Enumerating mathematical objects in the cloud

Andrew V. Sutherland

Massachusetts Institute of Technology

May 22, 2022

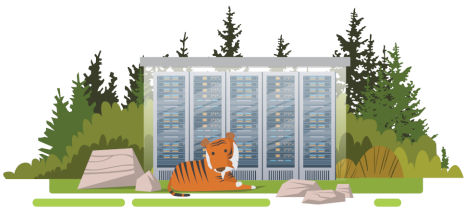

This work was supported by NSF grant DMS-1522526 and Simons Foundation grant 550033.

LuCaNT conference announcement

Next summer the very first conference on the LMFDB, Computation, and Number Theory (LuCaNT) will take place July 10–14, 2023 at the Institute for Computational and Experimental Research in Mathematics (ICERM) in Providence, Rhode Island.

This conference is broadly focused on the topics of the LMFDB, mathematical databases, computation, number theory, and arithmetic geometry. It will include invited talks, as well as presentations by authors of papers submitted to the conference and selected by the scientific committee following peer-review that will be published in a proceedings volume for the conference. The call for papers will appear in August.

You are all cordially invited to submit papers, and to attend!

The L-functions and modular forms database (LMFDB)

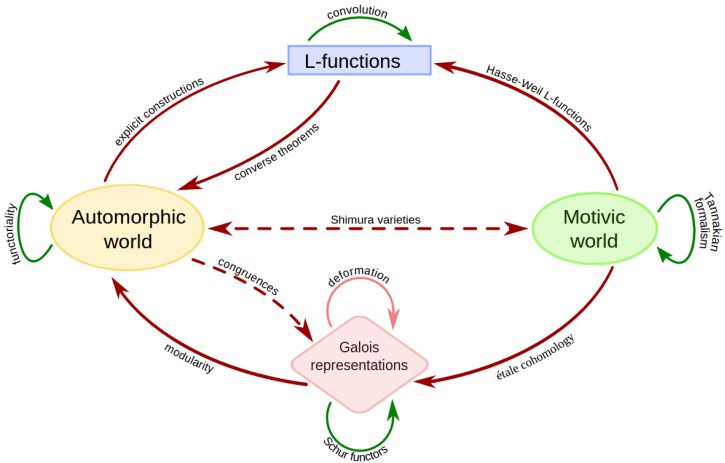

For a gentle introduction, see Cremona-Jones-S-Voight, *The L-functions and modular forms database*, Notices of the AMS **68** (2021), 1520–1522.

The L-functions and modular forms database (LMFDB)

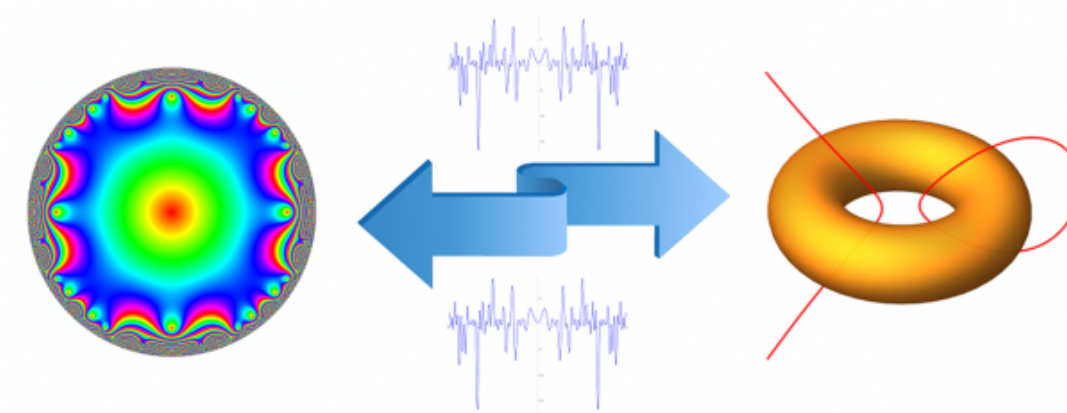

modular form
11.2.a.a

L-function
2-11-1.1-c1-0-0

elliptic curve
11.a

A brief history of LMFDB-related cloud computing

2015 Enumerated $\approx 10^{17.5}$ genus 2 curves using 72 000 vCPUs (≈ 50 vCPU-years)[1] to find 66,158 with discriminant $|\Delta| \leq 10^6$.

2017 Enumerated $\approx 10^{18}$ genus 3 curves using 580 000 vCPUs (≈ 300 vCPU-years)[2] to find 67,879 hyperelliptic and 82 240 nonhyperelliptic curves with $|\Delta| \leq 10^7$.

2018 Enumerated 14 417 694 modular forms (281 965 Galois orbits) with weights $1 \leq k \leq 316$ and levels $1 \leq N \leq 10,000$ (≈ 100 vCPU-years).[3]

2019 Enumerated 581 056 $E/\mathbb{Q}$ of conductor $(400\,000, 500\,000]$ (≈ 100 vCPU-years).[4]

2019 Enumerated $\approx 10^{18}$ genus 3 curves (hyperelliptic/nonhyperelliptic/Picard/covers) to find 21 188 748 curves with 7-smooth Δ (≈ 300 vCPU-years).[5]

[1] Booker-Sijsling-S-Voight-Yasaki, *A database of genus 2 curves*, ANTS XII, 2016.
[2] S, *A database of nonhyperelliptic genus 3 curves over* $\mathbb{Q}$, ANTS XIII, 2018.
[3] BBBCCDLLRSV, *Computing classical modular forms*, 2021.
[4] Cremona-S, in celebration of Stoll's Rational Points 2019 workshop.
[5] Fité-Kedlaya-S, *Sato-Tate groups of abelian threefolds*, 2021.

Why cloud computing?

- I'm not willing to wait 300 years for one core to finish the job.

- It gives you exactly the computer you need when you need it.

- I care about the planet.

An inconvenient truth

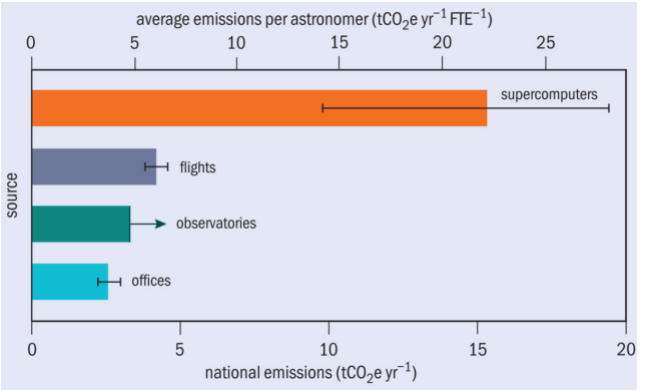

Annual carbon footprint of the average Australian astronomer 2018–2019.[6]

[6]Image taken from *The huge carbon footprint of large scale computing*, Physics World, March 22, by Michael Allen. Data source: Stevens-Bellstedt-Elahi-Murphy, *The imperative to reduce carbon emissions in astronomy*, Nature Astronomy **4** (2020), 843–851.

How much carbon does a 300 vCPU-year computation emit?

This is a question http://www.green-algorithms.org/ can help answer.[7]

300 vCPU-years is about 1 314 900 core-hours (2 vCPUs per core).

CPU	cores	platform	location	energy	carbon
i9-9900K	1	desktop	Massachusetts	78.34 MWh	27 420 Kg
i9-9900K	16	desktop	Massachusetts	19.57 MWh	6 850 Kg
Ryzen 3990X	16	desktop	Massachusetts	9.70 MWh	3 400 Kg
Ryzen 3990X	64	desktop	Massachusetts	6.77 MWh	2 370 Kg
Ryzen 3990X	64	server	Massachusetts	11.30 MWh	3 950 Kg
Ryzen 3990X	64	cloud	Virginia	7.51 MWh	2 660 Kg
Ryzen 3990X	64	cloud	Montreal	7.51 MWh	13 Kg

These figures exclude manufacturing and disposal (desktop $\approx$ 2x, cloud $\approx$ 1.1x).

[7] Lannelongue-Grealey-Inouye, *Green algorithms: Quantifying the carbon footprint of computation*, Advanced Science **8** (2021), article 2100707.

Discriminants and conductors

Every hyperelliptic curve $X/\mathbb{Q}$ of genus g has a minimal Weierstrass model

$$y^2 + h(x)y = f(x),$$

with $f, h \in \mathbb{Z}[x]$, $\deg f \leq 2g+2$, $\deg h \leq g+1$. We can assume without loss of generality that $h(x)$ has coefficients in $\{0,1\}$. The discriminant of X is then

$$\Delta(X) := 2^{-(4g+4)} \operatorname{disc}_{2g+2}(4f + h^2) \in \mathbb{Z}[f_0, \ldots, f_{2g+2}, h_0, \ldots h_{g+1}],$$

and the curve X has bad reduction at a prime p if and only if $p|\Delta(X)$.

This need not apply to the Jacobian of X, but if $p|N(X)$, then $p|\Delta(X)$.

L-functions are naturally ordered by their conductor N, which corresponds to the level of the corresponding automorphic form (under the Langlands correspondence).

Computing the discriminant

Using the standard (resultant-based) algorithm for computing the discriminant of a genus 2 curve with 8-bit coefficients yields the following timings:

implementation	time (μs)
SageMath	23.6
Magma	22.7
PARI/GP	7.24
C code (mpz_t)	2.45
C code (128-bit)	0.786
C code (64-bit)	0.319

4.4GHz Intel i9-9960X

Evaluating multivariate polynomials with monomial trees

Suppose we want to evaluate a polynomial $P(x_1, \ldots, x_n)$ at every point in a box $A_1 \times \cdots \times A_n \subset \mathbb{Z}^n$. We use a monomial tree with

- nodes at level n (leaves): monomials of $P(x_1, \ldots, x_n)$.
- nodes at level $n-1$: monomials of $P(x_1, \ldots, x_{n-1}, a_n)$.
- $\ldots$
- nodes at level 1: monomials of $P(x_1, a_2, \ldots, a_n) = P_1(x_1)$.

Nodes at level $m+1$ are connected to those at level m via an edge corresponding to the substitution $x_{m+1} = a_{m+1}$. We store a coefficient value at each node that is updated whenever we make a substitution.

At level 1 we evaluate a univariate polynomial $P'(x_1)$ of degree $\deg_{x_1}(P)$.

We can efficiently enumerate values of $P_1(x_1)$ using finite differences, or by using SIMD parallelism, such as AVX-512 instructions (but typically not both).

Monomial tree example

Consider the polynomial

$$g(a_0, a_1, a_2) \coloneqq a_0^3 a_2 + 3a_0^2 a_1^2 - 4a_0^2 a_1 a_2 - 5a_0 a_1^2 a_2 + 2a_1^4 + 7a_1^3 a_2.$$

A monomial tree for $g(a_0, a_1, a_2)$.

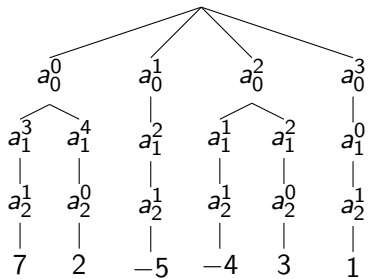

Monomial tree example

Consider the polynomial

$$g(a_0, a_1, a_2) := a_0^3 a_2 + 3a_0^2 a_1^2 - 4a_0^2 a_1 a_2 - 5a_0 a_1^2 a_2 + 2a_1^4 + 7a_1^3 a_2.$$

A better monomial tree for $g(a_0, a_1, a_2)$.

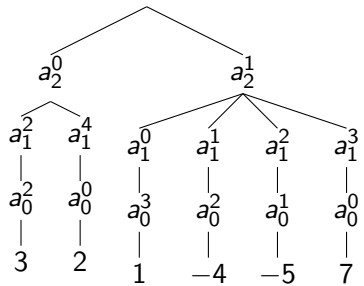

Monomial tree example

Consider the polynomial

$$g(a_0, a_1, a_2) \coloneqq a_0^3 a_2 + 3a_0^2 a_1^2 - 4a_0^2 a_1 a_2 - 5a_0 a_1^2 a_2 + 2a_1^4 + 7a_1^3 a_2.$$

Monomial tree for $g(2, a_1, a_2)$.

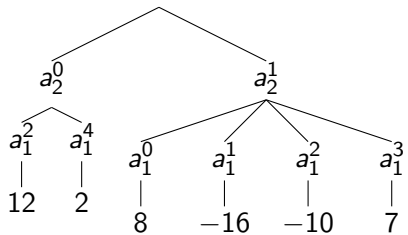

Monomial tree example

Consider the polynomial

$$g(a_0, a_1, a_2) \coloneqq a_0^3 a_2 + 3a_0^2 a_1^2 - 4a_0^2 a_1 a_2 - 5a_0 a_1^2 a_2 + 2a_1^4 + 7a_1^3 a_2.$$

Monomial tree for $g(2, -1, a_2)$.

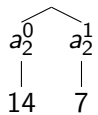

Monomial trees in practice

We compute discriminant monomial trees for hyperelliptic curves $y^2 + h(x)y = f(x)$ with $h(x)$ fixed (we assume coefficients of h are 0 or 1 and remove symmetries).

- For $g = 2$, we get 246 terms and 703 nodes in our monomial tree.
- For $g = 3$, we get 5247 terms and 19916 nodes in our monomial tree.

For nonhyperelliptic curves of genus 3 the monomial tree for Δ_4 has 50 767 957 terms and 246 798 254 nodes (for optimally ordered variables). But the inner loop simply enumerates values of a univariate quartic polynomial.

For genus 2 curves the inner loop enumerates values of a quintic polynomial using just five additions per step; in total, enumerating curves with 8-bit coefficients and their discriminants takes less than 15 (Skylake) clock cycles per curve.

Computing the discriminant

implementation	time (μs)
SageMath (resultant)	23.6
Magma (resultant)	22.7
PARI/GP (resultant)	7.24
C code (mpz_t resultant)	2.45
C code (128-bit resultant)	0.786
C code (64-bit resultant)	0.319
C code (64-bit monomial tree)	0.005

4.4GHz Intel i9-9960X

Smoothness testing takes more time, but it can also be extensively optimized. For 64-bit discriminants Δ we can test $\mathrm{rad}(\Delta) \leq 10^3$ in less than 5 nanoseconds, on average. Testing $\mathrm{rad}(\Delta) \leq 10^6$ takes 10-15 nanoseconds on average.

Tips for running computations in the cloud

Preemptible spot instances provide a cost effective way to access unused computational resources available on cloud computing platforms that are scaled for peak demand.[8]

- Compute instances may stop running at any time; this favors massive parallelization and short running times (a few hours is ideal).
- Use batch/bulk creation methods to efficiently launch instances in parallel.
- While you can restart stopped instances, I recommend auto-terminating and relaunching jobs that do not complete (this requires some administration).
- Configure a customized disk image or container that has exactly the software stack needed for the job (and nothing else).
- Be careful not to create communication bottlenecks – ideally each instance should run independently, and writes its output to a separate file in a shared storage.

[8]Major service providers include Amazon Web Services (33%), Microsoft Azure (20%), and Google Cloud Platform (10%), according to Statista's Global could infrastructure quarterly market share.

Enumerating genus 2 curves of small conductor

On January 9, 2022 we launched the last in a series of three computations exploring the Langlands correspondence in genus 2. We enumerated more than 10^{19} genus 2 curves looking for those with small conductor with the goal of expanding the LMFDB and matching L-functions of explicitly known automorphic forms of small level.

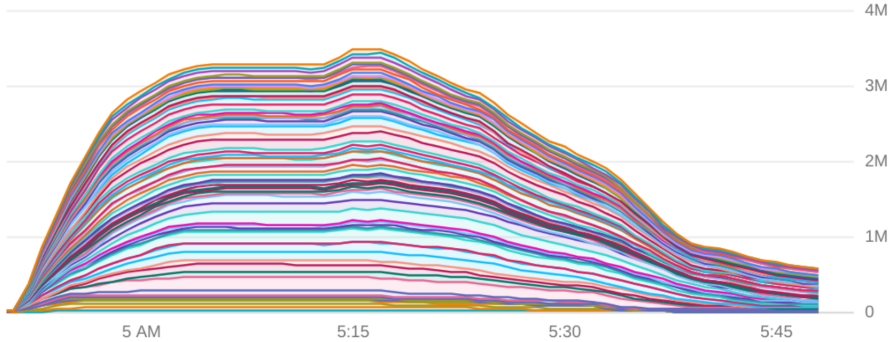

We used a total of 4,034,560 Intel/AMD cores in 73 data centers across the globe.

Enumerating genus 2 curves of small conductor

We found several million genus 2 curves of small conductor, including the curve

$$C_{903} : y^2 + (x^2 + 1)y = x^5 + 3x^4 - 13x^3 - 25x^2 + 61x - 28$$

of conductor 903 and whose L-function coefficients match those of the paramodular form of level 903 computed by Poor–Yuen. This was the last explicitly known paramodular form that had not been matched to a genus 2 curve or abelian surface.

conductor bound	1000	10000	100000	1000000
curves in LMFDB	159	3069	20265	66158
curves found	807	25438	447507	5151208
L-functions in LMFDB	109	2807	19775	65534
L-functions found	200	9409	212890	2426708