

Worst Cases and Lattice Reduction

Damien Stehlé

ENS Paris

45 rue d’Ulm

F-75005 Paris

damien.stehle@ens.fr

Vincent Lefèvre

LORIA/INRIA Lorraine

Technopôle de Nancy-Braboïs

615 rue du Jardin Botanique

F-54602 Villers-lès-Nancy Cedex

lefevre@loria.fr

Paul Zimmermann

LORIA/INRIA Lorraine

Technopôle de Nancy-Braboïs

615 rue du Jardin Botanique

F-54602 Villers-lès-Nancy Cedex

zimmerma@loria.fr

Abstract

We propose a new algorithm to find worst cases for correct rounding of an analytic function. We first reduce this problem to the real small value problem — i.e. for polynomials with real coefficients. Then we show that this second problem can be solved efficiently, by extending Copper-Smith’s work on the integer small value problem — for polynomials with integer coefficients — using lattice reduction [4, 5, 6].

For floating-point numbers with a mantissa less than N , and a polynomial approximation of degree d , our algorithm finds all worst cases at distance $< N^{\frac{d+1}{2d+1}}$ from a machine number in time $O(N^{\frac{d+1}{2d+1}+\varepsilon})$. For $d = 2$, this improves on the $O(N^{2/3+\varepsilon})$ complexity from Lefèvre’s algorithm [15, 16] to $O(N^{3/5+\varepsilon})$. We exhibit some new worst cases found using our algorithm, for double-extended and quadruple precision. For larger d , our algorithm can be used to check that there exist no worst cases at distance $< N^{-k}$ in time $O(N^{\frac{1}{2}+O(\frac{1}{k})})$.

1 Introduction

The IEEE-754 standard for binary floating-point arithmetic [11], approved in 1985 by the IEEE Standards Board and the American National Standards Institute, requires that all four basic arithmetic operations ($+$, $-$, \times , \div) and the square root are correctly rounded. For a given function, floating-point inputs for which it is difficult to guarantee correct rounding, called *worst cases*, are numbers for which the exact result — as computed in infinite precision — is near a machine number, or near the middle of two consecutive machine numbers. This is the famous “Table Maker’s Dilemma” problem (TMD for short). Several authors [13, 21, 12, 14, 19] have shown that for the class of algebraic functions, such worst cases cannot be too near

from a machine number or the middle of two consecutive machine numbers. Such bounds enable one to design some efficient algorithms that guarantee correct rounding for division and square root, and less efficient algorithms for other algebraic functions.

However, for non-algebraic functions, number theory bounds are not sharp enough, which makes correct rounding harder to implement. This is probably the reason why the IEEE-754 standard does not require correct rounding for those functions. Muller and other authors proposed in [20] to introduce different levels of quality for transcendental functions. This proposal was presented by Markstein at the May 2002 meeting of the IEEE-754 revision group, but the conclusion was that “we’re not yet ready to standardize”.

Systematic work on the Table Maker’s Dilemma was done by Lefèvre and Muller [16], who published worst cases for many elementary functions in double precision ($N = 2^{53}$), over the full range for some functions. Alas, their approach is too expensive to deal with the quadruple precision, which is included in the current revision of the IEEE-754 standard. Thus currently the only possible approaches for higher precisions are either to guess a reasonable bound on the precision required for the hardest to round cases and to write a library computing up to that precision, or to write a generic multiple-precision library. For instance, Ziv’s MathLib library does the former, where the guessed bound is 768 bits for double precision [22].

Having an efficient algorithm to find the hardest to round cases, for a given function and a given floating-point format, would help to replace guessed bounds — which are usually overestimated — by sharper and rigorous bounds. It would thus enable one to design very efficient libraries with correct rounding [7, 8]. Then there would be no good reason any more to exclude those functions from the correct rounding requirements of the IEEE-754 standard.

Exhaustive search methods consist in finding the hardest to round cases of the given function in the given range. They

give the best possible bound, but are very time-consuming. Moreover, a search for a given precision gives little knowledge for another precision. We propose here a new algorithm belonging to that class. It naturally extends the first algorithm proposed by Lefèvre [15], and is based on Coppersmith’s ideas.

Previous related work was done by Elkies, who gives in [9] a new algorithm using lattice reduction to find all rational points of small height near a plane curve; for example, his record:

$$5853886516781223^3 - 447884928428402042307918^2 = 1641843$$

corresponds to a worst case of the function $x^{3/2}$ for a 53-bit input and a 79-bit output; his other example

$$2220422932^3 - 283059965^3 - 2218888517^3 = 30$$

corresponds to a worst case of the function $(x^3 + y^3)^{1/3}$ in 32-bit arithmetic. More recently Gonnet [10] also used lattice reduction to find worst cases, however his approach seems equivalent to Lefèvre’s algorithm.

Our paper is organized as follows: Section 2 explains in mathematical terms the problem we want to solve, recalls Lefèvre’s algorithm and analyzes its complexity. Section 3 describes our new algorithm, after a short survey on lattice reduction and Coppersmith’s work, which we heavily use. Section 4 presents some new worst cases found with our algorithm for the 2^x function, in double-extended precision and quadruple precision. Section 5 discusses some ideas for possible improvements and open questions.

2 Preliminaries

2.1 Definitions and Notations

We assume we work here with floating-point numbers with a mantissa of n bits. Let $N = 2^n$; for instance, $N = 2^{53}$ corresponds to double precision, $N = 2^{64}$ corresponds to double-extended precision, and $N = 2^{113}$ corresponds to quadruple precision. A worst case for a function f is a floating-point number x such that $f(x)$ has m identical bits after the round bit. If all those m bits equal (resp. differ from) the round bit, x is a worst case for directed rounding (resp. rounding to nearest).

For sake of simplicity, we consider here directed rounding only (towards $-\infty$, towards $+\infty$, towards zero), since worst cases at precision n for all rounding modes are worst cases at precision $n + 1$ for directed rounding. (In general, we are interested in the worst cases for the inverse function too, in which case inputs are also chosen at precision $n + 1$ [16].) To find worst cases for directed rounding, we throw away the first n significant bits of the result




Figure 1. A function graph and the grid of machine numbers. Worst cases correspond to grid points with a small vertical distance to the curve.

mantissa. Then a worst case of length m corresponds to $|Nf(x) \bmod 1| < 2^{-m}$, where $x \bmod 1 := x - \lfloor x \rfloor$ denotes the “centered” fractional part (see Fig. 1).

We also consider that both argument x and result $y = f(x)$ are normalized, i.e. $\frac{1}{2} \leq x, f(x) < 1$. This is easy to achieve by multiplying x or $f(x)$ by some fixed powers of 2, unless the exponent of $f(x)$ varies a lot in the considered range. This excludes the case of numerically irregular functions like $\sin x$ for large x . Given a polynomial approximation $P(t)$ to $Nf(\frac{t}{N})$ (for example, a Taylor expansion), the Table Maker’s Dilemma can be reduced to the following problem.

REAL SMALL VALUE PROBLEM (REAL SVALP): Given positive integers M and T , and a polynomial P with real coefficients, find all integers $|t| < T$ such that

$$|P(t) \bmod 1| < \frac{1}{M}. \quad (1)$$

REMARK 1: The mantissa bound N does not appear explicitly in the real SVALP, however the polynomial $P(t)$ depends on N , and so does the error made in the polynomial approximation.

REMARK 2: If the fractional bits of the function behave randomly, we can expect $\approx \frac{T}{M}$ worst cases. Therefore we may assume $T \ll M$ if we want only few worst cases.¹

2.2 Lefèvre’s Algorithm

Lefèvre’s algorithm [15, 16] works as follows. One considers a linear approximation to the function f on small intervals. Those approximations are computed from higher order polynomial approximations on larger intervals, using an efficient scheme based on the “table of differences” method. On each small interval, worst cases are found using a modified version of the Euclidean algorithm, which gives a lower bound for $|Nf(\frac{t}{N}) \bmod 1|$ on that interval.

¹The notation $x \ll y$ is equivalent to $x = O(y)$.

In practice, Lefèvre’s algorithm is expensive but feasible for the double precision ($N^{2/3} \approx 4 \cdot 10^{10}$), near from the limits of the current processors for double-extended precision ($N^{2/3} \approx 7 \cdot 10^{12}$), and out of reach for quadruple precision ($N^{2/3} \approx 5 \cdot 10^{22}$).

³The $o()$ makes sense when α grows to infinity

vectors \mathbf{v}_1 and \mathbf{v}_2 of norms that are asymptotically less than $A^{\frac{2}{3}\alpha+o(\alpha)} \cdot X^{\frac{d}{3}\alpha+o(\alpha)} \cdot Y^{\frac{1}{3}\alpha+o(\alpha)}$. Those vectors \mathbf{v}_1 and \mathbf{v}_2 correspond to two polynomials $v_1(X\tau, Yv)$ and $v_2(X\tau, Yv)$. Moreover if $|x| \leq X$ and $|y| \leq Y$, then $|v_k(x, y)| \leq \sum_{i,j} |v_{i,j}^{(k)}| X^i Y^j \leq c \cdot \max |v_{i,j}^{(k)}| X^i Y^j \leq c \cdot \|\mathbf{v}_k\|$ for a certain constant c . Thus, to get $|v_k(x, y)| < A^\alpha$, it is sufficient that:

$$c \cdot A^{\frac{2}{3}\alpha+o(\alpha)} \cdot X^{\frac{d}{3}\alpha+o(\alpha)} \cdot Y^{\frac{1}{3}\alpha+o(\alpha)} < A^\alpha,$$

which asymptotically gives the bound $X^d Y \ll A$.

Using Coppersmith's technique, one can thus solve the integer SValP in polynomial time as long as $X^d Y < A^{1-\epsilon}$. In fact, this is not completely true because we used an argument we cannot prove: we assumed that $\text{Res}_y(v_1, v_2) \neq 0$. This heuristic has been made very often in cryptography (see [2, 3, 4]).

3.3 The SLZ Algorithm

The real SValP is the following problem: Given a polynomial P , find for which integers t , $P(t)$ is near an integer. We solve this problem by reducing it to the integer SValP. The difficulty is that $P(t)$ has real coefficients, and the LLL algorithm does not work well with real input. The following algorithm overcomes that difficulty (we present here a complete algorithm to solve the Table Maker's Dilemma, in which case $P(t) = Nf(t/N)$, but the sub-algorithm consisting of steps 3 to 11 may be of interest to solve the real SValP itself).

Input: a function f , positive integers N, T, M, d, α
Output: all worst cases at distance $< 1/M$ for $f(\frac{t}{N})$ for $|t| \leq T$

1. Let $P(t)$ be the Taylor expansion of $Nf(\frac{t}{N})$ up to order d , and $n = \frac{(\alpha+1)(d\alpha+2)}{2}$
2. Compute a bound ε such that $|P(t) - Nf(\frac{t}{N})| < \varepsilon$ for $|t| \leq T$
3. Let $M' = \lfloor \frac{1/2}{1/M+\varepsilon} \rfloor$, $C = (d+1)M'$, and $P'(\tau) = \lfloor CP(T\tau) \rfloor$, $Q(t) = T^d P'(\frac{t}{T})$
4. Let $\{e_1, \dots, e_n\} \leftarrow \{\tau^i v^j\}$ for $0 \leq i+dj \leq d\alpha$
5. Let $Y = (d+1)T^d$, $A = CT^d$ and $\{g_1, \dots, g_n\} \leftarrow \{(T\tau)^i (Q(T\tau) + Yv)^j A^{\alpha-j}\}$ for $0 \leq i+dj \leq d\alpha$
6. Form the $n \times n$ integral matrix L where $L_{k,l}$ is the coefficient of the monomial e_k in g_l
7. $V \leftarrow \text{LatticeReduce}(L)$
8. Let $\mathbf{v}_1, \mathbf{v}_2$ be the two smallest vectors from V , and $Q_1(T\tau, Yv)$ and $Q_2(T\tau, Yv)$ the corresponding polynomials

9. If there exists $(t, y) \in [-T, T] \times [-Y, Y]$ with $|Q_1(t, y)| \geq A^\alpha$ or $|Q_2(t, y)| \geq A^\alpha$, return(FAIL)
10. $p(t) \leftarrow \text{Res}_y(Q_1(t, y), Q_2(t, y))$;
if $p(t) = 0$ **then** return(FAIL)
11. **for** each t_0 in $\text{IntegerRoots}(p(t), [-T, T])$ **do**
if $|Nf(\frac{t_0}{N}) \bmod 1| < 1/M$ **then** output t_0 .

3.4 Correctness of the Algorithm

Theorem 2 *In case algorithm SLZ does not return FAIL, it behaves correctly, i.e. it prints exactly all integers $t \in [-T, T]$ such that $|Nf(\frac{t}{N}) \bmod 1| < 1/M$.*

PROOF. Because of the final check in step 11, we only have to verify that no worst case is missed. Suppose there is $t_0 \in [-T, T]$ with $|Nf(\frac{t_0}{N}) \bmod 1| < 1/M$. We first prove that $|Q(t_0) \bmod CT^d| \leq (d+1)T^d$. From the definition of P , $|P(t_0) \bmod 1| < 1/M + \varepsilon \leq \frac{1}{2M'}$. Since $|CP(T\tau) - P'(\tau)| \leq \frac{d+1}{2}$ for $|\tau| \leq 1$, by choosing $\tau = t_0/T$ we get $|Q(t_0) \bmod CT^d| < \frac{(d+1)T^d}{2} + \frac{(d+1)T^d}{2}$.

Whence $Q(t) + y = 0 \bmod CT^d$ has a root (t_0, y_0) with $|t_0| \leq T$ and $|y_0| < (d+1)T^d$. Since $Q_1(t, y)$ and $Q_2(t, y)$ are linear combinations of $Q(t) + y$ and its powers, then (t_0, y_0) is a common root of $Q_1(t, y)$ and $Q_2(t, y)$ modulo CT^d , and even over the reals since $|Q_1|, |Q_2| < CT^d$. Thus t_0 is an integer root of $\text{Res}_y(Q_1(t, y), Q_2(t, y))$, and will be found at step 11. \square

3.5 Choice of Parameters and Complexity Analysis

3.5.1 Coppersmith's Bound

Because of the use of Coppersmith's technique in our algorithm, to insure the algorithm does not return FAIL at step 9, the bound " $X^d Y \ll A$ " has to be verified. In our case, X corresponds to T , Y to $(d+1)T^d$ and A to CT^d , so we get:

$$T \ll M^{\frac{1}{d}}.$$

3.5.2 Choice of the Degree d With Respect to T

Let $(a_i)_i$ the Taylor coefficients of f . Since we neglect Taylor coefficients of degree $d+1$ and greater, the error made in the approximation to $Nf(\frac{t}{N})$ by $P(t)$ is $\approx a_{d+1}T^{d+1}N^{-d}$. Since we are looking for worst cases with $|P(t) \bmod 1| < 1/M$, we want $T^{d+1}N^{-d} \ll 1/M$, i.e. $T^{d+1} \ll N^d/M$.

3.5.3 Complexity Analysis

Thus we have two bounds for T : the first one $T \ll M^{1/d}$ comes from Coppersmith's method, the second one

⁴The notation $\lfloor CP(T\tau) \rfloor$ means that we round to the nearest integer each coefficient of $CP(T\tau)$. This provides an element of $\mathbb{Z}[\tau]$.

$T^{d+1} \ll N^d/M$ comes from the accuracy of the Taylor expansion. Therefore for $M \ll N^{\frac{d^2}{2d+1}}$, Coppersmith's bound wins and implies $T \ll M^{1/d}$, whereas for $M \gg N^{\frac{d^2}{2d+1}}$, Lagrange's bound gives $T^{d+1} \ll N^d/M$. The largest bound for T is obtained for $M \sim N^{\frac{d^2}{2d+1}}$, with $T \ll N^{\frac{d}{2d+1}}$. For $d = 1$, we find the constraint $T \ll N^{1/3}$ from Lefèvre's method; for $d = 2$, this gives $T \ll N^{2/5}$ with $M \sim N^{4/5}$; for $d = 3$, this gives $T \ll N^{3/7}$ with $M \sim N^{9/7}$. With $M \sim N^k$, we get a best possible interval length $T \sim N^{\frac{1}{2} - \frac{1}{8k} + o(\frac{1}{k})}$.

3.5.4 Working Precision

In step 1, we can use floating-point coefficients in the Taylor expansion $P(t)$ instead of symbolic coefficients, as long as it introduces no error in step 3 while computing $P'(\tau)$. Let a_i be the i th Taylor coefficient of f . Then to get $P'(\tau)$ correct at step 3, the error on $CN(T/N)^i a_i$ must be less than $1/2$, thus the error on a_i must be less than $1/(2CN)(N/T)^i$. Since $N \geq T$, it thus suffices to compute a_i with $\log_2(2CN)$ bits after the binary point.

REMARK 3: When searching worst cases with $M \ll N$, degree 2 is enough. Indeed, $N^{1-d}T^d \ll N^{1-d}M$ since $T^d \ll M$ (Coppersmith's bound), and for $d \geq 3$, $N^{1-d}T^d \ll N^{2-d} \ll 1/N \ll 1/M$. Thus all Taylor terms of degree ≥ 3 give a negligible contribution to $Nf(\frac{t}{N})$, and the largest value of T is $N^{2/5}$, giving a complexity of $N^{3/5}$ to search a whole range of $N/2$ values. More generally, for $M \ll N^k$, degree $2k$ is enough, giving a complexity of $N^{\frac{2k}{4k+1}}$.

4 Experimental Results

We have implemented algorithm SLZ in the Pari/GP system (version 2.2.4-alpha) [1] and experimented it on a Athlon XP 1600+ under Linux. We have chosen the 2^x function since it is the easiest one, with only one exponent range to study. Fig. 2 shows for each target precision (double, double-extended, quadruple), and for $M \approx N$ and $M \approx N^2$, the best parameters (T , d , and α) for our method, together with the estimated time to check the whole exponent range, i.e. $N/2$ floating-point numbers. For each precision, the first row gives the best parameters for the $d = \alpha = 1$ case, which is what Gonnet considers in [10]; comparing that first row to the following ones shows the speedup obtained. For $M \approx N$, the speedup increases from 3 to 17, whence is not dramatic. However for $M \approx N^2$, we get a speedup of about 1000 in quadruple precision with respect to the naive method ($d = \alpha = 1$), with $(d, \alpha) = (4, 2)$.

Fig. 3 shows a few worst cases found using algorithm SLZ for double-extended and quadruple precision. These experiments tend to show that with a carefully

	N	M	T	d	α	est. time
double	2^{53}	2^{28}	2^{15}	1	1	560 days
	2^{53}	2^{53}	2^{20}	2	2	120 days
precision	2^{53}	2^{106}	2^{25}	4	2	45 days
double	2^{64}	2^{32}	2^{19}	1	1	140 years
extended	2^{64}	2^{64}	2^{24}	2	2	43 years
precision	2^{64}	2^{128}	2^{30}	4	2	9 years
quadruple	2^{113}	2^{70}	2^{35}	1	1	1600 Gyears
	2^{113}	2^{113}	2^{43}	2	2	94 Gyears
precision	2^{113}	2^{226}	2^{53}	4	2	1.6 Gyears

Figure 2. Best experimental parameters for double, double-extended and quadruple precision, and estimated time for an exponent range of $N/2$ values.

N	t_0	$N2^{-1/2+t_0/N} \bmod 1$
2^{64}	586071771766963	0.1147001111...
2^{64}	594068190588573	0.0048100010...
2^{64}	891586182147388	0.1150001000...
2^{64}	9014384889202147	0.0153010011...
2^{64}	9602866023852631	0.0054111001...
2^{113}	1119374922072865495	0.0163000000...
2^{113}	8923960372306650064	0.0064101011...
2^{113}	43616445401128570224	0.0165011110...
2^{113}	53608038600996804036	0.0167000001...

Figure 3. Some worst cases found for the 2^x function in double-extended and quadruple precision.

tuned implementation, and several computers running a few months, solving the Table Maker's Dilemma for the double-extended precision is nowadays feasible for simple elementary functions.

5 Possible Improvements, Open questions

We have presented a new algorithm, based on lattice reduction, to search for worst cases for correct rounding of analytic functions. The first experimental results show that algorithm SLZ is quite efficient, especially to detect worst cases at distance much less than 2^{-n} , where n is the target precision. However the efficiency largely depends on the function considered, like in Lefèvre's algorithm.

Several open questions remain. Does this approach extend like in the modular case ([4]) to functions of two variables like x^y or $\arctan \frac{x}{y}$?

Our algorithm is complementary to that of Elkies [9],

which works well when $M \ll N$ (in our notation), i.e. when we expect many worst cases, whereas our algorithm is more efficient when $M \gg N$, i.e. when we expect only few worst cases, or none. However, in the case of $f(x) = x^{3/2}$, related to Hall's conjecture, Elkies proposes a special-purpose algorithm to find all worst cases at distance $< 1/N$ in $O(N^{1/2+\varepsilon})$. Does this algorithm generalize to other algebraic functions?

References

- [1] C. Batut, K. Belabas, D. Bernardi, H. Cohen, and M. Olivier. *User's Guide to PARI/GP*, 2000. <ftp://megrez.math.u-bordeaux.fr/pub/pari/manuals/users.pdf>.
- [2] D. Boneh and G. Durfee. Cryptanalysis of RSA with private key d less than $n^{0.292}$. In *Proceedings of Eurocrypt'99*, volume 1592 of *Lecture Notes in Computer Science*, pages 1–11. Springer-Verlag, 1999.
- [3] D. Boneh, G. Durfee, and N. Howgrave-Graham. Factoring $n = p^r q$ for large r . In *Proceedings of Eurocrypt'99*, volume 1592 of *Lecture Notes in Computer Science*, pages 326–337. Springer-Verlag, 1999.
- [4] D. Coppersmith. Finding a small root of a bivariate integer equation; factoring with high bits known. In *Proceedings of Eurocrypt'96*, volume 1070 of *Lecture Notes in Computer Science*, pages 178–189. Springer-Verlag, 1996.
- [5] D. Coppersmith. Finding a small root of a univariate modular equation. In *Proceedings of Eurocrypt'96*, volume 1070 of *Lecture Notes in Computer Science*, pages 155–165. Springer-Verlag, 1996.
- [6] D. Coppersmith. Finding small solutions to small degree polynomials. In *Proceedings of CALC'01*, volume 2146 of *Lecture Notes in Computer Science*, pages 20–31. Springer-Verlag, 2001.
- [7] D. Defour and F. de Dinechin. Software carry-save for fast multiple-precision algorithms. Research Report 2002-08, Laboratoire de l'Informatique du Parallélisme, 2002. 10 pages.
- [8] D. Defour, F. de Dinechin, and J.-M. Muller. Correctly rounded exponential function in double precision arithmetic. Research Report 2001-26, Laboratoire de l'Informatique du Parallélisme, 2001. 21 pages.
- [9] N. Elkies. Rational points near curves and small nonzero $|x^3 - y^2|$ via lattice reduction. In *Proceedings of ANTS-IV*, volume 1838 of *Lecture Notes in Computer Science*, pages 33–63. Springer-Verlag, 2000.
- [10] G. Gonnet. A note on finding difficult values to evaluate numerically. <http://www.inf.ethz.ch/personal/gonnet/FPAccuracy/NastyValues.ps>, Sept. 2002. 3 pages.
- [11] IEEE standard for binary floating-point arithmetic. Technical Report ANSI-IEEE Standard 754-1985, New York, 1985. approved March 21, 1985: IEEE Standards Board, approved July 26, 1985: American National Standards Institute, 18 pages.
- [12] C. S. Iordache and D. W. Matula. Infinitely precise rounding for division, square root, and square root reciprocal. In *Proceedings of the 14th IEEE Symposium on Computer Arithmetic*, pages 233–240. IEEE Computer Society, 1999.
- [13] W. Kahan. A test for correctly rounded SQRT. Lecture note, University of California at Berkeley, 1996. <http://www.cs.berkeley.edu/~wkahan/SQRTTest.ps>, 4 pages.
- [14] T. Lang and J.-M. Muller. Bounds on runs of zeros and ones for algebraic functions. In *Proceedings of the 15th IEEE Symposium on Computer Arithmetic*, pages 13–20. IEEE Computer Society, 2001.
- [15] V. Lefèvre. *Moyens arithmétiques pour un calcul fiable*. Thèse de doctorat, École Normale Supérieure de Lyon, Jan. 2000.
- [16] V. Lefèvre and J.-M. Muller. Worst cases for correct rounding of the elementary functions in double precision. In N. Burgess and L. Ciminiera, editors, *Proceedings of the 15th IEEE Symposium on Computer Arithmetic (ARITH'15)*, pages 111–118. IEEE Computer Society, 2001.
- [17] A. K. Lenstra, H. W. Lenstra, and L. Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261:515–534, 1982.
- [18] L. Lovász. An algorithmic theory of numbers, graphs and convexity. *SIAM lecture series*, 50, 1986.
- [19] L. D. McFearn and D. W. Matula. Generation and analysis of hard to round cases for binary floating point division. In *Proceedings of the 15th IEEE Symposium on Computer Arithmetic*, pages 119–127. IEEE Computer Society, 2001.
- [20] J.-M. Muller. Proposals for a specification of the elementary functions. In *Abstracts of SCAN'2002*, pages 54–55. Laboratory LIP6, Paris, France, 2002.
- [21] M. Parks. Number-theoretic test generation for directed rounding. In *Proceedings of the 14th IEEE Symposium on Computer Arithmetic*, pages 241–248. IEEE Computer Society, 1999.
- [22] A. Ziv. Fast evaluation of elementary mathematical functions with correctly rounded last bit. *ACM Trans. Math. Softw.*, 17(3):410–423, 1991.