$\frac{2}{3}$

MOHSEN IMANI, RICARDO GARCIA, SARANSH GUPTA, and TAJANA ROSING, University of

4 California San Diego

5 Many applications, such as machine learning and data sensing are statistical in nature and can tolerate some level of inaccuracy in their 6 computation. A variety of designs have been put forward exploiting the statistical nature of machine learning, through approximate 7 computing. With approximate multipliers being the main focus due to their high usage in machine learning designs. In this paper, we 8 propose a novel approximate floating point multiplier, called CMUL, which significantly reduces energy and improves performance of multiplication while allowing for a controllable amount of error. Our design approximately models multiplication by replacing the most 9 costly step of the operation with a lower energy alternative. In order to tune the level of approximation, CMUL dynamically identifies the inputs which produces the largest approximation error and processes them in precise mode. In order to use CMUL for deep neural network (DNN) acceleration, we propose a framework which modifies the trained DNN model to make it suitable for approximate 11 hardware. Our framework adjusts the DNN weights to a set of "potential weights" that are suitable for approximate hardware. Then, 12 it compensates the possible quality loss by iteratively retraining the network. Our evaluation with four DNN applications shows that, 13CMUL can achieve 60.3% energy efficiency improvement and 3.2× energy-delay product (EDP) improvement as compared to the baseline GPU, while ensuring less than 0.2% quality loss. These results are 38.7% and 2.0× higher than energy efficiency and EDP 14 improvement of the CMUL without using the proposed framework. 15

 $\begin{tabular}{l} \mbox{CCS Concepts: \bulletComputer systems organization \rightarrow Architectures; \bulletComputing methodologies \rightarrow Machine learning; $$16$ \end{tabular}$

Additional Key Words and Phrases: Approximate computing, Neural network, Floating point unit, Energy efficiency

18 ACM Reference format:

- Mohsen Imani, Ricardo Garcia, Saransh Gupta, and Tajana Rosing. 2010. Hardware-Software Co-design to Accelerate Neural Network
 Applications. ACM Comput. Entertain. 9, 4, Article 39 (March 2010), 16 pages.
- 20 DOI: 0000001.0000001
- 21

²² 1 INTRODUCTION

23In 2015, the number of smart devices around the world exceeded 25 billion. This number is expected to double by 2020 (Atzori et al. 2010; Gantz and Reinsel 2011). Many of these devices have batteries with strict power constraints, so $\mathbf{24}$ the need for systems that can efficiently handle the computing requirements of data-intensive workloads is undeniable (Ji 25et al. 2012; Khoshavi et al. 2016). Deep neural networks (DNNs) have been effectively used for diverse classifica-26 tion problems, such as image processing, video segmentation, speech recognition, computer vision, health-care, and 27 manufacturing (Hinton et al. 2012; Imani et al. 2018b,c; LeCun et al. 2010; Oquab et al. 2014; Salamat et al. 2018). Running DNNs on the general purpose processors is slow, energy hungry, and prohibitively expensive (Krizhevsky et al. 28 2012). Machine learning applications are stochastic in heart, thus they do not need highly accurate computation. So, by 29 accepting slight inaccuracy, instead of doing all computation precisely, we can get significant energy and performance 30 improvements (Han and Orshansky 2013; Imani et al. 2016d). Therefore, many traditional and state-of-the-art computing 31 systems use floating point units (FPUs) (Courbariaux et al. 2014; Razlighi et al. 2017). For such algorithms of high energy and performance high power is required. To cover the same dynamic range, the fixed point unit must be five times 32 larger and 40% slower than a corresponding floating point (Liang et al. 2003). Similarly, many DNN applications require 33 floating-point precision due to the fact that the iterative training algorithm often update the parameters using gradients 34 whose values are too small to sustain the additive quantization noise (Lin and Talathi 2016). Multiplication is one of the most common and costly FP operations, slowing down the computation in many applications 35

- such as signal processing, neural networks, and streaming processes (Imani et al. 2016c; Suhre et al. 2013). Multiplication
 cost can be reduced by designing an approximate multiplication unit. Most of prior work attempted to reduce the bit-size
- Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
- © 2010 ACM. Manuscript submitted to ACM

42 Manuscript submitted to ACM



Fig. 1. Approximate multiplication of proposed CMUL between A ad B operands.

of multiplication to enable approximation (Hashemi et al. 2015; Narayanamoorthy et al. 2015). However, either the lack
 of accuracy tuning or the large area requirements of the tuned designs, significantly reduce the advantages provided by such approximation.

13In this paper, we instead propose a configurable floating point multiplication, called CMUL, which significantly 14 improves the multiplication energy consumption by trading off accuracy. CMUL avoids the costly multiplication when calculating the fractional part of a floating point number by adding the input mantissas, instead of multiplying them. 15 To tune the level of accuracy, our design checks the number of consecutive 0's and 1's on the first N bits of both input 16 mantissas. The larger sequence of continuous bit, the higher accuracy CMUL multiplication can achieve. In order to use 17 CMUL for deep neural network (DNN) acceleration, we propose a framework which modifies the trained DNN model to 18 make it suitable for approximate hardware. Our framework adjusts the DNN weights to a set of "potential weights" that are suitable for approximate hardware. Then, it compensates the possible quality loss by iterative retraining the network 19based on the existing constraints. We evaluate the efficiency of the proposed approach on AMD GPU architecture by 20 replacing the conventional FPUs with the proposed CMUL. Our evaluations on four DNN applications show that, CMUL 21can achieve on average 60.3% energy efficiency and 3.2× energy-delay product (EDP) improvement as compared to the 22 baseline GPU, while ensuring less than 0.2% quality loss. These results are 38.7% and $2.0 \times$ higher energy efficiency and EDP improvement of the CMUL without using the proposed framework. 23

The rest of paper is organized as follow: Section 2 and Section 3 review the related work and background. Section 4 describes the proposed approximate multiplications. Section 6 describes the supported framework to accelerate neural network applications on approximate hardware. The hardware integration has been described in section 5. The experimental results are presented in Section 7. Finally, Section 8 concludes the paper.

2 RELATED WORK

2.1 Approximate Computing

There are several commonly examined approaches to approximate computing: voltage over scaling (VOS), use of approximate hardware blocks, and use of approximate memory units. VOS involves dynamically reducing the voltage supplied to a hardware component to save energy, but at the expense of accuracy. Error rates for VOS can be modeled to determine the trade-off between energy and accuracy for applications, allowing voltage to be lowered until an error threshold is reached (Imani et al. 2017c; Krause and Polian 2011). However, the circuit is sensitive to any variations, and if the operating voltage of a circuit is decreased too far, timing errors begin to appear.

Another strategy is the application of Non-volatile memories (NVM) to create approximate memory units, for energy efficient storage and computing purposes (Gnawali et al. 2018; Imani et al. 2016d; Kim et al. 2015). In computing, the goal of this approach is to store common inputs and their corresponding outputs. This style of associative memory can retrieve the closest output for given inputs in order to reduce power consumption (Imani et al. 2016a,b; Peroni et al. 2019). This approach does not work well in applications without a large number of the redundant calculations. Associative memory can be integrated into FPUs reduce these redundancies.

Approximate hardware involves redesigning basic component blocks to save energy, at the cost of accurate output (Camus et al. 2016; Hashemi et al. 2015; Lin and Lin 2013; Liu et al. 2014). Liu *et al.* utilize approximate adders to create an energy efficient approximate multiplier (Liu et al. 2014). Hashemi *et al.* designed a multiplier that selects a reduced

42 Manuscript submitted to ACM

2

9 10

number of bits used in the multiplication to conserve power (Hashemi et al. 2015). Camus et al. propose a speculative approximate multiplier combines gate-level pruning and an inexact speculative adder to lower power consumption and 2 shrink FPU area (Camus et al. 2016). All the methods adopt to operation accuracy needed at runtime. They only have 3 one level of approximation that is independent of the inputs. In contrast to previous work, we design a configurable 4 approximate floating point multiplier which approximately processes data using an input mantissa directly in the output. 5 In addition, we propose a framework which fixes one of the multiplication operands in neural network in order to

significantly reduce the error of approximate hardware. 6

7 2.2 Neural Network

8 Modern neural network algorithms are executed on diverse types of processors such as GPU, FPGAs, and ASIC chips (Ciresan et al. 2011; Han et al. 2016; Iandola et al. 2016; Imani et al. 2017b; Nazemi et al. 2018; Razlighi et al. 9 2017). Prior work tried to use fixed-point quantized numerals to improve the efficiency of DNN (Lin et al. 2016). Work in (Lin et al. 2015) exploited trained binary parameters to avoid multiplication. However, many applications require 11 floating-point precision since the iterative DNN training algorithm often update the parameters using gradients whose 12 values are too small to sustain the additive quantization noise (Lin and Talathi 2016). In contrast, our proposed design 13uses floating-point precision rather than confining the parameters to binary numerals.

Other efficient way to improve the DNN efficiency is model compression. For example, work in (Han et al. 2015) 14 trained sparse models with shared weights to compress the model. The compressed parameters of (Han et al. 2015) 15can be used to realize ASIC/FPGA accelerators (Han et al. 2016). However, compression does not help with execution 16 on general purpose processors, in which case the compressed parameters should be decompressed into the original 17 parameters. Dimensionality reduction is investigated for efficient execution of DNNs (Imani et al. 2018a). These methods

are orthogonal to our proposed CMUL, since CMUL only reduces the cost of hardware computation with minimal impact 18 of quality loss. 19

20 **3 DEEP NEURAL NETWORKS**

A DNN model consists of multiple layers which have multiple neurons. These layers are stacked on top of each other in a 21hierarchical formation, that is, the output of each layer is forwarded to the next layer. The output of the last layer is used 22 for inference. Figure 2 shows the structure of a fully connected layer in a neural network. The computation in a single 23 layer of neural network can be modeled as a vector-matrix multiplication, which involves large amount of multiplications. $\mathbf{24}$ However, floating point operations are costly and energy hungry. Multiplication is the most commonly used floating point 25operation in both learning and multimedia applications (Han and Orshansky 2013; Hashemi et al. 2015). For example, looking at image filters such as Sobel filter, Robert filter, we observed that about 85% of floating point arithmetic involve 26 multiplication. The neuron takes a vector of neuron values from the preceding layer $\mathbf{X} = \langle X_0, \cdots, X_n \rangle$, then computes its 27 output as follows: 28

$$Z_j = \varphi(\sum_{i=1}^n W_{ij} \dot{X}_i + b)$$

where W_i and X_i correspond to a weight and an input respectively, b is a bias parameter, and φ is a nonlinear activation 32 function. 33

343.1 IEEE 754 Standard

35In floating point notation, a number consists of three parts: a sign bit, an exponent, and a fractional value. In IEEE 754 floating point representation, the sign bit is the most significant bit, where bits 31 to 24 hold the exponent value, and 36 the remaining bits contain the fractional value, also known as the mantissa. The exponent bits represent a power of two 37 ranging from -127 to 128. The mantissa bits store a value between 1 and 2, which is multiplied by 2^{exp} to give the decimal 38 value.

39 FPU multiply follows the steps shown in Figure 1. First, the sign bit of $A \times B = C$ is calculated by XORing the sign 40 bit of the A and B operands. Second, the effective value of the exponential terms are added together. Finally, the two mantissa values are multiplied to provide the result's mantissa. Because the mantissa ranges from 1 to 2, the output of the 41 Manuscript submitted to ACM

42



4

increasing the exponent by 1. 25

26 3.2 Limitations

Recently, work in (Imani et al. 2017a) proposed a configurable floating point multiplier (CFPU), which adaptively 27 multiplies the input operands. CFPU decides to run the multiplication in exact or approximate mode depending on the 28 input mantissas. However, CFPU relies on one input to approximate which results in errors that range from 0% to 50% 29 works in approximate mode only when one of the input mantissa has N leading one or zero bits (N is a tuning bits). This 30 reduces the number of inputs that CFPU can process in approximate mode. In section 4, we explain the functionality 31 of the proposed approximate multiplier, then in Section 6, we explain the framework this approximate multiplier to accelerate DNN on GPU architecture. 32

33

4 PROPOSED APPROXIMATE MULTIPLIER

34 In floating point multiplication, the mantissa multiplication is the most costly operation which takes about 80% of the 35 total multiplication energy (Imani et al. 2017a). Here we propose CMUL to accelerate floating point multiplication by eliminating the costly mantissa multiplication. Our design XORs two input sign bits to get the output sign bit. The two 36 input exponents are added to calculate the output exponent. Finally, instead of multiplying the two mantissas, we add 37 two mantissas and used the result as the mantissa for the output. The result shows that when we replaced the mantissa 38 multiplication with addition, the error rate is less than or equal to 11.1%. Figure 4 shows the error distribution of 1 million 39 random approximations executed by CMUL and CFPU. The result shows that CMUL has higher accuracy than the CFPU 40 with no tuning. Due to a lower error rate without tuning, our design is able to approximate a large amount of numbers, resulting in speedup and energy efficiency improvement. 41



Fig. 4. Histogram of error distribution for proposed design and CFPU.



Fig. 5. Accuracy distribution of proposed design as the # of consecutive 1's or 0's changes.

23 4.1 Tuning Accuracy

2

3 4

6

9

11

12

13

14 15

17 18

19

20

21 22

Although proposed approximate multiplication provides high energy savings, the accuracy of computation is heavily impacted depending on the application. For some applications, with quantized inputs, e.g., *Sharpen filter*, the proposed design can work precisely with no quality loss. In addition, many recognition algorithms, such as motion tracking and plate detection applications, only need to quantify changes in the input data. Therefore, the approximate multiplication can be nearly exact for such applications.

In order to ensure the desired accuracy is achieved we design a tuning method that allows the design to operate only when the approximation is at the desired error rate. The tuning process consists of checking the *N* number of consecutive 0's or 1's in both of the input mantissas, if one of the inputs has the minimum required *N* value the design will operate in approximation mode. Figure 5 shows the error distribution of random approximations as the value of *N* changes. The data shows that as *N* increases exponentially the error rate decreases as the number of consecutive 1's or 0's found in one of the input mantissa.

In order to show the level of accuracy that can be achieved with the proposed design, random inputs with different N33 values were generated and input into the CFPU and CMUL in order to compare maximum and average error rates. Figure 6 34 shows that as N increases the error rate goes to zero for both maximum error and average error rate for both designs. The 35 data also shows that the proposed design is far better in both average error rate and maximum error. Comparing both designs, the CFPU has a larger maximum and average error rate for low N values whereas the proposed design has a 36 significantly lower maximum and average error rate for low N values. This is significant since the lower the N value the 37 more inputs the approximation design can approximate. Thus the proposed design can handle a greater number of inputs 38 than the CFPU with higher accuracy which will result in less energy and higher speeds since more multiplications would 39 be able to be approximated. 40

An example of CMUL multiplication is shown in Figure 7 for two 32-bit floating point numbers in precise FPU and proposed CMUL with N = 5. The conventional FPU finds the product of A = -65 and B = 10 by adding the exponents Manuscript submitted to ACM



23 24

and multiplying the two mantissa, while XORing the sign bit to find three parts of the output data. In contrast, our design 25first checks both of the input mantissas for N consecutive 0's or 1's, if one of the mantissas contains the desired or exceeds 26 the desired number of consecutive 0's or 1's. In this example, since the input operand mantissa of A has a leading zero in the mantissa the N number of consecutive 0's is checked. In this example N = 5 and the mantissa of A also has five 27 consecutive 0's the design will proceed with the approximation. However, if the set value of N was larger then five, the 28 design would run the exact mode instead of approximation. In application N is selected based on the maximum error 29 rate the application can tolerate, with accuracy increasing with higher N value, however if a lower or higher error rate is 30 required N can be changed accordingly. When one of the input operands meets the tuning condition, the multiplication processes in approximate mode. In the example shown in Figure 7, the approximation results is -648, while the exact 31 multiplication gets -650. If a higher accuracy is desired, increasing the value of N would allow the design to only 32 approximate values that are under a certain threshold. 33

34 5 CMUL INTEGRATION

35 5.1 AMD GPU Architecture

We integrated the proposed CMUL in a GPU southern Island architecture, Radeon HD 7970 device. The architecture of GPU has been shown in Figure 8. This GPU has 32 compute units, where each contains a scheduler and a set of four SIMD execution units. Each SIMD execution unit has 16 cores, which gives a total number of 64 cores per compute unit. Each streaming core consists of both integer and floating point units. We replace all multipliers in floating point multiplier (MUL), multiply-accumulator (MAC) and multiplier-addition (MAD) units with the proposed CMUL. Every time an application launches, all GPU cores are configured as approximation level. CMUL is a modified version of the standard floating point multiplier in GPUs which uses hardware modification to support approximation.



Fig. 8. (a) The architecture of AMD Radeon HD 7970 GPU. (b) circuitry to support tuning the level of approximation in CMUL.



Fig. 9. Framework to support tunable approximation.

23 5.2 CMUL Hardware Support

Figure 8b shows the circuitry to support CMUL accuracy tuning. Our design looks at the first N mantissa bits of both input 24operands to check the tuning condition. If the tuning condition is satisfied in either input mantissas, our design adds the 25mantissa of the input operands to generate the mantissa of the multiplication output. Otherwise, similar to conventional 26 FPUs the multiplication of the input operand mantissas generates the output multiplication mantissa. Similarly, to tune 27 the level of approximation, our design uses N bits (after the first mantissa bit) of the selected mantissa to decide when 28 to perform mantissa multiplication or approximate it. The number of tuning bits sets the level of approximation, with each additional bit reducing the maximum error by half. The goal is to check the value of the $A_{i-1}, ..., A_{i-N}$ to make 29 sure they are the same. As Figure 8b shows, the tuning circuitry is a simple transistor-resistor circuitry which samples 30 the match-line (ML) voltage to detect the $A_{i-1}, A_{i-2}, \dots, A_0$ input operand. In case of any 1-bit in a mantissa, the sense 31 amplifier will detect changes in the ML voltage (ML=1). The circuitry also needs to select the inverted values of the tuning bits for the circuitry to search. To detect the 1 bit on A_{i-1}^{th} , ..., A_0^{th} indices on CMUL, the sense amplifier Clk needs 32 to be set to 250ps. Based on the results, we can dynamically change the sampling time to balance the ratio of the running 33 input workload on the approximate CMUL core. For each application, this sampling time can individually set in order to 34 provide target accuracy.

For DNN application, CMUL hardware support do not need to use tuning circuitry, since the software framework always ensures that the DNN weights satisfy the tuning condition. Therefore, CMUL always works in approximate mode and adds the mantissa of the input operands to generate the mantissa of the multiplication output. The conventional 32-bit floating point multiplier takes 7690 μm^2 area. In order to enable CMUL functionality, the conventional multiplier needs to use extra 23-bit fixed-point adder and a tuning circuit. Our evaluation using *Synopsys Design Compiler* shows that the adder and the tuning logic consumes 101.5 μm^2 and 28.3 μm^2 area respectively. Thus, the CMUL has a 1.68% larger area as compared to the conventional floating point units. This area overhead is negligible considering the flexibility and efficiency that the CMUL can provide.

Manuscript submitted to ACM

42

12 13

14

15

18 19

20



Fig. 10. The overview of the proposed framework in adjusting the DNN weights to a set which is suitable for approximate hardware.

We propose an automated framework to fine tune the level of approximation and satisfy required required accuracy while providing the maximum energy savings. Figure 9 shows the proposed framework consisting of the accuracy tuning and accuracy measurement blocks. The framework starts by putting CMUL in the maximum level of approximation. Then, based on the user accuracy requirement, it dynamically decreases the level of approximation until computation accuracy satisfies the user desired quality. For each application, this framework returns the optimal number of CMUL tuning which provides maximum energy and performance efficiency.

23 6 DNN ACCELERATION FRAMEWORK

24 6.1 Overview

16

In this section, we describe a novel framework to accelerate DNN applications on the approximate GPU architecture. As 25we explained, the enhanced GPU is configurable, thus it can be used in a similar way as other applications to accelerate 26 DNN. However, we observe that using this method, there are a few numbers to satisfy the tuning condition. On the 27 other hand, DNNs during inference use a set of fixed weight values. Our framework ensures that DNNs use a weight 28 representation which is suitable for our new approximate hardware. Adjusting the DNN weights ensure the proposed CMUL has minimum error rate when multiplying input and weights. Figure 10 shows the overview of the proposed 29 framework. In the first step, we get the DNN model by training the network (1). Our framework adjusts the weights by 30 quantizing them to a closest value which satisfies the tuning condition $(\mathbf{2})$. The adjusted model inputs the DNN and 31 the accuracy of the new network checks over the validation dataset (3). This accuracy is compared with the baseline trained model ($\Delta e = e_{Adjusted} - e_{Baseline}$). If the quality loss due to model adjustment is less than ε , we use the adjusted 32 model for the rest of classification (**⑤**), otherwise we retrain the network using the adjusted weights (**⑥**). This iterative 33 process continues until the error condition is met or the algorithm runs for a pre-specified number of epochs. Note that 34 the retraining approach is general and it improves the classification accuracy regardless of approximate hardware. In other 35 words, the retraining framework adapts network to work with the existing constraints. In the following we explain the 36 details of the proposed framework.

³⁷ 6.2 Weight Modification

The DNN computation involves many of multiplications between the input vector and weight matrix. These multiplications can be accelerated by processing on approximate hardware. However, the error of the approximate hardware, described in Section 4, depends on the input operand values. As we showed in section 4, the multiplication of input and weight elements has low error rate when one of the input operands have a specific representation. In particular, when one of the



Fig. 11. An example of speech recognition accuracy during different retraining iterations (N = 5).

mantissas starts with a continues sequence of 0s or 1s, our approximate multiplication results in much lower error rate.
 The upper bound of the multiplication error rate can be controlled depending on the length of the sequence.

Here, we use the idea of weight modification to adjust the DNN weights such that they become suitable for underlying 15approximate hardware. The DNN training gives us weights which do not usually have our desired pattern. We modify 16 the trained DNN model to force the weights to follow a particular pattern. Our framework first generates a list of all "potential weights" which are suitable for approximate hardware. These numbers are all floating point values which have 17 N consecutive 0s or 1s in the start of their mantissa. Our framework looks at each trained weight in neural network and 18 assigns it to a closet value in *potential weights* list (2). In case, if the potential weights include large number of values 19(small N), there will be very small change in each DNN weight, so DNN model may work with the same accuracy as original DNN model. However, weights with small N, run on approximate hardware with larger error. Using a potential 20 weight with large N, the approximate hardware will have significantly low error rate. However, the modified weights will 21be far from the original DNN weights, so it will result in larger change in DNN accuracy. In fact, there is a trade-off 22 between hardware and software in enabling approximation. Our framework enables software approximation by limiting 23 the values that DNN weights can take. The more limitation on the weights to get patterns with large N, it results in higher software approximation. However, this reduces the level of approximation in hardware, as each multiplication can 24perform with lower error. To compensate for the software approximation error, a retraining of the network is done and 25then the optimal N value is selected such that the total hardware+software approximation error is minimized. 26

27 6.3 Error Compensation

2

3 4

6

9

11 12

42

Limiting the weight is often accompanied by some degree of additive error, $\Delta e = e_{Ad \ iusted} - e_{Baseline}(\textcircled{S})$. This error is a 28 difference of the DNN accuracy using baseline and adjusted model. After each model adjustment iteration, our framework 29 compares the Δe with the ε value (**4**). If the condition is not satisfied, our framework retrains the neural network to 30 find a new model adopted with the current constraints. After each retraining iteration, all DNN weights again map 31 to a closest value in "potential weight" list. This process continues for several iterations until Δe is less than ε or the number of iteration passes the maximum epochs (). Figure 11 shows an example of speech recognition (Dheeru and 32 Karra Taniskidou 2017) accuracy during retraining iterations when the mantissa of the DNN weights are forced to start 33 with N = 4 consecutive 0s or 1s. Our result shows that in the first iterations, the weight limitation has significantly impact 34 on the classification accuracy. However, our framework can completely compensate the possible quality loss by retraining 35 the network for several iterations ($\Delta e = 0\%$). Table 1 shows the error of different DNN applications, when we limit the *potential weights* to values which satisfy 36

required approximation, specified by user. We tested the impact of our framework on four different DNN applications including: handwritten digits recognition (MNIST) (LeCun et al. 1998), speech recognition (ISOLET) (Dheeru and Karra Taniskidou 2017), activity recognition (UCIHAR) (Anguita et al. 2013), and object recognition (CIFAR-10) (Krizhevsky and Hinton 2009). Our evaluation shows that for application such as MNIST and ISOLET, our framework can compensate the quality loss when using weights with *N* equal or less than 4. However, for applications such as CIFAR, 0% quality loss can be achieved using weight with N = 2.

Table 1. Error loss of different applications when the weight are adjusted to a list with a defined N tuning condition.

Ν	1	2	3	4	5	6	7
MNIST	0	0	0	0	0.40%	0.89%	1.93%
ISOLET	0	0	~0%	0.06%	0.53%	1.71%	2.86%
UCIHAR	0	0	0.07%	0.33%	0.42%	1.16%	2.49%
CIFAR-10	0	~0%	0.10%	0.45%	0.97%	2.33%	3.21%

Table 2. DNN models and baseline error rates for 4 applications (Input layer - IN, Fully connected layer - FC, Convolution layer - C, and Pooling layer - PL.)

Dataset	Network Topology	Error
MNIST	<i>IN</i> : 784, <i>FC</i> : 512, <i>FC</i> : 512, <i>FC</i> : 10	1.5%
ISOLET	<i>IN</i> : 617, <i>FC</i> : 512, <i>FC</i> : 512, <i>FC</i> : 26	3.6%
UCIHAR	<i>IN</i> : 561, <i>FC</i> : 512, <i>FC</i> : 512, <i>FC</i> : 19	1.7%
	$IN: 32 \times 32 \times 3, CV: 32 \times 3 \times 3, PL: 2 \times 2,$	
CIFAR-10	$CV: 64 \times 3 \times 3, CV: 64 \times 3 \times 3, FC: 512, FC: 10\ 100$	14.4%

Table 3. Comparing the energy, and performance of the CMUL and previous designs.

	Power(mW)	Delay(ns)	EDP (pJs)	Tuned Error	Tunable	No Tuning Error
CMUL 3	0.15	1.1	0.18	6.3%	Yes	11.1%
CFPU3 (Imani et al. 2017a)	0.17	1.6	0.44	6.3%	Yes	50%
DRUM6 (Hashemi et al. 2015)	0.29	1.9	1.04	6.3%	No	NA
ESSM8 (Narayanamoorthy et al. 2015)	0.28	2.1	1.2	11.1%	No	NA
Kulkarni (Kulkarni et al. 2011)	0.82	3.5	10.0	22.2%	No	NA

22 23 24

7 RESULTS

²⁵ 7.1 Experimental Setup

We integrated the proposed approximate CMUL on the floating point units of an AMD Southern Island GPU, Radeon HD 7970 device. We modified Multi2sim, a cycle accurate CPU-GPU simulator (Ubal et al. 2012) to model the CMUL functionality in three main floating point units in GPU architecture: multiplier, multiplier-accumulator (MAC) and Multiplier-then-adder (MAD). We evaluated power of conventional FPUs using Synopsys Design Compiler and optimized for power using Synopsys Prime Time for 1ns delay in 45-nm ASIC flow (Compiler 2000). The circuit level simulation of CMUL has been performed using HSPICE simulator in 45-nm TSMC technology. We test the efficiency of enhanced GPU on eleven general Open*CL* applications: *Sobel, Robert, Mean, Laplacian, Sharpen, Prewit, QuasiRandom, FFT, Mersenne, DwHaar1D* and *Blur*. In these applications, roughly 85% of the floating point operations involve multiplication.

³³ 7.2 Benchmarks and DNN Models

Table 2 lists the baseline neural network topologies running four applications and their error rates for train and test modes. For all four datasets, we compare the baseline accuracy of the train and inference phases with those when using the proposed CMUL framework. We compare the designs in terms of run time and power consumption. Stochastic gradient descent with momentum (Sutskever et al. 2013) is used for training. The momentum is set to 0.1, the learning rate is set to 0.001, and a batch size of 10 is used. Dropout (Srivastava et al. 2014) with drop rate of 0.5 is applied to hidden layers to avoid over-fitting. The activation functions are set to "Rectified Linear Unit" clamped at 6. A "Softmax" function is applied to the output layer.

Handwritten Image Recognition (MNIST): MNIST is a popular machine learning data set including images of handwritten digits (LeCun et al. 1998). The objective is to classify an input picture as one of the ten digits $\{0...9\}$.

42 Manuscript submitted to ACM

10

8

Transform hite	Sobel		Robert		Mean		Laplacian		FFT		Mersenne		DwtHaar1D		Blur	
Tuning bits	EDP	QL	EDP	QL	EDP	QL	EDP	QL	EDP	QL	EDP	QL	EDP	QL	EDP	QL
1 bit	0.11	2.43%	0.13	0.45%	0.15	0.27%	0.17	0.37%	0.11	9.18%	0.15	4.29%	0.14	11.09%	0.17	6.24%
2 bit	0.14	1.13%	0.15	0.17%	0.16	0.14%	0.18	0.21%	0.28	5.19%	0.23	2.37%	0.17	8.2%	0.26	2.93%
3 bits	0.16	0.21%	0.16	0.06%	0.19	0.03%	0.19	0.02%	0.37	3.1%	0.31	1.9%	0.25	4.1%	0.37	0.03%
4 bits	0.17	0.01%	0.17	~0%	0.23	~0%	0.20	0.01%	0.41	1.07%	0.36	0.62%	0.29	1.98%	0.42	0.09%
5 bits	0.18	~0%	0.17	~0%	0.25	~0%	0.21	~0%	0.46	0.43%	0.44	0.11%	0.36	0.30%	0.51	0.02%

Table 4. Normalized Energy-delay product (EDP) and quality loss (QL) of the GPU enhanced with CFPU in different tuning mode.

Fable 5.	Normalized	EDP a	ınd QL	of the G	PU enha	nced with	CMUL	in	different	tuning	mode.
----------	------------	-------	--------	----------	---------	-----------	------	----	-----------	--------	-------

Tuning hits	Sobel		Robert		Mean		Laplacian		FFT		Mersenne		DwtHaar1D		Blur	
Tuning bits	EDP	QL	EDP	QL	EDP	QL	EDP	QL	EDP	QL	EDP	QL	EDP	QL	EDP	QL
1 bit	0.08	2.09%	0.11	0.35%	0.14	0.09%	0.09	0.37%	0.10	7.26%	0.12	3.02%	0.10	8.42%	0.11	4.36%
2 bit	0.12	0.94%	0.13	0.07%	0.13	0.06%	0.14	0.09%	0.22	3.56%	0.18	1.33%	0.13	5.77%	0.20	1.05%
3 bits	0.13	0.35%	0.14	0.01%	0.16	0.02%	0.15	0.01%	0.30	1.17%	0.24	0.63%	0.21	1.8%	0.33	0.01%
4 bits	0.13	0.02%	0.14	~0%	0.17	~0%	0.16	0%	0.36	0.41%	0.32	0.12%	0.25	0.24%	0.27	~0%
5 bits	0.15	~0%	0.15	~0%	0.21	~0%	0.19	~0%	0.42	0.14%	0.39	0.03%	0.31	0.12%	0.42	~0%

Voice Recognition (ISOLET): Many mobile applications require online processing of vocal data. We evaluate lookNN with the Isolet dataset (Dheeru and Karra Taniskidou 2017) which consists of speech collected from 150 speakers. The goal of this task is to classify the vocal signal to one of the 26 English letters.

Human Activity Recognition (UCIHAR): For this data set, the objective is to recognize human activity based on 3-axial linear acceleration and 3-axial angular velocity that have been captured at a constant rate of 50Hz (Anguita et al. 2013).

Object Recognition (CIFAR): CIFAR-10 (Krizhevsky and Hinton 2009) are two datasets each of which includes 50000 training and 10000 testing images belonging to 10 classes, respectively. The goal is to classify an input image to the correct category, e.g., animals, airplane, automobile, ship, truck, etc. For the two datasets, we exploit similar topologies based on convolution layers (CV), but they have different numbers of neurons in the last FC layer according to the number of classes.

7.3 Approximate Multipliers

2 3 4

6

8 9

17

18

19

20

21

22

23

24

25 26

27

To understand the advantage of proposed design, we compare the energy consumption and delay of the proposed CMUL with the state-of-the-art approximate multipliers proposed in (Hashemi et al. 2015; Imani et al. 2017a; Kulkarni et al. 2011; Narayanamoorthy et al. 2015). The application of previous designs limits to a small range of robust and error tolerant applications, as they are not able to tune the level of accuracy in runtime. In contrast, the proposed CMUL dynamically finds the inaccurate data and processes them in precise mode. CMUL tunes the level of accuracy at runtime based on the user accuracy requirement. This makes the application of CMUL general. It should be noted that the proposed framework is general and can work properly on other large scale datasets. For example, as prior work in (Imani et al. 2018c) showed, CMUL can get minimal quality loss even for larger datasets such as ImageNet.

Table 3 lists the power consumption, critical path delay, and energy-delay product of CMUL alongside previous work in (Imani et al. 2017a), (Hashemi et al. 2015), (Narayanamoorthy et al. 2015) and (Kulkarni et al. 2011) in their best configurations. Our evaluation shows that at the same level of accuracy, the proposed design can achieve 2.4× EDP improvement compared to the state-of-the-art approximate multipliers.

³⁸ 7.4 Tunable CMUL

We show the efficiency of the CFPU by running different multimedia and general streaming applications on the enhanced GPU architecture. We consider 10% average relative error as an acceptable accuracy metric for all applications, verified by (Esmaeilzadeh et al. 2012). We tune the level of approximation by checking the *N* bits of mantissa in the input Manuscript submitted to ACM



Fig. 12. Quality loss of different DNN applications due to software/framework and hardware approximation using different tuning
 bits (*N*).



19operands. If all N tuning bits in one of the input mantissa is 0 or 1, the multiplication runs in approximate mode, otherwise it runs precisely by multiplying the mantissa of input operands. For each application, Table 4 and Table 5 20 show the normalized energy-delay product (EDP) and quality loss of different applications running on approximate GPU 21enhanced by CPU and CMUL respectively. For both designs, we change the number of tuning bit from 1 (none) to 5 22 bits. The results are normalized to the EDP of the GPU using conventional FPUs. Increasing the number of tuning bits 23 improves the computation accuracy by processing the far and inaccurate multiplications in precise mode. On the other hand, more number of tuning bits slows down the computation, because a larger portion of data is processed on precise. 24 Our experimental evaluation shows that running applications on proposed CFPU provides 3.1× EDP improvement as 25compared to a GPU using conventional FPUs, while ensuring less than 1% quality of loss. Our results in Table 5 shows 26 that CMUL can achieve $2.7 \times$ higher EDP improvement as compared to CFPU design while providing the same quality of 27 computation.

28 29

7.5 CMUL & DNN Acceleration

In order to provide large efficiency, we design a framework which adapts DNN to run on approximate hardware. Using 30 our framework, the DNN quality loss may be happened by both software and hardware. Figure 12 shows the quality loss 31 different DNN applications running on proposed approximate hardware. The x-axis in figure shows the N, the sequence of the 0s and 1s at the mantissa of the weight. For example, N = 4 ensures that all DNN weights have four consecutive 32 Os or 1s in the first N bits their mantissas. The lines in the figure show the breakdown of quality loss coming from 33 software framework and proposed approximate hardware. The results show that increasing the N parameter from 1 to 6, 34 the software framework approximation starts increasing due to weight constraint applied by the framework. Using large 35 N, the DNN do not have good flexibility to assign proper weights to DNN, thus it results in large quality loss. From other hand, the quality loss due to hardware approximation has reverse relation to N value. Using large N, the approximate 36 multiplier can achieve lower error. Figure 12 also shows the total DNN quality loss due to both software and hardware 37 approximation. Our result shows that applications provide optimum quality loss using different N values. For example, 38 MNIST can achieve to minimum 0% quality loss using N = 3 and 4, while ISOLET can achieve to 0.05% quality loss 39 using N = 4.

One major advantage of our proposed framework is that CMUL does not need to check the *N* for each DNN application.
 Regardless of the *N* value, CMUL always takes the same time/energy to run a DNN application. Figure 13 and Figure 14

42 Manuscript submitted to ACM

12

1

4

5

6

8

9

11

12

13

14



result shows that CMUL with supported framework can achieve 60.3% and energy efficiency and $3.2 \times \text{EDP}$ improvement as compare to the baseline GPU, while they ensure less than 0.2% quality loss for tested applications. At the same level of accuracy, these results are 38.7% and $2.0 \times$ higher than energy efficiency and EDP improvement of the CMUL with no supported framework.

27 Table 6 compares the EDP improvement of the CMUL using the proposed DNN framework with CMUL and CFPU (Imani et al. 2017a) with no framework support. Our evaluation shows that CMUL using framework has less 28 than 0.5% quality loss over all applications. To provide the same quality of classification, CMUL and CFPU with no 29 framework support need to run the computation in a configuration very close to the precise mode, thus they do not provide 30 much advantage as compared to conventional GPU. In addition, CMUL using framework ensures that all multiplications 31 can run in approximate mode, this results in significantly performance improvement. In contrast, in CMUL and CFPU with no framework support, the slowest thread with the least number of multiplications in approximate mode, bounds 32 the GPU performance. To further improve the EDP of the CMUL and CFPU, one can put the multiplications in deeper 33 approximate mode. However, this results in significantly quality loss. The results in Table 6 shows that even with 2% 34 (4%) quality loss, CMUL and CFPU without framework support provide $2.5 \times$ and $2.8 \times (1.9 \times$ and $2.2 \times)$ lower EDP 35 improvement as compared to CMUL which ensures less than 0.5% quality loss.

8 CONCLUSION

36

37

In this paper, we propose a configurable floating point multiplier which can approximately perform the computation with significantly lower energy and performance cost. The proposed approximate multiplication has tuning capability by adaptively process each new piece of data precisely. We also proposed a framework to accelerate DNN applications with our approximate FPU. Our framework modifies the training of the DNN to make it suitable for underlying approximate hardware. Our evaluations on four DNN applications show that, CMUL can achieve 60.3% and energy efficiency and Manuscript submitted to ACM



 $1.06 \times$

 $1.15 \times$

 $1.07 \times$

 $1.08 \times$

22 23

24 25

26

33

 $3.2 \times$ energy-delay product (EDP) improvement as compared to the baseline GPU, while they ensure less than 0.2%27 quality loss as compared to precise hardware. These results are 38.7% and $2.0\times$ higher than energy efficiency and EDP 28 improvement of the CMUL without using the proposed framework. Another main advantage of the proposed framework is its generality, as it can be applied on any approximate multiplier. 29

 $1.02 \times$

 $1.12 \times$

 $1.03 \times$

 $1.04 \times$

 $1.25 \times$

 $1.18 \times$

 $1.29 \times$

 $1.28 \times$

 $1.12 \times$

1.06×

 $1.17 \times$

1.15×

 $1.56 \times$

 $1.45 \times$

 $1.62 \times$

 $1.60 \times$

 $1.40 \times$

 $1.30 \times$

 $1.46 \times$

 $1.44 \times$

30 ACKNOWLEDGMENTS

ISOLET

UCIHAR

CIFAR-10

3.53×

3.01×

 $2.68 \times$

31 This work was partially supported by CRISP, one of six centers in JUMP, an SRC program sponsored by DARPA, and 32 also NSF grants #1730158 and #1527034.

REFERENCES

- 34 Davide Anguita, Alessandro Ghio, Luca Oneto, Xavier Parra, and Jorge Luis Reyes-Ortiz. 2013. A Public Domain Dataset for Human Activity Recognition 35 using Smartphones.. In ESANN.
- Luigi Atzori, Antonio Iera, and Giacomo Morabito. 2010. The internet of things: A survey. Computer networks 54, 15 (2010), 2787-2805. 36 Vincent Camus, Jeremy Schlachter, Christian Enz, Michael Gautschi, and Frank K Gurkaynak. 2016. Approximate 32-bit floating-point unit design with

53% power-area product reduction. In European Solid-State Circuits Conference, ESSCIRC Conference 2016: 42nd. Ieee, 465-468. 37 Dan C Ciresan, Ueli Meier, Jonathan Masci, Luca Maria Gambardella, and Jürgen Schmidhuber. 2011. Flexible, high performance convolutional neural

- networks for image classification. In IJCAI Proceedings-International Joint Conference on Artificial Intelligence, Vol. 22. Barcelona, Spain, 1237. 38 Design Compiler. 2000. Synopsys inc.
- 39 Matthieu Courbariaux, Jean-Pierre David, and Yoshua Bengio. 2014. Low precision storage for deep learning. arXiv preprint Arxiv: 1412.7024 (2014). Dua Dheeru and Efi Karra Taniskidou. 2017. UCI Machine Learning Repository. http://archive.ics.uci.edu/ml 40
- Hadi Esmaeilzadeh, Adrian Sampson, Luis Ceze, and Doug Burger. 2012. Neural acceleration for general-purpose approximate programs. In Proceedings of the 2012 45th Annual IEEE/ACM International Symposium on Microarchitecture. IEEE Computer Society, 449-460. 41

Manuscript submitted to ACM 42

John Gantz and David Reinsel. 2011. Extracting value from chaos. IDC iview 1142, 2011 (2011), 1–12.

- Self Prasad Gnawali, Seyed Nima Mozaffari, and Spyros Tragoudas. 2018. Low Power Spintronic Ternary Content Addressable Memory. IEEE Transactions on Nanotechnology (2018).
- Jie Han and Michael Orshansky. 2013. Approximate computing: An emerging paradigm for energy-efficient design. In Test Symposium (ETS), 2013 18th IEEE European. IEEE, 1–6.
- 4 Song Han, Xingyu Liu, Huizi Mao, Jing Pu, Ardavan Pedram, Mark A Horowitz, and William J Dally. 2016. EIE: efficient inference engine on compressed deep neural network. In *Computer Architecture (ISCA)*, 2016 ACM/IEEE 43rd Annual International Symposium on. IEEE, 243–254.
 5 Security 100 (100)
 - Song Han, Huizi Mao, and William J Dally. 2015. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. arXiv preprint arXiv:1510.00149 (2015).
- ⁶ County, arXiv preprint arXiv 1510,0149 (2015).
 ⁷ Soheil Hashemi, R Bahar, and Sherief Reda. 2015. DRUM: A dynamic range unbiased multiplier for approximate applications. In *Proceedings of the IEEE/ACM International Conference on Computer-Aided Design*. IEEE Press, 418–425.
- Geoffrey Hinton, Li Deng, Dong Yu, George E Dahl, Abdel-rahman Mohamed, Navdeep Jaitly, Andrew Senior, Vincent Vanhoucke, Patrick Nguyen, Tara N Sainath, et al. 2012. Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *IEEE Signal Processing Magazine* 29, 6 (2012), 82–97.
- Forrest N Iandola, Matthew W Moskewicz, Khalid Ashraf, and Kurt Keutzer. 2016. Firecaffe: near-linear acceleration of deep neural network training on compute clusters. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2592–2600.
- Farhad Imani, Changqing Cheng, Ruimin Chen, and Hui Yang. 2018a. Nested Gaussian Process Modeling for High-Dimensional Data Imputation in
 Healthcare Systems. In *IISE 2018 Conference & Expo, Orlando, FL, May.* 19–22.
- 12 Mahdi Imani, Seyede Fatemeh Ghoreishi, and Ulisses M Braga-Neto. 2018b. Bayesian Control of Large MDPs with Unknown Dynamics in Data-Poor Environments. In Advances in Neural Information Processing Systems.
- 13 Mohsen Imani, Yeseong Kim, Abbas Rahimi, and Tajana Rosing. 2016a. Acam: Approximate computing based on adaptive associative memory with online learning. In *Proceedings of the 2016 International Symposium on Low Power Electronics and Design*. ACM, 162–167.
- Mohsen Imani, Pietro Mercati, and Tajana Rosing. 2016b. ReMAM: low energy resistive multi-stage associative memory for energy efficient computing. In *Quality Electronic Design (ISQED), 2016 17th International Symposium on*. IEEE, 101–106.
 Mohsen Imani, Shruti Patil and Tajana S Rosing. 2016c. MASC: Ultra-low energy multiple-access single-charge TCAM for approximate computing. In
- 5 Mohsen Imani, Shruti Patil, and Tajana S Rosing. 2016c. MASC: Ultra-low energy multiple-access single-charge TCAM for approximate computing. In Proceedings of the 2016 Conference on Design, Automation & Test in Europe. EDA Consortium, 373–378.
- Mohsen Imani, Daniel Peroni, Yeseong Kim, Abbas Rahimi, and Tajana Rosing. 2017b. Efficient neural network acceleration on gpgpu using content addressable memory. In 2017 Design, Automation & Test in Europe Conference & Exhibition (DATE). IEEE, 1026–1031.
 Mohsen Imani, Daniel Peroni, and Tajana Rosing. 2017a. CFPU: Configurable floating point multiplier for energy-efficient computing. In Design Automation
- 18 Conference (DAC), 2017 54th ACM/EDAC/IEEE. IEEE, 1–6.
 Mohsen Imani, Abbas Rahimi, Pietro Mercati, and Tajana Rosing. 2017c. Multi-stage tunable approximate search in resistive associative memory. IEEE
- Transactions on Multi-Scale Computing Systems (2017).
 Mohsen Imani, Abbas Rahimi, and Tajana S Rosing. 2016d. Resistive configurable associative memory for approximate computing. In Design, Automation
- 20 Monsen Imani, Abbas Kanimi, and Tajana S Kosing. 2016d. Resistive configuratie associative memory for approximate computing. In Design, Automation & Test in Europe Conference & Exhibition (DATE), 2016. IEEE, 1327–1332.
- 21 Mohsen Imani, Mohammad Samragh, Yeseong Kim, Saransh Gupta, Farinaz Koushanfar, and Tajana Rosing. 2018c. RAPIDNN: In-Memory Deep Neural Network Acceleration Framework. arXiv preprint arXiv:1806.05794 (2018).
- Changqing Ji, Yu Li, Wenming Qiu, Uchechukwu Awada, and Keqiu Li. 2012. Big data processing in cloud computing environments. In *Pervasive Systems*, Algorithms and Networks (ISPAN), 2012 12th International Symposium on. IEEE, 17–23.
- Navid Khoshavi, Xunchao Chen, Jun Wang, and Ronald F DeMara. 2016. Read-Tuned STT-RAM and eDRAM Cache Hierarchies for Throughput and Energy Enhancement. arXiv preprint arXiv:1607.08086 (2016).
- Yeseong Kim et al. 2015. CAUSE: critical application usage-aware memory system using non-volatile memory for mobile devices. In *Computer-Aided Design (ICCAD), 2015 IEEE/ACM International Conference on.* IEEE, 690–696.
- Philipp Klaus Krause and Ilia Polian. 2011. Adaptive voltage over-scaling for resilient applications. In Design, Automation & Test in Europe Conference & Exhibition (DATE), 2011. IEEE, 1–6.
- Alex Krizhevsky and Geoffrey Hinton. 2009. Learning multiple layers of features from tiny images. (2009).
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. In Advances in neural information processing systems. 1097–1105.
- Parag Kulkarni, Puneet Gupta, and Milos Ercegovac. 2011. Trading accuracy for power with an underdesigned multiplier architecture. In VLSI Design (VLSI Design), 2011 24th International Conference on. IEEE, 346–351.
- Yann LeCun, Corinna Cortes, and Christopher JC Burges. 1998. The MNIST database of handwritten digits.
- Yann LeCun, Koray Kavukcuoglu, and Clément Farabet. 2010. Convolutional networks and applications in vision. In *Circuits and Systems (ISCAS)*, Proceedings of 2010 IEEE International Symposium on. IEEE, 253–256.
- Jian Liang, Russell Tessier, and Oskar Mencer. 2003. Floating point unit generation and evaluation for FPGAs. In *Field-Programmable Custom Computing* Machines, 2003. FCCM 2003. 11th Annual IEEE Symposium on. IEEE, 185–194.
- Chia-Hao Lin and Chao Lin. 2013. High accuracy approximate multiplier with error correction. In 2013 IEEE 31st International Conference on Computer Design (ICCD). IEEE, 33–38.
- Barryl Lin, Sachin Talathi, and Sreekanth Annapureddy. 2016. Fixed point quantization of deep convolutional networks. In International Conference on Machine Learning. 2849–2858.
- 35 Darryl D Lin and Sachin S Talathi. 2016. Overcoming challenges in fixed point training of deep convolutional networks. arXiv preprint arXiv:1607.02241 (2016).
- Zhouhan Lin, Matthieu Courbariaux, Roland Memisevic, and Yoshua Bengio. 2015. Neural networks with few multiplications. arXiv preprint arXiv:1510.03009 (2015).
 Core Lin, Lie Han and Exhibitio Lombardi. 2014. A low power high performance approximate multiplicar with configurable particle area recovery. In
- ³⁷ Cong Liu, Jie Han, and Fabrizio Lombardi. 2014. A low-power, high-performance approximate multiplier with configurable partial error recovery. In
 ³⁸ Design, Automation and Test in Europe Conference and Exhibition (DATE), 2014. IEEE, 1–4.
- Srinivasan Narayanamoorthy, Hadi Asghari Moghaddam, Zhenhong Liu, Taejoon Park, and Nam Sung Kim. 2015. Energy-efficient approximate
 multiplication for digital signal processing and classification applications. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* 23, 6 (2015), 1180–1184.
- Mahdi Nazemi, Amir Erfan Eshratifar, and Massoud Pedram. 2018. A hardware-friendly algorithm for scalable training and deployment of dimensionality reduction models on FPGA. arXiv preprint arXiv:1801.04014 (2018).

42

1	Maxime Oquab, Leon Bottou, Ivan Laptev, and Josef Sivic. 2014. Learning and transferring mid-level image representations using convolutional neural networks. In Computer Vision and Pattern Recognition (CVPR). 2014 IEEE Conference on. IEEE. 1717–1724.
2	Daniel Peroni et al. 2019. ALook: Adaptive Lookup for GPGPU Acceleration. In <i>IEEE ASP-DAC</i> . IEEE, 1–7.
3	Mohammad Samragh Razlighi, Mohsen Imani, Farinaz Koushanfar, and Tajana Rosing. 2017. Looknn: Neural network with no multiplication. In 2017 Design, Automation & Test in Europe Conference & Exhibition (DATE). IEEE, 1775–1780.
4	Sahand Salamat et al. 2018. RNSnet: In-Memory Neural Network Acceleration Using Residue Number System. In <i>Rebooting Computing (ICRC), 2018</i>
5	Nitish Srivastava, Geoffrey E Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. <i>Journal of Machine Learning Research</i> 15, 1 (2014), 1929–1958.
6	Alexander Suhre, Furkan Keskin, Tulin Ersahin, Rengul Cetin-Atalay, Rashid Ansari, and A Enis Cetin. 2013. A multiplication-free framework for signal
7	processing and applications in biomedical image analysis. In Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on. IEEE, 1123–1127.
8	Ilya Sutskever, James Martens, George E Dahl, and Geoffrey E Hinton. 2013. On the importance of initialization and momentum in deep learning. <i>ICML (3)</i> 28 (2013) 1139–1147
9	Rafael Ubal, Byunghyun Jang, Perhaad Mistry, Dana Schaa, and David Kaeli. 2012. Multi2Sim: a simulation framework for CPU-GPU computing. In <i>Parallel Architectures and Compilation Techniques (PACT). 2012 21st International Conference on</i> . IEEE, 335–344.
10	
11	
12	
13	
14	
15	
16	
17	
18	
19	
20	
21	
22	
23	
24	
25	
26	
27	
28	
29	
30	
31	
32	
33	
34	
35	
36	
37	
38	
39	
40	
41	Manuscript submitted to ACM
42	•