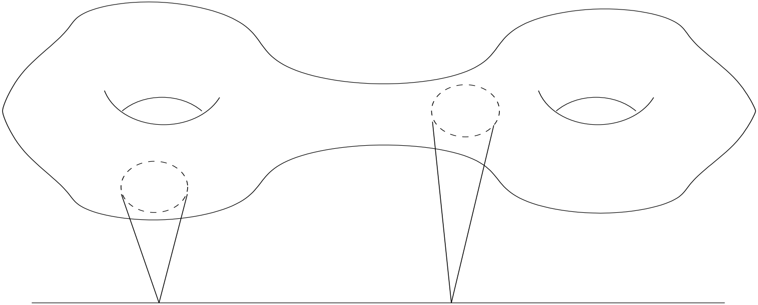

A *wide open* subset of Y is an open subset consisting of the complement of a nonempty finite union of closed discs, each contained in a residue disc.

Structure of the rigid analytification

By rigid GAGA, $H^1_{\mathrm{dR}}(\tilde{X}_{\mathbb{Q}}, \mathbb{Q}_q)$ is also the de Rham cohomology of the rigid analytification Y of $\tilde{X}_{\mathbb{Q}}$. There is a *reduction* map $\mathrm{red} : Y \to X$ whose inverse images are open *residue discs*.

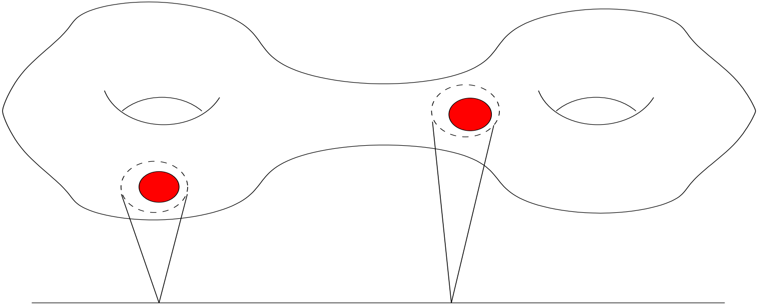

A *wide open* subset of Y is an open subset consisting of the complement of a nonempty finite union of closed discs, each contained in a residue disc.

The Monsky-Washnitzer Frobenius action

In general, there does not exist an endomorphism of $\tilde{X}_{\mathbb{Q}}$ (or Y) lifting the q-power Frobenius automorphism of X.

However, for any open affine subspace U of X, there exist wide open subsets V_1, V_2 of Y with $f(Y - V_i) = X - U$ and a morphism $\varphi : V_1 \to V_2$ lifting Frobenius.

Theorem (Monsky-Washnitzer, 1971)

Via the canonical isomorphism $H^1_{\mathrm{dR}}(V_1) \cong H^1_{\mathrm{dR}}(V_2)$, we have

$$\zeta(U, T) = \frac{\det(1 - q\varphi^{-1}T, H^i_{\mathrm{dR}}(V_1))}{1 - qT}.$$

Warning: from now on, we pretend $V_1 = V_2 = V$ for brevity.

The Monsky-Washnitzer Frobenius action

In general, there does not exist an endomorphism of $\tilde{X}_{\mathbb{Q}}$ (or Y) lifting the q-power Frobenius automorphism of X.

However, for any open affine subspace U of X, there exist wide open subsets V_1, V_2 of Y with $f(Y - V_i) = X - U$ and a morphism $\varphi : V_1 \to V_2$ lifting Frobenius.

Theorem (Monsky-Washnitzer, 1971)

Via the canonical isomorphism $H^1_{\mathrm{dR}}(V_1) \cong H^1_{\mathrm{dR}}(V_2)$, we have

$$\zeta(U, T) = \frac{\det(1 - q\varphi^{-1}T, H^i_{\mathrm{dR}}(V_1))}{1 - qT}.$$

Warning: from now on, we pretend $V_1 = V_2 = V$ for brevity.

The Monsky-Washnitzer Frobenius action

In general, there does not exist an endomorphism of $\tilde{X}_{\mathbb{Q}}$ (or Y) lifting the q-power Frobenius automorphism of X.

However, for any open affine subspace U of X, there exist wide open subsets V_1, V_2 of Y with $f(Y - V_i) = X - U$ and a morphism $\varphi : V_1 \to V_2$ lifting Frobenius.

Theorem (Monsky-Washnitzer, 1971)

Via the canonical isomorphism $H^1_{\mathrm{dR}}(V_1) \cong H^1_{\mathrm{dR}}(V_2)$*, we have*

$$\zeta(U, T) = \frac{\det(1 - q\varphi^{-1}T, H^i_{\mathrm{dR}}(V_1))}{1 - qT}.$$

Warning: from now on, we pretend $V_1 = V_2 = V$ for brevity.

The Monsky-Washnitzer Frobenius action

In general, there does not exist an endomorphism of $\tilde{X}_{\mathbb{Q}}$ (or Y) lifting the q-power Frobenius automorphism of X.

However, for any open affine subspace U of X, there exist wide open subsets V_1, V_2 of Y with $f(Y - V_i) = X - U$ and a morphism $\varphi : V_1 \to V_2$ lifting Frobenius.

Theorem (Monsky-Washnitzer, 1971)

Via the canonical isomorphism $H^1_{\mathrm{dR}}(V_1) \cong H^1_{\mathrm{dR}}(V_2)$, we have

$$\zeta(U, T) = \frac{\det(1 - q\varphi^{-1}T, H^i_{\mathrm{dR}}(V_1))}{1 - qT}.$$

Warning: from now on, we pretend $V_1 = V_2 = V$ for brevity.

A p-adic framework for computing ζ

- Choose the lift $\tilde{X}$ of X.
- Choose the open subset U, the wide open V, and the morphism φ. (If $q \neq p$, one can lift the p-power Frobenius and iterate to get φ.)
- Apply φ to 1-forms representing a basis of $H^1_{\mathrm{dR}}(V)$.
- Use known relations in $H^1_{\mathrm{dR}}(V)$ to write the results as exact 1-forms plus $\mathbb{Q}_q$-linear combinations of basis vectors.
- Recover $\zeta(U, T)$ from the characteristic polynomial of φ on $H^1_{\mathrm{dR}}(V)$.

Note that the characteristic polynomial has coefficients in $\mathbb{Z}$, but we compute it over $\mathbb{Q}_q$. This involves inexact (truncated) arithmetic: elements of $\mathbb{Q}_q$ are truncated to rational numbers, and functions on V_1 are truncated to meromorphic functions on Y. To provably recover $\zeta(U, T)$, one must be careful about these truncations (e.g., by identifying V).

A p-adic framework for computing ζ

- Choose the lift $\tilde{X}$ of X.
- Choose the open subset U, the wide open V, and the morphism φ. (If $q \neq p$, one can lift the p-power Frobenius and iterate to get φ.)
- Apply φ to 1-forms representing a basis of $H^1_{\mathrm{dR}}(V)$.
- Use known relations in $H^1_{\mathrm{dR}}(V)$ to write the results as exact 1-forms plus $\mathbb{Q}_q$-linear combinations of basis vectors.
- Recover $\zeta(U, T)$ from the characteristic polynomial of φ on $H^1_{\mathrm{dR}}(V)$.

Note that the characteristic polynomial has coefficients in $\mathbb{Z}$, but we compute it over $\mathbb{Q}_q$. This involves inexact (truncated) arithmetic: elements of $\mathbb{Q}_q$ are truncated to rational numbers, and functions on V_1 are truncated to meromorphic functions on Y. To provably recover $\zeta(U, T)$, one must be careful about these truncations (e.g., by identifying V).

A p-adic framework for computing ζ

- Choose the lift $\tilde{X}$ of X.
- Choose the open subset U, the wide open V, and the morphism φ. (If $q \neq p$, one can lift the p-power Frobenius and iterate to get φ.)
- Apply φ to 1-forms representing a basis of $H^1_{\mathrm{dR}}(V)$.
- Use known relations in $H^1_{\mathrm{dR}}(V)$ to write the results as exact 1-forms plus $\mathbb{Q}_q$-linear combinations of basis vectors.
- Recover $\zeta(U, T)$ from the characteristic polynomial of φ on $H^1_{\mathrm{dR}}(V)$.

Note that the characteristic polynomial has coefficients in $\mathbb{Z}$, but we compute it over $\mathbb{Q}_q$. This involves inexact (truncated) arithmetic: elements of $\mathbb{Q}_q$ are truncated to rational numbers, and functions on V_1 are truncated to meromorphic functions on Y. To provably recover $\zeta(U, T)$, one must be careful about these truncations (e.g., by identifying V).

A p-adic framework for computing ζ

- Choose the lift $\tilde{X}$ of X.
- Choose the open subset U, the wide open V, and the morphism φ. (If $q \neq p$, one can lift the p-power Frobenius and iterate to get φ.)
- Apply φ to 1-forms representing a basis of $H^1_{\mathrm{dR}}(V)$.
- Use known relations in $H^1_{\mathrm{dR}}(V)$ to write the results as exact 1-forms plus $\mathbb{Q}_q$-linear combinations of basis vectors.
- Recover $\zeta(U, T)$ from the characteristic polynomial of φ on $H^1_{\mathrm{dR}}(V)$.

Note that the characteristic polynomial has coefficients in $\mathbb{Z}$, but we compute it over $\mathbb{Q}_q$. This involves inexact (truncated) arithmetic: elements of $\mathbb{Q}_q$ are truncated to rational numbers, and functions on V_1 are truncated to meromorphic functions on Y. To provably recover $\zeta(U, T)$, one must be careful about these truncations (e.g., by identifying V).

A p-adic framework for computing ζ

- Choose the lift $\tilde{X}$ of X.
- Choose the open subset U, the wide open V, and the morphism φ. (If $q \neq p$, one can lift the p-power Frobenius and iterate to get φ.)
- Apply φ to 1-forms representing a basis of $H^1_{\mathrm{dR}}(V)$.
- Use known relations in $H^1_{\mathrm{dR}}(V)$ to write the results as exact 1-forms plus $\mathbb{Q}_q$-linear combinations of basis vectors.
- Recover $\zeta(U, T)$ from the characteristic polynomial of φ on $H^1_{\mathrm{dR}}(V)$.

Note that the characteristic polynomial has coefficients in $\mathbb{Z}$, but we compute it over $\mathbb{Q}_q$. This involves inexact (truncated) arithmetic: elements of $\mathbb{Q}_q$ are truncated to rational numbers, and functions on V_1 are truncated to meromorphic functions on Y. To provably recover $\zeta(U, T)$, one must be careful about these truncations (e.g., by identifying V).

A p-adic framework for computing ζ

- Choose the lift $\tilde{X}$ of X.
- Choose the open subset U, the wide open V, and the morphism φ. (If $q \neq p$, one can lift the p-power Frobenius and iterate to get φ.)
- Apply φ to 1-forms representing a basis of $H^1_{\mathrm{dR}}(V)$.
- Use known relations in $H^1_{\mathrm{dR}}(V)$ to write the results as exact 1-forms plus $\mathbb{Q}_q$-linear combinations of basis vectors.
- Recover $\zeta(U, T)$ from the characteristic polynomial of φ on $H^1_{\mathrm{dR}}(V)$.

Note that the characteristic polynomial has coefficients in $\mathbb{Z}$, but we compute it over $\mathbb{Q}_q$. This involves inexact (truncated) arithmetic: elements of $\mathbb{Q}_q$ are truncated to rational numbers, and functions on V_1 are truncated to meromorphic functions on Y. To provably recover $\zeta(U, T)$, one must be careful about these truncations (e.g., by identifying V).

A p-adic framework for computing ζ

- Choose the lift $\tilde{X}$ of X.
- Choose the open subset U, the wide open V, and the morphism φ. (If $q \neq p$, one can lift the p-power Frobenius and iterate to get φ.)
- Apply φ to 1-forms representing a basis of $H^1_{\mathrm{dR}}(V)$.
- Use known relations in $H^1_{\mathrm{dR}}(V)$ to write the results as exact 1-forms plus $\mathbb{Q}_q$-linear combinations of basis vectors.
- Recover $\zeta(U, T)$ from the characteristic polynomial of φ on $H^1_{\mathrm{dR}}(V)$.

Note that the characteristic polynomial has coefficients in $\mathbb{Z}$, but we compute it over $\mathbb{Q}_q$. This involves inexact (truncated) arithmetic: elements of $\mathbb{Q}_q$ are truncated to rational numbers, and functions on V_1 are truncated to meromorphic functions on Y. To provably recover $\zeta(U, T)$, one must be careful about these truncations (e.g., by identifying V).

Example: hyperelliptic curves (K, 2001)

Suppose $p \neq 2$. Let X be a hyperelliptic curve of the form $y^2 = P(x)$ where $\deg P(x) = 2g+1$, lifted to $y^2 = \tilde{P}(x)$ where $\deg \tilde{P}(x) = 2g+1$. Let $U \subseteq X$ be the subset where y is invertible. Then $H^1_{\mathrm{dR}}(V)$ admits the basis

$$\frac{x^i\,dx}{y}, \frac{x^j\,dx}{y^2} \qquad (i = 0, \dots, 2g-1; j = 0, \dots, 2g).$$

These span the eigenspaces $H^1_{\mathrm{dR}}(V)^{\mp}$ for the involution $y \mapsto -y$; the $-$ eigenspace coincides with $H^1_{\mathrm{dR}}(Y)$.

We may take φ sending x to x^q and y to $y^q(\tilde{P}(x^q)/\tilde{P}(x)^q)^{-1/2}$ (computed using a binomial expansion or better, Newton-Raphson iteration).

We may perform simplifications in $H^1_{\mathrm{dR}}(V)$ systematically: e.g., there are explicit relations converting $Q(x)\,dx/y^{2n+1}$ into $R(x)\,dx/y^{2n-1}$.

Example: hyperelliptic curves (K, 2001)

Suppose $p \neq 2$. Let X be a hyperelliptic curve of the form $y^2 = P(x)$ where $\deg P(x) = 2g+1$, lifted to $y^2 = \tilde{P}(x)$ where $\deg \tilde{P}(x) = 2g+1$. Let $U \subseteq X$ be the subset where y is invertible. Then $H^1_{\mathrm{dR}}(V)$ admits the basis

$$\frac{x^i\,dx}{y}, \frac{x^j\,dx}{y^2} \qquad (i = 0, \dots, 2g-1; j = 0, \dots, 2g).$$

These span the eigenspaces $H^1_{\mathrm{dR}}(V)^{\mp}$ for the involution $y \mapsto -y$; the $-$ eigenspace coincides with $H^1_{\mathrm{dR}}(Y)$.

We may take φ sending x to x^q and y to $y^q(\tilde{P}(x^q)/\tilde{P}(x)^q)^{-1/2}$ (computed using a binomial expansion or better, Newton-Raphson iteration).

We may perform simplifications in $H^1_{\mathrm{dR}}(V)$ systematically: e.g., there are explicit relations converting $Q(x)\,dx/y^{2n+1}$ into $R(x)\,dx/y^{2n-1}$.

Example: hyperelliptic curves (K, 2001)

Suppose $p \neq 2$. Let X be a hyperelliptic curve of the form $y^2 = P(x)$ where $\deg P(x) = 2g+1$, lifted to $y^2 = \tilde{P}(x)$ where $\deg \tilde{P}(x) = 2g+1$. Let $U \subseteq X$ be the subset where y is invertible. Then $H^1_{\mathrm{dR}}(V)$ admits the basis

$$\frac{x^i\,dx}{y}, \frac{x^j\,dx}{y^2} \qquad (i = 0, \dots, 2g-1; j = 0, \dots, 2g).$$

These span the eigenspaces $H^1_{\mathrm{dR}}(V)^{\mp}$ for the involution $y \mapsto -y$; the $-$ eigenspace coincides with $H^1_{\mathrm{dR}}(Y)$.

We may take φ sending x to x^q and y to $y^q(\tilde{P}(x^q)/\tilde{P}(x)^q)^{-1/2}$ (computed using a binomial expansion or better, Newton-Raphson iteration).

We may perform simplifications in $H^1_{\mathrm{dR}}(V)$ systematically: e.g., there are explicit relations converting $Q(x)\,dx/y^{2n+1}$ into $R(x)\,dx/y^{2n-1}$.

Coleman integration

As described in Gross's lecture, Coleman defined path integrals $\int_P^Q \omega$ for any meromorphic differential ω on a wide open subset V of Y and any points $P, Q \in V$ which are not poles of ω.

If P, Q lie in a single residue disc, this is easy: ω admits an analytic antiderivative F on the disc, so $\int_P^Q \omega = F(Q) - F(P)$.

But F is only locally analytic on V, and one is free to choose a different constant of integration on each residue disc. How to make coherent choices?

This was answered by Coleman in 1981. At Banff in 2007, Robert and I came up with an equivalent answer...

Coleman integration

As described in Gross's lecture, Coleman defined path integrals $\int_P^Q \omega$ for any meromorphic differential ω on a wide open subset V of Y and any points $P, Q \in V$ which are not poles of ω.

If P, Q lie in a single residue disc, this is easy: ω admits an analytic antiderivative F on the disc, so $\int_P^Q \omega = F(Q) - F(P)$.

But F is only locally analytic on V, and one is free to choose a different constant of integration on each residue disc. How to make coherent choices?

This was answered by Coleman in 1981. At Banff in 2007, Robert and I came up with an equivalent answer...

Coleman integration

As described in Gross's lecture, Coleman defined path integrals $\int_P^Q \omega$ for any meromorphic differential ω on a wide open subset V of Y and any points $P, Q \in V$ which are not poles of ω.

If P, Q lie in a single residue disc, this is easy: ω admits an analytic antiderivative F on the disc, so $\int_P^Q \omega = F(Q) - F(P)$.

But F is only locally analytic on V, and one is free to choose a different constant of integration on each residue disc. How to make coherent choices?

This was answered by Coleman in 1981. At Banff in 2007, Robert and I came up with an equivalent answer...

Coleman integration

As described in Gross's lecture, Coleman defined path integrals $\int_P^Q \omega$ for any meromorphic differential ω on a wide open subset V of Y and any points $P, Q \in V$ which are not poles of ω.

If P, Q lie in a single residue disc, this is easy: ω admits an analytic antiderivative F on the disc, so $\int_P^Q \omega = F(Q) - F(P)$.

But F is only locally analytic on V, and one is free to choose a different constant of integration on each residue disc. How to make coherent choices?

This was answered by Coleman in 1981. At Banff in 2007, Robert and I came up with an equivalent answer...

Banff, February 2007 (missing Robert)

Coleman integration and the M-W construction

One is supposed to have the Jacobian change of variables formula, particularly for Frobenius:

$$\int_P^Q \varphi^*(\omega) = \int_{\varphi(P)}^{\varphi(Q)} \omega.$$

Coleman used this by applying $R(\varphi)$ where R is the characteristic polynomial of φ on $H^1_{\mathrm{dR}}(V)$. By Cayley-Hamilton, $R(\varphi^*)(\omega)$ is exact.

Equivalently, let $\omega_1, \dots, \omega_n$ be forms representing a basis of $H^1_{\mathrm{dR}}(V)$. The Monsky-Washnitzer computation gave us formulas

$$\varphi^*(\omega_j) = \sum_i A_{ij}\omega_i + df_j,$$

in which we discarded f_j and evaluated the charpoly of A.

Coleman integration and the M-W construction

One is supposed to have the Jacobian change of variables formula, particularly for Frobenius:

$$\int_P^Q \varphi^*(\omega) = \int_{\varphi(P)}^{\varphi(Q)} \omega.$$

Coleman used this by applying $R(\varphi)$ where R is the characteristic polynomial of φ on $H^1_{\mathrm{dR}}(V)$. By Cayley-Hamilton, $R(\varphi^*)(\omega)$ is exact.

Equivalently, let $\omega_1, \dots, \omega_n$ be forms representing a basis of $H^1_{\mathrm{dR}}(V)$. The Monsky-Washnitzer computation gave us formulas

$$\varphi^*(\omega_j) = \sum_i A_{ij}\omega_i + df_j,$$

in which we discarded f_j and evaluated the charpoly of A.

Coleman integration and the M-W construction

One is supposed to have the Jacobian change of variables formula, particularly for Frobenius:

$$\int_P^Q \varphi^*(\omega) = \int_{\varphi(P)}^{\varphi(Q)} \omega.$$

Coleman used this by applying $R(\varphi)$ where R is the characteristic polynomial of φ on $H^1_{\mathrm{dR}}(V)$. By Cayley-Hamilton, $R(\varphi^*)(\omega)$ is exact.

Equivalently, let $\omega_1, \dots, \omega_n$ be forms representing a basis of $H^1_{\mathrm{dR}}(V)$. The Monsky-Washnitzer computation gave us formulas

$$\varphi^*(\omega_j) = \sum_i A_{ij}\omega_i + df_j,$$

in which we discarded f_j and evaluated the charpoly of A.

Coleman integration and the M-W construction (cont.)

Instead of discarding f_i, let's write

$$\int_{\varphi(P)}^{\varphi(Q)} \omega_j = \int_P^Q \varphi^*(\omega_j) = \sum_i A_{ij} \int_P^Q \omega_i + f_j(Q) - f_j(P).$$

By writing $\int_{\varphi(P)}^{\varphi(Q)} = \int_{\varphi(P)}^{P} + \int_P^Q + \int_Q^{\varphi(Q)}$, we end up with the equation

$$(A-1)^{-1}(\text{vector of } \int_P^Q \omega_i) = (\text{vector of computable quantities}).$$

Since A has no eigenvalues equal to 1 (by the Weil conjectures), this pins down all the constants of integration!

Variation: use "tiny" integrals (within a residue disc) to replace P and Q with φ-fixed points (*Teichmüller points*) P' and Q', and then get a linear equation directly on $\int_{P'}^{Q'} \omega_j$.

Coleman integration and the M-W construction (cont.)

Instead of discarding f_i, let's write

$$\int_{\varphi(P)}^{\varphi(Q)} \omega_j = \int_P^Q \varphi^*(\omega_j) = \sum_i A_{ij} \int_P^Q \omega_i + f_j(Q) - f_j(P).$$

By writing $\int_{\varphi(P)}^{\varphi(Q)} = \int_{\varphi(P)}^{P} + \int_P^Q + \int_Q^{\varphi(Q)}$, we end up with the equation

$$(A-1)^{-1}(\text{vector of } \textstyle\int_P^Q \omega_i) = (\text{vector of computable quantities}).$$

Since A has no eigenvalues equal to 1 (by the Weil conjectures), this pins down all the constants of integration!

Variation: use “tiny” integrals (within a residue disc) to replace P and Q with φ-fixed points (*Teichmüller points*) P' and Q', and then get a linear equation directly on $\int_{P'}^{Q'} \omega_j$.

Coleman integration and the M-W construction (cont.)

Instead of discarding f_i, let's write

$$\int_{\varphi(P)}^{\varphi(Q)} \omega_j = \int_P^Q \varphi^*(\omega_j) = \sum_i A_{ij} \int_P^Q \omega_i + f_j(Q) - f_j(P).$$

By writing $\int_{\varphi(P)}^{\varphi(Q)} = \int_{\varphi(P)}^{P} + \int_P^Q + \int_Q^{\varphi(Q)}$, we end up with the equation

$$(A - 1)^{-1}(\text{vector of } \textstyle\int_P^Q \omega_i) = (\text{vector of computable quantities}).$$

Since A has no eigenvalues equal to 1 (by the Weil conjectures), this pins down all the constants of integration!

Variation: use "tiny" integrals (within a residue disc) to replace P and Q with φ-fixed points (*Teichmüller points*) P' and Q', and then get a linear equation directly on $\int_{P'}^{Q'} \omega_j$.

Illustration

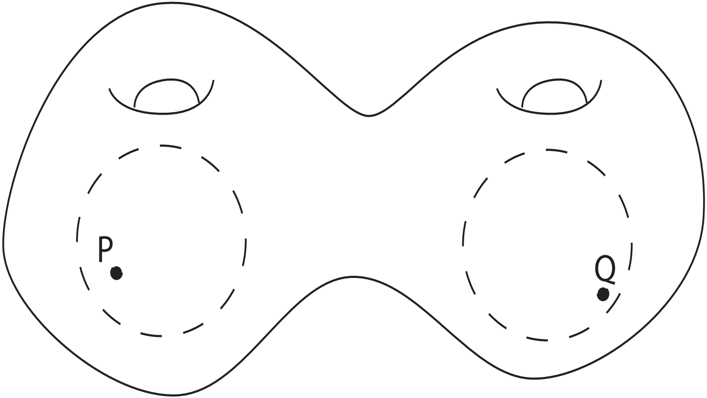

Illustration

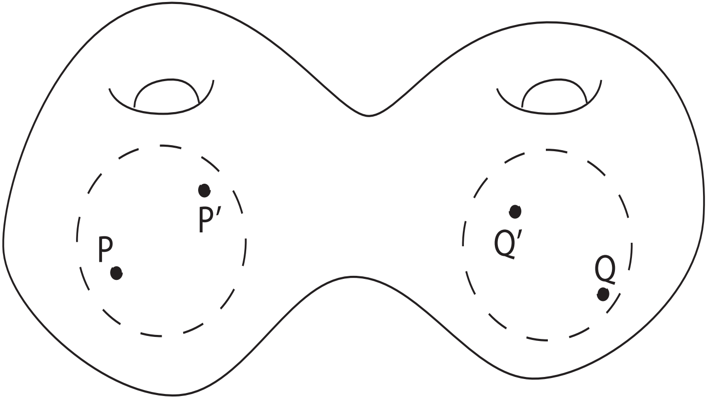

Illustration

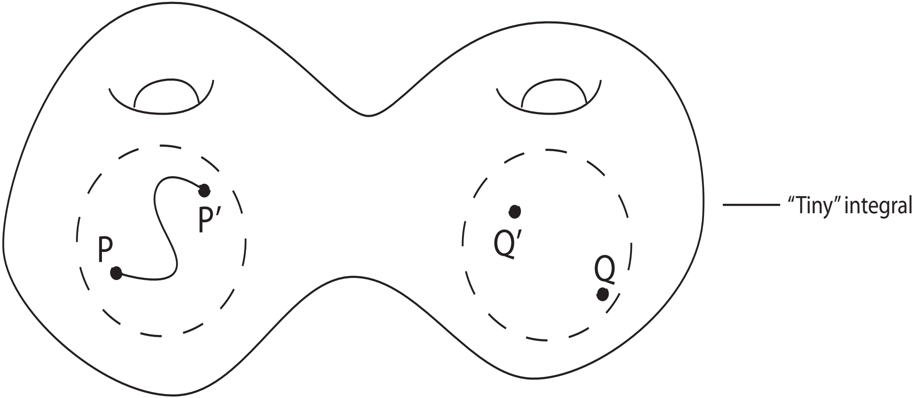

Illustration

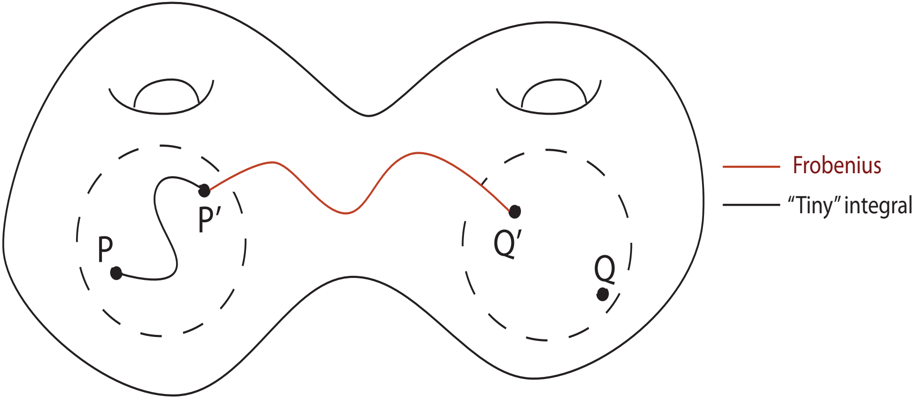

Illustration

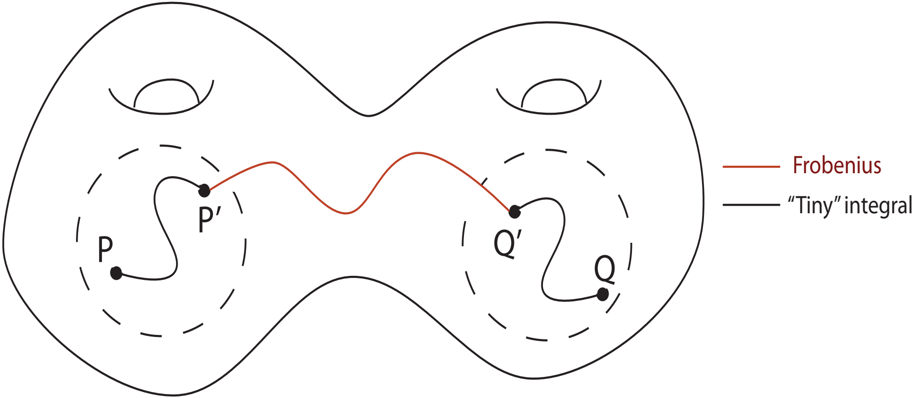

Numerical Coleman integration

For hyperelliptic curves, this strategy was implemented by Jennifer Balakrishnan in her PhD thesis (based on work of Robert Bradshaw at the 2007 Arizona Winter School).

Balakrishnan and Tuitman are currently extending this to more general curves. This should make it routine to compute the integrals arising in the Chabauty-Coleman method...

...and the *iterated Coleman integrals* arising in Kim's nonabelian Chabauty method. Applications to rational points on curves are in the works! (See Kim's talk for examples.)

Numerical Coleman integration

For hyperelliptic curves, this strategy was implemented by Jennifer Balakrishnan in her PhD thesis (based on work of Robert Bradshaw at the 2007 Arizona Winter School).

Balakrishnan and Tuitman are currently extending this to more general curves. This should make it routine to compute the integrals arising in the Chabauty-Coleman method...

...and the *iterated Coleman integrals* arising in Kim's nonabelian Chabauty method. Applications to rational points on curves are in the works! (See Kim's talk for examples.)

Numerical Coleman integration

For hyperelliptic curves, this strategy was implemented by Jennifer Balakrishnan in her PhD thesis (based on work of Robert Bradshaw at the 2007 Arizona Winter School).

Balakrishnan and Tuitman are currently extending this to more general curves. This should make it routine to compute the integrals arising in the Chabauty-Coleman method...

...and the *iterated Coleman integrals* arising in Kim's nonabelian Chabauty method. Applications to rational points on curves are in the works! (See Kim's talk for examples.)

Thank you!

To Robert, for demonstrating that mathematics is not just for those who color within the lines...

... and to all of you for helping us to honor his legacy!

Thank you!

To Robert, for demonstrating that mathematics is not just for those who color within the lines...

... and to all of you for helping us to honor his legacy!