

UC Irvine

UC Irvine Previously Published Works

Title

PLAM: a Posit Logarithm-Approximate Multiplier

Permalink

<https://escholarship.org/uc/item/2bg651zz>

Authors

Murillo, Raul
Barrio, Alberto A Del
Botella, Guillermo
[et al.](#)

Publication Date

2021-02-18

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

PLAM: A Posit Logarithm-Approximate Multiplier

Raul Murillo*, Alberto A. Del Barrio*, Guillermo Botella*, Min Soo Kim[†], HyunJin Kim[‡], Nader Bagherzadeh[§]

*Department of Computer Architecture and Automation, Complutense University of Madrid, 28040 Madrid, Spain

[†]NGD Systems, Irvine, California 92618, USA

[‡]School of Electronics and Electrical Engineering, Dankook University, Yongin-si, Gyeonggi-do 16890, Republic of Korea

[§]Department of Electrical Engineering and Computer Science, University of California, Irvine, California 92697, USA

Email: *{ramuri01, abarriog, gbotella}@ucm.es, [†]minsk1@uci.edu, [‡]hyunjin2.kim@gmail.com, [§]nader@uci.edu

Abstract—The Posit™ Number System was introduced in 2017 as a replacement for floating-point numbers. Since then, the community has explored its application in Neural Network related tasks and produced some unit designs which are still far from being competitive with their floating-point counterparts. This paper proposes a Posit Logarithm-Approximate Multiplication (PLAM) scheme to significantly reduce the complexity of posit multipliers, the most power-hungry units within Deep Neural Network architectures. When comparing with state-of-the-art posit multipliers, experiments show that the proposed technique reduces the area, power, and delay of hardware multipliers up to 72.86%, 81.79%, and 17.01%, respectively, without accuracy degradation.

I. INTRODUCTION

Deep Neural Networks (DNNs) dominate nowadays machine learning landscape due to the great advances in a large variety of applications, including computer vision, natural language processing or speech recognition. This continuous improvement of the state of the art has been accompanied by an increase in computational complexity and an overhead in hardware resources. While DNN models are commonly trained on high-end GPUs, reducing the computational complexity of DNNs to perform inference on resource-constraint devices has been a serious challenge and a lengthy line of research [1].

The IEEE 754 standard for floating-point arithmetic has been for decades the de facto implementation for the vast majority of real number-based applications. However, in the last years different computer arithmetic encodings and formats have been considered for DNN training and inference, including half-precision, fixed-point, bfloat16, or 8-bit integer quantization [2]. But probably, one of the most promising alternatives to the floating-point standard is the Posit Number System (PNS) [3]. Posit numbers provide a better trade-off than floating point between dynamic range and numerical precision, with a larger dynamic range under the same bitwidth and, most importantly, tapered accuracy around ± 1 , corresponding with the DNN weight distribution [4].

One of the most widespread techniques in DNNs is quantization [5], which allows the bitwidth reduction of the units deployed in the corresponding circuit. This contributes diminishing area, power or memory footprint. Furthermore, approximate computing techniques [6] have commonly been used in DNNs too, specially for inference in real time systems and resource-constraint devices, where the trade off between inference accuracy and execution time or resources is always




Fig. 1. Resource distribution of a Posit(32, 2) multiplier.

present. The inference stage mainly consists of addition and multiplication operations, the latter being the most expensive one from a hardware perspective. Despite the fact that the PNS has shown its potential in DNNs [4], [7]–[10], it is true that posit units are far from being competitive in terms of power with respect to FP or bfloat16 formats [11]–[16]. Without much details, which will be commented in Section III, a posit multiplier consists of similar stages to a IEEE 754 floating-point multiplier, i.e. unpacking/decoding of operands, fraction multiplication and packing/encoding of the result. As can be seen in Fig. 1, the fraction multiplier is, by far, the module with the highest consumption of resources. For this reason, as in the case of floating-point formats, reducing the complexity of the fraction multiplier is critical to optimize the power consumption of the whole unit.

To improve the efficiency of posit-based DNN inference, this work proposes the use of logarithm-approximate multipliers in combination with the PNS. Experimental results reveal that adopting the proposed Posit Logarithm-Approximate Multiplier (PLAM) allows to significantly reduce area, power, and delay with negligible accuracy degradation while taking advantage of the benefits of this novel arithmetic format. To the best of our knowledge, this is the first work proposing the use of approximate posit multipliers. Thus, the main contributions of this paper can be summarized as follows:

- Proposing an algorithm for performing logarithm-approximate multiplication in posit arithmetic, with a bounded error of 11.1%.
- Testing the proposed algorithm at inference for several




Fig. 2. Layout of an $\text{Posit}(n, es)$ number. The variable-length regime field may cause exponent be encoded with less than es bits, even no bits if regime is wide enough. Same occurs with the fraction.

DNNs including LeNet-5 and CifarNet, and well-known datasets as MNIST, SVHN and CIFAR-10, achieving negligible accuracy degradation with 16-bit posits.

- Implementing the proposed algorithm in the open-source FloPoCo framework, reducing area and power up to 72.86% and 81.79%, respectively, when comparing to cutting edge posit implementations.

The rest of the paper is organized as follows: Section II reviews strategies for using posits in NNs and implementations from related works. Section III introduces PLAM and discusses its approximation error. Experimental results of adopting PLAM for inference in different DNNs are presented in Section IV. Section V presents evaluations of area, power, and delay, and discusses the improvements of PLAM against previous posit and floating-point implementations. Finally, Section VI concludes this paper.

II. RELATED WORK

A. The Posit Number System

A posit format is defined as a tuple $\langle n, es \rangle$, where n is the total bitwidth and es is the maximum number of bits reserved for the exponent field. As Fig. 2 shows, posit numbers are encoded with four fields: a sign bit (s), several bits for encoding the regime value (k), up to es bits for the exponent (e), and the remaining bits for fraction (f). Thus, the numerical value X of a generic $\text{Posit}(n, es)$ is expressed by (1).

$$X = (-1)^s \times (2^{2^{es}})^k \times 2^e \times (1 + f), \quad (1)$$

$$k = -x_{n-2} + \sum_{i=n-2}^{x_i \neq x_{n-2}} (-1)^{1-x_i}. \quad (2)$$

The main differences with floating-point format are the utilization of an unsigned and unbiased exponent, if there exists such exponent field, and the existence of the regime field. This new field consists of a sequence of bits with the same value finished with the negation of such value, as shown in Fig. 2. Provided that $X = x_{n-1}x_{n-2}\dots x_1x_0$, this regime can be expressed as (2) shows. It is noteworthy that, while the new regime field provides important scaling capabilities that improve the dynamic range of posits, detecting the resulting varying-sized fields adds a hardware overhead.

B. The use of Posit Arithmetic in NNs

Posits were introduced by John Gustafson in 2017. Since then, multiple works have explored the benefits of this novel format against the standard floating-point, and most of them

focusing on NNs. In [7], J. Johnson designed an arithmetic unit for combining posit addition together with logarithmic multiplication for performing CNN inferences. Authors in [8] employ a posit DNN accelerator to represent weights and activations combined with an FPGA soft core for 8-bit posit exact-multiply-and-accumulate (EMAC) operations. In all these works, DNN training is performed in floating-point, while the inference stage is performed in low-precision posit format. Later works proposed different approaches for training NNs using the posit format, either directly training on this format with different precision [4], [9], [17], or with the help of a warmup training using floating-point format [10].

Several previous works have also proposed to introduce approximate multipliers into DNN inferences. Logarithmic multiplication has successfully been employed to fixed-point small DNN models in [18] and even to large ones in [19]. Finally, authors in [20] explore its use within floating-point multipliers to reduce the costs of small NNs during the training and inference stages. To the best of our knowledge, at the time of writing there are no previous works proposing the design of approximate posit multipliers and its use for deep learning tasks. With regard to other approximate computing techniques on posit arithmetic, Cococcioni et al. [17] explored the possibility to approximate simple operations and activation functions in DNNs using only ALU-based operations.

C. Posit Arithmetic-Based Implementations

Since the appearance of the PNS, several hardware implementations for this arithmetic format have been proposed. An open-source parameterized adder/subtractor was presented in [11], whose concepts were expanded in [12] to design a parameterized posit multiplier. These two works did not perform posit rounding, but fraction truncation, and used both a Leading One Detector (LOD) and a Leading Zero Detector (LZD) to determine regime value, which results in redundant area. This shortcomings were solved in [13], where only a LZD was used at the cost of inverting negative regimes, and results were correctly rounded using the round to nearest even scheme. The same idea was applied in [14], where authors expanded their previous works [11], [12] and presented an open-source posit core generator which included a parameterized divider based on the Newton-Raphson method. Another open-source posit core generator is presented [16], where parameterized adder and multiplier designs were integrated into the FloPoCo framework, allowing even to generate posit units with no exponent bits, in contrast with previous works. However, posits are still in development. As has been mentioned, in terms of delay, area, and power, the arithmetic units are not yet competitive against their floating-point counterparts, and although they have shown some promising improvements in the NNs field [4], [9] there is still some debate about their real improvement [21].

In this work, we propose the design of a log-based approximate posit multiplier to perform DNN inference without degrading the accuracy of the results. With this approach, it

is possible to drastically reduce the complexity of the posit multiplication.

III. PLAM: THE POSIT LOGARITHM-APPROXIMATE MULTIPLIER

Hardware multiplication is an expensive operation, specially between two real numbers, where all the state-of-the-art formats mandate multiplying two fixed-point values. As depicted in Fig. 1 for the posit format, this multiplication provokes the majority of the power consumption. Instead, logarithm multiplication [18], [22] avoids hardware multipliers entirely by approximating the multiplication as a fixed-point addition. This section introduces the Posit Logarithm-Approximate Multiplier (PLAM), and analyzes its approximation error.

A. Exact posit multiplication

As has been mentioned, while posit encoding may differ from usual floating-point, the core of the operations is quite similar between these number formats, with exception on the decoding and encoding of the posit fields [16]. In addition to this, in the PNS there are no special cases to be taken care of, as the denormal numbers in the case of floating-point based formats, a single rounding mode, i.e. round to nearest even, and unique representations for zero and infinite values. Provided that a posit number X is represented by the tuple (S_X, K_X, E_X, F_X) , where S_X, K_X, E_X, F_X , are the sign, regime, exponent and fraction values, respectively, the multiplication of two posit values $C = A \times B$ is depicted in Fig. 3. The computation of the different fields is defined by (3) to (6).

$$S = S_A \oplus S_B, \quad (3)$$

$$K = K_A + K_B, \quad (4)$$

$$E = E_A + E_B, \quad (5)$$

$$F = (1 + F_A) \times (1 + F_B). \quad (6)$$

Note that $(1 + F_A)$ (respectively $1 + F_B$) is obtained by appending a hidden bit with value 1 to the binary representation of F_A (respectively F_B). Therefore, the resulting posit C is obtained as described in (7) to (10).

$$S_C = S, \quad (7)$$

$$K_C = \begin{cases} K & \text{if } E_C \geq E, \\ K + 1 & \text{otherwise,} \end{cases} \quad (8)$$

$$E_C = \begin{cases} E \bmod 2^{es} & \text{if } F < 2, \\ (E + 1) \bmod 2^{es} & \text{otherwise,} \end{cases} \quad (9)$$

$$F_C = \begin{cases} F - 1 & \text{if } F < 2, \\ F/2 - 1 & \text{otherwise.} \end{cases} \quad (10)$$

B. Logarithm approximated posit multiplication

Once the posit multiplication is explained, the insights on how to approximate this operation taking advantage of the logarithmic multiplication will be described in detail.




Fig. 3. Exact posit multiplication.

Firstly, it is worthy to note that in the multiplication operation the computation of the sign is independent to the computation of the other fields. Therefore, let us focus on positive posit numbers from now on. In such case, (1) is simplified as shown in (11),

$$X = (2^{2^{es}})^k \times 2^e \times (1 + f). \quad (11)$$

Taking logarithms on both sides, (11) becomes

$$\log_2 X = 2^{es} \times k + e + \log_2(1 + f) \approx 2^{es} \times k + e + f, \quad (12)$$

where the approximation of the right term is based on the property

$$\log_2(1 + x) \approx x, \text{ for } 0 \leq x \leq 1. \quad (13)$$

Converting numbers to the logarithmic domain allows computing the multiplication as the addition of fixed point numbers. Thus, in order to compute the multiplication of two posits $C = A \times B$ as a logarithm-approximate multiplication, the different fields are processed as described in (14) to (17).

$$S = S_A \oplus S_B, \quad (14)$$

$$K = K_A + K_B, \quad (15)$$

$$E = E_A + E_B, \quad (16)$$

$$F = F_A + F_B, \quad (17)$$

where (14), (15) and (16) are identical to (3), (4) and (5). The only difference is in the computation of the fraction field, where (17) substitutes the product of (6) by an addition thanks to the log property shown in (13). Then, the multiplication

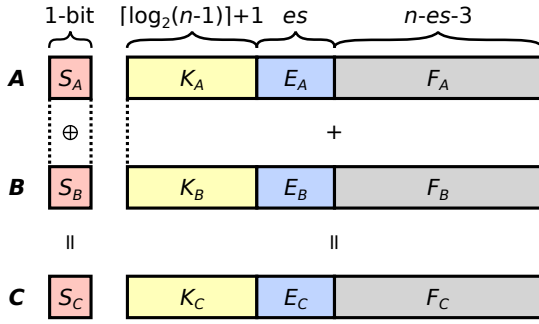


Fig. 4. Algorithm for implementing PLAM in hardware.

result is expressed by (18) to (21).

$$S_C = S, \quad (18)$$

$$K_C = \begin{cases} K & \text{if } E_C \geq E, \\ K + 1 & \text{otherwise,} \end{cases} \quad (19)$$

$$E_C = \begin{cases} E \bmod 2^{es} & \text{if } F < 1, \\ (E + 1) \bmod 2^{es} & \text{otherwise,} \end{cases} \quad (20)$$

$$F_C = \begin{cases} F & \text{if } F < 1, \\ F - 1 & \text{otherwise.} \end{cases} \quad (21)$$

Finally, let us focus on the hardware implementation of PLAM based on the algorithm presented in (14) to (21). According to (12), a posit in log-domain gets its regime value k multiplied by 2^{es} and then adds the exponent value. When implementing this in hardware, it is equivalent to concatenate both the regime and exponent bit fields. In this way, the condition in (19) can be efficiently computed in hardware, as the overflow coming from $E_A + E_B$ is directly added as a carry-in to $K_A + K_B$. Besides, conditions in (20) and (21) can be implemented in hardware in the same manner, as illustrated in Fig. 4.

C. Approximation error

To analyze the error of PLAM, which directly depends on the input operands, let us consider two positive posit numbers $A = s_A(1 + f_A)$ and $B = s_B(1 + f_B)$, where the scaling factor $s_i = 2^{(2^{es} \times k_i + e_i)}$. In such case, the results of exact and approximate multiplication $C = A \times B$ are given by (22) and (23), respectively.

$$C_{exact} = s_A s_B (1 + f_A)(1 + f_B) \quad (22)$$

$$C_{PLAM} = \begin{cases} s_A s_B (1 + f_A + f_B) & \text{if } f_A + f_B < 1, \\ 2s_A s_B (f_A + f_B) & \text{otherwise.} \end{cases} \quad (23)$$

The relative error of approximation, defined by (24), is a function of just f_A and f_B . As these parameters are restricted to the interval $[0, 1)$, the maximum error is 11.1%, which is obtained when both fractions are equal to 0.5, as Mitchell demonstrated in [22]. It is noteworthy that neither the exponents nor the novel regime fields affect the error value, it just depends on the fractions.

TABLE I
DNNs SETUP

Dataset	Architecture	Optimizer	Batch size	Epochs
ISOLET	(617, 128, 64, 26)	SGD	64	30
UCI HAR	(561, 512, 512, 6)	Nesterov	32	30
MNIST	LeNet-5	Adam	128	50
SVHN	LeNet-5	Adam	128	50
CIFAR-10	CifarNet	Adam	128	30

$$\begin{aligned} error &= \frac{C_{exact} - C_{PLAM}}{C_{exact}} \\ &= \begin{cases} \frac{f_A f_B}{(1+f_A)(1+f_B)} & \text{if } f_A + f_B < 1, \\ \frac{(1-f_A)(1-f_B)}{(1+f_A)(1+f_B)} & \text{otherwise.} \end{cases} \end{aligned} \quad (24)$$

The impact of PLAM error on the DNN inference will be investigated in Section IV, and its advantage in terms of the speed, power, and area will be discussed in Section V.

IV. EXPERIMENTAL RESULTS FOR DNN INFERENCE

This section demonstrates the effect of PLAM and posits in DNN inference.

A. Experimental setup

To demonstrate the effect of the proposed approximate posit multiplication in DNN inference, we evaluate accuracy results on different datasets and topologies. Table I lists the different datasets, network architectures and training configurations used for the experiments in this paper. For numeric datasets, fully connected networks with 2 hidden layers are employed (for which the number of neurons per layer is indicated in the table), while for image datasets, convolutional models as LeNet-5 [23] or CifarNet [24] are more appropriate. For all setups, ReLU is used as activation function of hidden layers, and softmax function is applied to the output layer.

Due to the lack of hardware support for the posit number system, computations on this format are performed via software emulation. The open-source framework Deep PeNSieve [4] allows to generate DNN models and perform inference and training entirely using posits. Deep PeNSieve relies on the reference library SoftPosit. In this work, both libraries are extended to support PLAM operation for both scalar and matrix multiplication using the algorithm explained in Section III. It must be noted that larger DNNs cannot be efficiently trained using this framework, since the lack of native support makes the training times too long. For instance, training CifarNet on an Intel® Core™ i7-9700K processor with 32 GB of RAM takes around 10 days.

B. Evaluation for DNNs

As previous works demonstrate [4], [10], 16-bit posits can be used for DNN training with no accuracy loss with respect the baseline 32-bit floating-point format. Accordingly, each model is trained under single-precision floating-point and

TABLE II
ACCURACY RESULTS FOR THE INFERENCE STAGE

Dataset	Float 32-bit		Posit(16, 1)		Posit(16, 1) _{PLAM}	
	Top-1	Top-5	Top-1	Top-5	Top-1	Top-5
ISOLET	0.9066	0.9568	0.9093	0.9585	0.9051	0.9585
UCI HAR	0.9383	0.9841	0.9307	0.9841	0.9282	0.9841
MNIST	0.9907	0.9999	0.9903	1.0	0.9898	1.0
SVHN	0.8624	0.9794	0.8513	0.9766	0.8489	0.9761
CIFAR-10	0.6933	0.9722	0.7247	0.9744	0.7251	0.9743

Posit(16, 1) formats. Trained posit models are then modified to use PLAM at inference stage. Three different models are trained for each dataset, and the averages of the results are presented in Table II. Approximate posit multiplication reaches similar accuracy on the inference stage as exact posit multiplication. 16-bit floating-point inference presented similar results as using single-precision, so it is omitted. It is quite remarkable that posits perform better than floats in CIFAR-10, but as authors mention in [4], this still need further evaluations in order to test larger architectures with regularization layers. In any case, the purpose of this paper is to prove that PLAM can perform as accurate as exact posit formats in posit-based DNNs.

V. EVALUATION OF HARDWARE IMPLEMENTATION

After evaluating the accuracy of PLAM when performing DNN inference, this section presents synthesis results for the proposed PLAM, demonstrating the advantage of approximate units from the hardware perspective.

The proposed architecture has been implemented into FloPoCo, an open-source C++ framework for the generation of arithmetic datapaths that provides a command-line interface that inputs operator specifications and outputs synthesizable VHDL [25]. This tool allows operators to be automatically generated with the specified parameters and, therefore, to obtain posit operators for arbitrary values of $\langle n, es \rangle$ with the same base design. The proposed PLAM design, which includes support for correct rounding, has been made publicly available in the FloPoCo git repository. Simulations with extensive testing vectors are performed to verify the functionality of the proposed design. The testing vectors are generated by extending SoftPosit library with support for logarithm-approximate multiplication.

To evaluate the impact of PLAM in terms of hardware resources, 16-bit and 32-bit models (without pipelining) of the proposed architecture have been generated. These operators will be then compared with state-of-the-art implementations of posit and floating-point exact multipliers. The latter ones have been generated using the FloPoCo library. However, it is important to mention that the designs provided by this tool do not include support for denormal numbers or full exception handling, and thus, use less resources compared with a fully IEEE-754 compliant implementation. To have a fair comparison with the works presented in [13], [15], the

TABLE III
FPGA RESOURCE UTILIZATION

Work	16-bits		32-bits	
	LUTs	DSP	LUTs	DSP
[12]	263	1	646	4
[13]	218	1	572	4
[14]	273	1	682	4
[15]	253	1	469	4
[16]	237	1	604	4
(prop.)	185	0	435	0

proposed operators, as well as the open-source implementations (Posit-HDL [12]¹, PACoGen [14]² and FloPoCo-Posit [16]³) have been synthesized on a Zedboard with a Zynq-7000 SoC, using Vivado 2020.1 with default settings. FPGA synthesis results, reported in Table III, show a clear reduction in resource utilization. It must be noted that PLAM uses less LUTs than the rest of posit multipliers as well as no DSPs.

Standard cell synthesis has been performed using Synopsys Design Compiler with a 45-nm library by TSMC for the proposed generated units and open-source implementations. The results for $es = 2$ are compared graphically in Fig. 5. Similar results are obtained for different exponent sizes. As can be seen, area and power savings are greater as the bitwidth increases, obtaining respective reductions of 69.06% and 63.63% in the 16-bit case and of 72.86% and 81.79% for 32-bit multipliers in comparison with units from [16]. Besides, significant resource savings are obtained compared to the single-precision floating-point operator. Under the same 32-bit length, the proposed posit approximate multiplier reduces area and power by 50.40% and 66.86%, respectively. These savings decrease in the 16-bit case, but the area and power usage is still lower than for the half-precision multiplier, closer to the bfloat16 one. Furthermore, it must be reminded that these FloPoCo-generated units do not consider special cases, while the posit-based units do not have to do this.

On the other hand, the delay-reduction is not as pronounced as in the previous case (up to a 17.01% with respect to the 32-bit multiplier from Posit-HDL), and it is still higher than the corresponding floating-point operator under the same bitwidth. This is due to the complexity of detecting the variable-length fields of posits, which is a design challenge for this format.

For a more thorough evaluation of the proposed design, area, power and energy are evaluated in different scenarios with a maximum delay constraint. As can be seen in the results depicted in Fig. 6, the approximate 32-bit posit multiplier is by far more efficient than exact posit units, and even better than the equivalent floating-point unit, while in the case of

¹source code accessed on November 1, 2020 from <https://github.com/manish-kj/Posit-HDL-Arithmetic>

²source code accessed on November 1, 2020 from <https://github.com/manish-kj/PACoGen>

³source code accessed on November 1, 2020 from <https://gitlab.inria.fr/fdupont/flopoco>




Fig. 5. Posit($n, 2$) and floating-point multiplier implementation results. The prefix ‘Flo’ stands for FloPoCo-generated.




Fig. 6. Results for time-constrained multiplier implementations. Implementations marked with ‘*’ do violate the maximum delay constraint.

16-bits, the resource consumption of PLAM is similar to that produced by floating-point multipliers. Only the Flopoco-generated bfloat16 implementation shows better figures.

VI. CONCLUSIONS

While the posit format has demonstrated to be a promising alternative to the IEEE-754 floating-point standard for DNNs, arithmetic units are still far from being competitive in terms of power. This paper aims to reduce such a gap by proposing a Posit Logarithm-Approximate Multiplication (PLAM) scheme to reduce posit multiplication complexity. The experimental results show that applying PLAM in DNN inference does not affect accuracy. When compared to other posit hardware solutions, the proposed implementation achieves area, power, and delay reduction of 72.86%, 81.79%, and 17.01%, respectively.

REFERENCES

- [1] S. Bianco, R. Cadene, L. Celona, and P. Napoletano, “Benchmark analysis of representative deep neural network architectures,” *IEEE Access*, vol. 6, pp. 64 270–64 277, 2018.
- [2] J. L. Hennessy and D. A. Patterson, “A new golden age for computer architecture,” *Communications of the ACM*, vol. 62, no. 2, pp. 48–60, jan 2019.
- [3] J. L. Gustafson and I. Yonemoto, “Beating Floating Point at its Own Game: Posit Arithmetic,” *Supercomputing Frontiers and Innovations*, vol. 4, no. 2, pp. 71–86, jun 2017.
- [4] R. Murillo, A. A. Del Barrio, and G. Botella, “Deep PeNSieve: A deep learning framework based on the posit number system,” *Digital Signal Processing*, vol. 102, p. 102762, jul 2020.
- [5] R. Krishnamoorthi, “Quantizing deep convolutional networks for efficient inference: A whitepaper,” *arXiv e-prints*, jun 2018.
- [6] U. Lotrič and P. Bulić, “Applicability of approximate multipliers in hardware neural networks,” *Neurocomputing*, vol. 96, pp. 57–65, 2012.
- [7] J. Johnson, “Rethinking floating point for deep learning,” *arXiv e-prints*, nov 2018. [Online]. Available: <http://arxiv.org/abs/1811.01721>
- [8] Z. Carmichael *et al.*, “Deep Positron: A Deep Neural Network Using the Posit Number System,” in *2019 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, mar 2019, pp. 1421–1426.
- [9] H. F. Langroudi, Z. Carmichael, and D. Kudithipudi, “Deep Learning Training on the Edge with Low-Precision Posits,” *arXiv e-prints*, pp. 1474–1479, jul 2019.
- [10] J. Lu *et al.*, “Evaluations on Deep Neural Networks Training Using Posit Number System,” *IEEE Transactions on Computers*, vol. 14, no. 8, pp. 1–1, 2020.
- [11] M. K. Jaiswal and H. K. So, “Architecture Generator for Type-3 Unum Posit Adder/Subtractor,” in *2018 IEEE International Symposium on Circuits and Systems (ISCAS)*, vol. 2018-May. IEEE, 2018, pp. 1–5.
- [12] M. K. Jaiswal and H. K. So, “Universal number posit arithmetic generator on fpga,” in *2018 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, vol. 2018-Janua. IEEE, mar 2018, pp. 1159–1162.
- [13] R. Chaurasiya *et al.*, “Parameterized Posit Arithmetic Hardware Generator,” in *2018 IEEE 36th International Conference on Computer Design (ICCD)*. IEEE, oct 2018, pp. 334–341.
- [14] M. K. Jaiswal and H. K. So, “PACoGen: A hardware posit arithmetic core generator,” *IEEE Access*, vol. 7, pp. 74 586–74 601, 2019.
- [15] Y. Uguen, L. Forget, and F. de Dinechin, “Evaluating the Hardware Cost of the Posit Number System,” in *2019 29th International Conference on Field Programmable Logic and Applications (FPL)*. IEEE, sep 2019, pp. 106–113.
- [16] R. Murillo, A. A. Del Barrio, and G. Botella, “Customized posit adders and multipliers using the FloPoCo core generator,” in *2020 IEEE International Symposium on Circuits and Systems (ISCAS)*. IEEE, oct 2020, pp. 1–5.
- [17] M. Cococcioni, F. Rossi, E. Ruffaldi, and S. Saponara, “Fast deep neural networks for image processing using posits and ARM scalable vector extension,” *Journal of Real-Time Image Processing*, vol. 17, no. 3, pp. 759–771, 2020.
- [18] M. S. Kim *et al.*, “Efficient Mitchell’s Approximate Log Multipliers for Convolutional Neural Networks,” *IEEE Transactions on Computers*, vol. 68, no. 5, pp. 660–675, may 2019.
- [19] M. S. Kim, A. A. Del Barrio, H. Kim, and N. Bagherzadeh, “Effects of approximate multiplication on convolutional neural networks,” *arXiv e-prints*, pp. 1–12, jul 2020. [Online]. Available: <http://arxiv.org/abs/2007.10500>

- [20] T. Y. Cheng *et al.*, “Logarithm-approximate floating-point multiplier is applicable to power-efficient neural network training,” *Integration*, vol. 74, no. November 2019, pp. 19–31, 2020.
- [21] F. de Dinechin, L. Forget, J.-M. Muller, and Y. Uguen, “Posits: the good, the bad and the ugly,” in *Proceedings of the Conference for Next Generation Arithmetic 2019*. New York, NY, USA: ACM, mar 2019, pp. 1–10.
- [22] J. N. Mitchell, “Computer multiplication and division using binary logarithms,” *IRE Transactions on Electronic Computers*, vol. EC-11, no. 4, pp. 512–517, 1962.
- [23] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, “Gradient-based learning applied to document recognition,” *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2323, 1998.
- [24] A. Krizhevsky, “Learning multiple layers of features from tiny images,” Ph.D. dissertation, University of Toronto, Canada, 2009.
- [25] F. de Dinechin and B. Pasca, “Designing custom arithmetic data paths with FloPoCo,” *IEEE Design & Test of Computers*, vol. 28, no. 4, pp. 18–27, jul 2011.