

Finding large Poisson polynomials using four-level variable precision

David H. Bailey

Lawrence Berkeley National Laboratory (retired)

University of California, Davis, Department of Computer Science

This talk is available at:

<http://www.davidhbailey.com/dhbtalks/dhb-sc21-correct.pdf>



ICERM report on reproducibility in mathematical and scientific computing

The field of high-performance computing has lagged efforts in other fields to ensure reproducibility. One key issue is numerical reproducibility:

Numerical reproducibility has emerged as a particularly important issue, since the scale of computations has greatly increased in recent years, particularly with computations performed on many thousands of processors and involving similarly large datasets. Large computations often greatly magnify the level of numeric error, so that numerical difficulties that were once of little import now are large enough to alter the course of the computation or to draw into question the overall validity of the results.

-
- ▶ D. H. Bailey, J. M. Borwein, O. Caprotti, U. Martin, B. Salvy and M. Taufer, "Opportunities and challenges in 21st century mathematical computation: ICERM workshop report," 31 Jul 2014, <https://www.davidhbailey.com/dhbpapers/ICERM-2014.pdf>.

Commonly used formats for floating-point computing

Name	Number of bits				Digits
	Sign	Exp.	Mant.	Hidden	
IEEE half	1	5	10	1	3
ARM half	1	5	10	1	3
bfloat16	1	8	7	1	2
IEEE single	1	8	23	1	7
IEEE double	1	11	52	1	15
IEEE extended	1	15	64	0	19
IEEE quad	1	15	112	1	34
double-double	1	11	104	2	31
quad-double	1	11	208	4	62
double-quad	1	15	224	1	68
arbitrary	varies	varies	varies	varies	varies

Advantages of variable precision

Compared with using a high fixed level of precision for the entire application, employing variable precision *usually* results in:

- ▶ Faster processing.
- ▶ Better cache utilization.
- ▶ Lower run-time memory usage.
- ▶ Lower offline data storage.
- ▶ Lower energy costs.

Challenge:

- ▶ Unless variable precision is carefully managed, there is significantly greater potential for loss of numerical accuracy and reproducibility.

Finding Poisson polynomials: Application of variable precision computing

Lattice sums related to the Poisson potential function, which naturally arise in studies of gravitational and electrostatic potentials, have been studied for many years in the mathematical physics community, and recently have been applied to image sharpening:

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

In 2012 researchers discovered that when x and y are **rational numbers**,

$$\phi_2(x, y) = \frac{1}{\pi} \log \alpha(x, y),$$

where $\alpha(x, y)$ is an **algebraic number**, namely the root of a degree- m integer polynomial. **What is the connection between x and y and the degree of α ?**

-
- ▶ 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), 115201.
 - ▶ D. H. Bailey, J. M. Borwein, J. Kimberley and W. Ladd, "Computer discovery and analysis of large Poisson polynomials," *Experimental Mathematics*, 27 Aug 2016, vol. 26, 349–363.

Finding Poisson polynomials, cont.

By computing high-precision numerical values of $\phi_2(x, y)$ for various specific rational x and y , using a three-level variable precision code (e.g., double/500-digit/10,000-digit), and applying a variant of the PSLQ algorithm, we found the following results:

s	Minimal polynomial of α corresponding to $x = y = 1/s$
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	$1 - 216\alpha + 860\alpha^2 - 744\alpha^3 + 454\alpha^4 - 744\alpha^5 + 860\alpha^6 - 216\alpha^7 + \alpha^8$

Kimberley's formula

Based on these preliminary results, Jason Kimberley conjectured that the degree $m(s)$ of the minimal polynomial associated with the case $x = y = 1/s$ is given by this rule:

Set $m(2) = 1/2$. Otherwise for primes p congruent to 1 mod 4, set $m(p) = \text{int}^2(p/2)$, where int denotes greatest integer, and for primes p congruent to 3 mod 4, set $m(p) = \text{int}(p/2)(\text{int}(p/2) + 1)$. Then for any other positive integer s whose prime factorization is $s = p_1^{e_1} p_2^{e_2} \cdots p_r^{e_r}$,

$$m(s) = 4^{r-1} \prod_{i=1}^r p_i^{2(e_i-1)} m(p_i).$$

Additional three-level precision computations, (e.g., double/2000-digit/64,000-digit), confirmed Kimberley's formula for all s up to 40, and for even s up to 52.

By examining the computed results, connections were found to a sequence of polynomials defined in a 2010 paper by Savin and Quarfoot. These observations ultimately led to a proof of Kimberley's formula and several other facts.

What about the more general case $x = p/s$ and $y = q/s$, for integers p, q, s ?

Thread safety in arbitrary precision libraries

Most existing arbitrary precision libraries today are not thread-safe. Thus applications written to use them do not work in a multi-thread shared-memory parallel environment.

Currently the GNU MPFR package is the only arbitrary precision library that is thread-safe. It is also very fast (esp. for transcendentals), and correctly rounds results.

However:

- ▶ MPFR is only a low-level library; it does not include high-level language bindings required to support full-scale application codes.
- ▶ MPFR requires a lengthy installation procedure with administrator permission (to install MPFR and GMP on the user's computer), which discourages many users.
- ▶ If a thread-safe build option is not properly invoked during installation, the resulting software is not thread-safe.

-
- ▶ L. Fousse, G. Hanrot, V. Lefevre, P. Pelissier and P. Zimmermann, "MPFR: A multiple precision binary floating-point library with correct rounding," *ACM Transactions on Mathematical Software*, vol. 33 (2007), <https://doi.org/10.1145/1236463.1236468>.

DHB's new MPFUN2020 arbitrary precision library

- ▶ An all-Fortran design that can be compiled in a few seconds with any Fortran-2008 compliant compiler, including the gfortran, Intel and NAG Fortran compilers, on a variety of systems, including Mac OS X and various Linux platforms.
- ▶ A 100% thread-safe design.
- ▶ A full-featured high-level Fortran language interface, so that only type statements and a few other minor changes are required to convert most double precision codes.
- ▶ Support for common transcendental functions (sin, cos, exp, etc.) and special functions.
- ▶ Support for arbitrary precision real and complex datatypes, plus double and quad.
- ▶ Support for both a full and a medium precision datatype, which in some applications (such as Poisson problems) results in lower memory usage and memory traffic.
- ▶ FFT-based multiplication for faster performance at very high precision.
- ▶ A comprehensive test suite, together with six full-scale application test codes.

A separate version calls MPFR for all lower-level operations. It is roughly 15% faster on most applications (several times faster for transcendental functions).

DHB's new four-level multipair PSLQ code for finding integer relations

Given an $(m + 1)$ -long input vector $X = (x_i)$ of high-precision floating-point values, an **integer relation algorithm** finds an $(m + 1)$ -long vector of integers (a_i) such that

$$a_0x_0 + a_1x_1 + a_2x_2 + \cdots + a_mx_m = 0,$$

to within the tolerance of the numeric precision being used. As a simple illustration, suppose one suspects that the real constant

$\alpha = 2.11959126982917513132984833493468711062807145783249543921443000325120624764609847 \dots$


is an algebraic number of degree eight. After computing the vector $(1, \alpha, \alpha^2, \dots, \alpha^8)$ and applying the multipair PSLQ integer relation algorithm, the relation $(1, -216, 860, -744, 454, -744, 860, -216, 1)$ is produced, so that α satisfies $1 - 216\alpha + 860\alpha^2 - 744\alpha^3 + 454\alpha^4 - 744\alpha^5 + 860\alpha^6 - 216\alpha^7 + \alpha^8 = 0$.

DHB's new multipair PSLQ code employs four levels of precision: (a) IEEE double; (b) IEEE quad; (c) medium precision, typically 100-1000 digits; and (d) full precision, typically 5,000 to 50,000 digits.

How can one tell if a computed polynomial is real and not an artifact?

Iterations of the multipair PSLQ algorithm develop a sequence of invertible integer matrices A_n and inverses B_n so that $y = B_n \cdot X$ has steadily smaller entries, until one entry of y is smaller than the epsilon specified for detection.

The size of the drop in $\min(|y_i|)$ when the relation is detected can be viewed as a confidence level that the discovered relation is a real mathematical relation.



High-level computational algorithm

1. Given rationals $x = p/s$ and $y = q/s$ (typically satisfying $1 \leq p, q < s/2 \leq 50$), select a conjectured minimal polynomial degree m , a medium precision level P_1 digits, a full precision level P_2 digits and other parameters for the run.
2. Calculate $\phi_2(x, y)$ to P_2 -digit precision using the following formula:

$$\phi_2(x, y) = \frac{1}{2\pi} \log \left| \frac{\theta_2(z, q)\theta_4(z, q)}{\theta_1(z, q)\theta_3(z, q)} \right|,$$

where $q = e^{-\pi}$ and $z = \frac{\pi}{2}(y + ix)$. Compute the four theta functions using rapidly convergent formulas. Example: $\theta_1(z, q) = 2 \sum_{k=1}^{\infty} (-1)^{k-1} q^{(2k-1)^2/4} \sin((2k-1)z)$.

3. Calculate $\alpha = \exp(8\pi\phi_2(x, y))$ and generate the $(m+1)$ -long vector $X = (1, \alpha, \alpha^2, \dots, \alpha^m)$ to P_2 -digit precision.
4. Apply the four-level multipair PSLQ algorithm to find an integer relation for X , if one exists, to the precision being used (P_2 digits).
5. If a numerically significant relation is not found, try again with a larger degree m or a higher full precision level P_2 . If a relation is found, employ the polynomial factorization facilities in *Mathematica* or *Maple* to ensure that the resulting polynomial is irreducible.

Performance: New vs old codes

Arbitrary precision package	Multipair PSLQ	Run time (CPU seconds)	
		$x = y = 1/29$	$x = y = 1/35$
Old all-Fortran	3-level	247819.68	220202.68
New all-Fortran	3-level	55691.66	49221.37
	4-level	52920.67	47093.43
New MPFR-based	3-level	45038.87	41442.57
	4-level	45411.23	39762.23

The two application cases displayed are:

- ▶ $x = y = 1/29$: requires 19,700-digit arithmetic (full precision); produces a polynomial of degree 196 with coefficients ranging from 1 up to roughly 10^{87} .
- ▶ $x = y = 1/35$: requires 18,900-digit arithmetic (full precision), producing a polynomial of degree 192 with coefficients ranging from 1 up to roughly 10^{85} .

Results: Degrees of minimal polynomials for $x = 1/s; y = q/s$

s	q: first row 1–10; second row, 11–20									
	1	2	3	4	5	6	7	8	9	10
	11	12	13	14	15	16	17	18	19	20
10	8	16	4	16						
11	30	30	30	30	30					
12	16	32	16	32	8					
13	36	36	36	36	18	36				
14	24	48	24	48	24	48				
15	32	32	32	16	32	32	32			
16	32	64	32	64	32	64	16			
17	64	64	64	32	64	64	64	64		
18	36	72	36	72	36	72	36	72		
19	90	90	90	90	90	90	90	90	90	
20	32	64	32	64	32	64	32	64	16	
21	96	96	96	96	96	96	96	48	96	96
22	60	120	60	120	60	120	60	120	60	120
23	132	132	132	132	132	132	132	132	132	132
	132									

Results: Degrees of minimal polynomials for $x = 1/s; y = q/s$

s	q: first row 1–10; second row, 11–20									
	1	2	3	4	5	6	7	8	9	10
	11	12	13	14	15	16	17	18	19	20
24	64	128	64	128	32	128	32	128	64	128
	32									
25	100	100	100	100	100	100	50	100	100	100
	100	100								
26	72	144	72	144	36	144	72	144	72	144
	72	144								
27	162	162	162	162	162	162	162	162	162	162
	162	162	162							
28	96	192	96	192	96	192	96	192	96	192
	96	192	48							
29	196	196	196	196	196	196	196	196	196	196
	196	98	196	196						
30	64	128	64	128	64	128	64	128	64	128
	32	128	64	128						
31	240	240	240	240	240	240	240	240	240	240
	240	240	240	240	240					

Results: Degrees of minimal polynomials for $x = 1/s; y = q/s$

s	q: first row 1–10; second row, 11–20									
	1 11	2 12	3 13	4 14	5 15	6 16	7 17	8 18	9 19	10 20
32	128	256	128	256	128	256	128	256	128	256
	128	256	128	256	64					
33	240	240	240	240	240	240	240	240	240	120
	240	240	240	240	240	240				
34	128	256	128	256	128	256	128	256	128	256
	128	256	64	256	128	256				
35	192	192	192	192	192	96	192	192	192	192
	192	192	192	192	192	192	192			
36	144	288	144	288	144	288	144	288	144	288
	144	288	144	288	144	288	72			
38	180	360	180	360	180	360	180	360	180	360
	180	360	180	360	180	360	180	360		
39	288	288	288	288	288	288	288	288	288	288
	288	288	288	144	288	288	288	288	288	
40	128	256	128	256	128	256	128	256	64	256
	64	256	128	256	128	256	128	256	64	

Research findings

Note that many of these results are **not consistent with Kimberley's formula**, which has been proven only for the cases $p = q = 1$ (i.e., $x = 1/s$; $y = 1/s$).

For example, note that when s is even, the degrees for even q are, in most cases, double the degrees for $q = 1$. But there are exceptions, when the degrees are only half as large. For instance when $s = 36$, degrees alternate between 144 and 288, yet when $q = 17$, the degree is 72.

There are also anomalies when s is odd. Note, for instance, that when $s = 35$, all degrees are 192, as given by Kimberley's formula, except the degree is 96 when $q = 6$.

Questions:

- ▶ Is there a generalization of Kimberley's formula for the much larger class $x = p/s, y = q/s, 1 \leq p, q < s/2$?
- ▶ How can we greatly expand the scope of these computations? A new integer relation algorithm?
- ▶ How can massively parallel computing be employed? Parallel implementations to date have yielded only modest speedups, compared with the best serial code.

Degree-100 minimal polynomial found for the case $x = y = 1/25$

```
-1  
-23300 x99  
-802070 x98  
-184237950 x97  
-2843477882 x96  
-280829540280 x95  
-62078135033800 x94  
-587223020214800 x93  
-13074431882111320 x92  
-13218486647388660 x91  
-10476466179461781320 x90  
-18641886747647170 x89  
-279488112471778262700 x88  
+813227944832288688000 x87  
-265415389839664333042000 x86  
-118670358728306421266400 x85  
+68881704042618941894106750 x84  
-73389054531176393694825000 x83  
-11734149574623881618073586000 x82  
-15783217362006003543768467641200 x81  
-11212160635965016386734323390 x80  
-8347129366578242288028860727280 x79  
-381218020102166607396166441748000 x78  
-1046227881170347271711813873078000 x77  
-740230695463646773187951381290500 x76  
-28955791532870578889863202133660 x75  
-175888888877520039459983521290200 x74  
-7866213228713888525131298178899600 x73  
-236884125690282976117756734267500 x72  
-58417611624222049212075730324848000 x71  
-1200419125962194317653170767089086480 x70  
-232791256796343654268686240874186000 x69  
-8868485064222739328899847883612375 x68  
-240349617176533227181569447208850051400 x67  
-1184632278986603841785779602545834390 x66  
+3843883231688887781178673789634618372 x65  
-10497711193112829126375484738640829758 x64  
-27423966616159757372945688868461827200 x63  
-37204887865477934685066274824741346000 x62  
-88772568818594615813670651262719480 x61  
+8482386793661736781468297960178866234824 x60  
-1996842817974852842754118517495517359480 x59  
+4713832070965411181279830319583488477200 x58  
-96652918643813977392482148484384208009820 x57  
-11844237501278021715651791587703858818240 x56  
-22707906681108421171257881748014670448216 x55  
-2401383861858881380777617941155846627860 x54  
-248781858450966360778681721338200518608 x53  
-252434946127646018463885161024436794136500 x52  
-18275358593186231187856652844410780800 x51  
-1172273954815846124838830548585836349876 x50  
-472081714818638183948888494888822187059800 x49  
+6492828201703835488184455889626738360 x48  
-3649496211871844632794315157048578840 x47  
-8274877652171386213074294888882020 x46  
-5285893921182224681127801896135681200380 x45  
-6548962388400803136703871413832420560 x44  
-3318961099725740968106282021046183420000 x43  
-201847778373833088146838997970027260 x42  
+18278246422268818742066113717647968480 x41  
-47166728041246875435483727936298849390 x40  
-1729561946183529787941988646887882660 x39  
-84012948040327232688266872976765101000 x38  
-67220353787394748818887947378858500 x37  
-46158454681204318878344662712763866055 x36  
-2685815737289481262187719396621808 x35  
-127258117894164862813622471272380730 x34  
-5242052688513778740796564621340542900 x33  
-18432168386824701291888738544754375 x32  
-6715970412251887480949730364829867320 x31  
-218843842813788657177412773898880 x30  
+1728737944609615071502185452252738800 x29  
-14142458894463697878747224894767800 x28  
-25729450410886117118818787879600 x27  
-326427861364762688172115843685000 x26  
+47328949188836184848319481178000 x25  
-17888427828487906868130791577000 x24  
+6628976489446464647964176288000 x23  
-157614048794612281346400860400 x22  
-24809138583135266367788710380 x21  
-289382573607655331818842080 x20  
-2793461624813078167577871860 x19  
-18654845738287314881818500 x18  
-19748506642088496687124000 x17  
+608828420867200209484330 x16  
-18868382286447766261280 x15  
-171431748938318284478400 x14  
-13461774808181562158500 x13  
+5713652131848124180900 x12  
-2811832788828259620 x11  
-1874817886811820436 x10  
+1849636549787120 x9  
-53845403571830 x8  
-6175821027600 x7  
-882098133760 x6  
-123880399624 x5  
-3364383685 x4  
-88947310 x3  
+138820 x2  
+2 x
```

Suggestions for future research in variable precision computing

1. Seamlessly incorporating an arbitrary precision facility, say based on the MPFR package, in future Fortran and C language standards.
2. Checking all available software supporting variable and arbitrary precision for thread safety, and fixing or deprecating libraries that are not strictly thread-safe.
3. Rethinking automatic type conversions and mixed-mode operations, which hinder the development of correct and reproducible variable precision computing, in DHB's opinion. At the least, could some of these features be optional?
4. Investigating novel numerical algorithms, error control and error analysis techniques, appropriate not only for single and double precision on a single processor system, but also for the realm of very high precision and very highly parallel computing.
5. Developing fundamentally new paradigms for variable precision computing, perhaps along the lines of Gustafson's "unum" framework.

This talk is available at:

<http://www.davidhbailey.com/dhbtalks/dhb-sc21-correct.pdf>