CPFloat: A C library for simulating low-precision arithmetic

Fasi, Massimiliano and Mikaitis, Mantas

2020

MIMS EPrint: **2020.22**

Manchester Institute for Mathematical Sciences
School of Mathematics

The University of Manchester

CPFloat: A C library for simulating low-precision arithmetic*

Massimiliano Fasi†   Mantas Mikaitis‡

Abstract

One can simulate low-precision floating-point arithmetic via software by executing each arithmetic operation in hardware and then rounding the result to the desired number of significant bits. For IEEE-compliant formats, rounding requires only standard mathematical library functions, but handling subnormals, underflow, and overflow demands special attention, and numerical errors can cause mathematically correct formulae to behave incorrectly in finite arithmetic. Moreover, the ensuing implementations are not necessarily efficient, as the library functions these techniques build upon are typically designed to handle a broad range of cases and may not be optimized for the specific needs of rounding algorithms. CPFloat is a C library for simulating low-precision arithmetics. It offers efficient routines for rounding, performing mathematical computations, and querying properties of the simulated low-precision format. The software exploits the bit-level floating-point representation of the format in which the numbers are stored, and replaces costly library calls with low-level bit manipulations and integer arithmetic. In numerical experiments, the new techniques bring a considerable speedup (typically one order of magnitude or more) over existing alternatives in C, C++, and MATLAB. To our knowledge, CPFloat is currently the most efficient and complete library for experimenting with custom low-precision floating-point arithmetic available in any language.

Key words. low-precision arithmetic, floating-point arithmetic, mixed precision, IEEE 754 standard, binary16, bfloat16, round-to-nearest, directed rounding, round-to-odd, stochastic rounding.

1   A plethora of floating-point formats and rounding modes

The 2019 revision of the IEEE 754 standard for floating-point arithmetic [29] specifies three basic binary formats for computation: binary32, binary64, and binary128. Most 64-bit CPUs equipped with a floating-point unit natively support binary32 and binary64 arithmetic, and 32-bit CPUs can simulate binary64 arithmetic very efficiently by relying on highly optimized software libraries. The binary128 format, introduced in the 2008 revision [28] of the original IEEE 754 standard [27], has not gained much popularity among hardware manufacturers, and over ten years after having been standardized

---

it is supported only by two IBM processor families: the POWER9, which implements version 3.0 of the Power Instruction Set Architecture [32], and the z/Architecture [35].

In fact, the low-precision requirements of artificial intelligence applications have steered the hardware market in the opposite direction, and a number of fewer-than-32-bit formats have been proposed in recent years. The first widely available 16-bit floating-point format was binary16. Despite having been defined in the last two revisions of the IEEE 754 standard only as an interchange format, binary16 has been supported as an arithmetic format by all NVIDIA microarchitectures since Pascal [39] and all AMD architectures since Vega [43]. Google has recently introduced the bfloat16 data type [30], a 16-bit format with approximately the same dynamic range as binary32. The latest Armv8 CPUs support a wide variety of floating-point formats, including binary32, binary64, bfloat16 [1, Sec. A1.4.5], binary16, and an alternative custom half-precision format [1, Sec. A1.4.2]. The latter is a 16-bit format that extends the dynamic range from $[-65{,}504, +65{,}504]$ to $[-131{,}008, +131{,}008]$ by reclaiming the 2,048 bit patterns (about 3%) that the binary16 format uses for infinities and NaNs (Not a Number). The latest NVIDIA graphics card microarchitecture, NVIDIA Ampere, features a 19-bit floating-point format, called TensorFloat-32, which combines the precision of binary16 with the exponent range of binary32 [40]. We discuss the general framework to which all these floating-point formats belong in Section 4.

Such a broad range of floating-point formats poses a major challenge to those developing mixed-precision algorithms for scientific computing, as studying the numerical behavior of an algorithm in different working precisions may require access to a number of high-end hardware platforms. To alleviate this issue, several software packages for simulating low-precision floating-point arithmetics have been proposed in recent years. We review the most popular options in Section 2.

Our contribution is two-fold. First, we discuss how the operations underlying the rounding of a floating-point number x to lower precision can be performed directly on the binary floating-point representation of x. We present the new algorithms in Section 5, and in Section 6 we explain how to implement them efficiently using bitwise operations and integer arithmetic. Second, we introduce CPFloat, a header-only C library that implements our algorithms and can be used to simulate low-precision binary floating-point arithmetic. The name of the library is a shorthand for *Custom-Precision Floats*. Section 3 lists the functions available in the library and provides a minimal, self-contained code snippet showing how CPFloat can be used as a full arithmetic library with custom precision floating-point formats. Section 7 describes the implemention and testing.

The floating-point representation in the simulated (target) format exists only implicitly: in practice, all numbers are represented in a (storage) format that is natively supported by the compiler. The storage formats currently available in CPFloat are binary32 and binary64.

CPFloat is not the first library for simulating low-precision arithmetic in C: the GNU MPFR library [19], for example, allows the programmer to work with arbitrarily few, as well as arbitrarily many, bits of precision. Unlike MPFR, CPFloat is intended only as an simulator for low-precision floating-point arithmetics. Internally, floating-point values are represented using binary32 or binary64 numbers, thus only formats with at most 53 bits of precision and an exponent range no broader than that of binary64 can be simulated. This narrower aim provides scope for the wide range of optimizations we discuss, which in turn yield more efficient implementations.

We provide a MEX interface that allows user to call CPFloat directly from their MATLAB and Octave code. We use this to compare the performance of the new library with that of the MATLAB function chop [25] and of the routines in the FLOATP_Toolbox [36]. Our experimental results, discussed in Section 8, show that the new codes bring a considerable speedup over these MATLAB alternatives, as long as the matrices being rounded are large enough to offset the overhead of calling C code from MATLAB.

The library can be used to prototype and test mixed-precision algorithms, as well as to simulate custom-precision hardware devices. Recently, for example, we have used CPFloat to develop an algorithm for computing matrix–matrix products on NVIDIA GPUs equipped with tensor cores [13]. These

Table 1: Synoptic table of available software packages for simulating low-precision floating-point arithmetic. The first three columns report the name of the package, the main programming language in which the software is written, and what storage formats are supported. The following three columns describe the parameters of the target formats: whether the number of bits of precision in the significand is arbitrary (A) or limited to the number of bits in the significand of the storage format (R); whether the exponent range can be arbitrary (A), must be a sub-range of the exponent range of the storage format (S), or can be a sub-range only for built-in types (B); whether the target format supports subnormal numbers (Y), does not support them (N), supports them only for built-in types (B), supports them by default but allows the user to switch the functionality off (F), or does not support them by default but allow the user to turn the functionality on (O). The following column lists the floating-point formats that are built into the system, if any. The last seven columns indicate what rounding modes the software supports: round-to-nearest with ties-to-even (RNE), ties-to-zero (RNZ), or ties-to-away (RNA), round-toward-zero (RZ), round-to-+∞ and round-to-−∞ (RUD), round-to-odd (RO), and the two variants of stochastic rounding discussed in Section 4 (SR). Fully supported rounding modes are denoted by ✓, while ! is used for rounding modes that can be simulated at a higher computational cost. The abbreviations bf16, tf32, fp16, fp32, and fp64 denote the formats bfloat16, TensorFloat-32, binary16, binary32, and binary64, respectively.

Package name	Primary language	Storage format	p	e	s	built-in	RNE	RNZ	RNA	RZ	RUD	RO	SR
GNU MPFR	C	custom	A	A	O		✓		!	✓	✓		
Sipe	C	multiple	R	S	Y		✓			✓			
rpe	Fortran	fp64	R	B	B	fp16	✓						
FloatX	C++	fp32/fp64	R	S	Y		✓						
chop	MATLAB	fp32/fp64	R	S	F	fp16/bf16	✓			✓	✓		✓
QPyTorch	Python	fp32	R	S	N		✓	✓	✓				✓
FLOATP	MATLAB	fp64	R	A	N		✓			✓	✓		✓
CPFloat	C	fp32/fp64	R	S	F	fp16/bf16/tf32	✓	✓	✓	✓	✓	✓	✓

mixed-precision matrix multiply–accumulate units internally use different precisions and rounding modes [14]. Using a software simulation, we could evaluate the numerical behavior of custom variations of the tensor cores, which we obtained by changing the rounding mode used in some key operations.

2   Related work

A number of packages to simulate low-precision floating-point arithmetic via software are available. These are summarised in Table 1, where we report the main programming language in which the software is written and detail what storage formats, target formats, and rounding modes are supported.

The most comprehensive software package for simulating arbitrary-precision floating-point arithmetic is the GNU MPFR library [19], which extends the formats in the IEEE 754 standard to any precision and to an arbitrarily large exponent range. The library is written in C, but interfaces for most programming languages are available. GNU MPFR is a de facto standard for working with arbitrary precision, and is typically used to perform computations that require high accuracy, rather than to simulate low-precision arithmetics.

Sipe is a header-only C mini-library designed to simulate low precision efficiently. This software supports round-to-nearest with ties-to-even and round-to-zero, and numbers can be stored either as a pair of signed integers representing the significand and the exponent of the floating-point value [34], or as a value in a native floating-point format. The latest version of the library [33] supports float,

double, `long double`, or `__float128` (from the GCC `libquadmath` library).

Dawson and Düben [11] recently developed a Fortran library, called rpe, for simulating reduced-precision floating-point arithmetics in large numerical simulations. In rpe, the reduced-precision floating-point values are stored as binary64 numbers, a solution that provides a very efficient technique for simulating floating-point formats with the same exponent range as binary64 and no more than 53 bits of precision.

FloatX [16] is a header-only C++ library for simulating low-precision arithmetic that supports binary32 and binary64 as storage formats. This software is more flexible than rpe, in that it allows the user to choose not only the number of significant digits, but also the number of bits used to represent the exponent in the target format.

The only rounding mode currently implemented by both rpe and FloatX is round-to-nearest with ties-to-even, which may be too restrictive when one wants to simulate hardware units where truncation or stochastic rounding are available.

Higham and Pranesh [25] have recently proposed chop, a MATLAB function for rounding arrays of binary32 or binary64 numbers to lower precision. This solution is more efficient and flexible than the `fp16` and `vfp16` MATLAB data types proposed by Moler [37], as the user can specify not only the boundaries of the dynamic exponent range and the number of bits in the significand, but also the rounding mode to be used and whether subnormals are supported. chop supports six rounding modes: the four default rounding modes prescribed by the IEEE 754-2019 for single operations and two variants of stochastic rounding. This function can be used only from within the MATLAB programming environment, and the underlying algorithms, which rely on mathematical operations involving the exponent and significand of the represented floating-point numbers, are not necessarily suitable for efficient implementations in a low-level language such as C. For example, chop uses the built-in MATLAB functions `abs`, `sign`, `ceil`, `floor`, `log2`, and `pow2`, which may not be optimized for the narrow range of inputs required in order disassemble and reassemble floating-point numbers.

The QPyTorch library is a low-precision arithmetic simulator, written in PyTorch, whose primary goal is to facilitate the training of neural networks in low precision without the overhead of building low-precision hardware [50]. The core design principles of this library are analogous to those of chop, in that numbers are stored in binary32 before as well as after rounding. QPyTorch supports custom floating-point formats that can fit into the binary32 format, arbitrary-precision fixed-point formats no wider than 24 bits, and block floating-point formats [49]. Infinities, NaNs, and subnormals are not supported for efficiency's sake and because, as the authors point out, these special values do not appear when training neural networks and thus may not be supported by the underlying low-precision hardware. The supported rounding modes are stochastic rounding and round-to-nearest with ties-to-even.

FLOATP_Toolbox [36] is a MATLAB toolbox for simulating reduced-precision fixed-point and floating-point arithmetics. This library uses binary64 as storage format, supports the same six rounding modes as chop [25], and implements a number of mathematical functions, such as `log`, `exp`, or `sin`, for example. The functionalities of the FLOATP_Toolbox can be used in two ways: either by calling the routines that work directly on the binary64 data structure, or by relying on the methods of the `floatp` class, which override a number of built-in MATLAB functions. A similar library for simulating posit arithmetic has been developed by the same author [24, 12].

VPREC-libm [4] is a related tool developed to evaluate the accuracy/performance trade-offs of numerical code. The library instruments the code so that calls to mathematical functions are intercepted at runtime and performed in simulated low precision. Each function is evaluated in binary128 arithmetic by using the GCC `libquadmath` library and the result is then rounded to the target precision.

The algorithms presented here complement existing software by proposing efficient techniques for implementing rounding using low-level bitwise instructions. Our library is intended as a software package that enables the use of these rounding functionalities in low-level languages such as C and C++, but high-level languages that allow the user to call C routines, either directly or indirectly, can

4

also benefit from it.

3 Features and examples of `CPFloat`

At the core of `CPFloat` are the efficient rounding routines described in Section 5. These convert a number from one floating-point format (the storage format, which will be either binary32 or binary64) to a second, custom format (the target format). The representation in the target format is implicit, as the rounded numbers are still stored as either binary32 or binary64 values. Relying on these rounding routines, `CPFloat` provides functions to simulate custom low-precision arithmetic when operating on arrays of scalars. We now describe its user interface, using Listing 1 as reference. In the example, we use `double` as storage format, thus we include the header file `cpfloat_binary64.h`; the header file `cpfloat_binary32.h` should be included when using `float` arrays.

The target format is specified by a C structure of type `optstruct`. The structure should be initialized by calling

```
optstruct *init_optstruct()
```

as done on line 8. This function allocates the memory underlying the `optstruct` correctly, but the target format that the structure represents is still undefined. The parameters of the target format should be initialized explicitly as shown on lines 11–17.

In Listing 1, we do not showcase all the functionalities of `CPFloat`, and in particular we do not inject soft errors in the low-precision results. The library supports two injection modes, which can be selected by changing the value of the `flip` field of the `optstruct` structure. If `flip` is set to `CPFLOAT_FRAC_SOFTERR`, then with probability p a single bit flip can strike the fraction of the low-precision rounded number. This is the injection mode used in chop [25].

If `flip` is set to `CPFLOAT_FP_SOFTERR`, on the other hand, the single bit flip can strike any bit of the low-precision floating-point representation. Soft errors are typically modelled in this way in high-performance computing studies [2] as well as in numerical simulations [44, Sec. V]. For reproducibility, errors are sometimes introduced by hands at specific points of the execution rather than at random [47]; in `CPFloat`, this can be achieved by setting p to 1.0 and switching the value of the `flip` field between `CPFLOAT_FP_SOFTERR` (soft errors on) and `CPFLOAT_NO_SOFTERR` (soft errors off).

The parameters in an `optstruct` structure can be validated with

```
int cpfloat_validate_optstruct(optstruct *fpopts)
```

if the storage format is binary64, or

```
int cpfloat_validate_optstructf(optstruct *fpopts)
```

if the storage format is binary32. The convention of appending an f to the function name when the storage format is binary32 is common in C, and is used throughout the library whenever a binary64 and a binary32 variant are offered. The two functions return a positive integer if the fields of the input do not represent a valid floating-point format, zero if they do, and a positive integer if they represent a valid format which may yield an unexpected behavior such as double rounding. Some of the fields of an `optstruct` are pointers, and the function

```
void free_optstruct(optstruct *fpopts)
```

should be used as shown on line 49 to free the memory correctly.

The functions that make up the library can be divided into three groups: functions for rounding, functions for computing in low-precision arithmetic, and functions for probing the low-precision floating-point format. We discuss these assuming arrays of binary64 values, for binary32 arrays it suffices to change all `double` to `float` and add an f at the end of the function name.

A `double` array can be rounded to low precision using the functions

```
int cpfloat(double *X, const double *A, const size_t n, optstruct *fpopts)
int cpf_fpround(double *X, const double *A, const size_t n, optstruct *fpopts)
```

Listing 1: CPFloat example

```c
#include <stdio.h>
#include "cpfloat_binary64.h"

#define N 3

int main () {
  // Allocate the data structure for target formats and rounding parameters.
  optstruct *fpopts = init_optstruct();

  // Set up the parameters for binary16 target format.
  fpopts->precision = 11;                 // Bits in the significand + 1.
  fpopts->emax = 15;                      // The maximum exponent value.
  fpopts->subnormal = CPFLOAT_SUBN_USE;   // Support for subnormals is on.
  fpopts->round = CPFLOAT_RND_TP;         // Round toward +infinity.
  fpopts->flip = CPFLOAT_NO_SOFTERR;      // Bit flips are off.
  fpopts->p = 0;                          // Bit flip probability (not used).
  fpopts->explim = CPFLOAT_EXPRANGE_TARG; // Limited exponent in target format.

  // Validate the parameters in fpopts.
  int retval = cpfloat_validate_optstruct(fpopts);
  printf("The validation function returned %d.\n", retval);

  // Initialize C array with arbitrary elements.
  double X[N] = { (double)5/3, M_PI, M_E };
  double Y[N] = { 1.5, 1.5, 1.5 };
  double Z[N];
  printf("X in binary64:\n  %.15e %.15e %.15e\n", X[0], X[1], X[2]);

  // Round the values of X to the binary16 format and store in Z.
  cpfloat(Z, X, N, fpopts);
  printf("X rounded to binary16:\n  %.15e %.15e %.15e\n", Z[0], Z[1], Z[2]);

  // Round the sum of X and Y.
  cpf_add(Z, X, Y, N, fpopts);
  printf("Sum rounded to binary16:\n  %.15e %.15e %.15e\n", Z[0], Z[1], Z[2]);

  // Round the product of X and Y.
  cpf_mul(Z, X, Y, N, fpopts);
  printf("Product rounded to binary16:\n  %.15e %.15e %.15e\n", Z[0], Z[1], Z[2]);

  // Round the logarithm of X.
  cpf_log(Z, X, N, fpopts);
  printf("Log rounded to binary16:\n  %.15e %.15e %.15e\n", Z[0], Z[1], Z[2]);

  // Round the 2-argument arctangent of X and Y.
  cpf_atan2(Z, X, Y, N, fpopts);
  printf("Angle rounded to binary16:\n  %.15e %.15e %.15e\n", Z[0], Z[1], Z[2]);

  free_optstruct(fpopts);
}
```

which convert the n elements in the array A to the format specified by the parameters in fpopts. We note that cpfloat is just an alias for cpf_fpround which ensures backward compatibility with previous versions of the library—the two functions rely on the same implementation. We also remark that X and A may point to the same memory location, in which case the rounding is performed in place.

The second group contains functions for simulating the four elementary arithmetic operations and many mathematical functions in low precision. The computation is first performed in the storage format, using the corresponding C operators or C mathematical library (math.h) functions, and the result is then rounded to the specified low-precision format.

The third and last group comprises functions that work on the implicit target-format representation. These can be used to extract the exponent and fraction of the number rounded to low precision, to compute the number that in the target format is closest to a given value in the storage format, and to check whether a value in the storage format would become subnormal, normal, infinite, or NaN once rounded to the target format.

We list all functions available in CPFloat in Table 2, and refer the reader to the package documentation for additional details. Most of the functions in the other two classes are inspired by the corresponding functions in the C mathematical library.

Table 2: CPFloat functions if double is used as storage format. The names of the corresponding functions for float are obtained by appending the suffix "f": for example, the function cpfloat_mulf can be used to multiply two float arrays elementwise. All functions require two extra arguments to define the number of elements in the input vectors (numelem) and the floating-point representation (fpopts). All mathematical functions return as result numbers in the target format.

Function	Description
cpfloat_validate_optstruct	Validate fields of optstruct struct variable.
cpfloat	Round double array to target format (legacy function name).
cpf_fpround	Round double array to target format.
cpf_add	Sum in target format.
cpf_sub	Difference in target format.
cpf_mul	Product in target format.
cpf_div	Ratio in target format.
cpf_cos	Trigonometric cosine.
cpf_sin	Trigonometric sine.
cpf_tan	Trigonometric tangent.
cpf_acos	Inverse trigonometric cosine.
cpf_asin	Inverse trigonometric sine.
cpf_atan	Inverse trigonometric tangent.
cpf_atan2	Two-argument arctangent.
cpf_cosh	Hyperbolic cosine.
cpf_sinh	Hyperbolic sine.
cpf_tanh	Hyperbolic tangent.
cpf_acosh	Inverse hyperbolic cosine.
cpf_asinh	Inverse hyperbolic sine.
cpf_atanh	Inverse hyperbolic tangent.
cpf_exp	Exponential.
cpf_frexp	Exponent and normalized fraction in target format.
cpf_ldexp	Scale number by power of 2.
cpf_log	Natural logarithm.
cpf_log10	Base-10 logarithm.

Table 2: (continued)

Function	Description
cpf_modf	Integral and fractional part.
cpf_exp2	Base-2 exponential.
cpf_expm1	$\exp(x) - 1$.
cpf_ilogb	Integral part of logarithm of absolute value.
cpf_log1p	Natural logarithm of number shifted by one.
cpf_log2	Base-2 logarithm.
cpf_scalbn	Scale number by power of FLT_RADIX.
cpf_scalbln	Scale number by power of FLT_RADIX.
cpf_pow	Real powers.
cpf_sqrt	Square root.
cpf_cbrt	Cube root.
cpf_hypot	Hypotenuse of a right-angle triangle.
cpf_erf	Error function.
cpf_erfc	Complementary error function.
cpf_tgamma	Gamma function.
cpf_lgamma	Natural logarithm of absolute value of gamma function.
cpf_ceil	Ceiling function.
cpf_floor	Floor function.
cpf_trunc	Integer truncation.
cpf_round	Closest integer (round-to-nearest).
cpf_lround	Closest integer (round-to-nearest).
cpf_llround	Closest integer (round-to-nearest).
cpf_rint	Closest integer with specified rounding mode.
cpf_lrint	Closest integer with specified rounding mode.
cpf_llrint	Closest integer with specified rounding mode.
cpf_nearbyint	Closest integer with specified rounding mode.
cpf_remainder	Remainder of floating point division.
cpf_remquo	Remainder and quotient of rounded numbers.
cpf_copysign	Compose number from magnitude and sign.
cpf_nextafter	Next number in specified direction in target format.
cpf_nexttoward	Next number in specified direction in target format.
cpf_fdim	Positive difference.
cpf_fmax	Element-wise maximum.
cpf_fmin	Element-wise minimum.
cpf_fpclassify	Categorize floating-point values.
cpf_isfinite	Check whether value is finite in target format.
cpf_isinf	Check whether value is infinite in target format.
cpf_isnan	Check if value is not a number in target format.
cpf_isnormal	Check whether value is normal in target format.
cpf_fabs	Absolute value.
cpf_fma	Fused multiply-add.

The example in Listing 1 is available in the CPFloat source code repository on GitHub.[1] As we explain in more details in Section 7, CPFloat can leverage OpenMP to run multiple threads in parallel, but the runtime overhead causes the parallel version of the code to be slower than the sequential

---

[1] https://github.com/mfasi/cpfloat

one on arrays with very few elements. To alleviate this issue, we use a sequential version of the code for small arrays and switch to the parallel variant only when the number of elements exceeds a machine-dependent threshold which we determine with an auto-tuning procedure. For the C library, the auto-tuning can be triggered with the command

```
$ make autotune
```

The example in Listing 1 can be compiled with

```
$ make example
```

which produces the binary `bin/example`. When executed, this program produces the following output.

```
The validation function returned 0.
X in binary64:
  1.666666666666667e+00  3.141592653589793e+00  2.718281828459045e+00
X rounded to binary16:
  1.666992187500000e+00  3.142578125000000e+00  2.718750000000000e+00
Sum rounded to binary16:
  3.167968750000000e+00  4.644531250000000e+00  4.218750000000000e+00
Product rounded to binary16:
  2.500000000000000e+00  4.714843750000000e+00  4.078125000000000e+00
Log rounded to binary16:
  5.112304687500000e-01  1.145507812500000e+00  1.000000000000000e+00
Angle rounded to binary16:
  7.329101562500000e-01  4.455566406250000e-01  5.043945312500000e-01
```

4   Floating-point storage formats and rounding

A family of binary floating-point numbers $\mathcal{F}\langle p, e_{\min}, e_{\max}, \mathfrak{s}_n \rangle$ is a finite set that includes a subset of the real line and a handful of values with a special meaning. In our notation, the three integer parameters p, $e_{\min}$, and $e_{\max}$ represent the number of binary digits of precision, the minimum exponent, and the maximum exponent, respectively, and the Boolean flag $\mathfrak{s}_n$ indicates whether subnormals are supported. A real number $x := (s, m, e)$ in $\mathcal{F}\langle p, e_{\min}, e_{\max}, \mathfrak{s}_n \rangle$ can be written as

$$x = (-1)^s \cdot m \cdot 2^{e-p+1}, \tag{4.1}$$

where s is the sign bit, set to 0 if x is positive and to 1 otherwise, the integral siginificand m is a natural number not greater than $2^p - 1$, and the exponent e is an integer between $e_{\min}$ and $e_{\max}$ inclusive.

In order to ensure a unique representation for all numbers in $\mathcal{F}\langle p, e_{\min}, e_{\max}, \mathfrak{s}_n \rangle \setminus \{0\}$, it is customary to normalize the system by assuming that if $x \geq 2^{e_{\min}}$ then $2^{p-1} \leq m \leq 2^p - 1$, that is, the number is represented using the smallest possible exponent and the largest possible significand. In such systems, the number $(s, m, e) \in \mathcal{F}\langle p, e_{\min}, e_{\max}, \mathfrak{s}_n \rangle \setminus \{0\}$ is normal if $m \geq 2^{p-1}$, and subnormal otherwise. The exponent of subnormals is always $e_{\min}$, and in a normalized system any number $x = (s, m, e) \neq 0$ has a unique p-digit binary expansion $(-1)^s \cdot \widetilde{m} \cdot 2^e$, where

$$\widetilde{m} = m \cdot 2^{1-p} = d_0 + \frac{d_1}{2} + \cdots + \frac{d_{p-1}}{2^{p-1}} = d_0.d_1 \ldots d_{p-1}, \tag{4.2}$$

for some $d_0, d_1, \ldots, d_{p-1} \in \{0, 1\}$, is called the normal significand of x. One can verify that if x is normal then $d_0 = 1$ and $1 \leq \widetilde{m} < 2$, whereas if x is subnormal then $d_0 = 0$ and $0 < \widetilde{m} < 1$. Conventionally, the signed zero values $+0$ and -0 are considered neither normal nor subnormal. In a normalized system, the smallest subnormal is $x_{\text{minsub}} := 2^{e_{\min}-p+1}$, whereas the smallest and largest positive normal numbers are $x_{\min} := 2^{e_{\min}}$ and $x_{\max} := 2^{e_{\max}}(2 - 2^{1-p})$, respectively. Their negative counterparts can be obtained by observing that a floating-point number system is symmetric with respect to 0. In our

notation, subnormals are kept if $\mathfrak{s}_n = \textbf{true}$, and rounded to either the closest smallest floating-point number or the zero of appropriate sign if $\mathfrak{s}_n = \textbf{false}$.

The floating-point numbers in $\mathcal{F}\langle p, e_{\min}, e_{\max}, \mathfrak{s}_n \rangle$ can be represented efficiently as binary strings. The sign is stored in the leftmost bit of the representation, and the following b_e bits are used to store the exponent. Under the IEEE 754 assumption that $e_{\min} = 1 - e_{\max}$, the most efficient way of representing the exponent is obtained by setting $e_{\max} = 2^{b_e - 1} - 1$ and using a representation biased by $e_{\max}$, so that $00\cdots01_2$ represents the smallest allowed exponent and $11\cdots10_2$ represents the largest. The trailing $p - 1$ bits are used to store the significand of x. It is not necessary to store the value of d_0 explicitly, as the IEEE standard uses an encoding often referred to as "implicit bit": d_0 is assumed to be 1 unless the binary encoding of the exponent is the all-zero string, in which case $d_0 = 0$ and x represents $+0$ or -0 if $m = 0$, or a subnormal number otherwise. If the exponent field contains the reserved all-one string, then the number represents $+\infty$ or $-\infty$ if the fraction field is set to 0, and a NaN otherwise. Infinities are needed to express values whose magnitude exceeds that of the largest positive and negative numbers that can be represented in $\mathcal{F}\langle p, e_{\min}, e_{\max}, \mathfrak{s}_n \rangle$, whereas NaNs represent the result of invalid operations, such as taking the square root of a negative number, dividing a zero by a zero, or multiplying an infinity by a zero. These are needed in order to ensure that the semantics of all floating-point operations is well specified and that the resulting floating-point number system is closed.

A rounding is an operator that maps a real number x to one of the floating-point numbers closest to x in a given floating-point family. The original IEEE 754 standard [27, Sec. 4] defines four rounding modes for binary formats: the default round-to-nearest with ties-to-even, and three directed rounding modes, round-toward-$+\infty$, round-toward-$-\infty$, and round-toward-zero. The 2008 revision of the standard [28] introduces a fifth rounding mode, round-to-nearest with ties-to-away, but states that it is not necessary for an implementation of a binary format to be IEEE compliant. Finally, the latest revision of the standard [29] recommends the use of round-to-nearest with ties-to-zero for augmented operations [29, Sec. 9.5]. Here we also consider two less common rounding strategies, round-to-odd [3] and stochastic rounding [9].

With round-to-odd, we have to set the least significant bit of the rounded number to 1 unless the infinitely-precise number is exactly representable in the target format. This rounding has applications in several domains. In computer arithmetic, for instance, it can be used to emulate a correctly rounded fused multiply–add (FMA) operator when a hardware FMA unit is not available [3]. In general, round-to-odd helps avoid issues associated with double rounding, and is used for this purpose in GCC,[2] to name one example. The fact that it can be implemented at low cost in hardware has recently prompted interest from the machine learning community [7].

Unlike all rounding modes discussed so far, stochastic rounding is non-deterministic, in the sense that the direction in which to round is chosen randomly and repeating the rounding may yield different results. The simplest choice is to round an infinitely-precise number that is not representable in the target format to either of the two closest representable floating-point numbers with equal probability. A more interesting flavor is obtained by rounding a non-representable number x to either of the closest floating-point numbers with probability proportional to the distance between x and the two rounding candidates. This rounding mode dates back to the fifties [17, 18], and has recently gained prominence owing to the surge of interest in low-precision floating-point formats. It is particularly effective at alleviating swamping in long sums [8] and ordinary differential equation solvers [26, 15], and at counteracting the loss of accuracy observed when the precision used to train neural networks is reduced [23, 48]. Stochastic rounding is not widely available in general-purpose hardware, but has started to appear in some specialized processors, such as the Intel Loihi neuromorphic chips [10] and the Graphcore IPU, an accelerator for machine learning [22]. We refer the reader to the recent survey [9] for more details on this rounding mode.

---

[2] https://gcc.gnu.org/bugzilla/show_bug.cgi?id=21718#c25

5 Efficient rounding to a lower-precision format

Now we discuss how to exploit the binary representation in the storage format to develop efficient algorithms for rounding $x \in \mathcal{F}^{(h)}$ to $\widetilde{x} \in \mathcal{F}^{(\ell)}$, where

$$\mathcal{F}^{(h)} := \mathcal{F}\langle p^{(h)}, e_{\min}^{(h)}, e_{\max}^{(h)}, \mathfrak{s}_n^{(h)} \rangle, \qquad \text{and} \qquad \mathcal{F}^{(\ell)} := \mathcal{F}\langle p, e_{\min}^{(\ell)}, e_{\max}^{(\ell)}, \mathfrak{s}_n^{(\ell)} \rangle, \tag{5.1}$$

and the same superscript notation is used for $x_{\min\text{sub}}$, $x_{\min}$, and $x_{\max}$. We make two assumptions:

A1. that $e_{\max}^{(\ell)} \leq e_{\max}^{(h)}$, which is a necessary—though not sufficient, as we will discuss later—condition for numbers in $\mathcal{F}^{(\ell)}$ to be representable exactly in $\mathcal{F}^{(h)}$; and

A2. that $p \leq p^{(h)}/2 - 1$, which guarantees that double rounding will be innocuous for the four elementary arithmetic operations ($+$, $-$, $\times$, and $\div$) and for the square root [45, 46].

We do not assume that the storage format supports subnormals, but we note that if that is not the case then one must additionally require that $e_{\min}^{(h)} \leq e_{\min}^{(\ell)} - p$, in order to ensure that the smallest subnormal number in $\mathcal{F}^{(\ell)}$ is representable as a normal number in $\mathcal{F}\langle p^{(h)}, e_{\min}^{(h)}, e_{\max}^{(h)}, \textbf{false} \rangle$.

We now list the high-level functions we need to operate on floating-point numbers. In the description, x denotes a floating-point number in $\mathcal{F}^{(h)}$, n denotes a positive integer, and i is an integer index between 0 and $p^{(h)} - 1$ inclusive. Unless otherwise stated, these functions are not defined for infinities and NaNs.

- ABS(x) returns the absolute value of x. For infinities, we use the definition ABS$(\pm\infty) = +\infty$.

- DIGIT(x, i) returns the ith digit of the significand of x from the left. The indexing starts from 0, so that DIGIT(x, i) is d_i in (4.2).

- EXPONENT(x) returns the exponent of x. This is the signed integer e in (4.1) if x is normal, and an integer $\gamma < e_{\min}$ otherwise.

- SIGNIFICAND(x) returns the integral significand of x, that is, the positive integer m in (4.1).

- RAND(n) returns a string of n randomly generated bits.

- SIGN(x) returns -1 if the floating-point number x is negative and $+1$ otherwise. This function behaves as expected for (signed) zeros and infinities, i.e., SIGN$(\pm 0) = \pm 1$ and SIGN$(\pm\infty) = \pm 1$.

- TAIL(x, i) returns the trailing $p^{(h)} - i$ bits of the significand of x as an unsigned integer. For infinities, this function returns 0 by convention, thus TAIL$(\pm\infty, i) = 0$ for any i.

- TRUNC(x, i) returns the number x with the last $p^{(h)} - i$ bits of the significand set to zero. Truncating an infinity leaves it unchanged, thus TRUNC$(\pm\infty, i) = \pm\infty$ for any i.

- ULP(x, i) returns the number $2^{\text{EXPONENT}(x) - i + 1}$, that is, the gap between x and its successor in a floating-point number system with i bits of precision. As noted by Muller [38], this function corresponds to the unit in the last place as defined by Overton [42, p. 14] and Goldberg [20].

How to implement these efficiently will be discussed in detail in Section 6.

Our rounding strategy is summarized in Algorithm 5.1. The function CPFLOAT computes the representation of the floating-point number $x \in \mathcal{F}^{(h)}$ in a lower-precision format $\mathcal{F}^{(\ell)}$. In the pseudocode, $\mathbb{N}^+$ denotes the set of positive integers. The input parameter ROUNDFUN is a pointer to one of the functions in Algorithm 5.2, 5.4, or 5.5. A call to ROUNDFUN(x, t, ζ) returns the floating-point number $\widetilde{x} \in \mathcal{F}^{(h)}$, that is, x rounded to $\mathcal{F}^{(\ell)}$. The function starts by setting the underflow threshold ζ, which corresponds to the smallest subnormal number if $\mathfrak{s}_n^{(\ell)} = \textbf{true}$ and to the smallest normal number if

---

Algorithm 5.1: Round a number from $\mathcal{F}^{(h)}$ to $\mathcal{F}^{(\ell)}$ in (5.1).

---

1 function $\text{CPFLOAT}(x \in \mathcal{F}^{(h)}, \mathcal{F}^{(\ell)}, \text{ROUNDFUN} : \mathcal{F}^{(h)} \times \mathbb{N}^+ \times \mathcal{F}^{(h)} \to \mathcal{F}^{(h)})$

2   **if** $\mathfrak{s}_n^{(\ell)}$ **then**

3 $\zeta \leftarrow x_{\text{minsub}}^{(\ell)}$

4   **else**

5 $\zeta \leftarrow x_{\text{min}}^{(\ell)}$

6   **if** $\text{ABS}(x) < x_{\text{min}}^{(\ell)}$ **and** $e_{\text{min}}^{(\ell)} > e_{\text{min}}^{(h)}$ **then**

7 $t \leftarrow p - (e_{\text{min}}^{(\ell)} - \text{EXPONENT}(x))$

8   **else**

9 $t \leftarrow p$

10 $\widetilde{x} \leftarrow \text{ROUNDFUN}(x, t, \zeta)$

$\mathfrak{s}_n^{(\ell)} = \textbf{false}$. This value will be used by ROUNDFUN to round numbers that are too small to be represented. Then the algorithm computes the number t of significant digits in the significand of the binary representation of $\widetilde{x}$. If x falls within the normal range of $\mathcal{F}^{(\ell)}$, then its significand has p significant digits and the algorithm sets $t = p$. If, on the other hand, $|x|$ is between $x_{\text{minsub}}^{(\ell)}$ and $x_{\text{min}}^{(\ell)}$, then the exponent of x is smaller than $e_{\text{min}}^{(\ell)}$ and the number t of significant binary digits may have to be reduced. If $e_{\text{min}}^{(\ell)} = e_{\text{min}}^{(h)}$, then x is subnormal and has the same number of leading zeros in both both storage and target format. Otherwise, the value of t is given by the difference between p and the number of leading zeros in the representation of $\widetilde{x}$, including the zero to the left of the binary point.

In the coming sections we discuss how the function ROUNDFUN can be implemented for the six rounding strategies we consider.

5.1 Round-to-nearest

Our algorithm for rounding a floating-point number in $\mathcal{F}^{(h)}$ to the closest floating-point number in the lower-precision format $\mathcal{F}^{(\ell)}$ is given in Algorithm 5.2. The pseudocode describes, in particular, a variant of round-to-nearest known as ties-to-even, whereby a number exactly in-between two floating-point numbers is rounded to the rounding candidate with even significand, that is, a number that has a 0 in the $t - 1$ position to the right of the binary point. Two other variants, ties-to-zero and ties-to-away, will be briefly examined at the end of this section.

Initially, the function checks if the number to be rounded is too small to be represented in $\mathcal{F}^{(\ell)}$. If subnormals are supported then the smallest representable number $x_{\text{minsub}}^{(\ell)}$ has an odd fraction, and a tie value $x \in \mathcal{F}^{(h)}$ such that $|x| = \zeta/2$ is rounded towards zero. If subnormals are not supported, then x is equidistant from two numbers with even significands, since the significand of the smallest normal number $x_{\text{min}}^{(\ell)}$ in $\mathcal{F}^{(\ell)}$ is even. We still round x to zero, as this behavior is consistent with that of the GNU MPFR library.

Next, a number $x \in \mathcal{F}^{(h)}$ that is too large to be rounded to zero but has absolute value below the threshold ζ is rounded to $\text{sign}(x) \cdot \zeta$. If $|x|$ is larger than the threshold, the algorithm checks whether x is within the dynamic range of $\mathcal{F}^{(\ell)}$: following the IEEE 754 standard [29, Sec. 4.3.1], x will overflow to infinity without changes in sign if $|x| < 2^{e_{\text{max}}^{(\ell)}}(2 - 2^{-p})$.

If x is within the range of numbers representable in $\mathcal{F}^{(\ell)}$, then the algorithm truncates $x \in \mathcal{F}^{(h)}$ to t significant digits to obtain $\widetilde{x} \in \mathcal{F}^{(\ell)}$, the largest number (in absolute value) that satisfies $\text{sign}(\widetilde{x}) = \text{sign}(x)$ and $|\widetilde{x}| \leq |x|$. In general, $\widetilde{x}$ is one of the two floating-point numbers in $\mathcal{F}^{(\ell)}$ closest to x, the other candidate being $\widetilde{x} + \text{sign}(x) \cdot \text{ULP}(\widetilde{x}, t)$. In order to choose a rounding direction, we examine the value of the discarded bits. The unsigned integer $d := \text{TAIL}(x, t)$ represents the trailing $p^{(h)} - t$ digits of the significand of x. Thus $0 \leq d \leq 2^{p^{(h)}-t} - 1$, and if $d < \gamma := 2^{p^{(h)}-t-1}$ then $\widetilde{x}$ is the result to be

---
Algorithm 5.2: Function for round-to-nearest with ties-to-even.
---
1 function ROUNDTONEAREST($x \in \mathcal{F}^{(h)}, t \in \mathbb{N}^+, \zeta \in \mathcal{F}^{(h)}$)

2   **if** ABS(x) $\leq \zeta$ **then**   ▷ Underflow.

3   **if** ABS(x) $\leq \zeta/2$ **then**

4 $\widetilde{x} \leftarrow$ SIGN(x) $\cdot 0$

5   **else**

6 $\widetilde{x} \leftarrow$ SIGN(x) $\cdot \zeta$

7   **else if** ABS(x) $\geq 2^{e_{\max}^{(\ell)}}(2 - 2^{-P})$ **then**   ▷ Overflow.

8 $\widetilde{x} \leftarrow$ SIGN(x) $\cdot +\infty$

9   **else**   ▷ In range.

10 $\widetilde{x} \leftarrow$ TRUNC(x, t)

11   **if** TAIL(x, p) $> 2^{p^{(h)}-t-1}$ **or** (TAIL(x, p) $= 2^{p^{(h)}-t-1}$ **and** DIGIT($x, t-1$) $= 1$) **then**

12 $\widetilde{x} \leftarrow \widetilde{x} +$ SIGN(x) $\cdot$ ULP($\widetilde{x}, t$)

13   **return** $\widetilde{x}$

---

returned, whereas if $d > \gamma$ then it is necessary to add (or subtract, if $\widetilde{x}$ is negative) ULP($\widetilde{x}, t$) in order to obtain the correctly rounded value. If $d = \gamma$, then we have a tie and we need to round to the nearest even number. Therefore we add sign(x) $\cdot$ ULP($\widetilde{x}, t$) if the bit in position $t - 1$ of the significand of $\widetilde{x}$ is a 1, and we leave $\widetilde{x}$ unchanged otherwise.

The latest revision of the IEEE 754 standard mentions two other tie-breaking rules for round-to-nearest: ties-to-zero, to be used for the recommended augmented operations, and ties-to-away, which is required for decimal formats. These can be implemented by changing the conditions of the if statements on lines 3 and 11 in Algorithm 5.2 to

  3 **if** ABS(x) $\leq \zeta/2$ **then**

     [...]

11 **if** TAIL(x, t) $> 2^{p^{(h)}-t-1}$ **then**

for ties-to-zero, and to

  3 **if** ABS(x) $< \zeta/2$ **then**

     [...]

11 **if** TAIL(x, t) $\geq 2^{p^{(h)}-t-1}$ **then**

for ties-to-away.

Note that this implementation preserves the sign of zero, maps infinities to infinities, and does not change the encoding of quiet and signaling NaN values. The same observation is true for the rounding functions in Algorithm 5.3, 5.4, and 5.5.

5.2   Round-to-odd

The function ROUNDTOODD in Algorithm 5.3 implements round-to-odd according to the definition in [3, Sec. 3.1], as this rounding mode is not part of the IEEE standard. Informally speaking, if x is exactly representable in $\mathcal{F}^{(\ell)}$, then the function returns it unchanged, otherwise it returns the number closest to x with an odd significand, that is, a significand with a trailing 1. In the spirit of the algorithms discussed so far, one could obtain $\widetilde{x}$ by truncating $|x|$ to the first t significant digits, checking the parity of the significand of $\widetilde{x}$, and adding ULP($\widetilde{x}, t$) to the result if $\widetilde{x}$ is even. However, in this case we know that the result of the truncation requires a correction only if the least significant digit of $\widetilde{x}$ is a 0, in which case adding ULP($\widetilde{x}, t$) amounts to setting that bit to 1. Therefore, we check the bits obliterated by the truncation, and if TAIL(x, t) $\neq 0$ then we conclude that x is not exactly representable in $\mathcal{F}^{(\ell)}$

Algorithm 5.3: Function for round-to-odd.

---

1 **function** ROUNDTOODD($x \in \mathcal{F}^{(h)}, t \in \mathbb{N}^+, \zeta \in \mathcal{F}^{(h)}$)

2   **if** ABS(x) $< \zeta$ **and** $x \neq 0$ **then**   ▷ Underflow.

3        $\widetilde{x} \leftarrow$ SIGN(x) $\cdot \zeta$

4   **else if** ABS(x) $> x_{\max}^{(\ell)}$ **and** ABS(x) $\neq +\infty$ **then**   ▷ Overflow.

5        $\widetilde{x} \leftarrow$ SIGN(x) $\cdot x_{\max}^{(\ell)}$

6   **else**   ▷ In range.

7        $\widetilde{x} \leftarrow$ TRUNC(x, t)

8   **if** TAIL(x, p) $\neq 0$ **then**

9   DIGIT($\widetilde{x}, t - 1$) $\leftarrow 1$

10   **return** $\widetilde{x}$

---

and that the significand of the rounded $\widetilde{x}$ must be odd. We can ensure this by setting the digit in position $t - 1$ of the significand of $\widetilde{x}$ to 1. In practice, this operation will have an effect only if that digit is not already set to 1, and in particular will have no effect when $t = 1$, as the leading digit of the significand is not stored explicitly. The core idea of this algorithm is the same as that of the second of the two methods for round-to-odd discussed in [3, Sec. 3.4].

The algorithm must round $\widetilde{x}$ to the closest odd number in $\mathcal{F}^{(\ell)}$ if it falls within the underflow or the overflow range. Since 0 is even, rounding to 0 is not an option when this rounding mode is used, and numbers too small to be represented must be rounded to the smallest floating-point number with an odd significand of corresponding sign. When subnormals are supported, the smallest representable number $x_{\text{minsub}}^{(\ell)}$ is odd, thus numbers smaller than $x_{\text{minsub}}^{(\ell)}$ in absolute value are simply rounded to $\text{sign}(x) \cdot x_{\text{minsub}}^{(\ell)}$. If subnormals are not supported, on the other hand, the smallest representable number $x_{\min}^{(\ell)}$ is even, and we are faced with a choice. We could round $\widetilde{x}$ such that $|\widetilde{x}| < x_{\min}^{(\ell)}$ to $\text{sign}(x) \cdot \left(x_{\min}^{(\ell)} + \text{ULP}(x_{\min}^{(\ell)}, t) \right)$, which is the closest number with odd significand. This choice, however, feels unnatural: the operator thus defined would not be rounding to either of the floating-point numbers closest to $\widetilde{x}$. In our pseudocode, we prefer to round $\widetilde{x}$ to the rounding candidate largest in magnitude, that is, $\text{sign}(x) \cdot x_{\min}^{(\ell)}$.

The definition given by Boldo and Melquiond cannot be applied directly to values in the overflow range, as in principle the significand of $\pm\infty$ is neither odd nor even. Since $-x_{\max}^{(\ell)}$ and $x_{\max}^{(\ell)}$ are necessarily odd, we prefer to round values outside the finite range of $\mathcal{F}^{(\ell)}$ to the closest finite number. Being exactly representable, infinities themselves represent an exception to this rule, and the algorithm leaves them unchanged.

5.3 Directed rounding

The functions in Algorithm 5.4 show how to implement the three directed rounding modes prescribed by the IEEE 754 standard. The idea underlying the three functions is similar to that discussed for the function ROUNDTONEAREST in Algorithm 5.2, the main differences being 1) the use of the sign, which is relevant when the rounding direction is not symmetric with respect to 0, and 2) the conditions under which a unit in the last place has to be added.

We start by discussing the function ROUNDTOWARDPLUSINFINITY. First, we check whether x is within the range of numbers that are representable in $\mathcal{F}^{(\ell)}$. Numbers that are too small to be represented are rounded up to 0 if negative and to ζ if positive. Finite positive numbers larger than the largest representable number $x_{\max}^{(\ell)}$ overflow to $+\infty$, whereas negative numbers smaller than the smallest representable number $-x_{\max}^{(\ell)}$ are rounded up to $-x_{\max}^{(\ell)}$, with the only exception of $-\infty$, which is

---
Algorithm 5.4: Functions for directed rounding modes.
---
1 **function** ROUNDTOWARDPLUSINFINITY($x \in \mathcal{F}^{(h)}, t \in \mathbb{N}^+, \zeta \in \mathcal{F}^{(h)}$)

2   **if** $x > 0$ **and** $x < \zeta$ **then**   ▷ Underflow.

3         $\widetilde{x} \leftarrow \zeta$

4   **else if** $x \leq 0$ **and** $x > -\zeta$ **then**

5         $\widetilde{x} \leftarrow \text{SIGN}(x) \cdot 0$

6   **else if** $x > x_{\max}^{(\ell)}$ **then**   ▷ Overflow.

7         $\widetilde{x} \leftarrow +\infty$

8   **else if** $x < -x_{\max}^{(\ell)}$ **and** $x \neq -\infty$ **then**

9         $\widetilde{x} \leftarrow -x_{\max}^{(\ell)}$

10   **else**   ▷ In range.

11         $\widetilde{x} \leftarrow \text{TRUNC}(x, t)$

12   **if** $\text{TAIL}(x, p) \neq 0$ **and** $x > 0$ **then**

13             $\widetilde{x} \leftarrow \widetilde{x} + \text{ULP}(\widetilde{x}, t)$

14   **return** $\widetilde{x}$

15 **function** ROUNDTOWARDMINUSINFINITY($x \in \mathcal{F}^{(h)}, t \in \mathbb{N}^+, \zeta \in \mathcal{F}^{(h)}$)

16   **if** $x < 0$ **and** $x > -\zeta$ **then**   ▷ Underflow.

17         $\widetilde{x} \leftarrow -\zeta$

18   **else if** $x \geq 0$ **and** $x < \zeta$ **then**

19         $\widetilde{x} \leftarrow \text{SIGN}(x) \cdot 0$

20   **else if** $\widetilde{x} > x_{\max}^{(\ell)}$ **and** $\widetilde{x} \neq +\infty$ **then**   ▷ Overflow.

21         $\widetilde{x} \leftarrow x_{\max}^{(\ell)}$

22   **else if** $\widetilde{x} < -x_{\max}^{(\ell)}$ **then**

23         $\widetilde{x} \leftarrow -\infty$

24   **else**   ▷ In range.

25         $\widetilde{x} \leftarrow \text{TRUNC}(x, t)$

26   **if** $\text{TAIL}(x, p) \neq 0$ **and** $x < 0$ **then**

27             $\widetilde{x} \leftarrow \widetilde{x} - \text{ULP}(\widetilde{x}, t)$

28   **return** $\widetilde{x}$

29 **function** ROUNDTOWARDZERO($x \in \mathcal{F}^{(h)}, t \in \mathbb{N}^+, \zeta \in \mathcal{F}^{(h)}$)

30   **if** $\text{ABS}(x) < \zeta$ **then**   ▷ Underflow.

31         $\widetilde{x} \leftarrow \text{SIGN}(x) \cdot 0$

32   **else if** $\text{ABS}(x) \geq x_{\max}^{(\ell)}$ **and** $\text{ABS}(x) \neq +\infty$ **then**   ▷ Overflow.

33         $\widetilde{x} \leftarrow \text{SIGN}(x) \cdot x_{\max}^{(\ell)}$

34   **else**   ▷ In range.

35         $\widetilde{x} \leftarrow \text{TRUNC}(x, t)$

36   **return** $\widetilde{x}$
---

handled below using the fact that $\text{TRUNC}(-\infty, t) = -\infty$ and $\text{TAIL}(-\infty, t) = 0$.

Next, the function computes $\widetilde{x}$, that is, the number x with significand truncated to t significant digits, and checks whether $\widetilde{x}$ is smaller than the smallest number representable in $\mathcal{F}^{(\ell)}$. The rounding can be easily performed by noting that the truncation $\widetilde{x}$ is the correct result if x is negative or exactly representable in $\mathcal{F}^{(\ell)}$. Otherwise, $\widetilde{x}$ is incremented by $\text{ULP}(\widetilde{x}, t)$.

The function $\text{ROUNDTOWARDMINUSINFINITY}$ is identical, modulo some sign adjustments to take the opposite rounding direction into account. The algorithm starts by checking that x is within the range of numbers representable in $\mathcal{F}^{(\ell)}$. Numbers between $-\zeta$ and 0 are rounded down to $-\zeta$, whereas those between 0 and ζ underflow and are flushed to 0. Numbers that are smaller than the smallest number representable in $\mathcal{F}^{(\ell)}$ overflow to $-\infty$, whereas finite numbers greater than the largest number representable in $\mathcal{F}^{(\ell)}$ are rounded to $x_{\max}^{(\ell)}$, the only exception in this case being $+\infty$. In order to round a number that falls within the range of $\mathcal{F}^{(\ell)}$, we first compute $\widetilde{x}$ by truncating the significand of x to t significant digits, and then subtract $\text{ULP}(\widetilde{x}, t)$ from $\widetilde{x}$ if x is negative and not exactly representable with a t-digit significand.

The function ROUNDTOWARDZERO is simpler than the other two considered in this section, as truncation is sufficient to correctly round the significand of x to t significant digits. Underflow and overflow are also easier to handle: finite numbers that are smaller than the smallest representable number in absolute value are flushed to 0, whereas numbers too large to be represented are rounded to the closest representable finite number.

5.4 Stochastic rounding

The functions in Algorithm 5.5 describe how to implement the two variants of stochastic rounding we are concerned with.

The function ROUNDSTOCHASTIC implements the strategy that rounds $x \in \mathcal{F}^{(h)}$ to one of the two closest floating-point numbers with probability proportional to the distance. First, the algorithm considers numbers in the underflow range, whose rounding candidates are 0 and the threshold value ζ, which equals $x_{\min\text{sub}}$ if subnormals are supported and $x_{\min}$ if they are not. The distance between $|x|$ and 0 depends not only on the significand but also on the magnitude of x, thus the algorithm starts by computing the two values e_x and m_x, which represent the exponent and the integral significand of x, respectively. Being the exponent of a floating-point number in $\mathcal{F}^{(h)}$, e_x can be much smaller than $e_{\min}^{(\ell)}$, in which case it may be necessary to rescale m_x in order to align its exponent to $e_{\min}^{(\ell)}$. This is achieved by multiplying m_x by $2^{e_x + 1 - e_{\min}^{(\ell)}}$. In the pseudocode we take the floor of the result in order to keep it integer, although this is not strictly necessary: we prefer to work with integer arithmetic here so that the integers generated by the random number generator can be used without any further post-processing. This is desirable not only from a performance point of view, but also because drawing floating-point numbers from the uniform distribution over an interval is not a trivial task, even when a good integer pseudo-random number generator is available [21]. Finally, the algorithm generates a $p^{(h)}$-digit random integer γ, which is used as a threshold to choose the rounding direction: if the discarded bits, interpreted as an unsigned integer, are larger than γ, then x is rounded away from zero, otherwise it is rounded towards zero.

The procedure for numbers in the representable range is easier. In this case it suffices to compute $\widetilde{x}$, the value of x truncated to t significant bits, and then generate a random integer r between 0 and $2^{p^{(h)} - t}$. Since $\text{TAIL}(x, t)$ represents the distance between x and $\widetilde{x}$, we increment the absolute value of $\widetilde{x}$ by $\text{ULP}(\widetilde{x}, t)$ if $\text{TAIL}(x, t) > r$, and leave it unchanged otherwise. For overflow, we use the threshold value that the IEEE 754 standard recommends for round-to-nearest, and round numbers whose absolute value after rounding is larger than the threshold $2^{e_{\max}^{(\ell)}}(2 - 2^{-p})$ to infinity, leaving the sign unchanged.

The function $\text{ROUNDSTOCHASTICEQUAL}$ deals with the simpler strategy that rounds x up or down with equal probability. Depending on the interval in which x falls, the function selects the two closest

---

Algorithm 5.5: Functions for stochastic rounding.

---

1 **function** ROUNDSTOCHASTIC($x \in \mathcal{F}^{(h)}, t \in \mathbb{N}^+, \zeta \in \mathcal{F}^{(h)}$)

2   **if** ABS(x) $< \zeta$ **then**   ▷ Underflow.

3 $e_{\min} \leftarrow$ EXPONENT(ζ)

4 $e_x \leftarrow$ EXPONENT(x)

5 $m_x \leftarrow$ SIGNIFICAND(x)

6 $t \leftarrow \lfloor m_x \cdot 2^{e_x+1-e_{\min}} \rfloor$

7   **if** $t >$ RAND($p^{(h)}$) **then**

8 $\widetilde{x} \leftarrow$ SIGN(x) $\cdot \zeta$

9   **else**

10 $\widetilde{x} \leftarrow$ SIGN(x) $\cdot 0$

11   **else**   ▷ In range or overflow.

12 $\widetilde{x} \leftarrow$ TRUNC(x, t)

13   **if** TAIL(x, t) $>$ RAND($p^{(h)} - p$) **then**

14 $\widetilde{x} \leftarrow \widetilde{x} +$ SIGN(x) $\cdot$ ULP($\widetilde{x}, t$)

15   **if** $\widetilde{x} \geq 2^{e^{(\ell)}_{\max}}(2 - 2^{-p})$ **then**   ▷ Overflow.

16 $\widetilde{x} \leftarrow$ SIGN(x) $\cdot +\infty$

17   **return** $\widetilde{x}$

18 **function** ROUNDSTOCHASTICEQUAL($x \in \mathcal{F}^{(h)}, t \in \mathbb{N}^+, \zeta \in \mathcal{F}^{(h)}$)

19   **if** ABS(x) $< \zeta$ **and** $x \neq 0$ **then**   ▷ Underflow.

20 $\widetilde{x} \leftarrow$ SIGN(x) $\cdot$ RANDSELECT($0, \zeta$)

21   **else if** ABS(x) $> x^{(\ell)}_{\max}$ **and** ABS(x) $\neq +\infty$ **then**   ▷ Overflow.

22 $\widetilde{x} \leftarrow$ SIGN(x) $\cdot$ RANDSELECT($x^{(\ell)}_{\max}, +\infty$)

23   **else**   ▷ In range.

24 $\widetilde{x} \leftarrow$ TRUNC(x, t)

25   **if** TAIL(x, t) $\neq 0$ **then**

26 $\widetilde{x} \leftarrow$ RANDSELECT($\widetilde{x}, \widetilde{x} +$ SIGN(x) $\cdot$ ULP($\widetilde{x}, t$))

27   **return** $\widetilde{x}$

28 **function** RANDSELECT($x \in \mathcal{F}^{(h)}, y \in \mathcal{F}^{(h)}$)

29   **if** RAND(1) $= 1$ **then**

30   **return** x

31   **else**

32   **return** y

---

representable numbers in $\mathcal{F}^{(\ell)}$ and calls the function RANDSELECT to select one of them with equal probability. In the pseudocode, we use a single bit generated randomly to discriminate between the two rounding directions.

6   Efficient implementation for IEEE-like representation formats

The subroutines used in Section 5 can be implemented efficiently if we assume that the numbers are represented using the floating-point format described in Section 4. First, we need to define the semantics of the operators for bit manipulation that we will rely on. These are available in most programming languages, although the notation varies greatly from language to language. For clarity, we use a prefix notation for all the operators.

Let a and b be strings of n bits. The bits are indexed from left to right, so that a_0 and a_{n-1} denote the leftmost and the rightmost bit of a, respectively. For $i \in \mathbb{N}$, we define the following operators.

- Conjunction: $c = \text{AND}(a, b)$ is an n-bit string such that $c_k = 1$ if a_k and b_k are both set to 1, and $c_k = 0$ otherwise.

- Disjunction: $c = \text{OR}(a, b)$ is an n-bit string such that $c_k = 1$ if at least one of a_k and b_k is set to 1, and $c_k = 0$ otherwise.

- Negation: $c = \text{NOT}(a)$ is an n-bit string such that $c_k = 1$ if $a_k = 0$, and $c_k = 0$ otherwise.

- Logical shift left: $c = \text{LSL}(a, i)$ is an n-bit string such that $c_k = 0$ if $k > (n-1) - i$, and $c_k = a_{k+i}$ otherwise.

- Logical shift right: $c = \text{LSR}(a, i)$ is an n-bit string such that $c_k = 0$ if $k < i$, and $c_k = a_{k-i}$ otherwise.

Most of the operations used in Section 5 require extracting a certain subset of the bits in the binary representation of the floating-point number $x \in \mathcal{F}\langle p, e_{\min}, e_{\max}, \mathfrak{s}_n \rangle$ while zeroing out the remaining ones. This can be achieved by taking the bitwise conjunction between the binary string that represents x and a bitmask, that is, a string as long as the binary representation of x that has ones in the positions corresponding to the bits to be extracted and zeros everywhere else. More generally, the functions in Section 5 can be implemented using the operators above as follows. In the descriptions, x denotes the binary floating-point representation of x, n denotes a positive integer, and i denotes an integer index between 0 and $p - 1$.

- $\text{ABS}(x)$ can be implemented as $\text{AND}(x, m_{\text{abs}})$, where m_{abs} is constituted by a single leading 0 followed by ones.

- $\text{DIGIT}(x, i)$ can be implemented as $\text{AND}(x, m_{\text{digit}})$, where m_{digit} has a 1 in the position corresponding to the digit to be extracted and 0 everywhere else. We note that checking whether this digit is 0 or 1 does not require any additional operations in programming languages such as C where 0 is interpreted as **false** and any other integer is interpreted as **true**.

- $\text{EXPONENT}(x)$ can be implemented as a sequence of logic and arithmetic operations. The raw bits of the exponents can be extracted with $c := \text{AND}(x, m_{\text{exp}})$, where m_{exp} has 1 in the positions corresponding to the exponent bits of the binary representation of x. This can be converted into the unsigned integer $\text{LSR}(c, p - 1)$, and the signed exponent can be obtained by subtracting the bias of the storage floating-point format. If x is subnormal in $\mathcal{F}\langle p, e_{\min}, e_{\max}, \mathfrak{s}_n \rangle$, then the value computed in this way is $-e_{\max} = e_{\min} - 1$, and the correct value to return in this case is $e_{\min} + \lambda$, where λ is the number of trailing zeros in the significand of x, including the implicit bit.

- SIGNIFICAND(x) can be implemented leveraging the function EXPONENT. The digits to the right of the radix point can be obtained as c := AND(x, m_{frac}) where m_{frac} is the bitmask that has the $p-1$ trailing bits set to 1 and the remaining bits set to 0. If $x_{min} \leq |x| \leq x_{max}$, then EXPONENT$(x) > e_{min}$ and the implicit bit must be set to 1. This can be achieved by using OR(c, LSL(1, $p-1$)), for instance.

- RAND(n) can be implemented by concatenating numbers produced by a pseudo-random number generator. Two m-bit strings a and b can be joined together by OR(LSL(a, m), b), and the unnecessary bits can be set to zero using a suitable bitmask.

- SIGN(x) is relatively expensive to implement by means of bit manipulation. However, note that we only need to compute the product sign$(x) \cdot y$ where y is a positive floating-point number. This operation can be implemented as OR(AND(x, m_{sign}), y), where m_{sign} is the bitmask with a leading 1 followed by zeros and the string y denotes the floating-point representation of y.

- TAIL(x, i) can be implemented as AND(x, m_{tail}), where the trailing $p-i$ bits of m_{tail} are set to 1 and the remaining bits are set to 0. This way, bits i to $p-1$ of the significand of x are preserved while the rest of the bits, including those representing the sign and exponent, are set to zero.

- TRUNC(x, i) can be implemented as AND(x, m'), where m' = NOT(m_{tail}). This way, bits i to $p-1$ of the significand of x are set to zero while the rest of the bits of x, including the exponent and sign bits, are preserved.

- ULP(x, p) is a rather expensive function to implement, because it requires extracting the exponent from the binary representation of x and then performing arithmetic operations on it. Increasing or decreasing x by ULP(x, p), on the contrary, can be achieved efficiently using only one bit shift and one integer arithmetic operation. In particular, it suffices to add to the binary representation of x, seen as an unsigned integer, a number that has 0 everywhere but in the pth digit, that is, the digit in position $p-1$ of the significand. We note that this technique could fail if $x = \pm\infty$, since adding ULP(x, p) in this fashion would turn infinities into NaNs. It is easy to check that this is not a problem in our setting, as we only add or subtract ULP(x, p) when x is finite and nonzero.

It is possible to implement some of the rounding routines even more efficiently by extending to other rounding modes the technique for round-to-nearest with ties-to-even developed on [31, p. 2–17], which manipulates the binary representation of the floating-point number by using only integer arithmetic. As a demonstration, here we show the concrete values of the bitmasks, expressed in hexadecimal notation, that one would use to round a binary32 number y to binary16. These methods generalize easily to other combinations of storage and target formats, and we describe this in general for the conversion of a floating-point number $x \in \mathcal{F}\langle p^{(h)}, e_{min}^{(h)}, e_{max}^{(h)}, \mathfrak{s}_n^{(h)} \rangle$ to $\mathcal{F}\langle p, e_{min}^{(\ell)}, e_{max}^{(\ell)}, \mathfrak{s}_n^{(\ell)} \rangle$.

We denote the 32-bit string containing the floating-point representation of y by y, and use the uppercase Latin letters X and Y to denote the unsigned integers that can be obtained by interpreting x and y as unsigned integers in radix 2. All the usual underflow and overflow checks are not included here—the aim is to demonstrate the core ideas for performing each type of rounding efficiently. We recall that the sign of a floating-point number can be determined by checking the leftmost bit, and that x (resp. y) is negative if x_0 (resp. y_0) is set and positive otherwise. The rounding modes amenable to this approach can be implemented as follows.

- Round-to-nearest with ties-to-even: isolate the bit in position $p-1$ of the significand of x, and then compute TRUNC$(X + \text{LSR}(m_{tail}, 1) + \text{DIGIT}(x, p-1), p)$. When rounding a binary32 number to binary16, the formula becomes TRUNC$(Y + \text{0x7FFF} + y_{15}, 16)$.

- Round-to-nearest with ties-to-away: return $\textsc{trunc}(X + \textsc{lsl}(0\text{x}1, p^{(h)} - p - 1), p)$, which in our example becomes $\textsc{trunc}(Y + 0\text{x}8000, 16)$.

- Round-to-nearest with ties-to-zero: return $\textsc{trunc}(X + \textsc{lsr}(m_{\text{tail}}, 1), p)$, which in our example becomes $\textsc{trunc}(Y + 0\text{x}7\text{FFF}, 16)$.

- Round-toward-$+\infty$: return $\textsc{trunc}(X, p)$ if x_0 is set and $\textsc{trunc}(X + m_{\text{tail}}, p)$ otherwise. For our example, return $\textsc{trunc}(Y, 16)$, if y_0 is set and $\textsc{trunc}(Y + 0\text{x}\text{FFFF}, 16)$ if not.

- Round-toward-$-\infty$: return $\textsc{trunc}(X, p)$ if x_0 is not set and $\textsc{trunc}(X + m_{\text{tail}}, p)$ otherwise. For our example, return $\textsc{trunc}(Y, 16)$ if y_0 is not set and $\textsc{trunc}(Y + 0\text{x}\text{FFFF}, 16)$ otherwise.

- Round-toward-zero: return $\textsc{trunc}(X, p)$. For our example, return $\textsc{trunc}(Y, 16)$.

- Stochastic rounding with proportional probabilities: return $\textsc{trunc}(X + \textsc{rand}(p^{(h)} - t), p)$. For our example, return $\textsc{trunc}(Y + \textsc{rand}(13), 13)$.

7   Implementation and validation of the code

Our C implementation of the algorithms discussed in Section 5 and Section 6 is available on GitHub. The code is provided as a header-only library to allow the user to take advantage of the inlining feature of the C language for maximum efficiency. This also enhances the portability of the code, as packaging of the binaries and installation of the library are not required. The main drawback of this approach is a longer compilation time, which we alleviate by dividing the library into two separate units, one for each supported storage format.

In order to achieve better performance on large data, our functions work directly on C arrays. All the algorithms discussed in Section 5 are embarrassingly parallel, and each element of an array can be rounded independently from all the others. Therefore, our code can leverage the OpenMP library, if available on the system in use.

In general, OpenMP brings significant gains in terms of performance, but greatly increases the execution time for arrays with a limited number of elements. This well-known phenomenon is caused by the additional overhead of synchronization and loop scheduling [5]. While negligible for large arrays, these can be significant when only a small amount of work is allocated to each OpenMP thread. The impact of this overhead is hard to quantify in general, as it depends on the hardware platform as well as the number of OpenMP threads and the compiler used [6]. Our library contains both a parallel and a sequential version of each function, but we were unable to provide a single threshold that would allow the code to switch automatically from one variant to the other for optimal performance. Thus we devised a simple auto-tuning strategy that tries to determine the optimal threshold for the system in use by timing the rounding function on several arrays of different lengths and performing a binary search.

For generating the pseudo-random numbers required for stochastic rounding, we rely on algorithms from the family of permuted congruential generators developed by O'Neill [41], who provides a pure C implementation available on GitHub.[3] In our code we use the functions `pcg32_random_r` and `pcg64_random_r` to generate 32-bit and 64-bit random numbers, respectively; we initialize the random number generators with `pcg32_srandom_r` and `pcg64_srandom_r`, and advance them with `pcg32_advance_r` and `pcg64_advance_r`, respectively. As initial state, we use the current time as returned by `time(NULL)`. Use of the—considerably slower—default C pseudo-random number generator is also supported.

In order to validate our code experimentally, we wrote a suite of extensive unit tests. We describe in detail how we tested the rounding routines—all other functions in the library are tested by relying

---

[3] https://github.com/imneme/pcg-c

on them. We considered two storage formats, binary32 and binary64, which are available in C via the native data types `float` and `double`, respectively, and two target formats, binary16 and bfloat16. For each combination of storage and target formats, we performed three types of tests.

First we checked that all the numbers that can be represented exactly in the target format, including subnormals and special values such as infinities and NaNs, are not altered by any of the rounding routines. As the target formats we consider do not have an unduly large cardinality, we can test that this property is true for all representable numbers.

To check the correctness of the code when rounding is necessary, exhaustive testing is not an option, as the storage formats contain too many distinct values. In this case, we opted for testing only a set of representative values. For deterministic rounding, the correctness of the function can be assessed by checking that the output of the rounding routine matches the value predicted by the definition. For each pair of numbers $x_1, x_2 \in \mathcal{F}^{(\ell)}$ such that x_1 and x_2 are consecutive in $\mathcal{F}^{(\ell)}$ and $x_1 < x_2$, we considered five values in $\mathcal{F}^{(h)}$: nextafter($x_1, +\infty$), nextafter($x_m, -\infty$), x_m, nextafter($x_m, +\infty$), and nextafter($x_2, -\infty$). Here nextafter(x, y) denotes the next number in $\mathcal{F}^{(h)}$ after x in the direction of y, and x_m denotes the mid point between x_1 and x_2. We used the same technique for numbers in the underflow range, whereas for testing the correctness of overflow we used the values nextafter($\pm x_{\max}^{(\ell)}, \pm\infty$), nextafter($\pm x_{\text{bnd}}^{(\ell)}, \mp\infty$), $\pm x_{\text{bnd}}^{(\ell)}$, nextafter($\pm x_{\text{bnd}}^{(\ell)}, \pm\infty$), where $x_{\text{bnd}}^{(\ell)} := 2^{e_{\max}^{(\ell)}}(2 - 2^{-p})$ is the IEEE 754 threshold for overflow in round-to-nearest.

This technique would not work for stochastic rounding, as each value that is not representable in $\mathcal{F}^{(\ell)}$ can be rounded to two different values. We produced a test set by taking, for each pair of numbers $x_1, x_2 \in \mathcal{F}^{(\ell)}$ such that x_1 and x_2 are consecutive in $\mathcal{F}^{(\ell)}$ and $x_1 < x_2$, the numbers $(3x_1 + x_2)/4$, $(x_1 + x_2)/2$, and $(x_1 + 3x_2)/4$. We rounded each number 1,000 times and confirmed that the rounding routines always return either x_1 or x_2, and that the empirical probability distribution matches the expected one. We validated the correctness of the implementation for values in the underflow range by using the same technique, whereas for inputs in the overflow range we repeated the test on the three values: $(3x_{\max}^{(\ell)} + x_{\text{bnd}}^{(\ell)})/4$, $(x_{\max}^{(\ell)} + x_{\text{bnd}}^{(\ell)})/2$, and $(x_{\max}^{(\ell)} + 3x_{\text{bnd}}^{(\ell)})/4$.

The Makefile target

```
$ make ctest
```

runs the test suite for the C implementations.

We designed a MEX interface for MATLAB and Octave which is in charge of parsing and checking the input, allocating the output, and calling our library to perform the rounding. In order to show that the interface is fully compatible with chop, we designed a set of tests by modifying the default test suite for chop.[4] These test can be run with

```
$ make mtest
```

in MATLAB, and with

```
$ make otest
```

in Octave.

8 Performance evaluation

The experiments were run on a machine equipped with a 12-core Intel Xeon CPU E5-2690 v3 running at 2.60GHz. Exclusive node access was used in order to avoid timing artifacts and to ensure that all 24 CPU threads were available for parallel runs. The C code was compiled with version 9.3.0 of the GNU Compiler Collection (GCC) with the optimization flag `-O3` and the architecture option `-march=native`. The MATLAB experiments were run using the 64-bit GNU/Linux version of MATLAB 9.10 (R2021a). For our parallel implementations, we used version 4.5 of the OpenMP library. We compare the three versions of our codes.

---

[4] https://github.com/higham/chop/blob/master/test_chop.m

- `cpfloat_seq` denotes the sequential C codes in our library. Rounding is performed using the algorithms in Section 6.

- `cpfloat` denotes the C codes in our library that leverage OpenMP and employ the auto-tuning technique discussed in Section 7 to switch between sequential and parallel implementations. Rounding is performed using the algorithms in Section 6.

- `cpfloat_ml` denotes our MEX interface to `cpfloat` compiled in MATLAB. For large matrices, this function relies on the parallel version of our C codes.

In the plots, we use the shorthand notation $\langle f_s \mid f_t \rangle$ to denote the conversion of numbers in the storage format f_s to the target format f_t. As the numerical validation of the code has already been discussed in Section 7, here we focus on timings. We time the C or C++ code by comparing the value returned by the function `clock_gettime` with `CLOCK_MONOTONIC` before and after the execution, and take the median of 1,000 repetitions in order to reduce the influence of possible outliers. For the MATLAB code, we rely on the function `timeit`, which runs a portion of code several times and returns the median of the measurements.

8.1 Performance of the C interface

In this section we compare the performance of `CPFloat`, GNU MPFR, and FloatX by considering the following implementations.

- `chop_mpfr` denotes the codes that rely on the GNU MPFR library.[5] As storage format, we use the GNU MPFR custom data type `mpfr_t`. For rounding, our implementation sets the precision and exponent range of MPFR to the precision and exponent range of the target format and then converts the binary64 input using the function `mpfr_set_d`. Arithmetic operations use `mpfr_t` arrays for both input and output.

- `floatx` denotes the codes that use the `floatx` class from the FloatX library.[6] Arrays of binary64 numbers are converted to a lower-precision target format by invoking the constructor of the `floatx` class. This code requires that the parameters of the target format be specified at compile time, as the `floatx` class uses C++ templates that are instantiated only for the low-precision formats declared in the source code. For arithmetic operations, we use arrays of `floatx` objects as input and output.

- `floatxr` denotes the codes that rely on the `floatxr` class in the FloatX library. Conversion is performed using the class constructor. This function is more flexible than `floatx` in that the number of bits of precision and the maximum exponent allowed for the target format can be specified at runtime. As with the other implementations, for arithmetic operations we assume that input and output are arrays of `floatxr` objects.

Figure 1 compares the time required by `chop_mpfr`, `floatxr`, `floatx`, and `cpfloat_seq` to perform two operations: converting a square matrix from the storage format to the target format and computing the sum of two matrices in the target format. For format conversion, we report timings both including (left) and excluding (middle) the time needed to allocate the output vector. For the sum we report only the time needed to execute the operation (right), assuming that the memory to store the result has already been allocated.

As storage and target formats, we use binary64 and binary16, respectively. We repeated the experiment simulating bfloat16 and TensorFloat-32 arithmetic—we do not reproduce these results here as

---

[5] https://www.mpfr.org/
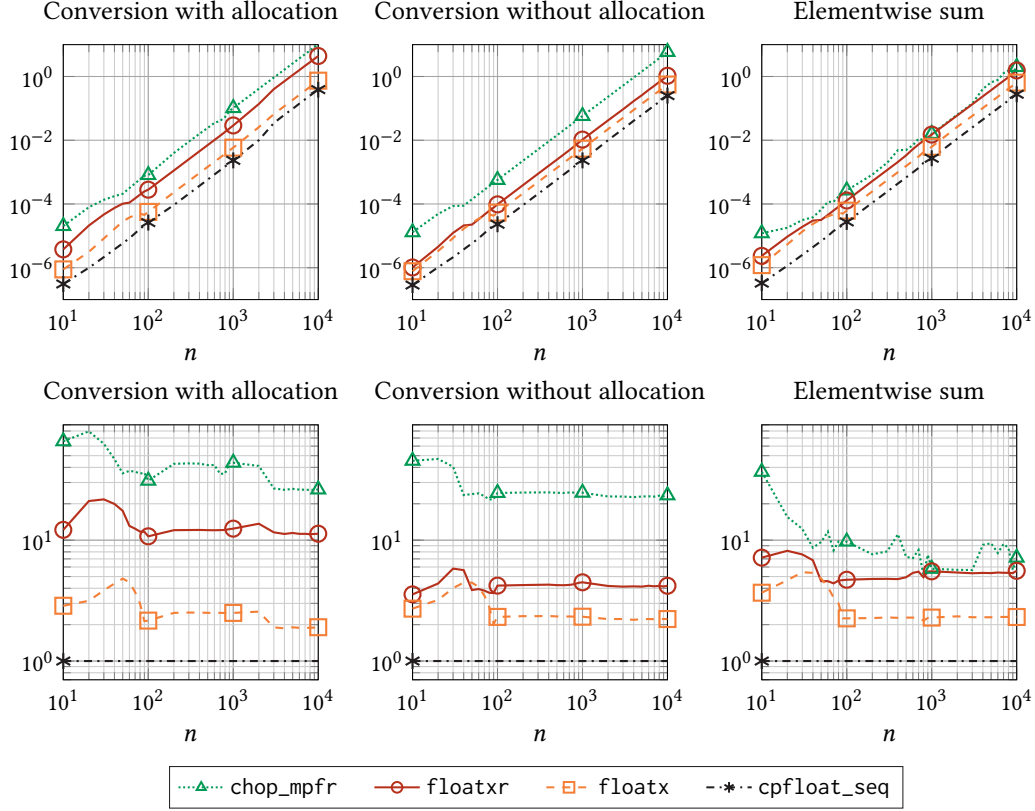[6] https://github.com/oprecomp/FloatX

Figure 1: Top: execution time, in seconds, of `chop_mpfr`, `floatxr`, `floatx`, and `cpfloat_seq` to convert matrices of order n from binary64 to binary16 (left and middle) and to compute the sum of two matrices of size n in binary16 (right) using round-to-nearest with even-on-ties. Bottom: Corresponding slowdown plots, where `cpfloat_seq` is taken as baseline. The timings include the allocation of the output vector only in the left-most panel.

they are indistinguishable from those for binary16. We use only round-to-nearest with ties-to-even, as the FloatX library currently does not support any other rounding modes. We observe, however, that `chop_mpfr` also supports directed rounding as prescribed by the IEEE 754 floating-point standard.

The plots in the top row report the median timing of 100 runs for each algorithm. The plots in the bottom row present the same data as slowdown with respect to the timings of `cpfloat_seq`. Unsurprisingly, the execution time of the four algorithms grows quadratically with the order of the matrix to be converted and thus linearly with the number of entries in the matrix. In all cases, `cpfloat_seq` is the fastest of the four implementations we consider, and its better performance is particularly evident for matrices of small size. Typically, `floatx` is two to three times slower than `cpfloat_seq`, but can reach a slowdown factor up to five for small matrices. For the two remaining algorithms, the performance varies depending on the operation being considered. The slowdown factor of `floatxr` is typically over ten for conversion if the time needed to allocate the output is factored in, but goes below five when allocations are not considered. The performance of `chop_mpfr`, on the other hand, seems to depend mostly on what operation is performed: the slowdown factor is always over 20 for data conversion, but generally below 10 for sums.

We remark that `chop_mpfr`, `floatxr`, and `cpfloat_seq` are more flexible than `floatx`, as the latter requires the parameters of the target format be known at compile time, in order for the compiler to instantiate the templates appropriately.
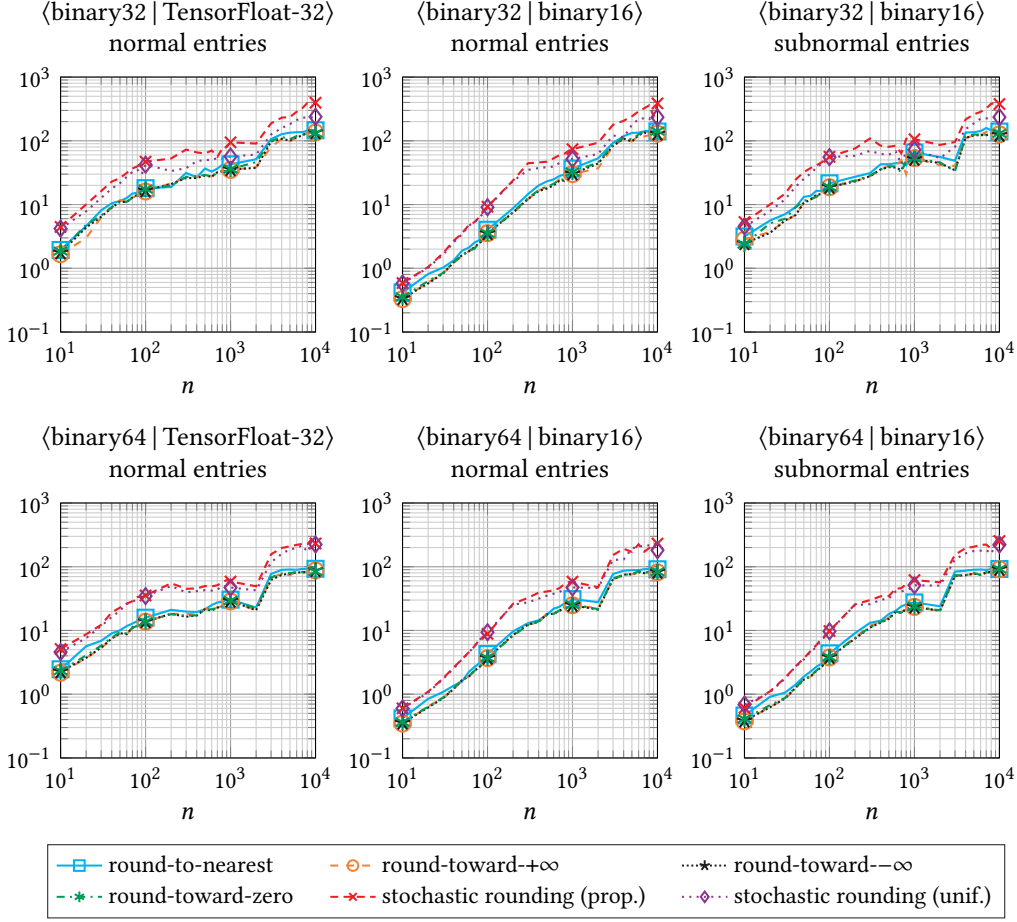
Figure 2: Ratio of the execution times of `chop` to that of `cpfloat_ml` on $n \times n$ matrices of normal (left and middle) and subnormal (right) floating-point numbers stored in binary32 (top row) and binary64 (bottom row).

8.2 Performance of the MATLAB interface

In this section we compare the performance of `cpfloat_ml` with that of existing MATLAB alternatives.

First, we consider the MATLAB function chop, which is available on GitHub.[7] Figure 2 reports the speedup of `cpfloat_ml` over chop. In each plot we consider the conversion of square matrices of size n between 10 and 10,000 using the six rounding modes implemented in chop. As target formats, we consider TensorFloat-32 (left column) and binary16 (middle and right columns). As storage format, we consider both binary32 (top row) and binary64 (bottom row).

The input data is obtained by manipulating the entries of an $n \times n$ matrix X of pseudo-random numbers uniformly distributed in $(0, 1)$. To generate a matrix of normal numbers, we add to each entry of X the constant value $x_{\min}^{(\ell)}$, which guarantees that x_{ij} is distributed uniformly in the interval $\left(x_{\min}^{(\ell)}, 1 + x_{\min}^{(\ell)}\right)$ for $i, j = 1, \ldots, n$. To produce a matrix with subnormal elements, on the other hand, we scale the entries of X by $x_{\min}^{(\ell)} - x_{\text{minsub}}^{(\ell)}$, and then add $x_{\text{minsub}}^{(\ell)}$ to the result, which guarantees that all the entries of X belong to the interval $(x_{\text{minsub}}^{(\ell)}, x_{\min}^{(\ell)})$, and are thus subnormal in the target precision. We do not report the results of the conversion to TensorFloat-32 for matrices of subnormals, as these did not differ significantly from those for normal entries.

In all six cases we consider, the speedup is greater than one for matrices of order 40 or more,

---
[7]https://github.com/higham/chop

and increases with the size of the input matrix. The low performance we register on small matrices depends on an unavoidable overhead caused by calling C code via a MEX interface—we examine this more in detail in Section 8.3. The two rounding modes for which the new algorithms bring the most significant gains are the two flavors of stochastic rounding. This is expected: generating pseudo-random numbers accounts for a large fraction of the execution time for this rounding mode, and by using a more efficient pseudo-random number generator the new algorithms have a great advantage over chop. Round-to-nearest and round-toward-zero show a very similar speedup, just below that of stochastic rounding. The performance gain is somewhat lower for round-toward-$-\infty$ and round-toward-$+\infty$, and the two rounding modes typically achieve very similar results. This is not surprising, as the underlying algorithms essentially perform the same operations.

We now look at the differences. When the target format is TensorFloat-32, the speedup is always above one, which implies that cpfloat_ml is always faster than chop in this case. We believe that this is due to the fact that chop treats TensorFloat-32 as a custom format, which implies some pre-processing overhead which is noticeable with inputs of small dimension. To confirm this, we repeated the experiments using as target format the other built-in low-precision format, bfloat16. The speedups were indistinguishable from those of binary16, which suggests that the fact that binary32 and TensorFloat-32 allocate the same number of bits to the exponent field does not play a role in this case.

We also note that when binary32 is used as storage format, chop is slower at converting matrices of subnormal than matrices of normals, but this phenomenon does not seem to affect matrices stored using binary64 format.

We compared the performance of cpfloat_ml and of the MATLAB function f_d_dec2floatp from the FLOATP_Toolbox[8] in a similar way. This library is less efficient than chop at the task we examine: it is typically over 100 times slower than cpfloat_ml at simulating binary16 and bfloat16 arithmetic, and always over 1,000 slower at simulating TensorFloat-32 arithmetic.

8.3 Overhead of the MATLAB interface

As a final test, we consider the overhead introduced by MATLAB when calling the underlying C implementation of the rounding algorithms. In Figure 3 we compare the execution time required to convert a matrix to binary16 with a direct call to the C code (left column) or with a call to the MEX interface in MATLAB (middle column). As the performance of the two algorithms is very similar and the data in the two series is hard to compare directly, we provide the speedup in the third column. As done in previous experiments, we repeat the experiment for both binary32 (top row) and binary64 (bottom row). We remark that the C functions were tuned by using the make autotune command, whereas for the MEX interface we used cpfloat_autotune, a MATLAB function included in the software package.

The raw timings show that in our implementations stochastic rounding is the slowest rounding mode. The performance of the other rounding modes is so similar that the lines are hard to distinguish for both the C and the MEX interface. The data in the right-most column shows that for both storage formats we consider, the overhead of the MEX interface is significant for small matrices, but becomes negligible for matrices of order 3,000 or more.

Our results seem to suggest that MATLAB code that requires the functionalities of cpfloat_ml should be translated into C in order to obtain the maximum efficiency. We would like to stress, however, that this translation might bring only a minor performance gain, and in fact not be worth the effort, unless the cpfloat function is used extensively on matrices of small size. In fact, the overhead of the MEX interface is modest in absolute terms, and is noticeable only in cases when the overall execution time of both cpfloat and cpfloat_ml is below 5 milliseconds. This suggests that, in most cases, switching to a pure C implementation would bring only a marginal benefit, if any.
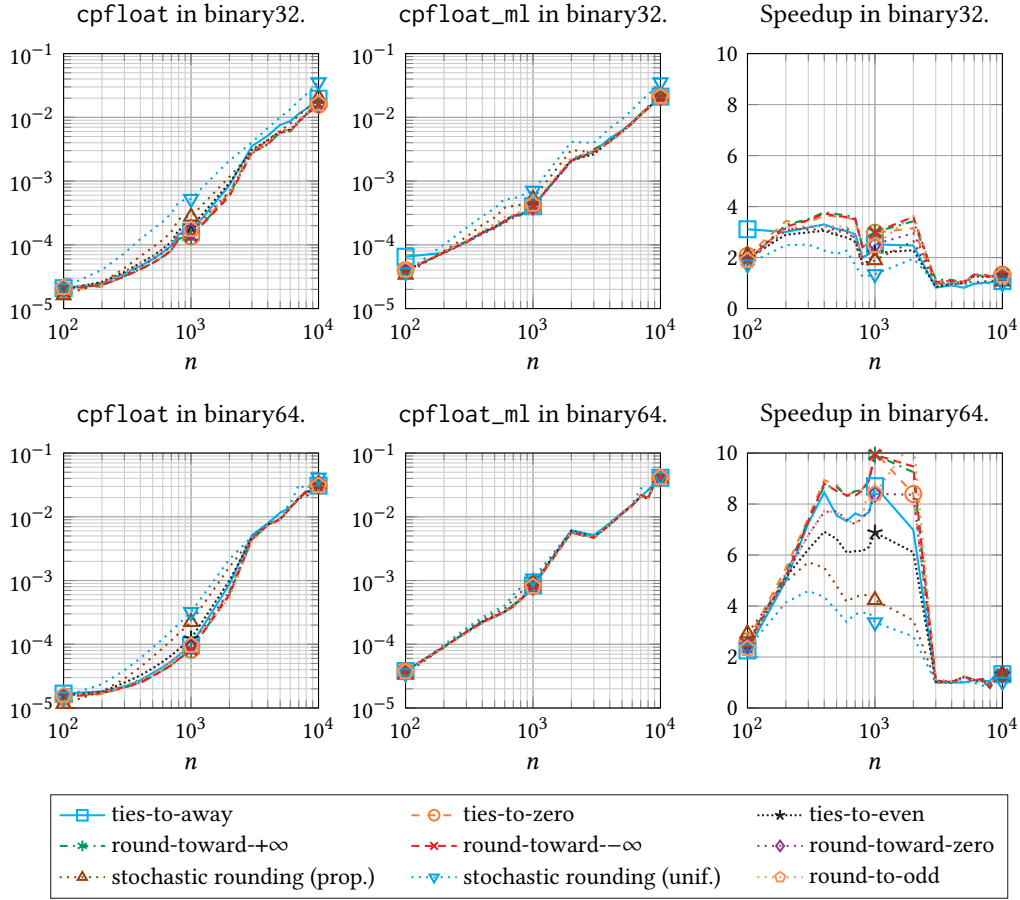
---

[8]https://gerard-meurant.pagesperso-orange.fr/floatp.zip

Figure 3: Execution time, in seconds, of `cpfloat` (first column) and `cpfloat_ml` (second column) on matrices of increasing order n and target format binary16. The third column represents the ratio of the execution time in the first column to that in the second.

9 Summary and future work

Motivated by the growing number of tools and libraries for simulating low-precision arithmetic, mainly driven by the explosion of floating-point formats made available in today's hardware, we considered the problem of rounding floating-point numbers to low precision in software. We developed low-level algorithms for a number of rounding modes, explained how to implement them efficiently using bit manipulation, and how to validate their behavior experimentally by means of exhaustive testing. We developed CPFloat, an efficient C library that implements all the algorithms discussed here from within MATLAB and Octave by means of a MEX interface we provide. When used in C, CPFloat can act as a full custom-precision floating-point arithmetic library as we have provided library calls for elementary arithmetic operations and mathematical functions such as the ones available for binary32 and binary64 in `math.h`. Our experimental results show that the new implementations outperform existing alternatives in C, C++, and MATLAB.

Traditionally, floating-point arithmetic has been the most widely adopted technique for working with non-integer numbers in high-performance scientific computing, but alternative methods have recently begun to gain popularity. In particular, we believe that the techniques we developed here could be adapted to posit arithmetic [24, 12], a generalization of the IEEE 754 floating-point number format, and to fixed-point arithmetic, a de facto standard technique for working with reals on systems that are not equipped with a floating-point unit. This will be the subject of future work.

Acknowledgments

References

[1] ARM LIMITED, *Arm architecture reference manual*, Tech. Report ARM DDI 0487F.c (ID072120), Mar. 2020.

[2] G. AUPY, A. BENOIT, A. CAVELAN, M. FASI, Y. ROBERT, H. SUN, AND B. UÇAR, *Coping with silent errors in HPC applications*, in https://doi.org/10.1007/978-3-319-46376-6, A. Adamatzky, ed., Springer-Verlag, Cham, Switzerland, Nov. 2016, pp. 269–292.

[3] S. BOLDO AND G. MELQUIOND, *Emulation of a FMA and correctly rounded sums: Proved algorithms using rounding to odd*, IEEE Trans. Comput., 57 (2008), p. 462–471.

[4] E. BRUN, D. DEFOUR, P. DE OLIVEIRA CASTRO, M. ISTOAN, D. MANCUSI, E. PETIT, AND A. VAQUET, *A study of the effects and benefits of custom-precision mathematical libraries for HPC codes*, IEEE Trans. Emerg. Topics Comput., 9 (2021), p. 1467–1478.

[5] J. M. BULL, *Measuring synchronisation and scheduling overheads in OpenMP*, in Proceedings of First European Workshop on OpenMP, Lund, Sweden, Sept. 1999, Wiley, pp. 99–105.

[6] J. M. BULL, F. REID, AND N. McDONNELL, *A microbenchmark suite for OpenMP tasks*, Lecture Notes in Computer Science, 7312 (2012), p. 271–274.

[7] N. BURGESS, J. MILANOVIC, N. STEPHENS, K. MONACHOPOULOS, AND D. MANSELL, *Bfloat16 processing for neural networks*, in Proceedings of the 26th IEEE Symposium on Computer Arithmetic, Kyoto, Japan, June 2019, Institute of Electrical and Electronics Engineers, pp. 88–91.

[8] M. P. CONNOLLY, N. J. HIGHAM, AND T. MARY, *Stochastic rounding and its probabilistic backward error analysis*, SIAM J. Sci. Comput., 43 (2021), pp. A566–A585.

[9] M. CROCI, M. FASI, N. J. HIGHAM, T. MARY, AND M. MIKAITIS, *Stochastic rounding: Implementation, error analysis, and applications*, MIMS EPrint 2021.17, Manchester Institute for Mathematical Sciences, The University of Manchester, UK, Oct. 2021. Revised January 2022. To appear in Roy. Soc. Open Sci.

[10] M. DAVIES, N. SRINIVASA, T.-H. LIN, G. CHINYA, Y. CAO, S. H. CHODAY, G. DIMOU, P. JOSHI, N. IMAM, S. JAIN, Y. LIAO, C.-K. LIN, A. LINES, R. LIU, D. MATHAIKUTTY, S. McCOY, A. PAUL, J. TSE, G. VENKATARAMANAN, Y.-H. WENG, A. WILD, Y. YANG, AND H. WANG, *Loihi: A neuromorphic manycore processor with on-chip learning*, IEEE Micro, 38 (2018), pp. 82–99.

[11] A. DAWSON AND P. D. DÜBEN, *rpe v5: An emulator for reduced floating-point precision in large numerical simulations*, Geosci. Model Dev., 10 (2017), pp. 2221–2230.

[12] F. DE DINECHIN, L. FORGET, J.-M. MULLER, AND Y. UGUEN, *Posits: The good, the bad and the ugly*, in Proceedings of the Conference for Next Generation Arithmetic, Singapore, Singapore, Mar. 2019, ACM Press, pp. 1–10.

[13] M. Fasi, N. J. Higham, F. Lopez, T. Mary, and M. Mikaitis, *Matrix multiplication in multiword arithmetic: Error analysis and application to GPU tensor cores*, MIMS EPrint 2022.3, Manchester Institute for Mathematical Sciences, The University of Manchester, UK, Jan. 2022.

[14] M. Fasi, N. J. Higham, M. Mikaitis, and S. Pranesh, *Numerical behavior of NVIDIA tensor cores*, PeerJ Comput. Sci., 7 (2021), pp. e330(1–19).

[15] M. Fasi and M. Mikaitis, *Algorithms for stochastically rounded elementary arithmetic operations in IEEE 754 floating-point arithmetic*, IEEE Trans. Emerg. Topics Comput., 9 (2021), p. 1451–1466.

[16] G. Flegar, F. Scheidegger, V. Novaković, G. Mariani, A. E. Tomás, A. C. I. Malossi, and E. S. Quintana-Ortí, *FloatX: A C++ library for customized floating-point arithmetic*, ACM Trans. Math. Software, 45 (2019), p. 1–23.

[17] G. E. Forsythe, *Round-off errors in numerical integration on automatic machinery*, Bull. Amer. Math. Soc., 56 (1950), pp. 55–64.

[18] G. E. Forsythe, *Reprint of a note on rounding-off errors*, SIAM Rev., 1 (1959), p. 66–67.

[19] L. Fousse, G. Hanrot, V. Lefèvre, P. Pélissier, and P. Zimmermann, *MPFR: A multiple-precision binary floating-point library with correct rounding*, ACM Trans. Math. Software, 33 (2007), pp. 13:1–13:15.

[20] D. Goldberg, *What every computer scientist should know about floating-point arithmetic*, ACM Comp. Surv., 23 (1991), p. 5–48.

[21] F. Goualard, *Generating random floating-point numbers by dividing integers: A case study*, in Computational Science – ICCS 2020, V. V. Krzhizhanovskaya, G. Závodszky, M. H. Lees, J. J. Dongarra, P. M. A. Sloot, S. Brissos, and J. Teixeira, eds., Lecture Notes in Computer Science, Cham, Switzerland, 2020, Springer-Verlag, p. 15–28.

[22] Graphcore Limited, *IPU programmer's guide*, 2020.

[23] S. Gupta, A. Agrawal, K. Gopalakrishnan, and P. Narayanan, *Deep learning with limited numerical precision*, in Proceedings of the 32nd International Conference on Machine Learning, F. Bach and D. Blei, eds., vol. 37 of Proceedings of Machine Learning Research, Lille, France, July 2015, PMLR, pp. 1737–1746.

[24] J. L. Gustafson and I. T. Yonemoto, *Beating floating point at its own game: Posit arithmetic*, Supercomputing Frontiers and Innovations, 4 (2017), pp. 71–86.

[25] N. J. Higham and S. Pranesh, *Simulating low precision floating-point arithmetic*, SIAM J. Sci. Comput., 41 (2019), pp. C585–C602.

[26] M. Hopkins, M. Mikaitis, D. R. Lester, and S. Furber, *Stochastic rounding and reduced-precision fixed-point arithmetic for solving neural ordinary differential equations*, Philos. Trans. R. Soc. A, 378 (2020).

[27] *IEEE Standard for Binary Floating-Point Arithmetic, ANSI/IEEE Standard 754-1985*, Institute of Electrical and Electronics Engineers, Piscataway, NJ, USA, Oct. 1985. Reprinted in SIGPLAN Notices, 22(2):9–25, 1987.

[28] *IEEE Standard for Floating-Point Arithmetic, IEEE Std 754-2008 (revision of IEEE Std 754-1985)*, Institute of Electrical and Electronics Engineers, Piscataway, NJ, USA, Aug. 2008.

[29] *IEEE Standard for Floating-Point Arithmetic, IEEE Std 754-2019 (revision of IEEE Std 754-2008)*, Institute of Electrical and Electronics Engineers, Piscataway, NJ, USA, July 2019.

[30] Intel Corporation, *BFLOAT16—Hardware Numerics Definition*, Nov. 2018. White paper. Document number 338302-001US.

[31] Intel Corporation, *Intel architecture instruction set extensions and future features programming reference*, Mar. 2020.

[32] International Business Machines Corporation, *Power ISA version 3.0 B*, 2017.

[33] V. Lefèvre, *SIPE: a Mini-Library for Very Low Precision Computations with Correct Rounding*. hal-00864580, version 1, Sept. 2013.

[34] V. Lefèvre, *SIPE: Small integer plus exponent*, in Proceedings of the 21st IEEE Symposium on Computer Arithmetic, Austin, TX, USA, Apr. 2013, Institute of Electrical and Electronics Engineers, pp. 99–106.

[35] C. Lichtenau, S. Carlough, and S. M. Mueller, *Quad precision floating point on the IBM z13*, in Proocedings of the 23rd IEEE Symposium on Computer Arithmetic, Silicon Valley, USA, July 2016, Institute of Electrical and Electronics Engineers, pp. 87–94.

[36] G. Meurant, *FLOATP_Toolbox*, 2020. Matlab software, variable precision floating point arithmetic.

[37] C. B. Moler, *"Half precision" 16-bit floating point arithmetic*. Blog post, Dec. 2017.

[38] J.-M. Muller, *On the definition of $\mathrm{ulp}(x)$*, Research Report RR-5504, LIP RR-2005-09, INRIA, LIP, May 2005.

[39] NVIDIA Corporation, *NVIDIA Tesla P100 GPU architecture*, Tech. Report WP-08019-001_v01.1, 2016.

[40] NVIDIA Corporation, *NVIDIA A100 tensor core GPU architecture*, 2020. NVIDIA whitepaper v1.0.

[41] M. E. O'Neill, *PCG: A family of simple fast space-efficient statistically good algorithms for random number generation*, Tech. Report HMC-CS-2014-0905, Harvey Mudd College, Claremont, CA, Sept. 2014.

[42] M. L. Overton, *Numerical Computing with IEEE Floating Point Arithmetic*, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2001.

[43] Radeon Technologies Group, *Radeon's next-generation Vega architecture*, tech. report, Advanced Micro Devices, 2017. File no longer available on the AMD website. Archived version at https://en.wikichip.org/w/images/a/a1/vega-whitepaper.pdf.

[44] A. Reinarz, J.-M. Gallard, and M. Bader, *Influence of a-posteriori subcell limiting on fault frequency in higher-order DG schemes*, in Proceedings of the 8th IEEE/ACM Workshop on Fault Tolerance for HPC at eXtreme Scale (FTXS), Piscataway, NJ, USA, Nov. 2018, Institute of Electrical and Electronics Engineers, pp. 79–86.

[45] P. Roux, *Innocuous double rounding of basic arithmetic operations*, J. Formaliz. Reason., 7 (2014), pp. 131–142.

[46] S. M. Rump, *IEEE754 precision-k base-β arithmetic inherited by precision-m base-β arithmetic for $k < m$*, ACM Trans. Math. Software, 43 (2017), p. 1–15.

[47] P. Samfass, T. Weinzierl, A. Reinarz, and M. Bader, *Doubt and redundancy kill soft errors—towards detection and correction of silent data corruption in task-based numerical software*, in Proceedings of the 11th IEEE/ACM Workshop on Fault Tolerance for HPC at eXtreme Scale, Piscataway, NJ, USA, Nov. 2021, Institute of Electrical and Electronics Engineers.

[48] N. Wang, J. Choi, D. Brand, C.-Y. Chen, and K. Gopalakrishnan, *Training deep neural networks with 8-bit floating point numbers*, in Advances in Neural Information Processing Systems 31, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, eds., Curran Associates, Montreal, Canada, Dec. 2018, pp. 7675–7684.

[49] J. H. Wilkinson, *Rounding Errors in Algebraic Processes*, Notes on Applied Science No. 32, Her Majesty's Stationery Office, London, 1963. Also published by Prentice-Hall, Englewood Cliffs, NJ, USA. Reprinted by Dover, New York, 1994.

[50] T. Zhang, Z. Lin, G. Yang, and C. De Sa, *QPyTorch: A low-precision arithmetic simulation framework*, in Proceedings of the 5th Workshop on Energy Efficient Machine Learning and Cognitive Computing - NeurIPS Edition, Piscataway, NJ, USA, Dec. 2019, Institute of Electrical and Electronics Engineers, pp. 11–13.