



HAL
open science

The Positive Effects of Stochastic Rounding in Numerical Algorithms

El-Mehdi El Arar, Devan Sohier, Pablo de Oliveira Castro, Eric Petit

► **To cite this version:**

El-Mehdi El Arar, Devan Sohier, Pablo de Oliveira Castro, Eric Petit. The Positive Effects of Stochastic Rounding in Numerical Algorithms. 29th IEEE Symposium on Computer Arithmetic ARITH 2022, Sep 2022, Virtual conference, France. hal-03716058

HAL Id: hal-03716058

<https://hal.science/hal-03716058>

Submitted on 7 Jul 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The Positive Effects of Stochastic Rounding in Numerical Algorithms

El-Mehdi El Arar¹, Devan Sohier¹, Pablo de Oliveira Castro¹, Eric Petit²

¹Université Paris-Saclay, UVSQ, LI-PaRAD

²Intel Corp

{el-mehdi.el-arar, devan.sohier, pablo.oliveira}@uvsq.fr

eric.petit@intel.com

Abstract – Recently, stochastic rounding (SR) has been implemented in specialized hardware but most current computing nodes do not yet support this rounding mode. Several works empirically illustrate the benefit of stochastic rounding in various fields such as neural networks and ordinary differential equations. For some algorithms, such as summation, inner product or matrix-vector multiplication, it has been proved that SR provides probabilistic error bounds better than the traditional deterministic bounds.

In this paper, we extend this theoretical ground for a wider adoption of SR in computer architecture. First, we analyze the biases of the two SR modes: SR-nearness and SR-up-or-down. We demonstrate on a case-study of Euler’s forward method that IEEE-754 default rounding modes and SR-up-or-down accumulate rounding errors across iterations and that SR-nearness, being unbiased, does not. Second, we prove a $O(\sqrt{n})$ probabilistic bound on the forward error of Horner’s polynomial evaluation method with SR, improving on the known deterministic $O(n)$ bound.

Keywords – Stochastic rounding, Floating-point arithmetic, Polynomial evaluation, Horner’s algorithm, Numerical integration.

I INTRODUCTION

In floating-point arithmetic, rounding errors occur because of the finite representation of floating-point numbers in computers. The accumulation of rounding errors can significantly reduce the accuracy of a computation [1], p. 8.

Stochastic arithmetic proposed in the 1950s by von Neumann and Goldstine [2] is a computing paradigm developed as a model for exact computation with imprecise data. However, hardware units proposing stochastic rounding (SR) are still unavailable in most computer machines. Various specialized processors have introduced it such as: Graphcore IPU which supports SR for binary32 and binary16 arithmetic [3]–[5], or Intel neuromorphic chip Loihi [6] to improve the accuracy of biological neuron and synapse models. Moreover, several related patents are owned by major chip designers such as AMD [7], NVIDIA [8], IBM [9], [10], and other computing companies [11]–[13]. These developments support the idea of hardware implementations using stochastic rounding becoming more available in the future.

Stochastic arithmetic has two main applications [14]. First, it can be used to estimate empirically the numerical error of complex programs. In this context, stochastic arithmetic

introduces a random noise in each floating-point operation to simulate the effect of rounding errors. To make this error analysis automatic, various tools such as Verificarlo [15], Cadna [16] and Verrou [17] have been developed. Second, stochastic arithmetic is sometimes used as a replacement for the default deterministic rounding mode in numerical simulations. Indeed, it has been demonstrated that in multiple domains, SR provides positive effects compared to the deterministic IEEE-754 [18] default rounding mode.

In particular, in the neural networks field, using SR instead of deterministic rounding [19] enables training with smaller data types and better accuracy. Indeed, the noise during training has been observed to help regularization and avoid model over-fitting [20], [21]. The positive effect of SR extends also to the calculation of the solution of ordinary differential equations (ODEs) in low precision [22], [23] where SR reduces the accumulation of rounding errors by avoiding stagnation phenomenon when the step decreases. Various other applications such as PDEs, Quantum mechanics, Quantum computing use SR to improve their results [14].

The IEEE-754 norm defines five rounding modes for floating-point arithmetic which are all deterministic: round to nearest ties to even (default), round to nearest ties away, round to zero, round to $+\infty$, and round to $-\infty$. In section II, we present a floating-point arithmetic background, and, we describe two stochastic rounding modes defined in [1], p. 34: SR-nearness and SR-up-or-down. Section III presents our first contribution: we study the bias and we compare the two stochastic rounding modes above and the default rounding mode in the IEEE-754 norm (RN-nearest32) on rectangular integration, which is at the basis of Euler’s Method for ODE. We show that, contrarily to SR-up-or-down, SR-nearness is unbiased. An exact expression and an estimation of the bias are given for SR-up-or-down. We show how the accumulation of errors with both SR-up-or-down and IEEE-754 modes leads to results significantly less accurate than with SR-nearness.

At the beginning of section IV, a probabilistic background, as well as some properties of SR-nearness rounding are presented. We recall some techniques that have been used in order to obtain a probabilistic bound for the inner product on \sqrt{n} instead of the deterministic bound on n . We explain two approaches proposed by Higham and Mary [24] and Ilse, Ipsen and Zhou [25]. We extend these techniques for our second contribution and showing that they can still be used in situations where one multiplication operand is affected by an error. In particular, using SR-nearness and without additional assumption, we provide a probabilistic bound in $O(\sqrt{n})$ rather than the deterministic bound which is in $O(n)$. We conclude this section with numerical experiments comparing the bounds

above for the Chebyshev polynomial.

II Background

A. Floating-point arithmetic

A normal floating-point number in such a format is a number x for which there exists a triple (s, m, e) such that $x = \pm m \times \beta^{e-p}$, where β is the basis, e is the exponent, p is the working precision, and m is an integer (the significand) such that $\beta^{p-1} \leq m < \beta^p$. This triple is unique. We only consider normal floating-point numbers; detailed information on the floating-point format most generally in use in current computer systems is defined in the IEEE-754 norm [18].

Let us denote $\mathcal{F} \subset \mathbb{R}$ the set of normal floating-point numbers and $x \in \mathbb{R}$. Upward rounding $\lceil x \rceil$ and downward rounding $\lfloor x \rfloor$ are defined by:

$$\lceil x \rceil = \min\{y \in \mathcal{F} : y \geq x\}, \quad \lfloor x \rfloor = \max\{y \in \mathcal{F} : y \leq x\},$$

clearly, $\lfloor x \rfloor \leq x \leq \lceil x \rceil$, with equalities if and only if $x \in \mathcal{F}$.

The floating-point approximation of a real number $x \neq 0$ is one of $\lfloor x \rfloor$ or $\lceil x \rceil$.

$$\text{fl}(x) = x(1 + \delta), \quad (1)$$

where $\delta = \frac{\text{fl}(x) - x}{x}$ is the relative error: $|\delta| \leq \beta^{1-p}$. In the following, we note $u = \beta^{1-p}$. IEEE-754 mode RN (round to nearest, ties to even) has the stronger property¹ that $|\delta| \leq \frac{1}{2}\beta^{1-p} = \frac{1}{2}u$.

For $x, y \in \mathcal{F}$, the considered rounding modes verify $\text{fl}(x \text{ op } y) \in \{\lfloor x \text{ op } y \rfloor, \lceil x \text{ op } y \rceil\}$ for $\text{op} \in \{+, -, *, /\}$. Moreover, for IEEE-754 RN [18] and stochastic rounding [26] the error in one operation is bounded:

$$\text{fl}(x \text{ op } y) = (x \text{ op } y)(1 + \delta), \quad |\delta| \leq u, \quad (2)$$

specifically for RN we have $|\delta| \leq \frac{1}{2}u$.

Assume that x is a real that is not representable: $x \in \mathbb{R} \setminus \mathcal{F}$. The machine-epsilon or the distance between the two floating-point numbers enclosing x is $\varepsilon(x) = \lceil x \rceil - \lfloor x \rfloor = \beta^{e-p}$. The fraction of $\varepsilon(x)$ rounded away, as shown in figure 1, is $\theta(x) = \frac{x - \lfloor x \rfloor}{\lceil x \rceil - \lfloor x \rfloor}$




Fig. 1. $\theta(x)$ is the fraction of $\varepsilon(x)$ to be rounded away.

We note $\llbracket x \rrbracket$ the integer part of x . The following lemma gives an important property of downward rounding.

Lemma II.1. Let $x \in \mathbb{R} \setminus \mathcal{F}$. $\beta^{p-e} \lfloor x \rfloor = \llbracket \beta^{p-e} x \rrbracket$.

Proof. We know that $\beta^{p-e} \lfloor x \rfloor, \beta^{p-e} \lceil x \rceil \in \mathbb{Z}$, and $\lfloor x \rfloor < x < \lceil x \rceil$, then $\beta^{p-e} \lfloor x \rfloor < \beta^{p-e} x < \beta^{p-e} \lceil x \rceil$. We thus have

$$\beta^{p-e} \lfloor x \rfloor \leq \llbracket \beta^{p-e} x \rrbracket < \beta^{p-e} \lceil x \rceil.$$

¹In many works focusing on IEEE-754 RN, u is chosen to be $\frac{1}{2}\beta^{1-p}$.

Since $\lceil x \rceil - \lfloor x \rfloor = \beta^{e-p}$, then $\beta^{p-e} \lceil x \rceil - \beta^{p-e} \lfloor x \rfloor = 1$ and

$$\beta^{p-e} \lfloor x \rfloor \leq \llbracket \beta^{p-e} x \rrbracket < \beta^{p-e} \lceil x \rceil + 1.$$

■

B. Stochastic arithmetic

In order to define the two principal stochastic rounding modes, consider \hat{x} the random variable of the distribution of results after random rounding of x . Then

$$\hat{x} = \mathbf{random_round}(x) = \mathbf{round}(x + \beta^{e-p} \xi),$$

where ξ is a random variable that can be discrete or continuous and \mathbf{round} is the default IEEE-754 rounding mode to the nearest.

- SR-nearness [1], p. 34 is defined for an ξ uniform random variable on $]-\frac{1}{2}; \frac{1}{2}[$, hence ξ has mean 0 and standard deviation $1/\sqrt{12}$.




Fig. 2. SR-nearness.

In other words, SR-nearness consists in rounding up $x \in \mathbb{R} \setminus \mathcal{F}$ with probability $\theta(x) = (x - \lfloor x \rfloor) / (\lceil x \rceil - \lfloor x \rfloor)$, we round down with probability $1 - \theta(x)$, proportional to the distances between x and the closest representable numbers.

Stott Parker shows that SR-nearness is unbiased [1], p. 34 so $\mathbb{E}(\hat{x}) = x$.

- SR-up-or-down is defined for a ξ uniform random variable on $]-\theta(x); 1 - \theta(x)[$, ξ is biased with mean $\frac{1}{2} - \theta(x)$ and standard deviation $1/\sqrt{12}$.




Fig. 3. SR-up-or-down.

In other words, SR-up-or-down consists in rounding x up or down with probability $\frac{1}{2}$. SR-up-or-down mode can be expressed [1], p. 34 in terms of $\theta(x)$: since the two outcomes of SR-up-or-down mode are equiprobable, we have $\mathbb{E}(\hat{x}) = \frac{\lceil x \rceil + \lfloor x \rfloor}{2}$, which allow us to write the bias as

$$\mathbb{E}(\hat{x} - x) = \frac{\lceil x \rceil + \lfloor x \rfloor}{2} - x,$$

because $\theta(x) = \frac{x - \lfloor x \rfloor}{\lceil x \rceil - \lfloor x \rfloor}$

$$\begin{aligned} \mathbb{E}(\hat{x} - x) &= (\lceil x \rceil - \lfloor x \rfloor) \left(\frac{1}{2} - \theta(x) \right) \\ &= \varepsilon(x) \left(\frac{1}{2} - \theta(x) \right). \end{aligned}$$

Thus, we conclude that SR-up-or-down is biased and the

expected value depends on $\theta(x)$ and $\varepsilon(x)$.

III Integrating a constant function

Rectangular integration rule is a classic approximation for performing numerical integration: the area under a curve is approximated by a sum of N rectangle areas:

$$\int_a^b f(t)dt \approx \sum_{k=0}^{N-1} hf(a+kh)$$

where $h = \frac{b-a}{N}$. In particular, rectangular rule is one of the resolution techniques for ODE using Euler's forward method.

Verrou's tutorial [27] integrates the cosinus function with the rectangular rule; with deterministic round to nearest or SR-up-or-down modes, the solution is biased. When the number of integration steps grows, this bias can become high and degrade the quality of the solution. In this section, we show why deterministic and SR-up-or-down modes are sometimes biased with rectangular rule.

We perform the analysis on a constant function $f(t) = 1$ for all $t \in [0; 1]$. With f constant, the evaluation error is zero, making it clear how the numerical error accumulates on the summation.

Denote $x = 1 = \sum_{k=0}^{N-1} h$, where $h = 1/N$. The distribution \hat{x} is produced by summing N times the integration step h . We note \hat{s}_k the random variable for the partial sum at step $0 \leq k \leq N-1$ and s_k the exact expected result, with $\hat{s}_{N-1} = \hat{x}$.

SR-up-or-down : As shown before, for each \hat{s}_k we introduce a bias corresponding to

$$\mathbb{E}(\hat{s}_k - s_k) = \varepsilon(s_k) \left(\frac{1}{2} - \theta(s_k) \right),$$

from the definition of $\theta(s_k)$, we have $0 < \theta(s_k) < 1$, then $-\frac{1}{2} < \frac{1}{2} - \theta(s_k) < \frac{1}{2}$ and

$$|\mathbb{E}(\hat{s}_k - s_k)| < \frac{1}{2} \varepsilon(s_k).$$

Table I shows these different values for $N = 20$.

Interestingly, in this table, we note that $\theta(s_k)$ is constant between two powers-of-the-base except for the first value. For example for $9 < k < 20$, s_k stays within $[2^{-1}; 2^0)$ and both $\theta(s_k)$ and $\mathbb{E}(S_k - s_k)$ are constant. Let us show why that is always the case.

Suppose $s_k \in [\beta^e; \beta^{e+1})$. Then $\varepsilon(s_k) = \beta^{e-p}$. At each step the next partial sum is computed as, $s_{k+1} = \text{fl}(s_k) + h$, in that case, using the lemma II.1, we have

$$\begin{aligned} \theta(s_{k+1}) &= \beta^{p-e}(s_{k+1} - \lfloor s_{k+1} \rfloor) \\ &= \beta^{p-e}s_{k+1} - \lfloor \beta^{p-e}s_{k+1} \rfloor \\ &= \beta^{p-e}\text{fl}(s_k) + \beta^{p-e}h - \lfloor \beta^{p-e}\text{fl}(s_k) + \beta^{p-e}h \rfloor. \end{aligned}$$

Since $\text{fl}(s_k) \in \mathcal{F}$, we have $\beta^{p-e}\text{fl}(s_k) \in \mathbb{Z}$ and

$$\lfloor \beta^{p-e}\text{fl}(s_k) + \beta^{p-e}h \rfloor = \beta^{p-e}\text{fl}(s_k) + \lfloor \beta^{p-e}h \rfloor.$$

TABLE I

s_k, θ, bias and ε for $N = 20$.

k	s_k	$\theta(s_k)$	$\mathbb{E}(S_k - s_k)$	$\varepsilon(s_k)$
2	0.150...	0.7500	-3.725290e-09	1.490116e-08
3	0.200...	0.2500	3.725290e-09	1.490116e-08
4	0.250...	0.6250	-3.725290e-09	2.980232e-08
5	0.300...	0.6250	-3.725290e-09	2.980232e-08
6	0.350...	0.6250	-3.725290e-09	2.980232e-08
7	0.400...	0.6250	-3.725290e-09	2.980232e-08
8	0.450...	0.6250	-3.725290e-09	2.980232e-08
9	0.500...	0.3125	1.117587e-08	5.960464e-08
10	0.550...	0.8125	-1.862645e-08	5.960464e-08
11	0.600...	0.8125	-1.862645e-08	5.960464e-08
12	0.650...	0.8125	-1.862645e-08	5.960464e-08
13	0.700...	0.8125	-1.862645e-08	5.960464e-08
14	0.749...	0.8125	-1.862645e-08	5.960464e-08
15	0.799...	0.8125	-1.862645e-08	5.960464e-08
16	0.849...	0.8125	-1.862645e-08	5.960464e-08
17	0.899...	0.8125	-1.862645e-08	5.960464e-08
18	0.949...	0.8125	-1.862645e-08	5.960464e-08
19	0.999...	0.8125	-1.862645e-08	5.960464e-08

Finally

$$\theta(s_{k+1}) = \beta^{p-e}h - \lfloor \beta^{p-e}h \rfloor = \theta(h).$$

Thus $\theta(s_{k+1})$ depends only on h and e . Recursively for all $l > 0$ satisfying $s_{k+l} \in [\beta^e; \beta^{e+1})$, $\theta(s_{k+l}) = \beta^{p-e}h - \lfloor \beta^{p-e}h \rfloor = \theta(h)$ is constant. The bias

$$\begin{aligned} \mathbb{E}(\hat{s}_k - s_{k+l}) &= \varepsilon(s_k) \left(\frac{1}{2} - \theta(s_{k+l}) \right) \\ &= \beta^{e-p} \left(\frac{1}{2} - \beta^{p-e}h - \lfloor \beta^{p-e}h \rfloor \right), \end{aligned}$$

is also constant in this interval.

Between two powers of the base, θ remain constant as well as the bias. Because the bias is constant (and, consequently, its sign too), it accumulates across iterations.

The total bias can be written as,

$$\mathbb{E}(\hat{x} - x) = E(\hat{s}_{N-1} - s_{N-1}) = \varepsilon(s_{N-1}) \left(\frac{1}{2} - \theta(s_{N-1}) \right),$$

and

$$|\mathbb{E}(\hat{x} - x)| < \frac{1}{2} \varepsilon(s_{N-1}).$$

We can neglect the effect of the first partial sum in each power-of-the-base interval. Since s_k increases arithmetically, and the size of the power of the base interval grows geometrically, the bias in the last interval usually dominates since it contains more summation terms and has a larger ε .

Numerical experiment. We have verified numerically that the above expression for the bias closely predicts the bias measured with SR-up-or-down stochastic rounding.

We consider a fixed number of iterations N . We ran one time the C program in listing 1 with each of the two previously defined stochastic rounding modes as well as round to nearest.

```

Listing 1: Fixed-step rectangle integration of a constant
float h = 1/N;
float s = 0.0;
for (int i=0 ; i < N ; i++) {
    s += h*1;
}
return s;

```

Figure 4 plots the three distributions over N .




Fig. 4. Round to nearest (RN-binary32) vs stochastic rounding SR-up-or-down and SR-nearness for $\int_0^1 1 dt$.

Figure 4 shows that SR-nearness mode samples is unbiased regardless of the number N of rectangles. The unbiased nature is unsurprising, since SR-nearness mode can be seen as a sub case of Monte Carlo Arithmetic (MCA). Stott Parker proves [1], p. 46 that the expectation of a sum of terms with MCA is the exact mathematical result.

On the other hand, SR-up-or-down mode and RN-binary-32 samples have a bias, which confirms the previous results for SR-up-or-down mode. The maximal amplitude of the bias for both SR-up-or-down and increases with N because of errors accumulation. The bias is reproducible and constant across different runs.

In conclusion, this example illustrates that SR-nearness is unbiased not only for one elementary operation, but, even in other numerical methods such as rectangular integration, it is much closer to the expected value than SR-up-or-down or RN-binary32, in particular for N large. In the remainder of this paper, we focus on SR-nearness.

IV Probabilistic bound

The aim of this section is to introduce a probabilistic bound in $O(\sqrt{nu})$ on the forward error of Horner algorithm, based on the Azuma–Hoeffding inequality and martingale properties. We first recall some probabilistic properties, then provide a review of the results made on the forward error in the numerically computed inner product, and we conclude with numerical experiments illustrating the previous results.

A. Probabilistic Background

A random variable Y is said to be mean independent of random variable X if and only if its conditional mean

$\mathbb{E}[Y/X = x]$ equals its unconditional mean $\mathbb{E}(Y)$ for all reals x such that the probability that $X = x$ is not zero and we write $\mathbb{E}[Y/X] = \mathbb{E}(Y)$. The random sequence (X_1, X_2, \dots) is mean independent if $\mathbb{E}[X_k/X_1, \dots, X_{k-1}] = \mathbb{E}(X_k)$ for all k .

Proposition IV.1. *Let X and Y two real random variables: X and Y are independents $\implies X$ is mean independent from $Y \implies X$ and Y are uncorrelated. The reciprocals of these two implications are false.*

Under stochastic rounding, the elementary operations are stochastically rounded, hence, for $x \in \mathbb{R}$, if $\text{fl}(x) = x(1 + \delta)$ is obtained by SR-nearness then δ is a random variable such that $\mathbb{E}(\delta) = 0$. For $x_1, x_2, x_3 \in \mathbb{R}$, such that $c = x_1 \text{ op } x_2 \text{ op } x_3$ where $\text{op} \in \{+, -, *, /\}$, and

$$\text{fl}(c) = ((x_1 \text{ op } x_2)(1 + \delta_1) \text{ op } x_3)(1 + \delta_2)$$

obtained from SR-nearness. δ_1, δ_2 are random variables such that $\mathbb{E}(\delta_1) = \mathbb{E}(\delta_2) = 0$.

The following lemma has been proven in [26], Lem 5.2 and shows that SR-nearness satisfies the property of mean independence.

Lemma IV.1. *For some $\delta_1, \delta_2, \dots$, in that order obtained from SR-nearness, the δ_k are random variables with mean zero such that $\mathbb{E}[\delta_k/\delta_1, \dots, \delta_{k-1}] = \mathbb{E}(\delta_k) = 0$.*

Finally, we need to recall the concept of a martingale and the Azuma-Hoeffding inequality for a martingale [28].

Definition IV.1. *A sequence of random variables M_1, \dots, M_n is a martingale with respect to the sequence X_1, \dots, X_n if, for all k ,*

- M_k is a function of X_1, \dots, X_k ,
- $\mathbb{E}(|M_k|) < \infty$, and
- $\mathbb{E}[M_k/X_1, \dots, X_{k-1}] = M_{k-1}$.

Lemma IV.2. (Azuma-Hoeffding inequality). *Let M_0, \dots, M_n be a martingale with respect to a sequence X_1, \dots, X_n . We assume that there exist $a_k < b_k$ such that $a_k \leq M_k - M_{k-1} \leq b_k$ for $k = 1 : n$. Then for any $A > 0$*

$$\mathbb{P}(|M_n - M_0| \geq A) \leq 2 \exp\left(-\frac{2A^2}{\sum_{k=1}^n (b_k - a_k)^2}\right).$$

In the particular case $a_k = -b_k$ we have

$$\mathbb{P}\left(|M_n - M_0| \geq \sqrt{\sum_{k=1}^n b_k^2} \sqrt{2 \ln(2/\lambda)}\right) \leq \lambda,$$

where $0 < \lambda < 1$.

B. Inner product bound

Under SR-nearness, the deterministic bound on the error of an inner product $y = a^\top b$, where $a, b \in \mathbb{R}^n$, is proportional to nu , where u is the unit roundoff of the floating-point arithmetic in use. Wilkinson [29], sec 1.33 had the intuition that the roundoff error accumulated in n operations is proportional to \sqrt{nu} rather than nu . Based on the mean independence of errors established in Lemma IV.1, Higham and Mary [24] and Ilse, Ipsen and Zhou [25] have proved this result for SR. Both works build on the mean independence property

of SR. This allows them to form a martingale, and then to apply the Azuma-Hoeffding concentration inequality. The difference between these two works is in the way they form the martingale. In [24], the martingale is built using the errors accumulated in the whole process, while [25] forms it by following step-by-step how the error accumulates. In particular, they distinguish between the multiplications and additions computed in the inner product, and carefully monitor their mean independences. In the following, we adapt this construction to Horner's algorithm.

C. Horner algorithm bound

Horner algorithm is an efficient way of evaluating polynomials. When performed in floating-point arithmetic this algorithm may suffer from catastrophic cancellations and yield a computed value with less accurate than expected.

In the following, we derive a probabilistic bound for the computed $\text{fl}(P(x))$ based on the previous methods applied for the inner product.

Model IV.1. Let $P(x) = \sum_{i=0}^n a_i x^i$, Horner's rule consists in writing this polynomial as

$$P(x) = (((a_n x + a_{n-1})x + a_{n-2})x \dots + a_1)x + a_0.$$

We define by induction the following sequence

$$\begin{aligned} r_0 &= a_n, \\ r_{2k-1} &= r_{2k-2}x, \\ r_{2k} &= r_{2k-1} + a_{n-k}, \end{aligned}$$

for all $1 \leq k \leq n$. Likewise, We define

$$\begin{aligned} \hat{r}_0 &= a_n, \\ \hat{r}_{2k-1} &= \hat{r}_{2k-2}x(1 + \delta_{2k-1}), \\ \hat{r}_{2k} &= (\hat{r}_{2k-1} + a_{n-k})(1 + \delta_{2k}), \end{aligned}$$

for all $1 \leq k \leq n$, with δ_{2k-1} and δ_{2k} represent the rounding errors from the products and the additions, respectively.

Let us define $Z_i := \hat{r}_i - r_i$ for all $1 \leq i \leq 2n$. The total forward error is $|Z_{2n}| = |\hat{r}_{2n} - r_{2n}| = |\text{fl}(P(x)) - P(x)|$. The computation of Z_{2n} introduces $\delta_1, \dots, \delta_{2n}$ such that $|\delta_k| \leq u$ for all $1 \leq k \leq 2n$. We prove by induction that the accumulation of the $(1 + \delta_k)$ errors on the $a_i x^i$ term for $0 \leq i \leq n-1$ reaches $\prod_{k=2n-2i}^{2n} (1 + \delta_k) := \varphi_i$, and for $i = n$, it reaches $\prod_{k=1}^{2n} (1 + \delta_k) := \varphi_n$. Hence

$$\begin{aligned} |\text{fl}(P(x)) - P(x)| &= \left| \sum_{i=0}^n a_i x^i (\varphi_i - 1) \right| \leq \sum_{i=0}^n |a_i x^i| |\varphi_i - 1| \\ &\leq \sum_{i=0}^n |a_i x^i| \gamma_{2n}, \end{aligned}$$

where $\gamma_{2n} = (1 + u)^{2n} - 1 = 2nu + O(u^2)$ (we recall that, $\forall k$, $|\delta_k| \leq u$). Finally

$$\frac{|\text{fl}(P(x)) - P(x)|}{|P(x)|} \leq \text{cond}_1(P, x) \gamma_{2n}, \quad (3)$$

where $\text{cond}_1(P, x) := \frac{\sum_{i=1}^n |a_i x^i|}{|P(x)|}$ is the condition number of the evaluation of P in x using the 1-norm. The deterministic bound is proportional to nu . In the following, we prove a probabilistic bound in $O(\sqrt{nu})$.

The partial sums forward errors satisfy

$$\begin{aligned} Z_{2k-1} &= \hat{r}_{2k-1} - r_{2k-1} \\ &= \hat{r}_{2k-2}x(1 + \delta_{2k-1}) - r_{2k-2}x \\ &= xZ_{2k-2} + \hat{r}_{2k-2}x\delta_{2k-1}, \\ Z_{2k} &= \hat{r}_{2k} - r_{2k} \\ &= (\hat{r}_{2k-1} + a_{n-k})(1 + \delta_{2k}) - r_{2k-1} - a_{n-k} \\ &= \hat{r}_{2k-1} + (\hat{r}_{2k-1} + a_{n-k})\delta_{2k} - r_{2k-1} \\ &= Z_{2k-1} + (\hat{r}_{2k-1} + a_{n-k})\delta_{2k}, \end{aligned}$$

for all $1 \leq k \leq n$. The sequence Z_1, \dots, Z_{2n} does not form a martingale with respect to $\delta_1, \dots, \delta_{2n}$ because, due to the multiplication in odd steps, $E[Z_{2k-1} | \delta_1, \dots, \delta_{2k-2}] = xZ_{2k-2}$. In order to form a martingale and use the Azuma-Hoeffding inequality, we define the following variable change

$$Y_i = \frac{Z_i}{x^{\lfloor (i+1)/2 \rfloor}},$$

where $\lfloor (i+1)/2 \rfloor$ is the integer part of $(i+1)/2$, we thus have

$$Y_{2k-1} = Y_{2k-2} + \frac{1}{x^{k-1}} \hat{r}_{2k-2} \delta_{2k-1}, \quad (4)$$

$$Y_{2k} = Y_{2k-1} + \frac{1}{x^k} (\hat{r}_{2k-1} + a_{n-k}) \delta_{2k}, \quad (5)$$

for all $1 \leq k \leq n$.

Theorem IV.1. The sequence of random variables Y_1, \dots, Y_{2n} is a martingale with respect to $\delta_1, \dots, \delta_{2n}$.

Proof. We check that the three conditions of definition IV.1 are satisfied. Throughout the proof, we note the set $\mathbb{F}_k = \{\delta_1, \dots, \delta_k\}$.

- The recursion in Model IV.1 shows that Y_i is a function of $\delta_1, \dots, \delta_i$ for all $1 \leq i \leq 2n$.
- $\mathbb{E}(|Y_i|)$ is finite because x and a_k are finite for all $n-i \leq k \leq n$ and $|\delta_j| \leq u$ for all $1 \leq j \leq i$.
- We prove that $\mathbb{E}[Y_i | \mathbb{F}_{i-1}] = Y_{i-1}$ by distinguishing the even and odd case. Firstly, using the mean independence of $\delta_1, \dots, \delta_{2k-1}$ and equation (4) we obtain

$$\begin{aligned} \mathbb{E}[Y_{2k-1} | \mathbb{F}_{2k-2}] &= \mathbb{E}[Y_{2k-2} | \mathbb{F}_{2k-2}] \\ &\quad + \mathbb{E}\left[\frac{1}{x^{k-1}} \hat{r}_{2k-2} \delta_{2k-1} | \mathbb{F}_{2k-2}\right] \\ &= Y_{2k-2} + \frac{1}{x^{k-1}} \hat{r}_{2k-2} \mathbb{E}[\delta_{2k-1} | \mathbb{F}_{2k-2}] \\ &= Y_{2k-2}. \end{aligned}$$

Secondly, using the mean independence of $\delta_1, \dots, \delta_{2k}$ and

equation (5) we obtain

$$\begin{aligned}
\mathbb{E}[Y_{2k}/\mathbb{F}_{2k-1}] &= \mathbb{E}[Y_{2k-1}/\mathbb{F}_{2k-1}] \\
&\quad + \mathbb{E}\left[\frac{1}{x^k}(\widehat{r}_{2k-1} + a_{n-k})\delta_{2k}/\mathbb{F}_{2k-1}\right] \\
&= Y_{2k-1} \\
&\quad + \frac{1}{x^k}(\widehat{r}_{2k-1} + a_{n-k})\mathbb{E}[\delta_{2k}/\mathbb{F}_{2k-1}] \\
&= Y_{2k-1}.
\end{aligned}$$

■

Lemma IV.3. *The above martingale Y_1, \dots, Y_{2n} satisfies*

$$|Y_i - Y_{i-1}| \leq C_i u, \quad 1 \leq i \leq 2n,$$

where

$$C_{2k-1} = |a_n|(1+u)^{2k-2} + \sum_{j=1}^{k-1} |a_{n-j}||x|^{-j}(1+u)^{2(k-j)-1},$$

$$C_{2k} = |a_n|(1+u)^{2k-1} + \sum_{j=1}^k |a_{n-j}||x|^{-j}(1+u)^{2(k-j)},$$

for all $1 \leq k \leq n$.

Proof.

Note that $Y_0 = 0$, then $|Y_1 - Y_0| = |Y_1| = |a_n|$ and the equality holds for C_1 . Using equation (4)

$$|Y_{2k-1} - Y_{2k-2}| \leq \frac{1}{|x|^{k-1}} |\widehat{r}_{2k-2}| u.$$

However

$$\begin{aligned}
|\widehat{r}_{2k-2}| &\leq |\widehat{r}_{2k-3}|(1+u) + |a_{n-k+1}|(1+u) \\
&\leq |\widehat{r}_{2k-4}| |x|(1+u)^2 + |a_{n-k+1}|(1+u),
\end{aligned}$$

by induction we obtain

$$|\widehat{r}_{2k-2}| \leq |a_n||x|^k(1+u)^{2k-2} + \sum_{j=1}^{k-1} |a_{n-j}||x|^{k-j}(1+u)^{2(k-j)-1}.$$

This completes the proof for C_{2k-1} for all $1 \leq k \leq n$. A similar approach can be applied to proving the same result for C_{2k} for all $1 \leq k \leq n$. ■

We now have all the tools to state and then demonstrate the main result of this section:

Theorem IV.2. *Under SR-nearness, for all $0 < \lambda < 1$ and with probability at least $1 - \lambda$*

$$\frac{|\text{fl}(P(x)) - P(x)|}{|P(x)|} \leq \text{cond}_1(P, x) \sqrt{u\gamma_{4n}} \sqrt{\ln(2/\lambda)}, \quad (6)$$

where $\text{cond}_1(P, x) = \frac{\sum_{i=1}^n |a_i x^i|}{|P(x)|}$ is the condition number of the polynomial evaluation and $\gamma_{4n} = (1+u)^{4n} - 1$.

Proof. Recall that $|\widehat{r}_{2n} - r_{2n}| = |Z_{2n}| = |x^n| |Y_{2n}|$. Therefore, Y_1, \dots, Y_{2n} is a martingale with respect to $\delta_1, \dots, \delta_{2n}$ and Lemma IV.3 implies $|Y_i - Y_{i-1}| \leq C_i u$ for all $1 \leq i \leq 2n$. Using the

Azuma-Hoeffding inequality yields

$$\mathbb{P}\left(|Y_{2n}| \leq u \sqrt{\sum_{i=1}^{2n} C_i^2} \sqrt{2 \ln(2/\lambda)}\right) \geq 1 - \lambda,$$

it follows that

$$|Z_{2n}| \leq u \sqrt{\sum_{i=1}^{2n} (|x|^n C_i)^2} \sqrt{2 \ln(2/\lambda)},$$

with probability at least $1 - \lambda$, where

$$\begin{aligned}
|x|^n C_{2k} &= |a_n| |x|^n (1+u)^{2k-1} + \sum_{j=1}^k |a_{n-j} x^{n-j}| (1+u)^{2(k-j)} \\
&\leq (1+u)^{2k-1} \sum_{j=0}^k |a_{n-j} x^{n-j}| \\
&\leq (1+u)^{2k-1} \sum_{j=0}^n |a_j x^j|,
\end{aligned}$$

for all $1 \leq k \leq n$. Hence,

$$(|x|^n C_{2k})^2 \leq (1+u)^{2(2k-1)} \left(\sum_{j=0}^n |a_j x^j|\right)^2.$$

In a similar way

$$(|x|^n C_{2k-1})^2 \leq (1+u)^{2(2k-2)} \left(\sum_{j=0}^n |a_j x^j|\right)^2.$$

Thus

$$\begin{aligned}
\sum_{i=1}^{2n} (|x|^n C_i)^2 &\leq \left(\sum_{j=0}^n |a_j x^j|\right)^2 \sum_{i=0}^{2n-1} ((1+u)^2)^i \\
&= \left(\sum_{j=0}^n |a_j x^j|\right)^2 \frac{((1+u)^2)^{2n} - 1}{(1+u)^2 - 1} \\
&= \left(\sum_{j=0}^n |a_j x^j|\right)^2 \frac{\gamma_{4n}}{u^2 + 2u}.
\end{aligned}$$

As a result

$$|\text{fl}(P(x)) - P(x)| = |Z_{2n}| \leq \sum_{j=0}^n |a_j x^j| \sqrt{\frac{u\gamma_{4n}}{2+u}} \sqrt{2 \ln(2/\lambda)},$$

with probability at least $1 - \lambda$. Finally

$$\frac{|\text{fl}(P(x)) - P(x)|}{|P(x)|} \leq \text{cond}_1(P, x) \sqrt{u\gamma_{4n}} \sqrt{\ln(2/\lambda)},$$

with probability at least $1 - \lambda$. ■

Remark IV.1. *The bounds (3) and (6) have the same condition number, but differ in another factor: γ_{2n} for (3) against $\sqrt{u\gamma_{4n}} \sqrt{\ln(2/\lambda)}$ for (6).*

For n such that $2nu < 1$, [30], Lemma 3.1 implies

$\gamma_{2n} \leq \frac{2nu}{1-2nu}$, it follows that for $4nu < 1$

$$\sqrt{u\gamma_{4n}} \leq \sqrt{\frac{4nu^2}{1-4nu}} = u\sqrt{n} \frac{2}{\sqrt{1-4nu}}.$$

For n large, Taylor's formula implies $\gamma_n = nu + O(u^2)$ and $\gamma_{2n} \approx 2nu$. This approach can't be used for $\sqrt{\gamma_{4n}}$ because it's indeterminate in 0. However,

$$\lim_{u \rightarrow 0} \frac{\sqrt{u\gamma_{4n}}}{\sqrt{nu}} = 2 \iff \sqrt{u\gamma_{4n}} \approx 2\sqrt{nu}.$$

Eventually, the probabilistic bound for the forward error of Horner's algorithm is in $O(\sqrt{nu})$.

D. Numerical experiments.

In this section, we illustrate that the probabilistic bound is tighter than the deterministic bound for SR-nearness forward error on a numerical application: the evaluation of the Chebyshev polynomial. We use Horner's method to evaluate the polynomial $P(x) = T_N(x) = \sum_{i=0}^{\lfloor \frac{N}{2} \rfloor} a_i(x^2)^i$ where T_N is the Chebyshev polynomial of degree N . Consider an even $N = 2n$. We use single-precision (binary32) for both SR-nearness and round to nearest ties to even. All SR computations are repeated 30 times with verifcarlo [15]; we plot all samples and the forward error of the average of the 30 SR instances. The following error bounds and evaluations apply:

$$\begin{aligned} \text{Probabilistic bound} &= \text{cond}_1(P, x) \sqrt{u\gamma_{4n}} \sqrt{\ln(2/\lambda)}, \\ \text{Deterministic bound} &= \text{cond}_1(P, x) \gamma_{2n}, \\ \text{SR-nearness} &= \frac{|\text{fl}(P(x)) - P(x)|}{|P(x)|}. \end{aligned}$$

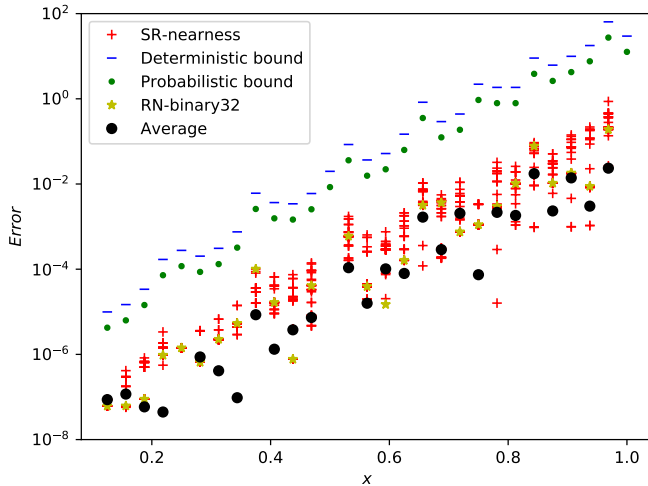


Fig. 5. Probabilistic bound with $\lambda = 0.5$ vs deterministic bound of the computed forward errors of Horner's rule for Chebyshev polynomial $T_{20}(x)$.

Chebyshev polynomial is ill-conditioned near 1: in figure 5, we evaluate $T_{20}(x)$ for $x \in [\frac{8}{64}; 1]$. As expected and due to catastrophic cancellations among the coefficients, the

condition number increases from 10^0 to 10^7 which explains the increase of numerical error from 10^{-7} to 10^{-1} . The probabilistic bound is closer to the forward error points than the deterministic bound even for a small $n = 10$. The average of SR-nearness stays below RN-binary32 for almost all points.

In Figure 6, the two previous bounds and the forward error are divided by the condition number $\text{cond}(P, x)$, and the evaluation in $x = 24/26 \approx 0.923$ is plotted for various polynomial degrees N .




Fig. 6. Forward errors/ $\text{cond}(P, x)$ of Horner's rule for Chebyshev polynomial $T_N(24/26)$.

The comparison between the two bounds is fairly visible in figure 6. By increasing N , the deterministic bound draws away from the forward error faster than the probabilistic bound and for N large, the gap becomes more interesting. SR-nearness points are between 10^{-8} and 10^{-10} , versus 10^{-7} for the probabilistic bound and 10^{-6} for the deterministic bound.

V Conclusion.

Stochastic rounding has drawn a lot of attention in various domains [14], [19], [22], [23] due to its efficiency compared to the default rounding mode. The fact that SR-nearness satisfies mean independence (a weaker property than independence) leads to an expected value that coincides with the exact value. We have shown that the bias in SR-up-or-down can significantly reduce the precision of the computation, even on simple algorithms such as rectangular integration section III, and that SR-nearness can remain unbiased and provide the full expected precision on them. We also discussed that using SR-nearness leads to having a probabilistic bound in $O(\sqrt{nu})$, compared to the $O(nu)$ deterministic bound for the inner product forward error. We have shown this property for Horner's method using Azuma–Hoeffding inequality and martingale properties. As opposed to the study made for the inner product, the issue of this algorithm is that the martingale does not appear explicitly; nevertheless, a change of variable shows that it is present, allowing the use of concentration inequalities. As future work we will investigate more complex algorithms with non-explicit martingales.

Acknowledgment

This research was supported by the French National Agency for Research (ANR) via the InterFLOP project (No. ANR-20-CE46-0009).

REFERENCES

- [1] D. S. Parker, *Monte Carlo Arithmetic: exploiting randomness in floating-point arithmetic*, University of California (Los Angeles). Computer Science Department, 1997.
- [2] J. von Neumann, H. H. Goldstine, “Numerical inverting of matrices of high order”, *Bulletin of the American Mathematical Society*, vol. 53, pp. 1021–1099, 1947.
- [3] “Graphcore Limited. 2021a IPU Programmer’s Guide. Version 2.0.0.”, .
- [4] “Graphcore Limited. 2021b Targeting the IPU from TensorFlow 1. Version 2.0.0-dc2.”, .
- [5] “Graphcore Limited. 2021c AI-Float™- Mixed Precision Arithmetic for AI: A Hardware Perspective. Version latest: Aug 25, 2021.”, .
- [6] M. Davies, N. Srinivasa, T.-H. Lin, G. China, Y. Cao, S. H. Choday, G. Dimou, P. Joshi, N. Imam, S. Jain, *et al.*, “Loihi: A neuromorphic manycore processor with on-chip learning”, *Ieee Micro*, vol. 38, no. 1, pp. 82–99, 2018.
- [7] “Loh GH. 2019 Stochastic rounding logic. Patent Status: Active.”, .
- [8] J. Alben, P. Micikevicius, H. Wu, M. Siu, N. Corporation, “Stochastic rounding of numerical values”, *Patent Status: Active*, 2019.
- [9] J. D. Bradbury, S. R. Carlough, B. R. Prasky, E. M. Schwarz, “Reproducible stochastic rounding for out of order processors”, US Patent 10,083,008, Sep. 25 2018.
- [10] J. D. Bradbury, S. R. Carlough, B. R. Prasky, E. M. Schwarz, “Stochastic rounding floating-point multiply instruction using entropy from a register”, US Patent 10,445,066, Oct. 15 2019.
- [11] G. G. Henry, D. R. Reed, “Processor with memory array operable as either cache memory or neural network unit memory”, US Patent 10,664,751, May 26 2020.
- [12] O. A. Kanter, I. Bar, “Apparatus and methods for hardware-efficient unbiased rounding”, US Patent 8,972,472, Mar. 3 2015.
- [13] S. Lifsches, “In-memory stochastic rounder”, US Patent 10,803,141, Oct. 13 2020.
- [14] M. Croci, M. Fasi, N. J. Higham, T. Mary, M. Mikaitis, “Stochastic Rounding: Implementation, Error Analysis, and Applications”, *Royal Society Open Science*, 2021.
- [15] C. Denis, P. de Oliveira Castro, E. Petit, “Verificarlo: Checking Floating Point Accuracy through Monte Carlo Arithmetic”, in *23rd IEEE Symposium on Computer Arithmetic*, pp. 55–62, 2016, doi:10.1109/ARITH.2016.31, URL: <http://dx.doi.org/10.1109/ARITH.2016.31>.
- [16] F. Jézéquel, J.-M. Chesneaux, “CADNA: a library for estimating round-off error propagation”, *Computer Physics Communications*, vol. 178, no. 12, pp. 933–955, 2008.
- [17] F. Févotte, B. Lathuiliere, “VERROU: a CESTAC evaluation without recompilation”, *SCAN 2016*, p. 47, 2016.
- [18] “IEEE Standard for Floating-Point Arithmetic”, *IEEE Std 754-2019 (Revision of IEEE 754-2008)*, pp. 1–84, 2019, doi:10.1109/IEEESTD.2019.8766229.
- [19] S. Gupta, A. Agrawal, K. Gopalakrishnan, P. Narayanan, “Deep Learning with Limited Numerical Precision”, , 02 2015.
- [20] M. Höhfeld, S. E. Fahlman, “Probabilistic rounding in neural network learning with limited precision”, *Neurocomputing*, vol. 4, no. 6, pp. 291–299, 1992, doi: [https://doi.org/10.1016/0925-2312\(92\)90014-G](https://doi.org/10.1016/0925-2312(92)90014-G).
- [21] M. Hoehfeld, S. Fahlman, “Learning with limited numerical precision using the cascade-correlation algorithm”, *IEEE Transactions on Neural Networks*, vol. 3, no. 4, pp. 602–611, 1992, doi:10.1109/72.143374.
- [22] M. Hopkins, M. Mikaitis, D. R. Lester, S. Furber, “Stochastic rounding and reduced-precision fixed-point arithmetic for solving neural ordinary differential equations”, *Philosophical Transactions of the Royal Society A*, vol. 378, no. 2166, p. 20190052, 2020.
- [23] M. Fasi, M. Mikaitis, “Algorithms for Stochastically Rounded Elementary Arithmetic Operations in IEEE 754 Floating-Point Arithmetic”, *IEEE Transactions on Emerging Topics in Computing*, vol. 9, no. 3, pp. 1451–1466, 2021, doi:10.1109/TETC.2021.3069165.
- [24] N. J. Higham, T. Mary, “A New Approach to Probabilistic Rounding Error Analysis”, *SIAM Journal on Scientific Computing*, vol. 41, no. 5, pp. A2815–A2835, 2019, doi:10.1137/18M1226312.
- [25] I. C. F. Ipsen, H. Zhou, “Probabilistic Error Analysis for Inner Products”, *SIAM Journal on Matrix Analysis and Applications*, vol. 41, no. 4, pp. 1726–1741, 2020, doi:10.1137/19M1270434.
- [26] N. J. H. M. P. Connolly, T. Mary, “Stochastic rounding and its probabilistic backward error analysis”, *SIAM Journal on Scientific Computing*, vol. 43, No. 1, pp. A566–A585, 2021, doi:<https://doi.org/10.1137/20M1334796>.
- [27] F. Févotte, B. Lathuilière, P. de Oliveira Castro, “Etudier la qualité numérique d’un code avec Verrou”, <https://github.com/edf-hpc/verrou/releases/download/tutorials/tp-verrou.tgz>, [Online; accessed 21-April-2021], 2018.
- [28] M. Mitzenmacher, E. Upfal, *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*, Cambridge University Press, 2005, doi: 10.1017/CBO9780511813603.
- [29] J. H. Wilkinson, *Rounding errors in algebraic processes*, Courier Corporation, 1994.
- [30] N. J. Higham, *Accuracy and stability of numerical algorithms*, SIAM, 2002.