

2.6. Fused Multiply-Add Operation

Some computers have a *fused multiply-add* (FMA) operation that enables a floating point multiplication followed by an addition or subtraction, $x * y + z$ or $x * y - z$, to be performed as though it were a single floating point operation. An FMA operation therefore commits just one rounding error:

$$fl(x \pm y * z) = (x \pm y * z)(1 + \delta), \quad |\delta| \leq u.$$

The Intel IA-64 architecture, as first implemented on the Intel Itanium chip, has an FMA instruction [273, 1999], as did the IBM RISC System/6000 and IBM Power PC before it. The Itanium's FMA instruction enables a multiply and an addition to be performed in the same number of cycles as one multiplication or addition, so it is advantageous for speed as well as accuracy.

An FMA is a welcome capability in that it enables the number of rounding errors in many algorithms to be approximately halved. Indeed, by using FMAs an inner product $x^T y$ between two n -vectors can be computed with just n rounding errors instead of the usual $2n - 1$, and any algorithm whose dominant operation is an inner product benefits accordingly.

Opportunities to exploit an FMA are frequent. Consider, for example, Newton's method for solving $f(x) = a - 1/x = 0$. The method is

$$\begin{aligned} x_{k+1} &= x_k - \frac{f(x_k)}{f'(x_k)} = x_k - \frac{a - 1/x_k}{x_k^{-2}} \\ &= x_k + (1 - x_k a)x_k, \end{aligned}$$

and its quadratic convergence can be seen from $1 - x_{k+1}a = (1 - x_k a)^2$. This method was used on early computers to implement reciprocation in terms of multiplication and thence division as $a/b = a * (1/b)$ [551, 1946], [1222, 1997], and this technique is still in use [714, 1997]. Since the computation of x_{k+1} from x_k can be expressed as two multiply-adds, the method is attractive when an FMA is available; an FMA also has the advantage of enabling a correctly rounded quotient to be achieved more easily [562, 1996, §A.7]. The floating point divide algorithms for the IA-64 architecture use this Newton-based approach [273, 1999].

An FMA makes it possible to compute an exact representation of the product of two floating point numbers x and y : by computing $\hat{a} = fl(xy)$ and $\hat{b} = fl(xy - \hat{a})$ with two FMAs, $\hat{a} + \hat{b} \equiv xy$ (see Problem 2.26). Furthermore, clever use of an FMA enables highly accurate computation of, for example, the determinant of a 2×2 matrix (see Problem 2.27).

However, the presence of an FMA introduces some tricky programming language issues [700, 1996]. If a programmer writes `a*d + c*b` how is this expression to be compiled for a machine with an FMA? There are three possibilities—two using the FMA and one without it—and they can all yield different answers. An example of where this question is particularly important is the product of complex numbers

$$(x + iy)*(x + iy) = x^2 + y^2 + i(xy - yx).$$

The product is obviously real, and the right-hand side evaluates as real in IEEE arithmetic, but if an FMA is employed in evaluating $xy - yx$ then a nonreal result will generally be obtained.

In the course of solving the quadratic $ax^2 - 2bx + c = 0$ for x , the expression $\sqrt{b^2 - ac}$ must be computed. Can the discriminant under the square root evaluate to a negative number when $b^2 \geq ac$? In correctly rounded arithmetic the answer is no: the monotonicity of rounding implies $fl(b^2) - fl(ac) \geq 0$ and the floating point subtraction of these incurs a small relative error and hence produces a nonnegative result. However, evaluation as $fl(fl(b^2) - ac)$ using an FMA can produce a negative result (for example, if $b^2 = ac$ and $fl(b^2) < b^2$).

In conclusion, as Kahan [700, 1996] puts it, “[FMAs] should not be used indiscriminately”. Unfortunately, compiler writers, in their efforts to obtain maximum performance, may not give programmers the capability to inhibit FMAs in those subexpressions where their use can have undesirable effects.

2.7. Choice of Base and Distribution of Numbers

What base β is best for a floating point number system? Most modern computers use base 2. Most hand-held calculators use base 10, since it makes the calculator easier for the user to understand (how would you explain to a naive user that 0.1 is not exactly representable on a base 2 calculator?). IBM mainframes traditionally have used base 16. Even base 3 has been tried—in an experimental machine called SETUN, built at Moscow State University in the late 1950s [1208, 1960].

Several criteria can be used to guide the choice of base. One is the impact of wobbling precision: as we saw at the end of §2.1, the spread of representation errors is smallest for small bases. Another possibility is to measure the worst-case representation error or the mean square representation error. The latter quantity depends on the assumed distribution of the numbers that are represented. Brent [160, 1973] shows that for the logarithmic distribution the worst-case error and the mean square error are both minimal for (normalized) base 2, provided that the most significant bit is not stored explicitly.

The logarithmic distribution is defined by the property that the proportion of base β numbers with leading significant digit n is

$$\log_{\beta}(n+1) - \log_{\beta} n = \log_{\beta}\left(1 + \frac{1}{n}\right).$$

It appears that in practice real numbers *are* logarithmically distributed. In 1938, Benford [102, 1938] noticed, as had Newcomb [889, 1881] before him, that the early pages of logarithm tables showed greater signs of wear than the later ones. (He was studying dirty books!) This prompted him to carry out a survey of 20,229 “real-life” numbers, whose decimal representations he found matched the logarithmic distribution closely.

The observed logarithmic distribution of leading significant digits has not been fully explained. Some proposed explanations are based on the assumption that the actual distribution is scale invariant, but this assumption is equivalent to the observation to be explained [1170, 1984]. Barlow [67, 1981], [68, 1981], [70, 1988] and Turner [1169, 1982], [1170, 1984] give useful insight by showing that if uniformly distributed numbers are multiplied together, then the resulting distribution converges to the logarithmic one; see also Boyle [157, 1994]. Furthermore, it is an interesting result that the leading significant digits of the numbers q^k ,

$k = 0, 1, 2, \dots$, are logarithmically distributed if q is positive and is not a rational power of 10; when $q = 2$ and the digit is 7 this is Gelfand's problem [939, 1981, pp. 50–51].

The nature of the logarithmic distribution is striking. For decimal numbers, the digits 1 to 9 are not equally likely to be a leading significant digit. The probabilities are as follows:

| | | | | | | | | |
|-------|-------|-------|-------|-------|-------|-------|-------|-------|
| 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 |
| 0.301 | 0.176 | 0.125 | 0.097 | 0.079 | 0.067 | 0.058 | 0.051 | 0.046 |

As an example, here is the leading significant digit distribution for the elements of the inverse of one random 100×100 matrix from the normal $N(0, 1)$ distribution:

| | | | | | | | | |
|-------|-------|-------|-------|-------|-------|-------|-------|-------|
| 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 |
| 0.334 | 0.163 | 0.100 | 0.087 | 0.077 | 0.070 | 0.063 | 0.056 | 0.051 |

For an entertaining survey of work on the distribution of leading significant digits see Raimi [968, 1976] (and also the popular article [967, 1969]).

2.8. Statistical Distribution of Rounding Errors

Most rounding error analyses, including all the ones in this book, are designed to produce worst-case bounds for the error. The analyses ignore the signs of rounding errors and are often the result of many applications of the triangle inequality and the submultiplicative inequality. Consequently, although the bounds may well give much insight into a method, they tend to be pessimistic if regarded as error estimates.

Statistical statements about the effect of rounding on a numerical process can be obtained from statistical analysis coupled with probabilistic models of the rounding errors. For example, a well-known rule of thumb is that a more realistic error estimate for a numerical method is obtained by replacing the dimension-dependent constants in a rounding error bound by their square root; thus if the bound is $f(n)u$, the rule of thumb says that the error is typically of order $\sqrt{f(n)}u$ (see, for example, Wilkinson [1232, 1963, pp. 26, 102]). This rule of thumb can be supported by assuming that the rounding errors are independent random variables and applying the central limit theorem. Statistical analysis of rounding errors goes back to one of the first papers on rounding error analysis, by Goldstine and von Neumann [501, 1951].

As we noted in §1.17, rounding errors are *not* random. See Problem 2.10 for an example of how two rounding errors cancel in one particular class of computations. Forsythe [424, 1959] points out that rounding errors do not necessarily behave like independent random variables and proposes a random form of rounding (intended for computer testing) to which statistical analysis is applicable.

Henrici [564, 1962], [565, 1963], [566, 1964] assumes models for rounding errors and then derives the probability distribution of the overall error, mainly in the context of difference methods for differential equations. Hull and Swenson [650, 1966] give an insightful discussion of probabilistic models, pointing out that “There is no claim that ordinary rounding and chopping are random processes, or that successive errors are independent. The question to be decided is whether or

not these particular probabilistic models of the processes will adequately describe what actually happens” (see also the ensuing note by Henrici [567, 1966]). Kahan [699, 1996] notes that “The fact that rounding errors are neither random nor uncorrelated will not in itself preclude the possibility of modelling them usefully by uncorrelated random variables. What will jeopardize the utility of such models is their failure to mimic important properties that actual rounding errors possess.” In the last sentence Kahan is referring to results such as Theorem 2.5.

Several authors have investigated the distribution of rounding errors under the assumption that the significands of the operands are from a logarithmic distribution and for different modes of rounding; see Barlow and Bareiss [72, 1985] and the references therein.

Other work concerned with statistical modelling of rounding errors includes that of Tienari [1137, 1970] and Linnainmaa [789, 1975]; see also the discussion of the CESTAC and PRECISE systems in §26.5.

2.9. Alternative Number Systems

The floating point format is not the only means for representing numbers in finite precision arithmetic. Various alternatives have been proposed, though none has achieved widespread use.

A particularly elegant system is the “level index arithmetic” of Clenshaw, Olver, and Turner, in which a number $x > 1$ is represented by $\xi = l + f$, where $f \in [0, 1]$ and

$$x = e^{e^{\dots e^f}}, \quad \text{or} \quad f = \ln(\ln(\dots(\ln x)\dots)),$$

where the exponentiation or logarithm is performed l times (l is the “level”). If $0 < x < 1$, then x is represented by the reciprocal of the representation for $1/x$. An obvious feature of the level index system is that it can represent much larger and smaller numbers than the floating point system, for similar word lengths. A price to be paid is that as the size of the number increases the precision with which it is represented shrinks. If l is sufficiently large then adjacent level index numbers are so far apart that even their exponents in base 10 differ. For very readable introductions to level index arithmetic see Clenshaw and Olver [240, 1984] and Turner [1171, 1991], and for more details see Clenshaw, Olver, and Turner [241, 1989]. Level index arithmetic is somewhat controversial in that there is disagreement about its advantages and disadvantages with respect to floating point arithmetic; see Demmel [310, 1987]. A number system involving levels has also been proposed by Matsui and Iri [826, 1981]; in their system, the number of bits allocated to the significand and exponent is allowed to vary (within a fixed word size).

Other number systems include those of Swartzlander and Alexopolous [1115, 1975], Matula and Kornerup [831, 1985], and Hamada [539, 1987]. For summaries of alternatives to floating point arithmetic see the section “Alternatives to Floating-Point—Some Candidates”, in [241, 1989], and Knuth [744, 1998, Chap. 4].

2.10. Elementary Functions

The term elementary functions refers to the base 2, e and 10 logarithms and exponentials, and the trigonometric and hyperbolic functions and their inverses. These functions are much harder to implement than the elementary operations $+$, $-$, $*$, $/$ and square root and, as noted earlier, the IEEE arithmetic standard does not have anything to say about elementary functions.

In the past, elementary function libraries were of variable quality, but with increased demand for high-quality floating point arithmetic, the ability now to build on IEEE standard elementary operations, and continuing research into elementary function evaluation, the quality of algorithms and libraries is improving.

Ideally, we would like the computed value of an elementary function at a machine number to be the correctly rounded result. The tablemaker's dilemma (see §1.2) casts doubt on whether this goal is achievable. However, Lindemann's result that the exponential of a nonzero algebraic number cannot be algebraic implies that the exponential of a nonzero *machine* number cannot be a machine number or half-way between two machine numbers. Therefore for a given machine number x and precision t there is a finite m so that computing m digits of e^x is sufficient to determine e^x correctly rounded to t digits. Similar arguments apply to the other elementary functions, and so the tablemaker's dilemma does not occur in the context of floating point arithmetic [778, 1998]. In extensive computations, Lefèvre, Muller, and Tisserand [778, 1998], [777, 2001] have determined the maximum value of m over all double precision IEEE arithmetic numbers ($t = 53$) for each of the elementary functions. They find that m is usually close to $2t = 106$. This knowledge of the worst-case m makes the goal of correctly rounded results achievable.

Other desirable properties of computed elementary functions are preservation of monotonicity (e.g., $x < y \Rightarrow e^x < e^y$), preservation of symmetries (e.g., $\sin(x) = -\sin(-x)$), and preservation of mathematical relations and identities (e.g., $\sin(x) \in [-1, 1]$ and $\cos^2(x) + \sin^2(x) = 1$). Unfortunately, these requirements can conflict. Lefèvre, Muller, and Tisserand [778, 1998] note that in IEEE single precision arithmetic the machine number closest to $\arctan(2^{30})$ is slightly greater than $\pi/2$, so a correctly rounded approximation to $\arctan(2^{30})$ necessarily violates the requirement that $\arctan(x) \in [-\pi/2, \pi/2]$.

Describing methods for evaluating elementary functions is outside the scope of this book. Details can be found in the excellent book by Muller [876, 1997], which is the best overall reference for this topic. Algorithms for evaluating elementary functions in IEEE arithmetic are developed by Tang [1121, 1989], [1123, 1990], [1125, 1992], Gal and Bachelis [451, 1991], and Ziv [1285, 1991]. Tang [1124, 1991] gives a very readable description of table lookup algorithms for evaluating elementary functions, which are used in a number of current chips and libraries.

Algorithms for evaluating complex elementary functions that exploit exception handling and assume the availability of algorithms for the real elementary functions are presented by Hull, Fairgrieve, and Tang [649, 1994]. For details of how elementary functions are evaluated on many of today's pocket calculators see Schelin [1016, 1983], as well as Muller [876, 1997].

Table 2.3. *Test arithmetics.*

| Hardware | Software | $ 3 \times (4/3 - 1) - 1 ^a$ |
|-----------------------------------|------------------------------------|-------------------------------|
| Casio fx-140 (≈ 1979) | | 1×10^{-9} |
| Casio fx-992VB (≈ 1990) | | 1×10^{-13} |
| HP 48G (1993) | | 1×10^{-11} |
| Sharp EL-5020 (1994) | | 0.0^b |
| Pentium III | MATLAB 6.1 (2001) | $2.2 \dots \times 10^{-16}$ |
| 486DX | WATFOR-77 ^c V3.0 (1988) | $2.2 \dots \times 10^{-16}$ |
| 486DX | FTN 90 ^d (1993) | $2.2 \dots \times 10^{-16}$ |
| 486DX | MS-DOS QBasic 1.1 | $1.1 \dots \times 10^{-19}^e$ |

^aIntegers in the test expression are typed as real constants 3.0, etc., for the Fortran tests.

^b 1×10^{-9} if $4/3$ is stored and recalled from memory.

^cWATCOM Systems Inc.

^dSalford Software/Numerical Algorithms Group, Version 1.2.

^e $2.2 \dots \times 10^{-16}$ if $4/3$ is stored and recalled from a variable.

2.11. Accuracy Tests

How can you test the accuracy of the floating point arithmetic on a computer or pocket calculator? There is no easy way, though a few software packages are available to help with the tasks in specific programming languages (see §27.6). There are, however, a few quick and easy tests that *may* reveal weaknesses. The following list is far from comprehensive and good performance on these tests does not imply that an arithmetic is correct. Results from the tests are given in Tables 2.4–2.5 for the selected floating point arithmetics described in Table 2.3, with incorrect digits in boldface and underlined>. Double precision was used for the compiled languages. The last column of Table 2.3 gives an estimate of the unit roundoff (see Problem 2.14). The estimate produced by QBasic indicates that the compiler used extended precision in evaluating the estimate.

- (Cody [249, 1982]) Evaluate $\sin(22) = -8.8513\ 0929\ 0403\ 8759\ 2169 \times 10^{-3}$ (shown correct to 21 digits). This is a difficult test for the range reduction used in the sine evaluation (which brings the argument within the range $[-\pi/2, \pi/2]$, and which necessarily uses an approximate value of π), since 22 is close to an integer multiple of π .
- (Cody [249, 1982]) Evaluate $2.5^{125} = 5.5271\ 4787\ 5260\ 4445\ 6025 \times 10^{49}$ (shown correct to 21 digits). One way to evaluate $z = x^y$ is as $z = \exp(y \log x)$. But to obtain z correct to within a few ulps it is not sufficient to compute \exp and \log correct to within a few ulps; in other words, the composition of two functions evaluated to high accuracy is not necessarily obtained to the same accuracy. To examine this particular case, write

$$w := y \log x, \quad z = \exp(w).$$

If $w \rightarrow w + \Delta w$ then $z \rightarrow z + \Delta z$, where $z + \Delta z = \exp(w + \Delta w) =$

Table 2.4. *Sine test.*

| Machine | sin(22) |
|----------------|--|
| Exact | $-8.8513\ 0929\ 0403\ 8759 \times 10^{-3}$ |
| Casio fx-140 | $-8.8513\ \underline{\mathbf{62}} \times 10^{-3}$ |
| Casio fx-992VB | $-8.8513\ 0929\ \underline{\mathbf{096}} \times 10^{-3}$ |
| HP 48G | $-8.8513\ 0929\ 040 \times 10^{-3}$ |
| Sharp EL-5020 | $-8.8513\ \underline{\mathbf{0915\ 4}} \times 10^{-3}$ |
| MATLAB 6.1 | $-8.8513\ 0929\ 0403\ 876 \times 10^{-3}$ |
| WATFOR-77 | $-8.8513\ 0929\ 0403\ \underline{\mathbf{880}} \times 10^{-3}$ |
| FTN 90 | $-8.8513\ 0929\ 0403\ 876 \times 10^{-3}$ |
| QBasic | $-8.8513\ 0929\ 0403\ 876 \times 10^{-3}$ |

Table 2.5. *Exponentiation test. No entry for last column means same value as previous column.*

| Machine | 2.5^{125} | $\exp(125 \log(2.5))$ |
|----------------|---|---|
| Exact | $5.5271\ 4787\ 5260\ 4446 \times 10^{49}$ | $5.5271\ 4787\ 5260\ 4446 \times 10^{49}$ |
| Casio fx-140 | $5.5271\ \underline{\mathbf{477}} \times 10^{49}$ | $5.5271\ \underline{\mathbf{463}} \times 10^{49}$ |
| Casio fx-992VB | $5.5271\ 4787\ 526 \times 10^{49}$ | |
| HP 48G | $5.5271\ 4787\ 526 \times 10^{49}$ | $5.5271\ 4787\ \underline{\mathbf{377}} \times 10^{49}$ |
| Sharp EL-5020 | $5.5271\ 4787\ \underline{\mathbf{3}} \times 10^{49}$ | $5.5271\ \underline{\mathbf{4796\ 2}} \times 10^{49}$ |
| MATLAB 6.1 | $5.5271\ 4787\ 5260\ \underline{\mathbf{444}} \times 10^{49}$ | $5.5271\ 4787\ 5260\ \underline{\mathbf{459}} \times 10^{49}$ |
| WATFOR-77 | $5.5271\ 4787\ 5260\ \underline{\mathbf{450}} \times 10^{49}$ | $5.5271\ 4787\ 5260\ \underline{\mathbf{460}} \times 10^{49}$ |
| FTN 90 | $5.5271\ 4787\ 5260\ 445 \times 10^{49}$ | $5.5271\ 4787\ 5260\ \underline{\mathbf{459}} \times 10^{49}$ |
| QBasic | $5.5271\ 4787\ 5260\ \underline{\mathbf{444}} \times 10^{49}$ | |

$\exp(w)\exp(\Delta w) \approx \exp(w)(1 + \Delta w)$, so $\Delta z/z \approx \Delta w$. In other words, the relative error of z depends on the absolute error of w and hence on the size of w . To obtain z correct to within a few ulps it is necessary to use extra precision in calculating the logarithm and exponential [256, 1980, Chap. 7].

3. (Karpinski [715, 1985]) A simple test for the presence of a guard digit on a pocket calculator is to evaluate the expressions

$$9/27 * 3 - 1, \quad 9/27 * 3 - 0.5 - 0.5,$$

which are given in a form that can be typed directly into most four-function calculators. If the results are equal then a guard digit is present. Otherwise there is probably no guard digit (we cannot be completely sure from this simple test). To test for a guard digit on a computer it is best to run one of the diagnostic codes described in §27.5.

2.12. Notes and References

The classic reference on floating point arithmetic, and on all aspects of rounding error analysis, is Wilkinson's *Rounding Errors in Algebraic Processes* (REAP) [1232,

1963]. Wilkinson was uniquely qualified to write such a book, for not only was he the leading expert in rounding error analysis, but he was one of the architects and builders of the Automatic Computing Engine (ACE) at the National Physical Laboratory [1226, 1954]. The Pilot (prototype) ACE first operated in May 1950, and an engineered version was later sold commercially as the DEUCE Computer by the English Electric Company. Wilkinson and his colleagues were probably the first to write subroutines for floating point arithmetic, and this enabled them to accumulate practical experience of floating point arithmetic much earlier than anyone else [395, 1976], [1243, 1980].

In REAP, Wilkinson gives equal consideration to fixed point and floating point arithmetic. In fixed point arithmetic, all numbers are constrained to lie in a range such as $[-1, 1]$, as if the exponent were frozen in the floating point representation (2.1). Preliminary analysis and the introduction of scale factors during the computation are needed to keep numbers in the permitted range. We consider only floating point arithmetic in this book. REAP, together with Wilkinson's second book, *The Algebraic Eigenvalue Problem* (AEP) [1233, 1965], has been immensely influential in the areas of floating point arithmetic and rounding error analysis.

Wilkinson's books were preceded by the paper "Error Analysis of Floating-Point Computation" [1228, 1960], in which he presents the model (2.4) for floating point arithmetic and applies the model to several algorithms for the eigenvalue problem. This paper has hardly dated and is still well worth reading.

Another classic book devoted entirely to floating point arithmetic is Sterbenz's *Floating-Point Computation* [1062, 1974]. It contains a thorough treatment of low-level details of floating point arithmetic, with particular reference to IBM 360 and IBM 7090 machines. It also contains a good chapter on rounding error analysis and an interesting collection of exercises. R. W. Hamming has said of this book, "Nobody should ever have to know that much about floating-point arithmetic. But I'm afraid sometimes you might" [942, 1988]. Although Sterbenz's book is now dated in some respects, it remains a useful reference.

A third important reference on floating point arithmetic is Knuth's *Seminumerical Algorithms* [744, 1998, Sec. 4.2], from his *Art of Computer Programming* series. Knuth's lucid presentation includes historical comments and challenging exercises (with solutions).

The first analysis of floating point arithmetic was given by Samelson and Bauer [1009, 1953]. Later in the same decade Carr [205, 1959] gave a detailed discussion of error bounds for the basic arithmetic operations.

An up-to-date and very readable reference on floating point arithmetic is the survey paper by Goldberg [496, 1991], which includes a detailed discussion of IEEE arithmetic. A less mathematical, more hardware-oriented discussion can be found in the appendix "Computer Arithmetic" written by Goldberg that appears in the book on computer architecture by Hennessy and Patterson [562, 1996].

A fascinating historical perspective on the development of computer floating point arithmetics, including background to the development of the IEEE standard, can be found in the textbook by Patterson and Hennessy [929, 1998, §4.12].

The idea of representing floating point numbers in the form (2.1) is found, for example, in the work of Forsythe [428, 1969], Matula [830, 1970], and Dekker [302, 1971].

An alternative definition of $fl(x)$ is the nearest $y \in G$ satisfying $|y| \leq |x|$. This operation is called *chopping*, and does not satisfy our definition of rounding. Chopped arithmetic is used in the IBM/370 floating point system.

The difference between chopping and rounding (to nearest) is highlighted by a discrepancy in the index of the Vancouver Stock Exchange in the early 1980s [963, 1983]. The exchange established an index in January 1982, with the initial value of 1000. By November 1983 the index had been hitting lows in the 520s, despite the exchange apparently performing well. The index was recorded to three decimal places and it was discovered that the computer program calculating the index was chopping instead of rounding to produce the final value. Since the index was recalculated thousands of times a day, each time with a nonpositive final error, the bias introduced by chopping became significant. Upon recalculation with rounding the index almost doubled!

When there is a tie in rounding, two possible strategies are to round to the number with an even last digit and to round to the one with an odd last digit. Both are stable forms of rounding in the sense that

$$fl(((x + y) - y) + y) = fl((x + y) - y),$$

as shown by Reiser and Knuth [982, 1975], [744, 1998, Sec. 4.2.2, Thm. D]. For other rules, such as round away from zero, repeated subtraction and addition of the same number can yield an increasing sequence, a phenomenon known as *drift*. For bases 2 and 10 rounding to even is preferred to rounding to odd. After rounding to even a subsequent rounding to one less place does not involve a tie. Thus we have the rounding sequence 2.445, 2.44, 2.4 with round to even, but 2.445, 2.45, 2.5 with round to odd. For base 2, round to even causes computations to produce integers more often [706, 1979] as a consequence of producing a zero least significant bit. Rounding to even in the case of ties seems to have first been suggested by Scarborough in the first edition (1930) of [1014, 1950]. Hunt [651, 1997] notes that in the presentation of meteorological temperature measurements round to odd is used as a tie-breaking strategy, and he comments that this may be to avoid the temperatures 0.5°C and 32.5°F being rounding down and the incorrect impression being given that it is freezing.

Predictions based on the growth in the size of mathematical models solved as the memory and speed of computers increase suggest that floating point arithmetic with unit roundoff $u \approx 10^{-32}$ will be needed for some applications on future supercomputers [57, 1989].

The model (2.4) does not fully describe any floating point arithmetic. It is merely a tool for error analysis—one that has been remarkably successful in view of our current understanding of the numerical behaviour of algorithms. There have been various attempts to devise formal models of floating point arithmetic, by specifying sets of axioms in terms of which error analysis can be performed. Some attempts are discussed in §27.7.4. No model yet proposed has been truly successful. Priest [955, 1992] conjectures that the task of “encapsulating all that we wish to know about floating point arithmetic in a single set of axioms” is impossible, and he gives some motivation for this conjecture.

Under the model (2.4), floating point arithmetic is not associative with respect to any of the four basic operations: $(a \overline{\text{op}} b) \overline{\text{op}} c \neq a \overline{\text{op}} (b \overline{\text{op}} c)$, $\text{op} = +, -, *, /$,

where $a \boxed{\text{op}} b := fl(a \text{ op } b)$. Nevertheless, floating point arithmetic enjoys some algebraic structure, and it is possible to carry out error analysis in the “ $\boxed{\text{op}}$ algebra”. Fortunately, it was recognized by Wilkinson and others in the 1950s that this laboured approach is unnecessarily complicated, and that it is much better to work with the exact equations satisfied by the computed quantities. As Parlett [925, 1990] notes, though, “There have appeared a number of ponderous tomes that do manage to abstract the computer’s numbers into a formal structure and burden us with more jargon.”

A draft proposal of IEEE Standard 754 is defined and described in [657, 1981]. That article, together with others in the same issue of the journal *Computer*, provides a very readable description of IEEE arithmetic. In particular, an excellent discussion of gradual underflow is given by Coonen [270, 1981]. A draft proposal of IEEE Standard 854 is presented, with discussion, in [253, 1984].

W. M. Kahan of the University of California at Berkeley received the 1989 ACM Turing Award for his contributions to computer architecture and numerical analysis, and in particular for his work on IEEE floating point arithmetic standards 754 and 854. For interesting comments by Kahan on the development of IEEE arithmetic and other floating point matters, see [1029, 1998], [1251, 1997],

An interesting examination of the implications of the IEEE standard for high-level languages such as Fortran is given by Fateman [405, 1982]. Topics discussed include trap handling and how to exploit NaNs. For an overview of hardware implementations of IEEE arithmetic, and software support for it, see Cody [251, 1988].

Producing a fast and correct implementation of IEEE arithmetic is a difficult task. Correctness is especially important for a microprocessor (as opposed to a software) implementation, because of the logistical difficulties of correcting errors when they are found. In late 1994, much publicity was generated by the discovery of a flaw in the floating point divide instruction of Intel’s Pentium chip. Because of some missing entries in a lookup table on the chip, the FPDIV instruction could give as few as four correct significant decimal digits for double precision floating point arguments with certain special bit patterns [1036, 1994], [380, 1997]. The flaw had been discovered by Intel in the summer of 1994 during ongoing testing of the Pentium processor, but it had not been publically announced. In October 1994, a mathematician doing research into prime numbers independently discovered the flaw and reported it to the user community. Largely because of the way in which Intel responded to the discovery of the flaw, the story was reported in national newspapers (e.g., the *New York Times* [816, 1994]) and generated voluminous discussion on Internet newsgroups (notably `comp.sys.intel`). Intel corrected the bug in 1994 and, several weeks after the bug was first reported, offered to replace faulty chips. For a very readable account of the Pentium FPDIV bug story, see Moler [866, 1995]. To emphasize that bugs in implementations of floating point arithmetic are far from rare, we mention that the Calculator application in Microsoft Windows 3.1 evaluates $fl(2.01 - 2.00) = 0.0$.

Computer chip designs can be tested in two main ways: by software simulations and by applying formal verification techniques. Formal verification aims to prove mathematically that the chip design is correct, and this approach has been in use by chip manufacturers for some time [491, 1995]. Some relevant references are

Barrett [79, 1989] for the Inmos T800 (or Shepherd and Wilson [1037, 1989] for a more informal overview), and Moore, Lynch, and Kaufmann [871, 1998] and Russinoff [1003, 1998] for AMD processors.

For low-level details of computer arithmetic several textbooks are available. We mention only the recent and very thorough book by Parhami [921, 2000].

The floating point operation op ($\text{op} = +, -, *, /$) is *monotonic* if $fl(a \text{ op } b) \leq fl(c \text{ op } d)$ whenever $a, b, c,$ and d are floating point numbers for which $a \text{ op } b \leq c \text{ op } d$ and neither $fl(a \text{ op } b)$ nor $fl(c \text{ op } d)$ overflows. IEEE arithmetic is monotonic, as is any correctly rounded arithmetic. Monotonic arithmetic is important in the bisection algorithm for finding the eigenvalues of a symmetric tridiagonal matrix; see Demmel, Dhillon, and Ren [320, 1995], who give rigorous correctness proofs of some bisection implementations in floating point arithmetic. Ferguson and Brightman [410, 1991] derive conditions that ensure that an approximation to a monotonic function preserves the monotonicity on a set of floating point numbers.

On computers of the 1950s, (fixed point) multiplication was slower than (fixed point) addition by up to an order of magnitude [773, 1980, Apps. 2, 3]. For most modern computers it is a rule of thumb that a floating point addition and multiplication take about the same amount of time, while a floating point division is 2–10 times slower, and a square root operation (in hardware) is 1–2 times slower than a division.

During the design of the IBM 7030, Sweeney [1116, 1965] collected statistics on the floating point additions carried out by selected application programs on an IBM 704. He found that 11% of all instructions traced were floating point additions. Details were recorded of the shifting needed to align floating point numbers prior to addition, and the results were used in the design of the shifter on the IBM 7030.

The word *bit*, meaning binary digit, first appeared in print in a 1948 paper of Claude E. Shannon, but the term was apparently coined by John W. Tukey [1160, 1984]. The word *byte*, meaning a group of (usually eight) bits, did not appear in print until 1959 [173, 1981].

The earliest reference we know for Theorem 2.5 is Sterbenz [1062, 1974, Thm. 4.3.1]. Theorem 2.4 is due to Ferguson [409, 1995], who proves a more general version of the theorem that allows for trailing zero digits in x and y . A variation in which the condition is $0 \leq y \leq x \leq y + \beta^e$, where $e = \min\{j : \beta^j \geq y\}$, is stated by Ziv [1285, 1991] and can be proved in a similar way.

For more on the choice of base, see Cody [255, 1973] and Kuki and Cody [754, 1973]. Buchholz's paper *Fingers or Fists?* [172, 1959] on binary versus decimal representation of numbers on a computer deserves mention for its clever title, though the content is only of historical interest.

The model (2.4) ignores the possibility of underflow and overflow. To take underflow into account the model must be modified to

$$fl(x \text{ op } y) = (x \text{ op } y)(1 + \delta) + \eta, \quad \text{op} = +, -, *, /. \quad (2.8)$$

As before, $|\delta| \leq u$. If underflow is gradual, as in IEEE arithmetic, then $|\eta| \leq \frac{1}{2}\beta^{e_{\min}-t} = \lambda u$, which is half the spacing between the subnormal numbers ($\lambda = \beta^{e_{\min}-1}$ is the smallest positive normalized floating point number); if underflows

are flushed to zero then $|\eta| \leq \lambda$. Only one of δ and η is nonzero: δ if no underflow occurs, otherwise η . With gradual underflow the *absolute error* of an underflowed result is no greater than the smallest (bound for the) absolute error that arises from an operation $fl(x \text{ op } y)$ in which the arguments and result are normalized. Moreover, with gradual underflow we can take $\eta = 0$ for $\text{op} = +, -$ (see Problem 2.19). For more details, and a thorough discussion of how error analysis of standard algorithms is affected by using the model (2.8), see the perceptive paper by Demmel [308, 1984]. Another relevant reference is Neumaier [884, 1985].

Hauser [553, 1996] gives a thorough and very readable treatment of exception handling, covering underflow, overflow, indeterminate or undefined cases, and their support in software.

An important problem not considered in this chapter is the conversion of numbers between decimal and binary representations. These conversions are needed whenever numbers are read into a computer or printed out. They tend to be taken for granted, but if not done carefully they can lead to puzzling behaviour, such as a number read in as 0.1 being printed out as 0.099...9. To be precise, the problems of interest are (a) convert a number represented in decimal notation to the best binary floating point representation of a given precision, and (b) given a binary floating point number, print a correctly rounded decimal representation, either to a given number of significant digits or to the smallest number of significant digits that allows the number to be re-read without loss of accuracy. Algorithms for solving these problems are given by Clinger [247, 1990] and Steele and White [1060, 1990]; Gay [470, 1990] gives some improvements to the algorithms and C code implementing them. Precise requirements for binary–decimal conversion are specified in the IEEE arithmetic standard. A program for testing the correctness of binary–decimal conversion routines is described by Paxson [930, 1991]. Early references on base conversion are Goldberg [497, 1967] and Matula [829, 1968], [830, 1970]. It is interesting to note that, in Fortran or C, where the output format for a “print” statement can be precisely specified, most compilers will, for an (in)appropriate choice of format, print a decimal string that contains many more significant digits than are determined by the floating point number whose value is being represented.

Other authors who have analysed various aspects of floating (and fixed) point arithmetic include Diamond [339, 1978], Urabe [1174, 1968], and Feldstein, Goodman, and co-authors [510, 1975], [407, 1982], [511, 1985], [408, 1986]. For a survey of computer arithmetic up until 1976 that includes a number of references not given here, see Garner [460, 1976].

Problems

*The exercise had warmed my blood,
and I was beginning to enjoy myself amazingly.*

— JOHN BUCHAN, *The Thirty-Nine Steps* (1915)

2.1. How many normalized numbers and how many subnormal numbers are there in the system F defined in (2.1) with $e_{\min} \leq e \leq e_{\max}$? What are the figures for IEEE single and double precision (base 2)?

2.2. Prove Lemma 2.1.

2.3. In IEEE arithmetic how many double precision numbers are there between any two adjacent nonzero single precision numbers?

2.4. Prove Theorem 2.3.

2.5. Show that

$$0.1 = \sum_{i=1}^{\infty} (2^{-4i} + 2^{-4i-1})$$

and deduce that 0.1 has the base 2 representation $0.000\overline{1100}$ (repeating last 4 bits). Let $\hat{x} = fl(0.1)$ be the rounded version of 0.1 obtained in binary IEEE single precision arithmetic ($u = 2^{-24}$). Show that $(x - \hat{x})/x = -\frac{1}{4}u$.

2.6. What is the largest integer p such that all integers in the interval $[-p, p]$ are exactly representable in IEEE double precision arithmetic? What is the corresponding p for IEEE single precision arithmetic?

2.7. Which of the following statements is true in IEEE arithmetic, assuming that a and b are normalized floating point numbers and that no exception occurs in the stated operations?

1. $fl(a \text{ op } b) = fl(b \text{ op } a)$, $\text{op} = +, *$.
2. $fl(b - a) = -fl(a - b)$.
3. $fl(a + a) = fl(2 * a)$.
4. $fl(0.5 * a) = fl(a/2)$.
5. $fl((a + b) + c) = fl(a + (b + c))$.
6. $a \leq fl((a + b)/2) \leq b$, given that $a \leq b$.

2.8. Show that the inequalities $a \leq fl((a + b)/2) \leq b$, where a and b are floating point numbers with $a \leq b$, can be violated in base 10 arithmetic. Show that $a \leq fl(a + (b - a)/2) \leq b$ in base β arithmetic, for any β .

2.9. What is the result of the computation $\sqrt{1 - 2^{-53}}$ in IEEE double precision arithmetic, with and without double rounding from an extended format with a 64-bit significand?

2.10. A theorem of Kahan [496, 1991, Thm. 7] says that if $\beta = 2$ and the arithmetic rounds as specified in the IEEE standard, then for integers m and n with $|m| < 2^{t-1}$ and $n = 2^i + 2^j$ (some i, j), $fl((m/n) * n) = m$. Thus, for example, $fl((1/3) * 3) = 1$ (even though $fl(1/3) \neq 1/3$). The sequence of allowable n begins 1, 2, 3, 4, 5, 6, 8, 9, 10, 12, 16, 17, 18, 20, so Kahan's theorem covers many common cases. Test the theorem on your computer.

2.11. Investigate the leading significant digit distribution for numbers obtained as follows.

1. k^n , $n = 0:1000$ for $k = 2$ and 3.
2. $n!$, $n = 1:1000$.

3. The eigenvalues of a random symmetric matrix.
4. Physical constants from published tables.
5. From the front page of the *London Times* or the *New York Times*.

(Note that in writing a program for the first case you can form the powers of 2 or 3 in order, following each multiplication by a division by 10, as necessary, to keep the result in the range $[1, 10]$. Similarly for the second case.)

2.12. (Edelman [379, 1994]) Let x be a floating point number in IEEE double precision arithmetic satisfying $1 \leq x < 2$. Show that $fl(x * (1/x))$ is either 1 or $1 - \epsilon/2$, where $\epsilon = 2^{-52}$ (the machine epsilon).

2.13. (Edelman [379, 1994]) Consider IEEE double precision arithmetic. Find the smallest positive integer j such that $fl(x * (1/x)) \neq 1$, where $x = 1 + j\epsilon$, with $\epsilon = 2^{-52}$ (the machine epsilon).

2.14. Kahan has stated that “an (over-)estimate of u can be obtained for almost any machine by computing $|3 \times (4/3 - 1) - 1|$ using rounded floating-point for every operation”. Test this estimate against u on any machines available to you.

2.15. What is 0^0 in IEEE arithmetic?

2.16. Evaluate these expressions in any IEEE arithmetic environment available to you. Are the values returned what you would expect? (None of the results is specified by the IEEE standard.)

1. 1^∞ .
2. 2^∞ .
3. $\exp(\infty)$, $\exp(-\infty)$.
4. $\text{sign}(\text{NaN})$, $\text{sign}(-\text{NaN})$.
5. NaN^0 .
6. ∞^0 .
7. 1^{NaN} .
8. $\log(\infty)$, $\log(-\infty)$, $\log(0)$.

2.17. Let x_{\max} denote the largest representable number in IEEE single or double precision arithmetic. In what circumstances does $2x_{\max}$ not overflow?

2.18. Can Theorem 2.4 be strengthened to say that $fl(x - y)$ is computed exactly whenever the exponents of $x \geq 0$ and $y \geq 0$ differ by at most 1?

2.19. (Hauser [553, 1996]) Show that with gradual underflow, if x and y are floating point numbers and $fl(x \pm y)$ underflows then $fl(x \pm y) = x \pm y$.

2.20. Two requirements that we might ask of a routine for computing \sqrt{x} in floating point arithmetic are that the identities $\sqrt{x^2} = |x|$ and $(\sqrt{x})^2 = |x|$ be satisfied. Which, if either, of these is a reasonable requirement?

2.21. Are there any floating point values of x and y (excepting values both 0, or so huge or tiny as to cause overflow or underflow) for which the computed value of $x/\sqrt{x^2 + y^2}$ exceeds 1?

2.22. (Kahan) A natural way to compute the maximum of two numbers x and y is with the code

```
% max(x, y)
if x > y then
    max = x
else
    max = y
end
```

Does this code always produce the expected answer in IEEE arithmetic?

2.23. Prove that Kahan's formula (2.7) computes the area of a triangle accurately. (Hint: you will need one invocation of Theorem 2.5.)

2.24. (Kahan) Describe the result of the computation $y = (x + x) - x$ on a binary machine with a guard digit and one without a guard digit.

2.25. (Kahan) Let $f(x) = (((x - 0.5) + x) - 0.5) + x$. Show that if f is evaluated as shown in single or double precision binary IEEE arithmetic then $f(x) \neq 0$ for all floating point numbers x .

2.26. Show that if x and y are floating point numbers and $\hat{a} = fl(xy)$ and $\hat{b} = fl(xy - \hat{a})$ are computed using fused multiply-add operations, then $\hat{a} + \hat{b} \equiv xy$.

2.27. (Kahan) Show that with the use of a fused multiply-add operation the algorithm

```
w = bc
e = w - bc
x = (ad - w) + e
```

computes $x = \det\begin{bmatrix} a & b \\ c & d \end{bmatrix}$ with high relative accuracy.

2.28. Suppose we have an iterative algorithm for computing $z = x/y$. Derive a convergence test that terminates the iteration (only) when full accuracy has been achieved. Assume the use of IEEE arithmetic with gradual underflow (use (2.8)).