Hypergeometric *L*-functions in average polynomial time

Edgar Costa, Kiran S. Kedlaya, and David Roe

Costa, Roe: Department of Mathematics, Massachusetts Institute of Technology
Kedlaya: Department of Mathematics, University of California, San Diego
edgarc@mit.edu, kedlaya@ucsd.edu, roed@mit.edu
paper: arXiv:2005.13640; slides: http://kskedlaya.org/slides/

(virtual) Algorithmic Number Theory Symposium (ANTS-XIV)
University of Auckland (Te Whare Wānanga o Tāmaki Makaurau)
July 2, 2020

Kedlaya was supported by NSF (grant DMS-1802161) and UC San Diego (Warschawski Professorship). Costa and Roe were supported by the Simons Collaboration on Arithmetic Geometry, Number Theory, and Computation.

The MIT campus sits on the traditional unceded territory of the Wampanoag Nation; we acknowledge the painful history of genocide and forced removal from this territory. The UCSD campus sits on the ancestral homelands of the Kumeyaay Nation; the Kumeyaay people continue to have an important and thriving presence in the region.

Arithmetic L-functions: examples

Given a smooth proper scheme X over a number field K, one can define **(incomplete) arithmetic L-functions**. These are Dirichlet series defined by products indexed by finite places of K at which X has good reduction.

Example

Take $X = \operatorname{Spec} K$. Then one gets the Dedekind zeta function

$$\zeta_K(s) = \prod_{\mathfrak{p}} (1 - \operatorname{Norm}(\mathfrak{p})^{-s})^{-1}.$$

Example

Let X be an elliptic curve over K. Then one of the L-functions is

$$L(X, s) = \prod_{\mathfrak{p}} (1 - a_{\mathfrak{p}} \operatorname{Norm}(\mathfrak{p})^{-s} + \operatorname{Norm}(\mathfrak{p})^{1-2s})^{-1}$$

where $a_{\mathfrak{p}}$ is the trace of Frobenius of $X_{\mathfrak{p}}$ (the mod-$\mathfrak{p}$ reduction of X).

Arithmetic L-functions: examples

Given a smooth proper scheme X over a number field K, one can define **(incomplete) arithmetic L-functions**. These are Dirichlet series defined by products indexed by finite places of K at which X has good reduction.

Example

Take $X = \operatorname{Spec} K$. Then one gets the Dedekind zeta function

$$\zeta_K(s) = \prod_{\mathfrak{p}} (1 - \operatorname{Norm}(\mathfrak{p})^{-s})^{-1}.$$

Example

Let X be an elliptic curve over K. Then one of the L-functions is

$$L(X, s) = \prod_{\mathfrak{p}} (1 - a_{\mathfrak{p}} \operatorname{Norm}(\mathfrak{p})^{-s} + \operatorname{Norm}(\mathfrak{p})^{1-2s})^{-1}$$

where $a_{\mathfrak{p}}$ is the trace of Frobenius of $X_{\mathfrak{p}}$ (the mod-$\mathfrak{p}$ reduction of X).

Arithmetic L-functions: examples

Given a smooth proper scheme X over a number field K, one can define **(incomplete) arithmetic L-functions**. These are Dirichlet series defined by products indexed by finite places of K at which X has good reduction.

Example

Take $X = \operatorname{Spec} K$. Then one gets the Dedekind zeta function

$$\zeta_K(s) = \prod_{\mathfrak{p}} (1 - \operatorname{Norm}(\mathfrak{p})^{-s})^{-1}.$$

Example

Let X be an elliptic curve over K. Then one of the L-functions is

$$L(X, s) = \prod_{\mathfrak{p}} (1 - a_{\mathfrak{p}} \operatorname{Norm}(\mathfrak{p})^{-s} + \operatorname{Norm}(\mathfrak{p})^{1-2s})^{-1}$$

where $a_{\mathfrak{p}}$ is the trace of Frobenius of $X_{\mathfrak{p}}$ (the mod-$\mathfrak{p}$ reduction of X).

Arithmetic L-functions: a general definition

In general, for $i \in \{0, \ldots, 2\dim(X)\}$, one gets an L-function whose factor at $\mathfrak{p}$ is $L_i(\mathrm{Norm}(\mathfrak{p})^{-s})^{-1}$, where L_i appears in the zeta function of $X_\mathfrak{p}$:

$$Z(X_\mathfrak{p}, T) = \frac{L_1(T) \cdots L_{2\dim(X)-1}(T)}{L_0(T) \cdots L_{2\dim(X)}(T)}.$$

On the previous slide, for $X = \mathrm{Spec}\, K$ we took $i = 0$; for X an elliptic curve we took $i = 1$.

These are expected to have analytic continuation/functional equation after completing the product so that it has one factor for each **finite or infinite** place of K. (Factors at infinite places involve the Gamma function.)

There is a rich theory of **special values** of arithmetic L-functions, including the Dirichlet class number formula, the Birch–Swinnerton-Dyer conjecture, and conjectures of Bloch–Kato, Deligne, Beilinson, etc.

Arithmetic L-functions: a general definition

In general, for $i \in \{0, \ldots, 2\dim(X)\}$, one gets an L-function whose factor at $\mathfrak{p}$ is $L_i(\mathrm{Norm}(\mathfrak{p})^{-s})^{-1}$, where L_i appears in the zeta function of $X_\mathfrak{p}$:

$$Z(X_\mathfrak{p}, T) = \frac{L_1(T) \cdots L_{2\dim(X)-1}(T)}{L_0(T) \cdots L_{2\dim(X)}(T)}.$$

On the previous slide, for $X = \operatorname{Spec} K$ we took $i = 0$; for X an elliptic curve we took $i = 1$.

These are expected to have analytic continuation/functional equation after completing the product so that it has one factor for each **finite or infinite** place of K. (Factors at infinite places involve the Gamma function.)

There is a rich theory of **special values** of arithmetic L-functions, including the Dirichlet class number formula, the Birch–Swinnerton-Dyer conjecture, and conjectures of Bloch–Kato, Deligne, Beilinson, etc.

Arithmetic L-functions: a general definition

In general, for $i \in \{0, \ldots, 2\dim(X)\}$, one gets an L-function whose factor at $\mathfrak{p}$ is $L_i(\mathrm{Norm}(\mathfrak{p})^{-s})^{-1}$, where L_i appears in the zeta function of $X_\mathfrak{p}$:

$$Z(X_\mathfrak{p}, T) = \frac{L_1(T) \cdots L_{2\dim(X)-1}(T)}{L_0(T) \cdots L_{2\dim(X)}(T)}.$$

On the previous slide, for $X = \operatorname{Spec} K$ we took $i = 0$; for X an elliptic curve we took $i = 1$.

These are expected to have analytic continuation/functional equation after completing the product so that it has one factor for each **finite or infinite** place of K. (Factors at infinite places involve the Gamma function.)

There is a rich theory of **special values** of arithmetic L-functions, including the Dirichlet class number formula, the Birch–Swinnerton-Dyer conjecture, and conjectures of Bloch–Kato, Deligne, Beilinson, etc.

Arithmetic L-functions: a general definition

In general, for $i \in \{0, \ldots, 2\dim(X)\}$, one gets an L-function whose factor at $\mathfrak{p}$ is $L_i(\mathrm{Norm}(\mathfrak{p})^{-s})^{-1}$, where L_i appears in the zeta function of $X_{\mathfrak{p}}$:

$$Z(X_{\mathfrak{p}}, T) = \frac{L_1(T) \cdots L_{2\dim(X)-1}(T)}{L_0(T) \cdots L_{2\dim(X)}(T)}.$$

On the previous slide, for $X = \operatorname{Spec} K$ we took $i = 0$; for X an elliptic curve we took $i = 1$.

These are expected to have analytic continuation/functional equation after completing the product so that it has one factor for each **finite or infinite** place of K. (Factors at infinite places involve the Gamma function.)

There is a rich theory of **special values** of arithmetic L-functions, including the Dirichlet class number formula, the Birch–Swinnerton-Dyer conjecture, and conjectures of Bloch–Kato, Deligne, Beilinson, etc.

Arithmetic L-functions in the LMFDB

In general, a single L-function can arise in various ways. E.g., isogenous elliptic curves, or abelian varieties, have the same L-function (and conversely by Tate–Faltings).

There are other ways to construct arithmetic L-functions for which there is not a distinguished "geometric origin". For example, any weight-2 rational eigenform for $\Gamma_0(N)$ has an L-function matching some elliptic curve over $\mathbb{Q}$ (Eichler–Shimura), but the latter is only determined up to isogeny.

A primary goal of the L-Functions and Modular Forms Database is to tabulate arithmetic L-functions with diverse discrete parameters (degree, weight, Hodge numbers). This paper is part of a project to add **hypergeometric L-functions**, which provide examples with assorted parameters; see the LMFDB beta site.

Arithmetic L-functions in the LMFDB

In general, a single L-function can arise in various ways. E.g., isogenous elliptic curves, or abelian varieties, have the same L-function (and conversely by Tate–Faltings).

There are other ways to construct arithmetic L-functions for which there is not a distinguished "geometric origin". For example, any weight-2 rational eigenform for $\Gamma_0(N)$ has an L-function matching some elliptic curve over $\mathbb{Q}$ (Eichler–Shimura), but the latter is only determined up to isogeny.

A primary goal of the L-Functions and Modular Forms Database is to tabulate arithmetic L-functions with diverse discrete parameters (degree, weight, Hodge numbers). This paper is part of a project to add **hypergeometric L-functions**, which provide examples with assorted parameters; see the LMFDB beta site.

Arithmetic L-functions in the LMFDB

In general, a single L-function can arise in various ways. E.g., isogenous elliptic curves, or abelian varieties, have the same L-function (and conversely by Tate–Faltings).

There are other ways to construct arithmetic L-functions for which there is not a distinguished "geometric origin". For example, any weight-2 rational eigenform for $\Gamma_0(N)$ has an L-function matching some elliptic curve over $\mathbb{Q}$ (Eichler–Shimura), but the latter is only determined up to isogeny.

A primary goal of the L-Functions and Modular Forms Database is to tabulate arithmetic L-functions with diverse discrete parameters (degree, weight, Hodge numbers). This paper is part of a project to add **hypergeometric L-functions**, which provide examples with assorted parameters; see the LMFDB beta site.

Hypergeometric data

A **hypergeometric datum** of degree r consists of two disjoint tuples $(\alpha_1, \ldots, \alpha_r), (\beta_1, \ldots, \beta_r)$ over $\mathbb{Q} \cap [0, 1)$ which are each **balanced**: the multiplicity of any reduced fraction depends only on its denominator. Later we will consider the example

$$\alpha = (\tfrac{1}{4}, \tfrac{1}{2}, \tfrac{1}{2}, \tfrac{3}{4}), \ \beta = (\tfrac{1}{3}, \tfrac{1}{3}, \tfrac{2}{3}, \tfrac{2}{3}).$$

To each such datum, we can define a family of arithmetic L-functions of degree r over $\mathbb{Q}$ parametrized by $z \in \mathbb{Q} \setminus \{0, 1\}$. The primes p of bad reduction have the following forms.

- p is **wild** if $\gamma \notin \mathbb{Z}_p$ for some $\gamma \in \alpha \cup \beta$ (e.g., 2 and 3 in our example).
- p is **tame** if it is not wild, and either $z \notin \mathbb{Z}_p^\times$ or $z - 1 \notin \mathbb{Z}_p^\times$.

This L-function is associated to a specific scheme defined in terms of $(\alpha, \beta), z$. However, there is no **distinguished** choice of this scheme.

Hypergeometric data

A **hypergeometric datum** of degree r consists of two disjoint tuples $(\alpha_1, \ldots, \alpha_r), (\beta_1, \ldots, \beta_r)$ over $\mathbb{Q} \cap [0, 1)$ which are each **balanced**: the multiplicity of any reduced fraction depends only on its denominator. Later we will consider the example

$$\alpha = (\tfrac{1}{4}, \tfrac{1}{2}, \tfrac{1}{2}, \tfrac{3}{4}),\ \beta = (\tfrac{1}{3}, \tfrac{1}{3}, \tfrac{2}{3}, \tfrac{2}{3}).$$

To each such datum, we can define a family of arithmetic L-functions of degree r over $\mathbb{Q}$ parametrized by $z \in \mathbb{Q} \setminus \{0, 1\}$. The primes p of bad reduction have the following forms.

- p is **wild** if $\gamma \notin \mathbb{Z}_p$ for some $\gamma \in \alpha \cup \beta$ (e.g., 2 and 3 in our example).
- p is **tame** if it is not wild, and either $z \notin \mathbb{Z}_p^\times$ or $z - 1 \notin \mathbb{Z}_p^\times$.

This L-function is associated to a specific scheme defined in terms of $(\alpha, \beta), z$. However, there is no **distinguished** choice of this scheme.

Hypergeometric data

A **hypergeometric datum** of degree r consists of two disjoint tuples $(\alpha_1, \dots, \alpha_r), (\beta_1, \dots, \beta_r)$ over $\mathbb{Q} \cap [0, 1)$ which are each **balanced**: the multiplicity of any reduced fraction depends only on its denominator. Later we will consider the example

$$\alpha = (\tfrac{1}{4}, \tfrac{1}{2}, \tfrac{1}{2}, \tfrac{3}{4}),\ \beta = (\tfrac{1}{3}, \tfrac{1}{3}, \tfrac{2}{3}, \tfrac{2}{3}).$$

To each such datum, we can define a family of arithmetic L-functions of degree r over $\mathbb{Q}$ parametrized by $z \in \mathbb{Q} \setminus \{0, 1\}$. The primes p of bad reduction have the following forms.

- p is **wild** if $\gamma \notin \mathbb{Z}_p$ for some $\gamma \in \alpha \cup \beta$ (e.g., 2 and 3 in our example).
- p is **tame** if it is not wild, and either $z \notin \mathbb{Z}_p^\times$ or $z - 1 \notin \mathbb{Z}_p^\times$.

This L-function is associated to a specific scheme defined in terms of $(\alpha, \beta), z$. However, there is no **distinguished** choice of this scheme.

Hypergeometric data

A **hypergeometric datum** of degree r consists of two disjoint tuples $(\alpha_1, \ldots, \alpha_r), (\beta_1, \ldots, \beta_r)$ over $\mathbb{Q} \cap [0, 1)$ which are each **balanced**: the multiplicity of any reduced fraction depends only on its denominator. Later we will consider the example

$$\alpha = (\tfrac{1}{4}, \tfrac{1}{2}, \tfrac{1}{2}, \tfrac{3}{4}),\ \beta = (\tfrac{1}{3}, \tfrac{1}{3}, \tfrac{2}{3}, \tfrac{2}{3}).$$

To each such datum, we can define a family of arithmetic L-functions of degree r over $\mathbb{Q}$ parametrized by $z \in \mathbb{Q} \setminus \{0, 1\}$. The primes p of bad reduction have the following forms.

- p is **wild** if $\gamma \notin \mathbb{Z}_p$ for some $\gamma \in \alpha \cup \beta$ (e.g., 2 and 3 in our example).
- p is **tame** if it is not wild, and either $z \notin \mathbb{Z}_p^\times$ or $z - 1 \notin \mathbb{Z}_p^\times$.

This L-function is associated to a specific scheme defined in terms of $(\alpha, \beta), z$. However, there is no **distinguished** choice of this scheme.

Hypergeometric data

A **hypergeometric datum** of degree r consists of two disjoint tuples $(\alpha_1, \ldots, \alpha_r), (\beta_1, \ldots, \beta_r)$ over $\mathbb{Q} \cap [0,1)$ which are each **balanced**: the multiplicity of any reduced fraction depends only on its denominator. Later we will consider the example

$$\alpha = (\tfrac{1}{4}, \tfrac{1}{2}, \tfrac{1}{2}, \tfrac{3}{4}),\ \beta = (\tfrac{1}{3}, \tfrac{1}{3}, \tfrac{2}{3}, \tfrac{2}{3}).$$

To each such datum, we can define a family of arithmetic L-functions of degree r over $\mathbb{Q}$ parametrized by $z \in \mathbb{Q} \setminus \{0,1\}$. The primes p of bad reduction have the following forms.

- p is **wild** if $\gamma \notin \mathbb{Z}_p$ for some $\gamma \in \alpha \cup \beta$ (e.g., 2 and 3 in our example).
- p is **tame** if it is not wild, and either $z \notin \mathbb{Z}_p^\times$ or $z - 1 \notin \mathbb{Z}_p^\times$.

This L-function is associated to a specific scheme defined in terms of $(\alpha, \beta), z$. However, there is no **distinguished** choice of this scheme.

Trace formulas

In order to add an L-function to the LMFDB, we need the first X coefficients of the Dirichlet series, for X on the order* of 2^{24}. It is sufficient to get the prime-power coefficients, as the others can be recovered using unique factorization.

The Euler factor at a prime p can be interpreted as the reverse charpoly of a matrix F_p. To get the desired Dirichlet coefficients, it suffices to compute the trace of F_p^f for all prime powers $q = p^f \leq X$. Note that for any fixed f, we need $q = p^f$ for $p \leq X^{1/f}$.

In similar situations, this is done by constructing F_p from a Weil cohomology theory (étale or p-adic). In this case, we instead use a direct trace formula based on **finite hypergeometric sums** (Greene, Katz, McCarthy, Beukers–Cohen–Mellit), plus the **Gross–Koblitz formula** for Gauss sums in terms of p-adic functions (Cohen–Rodriguez Villegas).

*The precise cutoff depends on the **conductor** of the L-function.

Trace formulas

In order to add an L-function to the LMFDB, we need the first X coefficients of the Dirichlet series, for X on the order* of 2^{24}. It is sufficient to get the prime-power coefficients, as the others can be recovered using unique factorization.

The Euler factor at a prime p can be interpreted as the reverse charpoly of a matrix F_p. To get the desired Dirichlet coefficients, it suffices to compute the trace of F_p^f for all prime powers $q = p^f \leq X$. Note that for any fixed f, we need $q = p^f$ for $p \leq X^{1/f}$.

In similar situations, this is done by constructing F_p from a Weil cohomology theory (étale or p-adic). In this case, we instead use a direct trace formula based on **finite hypergeometric sums** (Greene, Katz, McCarthy, Beukers–Cohen–Mellit), plus the **Gross–Koblitz formula** for Gauss sums in terms of p-adic functions (Cohen–Rodriguez Villegas).

*The precise cutoff depends on the **conductor** of the L-function.

Trace formulas

In order to add an L-function to the LMFDB, we need the first X coefficients of the Dirichlet series, for X on the order* of 2^{24}. It is sufficient to get the prime-power coefficients, as the others can be recovered using unique factorization.

The Euler factor at a prime p can be interpreted as the reverse charpoly of a matrix F_p. To get the desired Dirichlet coefficients, it suffices to compute the trace of F_p^f for all prime powers $q = p^f \leq X$. Note that for any fixed f, we need $q = p^f$ for $p \leq X^{1/f}$.

In similar situations, this is done by constructing F_p from a Weil cohomology theory (étale or p-adic). In this case, we instead use a direct trace formula based on **finite hypergeometric sums** (Greene, Katz, McCarthy, Beukers–Cohen–Mellit), plus the **Gross–Koblitz formula** for Gauss sums in terms of p-adic functions (Cohen–Rodriguez Villegas).

*The precise cutoff depends on the **conductor** of the L-function.

A preview of the formula

For $q = p^f$, the trace of F_p^f is given by

$$H_q\left(\begin{matrix}\alpha \\ \beta\end{matrix}\middle| z\right) := \frac{1}{1-q}\sum_{m=0}^{q-2}(-p)^{\eta_m(\alpha)-\eta_m(\beta)} q^{D+\xi_m(\beta)} \left(\prod_{j=1}^{r}\frac{(\alpha_j)_m^*}{(\beta_j)_m^*}\right)[z]^m,$$

where $(\gamma)_m^*$ is a p-adic variant of the Pochhammer symbol $(\gamma)_m = \gamma(\gamma+1)\cdots(\gamma+m-1)$ defined using the Morita p-adic Gamma function (see §2.1); $[z]$ is the multiplicative lift† of the reduction of z modulo p; and η_m, ξ_m are discrete invariants independent of z (see §2.2).

More discussion of this formula will take place in the live session. For the moment, note that the sum has $q-1$ terms.

†Commonly called the **Teichmüller lift**, but I recommend phasing out this eponym.

A preview of the formula

For $q = p^f$, the trace of F_p^f is given by

$$H_q\left(\begin{matrix}\alpha \\ \beta\end{matrix}\middle| z\right) := \frac{1}{1-q}\sum_{m=0}^{q-2}(-p)^{\eta_m(\alpha)-\eta_m(\beta)} q^{D+\xi_m(\beta)} \left(\prod_{j=1}^{r}\frac{(\alpha_j)^*_m}{(\beta_j)^*_m}\right)[z]^m,$$

where $(\gamma)^*_m$ is a p-adic variant of the Pochhammer symbol $(\gamma)_m = \gamma(\gamma+1)\cdots(\gamma+m-1)$ defined using the Morita p-adic Gamma function (see §2.1); $[z]$ is the multiplicative lift† of the reduction of z modulo p; and η_m, ξ_m are discrete invariants independent of z (see §2.2).

More discussion of this formula will take place in the live session. For the moment, note that the sum has $q-1$ terms.

†Commonly called the **Teichmüller lift**, but I recommend phasing out this eponym.

Amortization over primes

The trace formula is implemented in Magma and Sage. For each q its complexity is $O(q)$ (with small constants), so computing all Dirichlet coefficients up to X incurs complexity $O(X^2)$ (modulo log factors), dominated by the case $f = 1$. (The remaining cases add up to $O(X^{3/2})$.)

However, the shape of the formula makes it feasible to amortize this complexity over q, so that the complexity for each trace is $\mathrm{polylog}(X)$. We establish a partial result, restricting to $f = 1$ and reducing modulo p.

Theorem (Theorem 5.26 of the paper; details in §4, §5.1, §5.2)

We exhibit an algorithm to compute $H_p\left(\begin{matrix}\alpha\\ \beta\end{matrix}\middle| z\right) \pmod p$ *for all primes* $p \leq X$. *For fixed* α, β, z, *the complexity is* $O(X)$ *modulo log factors.*

We have implemented this in Sage/Cython (plus C code by Sutherland for remainder forests). The change from $O(X^2)$ to $O(X)$ appears clearly...

Amortization over primes

The trace formula is implemented in `Magma` and `Sage`. For each q its complexity is $O(q)$ (with small constants), so computing all Dirichlet coefficients up to X incurs complexity $O(X^2)$ (modulo log factors), dominated by the case $f = 1$. (The remaining cases add up to $O(X^{3/2})$.)

However, the shape of the formula makes it feasible to amortize this complexity over q, so that the complexity for each trace is $\mathrm{polylog}(X)$. We establish a partial result, restricting to $f = 1$ and reducing modulo p.

Theorem (Theorem 5.26 of the paper; details in §4, §5.1, §5.2)

We exhibit an algorithm to compute $H_p\left(\begin{matrix}\alpha\\ \beta\end{matrix}\middle| z\right) \pmod{p}$ *for all primes* $p \leq X$*. For fixed* α, β, z*, the complexity is* $O(X)$ *modulo log factors.*

We have implemented this in Sage/Cython (plus C code by Sutherland for remainder forests). The change from $O(X^2)$ to $O(X)$ appears clearly...

Amortization over primes

The trace formula is implemented in `Magma` and `Sage`. For each q its complexity is $O(q)$ (with small constants), so computing all Dirichlet coefficients up to X incurs complexity $O(X^2)$ (modulo log factors), dominated by the case $f = 1$. (The remaining cases add up to $O(X^{3/2})$.)

However, the shape of the formula makes it feasible to amortize this complexity over q, so that the complexity for each trace is $\mathrm{polylog}(X)$. We establish a partial result, restricting to $f = 1$ and reducing modulo p.

Theorem (Theorem 5.26 of the paper; details in §4, §5.1, §5.2)

We exhibit an algorithm to compute $H_p \left(\begin{matrix} \alpha \\ \beta \end{matrix} \middle| z \right) \pmod{p}$ *for all primes* $p \leq X$. *For fixed* α, β, z, *the complexity is* $O(X)$ *modulo log factors.*

We have implemented this in Sage/Cython (plus C code by Sutherland for remainder forests). The change from $O(X^2)$ to $O(X)$ appears clearly...

Amortization over primes

The trace formula is implemented in `Magma` and `Sage`. For each q its complexity is $O(q)$ (with small constants), so computing all Dirichlet coefficients up to X incurs complexity $O(X^2)$ (modulo log factors), dominated by the case $f = 1$. (The remaining cases add up to $O(X^{3/2})$.)

However, the shape of the formula makes it feasible to amortize this complexity over q, so that the complexity for each trace is $\mathrm{polylog}(X)$. We establish a partial result, restricting to $f = 1$ and reducing modulo p.

Theorem (Theorem 5.26 of the paper; details in §4, §5.1, §5.2)

We exhibit an algorithm to compute $H_p\left(\begin{matrix}\alpha\\ \beta\end{matrix}\middle| z\right) \pmod p$ *for all primes* $p \leq X$. *For fixed* α, β, z, *the complexity is* $O(X)$ *modulo log factors.*

We have implemented this in Sage/Cython (plus C code by Sutherland for remainder forests). The change from $O(X^2)$ to $O(X)$ appears clearly...

Timings

In this example $\alpha = (\frac{1}{4}, \frac{1}{2}, \frac{1}{2}, \frac{3}{4}), \beta = (\frac{1}{3}, \frac{1}{3}, \frac{2}{3}, \frac{2}{3}), z = \frac{1}{5}$. This L-function has weight 1, so $H_p\left(\begin{matrix}\alpha\\ \beta\end{matrix}\middle| z\right)$ is uniquely determined by its reduction mod p. (See §5.4 of the paper for more implementation details, and §5.5 for a worked example.)

X	Amortized	Sage	Magma
2^{10}	0.07s	0.39s	0.11s
2^{11}	0.05s	0.68s	0.35s
2^{12}	0.06s	2.12s	1.29s
2^{13}	0.08s	7.39s	4.83s
2^{14}	0.12s	26.0s	18.2s
2^{15}	0.18s	92.3s	68.4s
2^{16}	0.34s	343s	280s
2^{17}	0.80s	1328s	1190s

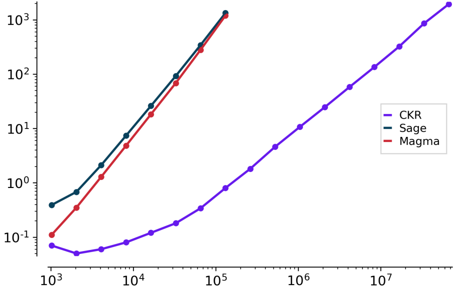

X	2^{18}	2^{19}	2^{20}	2^{21}	2^{22}	2^{23}	2^{24}	2^{25}	2^{26}
Amortized	1.81s	4.59s	10.7s	24.6s	58.0s	135s	322s	857s	1948s

Remainder trees

The key to amortizing is to reduce to subproblems of the following form: given a square matrix M over $\mathbb{Z}[x]$ and a function $k(p)$, compute

$$M(0) \cdots M(k(p) - 1) \pmod{p}$$

for all primes p in some arithmetic progression.

This can be done using **remainder trees/forests**, inspired by the fast Fourier transform. For more details, see Sutherland's talk (and §3).

We take M to be 2×2 triangular; the diagonal entries capture factorial-like products and the off-diagonal captures summation (see §4, §5.1, §5.2).

The mod-p restriction can probably be removed; this would simplify computing Dirichlet coefficients up to X from $O(X^2)$ to $O(X^{3/2})$. The restriction to prime Frobenius traces is subtler (see §2.2.2, §6.1, §6.2).

More details about these points will be given in the live session.

Remainder trees

The key to amortizing is to reduce to subproblems of the following form: given a square matrix M over $\mathbb{Z}[x]$ and a function $k(p)$, compute

$$M(0) \cdots M(k(p) - 1) \pmod p$$

for all primes p in some arithmetic progression.

This can be done using **remainder trees/forests**, inspired by the fast Fourier transform. For more details, see Sutherland's talk (and §3).

We take M to be 2×2 triangular; the diagonal entries capture factorial-like products and the off-diagonal captures summation (see §4, §5.1, §5.2).

The mod-p restriction can probably be removed; this would simplify computing Dirichlet coefficients up to X from $O(X^2)$ to $O(X^{3/2})$. The restriction to prime Frobenius traces is subtler (see §2.2.2, §6.1, §6.2).

More details about these points will be given in the live session.

Remainder trees

The key to amortizing is to reduce to subproblems of the following form: given a square matrix M over $\mathbb{Z}[x]$ and a function $k(p)$, compute

$$M(0) \cdots M(k(p) - 1) \pmod{p}$$

for all primes p in some arithmetic progression.

This can be done using **remainder trees/forests**, inspired by the fast Fourier transform. For more details, see Sutherland's talk (and §3).

We take M to be 2×2 triangular; the diagonal entries capture factorial-like products and the off-diagonal captures summation (see §4, §5.1, §5.2).

The mod-p restriction can probably be removed; this would simplify computing Dirichlet coefficients up to X from $O(X^2)$ to $O(X^{3/2})$. The restriction to prime Frobenius traces is subtler (see §2.2.2, §6.1, §6.2).

More details about these points will be given in the live session.

Remainder trees

The key to amortizing is to reduce to subproblems of the following form: given a square matrix M over $\mathbb{Z}[x]$ and a function $k(p)$, compute

$$M(0) \cdots M(k(p)-1) \pmod{p}$$

for all primes p in some arithmetic progression.

This can be done using **remainder trees/forests**, inspired by the fast Fourier transform. For more details, see Sutherland's talk (and §3).

We take M to be 2×2 triangular; the diagonal entries capture factorial-like products and the off-diagonal captures summation (see §4, §5.1, §5.2).

The mod-p restriction can probably be removed; this would simplify computing Dirichlet coefficients up to X from $O(X^2)$ to $O(X^{3/2})$. The restriction to prime Frobenius traces is subtler (see §2.2.2, §6.1, §6.2).

More details about these points will be given in the live session.

Remainder trees

The key to amortizing is to reduce to subproblems of the following form: given a square matrix M over $\mathbb{Z}[x]$ and a function $k(p)$, compute

$$M(0) \cdots M(k(p) - 1) \pmod{p}$$

for all primes p in some arithmetic progression.

This can be done using **remainder trees/forests**, inspired by the fast Fourier transform. For more details, see Sutherland's talk (and §3).

We take M to be 2×2 triangular; the diagonal entries capture factorial-like products and the off-diagonal captures summation (see §4, §5.1, §5.2).

The mod-p restriction can probably be removed; this would simplify computing Dirichlet coefficients up to X from $O(X^2)$ to $O(X^{3/2})$. The restriction to prime Frobenius traces is subtler (see §2.2.2, §6.1, §6.2).

More details about these points will be given in the live session.

Table of contents