Telescopes for mathematicians

Andrew V. Sutherland

Massachusetts Institute of Technology

January 9, 2019

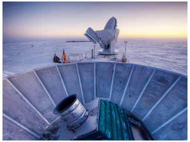

Simons Collaboration in Arithmetic Geometry, Number Theory, and Computation

What is arithmetic geometry?

Arithmetic geometers study solutions to polynomial equations like

$$y = 2x + 3, \qquad x^2 + y^2 = 1,$$

$$y^2 + y = x^3 - x^2,$$

$$y^2 + (x^3 + x + 1)y = x^5 + x^4, \qquad xy^3 + y^3z + z^3x = 0,$$

and even "cursed" examples like

$$\begin{aligned} y^4 + 5x^4 - 6x^2y^2 + 6x^3z + 26x^2yz + 10xy^2z - 10y^3z \\ -32x^2z^2 - 40xyz^2 + 24y^2z^2 + 32xz^3 - 16yz^3 = 0. \end{aligned}$$

Balakrishnan et al.
2017

Which solutions?

There is a robust theory (algebraic geometry) that addresses this problem over the complex numbers.

Number theorists are particularly interested in integer (or rational) solutions to these equations. These can be very difficult to find.

Indeed, this problem is unsolvable, in general, but it can be solved in many cases. Even when it cannot, we can simplify the problem by looking at solutions modulo primes, in other words, we can count points.

These point counts can be used to define an L-function that encodes the fundamental structure of the equation (or system of equations) in a canonical way.

Counting points modulo p

Let's count points (x, y) on the curve $C\colon x^2 + y^2 = 1$ modulo primes p:

p	2	3	5	7	11	13	17	19	23	29	...
	2	4	4	8	12	12	16	20	24	28	$p \pm 1$

Better, count points $(x, y, z) \sim (cx, cy, cz)$ on $x^2 + y^2 = z^2 \bmod p$:

p	2	3	5	7	11	13	17	19	23	29	...
	3	4	6	8	12	14	18	20	24	30	$p + 1$

We always get $p + 1$. The L-function of C is

$$L(C, s) = \prod (1 - p^{-s})^{-1} = \sum n^{-s} = \zeta(s).$$

We get the same L-function whenever C has genus 0.

Elliptic curves

Let E be an elliptic curve over $\mathbb{Q}$, which we can write as

$$E\colon y^2 = x^3 + ax + b.$$

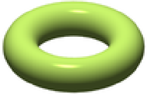

Every curve of genus 1 with rational points has this form.
You (via your phone/computer) use elliptic curves every day!

The number of points on E modulo p can be written as

$$\#E_p(\mathbb{F}_p) = p + 1 - a_p,$$

where the trace of Frobenius a_p satisfies $|a_p| \leq 2\sqrt{p}$.

H. Hasse

Let us now consider the sequence or real numbers $x_p := -t_p/\sqrt{p} \in [-2, 2]$.

Example: $y^2 = x^3 + x + 1$

p	t_p	x_p	p	t_p	x_p	p	t_p	x_p
3	0	**0.000000**	71	13	**−1.542816**	157	−13	**1.037513**
5	−3	**1.341641**	73	2	**−0.234082**	163	−25	**1.958151**
7	3	**−1.133893**	79	−6	**0.675053**	167	24	**−1.857176**
11	−2	**0.603023**	83	−6	**0.658586**	173	2	**−0.152057**
13	−4	**1.109400**	89	−10	**1.059998**	179	0	**0.000000**
17	0	**0.000000**	97	1	**−0.101535**	181	−8	**0.594635**
19	−1	**0.229416**	101	−3	**0.298511**	191	−25	**1.808937**
23	−4	**0.834058**	103	17	**−1.675060**	193	−7	**0.503871**
29	−6	**1.114172**	107	3	**−0.290021**	197	−24	**1.709929**
37	−10	**1.643990**	109	−13	**1.245174**	199	−18	**1.275986**
41	7	**−1.093216**	113	−11	**1.034793**	211	−11	**0.757271**
43	10	**−1.524986**	127	2	**−0.177471**	223	−20	**1.339299**
47	−12	**1.750380**	131	4	**−0.349482**	227	0	**0.000000**
53	−4	**0.549442**	137	12	**−1.025229**	229	−2	**0.132164**
59	−3	**0.390567**	139	14	**−1.187465**	233	−3	**0.196537**
61	12	**−1.536443**	149	14	**−1.146925**	239	−22	**1.423062**
67	12	**−1.466033**	151	−2	**0.162758**	241	22	**−1.417145**

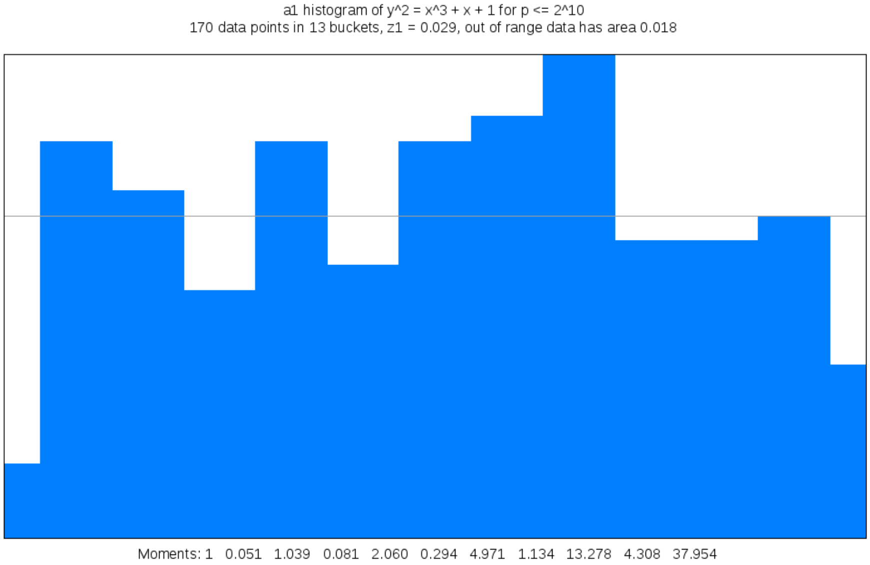

click histogram to animate (requires adobe reader)

a1 histogram of y^2 + xy + y = x^3 - x^2 - 2006776241557552658503320820933854275093023031217895 6502x
+ 34481611795030556467032985690390720374855944359319180361266008296291939448732243429 for p <= 2^10
172 data points in 13 buckets, z1 = 0.023, out of range data has area 0.250

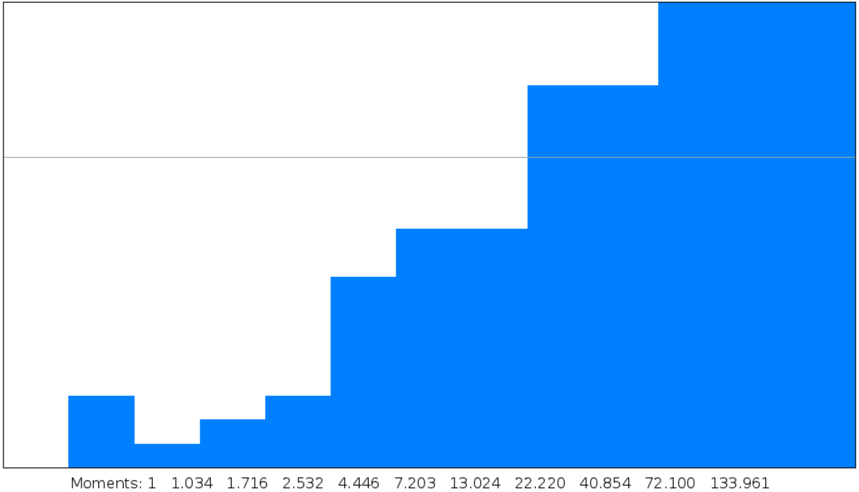

click histogram to animate (requires adobe reader)

The Sato-Tate conjecture

The Sato-Tate conjecture states that, except for certain families of well understood exceptions, we will always get the same limiting distribution as $p \to \infty$.

Mikio Sato

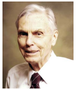

John Tate

Theorem (Taylor et al. 2008)

Let $E/\mathbb{Q}$ be an elliptic curve without extra endomorphisms.
The sequence x_p converges to the semi-circular distribution.

Richard Taylor

Richard Taylor received the 2014 Breakthrough Prize in Mathematics for this work.

a1 histogram of y^2 = x^3+1 for p <= 2^10
170 data points in 13 buckets, z1 = 0.518, out of range data has area 0.418

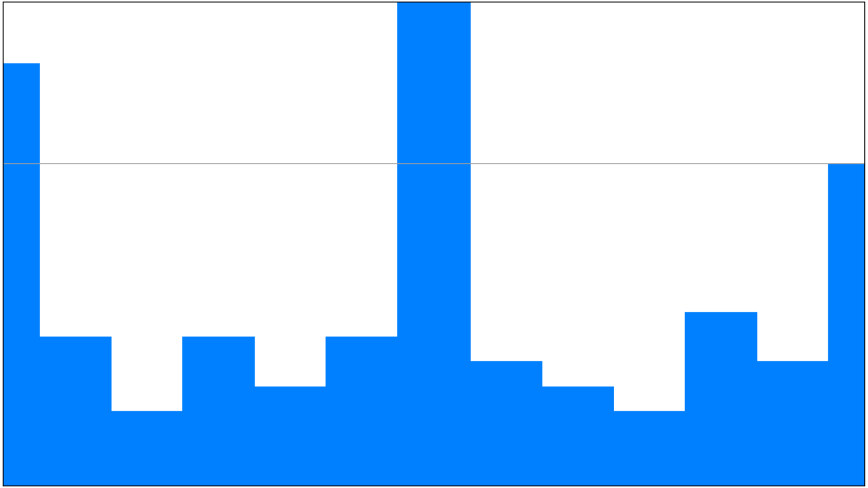

Moments: 1 -0.044 0.934 -0.160 2.754 -0.660 9.051 -2.655 31.232 -10.427 110.831

click histogram to animate (requires adobe reader)

Modularity

The proof of the Sato-Tate conjecture is built on the Modularity Theorem.

Theorem (Taylor-Wiles 1995, Breuil-Conrad-Diamond-Taylor 2001)

For every elliptic curve $E/\mathbb{Q}$ there is a modular form f_E for which $L(E,s) = L(f_E,s)$. The q-expansion $f_E(q) = \sum a_n q^n$ of f_E is determined by the Frobenius traces a_p of E.

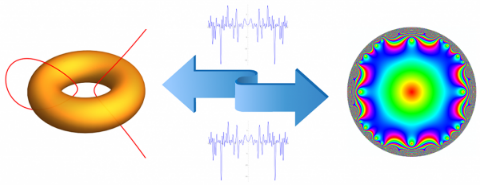

Corollary (Wiles 1995)

The equation $x^n + y^n = z^n$ has no nontrivial integer solutions for $n > 2$.

The L-functions and Modular Forms Database (LMFDB)

The relationship between elliptic curves and modular forms established by the Modularity Theorem was conjectured fifty years earlier.

Compelling evidence for this conjecture was obtained over many decades by tabulating elliptic curves and modular forms and computing their L-functions.

Extensive tables of these (and many other mathematical objects) are now available in the L-functions and Modular Forms Database.

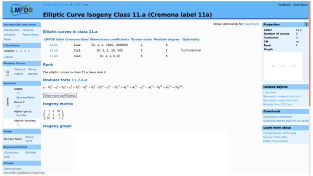

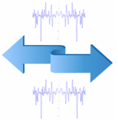

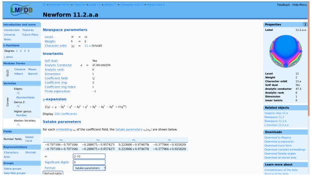

Ranks of elliptic curves

The rational points on an elliptic curve $E/\mathbb{Q}$ are generated by a finite set of points. The minimal number of infinite order generators is the rank r.

There are many things we do not know about r:

- Is there an algorithm that is guaranteed to compute r?
- Which values of r can occur? Is there an upper limit?
- How often does each possible value of r occur, on average?

Theorem (Elkies 1990)

The value of r can be as large as 28.

Theorem (Bhargava-Shankar 2012)

The average value of r lies between 0 and 1.

The Birch and Swinnerton-Dyer conjecture

Based on extensive computer experiments (in the early 1960s!), Bryan Birch and Sir Peter Swinnerton-Dyer made the following conjecture.

Conjecture (Birch and Swinnerton-Dyer)

Let $E/\mathbb{Q}$ be an elliptic curve of rank r. Then $L(E,s) = (s-1)^r g(s)$, for some $g(s)$ with $g(1) \neq 0$, in other words, r is the order of vanishing of $L(E,s)$ at 1.

Birch

EDSAC-2

Swinnerton-Dyer

They later made a more precise conjecture that gives the leading coefficient of $g(s)$.

The Langlands Program

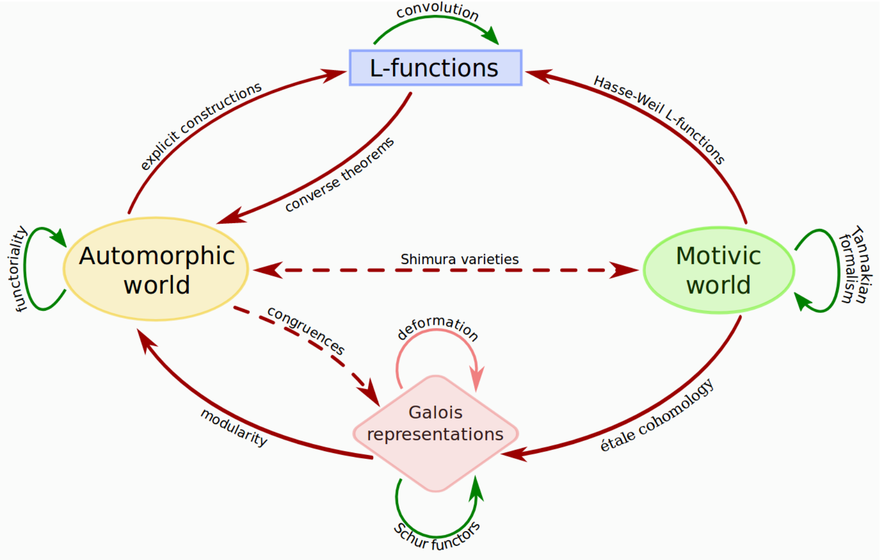

The L-function of a curve

The L-function of a (nice) curve $X/\mathbb{Q}$ can be written as

$$L(X,s) := \prod_p L_p(p^{-s})^{-1}.$$

For good primes p the polynomial $L_p \in \mathbb{Z}[T]$ is the numerator of the zeta function

$$Z(X_p;T) := \exp\left(\sum_{r\geq 1} \#X_p(\mathbb{F}_{p^r})\frac{T^r}{r}\right) = \frac{L_p(T)}{(1-T)(1-pT)}.$$

Zeta functions can be computed by counting points.

Under the Langlands philosophy, $L(X,s)$ is completely determined by point counts modulo "sufficiently many" good primes p.

How many we need depends on the conductor of $L(X,s)$.

Algorithms to compute zeta functions

Given $X/\mathbb{Q}$ of genus g, we want to compute $L_p(T)$ for all good $p \le B$.

	complexity per prime (ignoring factors of $O(\log\log p)$)		
algorithm	$g = 1$	$g = 2$	$g = 3$
point enumeration	$p \log p$	$p^2 \log p$	$p^3(\log p)^2$
group computation	$p^{1/4} \log p$	$p^{3/4} \log p$	$p(\log p)^2$
p-adic cohomology	$p^{1/2}(\log p)^2$	$p^{1/2}(\log p)^2$	$p^{1/2}(\log p)^2$
CRT (Schoof-Pila)	$(\log p)^5$	$(\log p)^8$	$(\log p)^{12}$
average poly-time	$(\log p)^4$	$(\log p)^4$	$(\log p)^4$

The bottom row is due to a 2014 breakthrough by David Harvey, followed by further refinements [Harvey-S 2016, 2018].

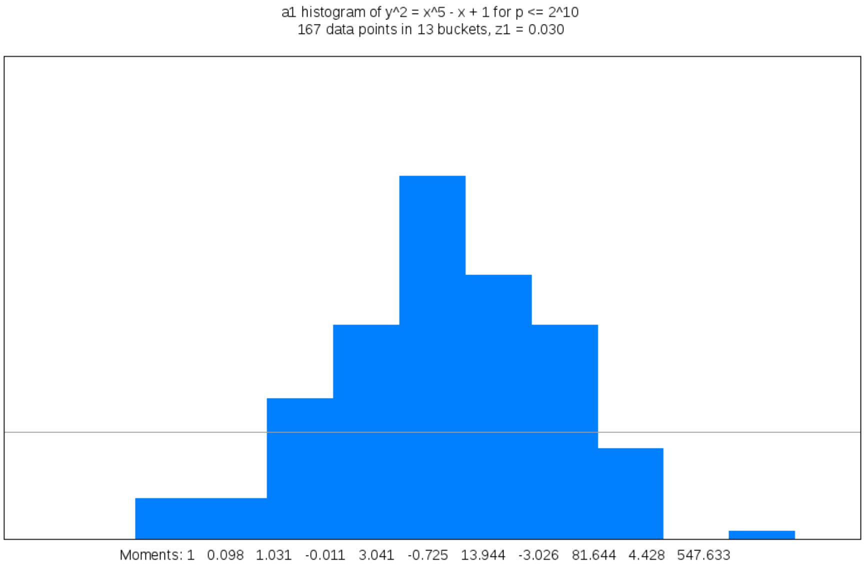

click histogram to animate (requires adobe reader)

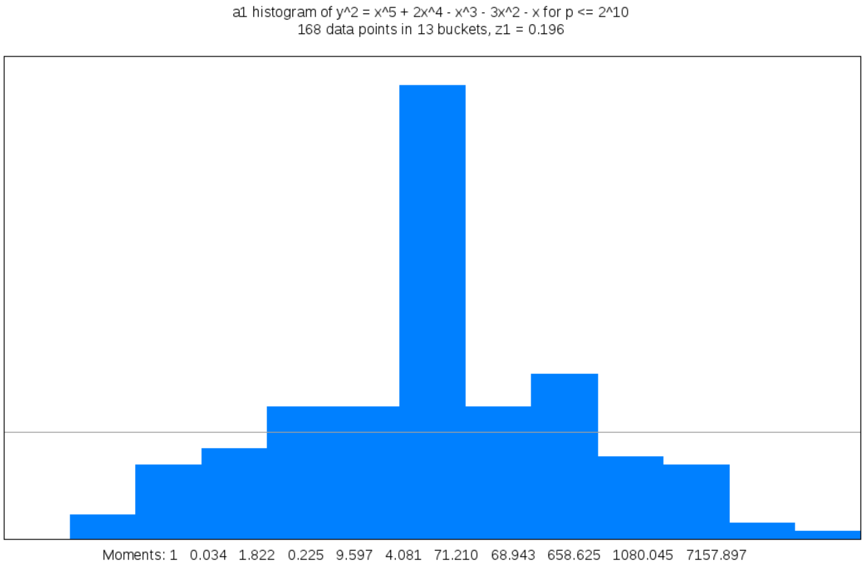

click histogram to animate (requires adobe reader)

Exceptional Sato-Tate distributions for genus 2 curves over $\mathbb{Q}$:

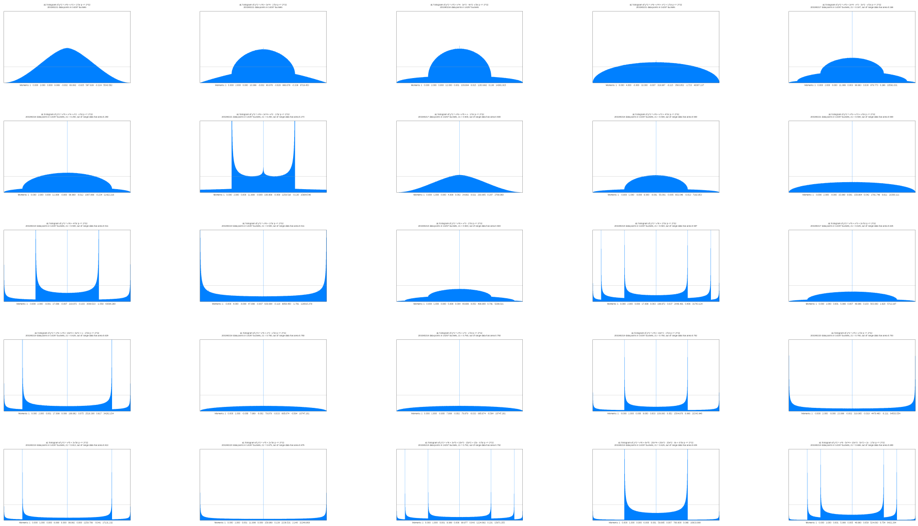

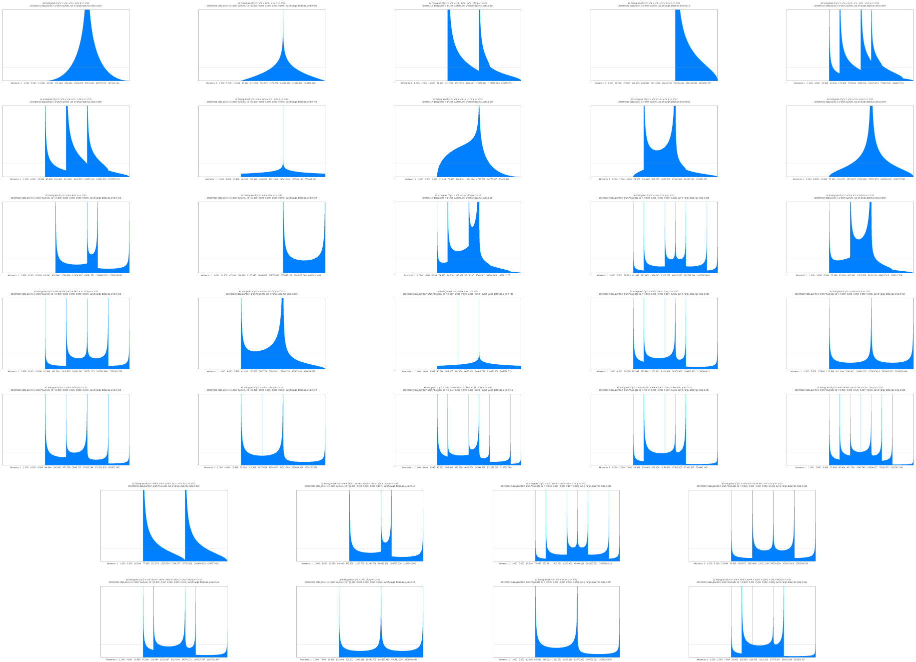

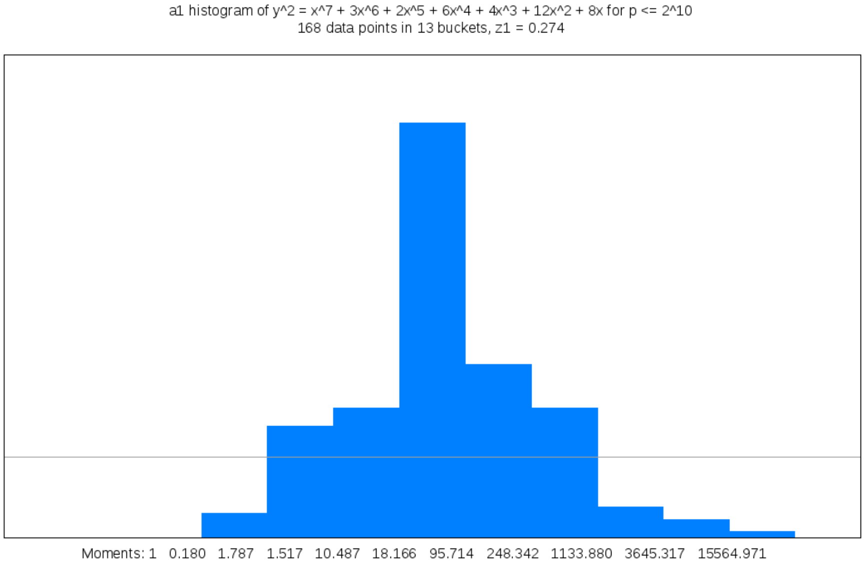

click histogram to animate (requires adobe reader)

Timings for genus 3 curves

Time to compute $L_p(T)$ mod p for all good $p \leq B$.

B	spq-Costa-AKR	hyp-Harvey	spq-HS	hyp-HS
2^{12}	18	1.3	1.4	0.1
2^{13}	49	2.6	2.4	0.2
2^{14}	142	5.4	4.6	0.5
2^{15}	475	12	9.4	1.0
2^{16}	1,670	29	21	2.1
2^{17}	5,880	74	47	5.3
2^{18}	22,300	192	112	14
2^{19}	78,100	532	241	37
2^{20}	297,000	1,480	551	97
2^{21}	1,130,000	4,170	1,240	244
2^{22}	4,280,000	12,200	2,980	617
2^{23}	16,800,000	36,800	6,330	1,500
2^{24}	66,800,000	113,000	14,200	3,520
2^{25}	244,000,000	395,000	31,900	8,220
2^{26}	972,000,000	1,060,000	83,300	19,700

(Intel Xeon E7-8867v3 3.3 GHz CPU seconds).

Building a database of low genus curves

To make it feasible to compute L-functions, and to facilitate investigation of the Langlands correspondence, we want to tabulate curves by conductor.

No one knows how to do this for curves of genus $g > 1$, not even in principle!

We instead "sieve the sky". We enumerate vast numbers of curves with small coefficients along with their discriminants (which bound the conductor).

In our genus 2 and genus 3 searches we enumerated a total of about 10^{18} curves. In each case we kept roughly 10^5 curves of interest.

To make such a computation feasible requires:

- extremely efficient enumeration algorithms;
- code optimization "down to the metal";
- massive parallelism.

Parallel computation

The genus 3 computation was parallelized and run on Google's Cloud Platform.
We spread the load across 24 data centers in nine geographic zones.

For the smooth plane quartic search we used $19,000$ pre-emptible 32-vCPU instances.
At peak usage we had 580,000 vCPUs running at full load (a new record).

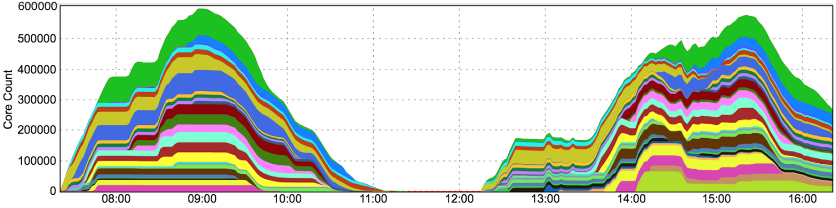

This 300 vCPU-year computation took about 10 hours.