

On Stochastic Rounding with Few Random Bits

Andrew Fitzgibbon
Graphcore
Cambridge, UK
Email: awf@ieee.org

Stephen Felix
Graphcore
Bristol, UK
Email: Stephen.Felix@graphcore.ai

Abstract—Large-scale numerical computations make increasing use of low-precision (LP) floating point formats and mixed precision arithmetic, which can be enhanced by the technique of *stochastic rounding* (SR), that is, rounding an intermediate high-precision value up or down randomly as a function of the value’s distance to the two rounding candidates. Stochastic rounding requires, in addition to the high-precision input value, a source of random bits. As the provision of high-quality random bits is an additional computational cost, it is of interest to require as few bits as possible while maintaining the desirable properties of SR in a given computation, or computational domain. This paper examines a number of possible implementations of few-bit stochastic rounding (FBSR), and shows how several natural implementations can introduce sometimes significant bias into the rounding process, which are not present in the case of infinite-bit, infinite-precision examinations of these implementations. The paper explores the impact of these biases in machine learning examples, and hence opens another class of configuration parameters of which practitioners should be aware when developing or adopting low-precision floating point.

I. INTRODUCTION

Large-scale numerical computations make increasing use of low-precision (LP) floating point formats, both for storage and computation. As compared to 32 and 64 bit formats, narrow storage formats (16, 8, and even fewer bits per element) allow larger arrays to be stored in high-bandwidth memory, close to accelerated arithmetic units; and allow more elements to be transported per second across communications links, whether within or between compute nodes. Narrow-format arithmetic can be implemented more efficiently, both in hardware and software, increasing the speed of computation. Of course, narrow formats imply decreased accuracy, requiring computations to be re-engineered in order to compensate. In areas such as machine learning, the training of large models has demonstrated that this re-engineering can yield significant improvement in the scale and speed at which a computational goal can be accomplished, while yielding models of equivalent or greater accuracy at reduced cost. In some instances, the low-precision models are superior in the sense that they can be deployed on lower-power devices after training.

This re-engineering, however, can prove difficult or impossible at lower format widths. There is, after all, a considerable loss of precision and dynamic range in LP formats, giving rise to the following challenges: careful dynamic range management, judicious deployment of mixed low and standard precision, a re-exploration of “hyper-parameter” settings (such as step sizes, number of iterations, etc). These challenges can

be mitigated by the use of stochastic rounding (SR), as has been shown in many application domains [1]–[5].

The focus of this paper is on implementations of SR which use small numbers of random bits, which we call *few-bit stochastic rounding* (FBSR). In particular, we consider the case where the number of random bits used is smaller than the difference in precision between the values to be rounded and the target precision. We illustrate (see Figure 1) that natural implementations may have sometimes significant biases, confounding any experimental investigations into FBSR. Mitigations for these biases are presented, and empirical results are presented to argue for the effectiveness of these mitigations, and hence for the utility of FBSR.

A. Related work

Croci *et al.* [6] provide an excellent survey of the theory and applications of SR in a variety of fields including machine learning and quantum computing. In almost all current implementations, whether hardware or software, the number of random bits used equals the number of trailing significand bits, so are not in the few-bit regime described in this paper. For future implementations, however, it is expected that our findings will be useful in making design decisions.

This paper’s primary contribution is a description of, and remediation of, the bias that some implementations of FBSR may introduce. This bias has been analysed theoretically in recent work by El Arar *et al.* [7], [8], but no mitigation is proposed. The current paper proposes two alternative mitigations, and provides a more application-focused derivation of the bias computation which may prove useful to some readers. Xia *et al.* [9] also consider bias in SR, but with the goal of introducing bias in order to reduce variance. Given current usage of SR in large-scale machine learning applications where millions or billions of roundings occur in one training iteration, it is suggested that bias is of greater concern than variance in current applications.

The paper does not discuss the quality of the supplied bits, merely their quantity. Some recent work [10] suggests that lower quality sources (e.g. nearby truncations) may serve effectively in practice. It is thus conceivable that more bits from a lower-quality source might allow a more efficient implementation than a bias-free implementation with fewer high-quality bits, although the extent of the bias is such that this appears unlikely, and the confluences involved in using

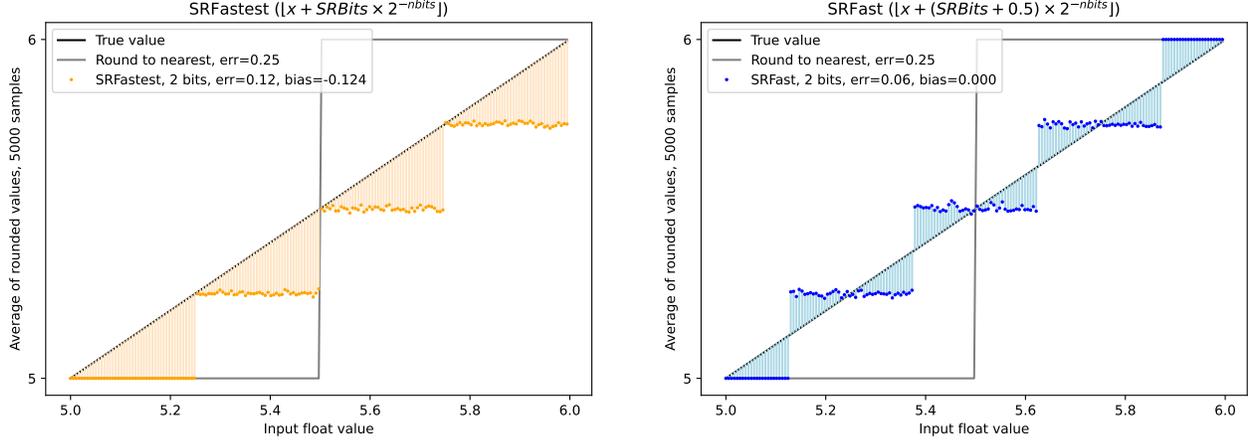


Fig. 1. Bias in stochastic rounding with finite randomness. For each X value, 5000 samples are rounded with 2 bits of randomness, and the mean of the rounded values is plotted. (Left, SRFF): Mean rounded values are below the line $Y = X$, indicating bias. The computed bias of $-0.124 \approx -2^{-3}$, as predicted by the calculation of §III-C. (Right, SRF): Mean rounded values are symmetric around $Y = X$, with bias ≈ 0 , following §III-D.

lower-quality sources might render theoretical analysis difficult or impossible.

As noted by several authors [1], [2], low-precision formats generally make use of a *scaled tensor* implementation, where blocks of low-precision values are accompanied by a scale factor, which may be a power of two (i.e. a pure-exponent format, or another floating point format). This somewhat expands the applicability of the approach described in this paper, as even if intermediate values are computed in mid-precision formats such as BFloat16, the scaled values may have larger effective precision, making the corrections in this paper more important.

B. Contributions and limitations

For the reader’s convenience, the contributions and limitations of this work are outlined:

- A simple but general mathematical framework for the description and analysis of FBSR, with a particular focus on practical implementation;
- A definition of several FBSR methods within this framework;
- Analysis of the bias properties of these methods, revealing biases that are not, to our knowledge, described in the existing literature;
- Experimental investigations of the practical impact of the discovered biases on representative computations from machine learning.

Primary limitations are:

- The experimental investigations are at a scale which is sufficient to confirm that the biases can have an effect on real-world use cases, but these conclusions may not extend to larger-scale deployments.
- The costs, in terms of area and power, of any hardware implementation are not discussed. This is at least in part because so many area/power decision tradeoffs exist even

when the parameters are fixed. It is hoped that the paper’s investigations are of utility in making such tradeoffs in any future hardware designs.

II. BACKGROUND AND NOTATION

We consider binary floating point numbers represented by the triple $(s, e, t) \in \mathbb{N}^3$, representing sign, (biased) exponent and the fractional part (trailing bits) of the significand. A floating point *format* is parameterized by width in bits K , precision in bits P , and exponent bias B . We shall also make use of M , the largest finite value in the format’s value set (from which it assumed symmetrically that $-M$ is the smallest finite value). The real value encoded by (s, e, t) in format $(K, P, B; M)$ is given by the function

$$F(s, e, t; P, B)$$

defined as

$$(-1)^s \times \begin{cases} (0 + t \times 2^{1-P}) \times 2^{1-B} & \text{if } e = 0 \\ (1 + t \times 2^{1-P}) \times 2^{e-B} & \text{if } e > 0 \end{cases} \quad (1)$$

where the $e = 0$ case encodes subnormal numbers. In this definition, the precision P is one more than the number of mantissa bits that need to be encoded in a packed binary representation of the number. Formats may additionally encode infinities, negative zero, and not-a-number (NaN) values, which shall largely not feature in the discussions of this paper (in implementations, it is assumed that values outside the representable finite float range are rounded deterministically as in IEEE-754 round to nearest).

A. Rounding

A strictly positive real value $\tilde{x} > 0$ is rounded as follows. We first compute the real-valued significand $\tilde{\sigma} \in [2^{P-1}, 2^P)$, which is then rounded to integer $\sigma \in \{2^{P-1}, \dots, 2^P\}$, of which the trailing bits are encoded (incrementing the exponent if $\sigma =$

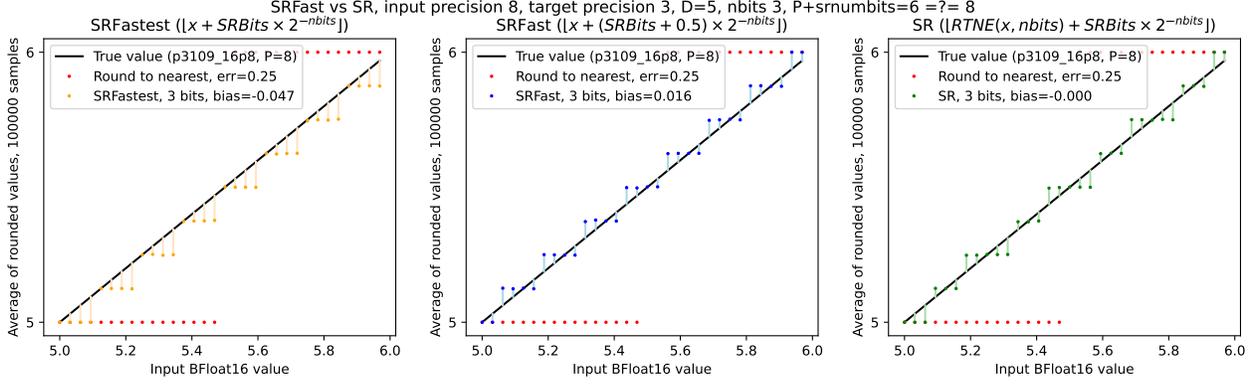


Fig. 2. Bias in stochastic rounding with finite randomness, finite-precision inputs. For each X value in BFloat16, 100,000 samples are rounded with 2 bits of randomness, and the mean of the rounded values is plotted. (Left, SRFF): Mean rounded values are below the line $Y = X$, indicating bias. The computed bias is ≈ -0.046875 , as predicted by the calculation of §III-E and §A-C. (Middle, SRF): Mean rounded values are not symmetric around $Y = X$, with bias ≈ 0.015625 , following §A-D. (Right, SRC): Mean rounded values are symmetric around $Y = X$, with empirical bias ≈ 0.000 , following the correction of §III-E.

2^P). The initial exponent is computed as follows, allowing for subnormals:

$$\tilde{e} = \max(\lfloor \log_2(\tilde{x}) \rfloor, 1 - B) \quad (2)$$

from which the (scaled) significand is

$$\tilde{\sigma} = \tilde{x} \times 2^{-e} \times 2^{P-1} \quad (3)$$

Then, rounding to integer with ties away from zero is represented by

$$\begin{aligned} \tilde{\delta} &= \tilde{\sigma} - \lfloor \tilde{\sigma} \rfloor \\ \sigma &= \begin{cases} \lfloor \tilde{\sigma} \rfloor & \text{if } \tilde{\delta} < 0.5 \\ \lfloor \tilde{\sigma} \rfloor + 1 & \text{if } \tilde{\delta} \geq 0.5 \end{cases} \end{aligned}$$

which may be written compactly in terms of a predicate

$$R_{TA}(\tilde{\delta}) = \tilde{\delta} \geq 0.5 \quad (4)$$

whence we can replace the above with

$$\sigma = \lfloor \tilde{\sigma} \rfloor + \mathbb{1}[R_{\text{RoundMode}}(\tilde{\sigma} - \lfloor \tilde{\sigma} \rfloor)] \quad (5)$$

For ties to nearest even or odd (TNE, TNO), we may define an alternative $R_{\text{RoundMode}}$, which has access to $\tilde{\sigma}$ in order to determine the direction in which to resolve ties. For example,

$$R_{TNE}(\tilde{\sigma}, \tilde{\delta}) = \tilde{\delta} > 0.5 \text{ or } (\tilde{\delta} = 0.5 \text{ and } \text{IsEven}(\lfloor \tilde{\sigma} \rfloor + 1))$$

where $\text{IsEven}(i)$ is true iff integer i is even. In practice, implementations may use comparison as indicated above, or an alternative implementation using addition, analogous to $\tilde{\delta} + 0.5 \geq 1.0$. The conclusions of this paper apply to both implementations, with possible inversion of the direction of any bias.

B. Stochastic Rounding

With the above notation, stochastic rounding is readily defined

$$R_{SR}(\tilde{\delta}) = \tilde{\delta} + \tilde{n} \geq 1$$

where the “noise” \tilde{n} is a random variable drawn from a uniform distribution between 0 and 1. As mentioned in the introduction, the quality of the random number generator is orthogonal to the concerns of this paper. In practical implementations of SR, the random numbers will be supplied as a finite sequence of random bits, denoted by an integer n such that

$$0 \leq n < 2^N,$$

where N is the number of bits in n . Random bits are passed to the R function, so that a natural analogue of the infinite precision definition is given by

$$R_{SRFF}(\tilde{\delta}, n) = \tilde{\delta} + n \times 2^{-N} \geq 1$$

As will be explored in the following, this definition proves to be biased, and variations which mitigate this bias shall be defined. One which is of common practical utility is the form we shall name “SRF”¹:

$$R_{SRF}(\tilde{\delta}, n) = \tilde{\delta} + \left(n + \frac{1}{2}\right) \times 2^{-N} \geq 1$$

III. BIAS

A property that one might require of any rounding scheme is that it is *unbiased*, or that the expected difference between inputs and outputs is zero. Written loosely, this might be

$$\mathbb{E}[x - \text{round}(x)] = 0.$$

As written, this definition is somewhat vacuous: the expectation can only be over values of x , and hence depends on

¹The naming scheme encodes speed (and concomitant reduced accuracy): corrected (SRC), fast (SERF), faster (SRFF).

a choice of prior probability distribution $p(x)$. Writing more explicitly, the requirement is

$$\mathbb{E}_{x \sim p(x)} [x - \text{round}(x)] = 0.$$

This prior should not, for example, simply be the set of real values which map to the format's finite range, e.g. $p(x)$ being the uniform distribution on $[-M, M]$. Under such a distribution, always rounding toward zero is unbiased, as the negative values will cancel the positives.

More meaningfully, the expectation may be over a family of *test distributions* \mathcal{P} , and we require unbiased rounding for all members of the family:

$$\forall p \in \mathcal{P} : \mathbb{E}_{x \sim p(x)} [x - \text{round}(x)] = 0.$$

In this paper, we choose to require that values between each pair of floats are unbiased on average, i.e., if $\mathcal{U}(l, h)$ is the uniform distribution between l and h , and \mathcal{F} is the set of finite floats in a format, then a family of test distributions is

$$\mathcal{P} = \{\mathcal{U}(f, \text{succ}(f)), f \in \mathcal{F} - \{M\}\} \quad (6)$$

hence we require that

$$\forall f \in \mathcal{F}, f < M : \int_f^{\text{succ}(f)} (x - \text{round}(x)) dx = 0$$

In terms of a rounding predicate R , as in (5), unbiased rounding then means

$$\int_0^1 \mathbb{1}[R(\tilde{\delta})] - \tilde{\delta} d\tilde{\delta} = 0$$

or, equivalently

$$\int_0^1 \mathbb{1}[R(\tilde{\delta})] d\tilde{\delta} - \frac{1}{2} = 0$$

A. Stochastic rounding

For SR, we must additionally compute the expectation over the random variable \tilde{n} :

$$\forall p \in \mathcal{P} : \mathbb{E}_{\tilde{n}} [\mathbb{E}_{x \sim p(x)} [x - \text{round}(x)]] = 0.$$

In the infinite-precision case, the family \mathcal{P} can be delta functions at all $x \in [-M, M]$, with zero bias at all points

$$\forall x \in [-M, M] : \int_0^1 [x - \text{round}(x, \tilde{n})] d\tilde{n} = 0$$

Equivalently, this is the requirement that

$$\forall \tilde{\delta} \in [0, 1] : \int_0^1 \mathbb{1}[\tilde{\delta} + \tilde{n} \geq 1] d\tilde{n} = \tilde{\delta}$$

which is straightforwardly true as for all $\tilde{\delta} \in [0, 1]$ we have

$$\begin{aligned} \int_0^1 \mathbb{1}[\tilde{\delta} + \tilde{n} \geq 1] d\tilde{n} &= \\ \int_0^1 \mathbb{1}[\tilde{n} \geq 1 - \tilde{\delta}] d\tilde{n} &= \\ \int_{1-\tilde{\delta}}^1 1 d\tilde{n} &= \\ 1 - (1 - \tilde{\delta}) &= \tilde{\delta} \end{aligned}$$

B. Stochastic rounding, finite randomness

In the finite-randomness case, the test distributions must revert to the family of (6), i.e. uniform between each pair of floats, so the unbiasedness condition is

$$\mathbb{E}_n \left[\int_0^1 \mathbb{1}[R(x, n)] dx \right] - \frac{1}{2} = 0$$

which is the discrete sum

$$2^{-N} \sum_{n=0}^{2^N-1} \int_0^1 \mathbb{1}[R(x, n)] dx - \frac{1}{2} = 0$$

C. Bias of SRFF

Let us now compute the bias of the SRFF mode, for a given number of bits N . Appendix A-A contains the complete calculation, key steps of which are as follows:

$$\begin{aligned} 2^{-N} \sum_{n=0}^{2^N-1} \left(\int_0^1 \mathbb{1}[R_{\text{SRFF}}(x, n)] dx - \frac{1}{2} \right) &= \\ 2^{-N} \sum_{n=0}^{2^N-1} \int_0^1 \mathbb{1}[x + n \times 2^{-N} \geq 1] dx - \frac{1}{2} &= \\ 2^{-N} \sum_{n=0}^{2^N-1} \int_{1-n \times 2^{-N}}^1 dx - \frac{1}{2} &= \\ -2^{-(N+1)} &\neq 0 \end{aligned}$$

Hence SRFF is biased, with bias decreasing with increased number of random bits. As shown later in experiments, even with 4 bits of randomness, the bias can significantly affect deep learning training.

D. Unbiasedness of SRF

A similar calculation for SRF mode becomes (again see appendix for detail)

$$\begin{aligned} 2^{-N} \sum_{n=0}^{2^N-1} \int_0^1 \mathbb{1}[R_{\text{SRF}}(x, n)] dx - \frac{1}{2} &= \\ 2^{-N} \sum_{n=0}^{2^N-1} \int_0^1 \mathbb{1}\left[x > \left(n + \frac{1}{2}\right) \times 2^{-N}\right] dx - \frac{1}{2} &= 0 \end{aligned}$$

showing that SRF is unbiased, in that the expectation computed over all real values between two floats is zero.

E. Bias, finite precision inputs

The above calculations assume that all real numbers are equally likely inputs to rounding. In real-world systems, the inputs will be available at some precision $Q > P$, which we shall write in terms of the number of excess bits D , i.e. $Q = P + D$. When D is large, for example, when converting from 32-bit or higher formats, we might expect the calculations above for infinite-precision inputs to pertain. However it is not unusual for the source format to be relatively close in precision to the target. For example, BFloat16 has a precision of 8 and might be rounded to an E5M2 format with a precision of 3, in which case $D = 5$.

Algorithm 1 Quantization-aware training of an 8-bit (F8) model using high-precision (F16) gradients.

```

 $w_8 : F8 \leftarrow$  initializers
while not done do
   $w_{16} : F16 \leftarrow w_8$ 
   $g_{16} : F16 \leftarrow \nabla \text{loss}(w_{16})$ 
   $u_{16} : F16 \leftarrow$  AdamW update in F16
   $w_8 \leftarrow \text{RoundTo8}(w_{16} + u_{16})$   $\triangleright$  SR applied here
end while

```

In such cases, the integral also becomes a discrete sum over the incoming values between 0 and 1. For simplicity, we consider initially only the case where the incoming range does not wrap binades in our target interval. Hence there are 2^D values (excluding 1). Then the bias of SRFF becomes

$$\text{bias}_{\text{SRFF},D} = 2^{-D} \sum_{i=0}^{2^D-1} \left(2^{-N} \sum_{n=0}^{2^N-1} \mathbb{1}[\text{R}_{\text{SRFF}}(2^{-D}i, n)] - 2^{-D}i \right) \quad (7)$$

As derived (with more verbose annotations) in the appendix, this calculation yields the bound (tight for $N \leq D$)

$$\text{bias}_{\text{SRFF},D} \leq \frac{2^{-D} - 2^{-N}}{2}$$

from which it is clear that the bias is zero for $N = D$.

We note that $N = D$ is the case commonly implemented in today’s hardware, however, as noted in the introduction, the research literature is increasingly exploring $N < D$, in which case correction of the bias will be required.

It is not sufficient to use SRF, which also has a bias, as shown in §A-D bounded by (and again, with the bound being tight for $N \leq D$)

$$\text{bias}_{\text{SRF},D} \leq 2^{-(D+1)}$$

F. Bias-corrected few-bit SR (SRC)

There is, however, a simple correction: first deterministically round the incoming precision- Q values to precision $Q' = P + N$ then apply SRFF. This reduces the problem to the $N = D'$ case, which is unbiased for SRFF as described above. More formally:

$$\text{R}_{\text{SRC}}(\tilde{\delta}, n) = \text{R}_{\text{SRFF}}(\text{Round}(\tilde{\delta} \times 2^N) \times 2^{-N}, n) \quad (8)$$

where Round is any unbiased rounding scheme, e.g. round to nearest, with ties to even or odd. Despite its simplicity, we are not aware of this correction having been described in any existing discussion of FBSR. As shall be shown in the sequel, empirical results support the theoretical calculations above.

IV. EXPERIMENTS

A current interest in the field of generative language modelling is in the training of models with low-precision weights, such that inference can be performed on low-power devices. While low-precision weights can be obtained by post-training

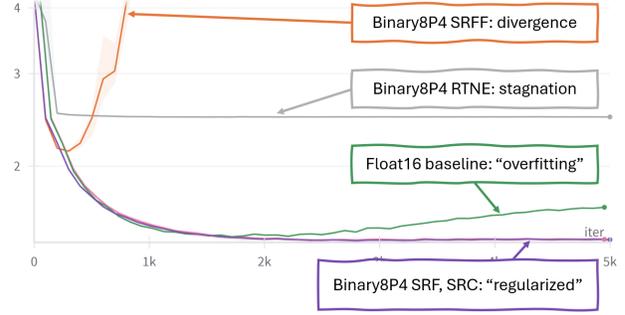


Fig. 3. Experiments on language model training (nanoGPT, small model). Binary16 baseline: achieves minimum validation loss at about 1900 iterations, after which training loss (not shown) reduces, but validation loss increases. Binary8P4 RTNE: initial loss reduction followed by “stagnation”; SRFF: initial loss reduction followed by “divergence”; SRF, SRC: convergence to stable values. Although SRF and SRC converge to better validation errors than Binary16, this regularization behaviour is expected to apply only to such small-scale experiments, and should not be taken as an indication of any superiority of Binary8+SR over Binary16.

quantization (PTQ), that is, simply casting learned weights to the low-precision format, it is increasingly understood that performing quantization-aware training (QAT) can yield more accurate models [11]. To this end, we have performed a simple simulation of QAT whereby weights are rounded to the target precision at initialization and after every gradient update (or equivalently, before every gradient calculation). A more precise definition in pseudocode is given at Algorithm 1.

Algorithms were implemented on the nanoGPT codebase [12], with stochastic rounding implemented via the gfloat library [13]. The implementation did not make use of SR implemented in hardware, in order to be able to freely vary the SR implementation.

The default model was trained on the Shakespeare dataset. At this small scale, the model is known to overfit training data, and hence validation loss is the figure of merit followed. Training baselines were run with gradient updates computed in BFloat16, and Binary16, yielding essentially identical validation curves. The target format was set to IEEE WG P3109 Binary8P4 [14] format whose precision corresponds to that of OCP E4M3 format. For these experiments, the number of SR bits used was 3. In principle, as gradient updates are in BFloat16, with a precision of 8, and the target has a precision of 4, the ideal number of SR bits to use is 4.

As illustrated in Figure 3, rounding to nearest with ties to even (RTNE) succeeds in initially reducing the validation loss, but then “stagnates”, where further gradient updates make no change to the weights and hence to the loss. Rounding to Binary8P4 under the (biased) SRFF implementation initially decreases the loss, but soon “diverges”, with the loss rapidly increasing. Both bias-corrected implementations (SRF, SRC) decrease the loss, initially to the minimum value achieved by the Binary16 baseline, and then continue to slightly reduce loss. However, as mentioned in the figure caption, this regularization behaviour is expected to apply only to

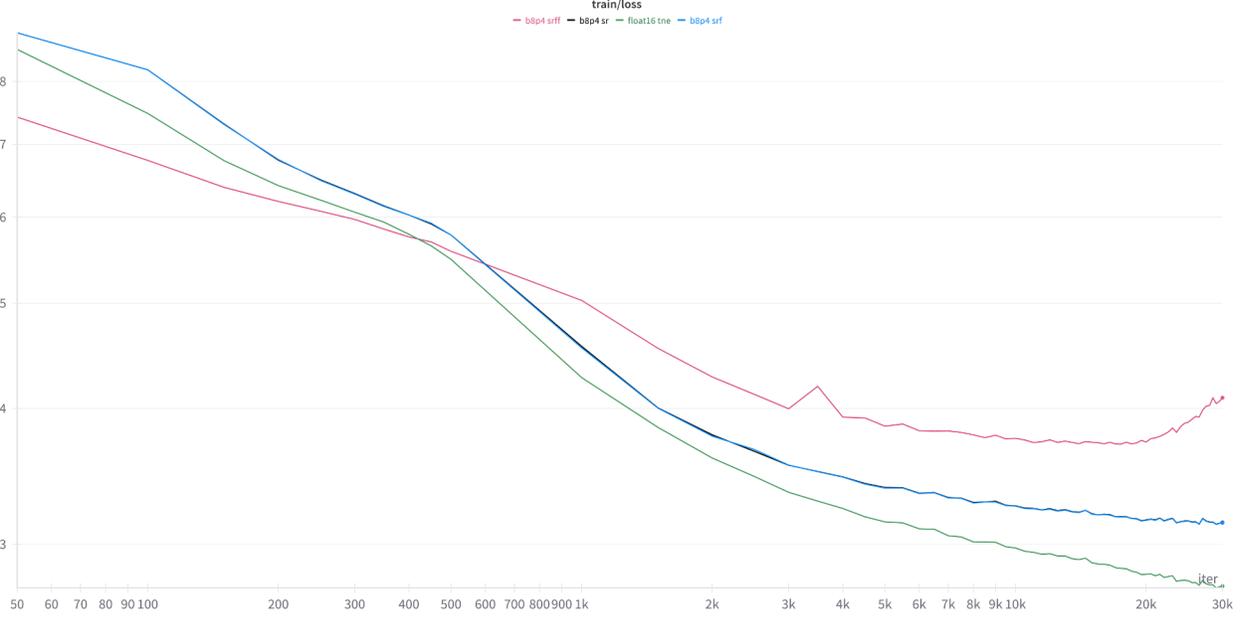


Fig. 4. Experiments on language model training (nanoGPT, GPT-2 (350M params)). Binary16 baseline: validation loss follows training loss (not shown) very closely. Binary8P4 SRFF: initial loss reduction followed by “divergence”; SRF, SRC: convergence to stable values. In this case SRF and SRC converge to better validation errors than SRFF, but not to the same loss as Binary16.

such small-scale experiments, and should not be taken as an indication of any superiority of Binary8+SR over Binary16. The model was then configured to GPT-2 settings (26 layers, 16 heads, embedding dimension 1024), amounting to 354 million parameters. In these experiments, training was stopped at 30,000 iterations, meaning that we expect under-training, so validation loss follows training loss (this was observed). The model was trained in Binary8P4 as above, under SRFF, SRF, SRC modes. In this instance, SRFF attained a final loss value of 4.06, compared to 3.14 for each of SRC, SRF, and 2.74 for Binary16. It should be noted that the training dynamics (and hence the optimal hyperparameter settings) of this simple QAT method are almost certainly quite different from the Binary16 defaults, so it is not necessarily expected that the loss achieved by SR will match Binary16. The key outcome of this experiment is the indication that the correction of SRFF’s bias can avoid divergence. Divergence was not as dramatic as for the smaller model, but was nevertheless significant, suggesting that alternatives to SRFF may be of value in future implementations.

V. CONCLUSION

This paper has illustrated the potential for systematic bias in some plausible implementations of few-bit stochastic rounding. The mitigations of these biases are also presented, with different cost/accuracy tradeoffs. Some experimental evidence is provided to suggest that these biases may prove significant in practice, and that further exploration into the costs of hardware or software implementations may be of value. Source

code implementing all experiments, and including symbolic algebra derivations of the various computations, will be made available.

APPENDIX A BIAS COMPUTATIONS

This section records a detailed derivation of the bias computations in §III.

A. Bias of SRFF, infinite-precision inputs

The calculation of §III-C is expanded here:

$$\begin{aligned}
 & 2^{-N} \sum_{n=0}^{2^N-1} \int_0^1 \mathbb{1}[\mathbb{R}_{\text{SRFF}}(x, n)] dx - \frac{1}{2} \\
 &= 2^{-N} \sum_{n=0}^{2^N-1} \int_0^1 \mathbb{1}[x + n \times 2^{-N} \geq 1] dx - \frac{1}{2} \\
 &= 2^{-N} \sum_{n=0}^{2^N-1} \int_{1-n \times 2^{-N}}^1 dx - \frac{1}{2} \\
 &= 2^{-N} \sum_{n=0}^{2^N-1} n \times 2^{-N} - \frac{1}{2} \\
 &= 2^{-2N} \sum_{n=0}^{2^N-1} n - \frac{1}{2} \\
 &= 2^{-2N} \frac{(2^N - 1)(2^N)}{2} - \frac{1}{2}
 \end{aligned}$$

$$\begin{aligned}
&= 2^{-2N} \frac{2^{2N} - 2^N}{2} - \frac{1}{2} \\
&= \frac{1 - 2^{-N}}{2} - \frac{1}{2} \\
&= -\frac{2^{-N}}{2} \\
&= -2^{-(N+1)} \neq 0
\end{aligned}$$

B. Bias of SRF, infinite-precision inputs

The calculation in §III-D is expanded as follows:

$$bias_{SRF} = I - \frac{1}{2}$$

where

$$\begin{aligned}
I &= 2^{-N} \sum_{n=0}^{2^N-1} \int_0^1 \mathbb{1}[\mathbf{R}_{SRF}(x, n)] dx \\
&= 2^{-N} \sum_{n=0}^{2^N-1} \int_0^1 \mathbb{1}\left[x + \left(n + \frac{1}{2}\right) \times 2^{-N} \geq 1\right] dx \\
&= 2^{-N} \sum_{n=0}^{2^N-1} \int_{1-(n+\frac{1}{2}) \times 2^{-N}}^1 dx \\
&= 2^{-N} \sum_{n=0}^{2^N-1} \left(\left(n + \frac{1}{2}\right) \times 2^{-N}\right)
\end{aligned}$$

so

$$\begin{aligned}
bias_{SRF} &= 2^{-N} \sum_{n=0}^{2^N-1} \left(n \times 2^{-N} + \frac{1}{2} \times 2^{-N}\right) - \frac{1}{2} \\
&= bias_{SRFF} + 2^{-N} \sum_{n=0}^{2^N-1} \left(\frac{1}{2} \times 2^{-N}\right) \\
&= -2^{-(N+1)} + 2^{-(N+1)} = 0
\end{aligned}$$

C. Bias of SRFF, finite-precision inputs

The bias at a single x value is

$$\begin{aligned}
bias_{SRFF}(x) &= 2^{-N} \sum_{n=0}^{2^N-1} \mathbb{1}[\mathbf{R}_{SRFF}(x, n)] - x \\
&= 2^{-N} \sum_{n=0}^{2^N-1} \mathbb{1}[x + 2^{-N}n \geq 1] - x
\end{aligned}$$

and its expectation over the 2^D input values is

$$\begin{aligned}
bias_{SRFF,D} &= 2^{-D} \sum_{i=0}^{2^D-1} bias_{SRFF}(x := 2^{-D}i) \\
&= 2^{-D} \sum_{i=0}^{2^D-1} \left(2^{-N} \sum_{n=0}^{2^N-1} \mathbb{1}[2^{-D}i + 2^{-N}n \geq 1] - 2^{-D}i\right) \\
&= 2^{-D} \sum_{i=0}^{2^D-1} \left(2^{-N} \sum_{n=0}^{2^N-1} \mathbb{1}[2^{-D}i + 2^{-N}n \geq 1] - 2^{-D}i\right)
\end{aligned}$$

$$\begin{aligned}
&= \left(2^{-D} \sum_{i=0}^{2^D-1} 2^{-N} \sum_{n=0}^{2^N-1} \mathbb{1}[2^{-D}i + 2^{-N}n \geq 1]\right) \\
&\quad - 2^{-D} \sum_{i=0}^{2^D-1} 2^{-D}i \\
&= \left(2^{-N} \sum_{n=0}^{2^N-1} 2^{-D} \sum_{i=0}^{2^D-1} \mathbb{1}[i \geq 2^D - 2^{D-N}n]\right) \\
&\quad - \frac{1 - 2^{-D}}{2} \\
&= \left(2^{-N} \sum_{n=0}^{2^N-1} 2^{-D} \sum_{i=[2^D-2^{D-N}n]}^{2^D-1} 1\right) - \frac{1 - 2^{-D}}{2} \\
&= 2^{-N} \sum_{n=0}^{2^N-1} 2^{-D} \left((2^D - 1) - [2^D - 2^{D-N}n] + 1\right) \\
&\quad - \frac{1 - 2^{-D}}{2} \\
&= 2^{-N} \sum_{n=0}^{2^N-1} 2^{-D} [2^{D-N}n] - \frac{1 - 2^{-D}}{2}
\end{aligned}$$

If $N \leq D$ then 2^{D-N} is an integer so the floor operation is a no-op. If $D < N$ then 2^{D-N} is a reciprocal power of two so the argument will be integral for some values of n . We may obtain a bound using $\lfloor x \rfloor \leq x$, which will be tight for $N \leq D$.

$$\begin{aligned}
bias_{SRFF,D} &= 2^{-N} \sum_{n=0}^{2^N-1} 2^{-D} [2^{D-N}n] - \frac{1 - 2^{-D}}{2} \\
&\leq 2^{-N} \sum_{n=0}^{2^N-1} 2^{-N}n - \frac{1 - 2^{-D}}{2} \\
&= \frac{1 - 2^{-N}}{2} - \frac{1 - 2^{-D}}{2} \\
&= \frac{2}{-2^{-(N+1)} + 2^{-(D+1)}} \\
&= \frac{2^{-D} - 2^{-N}}{2}
\end{aligned}$$

D. Bias of SRF, finite-precision inputs

The bias at a single x value is

$$\begin{aligned}
bias_{SRF,D}(x) &= 2^{-N} \sum_{n=0}^{2^N-1} \mathbb{1}[\mathbf{R}_{SRF}(x, n)] - x \\
&= 2^{-N} \sum_{n=0}^{2^N-1} \mathbb{1}\left[x + \left(n + \frac{1}{2}\right) \times 2^{-N} \geq 1\right] - x
\end{aligned}$$

and its expectation over the 2^D input values is

$$\begin{aligned}
&2^{-D} \sum_{i=0}^{2^D-1} bias_{SRF,D}(x := 2^{-D}i) \\
&= 2^{-D} \sum_{i=0}^{2^D-1} \left(2^{-N} \sum_{n=0}^{2^N-1} \mathbb{1}\left[2^{-D}i + \left(n + \frac{1}{2}\right) \times 2^{-N} \geq 1\right] - 2^{-D}i\right) \\
&= 2^{-D} \sum_{i=0}^{2^D-1} \left(2^{-N} \sum_{n=0}^{2^N-1} \mathbb{1}\left[2^{-D}i + \left(n + \frac{1}{2}\right) \times 2^{-N} \geq 1\right] - 2^{-D}i\right)
\end{aligned}$$

$$\begin{aligned}
&= \left(2^{-D} \sum_{i=0}^{2^D-1} 2^{-N} \sum_{n=0}^{2^N-1} \mathbb{1} \left[2^{-D}i + \left(n + \frac{1}{2} \right) \times 2^{-N} \geq 1 \right] \right) \\
&\quad - 2^{-D} \sum_{i=0}^{2^D-1} 2^{-D}i \\
&= \left(2^{-N} \sum_{n=0}^{2^N-1} 2^{-D} \sum_{i=0}^{2^D-1} \mathbb{1} \left[i \geq 2^D - \left(n + \frac{1}{2} \right) \times 2^{D-N} \right] \right) \\
&\quad - \frac{1-2^{-D}}{2} \\
&= \left(2^{-N} \sum_{n=0}^{2^N-1} 2^{-D} \sum_{i=\lceil 2^D - (n+\frac{1}{2}) \times 2^{D-N} \rceil}^{2^D-1} 1 \right) - \frac{1-2^{-D}}{2} \\
&= 2^{-N} \sum_{n=0}^{2^N-1} 2^{-D} \left((2^D-1) - \left\lfloor 2^D - \left(n + \frac{1}{2} \right) \times 2^{D-N} \right\rfloor + 1 \right) \\
&\quad - \frac{1-2^{-D}}{2} \\
&= 2^{-N} \sum_{n=0}^{2^N-1} 2^{-D} \left[\left(n + \frac{1}{2} \right) \times 2^{D-N} \right] - \frac{1-2^{-D}}{2} \\
&= 2^{-N} \sum_{n=0}^{2^N-1} 2^{-D} \left[n \times 2^{D-N} + \frac{1}{2} \times 2^{D-N} \right] - \frac{1-2^{-D}}{2} \\
&= 2^{-N} \sum_{n=0}^{2^N-1} 2^{-D} \left[n \times 2^{D-N} + 2^{D-N-1} \right] - \frac{1-2^{-D}}{2}
\end{aligned}$$

As before, the floor operation is a complication. If $D > N$ then $2^{D-N}n$ and 2^{D-N-1} are integers, so the bound using $\lfloor x \rfloor \leq x$ will be tight for $N < D$.

Hence $\text{bias}_{SRF,D} =$

$$\begin{aligned}
&2^{-N} \sum_{n=0}^{2^N-1} 2^{-D} \left[n \times 2^{D-N} + 2^{D-N-1} \right] - \frac{1-2^{-D}}{2} \\
&\leq 2^{-N} \sum_{n=0}^{2^N-1} (n \times 2^{-N} + 2^{-N-1}) - \frac{1-2^{-D}}{2} \\
&= 2^{-2N} \sum_{n=0}^{2^N-1} n + 2^{-N-1} - \frac{1-2^{-D}}{2} \\
&= 2^{-2N} \frac{(2^N-1)2^N}{2} + 2^{-N-1} - \frac{1-2^{-D}}{2} \\
&= \frac{1-2^{-N}}{2} + 2^{-N-1} - \frac{1-2^{-D}}{2} \\
&= -2^{-(N+1)} + 2^{-N-1} + 2^{-(D+1)} \\
&= 2^{-(D+1)}
\end{aligned}$$

so the bias depends on D but is independent of N (noting that $N \leq D$ is required for the computation to be exact.)

REFERENCES

- [1] M. Essam, T. B. Tang, E. T. W. Ho, and H. Chen, "Dynamic point stochastic rounding algorithm for limited precision arithmetic in deep belief network training," in *2017 8th International IEEE/EMBS Conference on Neural Engineering (NER)*. IEEE, 2017, pp. 629–632.
- [2] B. Nouné, P. Jones, D. Justus, D. Masters, and C. Luschi, "8-bit numerical formats for deep neural networks," *arXiv preprint arXiv:2206.02915*, 2022.
- [3] E. A. Paxton, M. Chantry, M. Klöwer, L. Saffin, and T. Palmer, "Climate modeling in low precision: Effects of both deterministic and stochastic rounding," *Journal of Climate*, vol. 35, no. 4, pp. 1215–1229, 2022.
- [4] S. Q. Zhang, B. McDanel, and H. Kung, "Fast: DNN training under variable precision block floating point with stochastic rounding," in *2022 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*. IEEE, 2022, pp. 846–860.
- [5] M. Croci and M. B. Giles, "Effects of round-to-nearest and stochastic rounding in the numerical solution of the heat equation in low precision," *IMA Journal of Numerical Analysis*, vol. 43, no. 3, pp. 1358–1390, 2023.
- [6] M. Croci, M. Fasi, N. J. Higham, T. Mary, and M. Mikaitis, "Stochastic rounding: implementation, error analysis and applications," *Royal Society Open Science*, vol. 9, no. 3, p. 211631, 2022.
- [7] E.-M. El Arar, D. Sohler, P. de Oliveira Castro, and E. Petit, "Stochastic rounding variance and probabilistic bounds: A new approach," *SIAM Journal on Scientific Computing*, vol. 45, no. 5, pp. C255–C275, 2023.
- [8] E.-M. El Arar, M. Fasi, S.-I. Filip, and M. Mikaitis, "Probabilistic error analysis of limited-precision stochastic rounding," *arXiv preprint arXiv:2408.03069*, 2024.
- [9] L. Xia, M. Anthonissen, M. Hochstenbach, and B. Koren, "Improved stochastic rounding," *arXiv preprint arXiv:2006.00489*, 2020.
- [10] G. Yuan, S.-E. Chang, Q. Jin, A. Lu, Y. Li, Y. Wu, Z. Kong, Y. Xie, P. Dong, M. Qin *et al.*, "You already have it: A generator-free low-precision DNN training framework using stochastic rounding," in *European Conference on Computer Vision*. Springer, 2022, pp. 34–51.
- [11] A. Or, J. Zhang, E. Smothers, K. Khandelwal, and S. Rao, "Quantization-aware training for large language models with PyTorch," The PyTorch Foundation, 2024. [Online]. Available: <https://pytorch.org/blog/quantization-aware-training>
- [12] A. Karpathy, "NanoGPT," <https://github.com/karpathy/nanoGPT>, 2024.
- [13] A. Fitzgibbon, "GFloat: Generic floating point formats in Python," <https://gfloat.readthedocs.io/>, 2024.
- [14] IEEE SA P3109 Working Group, "Interim report on binary floating-point formats for machine learning," The Institute of Electrical and Electronics Engineers, Inc., 2024. [Online]. Available: <https://github.com/P3109/Public>