

How to Read Floating Point Numbers Accurately

William D. Clinger

CIS-TR-90-01

June 5, 1990

Abstract

Consider the problem of converting decimal scientific notation for a number into the best binary floating point approximation to that number, for some fixed precision. This problem cannot be solved using arithmetic of any fixed precision. Hence the *IEEE Standard for Binary Floating-Point Arithmetic* does not require the result of such a conversion to be the best approximation.

This paper presents an efficient algorithm that always finds the best approximation. The algorithm uses a few extra bits of precision to compute an IEEE-conforming approximation while testing an intermediate result to determine whether the approximation could be other than the best. If the approximation might not be the best, then the best approximation is determined by a few simple operations on multiple-precision integers, where the precision is determined by the input. When using 64 bits of precision to compute IEEE double precision results, the algorithm avoids higher-precision arithmetic over 99% of the time.

The input problem considered by this paper is the inverse of an output problem considered by Steele and White: Given a binary floating point number, print a correctly rounded decimal representation of it using the smallest number of digits that will allow the number to be read without loss of accuracy. The Steele and White algorithm assumes that the input problem is solved; an imperfect solution to the input problem, as allowed by the IEEE standard and ubiquitous in current practice, defeats the purpose of their algorithm.

NOTE

This is a corrected version of a paper to appear in the proceedings of the ACM SIGPLAN '90 Conference on Programming Language Design and Implementation. The three corrections are listed on the last page of this report.

given an oracle that delivers the value of g rounded to the nearest 10. Given an input x , the obvious approach is to ask the oracle for the value of $g(x)$ to the nearest 10, and then to round that result to the nearest 10000. Unfortunately, this does not always work. If $g(x)$ is 11074996 and $g(y)$ is 11075004, for example, then the answers should be 11070000 and 11080000, respectively, but the oracle will deliver 11075000 for both x and y . This approach usually works, though. It fails only when the value delivered by the oracle ends in 5000.

To find the best approximation, the algorithm of this paper uses an extended precision to compute an approximation whose accuracy is significantly better than the accuracy required in the result. It then examines the low order bits of that intermediate result to bound the additional error that will be introduced when this particular value is rounded to the precision required of the final result. If the sum of the bounds for the intermediate approximation error and for the rounding error is less than the error allowed in the final result, then the rounded value must be correct. Otherwise the rounded value must be checked using arithmetic of greater precision.

By choosing a large enough extended precision, the probability that an even greater precision will be needed can be made as small as desired. For IEEE double precision numbers, an experiment reported in Section 9 found that IEEE extended precision is enough to find the best approximation over 99% of the time.

2. EXTERNAL AND INTERNAL RADIXES

For concreteness it is often appropriate to assume that the input is given in decimal scientific notation and that the output is a binary floating point number, but most results presented here hold for more general radices. Let $\Delta \geq 2$ be the radix used to express the input, and let $\beta \geq 2$ be the radix of the floating point output. Δ is the *external* radix, and β is the *internal* radix. In practice Δ is usually 10 and β is usually 2 or 16.

Definition 1. [Matula70] Radices Δ and β are *commensurable* if and only if both are integral powers of a common integral root. Equivalently, Δ and β are commensurable iff $\log_{\beta} \Delta$ is rational.

For example, 2 and 16 are commensurable radices, as are 8 and 16. On the other hand 10 and 2 are incommensurable, as are 10 and 16. As shown in Section 4, the problem of computing the best floating point approximation is trivial if Δ and β are commensurable radices. Most of this paper therefore assumes that Δ and β are incommensurable, and for simplicity the internal radix β is always assumed to be even.

For convenience "bit" will be used to refer to digits in the internal radix β , and "digit" will be used for digits in the external radix Δ .

3. FLOATING POINT NUMBERS

This paper deals exclusively with positive floating point numbers. Overflow and underflow are considered only in Section 8. Because the following definition refers to precision and not to a data structure, this paper considers an IEEE single precision number to be a 24-bit floating point number and an IEEE double to be a 53-bit floating point number.

Let n , the precision in bits of a floating point number, be a positive integer. For the purposes of this paper, an n -bit floating point number consists of an integer significand m and an integer exponent q with $0 < m < \beta^n$, representing the value $m \times \beta^q$. A floating point number is normalized iff $\beta^{n-1} \leq m < \beta^n$. For a fixed precision n , the significand m and exponent q of a normalized floating point number are uniquely determined by its value $m \times \beta^q$.

For any fixed precision n , a closest floating point approximation to a real number a is a normalized floating point number $m \times \beta^q$ such that $a = (m + \epsilon) \times \beta^q$ where $|\epsilon| \leq 1/2$, and where $m = \beta^{n-1}$ only if $-1/(2\beta) \leq \epsilon$. The closest approximation is uniquely determined unless $|\epsilon| = 1/2$, or $m = \beta^{n-1}$ and $\epsilon = -1/(2\beta)$, in which cases there are two closest approximations. The best approximation is a closest approximation where, in a case of two closest approximations, the tie is broken by a fixed rounding rule such as round to even.

For any real number a , the fractional part of a is defined to be $\{a\} = a - \lfloor a \rfloor$. The value of a rounded to the nearest integer, assuming ties round to even, is defined to be

$$\lfloor a \rfloor = \begin{cases} \lfloor a \rfloor & \text{if } \{a\} < 1/2 \\ \lceil a \rceil & \text{if } \{a\} > 1/2 \\ \lfloor a \rfloor & \text{if } \{a\} = 1/2 \text{ and } \lfloor a \rfloor \text{ is even} \\ \lceil a \rceil & \text{if } \{a\} = 1/2 \text{ and } \lfloor a \rfloor \text{ is odd} \end{cases}$$

This definition can easily be changed to accommodate other tie-breaking rules, but the algorithms presented later will assume that ties round to even. The significand of the best n -bit approximation to a is

$$\lfloor \beta^{n-1} \beta^{\lfloor \log_{\beta} a \rfloor} \rfloor$$

unless this value is β^n , in which case the significand is β^{n-1} .

The floating point product of normalized floating point numbers $x \times \beta^q$ and $y \times \beta^r$ is a best approximation to $xy \times \beta^{q+r}$. That is, the floating point multiplication operator is assumed to be reliably accurate, to round to nearest, and to resolve ties that result when the mathematical product is exactly halfway between two adjacent floating point numbers by the same rounding rule used to define the best approximation. These assumptions hold for the default rounding mode in IEEE arithmetic, but do not hold for many other implementations of floating point numbers.

Theorem 5. For $n \geq 5$, no finite automaton computes the significand of the best n -bit binary floating point approximation to $f \times 10^e$, where f and e are presented in base 10.

Proof: This is a special case of the lemmas below. ■

Lemma 6. If x and y are positive real numbers with

$$\frac{1}{2\beta^{n-1} \log \beta} < \{\log_{\beta} x - \log_{\beta} y\} < 1 - \frac{1}{2\beta^{n-1} \log \beta}$$

then the best n -bit floating point approximations to x and y have distinct significands.

Proof: By symmetry suppose $\{\log_{\beta} x\} < \{\log_{\beta} y\}$. Let $w = \{\log_{\beta} x\}$ and $w + \delta = \{\log_{\beta} y\}$. Then

$$\begin{aligned} \beta^{w+\delta} - \beta^w &= \beta^w (\beta^{\delta} - 1) \\ &\geq \beta^{\delta} - 1 \\ &> \delta \log \beta \\ &> \frac{1}{2\beta^{n-1}} \end{aligned}$$

so

$$|[\beta^{n-1} \beta^{w+\delta}] - [\beta^{n-1} \beta^w]| > \frac{1}{2}. \quad \blacksquare$$

Lemma 7. (Kronecker's Theorem in one dimension) If θ is irrational, then $\{\{n\theta\} \mid n \in \omega\}$ is dense in the interval $(0, 1)$.

Proof: See [HW60]. ■

Lemma 8. If Δ and β are incommensurable, $\beta \geq 3$ or $n \geq 2$, and e is presented with its least significant digit first, then no finite automaton computes the significand of the best n -bit approximation to Δ^e .

Proof: Suppose e is presented in radix γ , and let D be a DFA. Let i and j be integers such that D is in the same state after reading $i+j$ zeroes as after reading i zeroes. By Kronecker's theorem there exists an integer k such that

$$\frac{1}{2\beta^{n-1} \log \beta} < \{k(\gamma^{i+j} - \gamma^i) \log_{\beta} \Delta\} < 1 - \frac{1}{2\beta^{n-1} \log \beta}.$$

Take $e_1 = k\gamma^{i+j}$ and $e_2 = k\gamma^i$. D computes the same result for Δ^{e_1} as for Δ^{e_2} , but their significands are distinct by Lemma 6. ■

Lemma 9. If θ is irrational and $\gamma > 1$ is an integer, then there exist infinitely many nonnegative integers k such that

$$\frac{\gamma - 1}{\gamma^2} < \{\gamma^k \theta\} < \frac{\gamma^2 - \gamma + 1}{\gamma^2}.$$

Proof: If all digits to the right of the radix point in the γ -ary representation of θ are zero or $\gamma - 1$, then let k be such that the k th digit to the right of the radix point is $\gamma - 1$ and the following digit is zero. Otherwise let k be such that the k th digit is neither zero nor $\gamma - 1$.

Since $\gamma^k \theta$ is also irrational, a larger such k always exists. ■

Lemma 10. If Δ and β are incommensurable, e is presented in base γ with its most significant digit first, and $\gamma^2/(\gamma - 1) < 2\beta^{n-1} \log \beta$, then no finite automaton computes the significand of the best n -bit approximation to Δ^e .

Proof: Let D be a DFA, and let i and j be integers such that D is in the same state after reading γ^{i+j} as after reading γ^i . By Lemma 9 there exists an integer k such that

$$\frac{\gamma - 1}{\gamma^2} < \{\gamma^k (\gamma^{i+j} - \gamma^i) \log_{\beta} \Delta\} < \frac{\gamma^2 - \gamma + 1}{\gamma^2}$$

so take $e_1 = \gamma^{i+j+k}$ and $e_2 = \gamma^{i+k}$. ■

The lemma above is unpleasantly technical, for there is no apparent reason why the base in which the inputs are presented should affect the difficulty of the problem. The lemma probably holds without such assumptions, but a more sophisticated proof will be required.

Figure 1 shows a straightforward algorithm, **AlgorithmM**, that uses integer arithmetic of unlimited precision to compute the best n -bit floating point approximation to $f \times \Delta^e$. As written, the algorithm assumes ties are broken by rounding to even.

Like the other algorithms in this paper, **AlgorithmM** is expressed as a purely functional Scheme program, with the assumption that all integer arithmetic is exact, *i.e.* of unlimited precision [Rees86]. For integers m and k with $\beta^{n-1} \leq m < \beta^n$, (**make-float** m k) is assumed to return the n -bit floating point number $m \times \beta^k$. The **nextfloat** procedure, shown in Figure 2, returns the least normalized floating point number greater than its argument.

For most applications **AlgorithmM** is impractical because it uses too much high-precision arithmetic. The next section obtains a better algorithm by starting from a close but not necessarily closest approximation.

5. AN ITERATIVE ALGORITHM

It is quite easy, using a few extra bits of precision, to find an approximation that differs from the best approximation by only a few units in the last place of the significand. In fact, it is fairly easy to find an n -bit approximation that differs from the best approximation by no more than one unit.

AlgorithmR, in Figure 3, takes a good approximation $m \times \beta^k$ and checks it using integer arithmetic of unlimited precision. If the given approximation is too small or too large, it then repeats the process with the next larger or smaller floating point number.

The algorithm begins by finding positive integers x and y such that

$$\frac{x}{y} = \frac{f \times \Delta^e}{m \times \beta^k}.$$

```
; Given exact integers f and e, with f positive,
; and a floating point number z0 close to f * delta^e,
; returns the best floating point approximation to
; f * delta^e.
```

```
(define (AlgorithmR f e z0)

  (define (loop z)

    (define m (float-significand z))
    (define k (float-exponent z))

    ; Given exact positive integers x and y with
    ; x/y = (f*delta^e)/(m*beta^k), returns the best
    ; approximation to f*delta^e.

    (define (compare x y)
      (let* ((D (- x y))
             (D2 (* 2 m (abs D))))
        (cond ((< D2 y)
               (if (and (= m beta^n-1)
                        (negative? D)
                        (> (* beta D2) y))
                   (loop (prevfloat z))
                   z))
              ((= D2 y)
               (cond ((even? m)
                      (if (and (= m beta^n-1)
                              (negative? D))
                          (loop (prevfloat z))
                          z))
                    ((negative? D)
                     (prevfloat z))
                    ((positive? D)
                     (nextfloat z))))
              ((negative? D)
               (loop (prevfloat z)))
              ((positive? D)
               (loop (nextfloat z))))))

    (cond ((and (>= e 0) (>= k 0))
           (compare (* f (expt delta e))
                    (* m (expt beta k))))
          ((and (>= e 0) (< k 0))
           (compare (* f (expt delta e)
                      (expt beta (- k)))
                    m))
          ((and (< e 0) (>= k 0))
           (compare f (* m (expt beta k)
                        (expt delta (- e)))))
          ((and (< e 0) (< k 0))
           (compare (* f (expt beta (- k))
                        (* m (expt delta (- e)))))
                    (* m (expt delta (- e))))))

    (loop z0))

  (define beta^n+1 (expt beta (+ n 1)))
```

Figure 3. AlgorithmR.

seems to be a pre-computed table of powers of Δ , containing the range of powers that is apt to occur in practice. With practical floating point formats the range of floating point exponents is usually limited, so very large exponents will overflow and very small exponents will underflow unless the number of digits in the input is unreasonably large. AlgorithmM can be used when the input exponent is out of the table's range.

Even when limited to the range needed for reasonable inputs, the table of powers may be fairly large. The size of the table can be reduced, at the expense of accuracy, by factoring it into two smaller tables. One table contains values for small powers of Δ^e , with $0 \leq e < h$, while another table contains approximations to 10^{hj} for integral j . It is convenient to assume that h is small enough that the small powers are represented exactly as p -bit floating point numbers. If this is so, and all other table entries are best approximations, and $\beta = 2$, then Corollary 3 says that the error in the value calculated for Δ^e is strictly less than $\frac{3}{2}$ units in the least significant bit.

If the floating point approximation to f is the best possible, and the approximation to Δ^e is within $\frac{3}{2}$ units, and $\beta = 2$, then the error in the product is less than $\frac{9}{2}$ units. If this calculation is performed using $p \geq n + 4$ bits of precision, then rounding the product to the nearest n bits yields either the best n -bit binary floating point approximation to $f \times \Delta^e$ or a next-best approximation.

7. AN EFFICIENT, NON-ITERATIVE ALGORITHM

Algorithm Bellerophon, shown in Figure 4 for the special case of $\Delta = 10$ and $\beta = 2$, is a practical algorithm based on the idea explained in the introduction. The terms used in Figure 4 differ from those used to describe previous algorithms, in that "float" refers to p -bit floating point numbers and "shortfloat" refers to n -bit numbers.

Given integers f and e , Bellerophon dispatches on the error introduced when f and Δ^e are approximated by floating point numbers with p bits of precision, where p is large enough to ensure that the product of the approximations, rounded to n bits, is either the best or a next-best approximation to $f \times \Delta^e$.

If f and Δ^e can both be represented exactly using n bits, then an n -bit floating point multiplication yields the best approximation to their product. If f and Δ^{-e} can both be represented exactly using n bits, then an n -bit division yields the best approximation. (This assumes that floating point division, like multiplication, is reliably accurate.)

Otherwise Bellerophon approximates f and Δ^e by p -bit floating point numbers x and y , and computes their floating point product $z = m \times \beta^q$, where $\beta^{p-1} \leq m < \beta^p$. Hence

$$f \times \Delta^e = (z + \epsilon) \times \beta^q$$

	$ \epsilon_x $	$ \epsilon_y $	$ \epsilon $
$f < 2^p \wedge 0 \leq e < h$	0	0	$\leq \frac{1}{2}$
$f < 2^p \wedge (e < 0 \vee e \geq h)$	0	$< \frac{3}{2}$	$< \frac{7}{2}$
$f \geq 2^p \wedge 0 \leq e < h$	$\leq \frac{1}{2}$	0	$< \frac{3}{2}$
$f \geq 2^p \wedge (e < 0 \vee e \geq h)$	$\leq \frac{1}{2}$	$< \frac{3}{2}$	$< \frac{9}{2}$

Figure 5. Error bounds in units of the least significant bit.

where for $\beta = 2$ the error ϵ is bounded by the values shown in Figure 5.

Unless z lies about halfway between two adjacent n -bit floating point numbers, the error ϵ will be absorbed when z is rounded to n bits. The multiply-and-test procedure therefore tests to see if z could be within ϵ of the midpoint. If not, then the correct answer is obtained by rounding z to n bits. Otherwise the efficient part of the algorithm fails, and the rounded value of z is passed to AlgorithmR as a starting approximation. Since the rounded value of z is always either the best or a next-best approximation, AlgorithmR always converges in one loop.

Theorem 11. Algorithm Bellerophon computes the best n -bit approximation to $f \times \Delta^e$.

Proof: This proof deals with the generalizations of Figures 4 and 5 to any even internal radix β . In general, Algorithm Bellerophon computes a p -bit floating point number $z = m \times \beta^k$ such that $f \times \Delta^e = (z + \epsilon) \times \beta^k$ and

$$|\epsilon| \leq \text{slop} + \frac{1}{2}$$

where the value of slop is determined by numerical analysis.

Let z_1 and z_0 be integers such that

$$z = z_1 \times \beta^{p-n} + z_0$$

$$\beta^{n-1} \leq z_1 < \beta^n$$

$$0 \leq z_0 < \beta^{p-n}$$

There are three cases, depending on whether z_0 is well below, well above, or near $\frac{1}{2}\beta^{p-n}$.

Case 1: $z_0 + \text{slop} < \frac{1}{2}\beta^{p-n}$. Rounding z to n bits yields $z_1 \times \beta^{k+p-n}$, and $f \times \Delta^e = (z_1 + \epsilon') \times \beta^{k+p-n}$ where the error ϵ' is

$$\begin{aligned} |\epsilon'| &= \left| \frac{f \times \Delta^e}{\beta^{k+p-n}} - z_1 \right| \\ &= \left| \frac{(z + \epsilon) \times \beta^k}{\beta^{k+p-n}} - z_1 \right| \\ &= \left| \frac{z_0 + \epsilon}{\beta^{p-n}} \right| \end{aligned}$$

$$\begin{aligned} &\leq \beta^{n-p}(z_0 + |\epsilon|) \\ &\leq \beta^{n-p}(z_0 + \text{slop} + \frac{1}{2}) \\ &< \beta^{n-p} \times \frac{1}{2}\beta^{p-n} \\ &= \frac{1}{2} \end{aligned}$$

Case 2: $\frac{1}{2}\beta^{p-n} < z_0 - \text{slop}$. Rounding z to n bits yields $(z_1 + 1) \times \beta^{k+p-n}$, and $f \times \Delta^e = (z_1 + 1 + \epsilon') \times \beta^{k+p-n}$ where

$$\begin{aligned} |\epsilon'| &= \left| \frac{(z + \epsilon) \times \beta^k}{\beta^{k+p-n}} - z_1 - 1 \right| \\ &= \left| \frac{z_0 + \epsilon}{\beta^{p-n}} - 1 \right| \\ &\leq \beta^{n-p}(\beta^{p-n} - z_0 + |\epsilon|) \\ &\leq \beta^{n-p}(\beta^{p-n} - z_0 + \text{slop} + \frac{1}{2}) \\ &< \beta^{n-p} \times (\beta^{p-n} - \frac{1}{2}\beta^{p-n}) \\ &= \frac{1}{2} \end{aligned}$$

Case 3: $\frac{1}{2}\beta^{p-n} - \text{slop} \leq z_0 \leq \frac{1}{2}\beta^{p-n} + \text{slop}$. This is the failure case in which another algorithm is used. ■

8. OVERFLOW AND UNDERFLOW

Overflow and underflow become possible when the range of floating point exponents is restricted. Algorithm Bellerophon can be modified to deal with overflow and underflow by testing the n -bit result to see if it is an infinity, the largest representable floating point number, the smallest normalized floating point number, denormalized, or zero. In such cases the computation may need to be repeated using some other algorithm, depending on the policies that have been established for handling overflow and underflow within the particular floating point number system in question.

With IEEE arithmetic, for example, a denormalized result may be required. Denormalized results can be generated

precision IEEE arithmetic faster than double or single precision IEEE arithmetic. By default, therefore, some compilers that appear to support IEEE single or double precision arithmetic may actually perform single or double precision calculations using extended precision, rounding to single or double only when a result is stored in a variable. Somewhat counterintuitively, this makes individual floating point operations less accurate, and does not meet the error bounds required by the IEEE specification for the default rounding mode using single or double precision arithmetic.

Suppose, for example, that Algorithm Bellerophon is used to compute the best IEEE double precision approximation to 1.448997445238699. The correct result is

$$6525704354437805 \times 2^{-52} \doteq 1.448997445238699$$

obtained by dividing 1448997445238699 by 10^{15} using double precision arithmetic. If this division is performed using 64-bit extended precision arithmetic instead, and the extended precision result rounded to double precision, then the incorrect result

$$6525704354437806 \times 2^{-52} \doteq 1.4489974452386991$$

will be obtained.

10. RELATED WORK

Mathematical properties of the best approximation function have been investigated by Matula, who does not consider algorithms for computing it [Matula68, 70].

Theorem 4 strengthens an observation by Matula and others [Matula68]. Calculations similar to Lemma 2 and Corollary 3 appear in [Knuth81] and in most books on numerical analysis, though the results are seldom stated as they appear here. Theorem 5 expresses well-known folklore, but to my knowledge this is the first proof of it.

AlgorithmM is essentially the same as Method (2a) in Section 4.4 of [Knuth81]. The solution to Exercise 3 of that section contains a forward reference to [Steele90].

A draft of [Clinger90] required the standard routine for numerical output to print floating point numbers using the fewest digits that allow the number to be read back in without loss of accuracy. Although this can be done by extending an IEEE-conforming but imperfect implementation, reference was made to a draft of [Steele90], which assumes a perfect input routine.

On 1 November 1989 Chris Hanson expressed concern over this requirement in electronic mail sent to Steele, White, and myself. Hanson described AlgorithmM but noted its inefficiency and asked whether any other perfect algorithms, especially perfect and efficient algorithms, were published or known. The matter was urgent because Hanson was editing a draft IEEE standard for Scheme to be voted on in January. After checking with Steele to confirm that he did

not know of an efficient solution to the input problem, I set to work, keeping the others informed of my progress.

Jon L White was out of town and unable to read his mail. On 10 November 1989, after I had announced the basic idea of Algorithm Bellerophon, White reported that Lucid Common Lisp has for several years used a similar algorithm of his invention. This algorithm has not been published, and was known only to a handful of people at Lucid. From subsequent telephone conversations, it appears that the algorithm in use at Lucid is essentially the same as Bellerophon but uses twice as many bits, primarily because the error bounds were not calculated very tightly.

Bellerophon is so named because it inverts the Dragon3 and Dragon4 algorithms of [Steele90]. Unlike its namesake, the algorithm reads its fate and acts accordingly.

ACKNOWLEDGEMENTS

Chris Hanson helped in many ways, and I am indebted also to David Wise, Jon L White, Guy L Steele Jr, and Anne Hartheimer.

REFERENCES

- [Clinger90] Clinger, William, and Jonathan Rees [editors]. Revised⁴ report on the algorithmic language Scheme. Technical Report CIS-TR-90-02, Department of Computer and Information Science, University of Oregon, 1990.
- [Coonen80] Coonen, Jerome T. An implementation guide to a proposed standard for floating-point arithmetic. *Computer* 13, 1, January 1980, pages 68–79.
- [Goldberg67] Goldberg, I. B. 27 bits is not enough for 8-digit accuracy. *CACM* 10, 2, February 1967, pages 105–106.
- [HW60] Hardy, G. H., and E. M. Wright. *An Introduction to the Theory of Numbers, Fourth Edition*. Oxford University Press, 1960.
- [IEEE85] *IEEE Standard 754-1985. IEEE Standard for Binary Floating-Point Arithmetic*. IEEE, New York, 1985.
- [Knuth81] Knuth, Donald E. *The Art of Computer Programming, Second Edition, Volume 2, Seminumerical Algorithms*. Addison-Wesley, 1981.
- [Matula68] Matula, David W. In-and-out conversions. *CACM* 11, 1, January 1968, pages 47–50.
- [Matula70] Matula, David W. A formalization of floating-point numeric base conversion. *IEEE Transactions on Computers*, C-19, 8, August 1970, pages 681–692.
- [Rees86] Rees, Jonathan, and William Clinger [editors]. Revised³ report on the algorithmic language Scheme. *ACM SIGPLAN Notices* 21, 12, December 1986, pages 37–79.
- [Steele90] Steele Jr, Guy Lewis, and Jon L White. How to print floating point numbers accurately. Proceedings of this conference.