

High-precision arithmetic and PSLQ

David H. Bailey

<http://www.davidhbailey.com>

Lawrence Berkeley National Lab (retired)

University of California, Davis

October 2, 2017

Numerical reproducibility in scientific computing

A December 2012 workshop on reproducibility in computing, held at Brown University, USA, noted that

Science is built upon the foundations of theory and experiment validated and improved through open, transparent communication. ...

Numerical round-off error and numerical differences are greatly magnified as computational simulations are scaled up to run on highly parallel systems. As a result, it is increasingly difficult to determine whether a code has been correctly ported to a new system, because computational results quickly diverge from standard benchmark cases. And it is doubly difficult for other researchers, using independently written codes and distinct computer systems, to reproduce published results.

- ▶ V. Stodden, D. H. Bailey, J. M. Borwein, R. J. LeVeque, W. Rider and W. Stein, "Setting the default to reproducible: Reproducibility in computational and experimental mathematics," <http://www.davidhbailey.com/dhbpapers/icerm-report.pdf>.

Numerical reliability

Many applications routinely use either 32-bit or 64-bit IEEE arithmetic, and employ fairly simple algorithms, assuming that all is well. But problems can arise:

1. Large-scale, highly parallel simulations, running on systems with hundreds of thousands or millions of processors — numerical sensitivities are greatly magnified.
 2. Certain applications with highly ill-conditioned linear systems.
 3. Large summations, especially those involving $+/-$ terms and cancellations.
 4. Long-time, iterative simulations (such as molecular dynamics or climate models).
 5. Computations to resolve small-scale phenomena.
 6. **Studies in computational physics and experimental mathematics.**
- D. H. Bailey, R. Barrio, and J. M. Borwein, "High precision computation: Mathematical physics and dynamics," *Applied Mathematics and Computation*, vol. 218 (2012), pg. 10106–10121.

Analysis of collisions at the Large Hadron Collider

- ▶ The 2012 discovery of the Higgs boson at the ATLAS experiment in the LHC relied crucially on the ability to track charged particles with exquisite precision (10 microns over a 10m length) and high reliability (over 99% of roughly 1000 charged particles per collision correctly identified).
- ▶ Software: 5 millions line of C++ and Python code, developed by roughly 2000 physicists and engineers over 15 years.
- ▶ Recently, in an attempt to speed up the calculation, researchers found that merely changing the underlying math library resulted in some collisions being missed or misidentified.

Questions:

- ▶ How serious are these numerical difficulties?
- ▶ How can they be tracked down?
- ▶ How can the library be maintained, producing numerically reliable results?

Ways to enhance numerical reproducibility and solve accuracy problems

1. Employ an expert numerical analyst to re-examine every algorithm employed in the computation to ensure that only the most stable known schemes are being used.
2. Carefully analyze every step of the computation to ascertain the level of numerical sensitivity at each step.
3. Employ interval arithmetic for large portions of the application (which greatly increases run time and code complexity).
4. Employ higher-precision arithmetic, assisted with some “smart” tools to help determine where extra precision is needed and where it is not.

Item #4 is often the only practical solution.

Numerical analysis expertise among U.C. Berkeley graduates

Of the 2010 U.C. Berkeley graduating class, 870 were in disciplines likely to require technical computing:

- ▶ Division of Mathematical and Physical Sciences (Math, Physics, Statistics).
- ▶ College of Chemistry.
- ▶ College of Engineering (including Computer Science).

Other fields (not counted) that will likely involve significant computing:

- ▶ Biology, geology, medicine, economics, psychology, sociology.

Enrollment in numerical analysis courses:


- ▶ Math 128A (Introductory numerical analysis required of math majors): 219.
- ▶ Math 128B (A more advanced course, required to do serious work): 24.

Conclusion: Fewer than 2% of U.C. Berkeley graduates who will do technical computing in their careers have had rigorous training in numerical analysis.

Enhancing reproducibility with selective usage of high-precision arithmetic

Problem: Find the arc length of the irregular function $g(x) = x + \sum_{0 \leq k \leq 10} 2^{-k} \sin(2^k x)$, over the interval $(0, \pi)$ (using 10^7 abscissa points).

- ▶ If this computation is done with ordinary double precision arithmetic, the calculation takes 2.59 seconds and yields the result 7.073157029008510.
- ▶ If it is done using all double-double arithmetic (31-digit accuracy), it takes 47.39 seconds and yields the result 7.073157029007832 (correct to 15 digits).
- ▶ But if only the summation is changed to double-double, the result is identical to the double-double result (to 15 digits), yet the computation only takes 3.47 seconds.



Graph of $g(x) = x + \sum_{0 \leq k \leq 10} 2^{-k} \sin(2^k x)$, over $(0, \pi)$.

Aren't 64 bits enough?

There has been considerable resistance in the scientific computing community to the usage of high-precision arithmetic. Why?

- ▶ Many are persuaded that physical reality fundamentally does not require high precision, beyond, say, the limits of 64-bit IEEE arithmetic.
- ▶ Many believe that numerical sensitivity problems are always due to the usage of inferior algorithms.
- ▶ Some regard the usage of high-precision arithmetic as “cheating” or “sinful” (because it is too easy?).
- ▶ Some complain that easy-to-use high-precision arithmetic software facilities are not widely available, or the computational cost is too great.

But none of these excuses are valid.

- ▶ D. H. Bailey, R. Barrio, and J. M. Borwein, “High precision computation: Mathematical physics and dynamics,” *Applied Mathematics and Computation*, vol. 218 (2012), pg. 10106–10121.

Innocuous example where standard 64-bit precision is inadequate

Problem: Find a polynomial to fit the data (1, 1048579, 16777489, 84941299, 268501249, 655751251, 1360635409, 2523398179, 4311748609) for arguments 0, 1, \dots , 8. The usual approach is to solve the linear system:

$$\begin{bmatrix} 1 & \sum_{k=1}^n x_k & \cdots & \sum_{k=1}^n x_k^n \\ \sum_{k=1}^n x_k & \sum_{k=1}^n x_k^2 & \cdots & \sum_{k=1}^n x_k^{n+1} \\ \vdots & \vdots & \ddots & \vdots \\ \sum_{k=1}^n x_k^n & \sum_{k=1}^n x_k^{n+1} & \cdots & \sum_{k=1}^n x_k^{2n} \end{bmatrix} \begin{bmatrix} a_0 \\ a_1 \\ \vdots \\ a_n \end{bmatrix} = \begin{bmatrix} \sum_{k=1}^n y_k \\ \sum_{k=1}^n x_k y_k \\ \vdots \\ \sum_{k=1}^n x_k^n y_k \end{bmatrix}$$

A 64-bit computation (e.g., using Matlab, Linpack or LAPACK) fails to find the correct polynomial in this instance, even if one rounds results to nearest integer.

However, if Linpack routines are converted to use double-double arithmetic (31-digit accuracy), the above computation quickly produces the correct polynomial:

$$f(x) = 1 + 1048577x^4 + x^8 = 1 + (2^{20} + 1)x^4 + x^8$$

Innocuous example, cont.

The result on the previous page can be obtained with double precision using Lagrange interpolation or the Demmel-Koev algorithm. But few scientists, outside of expert numerical analysts, are aware of these schemes.

Besides, even these schemes fail for higher-degree problems.

For example: (1, 134217731, 8589938753, 97845255883, 549772595201, 2097396156251, 6264239146561, 15804422886323, 35253091827713, 71611233653971, 135217729000001, 240913322581691, 409688091758593)
is generated by:

$$f(x) = 1 + 134217729x^6 + x^{12} = 1 + (2^{27} + 1)x^6 + x^{12}$$

Lagrange interpolation and Demmel-Koev fail for this problem, but a straightforward Linpack scheme, implemented with double-double arithmetic, works fine.

Some applications where high precision is essential

1. Planetary orbit calculations (32 digits).
2. Supernova simulations (32–64 digits).
3. Climate modeling (32 digits).
4. Optimization problems in biology and other fields (32 digits).
5. Coulomb n-body atomic system simulations (32–120 digits).
6. Schrodinger solutions for lithium and helium atoms (32 digits).
7. Electromagnetic scattering theory (32–100 digits).
8. Scattering amplitudes of fundamental particles (32 digits).
9. Discrete dynamical systems (32 digits).
10. The Taylor algorithm for ODEs (100–600 digits).
11. Problems in experimental mathematics (100–50,000 digits).


Coulomb n-body atomic system simulations

- ▶ Alexei Frolov of Queen's University in Canada has used high-precision arithmetic to solve a generalized eigenvalue problem that arises in Coulomb n-body interactions.
- ▶ Matrices are typically $5,000 \times 5,000$ and are very nearly singular.
- ▶ Computations typically involve massive cancellation, and high-precision arithmetic must be employed to obtain numerically reproducible results.
- ▶ Frolov has also computed elements of the Hamiltonian matrix and the overlap matrix in four- and five-body systems.
- ▶ These computations typically require 120-digit arithmetic.

Frolov: "We can consider and solve the bound state few-body problems ... beyond our imagination even four years ago."

- ▶ A. M. Frolov and D. H. Bailey, "Highly accurate evaluation of the few-body auxiliary functions and four-body integrals," *Journal of Physics B*, vol. 36, no. 9 (14 May 2003), pg. 1857–1867.

Taylor's method for ODEs with high-precision arithmetic



Numerical integration of the L25-R25 unstable periodic orbit for the Lorenz model during 16 time periods using TIDES with 300 digits, versus 1 time period using DP.

What is needed for high-precision floating-point software?

1. A rock-solid-reliable arithmetic engine, with precision scalable to 1,000,000 digits or more — 1,000 or even 10,000 digits is no longer enough.
2. Highly efficient algorithms, so single-processor performance is as fast as possible, given constraints of portability and ease of installation.
3. A thread-safe design to facilitate parallel processing, both thread-based (low-level) and message-passing (high-level) parallelism.
4. A comprehensive library of tuned transcendentals and special functions: sin, cos, exp, gamma, incomplete gamma, polylogarithms, Bessel functions, etc.
5. High-level language interfaces for C++, Fortran-90 and other languages.
6. Interoperability with *Maple* and *Mathematica*.

Free software for high-precision computation

1. ARPREC: Arbitrary precision, with numerous algebraic and transcendental functions. High-level interfaces for C++ and Fortran-90.
<http://crd.lbl.gov/~dhbailey/mpdist>.
2. GFORTRAN: Now provides full REAL*16 (IEEE 128-bit, or 34-digit) support.
3. GMP: Produced by a volunteer effort and distributed under the GNU license.
<http://gmplib.org>.
4. MPFR: C library for multiple-precision floating-point computations with exact rounding, based on GMP. <http://www.mpfr.org>.
5. MPFR++: High-level C++ interface to MPFR.
<http://perso.ens-lyon.fr/nathalie.revol/software.html>.
6. MPFUN90: Similar to ARPREC, but is written entirely in Fortran-90 and provides only a Fortran-90 interface. <http://crd.lbl.gov/~dhbailey/mpdist>.
7. QD: Performs “double-double” (31 digits) and “quad-double” (62 digits) arithmetic. High-level interfaces for C++ and Fortran-90.
<http://crd.lbl.gov/~dhbailey/mpdist>.

How do multiprecision packages operate?

- ▶ A multiprecision floating-point number is represented as a string of DP floats or long integers, typically 53 or 64 bits per word, with the first few words holding sign, exponent and working precision (see below).
- ▶ Addition and subtraction are performed using relatively straightforward extensions of familiar schemes.
- ▶ Multiplication is performed using FFTs (see below).
- ▶ Division is performed using Newton iterations, with a level of working precision that approximately doubles with each iteration.
- ▶ Square roots are performed using Newton iterations, similar to division.
- ▶ Transcendental functions (cos, sin, exp, log, etc.) are performed using series expansions, after transforming the input into a small argument range. For higher precision levels, quadratically convergent iterations are used.

How to perform high-precision multiplication using FFTs

Let $A = (a_0, a_1, a_2, \dots, a_{n-1})$ and $B = (b_0, b_1, b_2, \dots, b_{n-1})$, where each element stores successive bits, and extend A and B to length $2n$ by padding with zeroes. Let $C = A \times B$. Then except for releasing carries,

$$C_k = \sum_{j=0}^{2n-1} a_j b_{k-j},$$

where we add $2n$ to the b subscript if it is negative. Then

$$\begin{aligned} C &= (C_0, C_1, C_2, \dots, C_{2n-2}) \\ &= (a_0 b_0, a_0 b_1 + a_1 b_0, a_0 b_2 + a_1 b_1 + a_2 b_0, \dots, a_{n-1} b_{n-1}). \end{aligned}$$

This is merely the *convolution* of A and B , which can be computed using FFTs:

$$C = F^{-1}(F(A) \cdot F(B)),$$

where $(F(A) \cdot F(B))$ means the element-wise product of the sequences $F(A)$ and $F(B)$.

For sufficiently large precision levels, FFT-based multiplication is many times faster than conventional grade-school multiplication — $5n \log_2 n$ instead of n^2 operations.

Thread safety: A major challenge for arbitrary precision software

Since run times are greatly magnified by high-precision arithmetic, parallel processing is increasingly essential.

However, almost all existing arbitrary-precision packages are not thread-safe at the application user level, and this kills thread-based parallelization

- ▶ Almost all high-level packages use global variables, and these variables are deadly for thread safety.
- ▶ Most global variables can be dispensed with, but not the working precision level, which is changed frequently within the package itself and often by users also.
- ▶ The only reasonable design for user programming is via operator overloading, available in languages such as Fortran-90 and C++, but overloading does not permit the working precision level to be passed.

What can be done?

MPFUN2015: A thread-safe arbitrary precision package

DHB has written a package (approx. 50,000 lines of code) for arbitrary precision floating-point computation. It is available in two versions:

- ▶ MPFUN-Fort: Completely self-contained, all-Fortran version. Compilation is a simple one-line command, which completes in a few seconds.
- ▶ MPFUN-MPFR: Calls the MPFR package for lower-level operations. Installation is significantly more complicated (since GMP and MPFR must first be installed), but performance is roughly 3X faster. We used MPFUN-MPFR in the Poisson polynomial study.

Both versions include a **high-level language interface**, using Fortran custom datatypes and operator overloading — for most applications, only a few minor changes to existing double precision code are required.

A C++ version is planned.

Full details of design, algorithms, installation and usage are given in

- ▶ D. H. Bailey, “MPFUN2015: A thread-safe arbitrary precision computation package,” <http://www.davidhbailey.com/dhbpapers/mpfun2015.pdf>.

Experimental mathematics: Discovering new mathematical results

Methodology:

1. Compute various mathematical entities (limits, infinite series sums, definite integrals, etc.) to high precision, typically 100–10,000 digits.
2. Use algorithms such as PSLQ to recognize these numerical values in terms of well-known mathematical constants.
3. When results are found experimentally, seek formal mathematical proofs of the discovered relations.

Many results have recently been found using this methodology, both in pure mathematics and in mathematical physics.

*If mathematics describes an objective world just like physics, there is no reason why inductive methods should not be applied in mathematics just the same as in physics.
–D Kurt Godel*

The PSLQ integer relation algorithm

Given a vector (x_n) of real numbers, an integer relation algorithm finds integers (a_n) such that

$$a_1x_1 + a_2x_2 + \cdots + a_nx_n = 0$$

(to within the precision of the arithmetic being used), or else finds bounds within which no relation can exist.

Helaman Ferguson's **PSLQ algorithm** is the most widely used integer relation algorithm, although variants of the LLL algorithm can also be used.

Integer relation detection requires very high numeric precision, both in the input data and in the operation of the algorithm — at least nd digits, where d is the size in digits of the largest $|a_i|$.

1. H. R. P. Ferguson, D. H. Bailey and S. Arno, "Analysis of PSLQ, an integer relation finding algorithm," *Mathematics of Computation*, vol. 68, no. 225 (Jan 1999), 351–369.
2. D. H. Bailey and D. J. Broadhurst, "Parallel integer relation detection: Techniques and applications," *Mathematics of Computation*, vol. 70, no. 236 (Oct 2000), 1719–1736.

Helaman Ferguson's "Umbilic Torus SC" sculpture at Stony Brook Univ.




PSLQ, continued

- ▶ PSLQ constructs a sequence of integer-valued matrices B_n that reduce the vector $y = x \cdot B_n$, until either the relation is found (as one of the columns of matrix B_n), or else precision is exhausted.
- ▶ A relation is detected when the size of smallest entry of the y vector suddenly drops to roughly “epsilon” (i.e. 10^{-p} , where p precision in digits).
- ▶ The size of this drop can be viewed as a “confidence level” that the relation is not a numerical artifact: a drop of 20+ orders of magnitude almost always indicates a real relation.

Efficient variants of PSLQ:

- ▶ 2-level and 3-level PSLQ perform almost all iterations with only double precision, updating full-precision arrays as needed. They are hundreds of times faster than the original PSLQ.
- ▶ Multi-pair PSLQ dramatically reduces the number of iterations required. It was designed for parallel systems, but runs faster even on 1 CPU.

Decrease of $\log_{10}(\min |y_i|)$ in multipair PSLQ run



Application of multipair PSLQ

One simple but important application of multipair PSLQ is to recognize a computed numerical value as the root of an integer polynomial of degree m .

Example: The following constant is suspected to be an algebraic number:

$$\alpha = 1.232688913061443445331472869611255647068988824547930576057634684778 \dots$$

What is its minimal polynomial?

Method: Compute the vector $(1, \alpha, \alpha^2, \dots, \alpha^m)$ for $m = 30$, then input this vector to multipair PSLQ.

Answer (using 250-digit arithmetic):

$$\begin{aligned} 0 = & 697 - 1440\alpha - 20520\alpha^2 - 98280\alpha^3 - 102060\alpha^4 - 1458\alpha^5 + 80\alpha^6 - 43920\alpha^7 \\ & + 538380\alpha^8 - 336420\alpha^9 + 1215\alpha^{10} - 80\alpha^{12} - 56160\alpha^{13} - 135540\alpha^{14} - 540\alpha^{15} \\ & + 40\alpha^{18} - 7380\alpha^{19} + 135\alpha^{20} - 10\alpha^{24} - 18\alpha^{25} + \alpha^{30} \end{aligned}$$

Apéry-like summations

Apéry's proof of the irrationality of $\zeta(3)$ has prompted considerable effort to extend these results to larger integer arguments. The formulas in question are

$$\zeta(2) = 3 \sum_{k=1}^{\infty} \frac{1}{k^2 \binom{2k}{k}}, \quad (1)$$

$$\zeta(3) = \frac{5}{2} \sum_{k=1}^{\infty} \frac{(-1)^{k+1}}{k^3 \binom{2k}{k}}, \quad (2)$$

$$\zeta(4) = \frac{36}{17} \sum_{k=1}^{\infty} \frac{1}{k^4 \binom{2k}{k}}. \quad (3)$$

Is there an analogous formula with a rational coefficient for $\zeta(5)$ or higher? A PSLQ computation with 10,000 digits ruled this out. What about multi-term relations?

Some multi-term relations found by PSLQ

$$\zeta(5) = 2 \sum_{k=1}^{\infty} \frac{(-1)^{k+1}}{k^5 \binom{2k}{k}} - \frac{5}{2} \sum_{k=1}^{\infty} \frac{(-1)^{k+1}}{k^3 \binom{2k}{k}} \sum_{j=1}^{k-1} \frac{1}{j^2},$$

$$\zeta(7) = \frac{5}{2} \sum_{k=1}^{\infty} \frac{(-1)^{k+1}}{k^7 \binom{2k}{k}} + \frac{25}{2} \sum_{k=1}^{\infty} \frac{(-1)^{k+1}}{k^3 \binom{2k}{k}} \sum_{j=1}^{k-1} \frac{1}{j^4}$$

$$\begin{aligned} \zeta(9) = & \frac{9}{4} \sum_{k=1}^{\infty} \frac{(-1)^{k+1}}{k^9 \binom{2k}{k}} - \frac{5}{4} \sum_{k=1}^{\infty} \frac{(-1)^{k+1}}{k^7 \binom{2k}{k}} \sum_{j=1}^{k-1} \frac{1}{j^2} + 5 \sum_{k=1}^{\infty} \frac{(-1)^{k+1}}{k^5 \binom{2k}{k}} \sum_{j=1}^{k-1} \frac{1}{j^4} \\ & + \frac{45}{4} \sum_{k=1}^{\infty} \frac{(-1)^{k+1}}{k^3 \binom{2k}{k}} \sum_{j=1}^{k-1} \frac{1}{j^6} - \frac{25}{4} \sum_{k=1}^{\infty} \frac{(-1)^{k+1}}{k^3 \binom{2k}{k}} \sum_{j=1}^{k-1} \frac{1}{j^4} \sum_{j=1}^{k-1} \frac{1}{j^2}, \end{aligned}$$

$$\begin{aligned} \zeta(11) = & \frac{5}{2} \sum_{k=1}^{\infty} \frac{(-1)^{k+1}}{k^{11} \binom{2k}{k}} + \frac{25}{2} \sum_{k=1}^{\infty} \frac{(-1)^{k+1}}{k^7 \binom{2k}{k}} \sum_{j=1}^{k-1} \frac{1}{j^4} \\ & - \frac{75}{4} \sum_{k=1}^{\infty} \frac{(-1)^{k+1}}{k^3 \binom{2k}{k}} \sum_{j=1}^{k-1} \frac{1}{j^8} + \frac{125}{4} \sum_{k=1}^{\infty} \frac{(-1)^{k+1}}{k^3 \binom{2k}{k}} \sum_{j=1}^{k-1} \frac{1}{j^4} \sum_{i=1}^{k-1} \frac{1}{i^4}. \end{aligned}$$

General formulas

Using bootstrapping and an application of the “Pade” function (which in both *Mathematica* and *Maple* produces Padé approximations to a rational function satisfied by a truncated power series), the following general results were discovered:

$$\sum_{k=1}^{\infty} \frac{1}{k^3(1 - x^4/k^4)} = \frac{5}{2} \sum_{k=1}^{\infty} \frac{(-1)^{k+1}}{k^3 \binom{2k}{k} (1 - x^4/k^4)} \prod_{m=1}^{k-1} \left(\frac{1 + 4x^4/m^4}{1 - x^4/m^4} \right)$$
$$\sum_{k=1}^{\infty} \frac{1}{k^2 - x^2} = 3 \sum_{k=1}^{\infty} \frac{1}{k^2 \binom{2k}{k} (1 - x^2/k^2)} \prod_{m=1}^{k-1} \left(\frac{1 - 4x^2/m^2}{1 - x^2/m^2} \right)$$

Power series expansions for a specific x yield all the results for $\zeta(n)$.

The second formula was proven by the Wilf-Zeilberger method, a computer-based scheme that yields a rigorous proof. **It was thus discovered and proven by computer.**

- D. H. Bailey and J. M. Borwein, “Computer-assisted discovery and proof,” T. Amdeberhan and V. H. Moll, ed., *Tapas in Experimental Mathematics*, AMS, May 2008, pg. 21–52.

The tanh-sinh numerical integration algorithm

Given $f(x)$ defined on $(-1, 1)$, define $g(t) = \tanh(\pi/2 \cdot \sinh t)$. Then setting $x = g(t)$ yields

$$\int_{-1}^1 f(x) dx = \int_{-\infty}^{\infty} f(g(t))g'(t) dt \approx h \sum_{j=-N}^N w_j f(x_j),$$

where $x_j = g(h_j)$ and $w_j = g'(h_j)$ can be precomputed for a given h . Since $g'(t)$ goes to zero very rapidly for large t , the integrand $f(g(t))g'(t)$ typically is a nice bell-shaped function, so that the simple summation above converges very rapidly. Reducing h by half typically doubles the number of correct digits.

We have found that tanh-sinh is the best general-purpose integration scheme for functions with vertical derivatives or singularities at endpoints, or for any function at very high precision (> 1000 digits).

1. D. H. Bailey, X. S. Li and K. Jeyabalan, "A comparison of three high-precision quadrature schemes," *Experimental Mathematics*, vol. 14 (2005), no. 3, pg. 317–329.
2. H. Takahasi and M. Mori, "Double exponential formulas for numerical integration," *Publications of RIMS*, Kyoto University, vol. 9 (1974), pg. 721–741.

Ising integrals from mathematical physics

We applied our methods to study three classes of integrals: C_n are connected to quantum field theory, D_n arise in the Ising theory of mathematical physics, while the E_n integrands are derived from D_n :

$$C_n := \frac{4}{n!} \int_0^\infty \cdots \int_0^\infty \frac{1}{\left(\sum_{j=1}^n (u_j + 1/u_j)\right)^2} \frac{du_1}{u_1} \cdots \frac{du_n}{u_n}$$

$$D_n := \frac{4}{n!} \int_0^\infty \cdots \int_0^\infty \frac{\prod_{i < j} \left(\frac{u_i - u_j}{u_i + u_j}\right)^2}{\left(\sum_{j=1}^n (u_j + 1/u_j)\right)^2} \frac{du_1}{u_1} \cdots \frac{du_n}{u_n}$$

$$E_n = 2 \int_0^1 \cdots \int_0^1 \left(\prod_{1 \leq j < k \leq n} \frac{u_k - u_j}{u_k + u_j} \right)^2 dt_2 dt_3 \cdots dt_n$$

where in the last line $u_k = t_1 t_2 \cdots t_k$.

D. H. Bailey, J. M. Borwein and R. E. Crandall, "Integrals of the Ising class," *Journal of Physics A: Mathematical and General*, vol. 39 (2006), pg. 12271–12302.

Limiting value of C_n : What is this number?

Key observation: The C_n integrals can be converted to one-dimensional integrals involving the modified Bessel function $K_0(t)$:

$$C_n = \frac{2^n}{n!} \int_0^\infty t K_0^n(t) dt$$

High-precision numerical values, computed using this formula and tanh-sinh quadrature, approach a limit. For example:

$$C_{1024} = 0.6304735033743867961220401927108789043545870787 \dots$$

What is this number? We copied the first 50 digits into the online Inverse Symbolic Calculator (ISC) at <https://isc.carma.newcastle.edu.au>. The result was:

$$\lim_{n \rightarrow \infty} C_n = 2e^{-2\gamma}.$$

where γ denotes Euler's constant. This is now proven.

Other Ising integral evaluations found using PSLQ

$$D_3 = 8 + 4\pi^2/3 - 27L_{-3}(2)$$

$$D_4 = 4\pi^2/9 - 1/6 - 7\zeta(3)/2$$

$$E_2 = 6 - 8\log 2$$

$$E_3 = 10 - 2\pi^2 - 8\log 2 + 32\log^2 2$$

$$E_4 = 22 - 82\zeta(3) - 24\log 2 + 176\log^2 2 - 256(\log^3 2)/3 \\ + 16\pi^2\log 2 - 22\pi^2/3$$

$$E_5 = 42 - 1984\text{Li}_4(1/2) + 189\pi^4/10 - 74\zeta(3) - 1272\zeta(3)\log 2 + 40\pi^2\log^2 2 \\ - 62\pi^2/3 + 40(\pi^2\log 2)/3 + 88\log^4 2 + 464\log^2 2 - 40\log 2$$

where ζ is the Riemann zeta function and $Li_n(x)$ is the polylog function. E_5 remained a “numerical conjecture” for several years, but was proven in March 2014 by Erik Panzer.

E_5 was reduced to a 3-D integral of a 60-line integrand, which was evaluated using tanh-sinh quadrature to 250-digit arithmetic using over 1000 CPU-hours on a highly parallel system. The PSLQ calculation required only seconds.

Algebraic numbers in Poisson potential functions

Lattice sums arising from the Poisson equation have been studied widely in mathematical physics and image processing. We numerically discovered, and then proved, that for rational (x, y) , the 2-D Poisson potential function satisfies

$$\phi_2(x, y) = \frac{1}{\pi^2} \sum_{m, n \text{ odd}} \frac{\cos(m\pi x) \cos(n\pi y)}{m^2 + n^2} = \frac{1}{\pi} \log \alpha$$

where α is *algebraic*, i.e., the root of an integer polynomial

$$0 = a_0 + a_1\alpha + a_2\alpha^2 + \cdots + a_n\alpha^n$$

The minimal polynomials for these α were found by PSLQ calculations, with the $(n+1)$ -long vector $(1, \alpha, \alpha^2, \cdots, \alpha^n)$ as input, where $\alpha = \exp(\pi\phi_2(x, y))$. PSLQ returns the vector of integer coefficients $(a_0, a_1, a_2, \cdots, a_n)$ as output.

- D. H. Bailey, J. M. Borwein, R. E. Crandall and J. Zucker, "Lattice sums arising from the Poisson equation," *Journal of Physics A: Mathematical and Theoretical*, vol. 46 (2013), pg. 115201, <http://www.davidhbailey.com/dhbpapers/PoissonLattice.pdf>.

Samples of minimal polynomials found by PSLQ

k	Minimal polynomial for $\exp(8\pi\phi_2(1/k, 1/k))$
5	$1 + 52\alpha - 26\alpha^2 - 12\alpha^3 + \alpha^4$
6	$1 - 28\alpha + 6\alpha^2 - 28\alpha^3 + \alpha^4$
7	$-1 - 196\alpha + 1302\alpha^2 - 14756\alpha^3 + 15673\alpha^4 + 42168\alpha^5 - 111916\alpha^6 + 82264\alpha^7$ $- 35231\alpha^8 + 19852\alpha^9 - 2954\alpha^{10} - 308\alpha^{11} + 7\alpha^{12}$
8	$1 - 88\alpha + 92\alpha^2 - 872\alpha^3 + 1990\alpha^4 - 872\alpha^5 + 92\alpha^6 - 88\alpha^7 + \alpha^8$
9	$-1 - 534\alpha + 10923\alpha^2 - 342864\alpha^3 + 2304684\alpha^4 - 7820712\alpha^5 + 13729068\alpha^6$ $- 22321584\alpha^7 + 39775986\alpha^8 - 44431044\alpha^9 + 19899882\alpha^{10} + 3546576\alpha^{11}$ $- 8458020\alpha^{12} + 4009176\alpha^{13} - 273348\alpha^{14} + 121392\alpha^{15}$ $- 11385\alpha^{16} - 342\alpha^{17} + 3\alpha^{18}$
10	$1 - 216\alpha + 860\alpha^2 - 744\alpha^3 + 454\alpha^4 - 744\alpha^5 + 860\alpha^6 - 216\alpha^7 + \alpha^8$

The minimal polynomial for $\exp(8\pi\phi_2(1/32, 1/32))$ has degree 128, with individual coefficients ranging from 1 to over 10^{56} (see next viewgraph). This computation required 10,000-digit precision. Most of the time was spent on PSLQ, not quadrature.

Other polynomials required up to 50,000-digit precision, with runs extending to several days. The ARPREC software failed on at least one of these problems.

192-degree minimal polynomial found by multipair PSLQ for $x = y = 1/35$

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269
270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
84

Summary

- ▶ High-precision arithmetic can now be performed rapidly by means of widely available, easy-to-use software.
- ▶ High-precision arithmetic is required by the PSLQ algorithm, which finds integer relations among a set of computed floating-point real numbers.
- ▶ This combined capability (HP arithmetic and PSLQ) is a tool of enormous power for experimental mathematics.

Thanks!

This talk is available at

<http://www.davidhbailey.com/dhbtalks/dhb-hiprec-2017.pdf>