

MPFUN2020: A new thread-safe arbitrary precision package (Full Documentation)

David H. Bailey *

February 16, 2021

Abstract

Numerous research studies have arisen, particularly in mathematical physics and experimental mathematics, that require extremely high numeric precision. Such precision greatly magnifies computer run times, so software packages to support high-precision computing must be designed for thread-based parallel processing.

This paper describes a new arbitrary precision software package (“MPFUN2020”) that features several significant improvements over an earlier package. It comes in two versions: a self-contained all-Fortran version, and a version based on the MPFR package, which is even faster. Both of these versions feature: (a) a completely thread-safe design, so user codes can be converted for parallel execution at the application level; (b) a full-featured high-level Fortran interface, so that most applications can be converted to multiprecision with relatively minor changes to source code; (c) full support for both real and complex datatypes; (d) run-time checking and other facilities to overcome nagging problems with converting double precision constants and data; (e) a medium precision datatype, which improves performance and reduces memory cost on large mixed-precision applications; and (f) interoperability — with certain simple restrictions, application codes written for one version can also be run with the other.

*Lawrence Berkeley National Laboratory (retired), 1 Cyclotron Road, Berkeley, CA 94720, USA, and University of California, Davis, Department of Computer Science. E-mail: dhbailey@lbl.gov.

Contents

1	Applications of high-precision computation	2
1.1	The PSLQ integer relation algorithm	2
1.2	High-precision numerical integration	3
1.3	Ising integrals	4
1.4	Algebraic numbers in Poisson potential functions	5
2	Current high-precision software	7
2.1	Thread safety and parallel implementations	8
2.2	MPFUN2015	9
3	MPFUN2020: A new thread-safe package	10
3.1	Data structure	11
3.2	Modules	12
3.3	The MPFUN2020 solution to thread safety	12
4	Installation, compilation and linking	13
5	Coding Fortran applications	14
5.1	Functions and subroutines	17
5.2	Input and output of multiprecision data	17
5.3	Handling double precision values	21
5.4	Dynamically changing the working precision	23
5.5	Medium precision datatype	25
6	Sample applications and performance	27
6.1	Timings	28
7	Appendix: Numerical algorithms	30
7.1	Algorithms for basic arithmetic	30
7.2	Basic algorithms for transcendental functions	31
7.3	Special functions	33
7.4	FFT-based multiplication	36
7.5	Advanced algorithm for division	37

1 Applications of high-precision computation

For many scientific calculations, particularly those that employ empirical data, IEEE 32-bit floating-point arithmetic is sufficiently accurate, and is preferred since it saves memory, run time and energy usage. For other calculations, 64-bit floating-point arithmetic is required to produce results of sufficient accuracy; still others switch between 32-bit and 64-bit. Software tools are being developed to help users determine which portions of a computation can be performed with lower precision and which must be performed with higher precision [28].

Other applications, particularly in the fields of mathematical physics and experimental mathematics, require even higher precision — tens, hundreds or even thousands of digits. Here is a brief summary of these applications:

1. Supernova simulations (32–64 digits).
2. Optimization problems in biology and other fields (32–64 digits).
3. Coulomb n -body atomic system simulations (32–120 digits).
4. Electromagnetic scattering theory (32–100 digits).
5. The Taylor algorithm for ODEs (100–600 digits).
6. Ising integrals from mathematical physics (100–1000 digits).
7. Problems in experimental mathematics (100–50,000 digits and higher).

These applications are described in greater detail in [4, 3], which provides detailed references. Here is a brief overview of a handful of these applications:

1.1 The PSLQ integer relation algorithm

Very high-precision floating-point arithmetic is now considered an indispensable tool in experimental mathematics and mathematical physics [4]. Many of these computations involve variants of Ferguson’s PSLQ integer relation detection algorithm [19, 10]. Suppose one is given an n -long vector (x_i) of real or complex numbers (presented as a vector of high-precision values). The PSLQ algorithm finds the integer coefficients (a_i) , not all zero, such that

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

(to available precision), or else determines that there is no such relation within a certain bound on the size of the coefficients. Alternatively, one can employ the Lenstra-Lenstra-Lovasz (LLL) lattice basis reduction algorithm to find integer relations [24], or the “HJLS” algorithm, which is based on LLL. Both PSLQ and HJLS can be viewed as schemes to compute the intersection between a lattice and a vector subspace [16]. Whichever algorithm is used, integer relation detection requires very high precision—at least $(n \times d)$ -digit precision, where d is the size in digits of the largest a_i and n is the vector length, or else the true relation will be unrecoverable.

1.2 High-precision numerical integration

One of the most fruitful applications of the experimental methodology and the PSLQ integer relation algorithm has been to identify classes of definite integrals, based on very high-precision numerical values, in terms of simple analytic expressions.

These studies typically employ either Gaussian quadrature or the “tanh-sinh” quadrature scheme of Takahasi and Mori [30, 2]. The tanh-sinh quadrature algorithm approximates the integral of a function on $(-1, 1)$ as

$$\int_{-1}^1 f(x) dx \approx h \sum_{j=-N}^N w_j f(x_j), \quad (1)$$

where the abscissas x_j and weights w_j are given by

$$\begin{aligned} x_j &= \tanh(\pi/2 \cdot \sinh(hj)) \\ w_j &= \pi/2 \cdot \cosh(hj) / \cosh(\pi/2 \cdot \sinh(hj))^2, \end{aligned} \quad (2)$$

and where N is chosen large enough that summation terms in (1) beyond N (positive or negative) are smaller than the “epsilon” of the numeric precision being used. Full details are given in [2]. An overview of applications of high-precision integration in experimental mathematics is given in [5].

1.3 Ising integrals

In one study, tanh-sinh quadrature and PSLQ were employed to study the following classes of integrals [8]. The C_n are connected to quantum field

theory, the D_n integrals arise in the Ising theory of mathematical physics, while the E_n integrands are derived from D_n :

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

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

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

In the last line $u_k = \prod_{i=1}^k t_i$.

In general, it is very difficult to compute high-precision numerical values of n -dimensional integrals such as these. But as it turns out, the C_n integrals can be converted to one-dimensional integrals, which are amenable to evaluation with the tanh-sinh scheme:

$$C_n = \frac{2^n}{n!} \int_0^\infty p K_0^n(p) dp.$$

Here K_0 is the *modified Bessel function* [26]. 1000-digit values of these sufficed to identify the first few instances of C_n in terms of well-known constants. For example, $C_4 = 7\zeta(3)/12$, where ζ denotes the Riemann zeta function. For larger n , it quickly became clear that the C_n approach the limit

$$\lim_{n \rightarrow \infty} C_n = 0.630473503374386796122040192710 \dots$$

This numerical value was quickly identified, using the Inverse Symbolic Calculator 2.0 (now available at <http://carma-lx1.newcastle.edu.au:8087>), as

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

where γ is Euler's constant. This identity was then proven [8].

Other specific results found in this study include the following:

$$\begin{aligned}
D_3 &= 8 + 4\pi^2/3 - 27L_{-3}(2) \\
D_4 &= 4\pi^2/9 - 1/6 - 7\zeta(3)/2 \\
E_2 &= 6 - 8\log 2 \\
E_3 &= 10 - 2\pi^2 - 8\log 2 + 32\log^2 2 \\
E_4 &= 22 - 82\zeta(3) - 24\log 2 + 176\log^2 2 - 256(\log^3 2)/3 \\
&\quad + 16\pi^2\log 2 - 22\pi^2/3 \\
E_5 &= 42 - 1984\text{Li}_4(1/2) + 189\pi^4/10 - 74\zeta(3) - 1272\zeta(3)\log 2 + 40\pi^2\log^2 2 \\
&\quad - 62\pi^2/3 + 40(\pi^2\log 2)/3 + 88\log^4 2 + 464\log^2 2 - 40\log 2,
\end{aligned}$$

where ζ is the Riemann zeta function and $\text{Li}_n(x)$ is the polylog function.

E_5 was computed by first reducing it 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 to produce the relation above. This formula remained a “numerical conjecture” for several years, but was proven in March 2014 by Erik Panzer, who mentioned that he relied on these computational results to guide his research.

1.4 Algebraic numbers in Poisson potential functions

The Poisson equation arises in contexts such as engineering applications, the analysis of crystal structures, and even the sharpening of photographic images. In two recent studies [6, 7], the present author and others explored the following class of sums:

$$\phi_n(r_1, \dots, r_n) = \frac{1}{\pi^2} \sum_{m_1, \dots, m_n \text{ odd}} \frac{e^{i\pi(m_1 r_1 + \dots + m_n r_n)}}{m_1^2 + \dots + m_n^2}. \quad (3)$$

After extensive high-precision numerical experimentation using (3), we discovered (then proved) the remarkable fact that when x and y are rational,

$$\phi_2(x, y) = \frac{1}{\pi} \log A, \quad (4)$$

where A is an *algebraic number*, namely the root of an algebraic equation with integer coefficients.

In our experiments we computed $\alpha = A^8 = \exp(8\pi\phi_2(x, y))$, using some rapidly convergent formulas found in [6], for various simple rationals x and y (as it turns out, computing A^8 reduces the degree of polynomials and so computational cost). Then we generated the vector $(1, \alpha, \alpha^2, \dots, \alpha^d)$ as input to a program implementing the three-level multipair PSLQ program [10]. When successful, the program returned the vector of integer coefficients $(a_0, a_1, a_2, \dots, a_d)$ of a polynomial satisfied by α as output. Table 1 shows some examples [6].

s	Minimal polynomial 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	$1 - 216\alpha + 860\alpha^2 - 744\alpha^3 + 454\alpha^4 - 744\alpha^5 + 860\alpha^6 - 216\alpha^7 + \alpha^8$

Table 1: Sample of polynomials produced in earlier study [6].

Using this data, Jason Kimberley of the University of Newcastle, Australia, conjectured a formula that gives the degree d as a function of k [6]. These computations required prodigiously high precision (up to 10,000 digits) and very long run times. Other runs were attempted, but failed.

In a subsequent study, the present author and three other researchers were able to dramatically extend these studies, confirming Kimberley's formula for almost all integers up to $s = 52$, and also for $s = 60$ and $s = 64$. These runs, which required precision levels up to 64,000 digits, were facilitated by a significantly improved PSLQ code, the much faster MPFUN2015 package described in [1], as well as parallel execution (facilitated by the thread-safe nature of this package). Examination of these results led to additional discoveries, and ultimately to a full proof of Kimberley's conjecture as well as other facts about these curious polynomials, such as the fact that the polynomials are palindromic (left-to-right symmetric) when s is even (see the examples in the table above) [9].

2 Current high-precision software

By far the most common form of extra-precision arithmetic is roughly twice the level of standard 64-bit IEEE floating-point arithmetic. One option is the IEEE standard for 128-bit binary floating-point arithmetic, with 113 mantissa bits; sadly it is not yet widely implemented in hardware, although it is supported, in software, by some compilers. Another option for this level of precision is “double-double” arithmetic (approximately 31 digits), which consists of two 64-bit IEEE floats, or even quad-double arithmetic (approximately 62 digits), which consists of four IEEE 64-bit floats (see the QD package in the list below) [23].

For higher-levels of precision, software packages typically represent a high-precision datum as a string of floats or integers, where the first few words contain bookkeeping information and the binary exponent, and subsequent words contain the mantissa.

Software for performing high-precision arithmetic has been available for quite some time, for example in the commercial packages *Mathematica* and *Maple*. However, until 10 or 15 years ago, those with applications written in more conventional languages, such as C++ or Fortran-90, often found it necessary to rewrite their codes, replacing each arithmetic operation with a subroutine call, which was a very tedious and error-prone process. Nowadays there are several freely available high-precision software packages, together with accompanying high-level language interfaces that make code conversions relatively painless.

Here are some packages for high-precision floating-point computation:

- ARPREC: Supports arbitrary precision integer, real, complex and transcendental functions, with high-level C++ and Fortran-90 interfaces; available at <https://www.davidhbailey.com/dhbsoftware/>.
- CLN: Supports arbitrary precision integer, real, complex and transcendental functions in C++; available at <http://www.ginac.de/CLN>.
- GMP: Supports low-level high-precision integer, rational and floating-point calculations; available at <http://gmplib.org>.
- Julia: High-level programming environment incorporating GMP and MPFR; available at <http://julialang.org>.

- MPFR: Supports (at a low level) multiple-precision floating-point computations with correct rounding, based on GMP; includes numerous algebraic and transcendental functions, and a thread-safe build option; very fast timings; available at <http://www.mpfr.org>.
- MPFR++: High-level C++ interface to MPFR (although currently available version is not up-to-date with MPFR); available at <http://perso.ens-lyon.fr/nathalie.revol/software.html>.
- MPFR C++: High-level C++ interface to MPFR with a thread-safe option; available at <http://www.holoborodko.com/pavel/mpfr>.
- mpmath: Python library for arbitrary precision floating-point arithmetic, with transcendentals; available at <https://code.google.com/p/mpmath>.
- NTL: C++ library for arbitrary precision integer and floating-point arithmetic; available at <http://www.shoup.net/ntl>.
- Pari/GP: Computer algebra system including high-precision arithmetic and transcendental functions; available at <http://pari.math.u-bordeaux.fr>.
- QD: Supports “double-double” (approx. 31 digits) and “quad-double” (approx. 62 digits) arithmetic, with transcendental functions; includes high-level interfaces for C++ and Fortran-90; available at <https://www.davidhbailey.com/dhbsoftware/>.
- Sage: Open-source symbolic computing system including high-precision facilities; available at <http://www.sagemath.org>.

2.1 Thread safety and parallel implementations

The scientific computing world is moving rapidly into multicore and multi-node parallel computing [31], so that future improvements in performance on high-precision computations will only be obtained by aggressively exploiting parallelism. It is difficult to achieve significant parallel speedup *within* a single high-precision arithmetic operation, but parallelization at the *application* level, e.g., parallelizing a DO or FOR loop in an application code that contains multiprecision operations, is an attractive option, given the high cost of many multiprecision applications.

On modern systems that feature multicore processors, parallel computing is most efficiently performed using a shared memory, multithreaded environment such as OpenMP [31] within a single node, even if a message passing system such as MPI is employed for parallelism between nodes.

Computations that use a thread-parallel environment such as OpenMP must be entirely “thread-safe.” One impediment to thread safety for multiprecision applications is the design of the operator overloading feature (i.e., extending $+$, $-$, \times , \div to multiprecision data) of modern computer languages, which typically does not permit one to carry information such as the current working precision level. Also, most high-precision packages generate a “context” of auxiliary data, such as the current working precision level and data to support transcendental function evaluation, which, if not handled properly, can ruin thread safety.

Of the packages listed in Section 2, only one is both thread-safe and supports arbitrary precision floating-point calculations with a high-level interface, namely the MPFR C++ package [25]. This package is built upon the lower-level MPFR package [21], which in turn is well-designed, features correct rounding to the last bit, includes numerous transcendental and special functions, and achieves the the fastest overall timings of any floating-point package in the above list [17].

There was, to this author’s knowledge, no high-level, thread-safe arbitrary precision package to support Fortran applications, prior to MPFUN2015.

2.2 MPFUN2015

In response to these challenges, in 2015 the present author developed a new software package for arbitrary-precision computation, named MPFUN2015, which is documented in an earlier paper [1]. It is offered in two versions: (a) MPFUN-Fort: an all-Fortran version based on floating-point arithmetic; and (b) MPFUN-MPFR: a version that calls the MPFR package for all low-level functions and operations. The run-time performance of MPFUN-MPFR is several times faster than MPFUN-Fort, but installation is significantly more complicated (because the GMP and MPFR packages must first be installed).

Both versions feature full support for both real and complex multiprecision datatypes, including transcendental functions. In most cases, only minor changes need be made to convert an existing double precision code. Both versions are thread-safe, so that user applications can be converted to parallel execution at the application level.

However, this software has some shortcomings, as have become evident in recent work by the present author, among others: (a) as mentioned above, the run-time performance of the MPFUN-Fort package is typically 3X-5X slower than the MPFUN-MPFR package; (b) with both versions of the package, significant memory is wasted for applications where part of multiprecision arrays require only a medium level of precision; and (c) for such applications, significant CPU time is wasted in data movement.

3 MPFUN2020: A new thread-safe package

To that end, the present author has written a new thread-safe arbitrary precision package. As before, it is offered in two versions: an all-Fortran version (MPFUN20-Fort), which runs several times faster than the earlier MPFUN-Fort, and also a new version of MPFUN-MPFR that is even faster and more complete than before.

Both of these versions feature: (a) a completely thread-safe design, so user codes can be converted for parallel execution at the application level; (b) a full-featured high-level Fortran interface, so that most applications can be converted to multiprecision with relatively minor changes to source code; (c) full support for both real and complex datatypes; (d) run-time checking and other facilities to overcome nagging problems with converting double precision constants and data; (e) a medium precision datatype, which improves performance and reduces memory cost on large mixed-precision applications; and (f) interoperability — with certain simple restrictions, application codes written for one version can also be run with the other.

The sharply improved performance of the all-Fortran MPFUN20-Fort version is due primarily to changing the underlying design of the package to be based on 64/128-bit integer operations, rather than on 64-bit floating-point operations as in the earlier MPFUN-Fort. Additionally, the MPFUN20-Fort version employs FFT-based arithmetic for significantly faster execution at very high precision levels — see Sections 7.4 and 7.5. The new MPFUN-MPFR version is faster than the MPFUN20-Fort version on most applications, although it is more complicated to install, because the GMP and MPFR packages must first be installed.

What follows is a brief description of the MPFUN2020 software, followed by instructions for installation and usage.

3.1 Data structure

For the MPFUN20-Fort version, the structure is a $(N + 6)$ -long vector of 64-bit integers, where N is the number of mantissa words:

- Word 0: Total space allocated for this array, in 64-bit integer words.
- Word 1: The working precision level (in words) associated with this data.
- Word 2: The number of mantissa words N ; the sign of word 2 is the sign of the value.
- Word 3: The multiprecision exponent, base 2^{60} .
- Word 4 to $N + 3$: N mantissa words (whole numbers between 0 and $2^{60} - 1$).
- Word $N + 4$ and $N + 5$: Scratch words for internal usage.

For the MPFUN-MPFR version, the structure is a $(N + 6)$ -long vector of 64-bit integers, where N is the number of mantissa words:

- Word 0: Total space allocated for this array, in 64-bit words.
- Word 1: The working precision level (in bits) associated with this data.
- Word 2: The sign of the value.
- Word 3: The exponent, base 2.
- Word 4: A pointer to the first word of the mantissa, which in MPFUN-MPFR always points to Word 5.
- Word 5 to $N + 4$: Mantissa words (unsigned integers between 0 and $2^{64} - 1$).
- $N + 5$: Not used at present.

Note that in the MPFUN-MPFR version, words 1 through $N + 4$ correspond exactly to the data structure of the MPFR package.

For each version, a complex multiprecision datatype is a contiguous pair of real multiprecision data. The imaginary member of the real-imaginary pair starts at an offset in the array equal to the value of word 0. Note that this offset is independent of the working precision.

3.2 Modules

The MPFUN20-Fort and MPFUN-MPFR versions both include the following separate modules, each in its own source file:

1. MPFUNA: Contains compile-time global data. In the MPFUN20-Fort version, this includes the binary values of $\log 2$ and π (up to 20,000-digit precision).
2. MPFUNB (only in MPFUN20-Fort): Handles basic arithmetic functions, rounding, normalization, square roots and n -th roots.
3. MPFUNC (only in MPFUN20-Fort): Handles binary-to-decimal conversion, decimal-to-binary conversion and input/output operations.
4. MPFUND (only in MPFUN20-Fort): Includes routines for all common transcendental constants and functions, as well as special routines, implementing advanced algorithms, for very high precision levels.
5. MPFUNE (only in MPFUN20-Fort): Includes routines for special functions, such as the BesselJ, gamma, incomplete gamma and zeta functions.
6. MPFUNF: Defines the default standard precision and medium precision levels, in digits, and also the equivalent levels in words. This is the only module that needs to be modified by the user in normal usage.
7. MPFUNG: A high-level user interface that provides support for the standard high-precision datatype.
8. MPFUNH: A high-level user interface that provides support for the medium precision datatype.
9. MPMODULE: The main module that references the others; in normal usage it is the only module that is directly referenced by the user.

3.3 The MPFUN2020 solution to thread safety

All of the software modules above are 100% thread safe. There are no global parameters or arrays, except for static, compile-time data, and no initialization is required unless extremely high precision is required. The working

precision level is passed as a subroutine argument, ensuring thread safety. The MPFUN-MPFR version is thread safe provided that the MPFR package is compiled with the thread-safe option.

Thread safety at the language interface or user level in both versions is achieved by assigning a working precision level to *each multiprecision datum*, which then is passed through the multiprecision software. Note, in the data structures given in Section 3.1 above, that word 1 (the second word of the array) is the working precision level associated with that datum. This solves the thread safety problem when precision is dynamically changed in the application, although it requires a somewhat different programming style, as we shall briefly explain below (see Section 5.4).

All computations are performed to a fixed precision (set by the user in module MPFUNF) unless the user specifies a lower value for certain calculations (see Section 5.4 for details). The result of any operation involving multiprecision variables or array elements inherits the working precision level of the input operands (if the operands have different working precision levels, the higher precision level is chosen for the result). When assigning a multiprecision variable or array element to double precision constant or variable, or when reading multiprecision data from a file, the result is assigned the default precision unless a lower precision level is specified.

4 Installation, compilation and linking

The two versions of the MPFUN2020 package, together with installation instructions, are available at <https://www.davidhbhailey.com/dhbsoftware>.

Installation, compilation and linking is relatively straightforward, provided that one has a Unix-based system, such as Linux or Apple OS X, and a Fortran-2008-compliant compiler that includes support for the `integer(16)` datatype). See the file `README.txt` in the distribution package for details.

The gfortran compiler (highly recommended for both versions of the package) is available for a variety of systems at <https://gcc.gnu.org/wiki/GFortranBinaries>. Both versions should also work with IBM's `xlf_r` compiler, Intel's `ifort` compiler and Portland Group's `pgf90` compiler, although the author has not recently checked these other compilers. Sample scripts to compile the library and compile and link user codes are included for the gfortran and Intel ifort compilers.

Each version of the software comes in four variants:

- Variant 1: This is recommended for basic applications that do not dynamically change the precision level (or do so only rarely).
- Variant 2: This is recommended for more sophisticated applications that dynamically change the precision level (see Section 5.4 below).
- Variant Q1: This is the same as variant 1, except that it includes some support for the `real(16)` “quad” datatype (provided it is supported by the Fortran compiler).
- Variant Q2: This is the same as variant 2, except that it includes some support for the `real(16)` “quad” datatype (provided it is supported by the Fortran compiler).

The specific source files required for each of the above variants are given in the REAMDE file of the distribution package.

Compile/link scripts are available in the `fortran` directory for the `gfortran` and Intel `ifort` compilers. These scripts automatically select the proper source files from the package for compilation. For example, to compile variant 1 of either the MPFUN20-Fort or MPFUN-MPFR library using the GNU `gfortran` compiler, go to the `fortran` directory and type

```
./gnu-complib1.scr
```

To compile and link the application program `tpslq1.f90` for variant 1, using the GNU `gfortran` compiler, producing the executable file `tpslq1`, type

```
./gnu-complink1.scr tpslq1
```

To execute the program, with output to `tpslq1.txt`, type

```
./tpslq1 > tpslq1.txt
```

These scripts assume that the user program is in the same directory as the library files; this can easily be changed by editing the script files.

Several sample test programs are included in the `fortran` directory of the packages, together with output files — see Section 6.

5 Coding Fortran applications

A high-level Fortran interface is provided for each version of the package (MPFUN20-Fort and MPFUN-MPFR).

To use the software in a user program, first set the parameter `mpipl`, the default standard precision level in digits, which is the maximum precision level to be used for subsequent computation, and is used to specify the amount of storage required for multiprecision data. `mpipl` is set in a parameter statement in file `mpfunf.f90` in the fortran directory of the software. In the code as distributed, `mpipl` is set to 2500 digits (sufficient to run the six test programs of Section 6), but it can be set to any level greater than 50 digits. `mpipl` is automatically converted to mantissa words by the formula

$$\text{mpwds} = \text{int} (\text{mpipl} / \text{mpdpw} + 2),$$

where `mpdpw` is a system parameter set in file `mpfuna.f90` ($\text{mpdpw} = \log_{10} 2^{60} \approx 18.0617997 \dots$ for MPFUN20-Fort and $\log_{10} 2^{64} \approx 19.2659197 \dots$ for MPFUN-MPFR). The resulting parameter `mpwds` is the internal default precision level, in words. All subsequent computations are performed to `mpwds` words precision unless the user, in an application code, specifies a lower precision.

After setting the value of `mpipl`, if needed, compile the appropriate version of the library, using one of the scripts mentioned in the previous section.

Next, place the following line in every subprogram of the user's application code that contains a multiprecision variable or array, at the beginning of the declaration section, before any implicit or type statements:

```
use mpmodule
```

To designate a variable or array as multiprecision real in an application program, use a Fortran-90 type statement with the type `mp_real`, as in this example:

```
type (mp_real) a, b(m), c(m,n)
```

Similarly, to designate a variable or array as multiprecision complex, use a type statement with the type `mp_complex`. Thereafter when one of these variables or arrays appears, as in the code

```
d = a + b(i) * sqrt(3.d0 - c(i,j))
```

the proper underlying multiprecision routines are automatically called.

Most common mixed-mode combinations (arithmetic operations, comparisons and assignments) involving multiprecision real (MPR), multiprecision complex (MPC), double precision (DP), double complex (DC), and integer operands are supported. A complete list of supported mixed-mode operations is given in Table 2. See Section 5.3 below about DP and DC constants and expressions.

Input and output of MP variables or array elements are done by using

Operator	Arg 1	Arg 2	Operator	Arg 1	Arg 2
a = b (assignment)	MPR	MPR	+, -, *, / (+, -, ×, ÷)	MPR	MPR
	DP	MPR		DP	MPR
	Int	MPR		MPR	DP
	MPR	MPC		Int	MPR
	MPC	MPR		MPR	Int
	MPC	MPC		MPC	MPC
	DP	MPC		DP	MPC
	DC	MPC		MPC	DP
	MPR	DP [1]		DC	MPC
	MPR	Int [1]		MPC	DC
	MPR	Char [1]		MPR	MPC
	MPC	DP [1]		MPC	MPR
	MPC	DC [1]			
a**b (a^b)	MPR	Int	==, /= (=, ≠ tests)	MPR	MPR
	MPR	MPR		DP	MPR
	MPC	Int		MPR	DP
	MPC	MPC		Int	MPR
	MPR	MPC		MPR	Int
	MPC	MPR		MPC	MPC
<=, >=, <, > (≤, ≥, <, > tests)	MPR	MPR		DP	MPC
	DP	MPR		MPC	DP
	MPR	DP		DC	MPC
	Int	MPR		MPC	DC
	MPR	Int		MPR	MPC
				MPC	MPR

Table 2: Supported mixed-mode operator combinations. MPR denotes multiprecision real, MPC denotes multiprecision complex, DP denotes double precision, DC denotes double complex, Int denotes integer and Char denotes arbitrary-length character string. Note:

[1] These operations are not allowed in variant 2 or Q2 — see Section 5.4.

the special subroutines `mpread` and `mpwrite`. See Table 4 and Section 5.2.

In MPFUN20-Fort, the above instructions apply if the precision level, namely `mpip1`, is 12,600 digits or less. If one requires a precision level greater than 12,600 digits, then in addition to changing the value of `mpip1` (and, thus, `mpwds`) in module MPFUNF in file `mpfunf.f90`, one must include a call to `mpinit` at the start of execution, in a single-threaded section of code, before any multiprecision operations are performed:

```
call mpinit (mpwds)
```

See Table 4 and Section 5.4 for additional details.

5.1 Functions and subroutines

Most Fortran-2008 intrinsic functions [20] are supported with MPR and MPC arguments, as appropriate. A full listing of these functions is shown in Table 3. In each case, these functions represent a straightforward extension to MPR or MPC arguments, as indicated. Tables 4 and 5 present a list of additional functions and subroutines provided in this package. In these tables, “F” denotes function, “S” denotes subroutine, “MPR” denotes multiprecision real, “MPC” denotes multiprecision complex, “DP” denotes double precision, “DC” denotes double complex, “Int” denotes integer and “Q” denotes IEEE quad precision, i.e. `real(16)`, if supported by the compiler. The variable names `r1,r2,r3` are MPR; `z1` is MPC; `d1` is DP; `dc1` is DC; `i1,i2,i3` are integers; `s1` is `character(1)`; `sn` is `character(n)` for any `n`; and `rr` is a vector of MPR of length `i1`.

5.2 Input and output of multiprecision data

Binary-decimal conversion and input or output of multiprecision data are handled by special subroutines, as briefly mentioned in Table 4:

1. subroutine `mpeform (r1,i1,i2,s1)`. This converts the MPR number `r1` into decimal character form in the `character(1)` array `s1`. The argument `i1` (input) is the length of the output string, and `i2` (input) is the number of digits after the decimal point. The format is analogous to Fortran E format. The result is left-justified among the `i1` cells of `s1`. The condition `i1 >= i2 + 20` must hold.
2. subroutine `mpfform (r1,i1,i2,s1)`. This converts the MPR number `r1` into decimal character form in the `character(1)` array `s1`. The

Type	Name	Description
MPR	<code>abs(r1)</code>	Absolute value
MPR	<code>abs(z1)</code>	Absolute value of complex arg
MPR	<code>acos(r1)</code>	Inverse cosine
MPR	<code>acosh(r1)</code>	Inverse hyperbolic cosine
MPR	<code>aimag(r1)</code>	Imaginary part of complex arg
MPR	<code>aint(r1)</code>	Truncate to integer
MPR	<code>anint(r1)</code>	Round to closest integer
MPR	<code>asin(r1)</code>	Inverse sine
MPR	<code>asinh(r1)</code>	Inverse hyperbolic sine
MPR	<code>atan(r1)</code>	Inverse tangent
MPR	<code>atan2(r1,r2)</code>	Arctangent with two args
MPR	<code>atanh(r1)</code>	Inverse hyperbolic tangent
MPR	<code>bessel_j0(r1)</code>	Bessel function of the first kind, order 0
MPR	<code>bessel_j1(r1)</code>	Bessel function of the first kind, order 1
MPR	<code>bessel_jn(i1,r1)</code>	Bessel function of the first kind, order <code>i1</code>
MPR	<code>bessel_y0(r1)</code>	Bessel function of the second kind, order 0 [1]
MPR	<code>bessel_y1(r1)</code>	Bessel function of the second kind, order 1 [1]
MPR	<code>bessel_yn(i1,r1)</code>	Bessel function of the second kind, order <code>i1</code> [1]
MPC	<code>conjg(z1)</code>	Complex conjugate
MPR	<code>cos(r1)</code>	Cosine of real arg
MPC	<code>cos(z1)</code>	Cosine of complex arg
MPR	<code>cosh(r1)</code>	Hyperbolic cosine
DP	<code>dble(r1)</code>	Convert MPR argument to DP
DC	<code>dcmplx(z1)</code>	Convert MPC argument to DC
MPR	<code>erf(r1)</code>	Error function
MPR	<code>erfc(r1)</code>	Complementary error function
MPR	<code>exp(r1)</code>	Exponential function of real arg
MPR	<code>exp(z1)</code>	Exponential function of complex arg
MPR	<code>gamma(r1)</code>	Gamma function
MPR	<code>hypot(r1,r2)</code>	Hypotenuse of two args
MPR	<code>log(r1)</code>	Natural logarithm of real arg
MPR	<code>log(z1)</code>	Natural logarithm of complex arg
MPR	<code>log_gamma(r1)</code>	Log gamma function [1]
MPR	<code>max(r1,r2)</code>	Maximum of two (or three) args
MPR	<code>min(r1,r2)</code>	Minimum of two (or three) args
MPR	<code>mod(r1,r2)</code>	Mod function = <code>r1 - r2*aint(r1/r2)</code>
MPR	<code>sign(r1,r2)</code>	Transfer of sign from <code>r2</code> to <code>r1</code>
MPR	<code>sin(r1)</code>	Sine function of real arg
MPC	<code>sin(z1)</code>	Sine function of complex arg
MPR	<code>sqrt(r1)</code>	Square root of real arg
MPC	<code>sqrt(z1)</code>	Square root of complex arg
MPR	<code>sinh(r1)</code>	Hyperbolic sine
MPR	<code>tan(r1)</code>	Tangent function
MPR	<code>tanh(r1)</code>	Hyperbolic tangent function

Table 3: Fortran-2008 intrinsic functions extended to multiprecision. Notes: [1]: Only available with MPFUN-MPFR.

Type	Name	Description
F(MPC)	<code>mpcplx(r1,r2)</code>	Converts (r1,r2) to MPC [1]
F(MPC)	<code>mpcplx(dc1)</code>	Converts DC arg to MPC [1]
F(MPC)	<code>mpcplx(z1)</code>	Converts MPC arg to MPC [1]
F(MPC)	<code>mpcplxdc(dc1)</code>	Converts DC to MPC, without checking [1, 2]
S	<code>mpcssh(r1,r2,r3)</code>	Returns both cosh and sinh of r1, in the same time as calling just cosh or just sinh
S	<code>mpcssn(r1,r2,r3)</code>	Returns both cos and sin of r1, in the same time as calling just cos or just sin
S	<code>mpdecmd(r1,d1,i1)</code>	Converts r1 to DP with power of 10
S	<code>mpeform(r1,i1,i2,s1)</code>	Converts r1 to char(1) string in Ei1.i2 format, suitable for output (Sec. 5.2)
S	<code>mpfform(r1,i1,i2,s1)</code>	Converts r1 to char(1) string in Fi1.i2 format, suitable for output (Sec. 5.2)
F(MPR)	<code>mpegamma()</code>	Returns Euler's γ constant [1]
S	<code>mpinit()</code>	Initializes for extra-high precision (Sec. 5) [1]
F(MPR)	<code>mplog2()</code>	Returns $\log(2)$ [1]
F(MPR)	<code>mpnrt(r1,i1)</code>	Returns the i1-th root of r1
F(MPR)	<code>mppi()</code>	Returns π [1]
F(MPR)	<code>mpprodd(r1,d1)</code>	Returns $r1*d1$, without checking [2]
F(MPR)	<code>mpquodd(r1, d1)</code>	Returns $r1/d1$, without checking [2]
F(MPR)	<code>mprand(r1)</code>	Returns pseudorandom number, based on r1 Start with an irrational, say $r1 = mplog2()$ Typical iterated usage: $r1 = mprand(r1)$
S	<code>mpread(i1,r1)</code>	Inputs r1 from Fortran unit i1; up to five MPR args may be listed (Sec. 5.2) [1]
S	<code>mpread(i1,z1)</code>	Inputs z1 from Fortran unit i1; up to five MPC args may be listed (Sec. 5.2) [1]
F(MPR)	<code>mpreal(r1)</code>	Converts MPR arg to MPR [1]
F(MPR)	<code>mpreal(z1)</code>	Converts MPC arg to MPR [1]
F(MPR)	<code>mpreal(d1)</code>	Converts DP arg to MPR [1, 2]
F(MPR)	<code>mpreal(q1)</code>	Converts real(16) (quad) to MPR [1, 2]
F(MPR)	<code>mpreal(s1,i1)</code>	Converts char(1) string to MPR (Sec. 5.2) [1]
F(MPR)	<code>mpreal(sn)</code>	Converts char(n) string to MPR (Sec. 5.2) [1]
F(MPR)	<code>mpreald(d1)</code>	Converts DP to MPR, without checking [1, 2]
F(Int)	<code>mpwprec(r1)</code>	Returns precision in words assigned to r1
F(Int)	<code>mpwprec(z1)</code>	Returns precision in words assigned to z1
S	<code>mpwrite(i1,i2,i3,r1)</code>	Outputs r1 in Ei2.i3 format to unit i1; up to five MPR args may be listed (Sec. 5.2)
S	<code>mpwrite(i1,i2,i3,z1)</code>	Outputs z1 in Ei2.i3 format to unit i1; up to five MPC args may be listed (Sec. 5.2)
F(Q)	<code>qreal(r1)</code>	Converts MPR to real(16) (quad).

Table 4: Additional general routines (F: function, S: subroutine). Notes:
[1]: In variant 1 or Q1, an integer precision level argument (mantissa words) may optionally be added as the final argument; this argument is required in variant 2 or Q2. See Sec. 5.4.
[2]: These do not check DP or DC values. See Sec. 5.3.

Type	Name	Description
F(MPR)	<code>agm(r1,r2)</code>	Arithmetic-geometric mean
F(MPR)	<code>airy r1</code>	Airy function [3]
S	<code>berne(i1,rr)</code>	Array of first <code>i1</code> even Bernoulli numbers [1, 2]
F(MPR)	<code>besselj(r1,r2)</code>	BesselJ function with two MPR arguments [2]
F(MPR)	<code>digamma(r1)</code>	Digamma function [3]
F(MPR)	<code>expint(r1)</code>	Exponential integral function [3]
F(MPR)	<code>gammainc(r1,r2)</code>	Incomplete gamma function [2]
F(MPR)	<code>polylog(n1,r1)</code>	Polylogarithm function [3, 4]
F(MPR)	<code>zeta(r1)</code>	Zeta function
F(MPR)	<code>zetaem(n1,rr,r1)</code>	Zeta function with precomputed even Bernoulli numbers [2]

Table 5: Additional special functions. Notes:

[1]: In variant 1 or Q1, an integer precision level argument (mantissa words) may optionally be added as the final argument; this argument is required in variant 2 or Q2. See Sec. 5.4.

[2]: Only available with MPFUN20-Fort.

[3]: Only available with MPFUN-MPFR.

[4]: Currently restricted to `n1 = 2`.

argument `i1` (input) is the length of the output string, and `i2` (input) is the number of digits after the decimal point. The format is analogous to Fortran F format. The result is right-justified among the `i1` cells of `s1`. The condition `i1 >= i2 + 20` must hold.

3. `subroutine mpread (i1,r1)`. This reads the MPR number `r1`, presumed in decimal format, from Fortran logical unit `i1`. The digits of `r1` may span more than one line, provided that a backslash appears at the end of a line to be continued. Individual input lines may not exceed 2048 characters in length. Format: The input number must have a leading digit (possibly zero), must have a period somewhere; may include an `e`, `d`, `E` or `D` followed by an integer exponent; but must *not* have embedded blanks. Up to five MPR arguments may be included in argument list. See item 9 below on an additional precision argument.
4. `subroutine mpread (i1,z1)`. This is the same as the previous item (3), except that the input argument `z1` is of type MPC (a pair of MPR). Up to five MPC arguments may be included in argument list. See item 9 below on an additional precision argument.
5. `function mpreal (s1,i1)`. This converts the string `s1`, which is of type `character(1)` and length `i1`, to MPR. See item 3 for format. See

item 9 below on an additional precision argument.

6. **function mpreal (sn)**. This converts the string `sn`, which may be of type `character(n)` for any `n`, to MPR. See item 3 for format. On some systems, `n` may be limited, say to 2048; if this is a problem, use item 5. See item 9 below on an additional precision argument.
7. **subroutine mpwrite (i1,i2,i3,r1)**. This writes the MPR number `r1` to Fortran logical unit `i1`, in a format analogous to Fortran E format, left-justified in the field. The argument `i2` (input) is the length of the output field, and `i3` (input) is the number of digits after the decimal point. Up to five MPR arguments may be included in argument list.
8. **subroutine mpwrite (i1,i2,i3,z1)**. This is the same as the previous item (7), except that the argument `z1` is of type MPC (a pair of MPR). Up to five MPC arguments may be included in argument list.
9. Note: For `mpread` (items 3 and 4) and `mpreal` (items 5 and 6), when using variant 1 or Q1, an integer working precision level argument (mantissa words) may optionally be added as the final argument; this argument is required in variant 2 or Q2. See Section 5.4.

5.3 Handling double precision values

Double precision constants and expressions are indispensable in high-precision applications. For one thing, the product, say, of a multiprecision value times a double precision value is more rapidly computed with a routine dedicated to this task than converting the double precision value to multiprecision and then calling the full multiplication routine. Certainly the usage of double precision constants such as modest-sized whole numbers and exact binary fractions (e.g., 0., 12345., 2.5, 6.125), which are entirely safe in a multiprecision application, should be allowed.

However, problems can arise, which are inherent in how modern computer languages handle double precision constants and expressions. For example, the code

```
r1 = 3.14159d0
```

where the left-hand side is multiprecision, does *not* produce the full-precision equivalent of 3.14159, since by standard rules of precedence, the right-hand side is evaluated in double precision, then converted (by zero extension) to

the left-hand side, yielding a result accurate to only 15-digits or so. When using the MPFUN2020 package, one can avoid this problem by writing

```
r1 = '3.14159'
```

By enclosing the constant in apostrophes (and changing it to a literal), this indicates to the MPFUN2020 software that the constant is to be converted to full multiprecision accuracy.

A closely related problem is the following: The code

```
r2 = r1 + 3.d0 * sqrt (2.d0)
```

does *not* produce a fully accurate result, since the subexpression `3.d0 * sqrt (2.d0)` is performed in double precision (again, according to standard rules of operator precedence in all major programming languages). The solution here is to write this as

```
r2 = r1 + 3.d0 * sqrt (mpreal (2.d0))
```

or, if using variant 2 or Q2, as

```
r2 = r1 + 3.d0 * sqrt (mpreal (2.d0, nwd))
```

where `nwd` is the precision level in words to be assigned to the constant 2 (see Section 5.4). This forces all operations to be done using MP routines.

To help avoid such problems, both versions of the MPFUN2020 software check *every* double precision value (constants, variables and expression values) in a multiprecision statement at *execution time* to see if it has more than 40 significant bits. If so, it is flagged as an error, since very likely such usage represents an unintended loss of precision in the application program. This feature catches 99.99% of accuracy loss problems due to the usage of inexact double precision values.

On the other hand, some applications (including several of the sample test codes mentioned in Section 6) contain legitimate double precision constants that are trapped by this test. In order to permit such usage, four special functions have been provided: `mpprodd`, `mpquotd`, `mpreald`, `mpcmlxdc` (see Table 4). The first and second return the product and quotient, respectively, of a MPR argument and a DP argument; the third converts a DP value to MPR (with an optional precision level parameter — see Section 5.4); and the fourth converts a DC value to MPC (with an optional precision level parameter — see Section 5.4). These routines do *not* check the double precision argument to see if it has more than 40 significant bits.

5.4 Dynamically changing the working precision

Some multiprecision applications run fine with a fixed precision level, but others are more efficiently implemented with a working precision level that is changed frequently. Accordingly, for each version of the package, there are four variants of the package (see Section 4):

- Variant 1 or Q1: This is recommended for basic applications that do not dynamically change the precision level (or do so only rarely).
- Variant 2 or Q2: This is recommended for more sophisticated applications that dynamically change the precision level.

In particular, with variant 1 or Q1:

1. Assignments of the form $R = X$, where R is MPR and X is DP, integer or literal, are permitted. Assignments of the form $Z = Y$, where Z is MPC and Y is DP or DC, are permitted.
2. The routines `mpcplx`, `mpcplxdc`, `mpegamma`, `mpinit`, `mplog2`, `mppi`, `mpread`, `mpreal` and `mpreald` each have an (optional) integer argument as the final argument in the list. This argument is the working precision level, in words, to be assigned to the result(s). If this argument is not present, the default precision level (`mpwds` words, corresponding to `mpipl` digits) is assumed.

In contrast, with variant 2 or Q2:

1. The assignments mentioned in item 1 above are *not permitted*. If any of these appears in code, compile-time errors will result. Instead, one must use `mpreal` and `mpcplx`, as appropriate, with the precision level (in mantissa words) as the final argument, to perform these conversions.
2. The optional working precision level arguments mentioned in item 2 above are *required* in all cases. For example, if the full default precision level (`mpwds`, corresponding to `mpipl` digits) is required, then a call to one of the functions or subroutines in item 2 above must have `mpwds` as the final argument.

Note that the `mpreal` function, with the precision level (in words) as the second argument, can be used to assign an MPR argument with one precision

level to an MPR variable or array element with a different working precision level. The same is true of `mpcplx`. The working precision currently assigned to any MP variable or array element may be obtained by using the function `mpwprec` — see Table 4.

Along this line, when one uses the precision level arguments, a precision level of `ndig` digits can be converted to words by the formula `nwds = int (ndig / mpdpw + 2)`. By using the global built-in variable `mpdpw` (which is different between MPFUN20-Fort and MPFUN-MPFR) in this way, the user code remains portable between the two versions.

As it turns out, in most applications, even those that frequently require the working precision to be changed, only a few changes need to be made to the source code to implement variable precision. Consider, for example, the following user code, where the default precision is set in file `mpfunf.f90` to 2500 digits:

```
integer k, nx
parameter (nx = 128)
type (mp_real) x(nx)
x(1) = 1.d0
do k = 2, nx
    x(k) = 2.d0 * x(k-1) + 1.d0
enddo
```

This code, as written, is permissible with variant 1 or Q1, but not with variant 2 or Q2, because the assignment `x(k) = 1.d0` is not allowed. Furthermore, all operations are performed with the default (maximum) precision level of 2500 digits. So with variant 2 or Q2, where one wishes to perform this loop with a precision level of, say, 500 digits, this should be written as:

```
integer k, ndig, nwds, nx
parameter (nx = 128, ndig = 500, nwds = int (ndig / mpdpw + 2))
type (mp_real) x(nx)
x(1) = mpreal (1.d0, nwds)
do k = 2, nx
    x(k) = 2.d0 * x(k-1) + 1.d0
enddo
```

Note that by changing `x(1) = 1.d0` to `x(1) = mpreal (1.d0, nwd)`, the array element `x(1)` is assigned the value 1.0, with a working precision of `nwd` words (i.e., 500 digits). In the loop, when `k` is 2, `x(2)` also inherits the working precision level `nwd` words, since it is computed with an expression that involves `x(1)`. By induction, all elements of the array `x` inherit the working precision level `nwd` words (i.e., 500 digits).

This scenario is entirely typical of using variable precision — in most cases, it is only necessary to make a few code changes, such as in assignments to double precision values before a loop, to completely control dynamic precision. It is highly recommended, though, that the user frequently employ the system function `mpwprec`, which returns the working precision level currently assigned to an multiprecision variable or array element (see Table 4), to ensure that the working precision level the user thinks is assigned to a variable is indeed the level being used by the program.

Using variant 2 or Q2, with its stricter coding standards, requires a bit more programming effort, but in the present author's experience, when dealing with applications that dynamically change the precision level, this additional effort is more than repaid by fewer debugging and performance problems in actual usage. A code written for variant 2 or Q2 also works with variant 1 or Q1, but not vice versa. See the sample test codes mentioned in Section 6, all of which are written to conform to the stricter standards of variant 2 or Q2.

5.5 Medium precision datatype

In many high-precision applications, only part of the multiprecision variables and arrays contain full precision data; others contain data with only modest precision, typically only a fraction as high (although still higher than double precision). Since all multiprecision data are allocated sufficient space to accommodate full precision values, much of the storage and data movement costs for modest precision data are wasted.

To reduce memory usage and improve performance in such applications, a medium precision real and a medium precision complex datatype have been defined, and are available in both the MPFUN20-Fort and MPFUN-MPFR versions of the software. To use these datatypes, do the following:

First, set the medium precision level `mpiplm` in file `mpfun.f90`; by default it is set to 250 digits. Then in each user subprogram that will include either full precision or medium precision data, insert the following line:

```
use mpmodule
```

as with standard full precision. To designate a variable or array as medium precision real, use a Fortran-90 type statement with the type `mp_realm`, as in this example:

```
type (mp_realm) a, b(m), c(m,n)
```

Similarly, use the type `mp_cmplx` for medium precision complex data.

Direct assignments between full precision and medium precision variables, as well as mixed-mode arithmetic operations between full precision and medium precision variables, are *not* allowed. If one needs to convert a regular full precision value to a medium precision value, use the function `mprealm`, as in this example

```
type mp_real a; mp_realm b
b = mprealm (a)
```

Similarly, to convert a medium precision value to a full precision value, use the function `mpreal`, as in

```
type mp_real a; mp_realm b
a = mpreal (b)
```

Note however, that with either `mpreal` or `mprealm`, in variant 1 or Q1, an integer working precision level argument (mantissa words) may optionally be added as the final argument; this argument is required in variant 2 or Q2. See Section 5.4 for details.

Each of the intrinsic functions listed in tables 2, 3, 4 and 5 is also supported with the medium precision datatype, except for the following: `mpreald`, `mpcmplx`, `mpcmplxd`, `mpegamma`, `mppi` and `mplog2`. The medium precision equivalents of these functions are: `mprealdm`, `mpcmplx`, `mpcmplxdm`, `mpegammam`, `mppim` and `mplog2m`, respectively. Also, the comment regarding the working precision argument in the previous paragraph applies to each of these functions. For example, if one has set the medium precision level `mpiplm` to 250 digits in file `mpfunf.f90`, and if one is using variant 2 or Q2, then one can retrieve the precomputed value of π to a medium precision variable `b` by typing

```
b = mppim (mpiplm)
```

Usage of the medium precision datatype is illustrated in the programs `tpslqm3.f90` and `tpphix3.f90` in the set of sample application programs. See Section 6 below for details.

6 Sample applications and performance

Numerous full-scale multiprecision applications have been implemented using the MPFUN2020 software, including some that dynamically vary the working precision level. In most cases, only very minor modifications needed to be made to existing double precision source code.

The current release of the software includes a set of sample application programs in the fortran directory of each version of the software. These codes are identical across the two versions (MPFUN20-Fort and MPFUN-MPFR). These codes are:

1. `tps1q1.f90`: A one-level standard PSLQ program; finds the coefficients of the degree-30 polynomial satisfied by $3^{1/5} - 2^{1/6}$. Size of code: 803 lines. Precision level: 240 digits.
2. `tps1qm1.f90`: A one-level multipair PSLQ program; finds the coefficients of the degree-30 polynomial satisfied by $3^{1/5} - 2^{1/6}$. Size of code: 954 lines. Precision level: 240 digits.
3. `tps1qm2.f90`: A two-level multipair PSLQ program; finds the coefficients of the degree-56 polynomial satisfied by $3^{1/7} - 2^{1/8}$. Size of code: 1777 lines. Precision level: 640 digits; switches frequently between multiprecision and double precision.
4. `tps1qm3.f90`: A three-level multipair PSLQ program; finds the coefficients of the degree-72 polynomial satisfied by $3^{1/8} - 2^{1/9}$. Size of code: 2242 lines. Precision level: 1100 digits; switches frequently between full precision (1100 digits), medium precision (varies from 50 to 110 digits) and double precision.
5. `tpphix3.f90`: A Poisson phi program; computes the value of $\phi_2(x, y)$ and then employs a three-level multipair PSLQ algorithm to find the minimal polynomial of degree m satisfied by $\exp(8\pi\phi_2(1/k, 1/k))$ for a given k (see Section 1.4). In the code as distributed, $k = 28$, $m = 96$, and a palindromic option is employed so that the multipair PSLQ routines (which are part of this application) search for a relation of size 49 instead of 97. This computation involves transcendental functions and both real and complex multiprecision arithmetic. Size of code: 2576 lines. Precision level: 2500 digits; switches frequently between

full precision (2500 digits), medium precision (varies from 50 to 250 digits) and double precision.

6. `tquad.f90`: A quadrature program; performs the tanh-sinh, the exp-sinh or the sinh-sinh quadrature algorithm, as appropriate, on a suite of 18 problems involving numerous transcendental function references, producing results correct to 500-digit accuracy. Size of code: 1448 lines. Precision level: 1000 digits, but most computation is done to 500 digits; switches frequently between 500 and 1000 digits.
7. `tquadgs.f90`: A quadrature program; performs the Gaussian quadrature algorithm on many of the same suite of 18 problems as in `tquad`, producing results correct to 500-digit accuracy. This code runs much longer than `tquad`, due to the expense of computing weights and abscissas. Once the weights and abscissas are computed, they are written to a file, so they can be reused in future runs. Size of code: 706 lines. Precision level: 1000 digits, but most computation is done to 500 digits; switches frequently between 500 and 1000 digits.

Corresponding reference output files, e.g. `tpphix3.ref.txt` are also included for each of the above programs, together with the sample scripts `gnu-mpfun-tests.scr` and `intel-mpfun-tests.scr`, which compile the library and run each of the above sample programs above except `tquadgs.f90`, which takes considerably more run time. If, after running this script, the results in the output files with suffix `.txt` match those in the reference output files with suffix `ref.txt` (except for timings, etc.), then one can be fairly confident that the software is working properly.

These programs are provided, in part, as examples of programming techniques when using the MPFUN2020 package. Users may feel free to adapt these codes, although the present author asks to be notified and credited when this is done. All application programs and library codes are publicly available but are subject to copyright and other legal conditions. For details, see the file `disclaimer.txt` in the distribution package.

6.1 Timings

Table 6 presents some performance timings comparing the two versions of the package for the first six test programs listed above, plus an additional

Code name	Precision (digits)	param.		MPFUN20- Fort	MPFUN- MPFR
		k	m		
<code>tpslq1</code>	240			9.37	3.70
<code>tpslqm1</code>	240			4.31	3.44
<code>tpslqm2</code>	650			4.77	4.98
<code>tpslqm3</code>	1100			19.14	21.57
<code>tquad</code>	1000			47.47	4.87
<code>tpphix3</code>	2500	28	96	16.69	14.97
<code>tpphix3</code>	5500	25	100	480.82	427.29

Table 6: Timings on a suite of test programs (seconds).

run using the `tpphix3.f90` code, with different parameters k and m , and without the palindromic option, which is not available when k is odd.

These runs were performed using the GNU gfortran compiler (version 10.2.0). For the MPFUN-MPFR version, the GNU C compiler (version 10.2.0) was used to build the GMP and MPFR libraries. These runs were made on an Apple MacPro system with a 2.4 GHz Quad-Core Intel Core i5 processor and 16 Gbyte of main memory. For uniformity, the timings are listed in the Table 6 to two decimal place accuracy, but, as with all computer run time measurements, they should not be considered repeatable beyond about two or three significant digits.

7 Appendix: Numerical algorithms

This section presents a brief overview of the algorithms used in MPFUN20-Fort. The algorithms used in MPFUN-MPFR (except for pseudorandom number generation, which is given below) are described in [21].

7.1 Algorithms for basic arithmetic

Multiplication. For modest levels of precision, MPFUN20-Fort employ adaptations of the usual schemes we all learned in grade school, where the number base is $2^{60} = 1152921504606846976$. Note that if two n -word arguments are multiplied, and the working precision is also n words, then since only an n -word result is returned, only slightly more than half of the “multiplication pyramid” need be calculated.

Division. A similar approach, which involves obtaining an accurate trial divisor, is employed for division.

Square roots. Square roots are calculated by the following Newton-Raphson iteration, which converges to $1/\sqrt{a}$ [12, pg. 227]:

$$x_{k+1} = x_k + 1/2 \cdot (1 - x_k^2 \cdot a) \cdot x_k, \quad (5)$$

where the multiplication $() \cdot x_k$ is performed with only half of the normal level of precision. These iterations are performed with a working precision level that approximately doubles with each iteration, except that at three iterations before the final iteration, the iteration is repeated without doubling the precision, in order to enhance accuracy. The final iteration is performed as follows (due to A. Karp):

$$\sqrt{a} \approx (a \cdot x_n) + 1/2 \cdot [a - (a \cdot x_n)^2] \cdot x_n, \quad (6)$$

where the multiplications $(a \cdot x_n)$ and $[] \cdot x_n$ are performed with only half the final level of precision. If this is done properly, the total cost of the calculation is only about three times the cost of a single full-precision multiplication.

n -th roots. A similar scheme is used to compute n -th roots for any integer n . Computing x_k^n , which is required here, can be efficiently performed using the binary algorithm for exponentiation. This is merely the observation that exponentiations can be accelerated based on the binary expansion of the exponent: for example, 3^{17} can be computed as $((((3)^2)^2)^2)^2 \cdot 3 = 129140163$.

Pseudorandom number generation. The pseudorandom number included in the package (both in the MPFUN20-Fort and MPFUN-MPFR) versions is the following:

$$x_{n+1} = \text{frac}(5501758857861179 \cdot x_n), \quad (7)$$

where the calculation is performed with slightly more than the standard precision for full accuracy. Here the large prime 5501758857861179 is specified as the double precision constant `mprandx`, set in module MPFUNA in file `mpfuna.f90`. Its value may be changed if desired, but must not exceed 2^{53} . To use the generator, first set a MPR variable `r1` to some irrational value between 0 and 1, such as $\sqrt{1/2}$, $\log(2)$ or $\pi/4$. Then successive iterates can be generated by typing

```
r1 = mprand (r1)
```

The same scheme is used for both MPFUN20-Fort and MPFUN-MPFR. However, the specific sequence is *not* the same between the two versions, nor is it same if the precision level is varied.

Note that the above algorithms are trivially thread-safe, since no auxiliary data is involved.

7.2 Basic algorithms for transcendental functions

Most arbitrary precision packages require a significant “context” of data to support transcendental function evaluation at a particular precision level, and this data is often problematic for both thread safety and efficiency. For example, if this context data must be created and freed within each running thread, this limits the efficiency in a multithreaded environment. With this in mind, the transcendental function routines in MPFUN20-Fort were designed to require only a minimum of context, which context is provided in static data statements, except when extremely high precision is required.

Exponential and logarithm. In the current implementation, the exponential function routine in MPFUN20-Fort first reduces the input argument to within the interval $(-\log(2)/2, \log(2)/2]$. Then it divides this value by 2^q , producing a very small value, which is then input to the Taylor series for $\exp(x)$. The working precision used to calculate the terms of the Taylor series is reduced as the terms get smaller, thus saving approximately one-half of the total run time. When complete, the result is squared q times, and then corrected for

the initial reduction. In the current implementation, q is set to the nearest integer to $(48n)^{2/5}$, where n is the number of words of precision.

Since the Taylor series for the logarithm function converges much more slowly than that of the exponential function, the Taylor series is not used for logarithms unless the argument is extremely close to one. Instead, logarithms are computed based on the exponential function, by employing the following Newton iteration with a level of precision that approximately doubles with each iteration:

$$x_k = x_k - \frac{e^x - a}{e^x}. \quad (8)$$

Trigonometric functions. The sine/cosine routine first reduces the input argument to within the interval $(-\pi, \pi]$. This value is then divided by 2^q and then input to the Taylor series for $\sin(x)$, with a linearly varying precision level as above. Then the double-angle formulas

$$\cos(2x) = 1 - 2\sin^2(x) \quad (9)$$

$$\cos(2x) = 2\cos^2(x) - 1, \quad (10)$$

are applied q times (formula (9) is used once, and (10) thereafter). In the current implementation, q is set to the greatest integer in $\sqrt{2N}$, where N is the precision in bits, unless the reduced argument is very close to one, in which case $q = 0$. When complete, $\sin(x)$ is computed as $\sqrt{1 - \cos^2(x)}$, with corrected sign, except for the case $q = 0$, when $\cos(x)$ is computed as $\sqrt{1 - \sin^2(x)}$.

The inverse \cos/\sin function is based on the sine routine, by employing a Newton iteration with a level of numeric precision that roughly doubles with each iteration.

Power function. The power function, namely a^b for real $a > 0$ and b , can be computed as $e^{b \log a}$. To further accelerate this operation, the software first examines the value of b to see if it is a rational number with numerator and denominator up to 10^7 size, using the extended Euclidean algorithm performed in double precision. If it is, a^b is performed using a combination of the binary algorithm for exponentiation for the numerator, and the n -th root function for the denominator.

Euler's gamma constant. A function is also provided to compute Euler's constant γ . Euler's constant γ was calculated using the following formula,

which is an improvement of a technique previously used by Sweeney [29]. If a result accurate to at least B bits is desired, first select the integer $N = \lceil \log_2(B \log 2) \rceil$. Then

$$\gamma \approx \frac{2^N}{e^{2^N}} \sum_{m=0}^{\infty} \frac{2^{mN}}{(m+1)!} \sum_{t=0}^m \frac{1}{t+1} - N \log 2. \quad (11)$$

The error in this approximation is less than $1/(2^N e^{2^N})$.

Note that for all of the above algorithms, the only context required is the numerical values of $\log(2)$, π , which values are stored as compile-time data in module MPFUNA to support precision up to 20,000 digits. For higher precision levels, an initialization routine must be called.

7.3 Special functions

Modern mathematical and scientific computing frequency often involves other, more sophisticated functions, which collectively are termed “special functions” [18]. A number of these functions have been implemented in the MPFUN20-Fort package, and others will be added as they are developed. Here is a brief description of the functions that have been implemented and the algorithms employed. In each case, care is taken to preserve thread safety, and to avoid, as far as possible, any need to precalculate auxiliary data.

BesselJ function. The BesselJ function, or, more formally, the Bessel function of the first kind, is defined as [18, 10.2.2]:

$$J_\nu(z) = \left(\frac{z}{2}\right)^\nu \sum_{k=0}^{\infty} \frac{(-1)^k \left(\frac{z^2}{4}\right)^k}{k! \Gamma(\nu + k + 1)}. \quad (12)$$

For modest-sized values of z (the present author uses the condition $z < 1.73d$, where d is the precision level in digits), the MPFUN20-Fort BesselJ function implements formula (12). Note that although (12) involves the gamma function (see below), this is only called once to compute $\Gamma(\nu + 1)$, after which the recursion $\Gamma(t + 1) = t\Gamma(t)$ can be applied for the other terms.

For larger values of z , the following asymptotic formula is employed [18, 10.17.3]. Define $\omega = z - \nu\pi/2 - \pi/4$, and

$$a_k(\nu) = \frac{(4\nu^2 - 1^2)(4\nu^2 - 3^2) \cdots (4\nu^2 - (2k - 1)^2)}{k! 8^k}. \quad (13)$$

Then

$$J_\nu(z) = \left(\frac{2}{\pi z}\right)^{1/2} \left(\cos \omega \sum_{k=0}^{\infty} \frac{(-1)^k a_{2k}(\nu)}{z^{2k}} - \sin \omega \sum_{k=0}^{\infty} \frac{(-1)^k a_{2k+1}(\nu)}{z^{2k+1}} \right). \quad (14)$$

One important detail omitted from the above discussion is that large amounts of cancellation occurs in these formulas. Thus when evaluating these formulas, a working precision of 1.5 times the normal working precision is employed.

No precalculated auxiliary data is needed for either of these algorithms, so they are thread safe.

Gamma function. The gamma function employs a very efficient but little-known formula due to Ronald W. Potter [27], as follows. If the input t is a positive integer, then $\Gamma(t) = (t-1)!$. If not, use the recursion $\Gamma(t+1) = t\Gamma(t)$ to reduce the argument (positive or negative) to the interval $(0, 1)$. Then define $\alpha = \text{nint}(n/2 \cdot \log 2)$, where n is the number of bits of precision and nint means nearest integer, and set $z = \alpha^2/4$. Define the Pochhammer function as

$$(\nu)_k = \nu(\nu+1)(\nu+2)\cdots(\nu+k-1). \quad (15)$$

Then define the functions

$$\begin{aligned} A(\nu, z) &= \left(\frac{z}{2}\right)^\nu \nu \sum_{k=0}^{\infty} \frac{(z^2/4)^k}{k!(\nu)_{k+1}} \\ B(\nu, z) &= \left(\frac{z}{2}\right)^{-\nu} (-\nu) \sum_{k=0}^{\infty} \frac{(z^2/4)^k}{k!(-\nu)_{k+1}}. \end{aligned} \quad (16)$$

With these definitions, the gamma function can then be computed as

$$\Gamma(\nu) = \sqrt{\frac{A(\nu, z)}{B(\nu, z)} \frac{\pi}{\nu \sin(\pi\nu)}}. \quad (17)$$

No auxiliary data is needed for this algorithm, so it is thread-safe.

Incomplete gamma function. For modest-sized positive arguments (the author uses the condition $z < 2.768d$, where d is the precision level in digits),

the MPFUN20-Fort incomplete gamma function is evaluated using the following formula [18, 8.7.3]:

$$\Gamma(a, z) = \Gamma(a) \left(1 - \frac{z^a}{e^z} \sum_{k=0}^{\infty} \frac{z^k}{\Gamma(a+k+1)} \right). \quad (18)$$

Note, as with the BesselJ function, that although formula (18) involves the gamma function, this is only called once to compute $\Gamma(a+1)$, after which the recursion $\Gamma(t+1) = t\Gamma(t)$ can be applied for all other terms.

For large values of z , the following asymptotic formula is used [18, 8.11.2]:

$$\Gamma(a, z) \approx \frac{z^{a-1}}{e^z} \sum_{k=0}^{\infty} \frac{(-1)^k (1-a)_k}{z^k} \quad (19)$$

No auxiliary data is needed for this algorithm, so it is thread-safe.

Riemann zeta function. For large positive arguments s (the present author uses the condition $s > 2.303d/\log(2.215d)$, where d is the precision in digits), it suffices to use the definition of zeta, namely

$$\zeta(s) = \sum_{n=1}^{\infty} \frac{1}{n^s}. \quad (20)$$

For modest-sized arguments, the zeta function can be evaluated by means of this formula, due to Peter Borwein [13]. Select n to be the number of digits of precision required for the result. Define

$$e_j = (-1)^j \left(\sum_{k=0}^{j-n} \frac{n!}{k!(n-k)!} - 2^n \right), \quad (21)$$

where the empty sum is zero. Then

$$\zeta(s) \approx \frac{-1}{2^n(1-2^{1-s})} \sum_{j=0}^{2n-1} \frac{e_j}{(j+1)^s}. \quad (22)$$

The above formulas are used for positive real arguments (except $s = 1$, for which the zeta function is undefined). For negative s , the following “reflection formula” is used to convert the calculation to a positive argument:

$$\zeta(s) = \frac{2 \cos(\pi(1-s)/2) \Gamma(1-s) \zeta(1-s)}{(2\pi)^{1-s}}. \quad (23)$$

Formulas (21), (22) and (23) are implemented as the zeta function. No auxiliary data for this algorithm required, so it is thread-safe.

A even faster algorithm, based on the Euler-Maclaurin summation formula, can be derived from the following [18, 25.2.9]: Select an integer parameter $N > 0$ (the present author uses $N = 0.6d$, where d is the number of digits of precision). Then

$$\zeta(s) \approx \sum_{k=1}^N \frac{1}{k^s} + \frac{1}{(s-1)N^{s-1}} - \frac{1}{2N^s} + \sum_{k=1}^{\infty} \binom{s+2k-2}{2k-1} \frac{B_{2k}}{2kN^{s-1+2k}}, \quad (24)$$

where B_{2k} are the even Bernoulli numbers

$$B_{2k} = \frac{(-1)^{k-1} 2(2k)! \zeta(2k)}{(2\pi)^{2k}}. \quad (25)$$

Since the zeta function evaluations in (25) are for positive even integer arguments, they can be calculated quickly using (22). Once the requisite even Bernoulli numbers B_{2k} (up to index k approximately matching the precision level in digits) are computed, the function that implements formula (24) is thread-safe.

7.4 FFT-based multiplication

Although the multiplication algorithm described above is very efficient, for higher levels of precision (above 700 words, or approximately 12,600 digits, based on the present author's implementation), significantly faster performance can be achieved by employing an FFT-convolution approach [15][12, pg. 223–224].

Suppose one wishes to multiply two n -precision values whose mantissa words are given by $a = (a_0, a_1, a_2, \dots, a_{n-1})$ and $b = (b_0, b_1, b_2, \dots, b_{n-1})$. It is easy to see that the desired result, except for releasing carries, is an acyclic convolution. In particular, assume that a and b are extended to $2n$ words each by padding with zeroes. Then the product $c = (c_k)$ is given by

$$c_k = \sum_{j=0}^{2n-1} a_j b_{k-j}, \quad 0 \leq k < 2n, \quad (26)$$

where b_{k-j} is read as b_{k-j+2n} when $k-j$ is negative. This convolution can be calculated as

$$(c) = F^{-1}[F(a) \cdot F(b)], \quad (27)$$

where $F(a)$ and $F(b)$ denote a real-to-complex discrete Fourier transform (computed using an FFT algorithm), the dot means element-by-element complex multiplication, and $F^{-1}[\]$ means an inverse complex-to-real FFT. The c_k results from this process are floating-point numbers. Rounding these values to the nearest integer, and then releasing carries beginning at c_{2n-1} gives the desired multiplication result.

One important detail was omitted above: to avoid loss of numerical significance, the 60-bit mantissa words of the input multiprecision values are first divided into 15-bit sections. Then the FFT-convolution procedure is performed, and the result is suitably reconstituted to an output multiprecision value of 60-bit words at the end.

In contrast to the basic arithmetic algorithms, FFT-based multiplication requires precomputed FFT root-of-unity data. Thus, if one requires a precision level greater than 12,600 digits, one must call `mpinit` at the start of the user's program, in a single-threaded initialization section, before any parallel execution — see Table 4 and Section 5.4 for details.

7.5 Advanced algorithm for division

With an FFT-based multiplication facility in hand, division of two extra-high-precision arguments a and b can be performed by the following scheme. This Newton-Raphson algorithm iteration converges to $1/b$ [12, pg. 226]:

$$x_{k+1} = x_k + (1 - x_k \cdot b) \cdot x_k, \quad (28)$$

where the multiplication $() \cdot x_k$ is performed with only half of the normal level of precision. These iterations are performed with a working precision level that is approximately doubles with each iteration, except that at three iterations before the final iteration, the iteration is repeated without doubling the precision, in order to enhance accuracy. The final iteration is performed as follows (due to A. Karp):

$$a/b \approx (a \cdot x_n) + [a - (a \cdot x_n) \cdot b] \cdot x_n, \quad (29)$$

where the multiplications $a \cdot x_n$ and $[\] \cdot x_n$ are performed with only half of the final level of precision. The total cost of this procedure is only about three times the cost of a single full-precision multiplication. It is invoked if the working precision exceeds 12,600 digits, as with the FFT-based multiplication routine.

References

- [1] David H. Bailey, “MPFUN2015: A thread-safe arbitrary precision package (full documentation),” manuscript, 18 Nov 2020, <https://www.davidhbailey.com/dhbpapers/mpfun2015.pdf>
- [2] 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.
- [3] 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.
- [4] D. H. Bailey and J. M. Borwein, “High-precision arithmetic in mathematical physics,” *Mathematics*, vol. 3 (2015), pg. 337–367. <http://www.mdpi.com/2227-7390/3/2/337/pdf>.
- [5] D. H. Bailey and J. M. Borwein, “Hand-to-hand combat with thousand-digit integrals,” *Journal of Computational Science*, vol. 3 (2012), pg. 77-86.
- [6] 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>.
- [7] D. H. Bailey and J. M. Borwein, “Compressed lattice sums arising from the Poisson equation: Dedicated to Professor Hari Sirvastava,” *Boundary Value Problems*, vol. 75 (2013), DOI: 10.1186/1687-2770-2013-75, <http://www.boundaryvalueproblems.com/content/2013/1/75>.
- [8] 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.
- [9] 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 (2016), pg. 349-363, preprint at <https://www.davidhbailey.com/dhbpapers/poisson-res.pdf>.

- [10] D. H. Bailey and D. J. Broadhurst, “Parallel integer relation detection: Techniques and applications,” *Mathematics of Computation*, vol. 70, no. 236 (Oct 2000), pg. 1719–1736.
- [11] D. H. Bailey, X. S. Li and B. Thompson, “ARPREC: An arbitrary precision computation package,” Sep 2002, <http://crd.lbl.gov/~dhbailey/dhbpapers/arprec.pdf>.
- [12] J. M. Borwein and D. H. Bailey, *Mathematics by Experiment: Plausible Reasoning in the 21st Century*, 2nd ed., A. K. Peters, Natick, MA, 2008.
- [13] P. Borwein, “An efficient algorithm for the Riemann zeta function,” 1995, <http://www.cecm.sfu.ca/~pborwein/PAPERS/P155.pdf>.
- [14] R. P. Brent, “Fast multiple-precision evaluation of elementary functions,” *J. of the ACM*, vol. 23 (1976), 242–251.
- [15] R. P. Brent and P. Zimmermann, *Modern Computer Arithmetic*, Cambridge Univ. Press, 2010.
- [16] Jingwei Chen, Damien Stehle and Gilles Villard, “A new view on HJLS and PSLQ: Sums and projections of lattices,” *Proc. of ISSAC’13*, 149–156.
- [17] “Comparison of multiple-precision floating-point software,” <http://www.mpfr.org/mpfr-3.1.0/timings.html>.
- [18] “Digital library of mathematical functions,” National Institute of Standards and Technology, 2015, <http://dlmf.nist.gov>.
- [19] 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), pg. 351–369.
- [20] “Fortran 2008,” <http://fortranwiki.org/fortran/show/Fortran+2008>.
- [21] “The GNU MPFR library,” <http://www.mpfr.org>.
- [22] “GNU MPFR library: Comparison of multiple-precision floating-point software,” <http://www.mpfr.org/mpfr-current/timings.html>.

- [23] Y. Hida, X. S. Li and D. H. Bailey, “Algorithms for Quad-Double Precision Floating Point Arithmetic,” *Proc. of the 15th IEEE Symposium on Computer Arithmetic* (ARITH-15), 2001.
- [24] A. K. Lenstra, H. W. Lenstra, and L. Lovasz, “Factoring polynomials with rational coefficients,” *Mathematische Annalen*, vol. 261 (1982), pg. 515-534.
- [25] “MPFR C++,” <http://www.holoborodko.com/pavel/mpfr>.
- [26] *NIST Digital Library of Mathematical Functions*, version 1.0.6 (May 2013), <http://dlmf.nist.gov>.
- [27] R. W. Potter, *Arbitrary Precision Calculation of Selected Higher Functions*, Lulu.com, San Bernardino, CA, 2014.
- [28] C. Rubio-Gonzalez, C. Nguyen, H. D. Nguyen, J. Demmel, W. Kahan, K. Sen, D. H. Bailey, C. Iancu and D. Hough, “Precimonious: Tuning assistant for floating-point precision,” *Proceedings of SC13*, 26 Apr 2013, <http://www.davidhbailey.com/dhbpapers/precimonious.pdf>.
- [29] D. W. Sweeney, “On the computation of Euler’s constant”, *Mathematics of Computation*, vol. 17 (1963), pg. 170–178.
- [30] H. Takahasi and M. Mori, “Double exponential formulas for numerical integration,” *Publications of RIMS*, Kyoto University, vol. 9 (1974), pg. 721-741.
- [31] S. W. Williams and D. H. Bailey, “Parallel computer architecture,” in David H. Bailey, Robert F. Lucas and Samuel W. Williams, ed., *Performance Tuning of Scientific Applications*, CRC Press, Boca Raton, FL, 2011, 11–33.