

MPFUN2020: A thread-safe arbitrary precision package with special functions (Full Documentation)

David H. Bailey*

May 30, 2024

Abstract

Numerous research studies have arisen, particularly in mathematical physics and experimental mathematics, that require extremely high numeric precision — hundreds or thousands of digits in some cases. This paper describes a new arbitrary precision software package (“MPFUN2020”) to support such applications. It comes in two versions: a self-contained all-Fortran version that is easy to install, and a version based on the MPFR package that runs faster although it is significantly more complicated to install.

Both 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 only minor changes to source code; (c) full support for both real and complex datatypes; (d) an extensive set of transcendental and special functions; (e) run-time checking to guard against problems with converting double precision data; (f) a medium precision datatype that improves performance on large variable precision applications; and (g) interoperability, so that with a simple restriction, application codes written for one version can be run using the other without any code changes.

*Lawrence Berkeley National Laboratory (retired), 1 Cyclotron Road, Berkeley, CA 94720, USA. E-mail: dhbailey@lbl.gov.

Contents

1	Applications of high-precision computation	3
1.1	The PSLQ integer relation algorithm	3
1.2	High-precision numerical integration	4
1.3	Ising integrals	4
1.4	Algebraic numbers in Poisson potential functions	6
2	Current high-precision software	8
2.1	Thread safety and parallel implementations	10
3	MPFUN2020: A new thread-safe package	11
3.1	Data structure	11
3.2	Modules	12
3.3	The MPFUN2020 solution to thread safety	13
4	Installation, compilation and linking	14
4.1	The two variants	14
5	Coding Fortran applications	15
5.1	Basic instructions for using the package	15
5.2	Extra-high precision	16
5.3	Functions and subroutines	18
5.4	Special functions	18
5.5	Polynomial evaluation and solution	19
5.6	Input and output of multiprecision data	19
5.7	Handling double precision values	23
5.8	Support for quad precision	26
5.9	Dynamically changing the working precision	27
5.10	Medium precision datatype	30
6	Sample applications and performance	31
6.1	Timings	34
7	Appendix: Numerical algorithms	36
7.1	Algorithms for basic arithmetic	36
7.2	FFT-based multiplication	37
7.3	Transcendental constants and functions	38
7.4	Special functions	41

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.

Another group of applications, particularly in the fields of mathematical physics and experimental mathematics, require even higher precision — tens, hundreds or even thousands of digits:

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 [7, 5], which provides detailed references. Here is a brief overview of several key 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 [7]. Many of these computations involve variants of Ferguson’s PSLQ integer relation detection algorithm [25, 15]. 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 attempts to find nontrivial integer coefficients (a_i) , if they exist, such that

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

to available precision. Some researchers adapt variations of the Lenstra-Lenstra-Lovasz (LLL) lattice basis reduction algorithm [32], or the “HJLS”

algorithm, which is based on LLL, for integer relation detection [21]. Whichever algorithm is used, integer relation detection requires very high precision—at least $n \cdot \max_i \log_{10} |a_i|$ digits, or else the true relation will be lost in a sea of numerical artifacts.

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 [38, 4]. 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 and negative) are smaller than the “epsilon” of the numeric precision being used. Full details are given in [4]. An overview of applications of high-precision integration in experimental mathematics is given in [9].

1.3 Ising integrals

In one study, tanh-sinh quadrature and PSLQ were employed to study the following classes of integrals [11]. 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* (see Section 7.4). 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 (an online tool that unfortunately is no longer available), as

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

where γ is Euler's constant. This identity was then proven in [11].

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 *polylogarithm function* (see Section 7.4).

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 [35].

1.4 Algebraic numbers in Poisson potential functions

The Poisson potential function arises in several contexts of mathematical physics and engineering, including the analysis of crystal structures and even the sharpening of photographic images. In two recent studies [13, 12], researchers explored a simple two-dimensional instance:

$$\phi_2(x, y) = \frac{1}{\pi^2} \sum_{m, n \text{ odd}} \frac{\cos(m\pi x) \cos(n\pi y)}{m^2 + n^2}. \quad (3)$$

After detailed numerical experimentation, these researchers discovered and then proved the remarkable fact that when x and y are rational,

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

where $\beta(x, y)$ is an *algebraic number*, namely the root of an algebraic equation with integer coefficients of some degree m .

These experiments computed $\alpha = (\beta(x, y))^8 = \exp(8\pi\phi_2(x, y))$, using some rapidly convergent formulas found in [13], for various simple rationals x and y . Then the vector $(1, \alpha, \alpha^2, \dots, \alpha^m)$ was input to a program implementing the three-level multipair PSLQ program [15]. When successful, the program returned the vector of integer coefficients $(a_0, a_1, a_2, \dots, a_m)$ of a polynomial satisfied by α as output. Table 1 shows some examples [13].

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 [13].

Using this and other data, Jason Kimberley of the University of Newcastle, Australia, conjectured a formula that gives the degree m as a function of s [13]. These computations required prodigiously high precision (up to 10,000 digits) and very long run times.

In a subsequent study [14], the present author and three other researchers extended these studies, confirming Kimberley's formula for all $10 \leq s \leq 40$ and also for $s = 42, 44, 46, 48, 50, 52, 60$ and 64 . These runs, which required precision levels up to 64,000 digits, were facilitated by an improved PSLQ code, the faster MPFUN2015 package described below in Section 2, as well as parallel execution facilitated by the thread-safe feature of this package. Examination of these results led to additional discoveries, and ultimately to a full proof of Kimberley's conjecture, at least for the special case $x = y = 1/s$ for integer s , as well as other facts about these curious polynomials, such as the fact that when s is even, the polynomials for the case $x = y = 1/s$ are palindromic (left-to-right symmetric) [14].

In two more recent studies [2, 3], these results were further extended, using the arbitrary precision package described in this paper, to include the vastly

larger class of rationals $x = p/s$ and $y = q/s$, for $1 \leq p, q < s/2 \leq 38$, and also for $s = 40, 42$ and 50 . These computations revealed that Kimberley's formula does *not* give the correct degree for the larger class, but that a generalization of Kimberley's rule does hold. Two other new general features of these polynomials were also identified. See [3] for these results and additional details.

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) [30].

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 binary mantissa.

Software for performing high-precision arithmetic has been available for quite some time, for example in the commercial packages *Mathematica* and *Maple*, but until recently those using standard languages for high-performance computing, such as C, C++ or Fortran, often found necessary to rewrite their applications, replacing each arithmetic operation with a subroutine call to a custom-coded library. Today there are several freely available high-precision software packages for high-precision computing, and even some high-level language interfaces that make code conversions relatively painless:

- ARPREC: [By the present author; superseded by MPFUN2015 and MPFUN2020] 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; achieves 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: [By the present author] 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>.
- MPFUN2015: [By the present author; superseded by MPFUN2020] A thread-safe arbitrary precision package for Fortran application, with an all-Fortran version and a version based on MPFR for faster performance; available at <https://www.davidhbailey.com/dhbsoftware/>.

2.1 Thread safety and parallel implementations

The scientific computing world is moving rapidly into multicore and multi-node parallel computing, so that future improvements in performance on high-precision computations will only be obtained by aggressively exploiting parallelism. On modern multicore processors, parallel computing is most efficiently performed using a shared memory, multithreaded environment such as OpenMP [39] 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. More importantly, most existing high-precision packages generate a “context” of auxiliary data, such as the current working precision level and data to support transcendental function evaluation, which context typically ruins thread safety unless special care is taken.

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 [33]. This package is built upon the lower-level MPFR package [27], 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 [22]. (However, the MPFR package is only thread-safe when a special option is enabled in the installation step.)

There was, to this author’s knowledge, no high-level, thread-safe arbitrary precision package to support Fortran applications, prior to MPFUN2015, which was mentioned in the above list. However, the present author will be the first to acknowledge that MPFUN2015 has some shortcomings, notably that the run-time performance of the all-Fortran version is typically five times or more slower than the version based on MPFR. Another weakness is that in both versions of the package, significant memory is wasted for applications where part of multiprecision arrays require only a medium level of precision. Further, the MPFUN2015 package lacks a full package of special functions and polynomial solutions routines.

3 MPFUN2020: A new thread-safe package

To that end, the present author has written a new thread-safe arbitrary precision package. As with the MPFUN2015 package, it is offered in two versions: an easy-to-install all-Fortran version (MPFUN20-Fort), which runs several times faster than the all-Fortran version of MPFUN2015, and also a version based on the MPFR package (MPFUN20-MPFR), which is even faster and more complete than the MPFR version of MPFUN2015.

Both 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 only minor changes to source code; (c) full support for both real and complex datatypes; (d) an extensive set of transcendental and special functions; (e) run-time checking to guard against problems with converting double precision data; (f) a medium precision datatype that improves performance on large variable precision applications; and (g) interoperability, so that with a simple restriction, application codes written for one version can be run using the other without any code changes.

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-bit integer operations, rather than on 64-bit floating-point operations as in the earlier package. Additionally, the MPFUN20-Fort version employs FFT-based arithmetic for significantly faster execution at very high precision levels — see Sections 7.2. The new MPFUN20-MPFR version is faster than the MPFUN20-Fort version on most applications, although it is significantly more complicated to install, because the GMP and MPFR packages must first be installed, usually requiring administrator privilege.

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

1. MPFUNA: Contains compile-time global data, including the binary values of $\log(2)$, π and γ (Euler's constant), to 20,000-digit precision.
2. MPFUNB (MPFUN20-Fort only): Handles basic arithmetic functions, rounding, normalization, square roots and n -th roots.
3. MPFUNC (MPFUN20-Fort only): Handles binary-to-decimal conversion, decimal-to-binary conversion and input/output operations.
4. MPFUND: Includes routines for common transcendental constants and functions (see Table 3 and Section 7.3).
5. MPFUNE: Includes routines for special functions, including Bessel functions, the error function, the gamma function, the zeta function and numerous others (see Table 5 and Section 7.4).
6. MPFUNF: Defines the default standard precision and medium precision levels, in digits, and also the equivalent levels in words.
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 (over 20,000 digits). The working precision level is always passed as a subroutine argument, ensuring thread safety. The MPFUN20-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 slightly different programming style for some applications, as we shall explain below (see Section 5.9).

All computations are performed to a fixed precision set by the user in module MPFUNF (default is 4,000 digits), unless the user specifies a lower value for certain calculations (see Section 5.9 for details) or otherwise employs a variable precision design. 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 double precision constant or variable to a multiprecision variable or array element, 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, such as GNU gfortran, Intel's ifort or NAG's nagfor. Full details on installing the MPFUN2020 software (and also for installing gfortran, if needed) are available in the `README` file for the distribution, which is in the main directory of the distribution, or at <https://www.davidhbhailey.com/dhbsoftware>.

4.1 The two variants

Each of the two versions of the software comes in two variants:

- Variant 1: This is recommended for beginning users and for basic applications that do not dynamically change the working precision level (or do so only rarely).
- Variant 2: This is recommended for more sophisticated applications that dynamically change the working precision level. It does not allow some mixed-mode combinations, and requires one to explicitly specify

a working precision parameter for some functions. However, in the present author's experience, these restrictions result in less overall effort to produce a debugged, efficient application code. See Section 5.9.

The Fortran source files and scripts required for each of these variants are in the respective directories `fortran-var1` and `fortran-var2`.

Compile/link scripts are available in the `fortran-var1` and `fortran-var2` directories for the gfortran and Intel ifort compilers, and, with MPFUN20-Fort, for the NAG nagfor compiler. For example, to compile variant 1 of either the MPFUN20-Fort or MPFUN20-MPFR library using the GNU gfortran compiler, go to the `fortran-var1` 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, together with reference output files, are included in the `fortran-var1` and `fortran-var2` directories — see Section 6.

5 Coding Fortran applications

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

5.1 Basic instructions for using the package

To use either MPFUN20-Fort or MPFUN20-MPFR 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; this 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-var1` or `fortran-var2` directory of the software. In the code as distributed, `mpipl` is set to 4,000 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

```
mpwds = int (mpipl / mpdpw + 2)
```

where `mpdpw` is a double precision system parameter set in file `mpfuna.f90` (in MPFUN20-Fort, $\text{mpdpw} = \log_{10} 2^{60} = 18.0617997\dots$; in MPFUN20-MPFR, $\text{mpdpw} = \log_{10} 2^{64} = 19.2659197\dots$). 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 5.7 below about DP and DC constants and expressions.

Input and output of MP variables or array elements are done by using the special subroutines `mpread` and `mpwrite`. See Table 4 and Section 5.6.

5.2 Extra-high precision

The above instructions apply if the default precision level, namely `mpipl`, is less than or equal to 20,000 digits. If one requires a precision level greater

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

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 — see Section 5.9.

than this, then in addition to changing the value of `mpipl` (and, thus, `mpwds`) 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 (nwds)
```

where `nwds` is the requested precision level in words (normally `nwds = mpwds`). A call to `mpinit` initializes FFT tables (in MPFUN20-Fort) and calculates $\log(2)$, π and γ . If the requested precision (`mpwds` or `nwds`) is less than `mpl2pi`, a call to `mpinit` has no effect.

It should be noted, however, that some special functions internally require $\log(2)$, π and γ to higher precision. If one sees a run-time error message such as “Pi must be precomputed to precision 1200 words,” insert a call to `mpinit` at start of execution with the specified precision as the argument.

5.3 Functions and subroutines

As mentioned above, the MPFUN20-Fort and MPFUN20-MPFR packages support most Fortran-2008 intrinsics, including all the well-known transcendentals (e.g., `sin`, `exp`, `log`, etc.), and, in addition, an extensive set of special functions and a set of routines for polynomial solution. Further, the package includes a set of I/O and conversion functions, such as functions to convert between double precision and multiprecision real or between double complex and multiprecision complex. These functions and subroutines are listed below in tables 3, 4, 5, 6 and 7 (some are listed in more than one table).

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 “QP” denotes IEEE quad precision (if supported by the compiler). The variable names `r1`, `r2`, `r3` are MPR; `z1` is MPC; `d1` is DP; `dc1` is DC; `q1` is QP; `i1`, `i2`, `i3`, `n`, `nb`, `np`, `nq` are integers; `s1` is character(1); `sn` is character(`n`); `rb` is an MPR vector of length `nb`; `ss` is an MPR vector of length `n`; `aa` is an MPR vector of length `np`; and `bb` is an MPR vector of length `nq`.

5.4 Special functions

As briefly mentioned above, the package includes numerous special functions: the BesselI, BesselJ, BesselK and BesselY functions (each with integer

or MPR order), the digamma function, the erf and erfc functions, the exponential integral (expint) function, the gamma function, the Hurwitz zeta function, the hypergeometric pFq function, the incomplete gamma function, the polygamma function, the polylogarithm function, the StruveH function and the zeta function.

A complete listing of these functions and their arguments is presented in Tables 5 and 6. The algorithms employed in these routines are described in the Appendix (Section 7.4).

5.5 Polynomial evaluation and solution

As briefly mentioned above, the package includes several routines for polynomial evaluation and solution. Two routines solve an input MPR polynomial, given an MPR or MPC initial approximation to the root, by means of Newton iterations with a dynamically varying level of precision. A third routine returns the greatest common divisor (GCD) of two input MPR polynomials. A fourth routine divides two MPR polynomials, producing the quotient and remainder polynomials. Two other routines evaluate a polynomial at an MPR or MPC argument.

Polynomials with repeated roots often defeat Newton iteration schemes for solution. Such polynomials can be handled as follows: first apply the polynomial GCD routine to the input polynomial and its derivative, and then divide the input polynomial by the GCD polynomial. This produces a reduced polynomial that has all the roots of the original polynomial, but no repeated roots. Then the Newton iteration-based solution routines can quickly produce high-precision roots.

Table 7 lists the polynomial evaluation and solution routines. Usage of these routines is demonstrated by the `tpolysolve` sample program; see Section 6.

5.6 Input and output of multiprecision data

As briefly mentioned above in Table 4, the package includes routines for binary-decimal conversion and input or output of multiprecision data. These routines and their usage are described in Table 8.

Type	Name	Description
F:MPR	<code>abs(r1)</code>	Absolute value
F:MPR	<code>abs(z1)</code>	Absolute value of complex arg
F:MPR	<code>acos(r1)</code>	Inverse cosine
F:MPR	<code>acosh(r1)</code>	Inverse hyperbolic cosine
F:MPR	<code>aimag(z1)</code>	Imaginary part of complex arg
F:MPR	<code>aint(r1)</code>	Truncates to integer
F:MPR	<code>anint(r1)</code>	Rounds to closest integer
F:MPR	<code>asin(r1)</code>	Inverse sine
F:MPR	<code>asinh(r1)</code>	Inverse hyperbolic sine
F:MPR	<code>atan(r1)</code>	Inverse tangent
F:MPR	<code>atan2(r1,r2)</code>	Arctangent with two args
F:MPR	<code>atanh(r1)</code>	Inverse hyperbolic tangent
F:MPR	<code>bessel_j0(r1)</code>	BesselJ function, order 0
F:MPR	<code>bessel_j1(r1)</code>	BesselJ function, order 1
F:MPR	<code>bessel_jn(n,r1)</code>	BesselJ function, order n
F:MPR	<code>bessel_y0(r1)</code>	BesselY function, order 0
F:MPR	<code>bessel_y1(r1)</code>	BesselY function, order 1
F:MPR	<code>bessel_yn(n,r1)</code>	BesselY function, order n
F:MPC	<code>conjg(z1)</code>	Complex conjugate
F:MPR	<code>cos(r1)</code>	Cosine of real arg
F:MPC	<code>cos(z1)</code>	Cosine of complex arg
F:MPR	<code>cosh(r1)</code>	Hyperbolic cosine
F:DP	<code>dbl(r1)</code>	Converts MPR argument to DP
F:DC	<code>dcmplx(z1)</code>	Converts MPC argument to DC
F:MPR	<code>erf(r1)</code>	Error function
F:MPR	<code>erfc(r1)</code>	Complementary error function
F:MPR	<code>exp(r1)</code>	Exponential function of real arg
F:MPC	<code>exp(z1)</code>	Exponential function of complex arg
F:MPR	<code>gamma(r1)</code>	Gamma function
F:MPR	<code>hypot(r1,r2)</code>	Hypotenuse of two args
F:MPR	<code>log(r1)</code>	Natural logarithm of real arg
F:MPC	<code>log(z1)</code>	Natural logarithm of complex arg
F:MPR	<code>log10(r1)</code>	Base-10 logarithm of real arg
F:MPR	<code>max(r1,r2)</code>	Maximum of two (or three) args
F:MPR	<code>min(r1,r2)</code>	Minimum of two (or three) args
F:MPR	<code>mod(r1,r2)</code>	Mod function = $r1 - r2 * \text{aint}(r1/r2)$
F:MPR	<code>sign(r1,r2)</code>	Transfers sign from r2 to r1
F:MPR	<code>sin(r1)</code>	Sine function of real arg
F:MPC	<code>sin(z1)</code>	Sine function of complex arg
F:MPR	<code>sinh(r1)</code>	Hyperbolic sine
F:MPR	<code>sqrt(r1)</code>	Square root of real arg
F:MPC	<code>sqrt(z1)</code>	Square root of complex arg
F:MPR	<code>tan(r1)</code>	Tangent function
F:MPR	<code>tanh(r1)</code>	Hyperbolic tangent function

Table 3: Fortran-2008 intrinsic functions extended to multiprecision.

Type	Name	Description
F:MPC	<code>mpcplx(r1,r2)</code>	Converts <code>(r1,r2)</code> 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 <code>r1</code> , in the same time as calling just cosh or just sinh
S	<code>mpcssn(r1,r2,r3)</code>	Returns both cos and sin of <code>r1</code> , in the same time as calling just cos or just sin
S	<code>mpdecmd(r1,d1,i1)</code>	Converts <code>r1</code> to the form <code>d1*10ⁱ¹</code>
S	<code>mpeform(r1,i1,i2,s1)</code>	Converts <code>r1</code> to char(1) string in Ei1.i2 format, suitable for output (Sec. 5.6)
S	<code>mpfform(r1,i1,i2,s1)</code>	Converts <code>r1</code> to char(1) string in Fi1.i2 format, suitable for output (Sec. 5.6)
F:MPR	<code>mpegamma()</code>	Returns Euler's γ constant [1]
S	<code>mpinit()</code>	Initializes for extra-high precision (Sec. 5.2) [1]
F:MPR	<code>mplog2()</code>	Returns $\log(2)$ [1]
F:MPR	<code>mpnrt(r1,i1)</code>	Returns the <code>i1</code> -th root of <code>r1</code>
F:MPR	<code>mppi()</code>	Returns π [1]
F:MPR	<code>mpprod(r1,d1)</code>	Returns <code>r1*d1</code> , without checking [2]
F:MPR	<code>mpquot(r1,d1)</code>	Returns <code>r1/d1</code> , without checking [2]
F:MPR	<code>mprand(r1)</code>	Returns pseudorandom number, based on <code>r1</code> Start with an irrational, say <code>r1 = mplog2()</code> Typical iterated usage: <code>r1 = mprand(r1)</code>
S	<code>mpread(i1,r1)</code>	Inputs <code>r1</code> from Fortran unit <code>i1</code> ; up to five MPR args may be listed (Sec. 5.6) [1]
S	<code>mpread(i1,z1)</code>	Inputs <code>z1</code> from Fortran unit <code>i1</code> ; up to five MPC args may be listed (Sec. 5.6) [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]
F:MPR	<code>mpreal(s1,i1)</code>	Converts char(1) string of length <code>i1</code> to MPR (Sec. 5.6) [1]
F:MPR	<code>mpreal(sn)</code>	Converts char(n) string to MPR (Sec. 5.6) [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 <code>r1</code>
F:Int	<code>mpwprec(z1)</code>	Returns precision in words assigned to <code>z1</code>
S	<code>mpwrite(i1,i2,i3,r1)</code>	Outputs <code>r1</code> in Ei2.i3 format to unit <code>i1</code> ; up to five MPR args may be listed (Sec. 5.6)
S	<code>mpwrite(i1,i2,i3,z1)</code>	Outputs <code>z1</code> in Ei2.i3 format to unit <code>i1</code> ; up to five MPC args may be listed (Sec. 5.6)

Table 4: Additional general routines (F: function, S: subroutine). Notes:
[1]: In variant 1, an integer precision level argument (mantissa words) may optionally be added as the final argument; this argument is required in variant 2. See Sec. 5.9.
[2]: These do not check DP or DC values. See Sec. 5.7.

Type	Name	Description
F:MPR	<code>agm(r1,r2)</code>	Arithmetic-geometric mean
F:MPR	<code>airy(r1)</code>	Airy function [1]
S	<code>mpberne(nb,rb)</code>	Initialize array <code>rb</code> of length <code>nb</code> with first <code>nb</code> even Bernoulli numbers [2][3]
F:MPR	<code>bessel_i(r1,r2)</code>	BesselI function, order <code>r1</code> , of <code>r2</code>
F:MPR	<code>bessel_in(n,r1)</code>	BesselI function, order <code>n</code> , of <code>r1</code>
F:MPR	<code>bessel_j(r1,r2)</code>	BesselJ function, order <code>r1</code> , of <code>r2</code>
F:MPR	<code>bessel_j0(r1)</code>	BesselJ function, order 0, of <code>r1</code>
F:MPR	<code>bessel_j1(r1)</code>	BesselJ function, order 1, of <code>r1</code>
F:MPR	<code>bessel_jn(n,r1)</code>	BesselJ function, order <code>n</code> , of <code>r1</code>
F:MPR	<code>bessel_k(r1,r2)</code>	BesselK function, order <code>r1</code> , of <code>r2</code>
F:MPR	<code>bessel_kn(n,r1)</code>	BesselK function, order <code>n</code> , of <code>r1</code>
F:MPR	<code>bessel_y(r1,r2)</code>	BesselY function, order <code>r1</code> , of <code>r2</code>
F:MPR	<code>bessel_y0(r1)</code>	BesselY function, order 0, of <code>r1</code>
F:MPR	<code>bessel_y1(r1)</code>	BesselY function, order 1, of <code>r1</code>
F:MPR	<code>bessel_yn(n,r1)</code>	BesselY function, order <code>n</code> , of <code>r1</code>
F:MPR	<code>digamma(r1)</code>	Digamma function of <code>r1</code> [1]
F:MPR	<code>digamma.be(nb,rb,r1)</code>	Digamma function of <code>r1</code> , using <code>nb</code> even Bernoulli numbers in <code>rb</code> [3]
F:MPR	<code>erf(r1)</code>	Error function
F:MPR	<code>erfc(r1)</code>	Complementary error function
F:MPR	<code>expint(r1)</code>	Exponential integral function
F:MPR	<code>gamma(r1)</code>	Gamma function
F:MPR	<code>hurwitz_zetan(n,r1)</code>	Hurwitz zeta function, order $n \geq 2$, of <code>r1</code> [4]
F:MPR	<code>hurwitz_zetan.be(nb,rb,n,r1)</code>	Hurwitz zeta function, order $n \geq 2$, of <code>r1</code> , using <code>nb</code> even Bernoulli numbers in <code>rb</code> [3]
F:MPR	<code>hypergeom_pfq(np,nq,aa,bb,r1)</code>	Hypergeometric pFq function of <code>aa</code> , <code>bb</code> and <code>r1</code> ; dimensions are <code>aa(np)</code> and <code>bb(nq)</code> [5]
F:MPR	<code>incgamma(r1,r2)</code>	Incomplete gamma function [6]
F:MPR	<code>polygamma(n,r1)</code>	Polygamma function, order $n \geq 1$, of <code>r1</code> [4]
F:MPR	<code>polygamma.be(nb,rb,n,r1)</code>	Polygamma function, order $n \geq 1$, of <code>r1</code> , using <code>nb</code> even Bernoulli numbers in <code>rb</code> [3]

Table 5: Special functions (F: function, S: subroutine) (continued in Table 6). Notes:
[1]: Only available with MPFUN20-MPFR.
[2]: In variant 1, an integer precision level argument (mantissa words) may optionally be added as the final argument; this argument is required in variant 2. See Sec. 5.9.
[3]: For most applications, set `nb` > 2X precision in decimal digits; see `mpberne` above.
[4]: For `hurwitz_zetan` and `polygamma`, the argument `r1` is limited to the range (0, 1].
[5]: For `hypergeom_pfq`, the integers `np` and `nq` must not exceed 10.
[6]: For `incgamma`, `r1` must not be zero, and must not be negative unless `r2` = 0.

S	<code>polylog_ini(n,ss)</code>	Initialize array <code>ss</code> , of size <code> n </code> , for computing polylogarithms of order <code>n</code> when <code>n < 0</code> [2][6]
F:MPR	<code>polylog_neg(n,ss,r1)</code>	Polylogarithm function of <code>r1</code> , for <code>n < 0</code> , using precomputed data in <code>ss</code> [7]
F:MPR	<code>polylog_pos(n,r1)</code>	Polylogarithm function, order <code>n ≥ 0</code> , of <code>r1</code> [8]
F:MPR	<code>struve_hn(n,r1)</code>	StruveH function, order <code>n ≥ 0</code> , of <code>r1</code> [9]
F:MPR	<code>zeta(r1)</code>	Zeta function of <code>r1</code>
F:MPR	<code>zeta_be(nb,rb,r1)</code>	Zeta function of <code>r1</code> , using <code>nb</code> even Bernoulli numbers in <code>rb</code> [3]
F:MPR	<code>zeta_int(n)</code>	Zeta function of integer argument <code>n</code> [2]

Table 6: Special functions, continued (F: function, S: subroutine). Notes:
[2]: In variant 1, an integer precision level argument (mantissa words) may optionally be added as the final argument; this argument is required in variant 2. See Sec. 5.9.
[6]: For `incgamma`, `r1` must not be zero, and must not be negative unless `r2 = 0`.
[7]: For `polylog_ini` and `polylog_neg`, the integer `n` is limited to the range `[-1000, -1]`.
[8]: For `polylog_pos`, the argument `r1` is limited to the range `(-1, 1)`.
[9]: For `struve_hn`, the argument `r1` is limited to the range `[-1000, 1000]`.

5.7 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., 3., 12345., 2.5, 6.125), which are entirely safe in a multiprecision application, should be allowed.

Users should be aware, however, that there are some hazards in this type of programming, inherent in conventions adopted by all Fortran compilers and most other languages as well. For example, the code

```
r1 = 3.14159d0
```

where `r1` is MPR, does NOT produce the true multiprecision equivalent of 3.14159. In fact, the above line of code will result in a run-time error with either version of the MPFUN2020 software. To obtain the full MPR converted value, write this as

```
r1 = '3.14159'
```

or, if using variant 2, as

```
r1 = mpreal ('3.14159', nwd)
```

S	<code>mprpolysolve (nq, polyq, rootapp, root)</code> Given an input MPR polynomial with coefficients <code>polyq(0:nq)</code> (degree <code>nq</code>), and an input MPR approximation in <code>rootapp</code> , this returns the full precision MPR root in <code>root</code> ; if no root is found, MPR zero is returned in <code>root</code> .
S	<code>mpcpolysolve (nq, polyq, rootapp, root)</code> Given an input MPR polynomial with coefficients <code>polyq(0:nq)</code> (degree <code>nq</code>), and an input MPC approximation in <code>rootapp</code> , this returns the full precision MPC root in <code>root</code> ; if no root is found, MPC zero is returned in <code>root</code> .
S	<code>mprpolygcd (n1, poly1, n2, poly2, ng, polyg)</code> Given two input MPR polynomials <code>poly1(0:n1)</code> (degree <code>n1</code>) and <code>poly2(0:n2)</code> (degree <code>n2</code>), with <code>n1 >= n2</code> , this returns the MPR greatest common divisor (GCD) polynomial in <code>polyg(0:ng)</code> (degree <code>ng</code>). <code>polyg</code> must be dimensioned <code>(0:n1)</code> in calling program.
S	<code>mprpolydiv (n1, poly1, n2, poly2, n3, poly3, n4, poly4)</code> Given two input MPR polynomials <code>poly1(0:n1)</code> (degree <code>n1</code>) and <code>poly2(0:n2)</code> (degree <code>n2</code>), with <code>n1 >= n2</code> , this returns the MPR polynomial quotient in <code>poly3(0:n3)</code> (degree <code>n3</code>) and the MPR remainder in <code>poly4(0:n4)</code> (degree <code>n4</code>). <code>poly3</code> and <code>poly4</code> must be dimensioned <code>(0:n1)</code> in calling program.
S	<code>mprpolyeval (np, poly, arg, result)</code> Given an input MPR polynomial with coefficients <code>poly(0:np)</code> (degree <code>np</code>), and an input MPR argument in <code>arg</code> , this returns the MPR result of polynomial evaluation in <code>result</code> .
S	<code>mcpolyeval (np, poly, arg, result)</code> Given an input MPR polynomial with coefficients <code>poly(0:np)</code> (degree <code>np</code>), and an input MPC argument in <code>arg</code> , this returns the MPC result of polynomial evaluation in <code>result</code> .

Table 7: Polynomial evaluation and solution routines

S	<p><code>mpeform (r1,i1,i2,s1)</code></p> <p>This converts the MPR number <code>r1</code> into decimal character form in the <code>character(1)</code> array <code>s1</code>. The argument <code>i1</code> (input) is the length of the output string, and <code>i2</code> (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 <code>i1</code> cells of <code>s1</code>. The condition <code>i1 >= i2 + 20</code> must hold.</p>
S	<p><code>mpfform (r1,i1,i2,s1)</code></p> <p>This converts the MPR number <code>r1</code> into decimal character form in the <code>character(1)</code> array <code>s1</code>. The argument <code>i1</code> (input) is the length of the output string, and <code>i2</code> (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 <code>i1</code> cells of <code>s1</code>. The condition <code>i1 >= i2 + 20</code> must hold.</p>
S	<p><code>mpread (i1,r1)</code></p> <p>This reads the MPR number <code>r1</code>, presumed in decimal format, from Fortran logical unit <code>i1</code>. The digits of <code>r1</code> 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 <code>e</code>, <code>d</code>, <code>E</code> or <code>D</code> followed by an integer exponent; but must <i>not</i> have embedded blanks. Up to five MPR arguments may be included in argument list. See the Note below on an additional precision argument.</p>
S	<p><code>mpread (i1,z1)</code></p> <p>This is the same as the previous item, except that the input argument <code>z1</code> is of type MPC (a pair of MPR). Up to five MPC arguments may be included in argument list. See the Note below on an additional precision argument.</p>
F:MPR	<p><code>mpreal (s1,i1)</code></p> <p>This converts the string <code>s1</code>, which is of type <code>character(1)</code> and length <code>i1</code>, to MPR. See <code>mpread</code> above for format. See the Note below on an additional precision argument.</p>
F:MPR	<p><code>mpreal (sn)</code></p> <p>This converts the string <code>sn</code>, which may be of type <code>character(n)</code> for any <code>n</code>, to MPR. See <code>mpread</code> above for format. On some systems, <code>n</code> may be limited, say to 2048; if this is a problem, use the previous item. See the Note below on an additional precision argument.</p>
S	<p><code>mpwrite (i1,i2,i3,r1)</code></p> <p>This writes the MPR number <code>r1</code> to Fortran logical unit <code>i1</code>, in a format analogous to Fortran E format, left-justified in the field. The argument <code>i2</code> (input) is the length of the output field, and <code>i3</code> (input) is the number of digits after the decimal point; <code>i3</code> must exceed <code>i2</code> by 20. Up to five MPR arguments may be included in argument list.</p>
S	<p><code>mpwrite (i1,i2,i3,z1)</code></p> <p>This is the same as the previous item, except that the argument <code>z1</code> is of type MPC (a pair of MPR). Up to five MPC arguments may be included in argument list.</p>
Note	<p>For the <code>mpread</code> and <code>mpreal</code> routines above, when using variant 1, an integer working precision level argument (mantissa words) may optionally be added as the final argument; this argument is required in variant 2. See Section 5.9.</p>

Table 8: Input/output and conversion routines

where `nwds` is the working precision (in words) to be assigned to `r1`. Similarly, the code

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

where `r1` and `r2` are MPR, does NOT produce the true multiprecision value one might expect, since the expression `3.d0 * sqrt (2.d0)` will be performed in double precision, according to Fortran precedence rules. In fact, the above line of code will result in a run-time error with either version of the MPFUN2020 software. To obtain the fully accurate result, write this as

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

or, if using variant 2, as

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

where `nwds` is the working precision in words.

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: `mpprod`, `mpquot`, `mpreald`, `mpcmlpxdc` (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 5.9); and the fourth converts a DC value to MPC (with an optional precision level parameter — see Section 5.9). These routines do *not* check the double precision argument to see if it has more than 40 significant bits.

5.8 Support for quad precision

Several processor/compiler systems now support IEEE quad (128-bit) floating-point arithmetic. If one's platform does support IEEE quad, limited support is provided in the MPFUN2020 package. To determine whether or not IEEE quad is supported on one's platform, compile and execute the simple program

Type	Name	Description
F:MPR	<code>mpprod(r1,q1)</code>	Returns $r1*q1$, without checking [2]
F:MPR	<code>mpquot(r1,q1)</code>	Returns $r1/q1$, without checking [2]
F:MPR	<code>mpreal(q1)</code>	Converts quad real to MPR [1]
F:MPR	<code>mprealq(d1)</code>	Converts quad to MPR, without checking [1][2]
F:Q	<code>qreal(r1)</code>	Converts MPR to quad.

Table 9: Support for IEEE quad precision (if available) (F: function, S: subroutine).
Notes:

[1]: In variant 1, an integer precision level argument (mantissa words) may optionally be added as the final argument; this argument is required in variant 2. See Sec. 5.9.

[2]: These do not check quad values. See Sec. 5.8.

```

program testq
integer, parameter:: mprknd2 = selected_real_kind (33, 4931)
real (mprknd2) q1, q2
q1 = real (-1.d0, kind = mprknd2)
q2 = acos (q1)
write (6, '(1p,d42.34)') q2
stop
end

```

If this code is successfully compiled and produces the output

```
3.1415926535897932384626433832795028D+00
```

(except possibly for the final digit), then one's system supports IEEE quad.

The MPFUN2020 function `mpreal` function accepts a quad argument for conversion to multiprecision, and the `qreal` function accepts a multiprecision argument for conversion to quad. Other routines available are `mpprod`, `mpquot` and `mprealq`. Note that the `mpreal` function checks the quad argument to see if it has more than 90 significant bits; if so, an error is flagged, as mentioned above for double precision conversions. To convert without checking, use `mprealq`. See Table 9.

When using variant 1, an integer working precision level argument (mantissa words) may optionally be added as the final argument for the functions `mpreal` and `mprealq`; this argument is required in variant 2. See Section 5.9.

5.9 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 two variants of the package, as mentioned above in Section 4:

- Variant 1: This is recommended for beginning users and for basic applications that do not dynamically change the working precision level (or do so only rarely).
- Variant 2: This is recommended for more sophisticated applications that dynamically change the working precision level. It does not allow some mixed-mode combinations, and requires one to explicitly specify a working precision parameter for some functions. However, in the present author's experience, these restrictions result in less overall effort to produce a debugged, efficient application code.

In particular, with variant 1:

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`, `mpreald`, `mprealq` and `zeta_int` 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 `mpi1` digits) is assumed.

In contrast, with variant 2:

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 `mpi1` 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 MPFUN20-MPFR) in this way, the user code remains completely 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 has been set to 2,500 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, but not with variant 2, because the assignment `x(k) = 1.d0` is not allowed. Furthermore, all operations are performed with the default precision level of 2,500 digits. So with variant 2, where one wishes to perform this loop with a precision level of, say, only 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, nwds)`,

the array element `x(1)` is assigned the value 1.0, with a working precision of `nwds` words (i.e., 500 digits). In the loop, when `k` is 2, `x(2)` also inherits the working precision level `nwds` 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 `nwds` 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. However, it is *highly recommended* 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, 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 also works with variant 1, 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.

5.10 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 MPFUN20-MPFR versions of the software. To use these datatypes, do the following:

First, set the medium precision level `mpi1m` in file `mpfun.f90`; by default it is set to 400 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_complexm` 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; type (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; type (mp_realm) b
a = mpreal (b)
```

Note however, that with either `mpreal` or `mprealm`, in variant 1, an integer working precision level argument (mantissa words) may optionally be added as the final argument; this argument is required in variant 2. See Section 5.9 for details.

Each of the intrinsic functions and subroutines listed in tables 2, 3, 4, 5, 6, 7, 8 and 9 is also supported with medium precision arguments, except for the following: `mpreald`, `mprealq`, `mpcplx`, `mpcplxq`, `mpegamma`, `mplog2`, `mppi` and `zeta_int`. The medium precision equivalents of these functions are: `mprealdm`, `mprealqm`, `mpcplx`, `mpcplxq`, `mpegammam`, `mplog2m`, `mppim` and `zeta_intm`, respectively. See Table 10. Note that the comment regarding the working precision argument in the previous paragraph applies to each of these functions.

Usage of the medium precision datatype is illustrated in the programs `tps1qm3.f90` and `tpphix3.f90` in the application programs in Section 6.

6 Sample applications and performance

The current release of the software includes a set of sample application programs in the `fortran-var1` and `fortran-var2` directories:

1. `testmpfun.f90`: An overall test suite: This briefly tests most individual MPFUN2020 operations and functions (including mixed mode

Type	Name	Description
F:MPRM	<code>mprealm(r1)</code>	Converts MPR to med prec real [1]
F:MPRM	<code>mprealm(d1)</code>	Converts DP to med prec real [1]
F:MPRM	<code>mprealdm(d1)</code>	Converts DP to med prec real, without checking [1][2]
F:MPCM	<code>mprealqm(q1)</code>	Converts quad to med prec real, without checking [1][2]
F:MPCM	<code>mpcmplxm(dc1)</code>	Converts DC to med prec complex [1]
F:MPCM	<code>mpcmplxdm(dc1)</code>	Converts DC to med prec complex, without checking [1][2]
F:MPRM	<code>mppegammam()</code>	Returns Euler's gamma constant to med prec real [1]
F:MPRM	<code>mplog2m()</code>	Returns $\log(2)$ to med prec real [1]
F:MPRM	<code>mp pim()</code>	Returns π to med prec real [1]
F:MPRM	<code>zeta_intm(n)</code>	Returns $\zeta(n)$ to med precision real [1]

Table 10: Routines to support medium precision datatype (F: function, S: subroutine), in addition to the functions listed in Tables 2, 3, 4, 5 and 9. Notes:

[1]: In variant 1, an integer precision level argument (mantissa words) may optionally be added as the final argument; this argument is required in variant 2. See Sec. 5.9.

[2]: These do not check double or quad values. See Sec. 5.7 and Sec. 5.8.

arithmetic, comparison operations, transcendental functions, special functions and polynomial routines), by comparing each result with reference results in the file `testmpfun.ref.txt`, which must be present in the same directory. This is not an exhaustive test of all possible scenarios, but it often detects bugs and compiler issues. Size of code: 1,093 lines. Precision: 500 digits.

2. `tpolysolve.f90`: A polynomial solution test program: This computes each of the 12 real and 26 complex roots of an input degree-40 real polynomial with large integer coefficients and two repeated roots. The computation includes a polynomial greatest common divisor (GCD) calculation to produce a reduced polynomial with no repeated roots; this reduced polynomial is then solved using a Newton iteration scheme with dynamically varying precision. This code employs the polynomial evaluation and solution routines described above in Table 7. Size of code: 274 lines. Precision level: 4,000 digits.
3. `tps1q1.f90`: A one-level standard PSLQ program: This finds the coefficients of the degree-30 polynomial satisfied by $3^{1/5} - 2^{1/6}$. Size of code: 866 lines. Precision level: 250 digits.
4. `tps1qm1.f90`: A one-level multipair PSLQ program: This finds the coefficients of the degree-30 polynomial satisfied by $3^{1/5} - 2^{1/6}$. Size of code: 946 lines. Precision level: 250 digits.

5. `tps1qm2.f90`: A two-level multipair PSLQ program: This finds the coefficients of the degree-56 polynomial satisfied by $3^{1/7} - 2^{1/8}$. Size of code: 1,796 lines. Precision level: 750 digits; switches frequently between multiprecision and double precision.
6. `tps1qm3.f90`: A three-level multipair PSLQ program: This finds the coefficients of the degree-72 polynomial satisfied by $3^{1/8} - 2^{1/9}$. Size of code: 2,365 lines. Precision level: 1200 digits; switches frequently between full precision (1200 digits), medium precision (varies from 50 to 110 digits) and double precision.
7. `tpphix3.f90`: A Poisson phi program: This 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: 2,728 lines. Precision level: 2500 digits; switches frequently between full precision (2500 digits), medium precision (varies from 50 to 250 digits) and double precision.
8. `tquad.f90`: A quadrature (numerical integration) program: This 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: 1,565 lines. Precision level: 1000 digits, but most computation is done to 500 digits; switches frequently between 500 and 1000 digits.
9. `tquadgs.f90`: A quadrature (numerical integration) program: This performs the Gaussian quadrature algorithm on many of the same 18 problems as in `tquad`, producing results correct to 500-digit accuracy. This code runs much longer than `tquad`, due to the much greater 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: 759 lines. Precision level: 1000 digits, but

most computation is done to 500 digits; switches frequently between 500 and 1000 digits.

In addition, the `fortran-var1` and `fortran-var2` directories include test scripts that compile the library and run each of the above sample programs above (except `tquadgs.f90`): In directory `fortran-var1`, these scripts are:

- `gnu-mpfun-tests1.scr`
- `intel-mpfun-tests1.scr`
- `nag-mpfun-tests1.scr` (MPFUN20-Fort only)

while in directory `fortran-var2` they are:

- `gnu-mpfun-tests2.scr`
- `intel-mpfun-tests2.scr`
- `nag-mpfun-tests2.scr` (MPFUN20-Fort only)

For each of the eight test programs, the script outputs either TEST PASSED or TEST FAILED. If all tests pass, then one can be fairly confident that the MPFUN2020 software and underlying compilers are 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 11 presents some performance timings comparing the two versions of the package for the first seven test programs listed above, plus an additional 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. All runs except the last were performed with default precision level `mpipl` set to 4,000 digits, even though most of these runs required less.

These runs were performed using the GNU `gcc/gfortran` compiler (version 13.2.0). For the MPFUN20-MPFR version, the GNU `gcc` compiler (version 13.2.0) was used to build the GMP and MPFR libraries. These runs were

Code name	Precision (digits)	k	m	MPFUN20- Fort	MPFUN- MPFR
testmpfun	500			11.57	3.16
tpolysolve	4000			1.51	0.28
tpslq1	250			4.38	2.52
tpslqm1	250			2.31	2.12
tpslqm2	750			3.20	3.08
tpslqm3	1200			11.30	11.09
tquad	1000			18.68	3.56
tpphix3	2500	28	96	10.51	8.13
tpphix3	5200	25	100	288.99	219.29

Table 11: Timings on a suite of test programs (seconds).

executed on an Apple Mac Studio system, with OS version 14.4.1, an Apple M2 Max processor and 32 Gbyte of main memory. For uniformity, the timings are listed in the Table 11 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, and, in some cases, also in MPFUN20-MPFR; other algorithms used in MPFUN20-MPFR are described in [27].

7.1 Algorithms for basic arithmetic

Multiplication. [MPFUN20-Fort.] For modest levels of precision, MPFUN20-Fort employs 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. [MPFUN20-Fort.] Division of two extra-high-precision arguments a and b can be performed by the following scheme. First, note that this Newton-Raphson algorithm iteration converges to $1/b$ [17, pg. 226]:

$$x_{k+1} = x_k + (1 - x_k \cdot b) \cdot x_k, \quad (5)$$

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 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):

$$a/b \approx (a \cdot x_n) + [a - (a \cdot x_n) \cdot b] \cdot x_n, \quad (6)$$

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.

Square roots. [MPFUN20-Fort.] Square roots are calculated by the following Newton-Raphson iteration, which converges to $1/\sqrt{a}$ [17, pg. 227]:

$$x_{k+1} = x_k + 1/2 \cdot (1 - x_k^2 \cdot a) \cdot x_k, \quad (7)$$

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):

$$\sqrt{a} \approx (a \cdot x_n) + 1/2 \cdot [a - (a \cdot x_n)^2] \cdot x_n, \quad (8)$$

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. [MPFUN20-Fort.] 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. [MPFUN20-Fort and MPFUN-MPFR.] The pseudorandom number included in the package versions is the following:

$$x_{n+1} = \text{frac}(5501758857861179 \cdot x_n), \quad (9)$$

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, say, by typing

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

The same scheme is used for both MPFUN20-Fort and MPFUN20-MPFR. However, the specific sequence is *not* the same between the two versions, nor is it same if the precision level is changed. Note that the above algorithms are trivially thread-safe, since no auxiliary data or saved context is involved.

7.2 FFT-based multiplication

[MPFUN20-Fort only.] Although the multiplication algorithm described above in Section 7.1 is very efficient, for higher levels of precision (above

approximately 20,000 digits, based on the present author’s implementation), significantly faster performance can be achieved by employing a convolution scheme based on fast Fourier transforms [20][17, 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

$$c_k = \sum_{j=0}^{2n-1} a_j b_{k-j}, \quad 0 \leq k < 2n, \quad (10)$$

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

$$(c) = F^{-1}[F(a) \cdot F(b)], \quad (11)$$

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 20,000 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 5.2 for details.

7.3 Transcendental constants and functions

Most arbitrary precision packages require a significant “context” of global data to support transcendental function evaluation at a particular precision

level, and this data is often problematic for thread safety. Further, 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 functions in MPFUN2020 were designed to require only a minimum of global data, namely the three constants $\log(2)$, π and γ , which are provided in compile-time global data statements unless extremely high precision is required. If a precision over 20,000 digits is required, `mpinit` must be called at the start of execution, as noted above in Section 5.2, to compute these three constants.

π is computed in MPFUN20-Fort using a quadratically convergent scheme discovered by Richard Brent and Eugene Salamin in 1976 [8, 109–110]. Set $a_0 = 1$, $b_0 = 1/\sqrt{2}$ and $s_0 = 1/2$. Iterate, beginning with $k = 0$,

$$\begin{aligned} a_{k+1} &= \frac{a_k + b_k}{2}, & b_{k+1} &= \sqrt{a_k b_k}, \\ c_{k+1} &= a_{k+1}^2 - b_{k+1}^2, & s_{k+1} &= s_k - 2^{k+1} c_{k+1}, \\ p_{k+1} &= \frac{2a_{k+1}^2}{s_{k+1}}. \end{aligned} \tag{12}$$

Then p_k converges quadratically to π : each iteration approximately *doubles* the number of correct digits.

$\log(2)$ is computed in MPFUN20-Fort using the following quadratically convergent algorithm, also due to Salamin and Brent [8, 228]: Select $N > 2^B$, where B is the number of bits of desired precision. Then

$$\log(2) = \frac{\pi}{2 \operatorname{AGM}(1, 4/2^N)}, \tag{13}$$

where $\operatorname{AGM}(a, b)$ denotes the arithmetic-geometric mean: Set $a_0 = a$ and $b_0 = b$, then iterate

$$\begin{aligned} a_{k+1} &= (a_k + b_k)/2 \\ b_{k+1} &= \sqrt{a_k \cdot b_k} \end{aligned} \tag{14}$$

until convergence, i.e., until $a_k = b_k$ to available precision.

Euler’s constant $\gamma = 0.57721566490153286\dots$ is calculated in MPFUN20-Fort using the following formula, which is an improvement of a technique previously used by Sweeney [37]. If a result accurate to at least B bits is

desired, first select the integer $N = \lceil \log_2(B \log 2) \rceil$. Then

$$\gamma \approx \frac{2^N}{e^{2^N}} \sum_{m=0}^{\infty} \frac{2^{mN}}{(m+1)!} \sum_{t=0}^m \frac{1}{t+1} - N \log 2. \quad (15)$$

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

The binary values of $\log(2)$, π and γ are stored to 20,000 digit precision in data statements in module MPFUNA, in both versions. If higher precision is required, `mpinit` must be called at the start of execution in a single-threaded section of code, as explained in Section 5.2.

Exponential and logarithm. [MPFUN20-Fort only; MPFUN20-MPFR uses the MPFR routines.] The exponential function routine first reduces the input argument to 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 $\exp(x)$:

$$\exp(x) = x + \frac{x^2}{2!} + \frac{x^3}{3!} + \cdots, \quad (16)$$

where the terms are computed with a precision level that decreases as the terms decrease in size, 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 $B^{2/5}$, where B is the number of bits 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:

$$x_{k+1} = x_k - \frac{e^{x_k} - a}{e^{x_k}}. \quad (17)$$

Trigonometric functions. [MPFUN20-Fort only; MPFUN20-MPFR uses the MPFR routines.] 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 $\sin(x)$:

$$\sin(x) = x - \frac{x^3}{3!} + \frac{x^5}{5!} + \cdots, \quad (18)$$

where, as with $\exp(x)$, the terms are computed with a precision level that decreases as the terms decrease in size. Then the double-angle formulas

$$\cos(2x) = 1 - 2 \sin^2(x) \tag{19}$$

$$\cos(2x) = 2 \cos^2(x) - 1, \tag{20}$$

are applied q times (formula (19) is used once, and (20) thereafter). In the current implementation, q is set to the greatest integer in $\sqrt{2B}$, where B is the precision in bits, unless the reduced argument is very close to zero, 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. [MPFUN20-Fort only; MPFUN20-MPFR uses the MPFR routine.] 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.

7.4 Special functions

Modern mathematical and scientific computing frequently often involves more sophisticated functions, collectively termed “special functions” [23]. Many of these functions have been implemented in the MPFUN20-Fort and MPFUN20-MPFR packages, 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.

Even Bernoulli coefficients. [MPFUN20-Fort and MPFUN20-MPFR.] The subroutine `mpberne` returns an array containing Bernoulli numbers B_{2k} for $k = 1$ to n . It employs an advanced algorithm, due to the present author [10]: First, compute $\{\zeta(2k), 0 \leq k \leq n\}$ to P words precision based on the

following well-known formulas:

$$\begin{aligned}
\coth(\pi x) &= \cosh(\pi x) / \sinh(\pi x) \\
&= \frac{1}{\pi x} \cdot \frac{1 + (\pi x)^2/2! + (\pi x)^4/4! + \dots}{1 + (\pi x)^2/3! + (\pi x)^4/5! + \dots} \\
&= \frac{1}{\pi x} \cdot (1 + (\pi x)^2/3 - (\pi x)^4/45 + 2(\pi x)^6/945 - \dots) \\
&= \frac{2}{\pi x} \sum_{k=1}^{\infty} (-1)^{k+1} \zeta(2k) x^{2k}.
\end{aligned} \tag{21}$$

The strategy is to calculate the coefficients of the series by polynomial operations. Polynomial division is performed by computing the reciprocal of the denominator polynomial, by a polynomial Newton iteration, as follows. Let $N(x)$ be the polynomial approximation to the numerator series; let $D(x)$ be a polynomial approximation to the denominator series; and let $Q_k(x)$ be polynomial approximations to $R(x) = 1/D(x)$. Then iterate:

$$Q_{k+1} = Q_k(x) + [1 - D(x)Q_k(x)]Q_k(x). \tag{22}$$

In these iterations, both the degree of the polynomial $Q_k(x)$ and the precision level in words are initially set to 4. When convergence is achieved at this precision level, the degree is doubled, and iterations are continued, etc., until the final desired degree is achieved. Then the precision level is doubled and iterations are performed in a similar way, until the final desired precision level P is achieved. The reciprocal polynomial $R(x)$ produced by this process is then multiplied by the numerator polynomial $N(x)$ to yield an approximation to the quotient series. The even zeta values are then the coefficients of this series, scaled according to the formula (21) above.

Once the even integer zeta values have been computed, the even Bernoulli coefficients are computed via this well-known formula (for $n > 0$):

$$B_{2n} = \frac{(-1)^{n-1} 2(2n)! \zeta(2n)}{(2\pi)^{2n}}. \tag{23}$$

Aside from requiring π , which is a compile-time constant, the code implementing this function does not rely on any stored data, and so is thread-safe. The library subroutine name is `mpberne` (see Table 5).

BesselI function. [MPFUN20-Fort and MPFUN20-MPFR.] The BesselI function, or, more formally, the modified Bessel function of the first kind, is

defined as [23, 10.25.2]:

$$I_\nu(z) = \left(\frac{z}{2}\right)^\nu \sum_{k=0}^{\infty} \frac{\left(\frac{z^2}{4}\right)^k}{k! \Gamma(\nu + k + 1)} \quad (24)$$

For modest-sized z , this function can be evaluated directly according to this definition, terminating the summation when terms are smaller than epsilon times the current sum. For larger z (the author uses the condition $|z| > 0.4B$, where B is the precision in bits), the following asymptotic formula is faster [23, 10.40.1]:

$$I_\nu(z) \approx \frac{e^z}{\sqrt{2\pi z}} \sum_{k=0}^{\infty} \frac{(-1)^k a_k(\nu)}{z^k}, \quad (25)$$

where

$$a_k(\nu) = \frac{(4\nu^2 - 1^2)(4\nu^2 - 3^2) \cdots (4\nu^2 - (2k - 1)^2)}{8^k k!}. \quad (26)$$

Aside from requiring π , which is a compile-time constant, the code implementing this algorithm does not rely on any stored data, and so is thread-safe. The library function names are `bessel_i` and `bessel_in` (see Table 5).

BesselJ function. [MPFUN20-Fort only; MPFUN20-MPFR uses the MPFR routine.] The BesselJ function, or, more formally, the Bessel function of the first kind, is defined as [23, 10.2.2]:

$$J_\nu(z) = \left(\frac{z}{2}\right)^\nu \sum_{k=0}^{\infty} \frac{(-1)^k \left(\frac{z^2}{4}\right)^k}{k! \Gamma(\nu + k + 1)}. \quad (27)$$

For modest-sized z , this function can be evaluated directly according to this definition, terminating the summation when terms are smaller than epsilon times the current sum. However, large amounts of cancellation occurs in this formula, so the present author's implementation employs a working precision $P = B + 1.4|z|$ bits, where B is desired result precision in bits, up to 3X the requested precision.

For larger z (the author uses the condition $|z| > 0.4B$, where B is the precision in bits), the following asymptotic formula [23, 10.17.3] is faster:

$$J_\nu(z) \approx \sqrt{\frac{2}{\pi z}} \left(\cos(\omega) \sum_{k=0}^{\infty} \frac{(-1)^k a_{2k}(\nu)}{z^{2k}} - \sin(\omega) \sum_{k=0}^{\infty} \frac{(-1)^k a_{2k+1}(\nu)}{z^{2k+1}} \right), \quad (28)$$

where $\omega = z - \pi\nu/2 - \pi/4$, and where $a_k(\nu)$ is given by (26). Aside from requiring π , which is a compile-time constant, the code implementing these algorithms does not rely on any stored data, and so is thread-safe. The library function names are `bessel_j`, `bessel_jn`, `bessel_j0` and `bessel_j1` (see Tables 3 and 5).

BesselK function. [MPFUN20-Fort only; MPFUN20-MPFR uses the MPFR routine.] The BesselK function, or, more formally, the modified Bessel function of the second kind, is defined as [23, 10.31.1]:

$$K_n(z) = \frac{1}{2} \left(\frac{z}{2}\right)^{-n} \sum_{k=0}^{n-1} \frac{(n-k-1)!}{k!} \left(-\frac{z^2}{4}\right)^k + (-1)^{n+1} \log\left(\frac{z}{2}\right) I_n(z) \\ + (-1)^n \frac{1}{2} \left(\frac{z}{2}\right)^n \sum_{k=0}^{\infty} (\psi(k+1) + \psi(n+k+1)) \frac{\left(\frac{z^2}{4}\right)^k}{k!(n+k)!}, \quad (29)$$

where the $\psi(n)$ (digamma) function for positive integer arguments can be evaluated from the simple formula $\psi(n) = -\gamma + \sum_{k=1}^{n-1} 1/k$, where γ is Euler's constant. For modest-sized z , this function can be evaluated directly according to this definition, terminating the summation when terms are smaller than epsilon times the current sum. However, large amounts of cancellation occurs in this formula, so the present author's implementation employs a working precision $B + 2.8|z|$ bits, where B is desired result precision in bits, up to 3X the requested precision.

For larger z (the author uses the condition $|z| > 0.4B$, where B is the precision in bits), the following asymptotic formula [23, 10.40.1] is faster:

$$K_\nu(z) \approx \sqrt{\frac{\pi}{2z}} e^{-z} \sum_{k=0}^{\infty} \frac{a_k(\nu)}{z^k}, \quad (30)$$

where $a_k(\nu)$ is given by (26).

In the case when the order ν is not an integer, this formula is used instead:

$$K_\nu(z) = \frac{\pi(I_{-\nu} - I_\nu)}{2 \sin(\nu\pi)}. \quad (31)$$

Aside from requiring γ and π , which are compile-time constants, the code implementing these algorithms does not rely on any stored data, and so is

thread-safe. The library function names are `bessel_k` and `bessel_kn` (see Table 5).

BesselY function. [MPFUN20-Fort and MPFUN20-MPFR.] The BesselY function, or, more formally, the Bessel function of the second kind, is defined as [23, 10.8.1]:

$$Y_n(z) = \frac{1}{\pi} \left[\left(\frac{z}{2}\right)^{-n} \sum_{k=0}^{n-1} \frac{(n-k-1)!}{k!} \left(\frac{z^2}{4}\right)^k + 2 \log\left(\frac{z}{2}\right) J_n(z) + \left(\frac{z}{2}\right)^n \sum_{k=0}^{\infty} (\psi(k+1) + \psi(n+k+1)) \frac{\left(-\frac{z^2}{4}\right)^k}{k!(n+k)!} \right], \quad (32)$$

where the $\psi(n)$ (digamma) function for positive integer arguments can be evaluated from the simple formula $\psi(n) = -\gamma + \sum_{k=1}^{n-1} 1/k$, where γ is Euler's constant. For modest-sized z , this function can be evaluated directly according to this definition, terminating the summation when terms are smaller than epsilon times the current sum. However, large amounts of cancellation occurs in this formula, so the present author's implementation employs a working precision $P = B + 1.4|z|$ bits, where B is desired result precision in bit, up to 3X the requested precision.

In the case when the order ν is not an integer, this formula is used instead:

$$Y_\nu(z) = \frac{J_\nu \cos(\nu\pi) - J_{-\nu}}{\sin(\nu\pi)}. \quad (33)$$

For larger z (the author uses the condition $|z| > 0.4B$, where B is the precision in bits), the following asymptotic formula [23, 10.17.4] is faster:

$$J_\nu(z) \approx \sqrt{\frac{2}{\pi z}} \left(\sin(\omega) \sum_{k=0}^{\infty} \frac{(-1)^k a_{2k}(\nu)}{z^{2k}} + \cos(\omega) \sum_{k=0}^{\infty} \frac{(-1)^k a_{2k+1}(\nu)}{z^{2k+1}} \right), \quad (34)$$

where $\omega = z - \pi\nu/2 - \pi/4$, and where $a_k(\nu)$ is given by (26).

Aside for requiring γ and π , which are compile-time constants, the code implementing these algorithms does not rely on any stored data, and so is thread-safe. The library function names are `bessel_y`, `bessel_y0`, `bessel_y1` and `bessel_yn` (see Tables 3 and 5).

Digamma function. [MPFUN20-Fort and MPFUN20-MPFR; MPFUN20-MPFR also has a separate routine from MPFR.] This evaluates the digamma function, using the following asymptotic formula [23, 5.11.2]:

$$\psi(x) \approx \log(x) - \frac{1}{2x} - \sum_{k=1}^{\infty} \frac{B_{2k}}{2kx^{2k}}, \quad (35)$$

where B_{2k} are the even Bernoulli coefficients (see *Even Bernoulli coefficients* above). The code implementing this algorithm is entirely thread-safe, provided the requisite even Bernoulli numbers are precomputed.

The library function name is `digamma_be`. Before calling `digamma_be`, subroutine `mpberne` must be called to compute the even Bernoulli numbers, where the parameter `nb` is at least twice the desired precision in digits, and the resulting array and size must be included in subsequent calls to `digamma_be`. See Table 5.

Error functions. [MPFUN20-Fort only; MPFUN20-MPFR uses the MPFR routines.] The error function is defined as [23, 7.6.2]:

$$\operatorname{erf}(z) = \frac{2}{\sqrt{\pi}e^{z^2}} \sum_{n=0}^{\infty} \frac{2^n z^{2n+1}}{1 \cdot 3 \cdots (2n+1)}. \quad (36)$$

This formula is satisfactory for computation unless the argument z is large, in which case the following asymptotic expression for the error function complement works better:

$$\operatorname{erfc}(z) \approx \frac{1}{\sqrt{\pi}e^{z^2}} \sum_{n=0}^{\infty} \frac{(-1)^n 1 \cdot 3 \cdots (2n-1)}{2^n z^{2n+1}}. \quad (37)$$

where $\operatorname{erfc}(z) = 1 - \operatorname{erf}(z)$, and where the product in the numerator is 1 when $n = 0$. In the author's code, (36) is used when $|z| < B/100$, and (37) is used when $|z| \geq B/100$, where B is the number of bits of precision. Aside from requiring π , which is a compile-time constant, the code implementing this algorithm does not rely on any stored data, and so is thread-safe. The library function names are `erf` and `erfc` (see Tables 3 and 5).

Exponential integral function. [MPFUN20-Fort only; MPFUN20-MPFR uses the MPFR routine.] The exponential integral function is evaluated using the

following formula due to Ramanujan [24]:

$$\operatorname{Ein}(z) = \gamma + \log(|z|) + \exp(z/2) \sum_{k=1}^{\infty} \frac{(-1)^{k-1} z^k}{2^{k-1} k!} \sum_{j=0}^{\lfloor (k-1)/2 \rfloor} \frac{1}{2j+1} \quad (38)$$

$$= -\Gamma(0, -z) \quad (39)$$

See *incomplete gamma function* below for $\Gamma(0, -z)$.

Gamma function. [MPFUN20-Fort only; MPFUN20-MPFR uses the MPFR routine.] The gamma function employs a very efficient but little-known formula due to Ronald W. Potter [36] for most values, combined with a separate formula for values close to one. 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 $(0, 1)$. If $1 - 1/10^{10} < t < 1$, then employ this formula from the DLMF [23, 5.7.3]:

$$\log(\Gamma(1+x)) = -\log(1+x) + x(1-\gamma) + \sum_{k=2}^{\infty} (-1)^k (\zeta(k) - 1) \frac{x^k}{k}, \quad (40)$$

where $x = t-1$ and γ is Euler's constant, and where $\zeta(k)$ is computed using Peter Borwein's algorithm, as described below under *Zeta*. For other values of t , define $\alpha = 2 \operatorname{int}(b/4 \cdot \log 2 + 1)$, where b is the number of bits of precision and int is greatest integer. Define the Pochhammer symbol as

$$(t)_k = t(t+1)(t+2) \cdots (t+k-1). \quad (41)$$

Then define the functions

$$\begin{aligned} A(t, \alpha) &= \left(\frac{\alpha}{2}\right)^t t \sum_{k=0}^{\infty} \frac{(\alpha^2/4)^k}{k!(t)_{k+1}} \\ B(t, \alpha) &= \left(\frac{\alpha}{2}\right)^{-t} (-t) \sum_{k=0}^{\infty} \frac{(\alpha^2/4)^k}{k!(-t)_{k+1}}. \end{aligned} \quad (42)$$

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

$$\Gamma(t) = \sqrt{\frac{\pi A(t, \alpha)}{t \sin(\pi t) B(t, \alpha)}}. \quad (43)$$

Aside from requiring π and γ , which are compile-time constants, the code implementing this function does not rely on any stored data, and so is thread-safe. The library function name is `gamma` (see Tables 3 and 5).

Hurwitz zeta function. [MPFUN20-Fort and MPFUN20-MPFR.] The Hurwitz zeta function is computed by the following formula, which is a modification of a scheme originally due to Richard Crandall [6]: Let the notation $\sigma(x)$ denote $\text{sign}(x)$. Then for $0 < a < 1$ and for $s \neq 1, 0, -1, -2, \dots$,

$$\begin{aligned} \zeta(s, a) = & \frac{\sqrt{\pi} \pi^{(s-1)/2}}{(s-1) \Gamma(s/2)} \\ & + \frac{1}{2} \sum_{n=-\infty}^{\infty} \left(\frac{\Gamma(\frac{s}{2}, \pi(n+a)^2)}{\Gamma(s/2)} + \sigma(n+a) \frac{\Gamma(\frac{s+1}{2}, \pi(n+a)^2)}{\Gamma((s+1)/2)} \right) \frac{1}{|n+a|^s} \\ & + \pi^{s-1/2} \sum_{m=1}^{\infty} \frac{1}{m^{1-s}} \left(\frac{\Gamma(\frac{1-s}{2}, \frac{m^2 \pi^2}{\pi})}{\Gamma(s/2)} \cos(2\pi m a) + \frac{\Gamma(1-\frac{s}{2}, \frac{m^2 \pi^2}{\pi})}{\Gamma((s+1)/2)} \sin(2\pi m a) \right). \end{aligned} \quad (44)$$

Although each of the individual terms are relatively expensive to compute due to the incomplete gamma functions, convergence is quite rapid: 30 terms produce over 1000 digits. Note that when $a = 1$, $\zeta(s, a) = \zeta(s)$. Aside from requiring π , which is a compile-time constant, the code implementing this function does not rely on any stored data, and so is thread-safe. The library function name is `hurwitz_zetan` (see Table 5).

[MPFUN20-Fort and MPFUN20-MPFR.] An even faster formula, which relies on precomputed even Bernoulli numbers (see *Even Bernoulli numbers* above), is as follows. Select an integer $q > 0$ (the present author uses $q = \text{nint}(0.6d)$, where d is the precision in digits and nint is nearest integer). Then

$$\begin{aligned} \zeta(s, a) = & \sum_{k=0}^{q-1} \frac{1}{(a+k)^s} + \frac{1}{(s-1)(a+q)^{s-1}} + \frac{1}{2(a+q)^s} \\ & + \sum_{k=1}^{\infty} \frac{B_{2k} s(s+1) \cdots (s+2k-2)}{(2k)!(a+q)^{s+2k-1}}. \end{aligned} \quad (45)$$

The code implementing this function (`hurwitz_zetan_be`) relies only on pre-computed even Bernoulli numbers, and thus is thread-safe. Before calling `hurwitz_zetan_be`, subroutine `mpberne` must be called to compute the even

Bernoulli numbers, where the parameter `nb` is at least twice the desired precision in digits, and the resulting array and size must be included in subsequent calls to `hurwitz_zetan_be` (see Table 5).

Hypergeometric pFq function. [MPFUN20-Fort and MPFUN20-MPFR.] The hypergeometric pFq function is defined as:

$${}_pF_q \left(\begin{matrix} a_1, a_2, \dots, a_p; x \\ b_1, b_2, \dots, b_q \end{matrix} \right) = \sum_{k=0}^{\infty} \frac{(a_1)_k (a_2)_k \dots (a_p)_k x^k}{(b_1)_k (b_2)_k \dots (b_q)_k k!}, \quad (46)$$

where $(s)_n = s(s+1)(s+2)\dots(s+n-1)$ is the Pochhammer symbol as before. This formula is evaluated directly as stated, terminating the summation when terms are smaller than epsilon times the current sum. Note that the terms satisfy a relatively simple recursion. The code implementing this function does not rely on any stored data, and so is thread-safe. The library function name is `hypergeom_pfq` (see Table 5).

Incomplete gamma function. [MPFUN20-Fort only; MPFUN20-MPFR uses the MPFR routine.] For modest-sized positive arguments (the author uses the condition $|z| < 0.833B$, where B is the precision level in bits), the MPFUN20-Fort incomplete gamma function is evaluated using the following formula [23, 8.7.3]:

$$\Gamma(a, z) = \Gamma(a) \left(1 - \frac{z^a}{e^z} \sum_{k=0}^{\infty} \frac{z^k}{\Gamma(a+k+1)} \right). \quad (47)$$

Note, as with the BesselJ function, that although formula (47) 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 [23, 8.11.2]:

$$\Gamma(a, z) \approx \frac{z^{a-1}}{e^z} \sum_{k=0}^{\infty} \frac{(-1)^k (1-a)_k}{z^k}. \quad (48)$$

For the case $a = 0$, the following formula is employed, which is a modification of a formula due to Ramanujan for the exponential integral function [24]:

$$\Gamma(0, z) = -\gamma - \log(|z|) + \exp(-z/2) \sum_{k=1}^{\infty} \frac{z^k}{2^{k-1} k!} \sum_{j=0}^{\lfloor (k-1)/2 \rfloor} \frac{1}{2j+1}. \quad (49)$$

For $z < 0$, formula (49) returns the correct real part of the complex value. For $a = -1, -2, \dots$, the following formula is employed [31]:

$$\Gamma(-n, z) = \frac{1}{n!} \left((-1)^n \Gamma(0, z) + \frac{1}{e^z z^n} \sum_{k=0}^{n-1} (-1)^k (n-k-1)! z^k \right). \quad (50)$$

One difficulty is that cancellation may occur in formulas (49) and (50), so for certain arguments the working precision is increased, up to 3X the requested precision. Aside from requiring Euler's constant γ in the formula for $\Gamma(0, z)$, the code implementing this function does not require any saved data and thus is thread-safe. The library function name is `incgamma` (see Table 5).

Polygamma function. [MPFUN20-Fort and MPFUN20-MPFR.] The polygamma function is evaluated using the Hurwitz zeta function via the formula [29]

$$\psi^{(m)} = (-1)^{m+1} m! \zeta(m+1, z). \quad (51)$$

See *Hurwitz zeta function* above for details. The library function name is `polygamma`; an alternate function name is `polygamma_be`, requiring precomputed even Bernoulli numbers (see Table 5).

Polylogarithm function. [MPFUN20-Fort and MPFUN20-MPFR.] The polylogarithm function for integer order n is defined as [23, 25.12.10]:

$$\text{Li}_n(z) = \sum_{k=1}^{\infty} \frac{z^k}{k^n} \quad (52)$$

This is satisfactory for computation for $n \geq 1$ and $|z| < 1$. When $n = 0$, one uses the formula $\text{Li}_0(z) = z/(1-z)$. For negative order n , one can use the formula [40]:

$$\text{Li}_{-n}(z) = \frac{1}{(1-z)^{n+1}} \sum_{k=1}^n a_{n,k} z^k, \quad (53)$$

where the coefficients $a_{n,k}$ are given by the recurrence

$$a_{n,k} = (n+1-k)a_{n-1,k-1} + ka_{n-1,k}, \quad (54)$$

with $a_{n,1} = 1$. One difficulty is that cancellation may occur in formula (53), so for certain arguments the working precision is increased, up to 3X the

requested precision. Since the requisite auxiliary data is precomputed and passed via argument, the code implementing these algorithms is thread-safe.

Because of the different algorithms and data requirements, MPFUN2020 includes three routines for polylogarithms, namely `polylog_ini`, `polylog_neg` and `polylog_pos`. Before calling `polylog_neg`, subroutine `polylog_ini` must be called to compute the coefficients $a_{n,k}$ for the given value of n , and the resulting array and size must be included in subsequent calls to `polylog_neg` (see Table 6).

StruveH fuction. [MPFUN20-Fort and MPFUN20-MPFR.] The StruveH function is defined for positive integer order n as [23, 11.2.1]:

$$H_n(z) = \left(\frac{z}{2}\right)^{n+1} \sum_{k=0}^{\infty} \frac{(-1)^k (z/2)^{2k}}{\Gamma(k+3/2)\Gamma(n+k+3/2)}. \quad (55)$$

This formula is evaluated as stated. As with the BesselJ function, the gamma function references here can be computed iteratively, starting with the fact that $\Gamma(3/2) = \sqrt{\pi}/2$. One difficulty is that due to the alternating sign in the numerator, precision is lost due to cancellations. In the author's code, this is overcome by employing $B + B|z|/1000$ bits of precision, where B is the number of bits desired for the result, up to 3X the requested precision. Aside from requiring π , which is a compile-time constant, the code implementing this function does not rely on any stored data, and so is thread-safe. The library function name is `struve_hn` (see Table 6).

Zeta function. [MPFUN20-Fort only; MPFUN20-MPFR uses the MPFR routine.] For large positive arguments s (the present author uses the condition $s > B \log(2)/\log(2B/3)$), where B is the precision in digits), it suffices to use the definition of zeta, namely

$$\zeta(s) = \sum_{n=1}^{\infty} \frac{1}{n^s}. \quad (56)$$

For modest-sized arguments, the zeta function can be evaluated by means of this formula, due to Peter Borwein [18]. Select n to be the number of digits of precision required for the result. Define

$$e_j = (-1)^j \left(\sum_{k=0}^{j-n} \frac{n!}{k!(n-k)!} - 2^n \right), \quad (57)$$

where the empty sum is zero. Then

$$\zeta(s) \approx \frac{-1}{2^n(1-2^{1-s})} \sum_{j=0}^{2^n-1} \frac{e_j}{(j+1)^s}. \quad (58)$$

There is some loss of precision due to cancelation with these formulas, so the author's code employs a working precision 12.5% higher than the requested precision.

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

$$\zeta(s) = \frac{2 \cos(\pi(1-s)/2) \Gamma(1-s) \zeta(1-s)}{(2\pi)^{1-s}}. \quad (59)$$

Aside from the value of π , no auxiliary data for this algorithm required, so the code implementing these algorithms is thread-safe. The library formula names are `zeta` and `zeta_int` (see Table 6).

[MPFUN20-Fort and MPFUN20-MPFR.] An even faster algorithm for the zeta function, based on the Euler-Maclaurin summation formula, can be derived from the following [23, 25.2.9]: Select an integer parameter $N > 0$ (the present author uses $N = \text{nint}(0.6d)$, where d is the precision in digits and `nint` is nearest integer). Then

$$\zeta(s) = \sum_{k=1}^N \frac{1}{k^s} + \frac{1}{(s-1)N^{s-1}} - \frac{1}{2N^s} + \sum_{k=1}^{\infty} \binom{s+2k-2}{2k-1} \frac{B_{2k}}{2kN^{s-1+2k}}, \quad (60)$$

where B_{2k} are the even Bernoulli numbers (see *Even Bernoulli numbers* above). The library function name is `zeta_be`. Before calling `zeta_be`, the subroutine `mpberne` must be called to compute the even Bernoulli numbers, where the parameter `nb` is greater than twice the desired precision in digits, and the resulting array and size must be included in subsequent calls to `zeta_be` (see Table 6).

Polynomial solution routines. As briefly mentioned above, the package includes several routines for polynomial evaluation and solution. `mprpolysolve` and `mpcpolysolve`, respectively, solve an input MPR polynomial, given an MPR or MPC initial approximation. Both routines employ Newton iterations with a working precision that approximately doubles with each iteration.

Polynomials with repeated roots often defeat Newton iteration schemes for solution. Such polynomials may be solved by first applying the polynomial greatest common divisor (GCD) routine, provided in the package, to the input polynomial and its formal derivative, and then divide the input polynomial by the GCD polynomial, using the polynomial division routine. This produces a reduced polynomial that has all the roots of the original polynomial, but no repeated roots. Then the Newton iteration-based solution routines can quickly produce high-precision roots.

The polynomial GCD routine is `mprpolygcd`, and the polynomial division routine, which produces the quotient and remainder polynomials, is `mprpolydiv`. These routines employ straightforward implementations of the polynomial GCD and division algorithms. Because the polynomial GCD process is highly sensitive numerically, the author's code for this routine employs a working precision 12.5% higher than the requested precision.

Two other routines (`mprpolyeval`) and (`mpcpolyeval`) evaluate a polynomial at an MPR or MPC argument, respectively. See Table 7 for details.

References

- [1] D. H. Bailey, “MPFUN2015: A thread-safe arbitrary precision package (full documentation),” manuscript, 18 Nov 2020, preprint at <https://www.davidhbailey.com/dhbpapers/mpfun2015.pdf>
- [2] D. H. Bailey, “Finding large Poisson polynomials using four-level variable precision,” *Correctness 2021: 5th International Workshop on Software Correctness for HPC Applications*, 20 Nov 2021, preprint at <https://www.davidhbailey.com/dhbpapers/correct-2021.pdf>.
- [3] D. H. Bailey, “Large Poisson polynomials: Computation, results and analysis,” manuscript, 5 Mar 2023, available at <https://www.davidhbailey.com/dhbpapers/poisson-2023.pdf>.
- [4] D. H. Bailey, X. S. Li and K. Jeyabalan, “A comparison of three high-precision quadrature schemes,” *Experimental Mathematics*, vol. 14 (2005), no. 3, 317–329.
- [5] D. H. Bailey, R. Barrio and J. M. Borwein, “High precision computation: Mathematical physics and dynamics,” *Applied Mathematics and Computation*, vol. 218 (2012), 10106–10121.
- [6] D. H. Bailey and J. M. Borwein, “Crandall’s computation of the incomplete gamma function and the Hurwitz zeta function with applications to Dirichlet L-series,” *Applied Mathematics and Computation*, vol. 268C (Oct 2015), 462–477, preprint at <https://www.davidhbailey.com/dhbpapers/lerch.pdf>.
- [7] D. H. Bailey and J. M. Borwein, “High-precision arithmetic in mathematical physics,” *Mathematics*, vol. 3 (2015), 337–367, <http://www.mdpi.com/2227-7390/3/2/337/pdf>.
- [8] Jonathan M. Borwein and David H. Bailey, *Mathematics by Experiment: Plausible Reasoning in the 21st Century*, AK Peters, Natick, MA, 2008.
- [9] D. H. Bailey and J. M. Borwein, “Hand-to-hand combat with thousand-digit integrals,” *Journal of Computational Science*, vol. 3 (2012), 77–86.

- [10] D. H. Bailey, J. M. Borwein and R. E. Crandall, “On the Khintchine constant,” *Mathematics of Computation*, vol. 66 (1997), 417–431, preprint at <https://www.davidhbailey.com/dhbpapers/khinchine.pdf>.
- [11] 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), 12271–12302.
- [12] 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>.
- [13] 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, preprint at <http://www.davidhbailey.com/dhbpapers/PoissonLattice.pdf>.
- [14] 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), 349–363, preprint at <https://www.davidhbailey.com/dhbpapers/poisson-res.pdf>.
- [15] D. H. Bailey and D. J. Broadhurst, “Parallel integer relation detection: Techniques and applications,” *Mathematics of Computation*, vol. 70, no. 236 (Oct 2000), 1719–1736.
- [16] D. H. Bailey, X. S. Li and B. Thompson, “ARPREC: An arbitrary precision computation package,” Sep 2002, preprint at <http://crd.lbl.gov/~dhbailey/dhbpapers/arprec.pdf>.
- [17] J. M. Borwein and D. H. Bailey, *Mathematics by Experiment: Plausible Reasoning in the 21st Century*, 2nd ed., A. K. Peters, Natick, MA, 2008.
- [18] P. Borwein, “An efficient algorithm for the Riemann zeta function,” 1995, <http://www.cecm.sfu.ca/~pborwein/PAPERS/P155.pdf>.

- [19] R. P. Brent, “Fast multiple-precision evaluation of elementary functions,” *J. of the ACM*, vol. 23 (1976), 242–251.
- [20] R. P. Brent and P. Zimmermann, *Modern Computer Arithmetic*, Cambridge Univ. Press, 2010.
- [21] 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.
- [22] “Comparison of multiple-precision floating-point software,” <http://www.mpfr.org/mpfr-3.1.0/timings.html>.
- [23] “Digital library of mathematical functions,” National Institute of Standards and Technology, 2015, <http://dlmf.nist.gov>.
- [24] “Exponential integral,” Wikipedia article, viewed 3 Feb 2022, https://en.wikipedia.org/wiki/Exponential_integral.
- [25] H. R. P. Ferguson, D. H. Bailey and S. Arno, “Analysis of PSLQ, an integer relation finding algorithm,” *Mathematics of Computation*, vol. 68, no. 225 (Jan 1999), 351–369.
- [26] “Fortran 2008,” <http://fortranwiki.org/fortran/show/Fortran+2008>.
- [27] “The GNU MPFR library,” <http://www.mpfr.org>.
- [28] “GNU MPFR library: Comparison of multiple-precision floating-point software,” <http://www.mpfr.org/mpfr-current/timings.html>.
- [29] “Polygamma function,” Wikipedia page, viewed 28 Jan 2022, https://en.wikipedia.org/wiki/Polygamma_function.
- [30] 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.
- [31] “Incomplete gamma function,” Wikipedia article, viewed 28 Jan 2022, https://en.wikipedia.org/wiki/Incomplete_gamma_function.

- [32] A. K. Lenstra, H. W. Lenstra, and L. Lovasz, “Factoring polynomials with rational coefficients,” *Mathematische Annalen*, vol. 261 (1982), 515–534.
- [33] “MPFR C++,” <http://www.holoborodko.com/pavel/mpfr>.
- [34] *NIST Digital Library of Mathematical Functions*, version 1.0.6 (May 2013), <http://dlmf.nist.gov>.
- [35] E. Panzer, “Algorithms for the symbolic integration of hyperlogarithms with applications to Feynman integrals,” manuscript, 2014, <https://arxiv.org/abs/1403.3385>.
- [36] R. W. Potter, *Arbitrary Precision Calculation of Selected Higher Functions*, Lulu.com, San Bernardino, CA, 2014.
- [37] D. W. Sweeney, “On the computation of Euler’s constant”, *Mathematics of Computation*, vol. 17 (1963), 170–178.
- [38] H. Takahasi and M. Mori, “Double exponential formulas for numerical integration,” *Publications of RIMS*, Kyoto University, vol. 9 (1974), 721–741.
- [39] 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.
- [40] D. C. Wood, “The computation of polylogarithms,” undated manuscript, <https://www.cs.kent.ac.uk/pubs/1992/110/content.pdf>.