

Received 28 January 2025, accepted 4 March 2025, date of publication 12 March 2025, date of current version 21 March 2025.

Digital Object Identifier 10.1109/ACCESS.2025.3550520

## RESEARCH ARTICLE

# Leveraging Highly Approximated Multipliers in DNN Inference

GEORGIOS ZERVAKIS<sup>1</sup>, (Member, IEEE), FABIO FRUSTACI<sup>2</sup>, (Senior Member, IEEE),  
OURANIA SPANTIDI<sup>3</sup>, IRAKLIS ANAGNOSTOPOULOS<sup>4</sup>, (Member, IEEE),  
HUSSAM AMROUCH<sup>5</sup>, (Member, IEEE), AND JÖRG HENKEL<sup>6</sup>, (Fellow, IEEE)

<sup>1</sup>Computer Engineering and Informatics Department, University of Patras, 265 04 Patras, Greece

<sup>2</sup>DIMES, University of Calabria, 87036 Rende, Italy

<sup>3</sup>Department of Computer Science, Eastern Michigan University, Ypsilanti, MI 48197, USA

<sup>4</sup>School of Electrical, Computer and Biomedical Engineering, Southern Illinois University, Carbondale, IL 62901, USA

<sup>5</sup>Chair of AI Processor Design, School of Computation, Information and Technology, Technical University of Munich (TUM) and Munich Institute of Robotics and Machine Intelligence, Technical University of Munich (TUM), 80992 Munich, Germany

<sup>6</sup>Chair for Embedded Systems, Karlsruhe Institute of Technology, 76131 Karlsruhe, Germany

Corresponding authors: Georgios Zervakis (zervakis@ceid.upatras.gr) and Hussam Amrouch (amrouch@tum.de)

This work was supported in part by the German Research Foundation (project ACCROSS).

**ABSTRACT** In this work, we present our control variate approximation technique that enables the exploitation of highly approximate multipliers in Deep Neural Network (DNN) accelerators. Our approach does not require retraining and significantly decreases the induced error due to approximate multiplications, improving the overall inference accuracy. As a result, control variate approximation enables satisfying tight accuracy loss constraints while boosting the power savings. Our experimental evaluation, across six different DNNs and several approximate multipliers, demonstrates the versatility of control variate technique and shows that compared to the accurate design, it achieves the same performance, 45% power reduction, and less than 1% average accuracy loss. Compared to the corresponding approximate designs without using our technique, the error-correction of the control variate method improves the accuracy by 1.9x on average.

**INDEX TERMS** Approximate computing, approximate multipliers, control variate, deep neural networks, error correction, low power.

## I. INTRODUCTION

Deep Neural Networks (DNNs) have become one of the main methodologies to enable artificial intelligence (AI) in many applications fields [1]. The promising results offered by DNNs derive from the processing of a huge amount of data that, in many cases where a real time response is crucial, prevents their software-based execution [1]. Customized hardware DNN accelerators meet the demand for high inference speed. This is particularly emphasized when the DNN is deployed into IoT edge devices, where the computation has to be performed locally with a reduced resources budget. Multiply-accumulate (MAC) operation is the most intensive computational task performed by a DNN during the inference phase. As an example, the most popular Convolutional Neural Networks (CNNs) perform millions of

MAC operations in their convolutional and fully-connected layers. In order to speed-up the inference phase, DNN accelerators are typically equipped with thousands of MAC units (e.g., 4K MAC units in the Google Edge TPU) operating in parallel, thus leading to high power requirements [2], [3]. Since the MAC operations are responsible for most of the power consumption, the research effort has focused on optimizing the multiplication, which is the most complex operation in a MAC.

Recently, approximate computing has emerged as a powerful design paradigm that relaxes the constraint of an exact computation in error-resilient applications, in order to trade the quality of the result with speed, area, and power consumption [4], [5], [6]. Due to their inherent error-tolerance, DNNs have become appropriate candidates for approximate computations [3], [4], [7], [8], [9], [10], [11], [12], [13]. However, approximating a DNN imposes several challenges. The layers within the same DNN can


The associate editor coordinating the review of this manuscript and approving it for publication was Alba Amato<sup>id</sup>.

have a significantly different error-resiliency [9], [10], [12]. Moreover, the errors due to approximate circuits are not constant but they are highly input dependent [14]. Finally, it has been shown that the deeper the neural network, the more sensitive it becomes to even slight approximation [9]. State-of-the-art applies retraining to mitigate the accuracy loss due to approximation [7], [8]. However, retraining may be infeasible in many cases because it is either time consuming or the training set might not be available (e.g., proprietary models) [10]. Many approximate multipliers have been proposed in the past [6], such as those based on column truncation [15], [16], [17], approximate compressors [18], [19], partial product perforation [20], and recursive multipliers with approximate blocks [21], [22]. Nevertheless, employing such multipliers in DNN inference, without retraining, results to unacceptable accuracy degradation [12], even when the slightest approximation is selected.

To enable effective exploitation of approximate multipliers in DNN inference and maximize power reduction, we propose a *control variate approximation* method [23]. Our technique improves the accuracy of approximate DNN accelerators by estimating and mitigating at runtime the error caused by approximate multiplications, without the need to perform time-overwhelming retraining. Leveraging the accuracy improvement of our control variate approximation, we are able to integrate highly approximate multipliers in DNN inference, and thus maximize the achieved power reduction. Our extensive experimentation demonstrates that our technique improves the inference accuracy by 1.9x on average, when compared to exactly the same approximate DNN accelerator without our proposed control variate approximation. This is an expanded version of our work [23].

*Extending our prior publication:* In our prior conference publication [23], control variate was bound to only a specific approximate multiplier [20]. In this work, we extend [23] and we show how our control variate approximation can be employed with a diversity of approximate multipliers, demonstrating the versatility of our method. Importantly, we demonstrate that by slightly adjusting accordingly our rigorous mathematical formulation of the induced convolution error, our approach can be applied to any approximate multiplier with predictable error. The latter refers to approximate multipliers that we know apriori if an error occurred and we can calculate/estimate it by a linear function. Finally, we show that the overheads of our runtime error correction are negligible, merely affecting the hardware gains due to approximation. Our control variate method demonstrates a consistent behavior across all the examined approximate multipliers, as it always significantly improves the inference accuracy while the hardware gains are solely determined by the approximate multiplier.

Overall, we evaluate the efficiency and versatility of our technique when considering the approximate perforated, recursive, and truncated multipliers. Such multipliers, apply aggressive approximation by inducing high error, but also achieving high power reduction. We designed approximate



**FIGURE 1.** Partial product reduction stages for: a) the accurate multiplier b) the approximate perforated multiplier with  $m = 3$  and  $s = 0$ .

DNN accelerators based on the developed control variate equations and we performed a complete comparison among the investigated approximate multipliers, demonstrating that our control variate method is able to mitigate most of the accuracy loss while maintaining high power gains.

## II. DESCRIPTION OF APPROXIMATE MULTIPLIERS

This section provides a brief description of three approximate multipliers that will be used in our work to generate approximate DNN accelerators and evaluate the efficiency of our control variate method in mitigating the error due to the approximate multiplications. In our analysis we consider partial product perforated multipliers [20], truncated multipliers [15], [16], [17], and approximate recursive multipliers [21], [22]. The examined circuits are representative paradigms of approximate multipliers and achieve very high power reduction, at the cost, however, of high error.

### A. APPROXIMATE PERFORATED MULTIPLIERS

The partial product perforation technique is based on the following principle. Let us consider the multiplication operation between two  $n$ -bit inputs,  $W$  and  $A$ . The accurate result  $W \cdot A$  is the sum of all the partial products, obtained by multiplying  $W$  by each bit  $a_i$  of  $A$ , with  $a_i = 0 \dots n - 1$ :

$$W \cdot A = \sum_{i=0}^{n-1} W \cdot a_i \cdot 2^i, \quad (1)$$

The partial product perforation technique approximates the multiplication operation by omitting  $m$  consecutive partial products starting from the  $s$ -th one, with  $s < n$  and  $m < n - s$ . Hence, the perforated product equals:

$$\text{AMP}(W, A) = \sum_{\substack{i=0, \\ i \notin [s, s+m)}}^{n-1} W \cdot a_i \cdot 2^i. \quad (2)$$

The authors in [20] deduced that when the distribution of one of the multiplicands is unknown,  $s = 0$  should be preferred. Hence, in our work we set  $s = 0$  and we examine varying values for  $m$ . The error of the Perforated Approximate Multiplication, when the  $m$  least partial products are perforated, is therefore calculated as follows:

$$\begin{aligned}\epsilon &= W \cdot A - \text{AMp}(W, A) \\ &= W \cdot A - W \cdot (A - A \bmod 2^m) \\ &= W \cdot p, \quad p = A \bmod 2^m\end{aligned}\quad (3)$$

As an example, Fig. 1 depicts the partial product reduction process of an unsigned  $8 \times 8$  multiplier without (left) and with (right) perforation. Each dot represents a partial product bit that needs to be accumulated. For the perforated multiplier  $m = 3$  (and  $s = 0$ ) is considered. Perforation removes bits from tree rows. Since  $m$  is 3 and the input size is 8, the red dots (eight LSBs in each of the first three rows ( $s = 0$ )) are removed, simplifying the accumulation tree. Removing the red dots and shifting the remaining black dots upward in each column transforms the exact accumulation tree (left) into the approximate one (right). The accumulation strategy remains unchanged. As shown, in Fig. 1, compared to the exact counterpart (i.e., without perforation), the perforated multiplier has shallower partial product reduction stages and, consequently, a lower number of compressors entailing lower energy consumption and potentially higher speed.


### B. APPROXIMATE RECURSIVE MULTIPLIERS

The principle of the Recursive Approximate Multiplier is illustrated in Fig. 2 (left). Overall, a recursive multiplier decomposes the large multiplication to smaller products and accumulates them to obtain the final result. Such a process is possibly iterated by dividing the obtained blocks into smaller ones. One way to trade the energy consumption with the accuracy of the result is to simplify the design of the smaller blocks, employed to calculate the least significant bits of the result, by approximating their logic function. In that way, the number of employed logic gates is reduced. As an example, [21] proposes to use an approximate  $2 \times 2$  multiplying block where only the multiplication between the binary inputs “11” and “11” is approximated to the inaccurate result “111”. Nevertheless, several partitionings can be employed for the decomposition [22].

In our work, we consider a recursive multiplier in which each input is divided into two sub-words: the low-part, composed of  $m$  bits, and the high-part, composed of  $n - m$  bits, with  $m < n$ . In this way, the  $n \times n$  multiplier is obtained by accumulating the four sub-products as follows:

$$W \cdot A = W_H \cdot A_H \cdot 2^{2m} + (W_H \cdot A_L + W_L \cdot A_H) \cdot 2^m + W_L \cdot A_L. \quad (4)$$

Applying coarse approximation, we generate approximate recursive multipliers by pruning the entire sub-product of the two low-parts (Fig. 2 (right)). Hence, the approximate product



**FIGURE 2.** The principle of the recursive approximate multiplier: composing a large multiplier by using smaller inaccurate building blocks.

is given by:

$$\text{AM}_R(W, A) = (W_H \cdot A_H \cdot 2^m + W_H \cdot A_L + W_L \cdot A_H) \cdot 2^m \quad (5)$$

and thus, the multiplication error when the size of the low-part is  $m$  bits, can be calculated as follows:

$$\begin{aligned}\epsilon &= W \cdot A - \text{AM}_R(W, A) \\ &= (W_H \cdot A_H \cdot 2^m + W_H \cdot A_L + W_L \cdot A_H) \cdot 2^m + W_L \cdot A_L \\ &\quad - (W_H \cdot A_H \cdot 2^m + W_H \cdot A_L + W_L \cdot A_H) \cdot 2^m \\ &= W_L \cdot A_L\end{aligned}\quad (6)$$

### C. APPROXIMATE TRUNCATED MULTIPLIERS

One of the most widely used approximation techniques is truncation [15], [16], [17]. The Truncated Approximate Multiplier is composed by pruning the hardware resources that form the  $m$  least significant columns of the multiplier. At the hardware level, the AND gates computing the bits  $w_i \cdot a_j$ , with  $i, j \in [0, m)$  and  $i + j < m$ , are not implemented in the partial product generation stage. Consequently, all the compressors belonging to the least-significant  $m$  columns of the partial product reduction stages are removed. The application of the truncation technique with  $m = 7$  on an unsigned  $8 \times 8$  multiplier is shown in Fig. 3. Truncation is one of the most widely used approximation techniques also due to its simplicity. For example, in Fig. 3, where  $m = 7$ , all dots in the seven LSB columns are truncated, i.e., removed from the accumulation tree. The corresponding product bits are set to zero. Removing the red dots and nullifying the output bits in these columns transforms the exact accumulation tree (left) into the approximate one (right). The reduced number of partial product bits entails a lower number of compressors w.r.t. the accurate multiplier, thus leading to lower area, lower energy consumption, and potentially lower delay. The approximate result of the Truncated Approximate Multiplier with  $m$  truncated columns can be obtained by:

$$\text{AM}_T(W, A) = \sum_{i=0}^{n-1} \sum_{j=\max(m-i, 0)}^{n-1} w_j \cdot a_i \cdot 2^{(i+j)} \quad (7)$$




FIGURE 3. The truncated multiplier with  $m = 7$ .

whereas its error is computed by:

$$\begin{aligned}
 \epsilon &= W \cdot A - \text{AM}_T(W, A) \\
 &= \sum_{i=0}^{n-1} \sum_{j=0}^{n-1} w_j \cdot a_i \cdot 2^{(i+j)} \\
 &\quad - \sum_{i=0}^{n-1} \sum_{j=\max(m-i, 0)}^{n-1} w_j \cdot a_i \cdot 2^{(i+j)} \\
 &= \sum_{i=0}^{m-1} (W \bmod 2^{m-i}) \cdot a_i \cdot 2^i \quad (8)
 \end{aligned}$$

#### D. ERROR ANALYSIS

The three typologies of approximate multipliers, analyzed above, have their own peculiarities in terms of accuracy. In Table 1, we assess the error characteristics of the examined approximate multipliers considering varying approximation levels (i.e.,  $m$  values) and 8-bit unsigned operands with uniform and normal distribution. For the normal distribution, we arbitrarily chose, without loss of generality, a mean value and a standard variation of 125 and 24, respectively. Table 1 collects the error results obtained for 1M input operands couples. Note that for all the approximate multipliers, as  $m$  increases the applied approximation increases. However, there is not a direct comparison between the  $m$  value of different approximate multipliers. As shown in Table 1, the three multipliers behave differently as their configuration knob ( $m$ ) varies. For example, the Approximate Perforated multiplier shows the highest error mean value ( $\mu$ ) and highest standard deviation ( $\sigma$ ) for both input distributions. The  $\mu$  and  $\sigma$  of the Truncated Approximate Multiplier are low and scale more gracefully compared to the other multipliers. Moreover, the latter multiplier also exhibits the lowest value of the  $\sigma/\mu$  ratio, thus resulting in being the multiplier with the lowest

TABLE 1. Error analysis on the examined approximate multipliers.

Approximate Perforated Multiplier					
Unif. Distr. $U(0, 255)$			Norm. Distr. $\mathcal{N}(125, 24^2)$		
$m$	$\mu$	$\sigma$	$m$	$\mu$	$\sigma$
1	63.7	82	1	62.4	64.7
2	191	198	2	187	146
3	447	425	3	435	302
Approximate Recursive Multiplier					
Unif. Distr. $U(0, 255)$			Norm. Distr. $\mathcal{N}(125, 24^2)$		
$m$	$\mu$	$\sigma$	$m$	$\mu$	$\sigma$
2	2.24	2.67	2	2.25	2.68
3	12.26	12.51	3	12.24	12.47
4	56	53.4	4	56.2	53.4
5	239	219	5	239	219
Approximate Truncated Multiplier					
Unif. Distr. $U(0, 255)$			Norm. Distr. $\mathcal{N}(125, 24^2)$		
$m$	$\mu$	$\sigma$	$m$	$\mu$	$\sigma$
4	12	9.9	4	12.6	9.9
5	32	23	5	32.2	23
6	80	52	6	80.6	52.8
7	192	115	7	192	127

coefficient of variation. Hence, the truncated multipliers feature the lowest error dispersion overall. Interestingly, the error performance of the truncated and the recursive approximate multipliers does not show a sensible variation as the distribution of the inputs changes. Indeed, the values of the error  $\mu$  and  $\sigma$  are practically the same for both distributions in Table 1.

#### III. CONTROL VARIATE APPROXIMATION

This section presents our control variate approximation technique and describes how it is applied with different approximate multipliers. The approximate multipliers of Section II are considered and our rigorous error analysis demonstrates how the control variate parameters can be tuned with respect to the multiplier used to minimize the error at convolution level.

The core operation of a convolution is given by:

$$G = B + \sum_{j=1}^k W_j \cdot A_j, \quad (9)$$

where  $B$  is the bias of the neuron,  $W_j$  are the weights, and  $A_j$  are the input activations.

Aiming for low-power operation, we replace all the accurate multipliers of the DNN accelerator with approximate ones. We denote  $\epsilon_j$  the multiplication error of the product  $W_j \cdot A_j$ . Hence,  $\epsilon_j$  equals the difference between the accurate and the approximate products:

$$\epsilon_j = W_j \cdot A_j - \text{AM}(W_j, A_j). \quad (10)$$

For example, assuming the approximate multipliers described in Section II,  $\epsilon_j$  can be computed using (2), (5), or (7).

Given (9) and (10), the convolution error  $\epsilon_G$  equals:

$$\begin{aligned}\epsilon_G &= B + \sum_{j=1}^k W_j \cdot A_j - B - \sum_{j=1}^k \text{AM}(W_j, A_j) \\ &= \sum_{j=1}^k \epsilon_j.\end{aligned}\quad (11)$$

The error value of an approximate multiplier can be considered as a random variable, and is therefore defined by its mean value and variance [24]. Denoting by  $\mu_{AM}$  and  $\sigma_{AM}^2$  the mean error and the error variance of the approximate multiplier AM, the mean and variance of the approximate convolution error operation are given by:

$$\begin{aligned}\mathbb{E}[\epsilon_G] &= \mathbb{E}\left[\sum_{j=1}^k \epsilon_j\right] = k\mu_{AM} \\ \text{Var}(\epsilon_G) &= \text{Var}\left(\sum_{j=1}^k \epsilon_j\right) = k\sigma_{AM}^2.\end{aligned}\quad (12)$$

Note that, the error values  $\epsilon_j$  are independent variables and thus their covariance is zero [9], [24].

Hence, even if the approximate multiplier features small error (small  $\mu_{AM}$  and  $\sigma_{AM}^2$ ), the convolution error is significantly higher since it is proportional to the filter's size as (12) demonstrates. In [9], approximate multipliers with systematic error are employed and a constant correction term is used to compensate for the mean error (i.e.,  $\mathbb{E}[\epsilon_G]$ ). However, even in this case, the error of the convolution is still high, since it is defined by its high variance ( $\text{Var}(\epsilon_G)$ ).

In our work, we propose the utilization of a control variate technique to reduce the convolution error. A control variate is an easily evaluated random variable, with known mean, that is highly correlated with our variable of interest. To implement our control variate approximation [23], to perform the convolution we compute (13) instead of (9).

$$G^* = B + \sum_{j=1}^k \text{AM}(W_j, A_j) + V, \quad (13)$$

where  $V$  is the control variate.

Inspired by the convolution error  $\epsilon_G$  in (11) and targeting low computational complexity, we express the control variate  $V$  as a first order polynomial:

$$\begin{aligned}V &= \sum_{j=1}^k v_j + C_0 \\ &= \sum_{j=1}^k x_j \cdot C_j + C_0, \quad v_j = x_j \cdot C_j.\end{aligned}\quad (14)$$

where  $C_j, \forall j \geq 0$ , are constants and  $x_j$  is an input-dependent variable (i.e., obtained at runtime). Obviously, setting  $v_j = \epsilon_j$  and  $C_0 = 0$  would deliver accurate results. Nevertheless, this would neglect any hardware gains originating by the approximate multiplications due to the high computational

complexity of precisely computing  $\epsilon_j$  at runtime. For example, assume that the perforated approximate multiplier [20] is used. Considering (3), if we set  $C_j = W_j, \forall j > 0, C_0 = 0$ , and  $x_j = p_j$ , then  $v_j = \epsilon_j$  and thus  $\epsilon_G = V$ . Hence, in (13) the error of the approximate multiplications is cancelled out by the control variate  $V$ , leading to accurate computation of the convolution operation. However, calculating  $p_j \cdot W_j$  is computationally expensive and neglects the gains (area, power) of the perforated multiplier, since calculating  $V$  requires  $k$  multiplications and  $k - 1$  additions, and computing  $p_j \cdot W_j$  requires the generation and addition of the partial products that were initially omitted.

However, since a control variate must be easily evaluated, we simplify  $V$  in (14) by setting  $C_j = C, \forall j > 1$ . As a result,  $V$  is given by:

$$V = C \cdot \sum_{j=1}^k x_j + C_0 \text{ and } v_j = x_j \cdot C. \quad (15)$$

Note that to calculate  $V$  in (15), only  $k - 1$  additions and 1 multiplication are required.

Given (15), the approximate convolution using our control variate method (13) is written as:

$$\begin{aligned}G^* &= B + \sum_{j=1}^k (W_j \cdot A_j - \epsilon_j) + \sum_{j=1}^k v_j + C_0 \\ &= G - \sum_{j=1}^k (\epsilon_j - v_j) + C_0\end{aligned}\quad (16)$$

and thus, the approximate convolution error equals:

$$\epsilon_{G^*} = G - G^* = \sum_{j=1}^k (\epsilon_j - v_j) - C_0. \quad (17)$$

#### A. CONTROL VARIATE WITH APPROXIMATE PERFORATED MULTIPLIERS

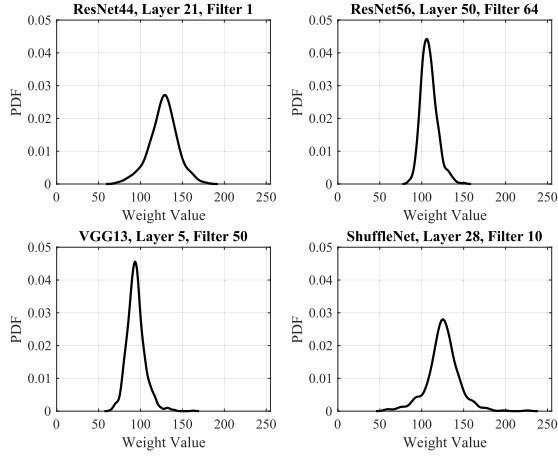
First, we examine the perforated multipliers [20] to perform the approximate multiplication, i.e.,  $\text{AM}(W_j, A_j) = \text{AMP}(W_j, A_j)$ . Considering the error of an approximate perforated multiplication that is given by (3), we set  $x_j$  in our control variate (15) equal to:

$$x_j = A_j \bmod 2^m = A_j \& (2^m - 1) \quad (18)$$

and thus, (17) becomes:

$$\begin{aligned}\epsilon_{G^*} &= \sum_{j=1}^k (\epsilon_j - v_j) - C_0 \\ &= \sum_{j=1}^k (x_j \cdot (W_j - C)) - C_0.\end{aligned}\quad (19)$$





**FIGURE 4.** Weight distribution of randomly selected filters of various NNs. Four examples are depicted. Figure obtained from [23].

Therefore, the variance  $\text{Var}(\epsilon_{G^*})$  of the error of the approximate convolution, i.e.,  $\epsilon_{G^*}$ , is calculated as:

$$\begin{aligned} \text{Var}(\epsilon_{G^*}) &= \sum_{j=1}^k \text{Var}(\epsilon_j - v_j) \\ &= \sum_{j=1}^k ((W_j - C)^2 \cdot \text{Var}(x_j)) \\ &= \underbrace{\frac{(2^m - 1)(2^m + 1)}{12}}_{\text{Var}(x_j)} \sum_{j=1}^k (W_j - C)^2. \end{aligned} \quad (20)$$

As a result,  $\text{Var}(\epsilon_{G^*})$  is minimized when:

$$\frac{d}{dC} \text{Var}(\epsilon_{G^*}) = 0 \Rightarrow C = E[W_j] = \frac{1}{k} \sum_{j=1}^k W_j. \quad (21)$$

Note that  $C \neq 0$ , i.e., variance without our control variate (as in (12)). In addition, note that the more squeezed the distribution of the weights is (i.e., concentrated close to  $E[W_j]$ ) the closer  $\text{Var}(\epsilon_{G^*})$  is to zero. Fig. 4, shows the distribution of weights for four different examples. In Fig. 4, the neural networks and the respective filters and layers, were randomly selected out of the neural networks we consider in Section V. Similar results are obtained for the rest of the filters and neural networks. As shown in Fig. 4, for all the examined filters, the majority of the weights is well concentrated in a closed region (squeezed dispersion in Fig. 4). Hence, this feature boosts the efficiency of our variance reduction method, as explained above.

Using the  $C$  value obtained in (21), that minimizes the variance, we compute the mean convolution error  $E[\epsilon_{G^*}]$ :

$$\begin{aligned} E[\epsilon_{G^*}] &= \sum_{j=1}^k E[\epsilon_j - v_j] - C_0 \\ &= \sum_{j=1}^k E[x_j] \cdot W_j - \sum_{j=1}^k E[x_j] \cdot E[W_j] - C_0 \end{aligned}$$

$$\begin{aligned} &= \underbrace{\frac{(2^m - 1)}{2}}_{E[x_j]} \left( \sum_{j=1}^k W_j - k \cdot E[W_j] \right) - C_0 \\ &= -C_0 \end{aligned} \quad (22)$$

Therefore, by setting  $C_0 = 0$ , (22) becomes zero. As a result, the proposed control variate approximation method with  $V = E[W_j] \sum_{j=1}^k x_j$ , effectively nullifies the mean error of the approximate convolution and also manages to decrease its variance. Hence, the error distribution is constrained in a squeezed region around zero and high convolution accuracy is expected. However, as (20) shows, the larger  $m$  is, the larger the error variance will be and thus the accuracy loss.

## B. CONTROL VARIATE WITH APPROXIMATE TRUNCATED MULTIPLIERS

Next, we examine the application of our proposed control variate approximation when the truncated multiplier is used for the approximate multiplication, i.e.,  $\text{AM}(W_j, A_j) = \text{AM}_T(W_j, A_j)$ . Although the error of a perforated multiplication can be easily obtained simply by the  $n \times m$  product, this is not the case for the truncated multipliers. As (7) shows, precise error estimation of a truncated multiplication requires  $m$  multiplications ( $m \times 1$  down to  $1 \times 1$ ) and  $m - 1$  additions. Hence, precise error estimation at run-time is very computationally expensive. On the other hand, as discussed in Section II, unlike the perforated multipliers, the truncated multipliers feature low error variance since it is constrained by the number of truncated columns (i.e.,  $m$ ). Leveraging the low error variance, we can efficiently estimate the truncated multiplication error by its average value. Considering (8), the average error of the truncated multiplication  $\text{AM}_T(W_j, A_j)$ ,  $\forall A_j$  is calculated by:

$$\begin{aligned} E[\text{AM}_T(W_j, A_j)] &= \sum_{i=0}^{m-1} E[(W_j \bmod 2^{m-i}) \cdot a_i \cdot 2^i] \\ &= \frac{1}{2} \sum_{i=0}^{m-1} (W_j \bmod 2^{m-i}) \cdot 2^i. \end{aligned} \quad (23)$$

Hence, by denoting

$$\hat{W}_j = \frac{1}{2} \sum_{i=0}^{m-1} (W_j \bmod 2^{m-i}) \cdot 2^i, \quad (24)$$

the error  $e_j$  of the truncated multiplication  $\text{AM}_T(W_j, A_j)$  is estimated by:

$$\begin{aligned} e_j &\approx \tilde{e}_j = x_j \cdot \hat{W}_j \\ \text{with } x_j &= (1 - \delta_{0,y_j}), \quad y_j = A_j \bmod 2^m \end{aligned} \quad (25)$$

where  $\delta_{0,y_j}$  is the Kronecker delta and thus  $x_j$  is easily calculated by a logic OR of the  $m$  LSB of  $A_j$ . In other words, if a multiplication error occurs,  $x_j$  is 1 while  $x_j$  is 0 when the error of the truncated multiplication is zero. The former results to  $\tilde{e}_j = \hat{W}_j$  and the latter to  $\tilde{e}_j = 0$ .

Given  $\tilde{e}_j$  in (25), we set  $V$  similarly to Section III-A:

$$\begin{aligned} V &= C \cdot \sum_{j=1}^k x_j + C_0, \quad v_j = x_j \cdot C, \\ x_j &= (1 - \delta_{0,y_j}) \text{ with } y_j = A_j \bmod 2^m, \\ C &= E[\hat{W}_j] = \frac{1}{k} \sum_{j=1}^k \hat{W}_j. \end{aligned} \quad (26)$$

Thus, in the case of the approximate truncated multiplier, the convolution error (17) becomes:

$$\begin{aligned} \epsilon_{G^*} &= \sum_{j=1}^k (\epsilon_j - v_j) - C_0 \\ &= \sum_{j=1}^k \left( \sum_{i=0}^{m-1} (W_j \bmod 2^{m-i}) \cdot a_i \cdot 2^i - x_j \cdot E[\hat{W}_j] \right) - C_0. \end{aligned} \quad (27)$$

Then, the mean convolution error  $E[\epsilon_{G^*}]$  is given by:

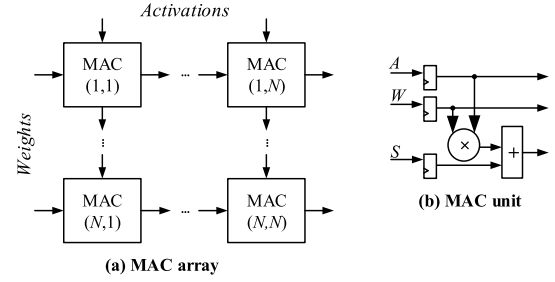
$$\begin{aligned} E[\epsilon_{G^*}] &= \sum_{j=1}^k E[\epsilon_j - v_j] - C_0 \\ &= \sum_{j=1}^k \left( \sum_{i=0}^{m-1} (W_j \bmod 2^{m-i}) \cdot 2^i \cdot E[a_i] \right) \\ &\quad - \sum_{j=1}^k E[x_j] \cdot E[\hat{W}_j] - C_0 \\ &= \sum_{j=1}^k \hat{W}_j - \underbrace{\frac{2^m - 1}{2^m}}_{E[x_j]} \sum_{j=1}^k \hat{W}_j - C_0 \\ &= \frac{1}{2^m} \sum_{j=1}^k \hat{W}_j - C_0 \end{aligned} \quad (28)$$

Therefore, by setting  $C_0 = \frac{1}{2^m} \sum_{j=1}^k \hat{W}_j$ , (28) becomes zero. As a result, in the case of the approximate truncated multiplier, the proposed control variate approximation with  $V = \sum_{j=1}^k E[\hat{W}_j](1 - \delta_{0,y_j}) + \frac{1}{2^m} \sum_{j=1}^k \hat{W}_j$  nullifies the mean error. Considering also that the error variance of the truncated multiplier is small, the convolution error variance will also be limited. Note that in this case  $C_0 \neq 0$ . However, the addition of  $C_0$  is performed with zero cost by just updating offline the bias value of the respective filter as in [9].

### C. CONTROL VARIATE WITH APPROXIMATE RECURSIVE MULTIPLIERS

Finally, we present the respective control variate analysis when the approximate recursive multipliers are employed, i.e.,  $AM(W_j, A_j) = AM_R(W_j, A_j)$ . Considering the error of an approximate perforated multiplication that is given by (6), we set  $x_j$  in our control variate (15) equal to:

$$x_j = A_j \bmod 2^m = A_j \& (2^m - 1) \quad (29)$$



**FIGURE 5.** The a) accurate systolic MAC array and b) MAC unit. Figure obtained from [23].

and thus, in the case of the approximate recursive multipliers (17) becomes:

$$\begin{aligned} \epsilon_{G^*} &= \sum_{j=1}^k (\epsilon_j - v_j) - C_0 \\ &= \sum_{j=1}^k (x_j \cdot ((W_j \bmod 2^m) - C)) - C_0. \end{aligned} \quad (30)$$

By denoting

$$W_j^m = W_j \bmod 2^m, \quad (31)$$

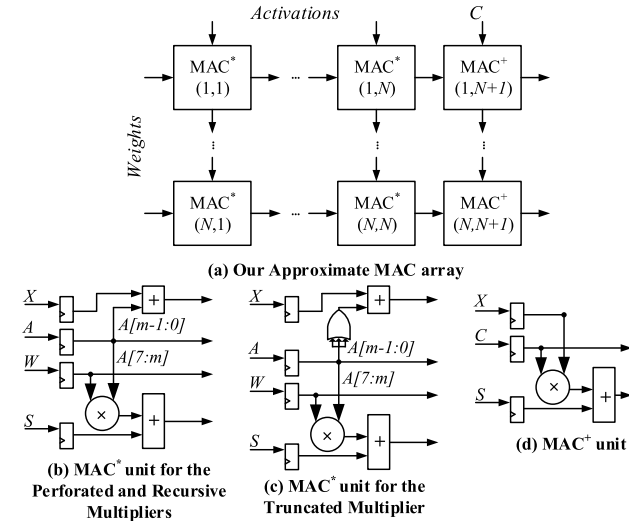
we set  $C = E[W_j^m] \cdot x_j$  and  $C_0 = 0$  in our control variate:

$$V = \sum_{j=1}^k E[W_j^m] \cdot x_j, \quad x_j = A_j \bmod 2^m. \quad (32)$$

Hence, the mean error of the approximate convolution ( $E(\epsilon_{G^*})$ ) is nullified and the variance of the approximate convolution error ( $\text{Var}(\epsilon_{G^*})$ ) is minimized. Proofs are identical to the perforated multiplier case and thus omitted.

### D. APPLICATION WITH OTHER APPROXIMATE MULTIPLIERS

As our analysis in Sections III-A to III-C demonstrates, our control variate technique can be effectively applied with three diverse approximate multipliers that exhibit varying error characteristics. The examined approximate multipliers apply high approximation leading to high power savings but also high error. Still, our rigorous error analysis shows that our method is able to mitigate the induced error at the convolution level. Our proposed control variate approximation is not limited to just the examined multipliers, but it can be employed with any approximate multiplier as long as its error can be expressed by an analytical model. Nevertheless, the cost-efficiency of our approach will depend on the complexity associated with computing the error model. However, as our analysis for the truncated multiplier demonstrates, our approach can still be efficiently applied, even in the case of a complex error model, as long as i) we can assess with low-cost if a multiplication error occurred and ii) the error of the approximate multiplier features low variance.



**FIGURE 6.** a) Our approximate systolic MAC array. b) MAC\* unit for the approximate perforated and recursive multipliers. c) MAC\* unit for the approximate truncated multiplier. d) MAC+ unit. Figure modified from [23].

#### IV. APPROXIMATE DNN ACCELERATOR

We employ our proposed control variate technique and design approximate DNN accelerators based on a micro-architecture similar to the Google TPU [1]. The latter is composed of a large  $N \times N$  systolic MAC array, as the one depicted in Fig. 5a. Fig. 5b presents a pipelined accurate MAC unit, i.e., the processing element replicated within the array. Each MAC unit comprises an 8-bit multiplier and a  $\lceil \log_2(N \times (2^{16} - 1)) \rceil$ -bit adder to avoid accumulation overflow [9]. As an example, for a  $64 \times 64$  MAC array, the size of the adder is 22-bit. In the approximate MAC array, depicted in Fig. 6a, each accurate MAC unit is substituted by its approximate version MAC\*, where the accurate multiplier is replaced with an approximate one from Section II. Moreover, the MAC\* unit is enhanced with the circuit computing the partial sum of  $\sum_{j=1}^k x_j$  needed to calculate the control variate  $V$  (see (15)) of each row. Finally, the approximate systolic array needs  $N + 1$  columns, i.e. one more column with respect to the accurate design. The extra column is composed by MAC+ units. The MAC+ unit calculates  $V = C \cdot \sum_{j=1}^k x_j$  and adds it to the convolution result generated by the MAC\* in the first  $N$  columns:  $B + \sum_{j=1}^k \text{AM}(W_j, A_j)$ . As discussed in the previous Section, the control variate approximation depends on the selected approximate multiplier, so the hardware implementation of the MAC\* unit is modified accordingly.

##### A. MAC\* UNIT WITH APPROXIMATE PERFORATED MULTIPLIERS

The product of the approximate perforated multiplier requires  $16 - m$  bits since  $m$  partial products are omitted. Hence, the adder of MAC\* can be simplified since its size can be reduced by  $m$  bits accordingly (compared with the adder of the MAC unit). The adder within the MAC\* units of the first column adds the first  $8 - m$  MSBs of  $B$ , i.e.  $B[7:m]$ ,

to the first approximate product  $\text{AMP}(W_1, A_1)$ . The addition of  $B[m-1:0]$  occurs in the MAC+ unit, as it will be explained later. Moreover, each MAC\* needs an extra adder to compute the partial sum required to calculate  $V$ . As explained in Section III-A, for the perforated multiplier  $x_j = A_j[m-1:0]$  and is  $m$ -bit wide. Thus, a  $\lceil \log_2(N \times (2^m - 1)) \rceil$ -bit adder is required. It is worth noting that such a size is considerably lower than the size of the main adder. As an example, for  $N = 64$  and  $m = 2$ , the size of the extra adder in the MAC\* unit is 8 bits while the main adder in MAC is 22 bits wide. Therefore, the associated hardware overhead is very small.

Overall, each MAC\* belonging to the  $h$ -th column of the approximate MAC array computes the following:

$$\begin{aligned} P_h^* &= \text{AMP}(W_h, A_h) \\ \text{sum}_h &= \text{sum}_{h-1} + P_h^*, \quad \text{sum}_0 = B[7:m] \\ \text{sum}X_h &= \text{sum}X_{h-1} + A_h[m-1:0], \quad \text{sum}X_0 = 0 \end{aligned} \quad (33)$$

##### B. MAC\* UNIT WITH APPROXIMATE TRUNCATED MULTIPLIERS

The output of the  $8 \times 8$  approximate truncated multiplier requires  $16 - m$  bits, since the  $m$  least significant columns are truncated. Consequently, as for the approximate perforated MAC\*, the main adder in the MAC\* unit can be again reduced by  $m$  bits with respect to the main adder of the MAC. Similarly, the main adder of the MAC\* units, belonging to the first column, receives, as one of the two inputs, the number  $B[7:m]$ . As described in Section III-B, in the case of the approximate truncated multiplier  $x_j$  is 1-bit wide and it is equal to  $(1 - \delta_{0,y_j})$ , where  $y_j = A_j \bmod 2^m$ . Hence, to calculate  $x_j$ , a simple  $m$ -bit OR gate is required that receives the bits  $A_j[m-1:0]$  as inputs. The output of the OR gate is sent to a small adder that calculates the partial sum of  $\sum_{j=1}^k x_j$ . Since a MAC\* can increment the latter by at most 1, the size of the small adder does not depend on  $m$  and it is equal to  $\lceil \log_2(N) \rceil$ . Overall, each MAC\* belonging to the  $h$ -th column of the approximate MAC array computes the following:

$$\begin{aligned} P_h^* &= \text{AM}_T(W_h, A_h) \\ \text{sum}_h &= \text{sum}_{h-1} + P_h^*, \quad \text{sum}_0 = B[7:m] \\ \text{sum}X_h &= \text{sum}X_{h-1} + \text{OR}(A_h[m-1:0]), \quad \text{sum}X_0 = 0 \end{aligned} \quad (34)$$

##### C. MAC\* UNIT WITH APPROXIMATE RECURSIVE MULTIPLIERS

In the approximate recursive multiplier the sub-product  $W_L \times A_L$  is omitted, with  $W_L$  and  $A_L$  being  $m$ -bit wide. Hence, again the approximate product requires  $16 - m$  bits. As described in III-C,  $x_j = A_j[m-1:0]$ , i.e., exactly the same in the MAC\* of the approximate perforated multiplier. Therefore, the design of the MAC\* unit of the approximate recursive multiplier is similar to Section IV-A and computes:

$$\begin{aligned} P_h^* &= \text{AM}_R(W_h, A_h) \\ \text{sum}_h &= \text{sum}_{h-1} + P_h^*, \quad \text{sum}_0 = B[7:m] \\ \text{sum}X_h &= \text{sum}X_{h-1} + A_h[m-1:0], \quad \text{sum}X_0 = 0 \end{aligned} \quad (35)$$



#### D. MAC<sup>+</sup> UNIT

Each MAC<sup>+</sup> unit (last column of the systolic array) calculates the control variate  $V = C \cdot \sum_{j=1}^N x_j$  by means of an exact multiplier. This is a common feature regardless the approximate multiplier employed in the MAC\* unit. The only difference is the size of the multiplier: it is equal to  $\lceil \log_2(N \times (2^m - 1)) \rceil \times 8$  when the approximate perforated and recursive multipliers are used in MAC\*, while it is equal to  $\lceil \log_2(N) \rceil \times 8$  for the truncated one. Furthermore, a  $\lceil \log_2(N \times (2^{16} - 1)) \rceil$ -bit adder is required to produce the final output  $G^*$ . Basically, the MAC<sup>+</sup> unit calculates the following:

$$V = C \cdot \text{sum}X_N \quad (36)$$


$$G^* = \{\text{sum}_N, B[m-1:0]\} + V \quad (37)$$

It is noteworthy that by concatenating  $\text{sum}_N$  and  $B[m-1:0]$  in (37) we manage to: i) shift left  $m$  places the partial sum of the MAC\* units ( $\text{sum}_N$ ) as required, and ii) add the  $m$ -LSBs of the bias  $B$  ( $B[m-1:0]$ ), which were not taken into account in the MAC\*. The additional  $(N+1)$ -th column composed of MAC<sup>+</sup> units increases the latency of the MAC array by one cycle. If the delay of the MAC<sup>+</sup> is higher than the delay of the accurate MAC, the MAC<sup>+</sup> unit can be further pipelined in order to sustain the same operating frequency. In this case, the latency may increase by two clock cycles per convolution layer. However, this overhead is completely negligible since the inference phase typically requires thousands of cycles for each convolution layer. Finally, the MAC<sup>+</sup> requires the value  $C$  to calculate  $V$ . This value can be transferred to the DNN accelerator along with the weights of the filters. Considering the size of DNNs, the data transfer overhead is again negligible.

As a final consideration, it is immediate to note that the computations of  $\text{sum}X_h$  and  $\text{sum}_h$  are independent so they are executed in parallel. Therefore, the adder required to calculate  $\text{sum}X_h$  is not on the critical path of MAC\* and thus, a slower and power-efficient ripple-carry adder can be employed. Moreover, the application of the selected approximate technique allows reducing the delay of the MAC\* with respect to the exact MAC. In addition, as explained above pipelining is used to make MAC<sup>+</sup> as fast as the exact MAC. Therefore, this delay slack enables downsizing the gates of the critical paths and boosts further the area and power savings [25].

#### V. EXPERIMENTAL RESULTS

In this section, we evaluate the efficiency of our proposed control variate approximation in terms of area, power, and accuracy. To achieve this, we design several  $N \times N$  exact as well as approximate MAC arrays. Four values are considered for  $N$ , from 16 up to 64. The accurate and approximate versions of the MAC array have been synthesized using Synopsys Design Compiler and mapped to a 14nm technology library. All the designs are implemented using the optimized components of the Synopsys DesignWare Library (i.e., reduction trees and adders), as typically done in commercial flows. During the synthesis, the `compile_ultra`



**FIGURE 7.** The a) power and b) area of our control variate approximation when using the approximate perforated multipliers for  $m \in [1, 3]$  and for varying MAC array sizes. The area and power values are normalized over the corresponding values of the respective accurate design.


command has been used to target performance optimization. The accurate MAC array has been synthesized at its minimum clock period. The latter value has been also used as a time constraint during the synthesis of the approximate MAC arrays. As described in the previous section, the MAC\* units are inherently faster than the accurate MAC, so the synthesis of the approximate array has been relaxed allowing optimizing area and power dissipation. Therefore, in the following evaluation, area and power analysis are performed at iso-delay. Power consumption is calculated with Synopsys PrimeTime on the basis of post-synthesis back-annotated simulations performed using Mentor Questasim. We simulate the analyzed MAC arrays for 10,000 inference cycles to obtain precise switching activity estimation. Such a value of simulated cycles is a good compromise between the accuracy of the obtained power results and the simulation duration time. Running post-synthesis timing simulations for the entire inference phase is infeasible due to the vast time required [9]. The inference accuracy is obtained by integrating the proposed control variate approximation into the approximate TensorFlow implementation of [26]. For the accuracy evaluation, we consider six popular CNNs of varying size, depth, and architecture, trained on the Cifar-10 and Cifar-100 datasets.

#### A. HARDWARE EVALUATION

First, we evaluate the power and area gains of our approximate MAC arrays compared to the accurate ones. For all the evaluated approximate designs, the critical path of the MAC<sup>+</sup> unit is shorter than the critical path of the exact MAC, so it does not need to be further pipelined. Therefore, compared to the accurate design, our approximate ones exhibit a latency overhead of only one clock cycle per convolution layer.

##### 1) CONTROL VARIATE & APPROXIMATE PERFORATED MULTIPLIERS

Fig. 7 presents the area and power savings delivered by the proposed control variate technique when applied with the perforated multiplier for increased values approximation ( $m$ ).



**FIGURE 8.** The a) power and b) area of our control variate approximation when using the approximate truncated multipliers for  $m \in [5, 7]$  and for varying MAC array sizes. The area and power values are normalized over the corresponding values of the respective accurate design.

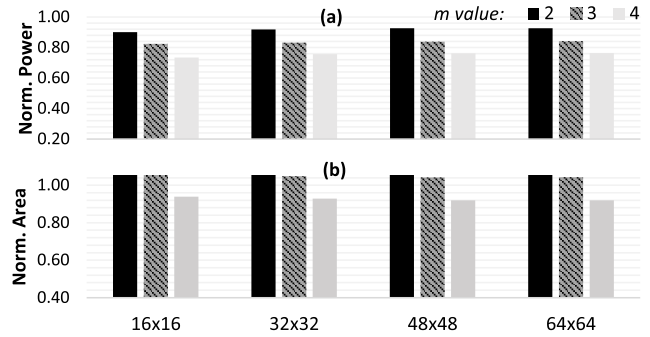
As shown in Fig. 7a, our approximate MAC arrays achieve large power reduction, ranging from 27.7% up to 46.1%. As expected, the power gain is directly proportional to the value of  $m$ . For  $m = 1$ , the power reduction ranges from 27.7% to 29.2%. For  $m = 2$ , the respective values range from 34.5% to 35.7%, while for  $m = 3$ , the power decreases from 44.4% to 46.1%. It is worth noting that the power reduction is almost insensitive to the array dimension  $N$ .

Fig. 7b shows that the area reduction entailed by the proposed approximate control variate technique is mainly defined by the  $m$  value and is slightly affected by the array size. As explained in the previous Section, the higher the value of  $m$  the higher the approximation and the lower the number of hardware resources, mainly Full Adders (FAs), needed by the MAC\* module. However, the MAC\* requires more flip-flops (FFs) than the accurate MAC due to the pipeline of the  $sumX$  path. For this reason, the area occupancy of the MAC\* unit for  $m = 1$  is almost the same with the one shown by the accurate MAC. Contrary, for higher values of  $m$  the area saving due to the reduced number of FAs overcomes the drawback of a higher number of FFs. Indeed, for  $m = 3$ , the area gain goes up to 22%. On average, the area reduction is 10%.

## 2) CONTROL VARIATE & APPROXIMATE TRUNCATED MULTIPLIERS

Fig. 8a shows the savings attained by our control variate technique when applied with the truncated multiplier. Similarly, the power dissipation decreases as the number of truncated columns  $m$  increases. For the considered range of  $m \in [5, 7]$ , the power gain goes up to a maximum of 41.9%, obtained for  $m = 1$  and  $N = 16$ . In this case, the power sensitivity of  $N$  is quite higher when compared to using the perforated multiplier. For  $m = 1$  the power gain ranges from 23.5% to 25.4% while for  $m = 2$  the respective values are from 28.6% to 35.0%. For  $m = 3$  the power savings range from 38.4% to 41.9%.

As shown in Fig. 8b, the area gain delivered by the control variate technique on the truncated array is considerably high,



**FIGURE 9.** The a) power and b) area of our control variate approximation when using the approximate recursive multipliers for  $m \in [2, 4]$  and for varying MAC array sizes. The area and power values are normalized over the corresponding values of the respective accurate design.

reaching a maximum value of 39% for  $m = 7$ . On average, the area reduction of the truncated multiplier is 31%. It is worth noting that the truncated multiplier entails a higher area saving with respect to the perforation approximation. This is explained by the fact that the additional adder that computes  $sumX$  is smaller and the associated FFs are fewer.

## 3) CONTROL VARIATE & APPROXIMATE RECURSIVE MULTIPLIERS

Finally, the power and area savings when using the approximate recursive multiplier are depicted in Fig. 9a-b. In this case, the approximate MAC array shows the lowest power gain since the power reduction is up to 26% (17% on average) compared to the exact design. Similarly, the maximum area gain is only 8%. Such limited gains are attributed to the fact that when  $m$  is small i) the hardware savings of omitting the least significant sub-product are constrained, and ii) the additional logic required by the control variate is significant compared to the approximation gains. It is noteworthy that for  $m = 2$  and  $N = 16$  there is a 14% area overhead.

## B. DNN ACCURACY EVALUATION

In this Section we evaluate the impact of our control variate technique on the delivered inference accuracy. Tables 2-4 report the accuracy loss of our control variate approximation compared to the exact design, for the examined approximate multipliers and for varying approximation values ( $m$ ), over six CNNs. Note that the accuracy does not depend on the size of the MAC array, since  $N$  only affects the number of operations performed in parallel by the array. Moreover, Tables 2-4 also report the accuracy loss when the approximate multiplier is used in the inference without our control variate technique (i.e. without adding  $V$ ). The latter highlights the accuracy improvement that is delivered by our method. Negative values in Tables 2-4 refer to higher accuracy compared to the baseline [10], [27].

As shown in Table 2, the average accuracy loss of our method when using the approximate perforated multiplier for the Cifar-10 dataset is 0.06%, 0.28%, and 4.12% for

**TABLE 2.** Accuracy evaluation when considering the approximate perforated multiplier. Six neural networks trained on Cifar-10 and Cifar-100 datasets are examined.

Accuracy Loss (%)						
NN on Cifar-10	$m = 1$		$m = 2$		$m = 3$	
	Ours <sup>+</sup>	w/o V <sup>*</sup>	Ours	w/o V	Ours	w/o V
GoogLeNet	-0.16	0.35	0.00	4.13	1.95	31.78
ResNet44	0.03	0.73	0.83	4.49	5.75	32.94
ResNet56	0.49	0.60	1.25	5.36	5.94	39.01
ShuffleNet	0.12	3.40	-0.48	7.60	6.53	31.74
VGG13	-0.01	0.23	-0.30	0.76	0.69	6.76
VGG16	-0.13	0.19	0.38	1.66	3.86	8.71
<b>Average</b>	0.06	0.92	0.28	4.00	4.12	25.16

NN on Cifar-100	$m = 1$		$m = 2$		$m = 3$	
	Ours	w/o V	Ours	w/o V	Ours	w/o V
GoogLeNet	0.05	2.43	1.47	13.19	7.02	44.52
ResNet44	0.77	1.02	1.82	14.17	11.27	43.40
ResNet56	-0.34	3.02	2.25	15.83	14.34	44.59
ShuffleNet	0.20	9.47	1.09	5.92	6.57	15.84
VGG13	0.89	3.01	1.91	5.89	5.52	14.70
VGG16	-0.03	3.96	0.03	2.41	2.80	11.54
<b>Average</b>	0.26	3.82	1.43	9.57	7.92	29.10

<sup>+</sup> Accuracy achieved when using the Approximate Perforated Multiplier with our proposed control-variate approximation.

<sup>\*</sup> Accuracy when using only the Approximate Perforated Multiplier.

$m = 1$ ,  $m = 2$ , and  $m = 3$  respectively. The corresponding values for the more challenging Cifar-100 dataset are 0.26%, 1.43%, and 7.92%. As a result, our technique achieves ~24% power reduction for negligible accuracy loss, i.e., 0.16% on average on both datasets for  $m = 1$ . The power gains rise to ~36% ( $m = 2$ ) for an average accuracy loss of only 0.85%. Finally, for 6.02% average accuracy loss ( $m = 3$ ), the power savings jump to ~46%. In addition, Table 2 highlights the efficiency of our control variate approximation in decreasing the convolution error. Indeed, compared with using the perforated multiplier standalone (i.e., without adding  $V$ ), our technique achieves 2%, 6%, and 21% higher accuracy, on average, for  $m = 1$ ,  $m = 2$ , and  $m = 3$  respectively. Interestingly, the higher the value of  $m$  the higher the accuracy improvement. This proves that our proposed control variate approach is also very effective when the approximation configuration of the employed approximate multiplier causes a large error value. As a consequence, the control variate technique can enable an aggressive approximation, and thus a significant energy reduction at a reduced accuracy loss.

Similarly, Table 3 presents the accuracy evaluation when using the approximate truncated multiplier. The average accuracy loss of our method for the Cifar-10 dataset is 0.3%, 3.46%, and 12.95% for  $m = 5$ ,  $m = 6$ , and  $m = 7$  respectively. The corresponding values for the Cifar-100 dataset are 0.09%, 2.43%, and 16.73%. Moreover, compared to the case of using the approximate truncated multiplier standalone (i.e., without adding  $V$ ), the proposed control variate technique improves the obtained accuracy by up to 22x for the same  $m$  value. On average, over all the examined cases, our control variate method improves the accuracy by

**TABLE 3.** Accuracy evaluation when considering the approximate truncated multiplier. Six neural networks trained on Cifar-10 and Cifar-100 datasets are examined.

Accuracy Loss (%)						
NN on Cifar-10	$m = 5$		$m = 6$		$m = 7$	
	Ours <sup>+</sup>	w/o V <sup>*</sup>	Ours	w/o V	Ours	w/o V
GoogLeNet	-0.44	1.54	0.19	10.38	0.79	33.78
ResNet44	0.20	0.98	0.48	4.95	3.93	27.70
ResNet56	-0.37	0.76	0.22	5.23	5.00	35.07
ShuffleNet	1.13	3.51	1.89	14.26	16.67	49.55
VGG13	0.07	5.14	4.14	35.19	17.33	72.42
VGG16	1.21	5.65	13.84	41.42	33.98	74.95
<b>Average</b>	0.30	2.93	3.46	18.57	12.95	48.91

NN on Cifar-100	$m = 5$		$m = 6$		$m = 7$	
	Ours	w/o V	Ours	w/o V	Ours	w/o V
GoogLeNet	0.58	11.84	1.13	24.62	9.45	53.04
ResNet44	-0.60	3.39	0.67	11.25	9.29	36.66
ResNet56	-0.70	2.50	1.00	13.96	8.84	40.29
ShuffleNet	-2.63	10.43	-0.85	23.41	8.22	37.50
VGG13	3.18	18.83	5.68	53.85	27.78	64.70
VGG16	0.70	20.31	6.96	56.32	36.81	63.52
<b>Average</b>	0.09	11.22	2.43	30.57	16.73	49.29

<sup>+</sup> Accuracy achieved when using the Approximate Truncated Multiplier with our proposed control-variate approximation.

<sup>\*</sup> Accuracy when using only the Approximate Truncated Multiplier.

**TABLE 4.** Accuracy evaluation when considering the approximate recursive multiplier. Six neural networks trained on Cifar-10 and Cifar-100 datasets are examined.

Accuracy Loss (%)						
NN on Cifar-10	$m = 2$		$m = 3$		$m = 4$	
	Ours <sup>+</sup>	w/o V <sup>*</sup>	Ours	w/o V	Ours	w/o V
GoogLeNet	-0.11	0.15	0.12	1.04	0.13	3.17
ResNet44	-0.10	-0.03	0.02	0.07	0.11	3.82
ResNet56	-0.12	0.74	0.07	0.13	-0.22	2.93
ShuffleNet	-0.28	-0.28	0.07	2.70	1.59	9.77
VGG13	-0.48	0.01	-0.81	0.54	2.18	8.11
VGG16	0.09	0.41	0.31	1.89	3.13	7.85
<b>Average</b>	-0.17	0.17	-0.04	1.06	1.15	5.94

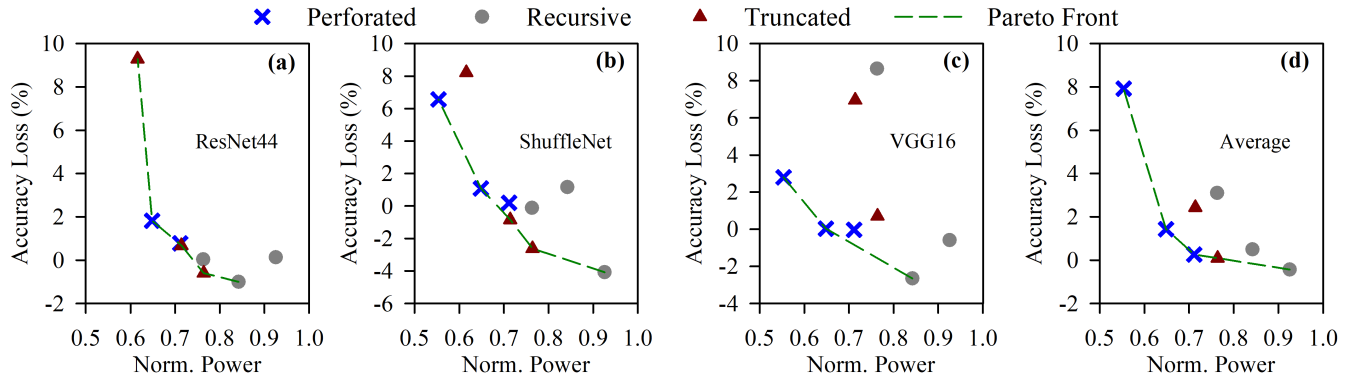
NN on Cifar-100	$m = 2$		$m = 3$		$m = 4$	
	Ours	w/o V	Ours	w/o V	Ours	w/o V
GoogLeNet	-0.09	-0.09	1.58	2.00	0.53	19.19
ResNet44	0.13	0.13	-1.00	1.87	0.03	9.58
ResNet56	-0.32	-0.32	-0.30	0.16	0.25	10.63
ShuffleNet	-4.08	-3.49	1.17	-1.10	-0.12	20.50
VGG13	2.33	8.75	4.18	8.43	9.26	38.65
VGG16	-0.59	-0.59	-2.65	3.67	8.65	35.74
<b>Average</b>	-0.44	0.73	0.50	2.51	3.10	22.38

<sup>+</sup> Accuracy achieved when using the Approximate Recursive Multiplier with our proposed control-variate approximation.

<sup>\*</sup> Accuracy when using only the Approximate Recursive Multiplier.

3.32x. As a result, the control variate technique applied to the truncated multiplier can reduce the power dissipation by ~25% at a negligible accuracy loss of only 0.2% on average for  $m = 5$ .

Finally, the same trend is observed in Table 4 for the approximate recursive multiplier. Similarly to the previous cases, the accuracy loss is proportional to the value of  $m$ .



**FIGURE 10.** The Accuracy Loss (%) – Normalized power pareto space for varying CNNs (subfigures a-c) as well as the average case (subfigure d). Cifar-100 and  $N = 64$  are used. Normalized power and accuracy loss are reported w.r.t. the baseline exact design.

Our control variate approach applied to the approximate recursive multipliers entails the lowest accuracy loss among the analyzed approximate multipliers. Indeed, the maximum accuracy loss is only 3.1% on average for  $m = 4$ . The average accuracy loss for the Cifar-10 dataset is  $-0.17\%$ ,  $-0.04\%$ , and  $1.15\%$  for  $m = 2$ ,  $m = 3$ , and  $m = 4$ , respectively. The corresponding values for the Cifar-100 dataset are  $-0.44\%$ ,  $0.5\%$ , and  $3.1\%$ . Compared to using the approximate recursive multiplier standalone (i.e., without adding  $V$ ), our method achieves a maximum accuracy improvement of 2.1x, obtained for the Cifar-100 dataset and  $m = 4$ . On average, over all the examined cases, our control variate method improves the accuracy by 4.78%. Finally, the power saving is  $\sim 17\%$  at a negligible average accuracy loss of  $0.23\%$  ( $m = 3$ ). For  $m = 4$ , the power saving is  $\sim 25\%$  at the cost of an accuracy loss of  $\sim 2.1\%$ .

As it can be inferred from the above discussion, our control variate method always improves the obtained accuracy and enables achieving high power reduction even under more constrained accuracy loss thresholds. Nevertheless, the power-accuracy trade-off depends on several parameters such as the approximate multiplier type, the applied level of approximation ( $m$ ), and the examined network. As an example, as shown in Table 2 and Table 3, our control variate method on the perforated multiplier mainly delivers, on average, better accuracy than when it is applied with the truncated one. Though, if we restrict our analysis on the ResNet family (ResNet44 and ResNet56), using the truncated multiplier outperforms the perforated one in most of the cases, for both the Cifar-10 and the Cifar-100. Fig. 10 depicts the respective Pareto space (accuracy vs. power) for a variety of representative cases: the a) ResNet44, b) ShuffleNet, c) VGG16, and d) average accuracy over all the examined CNNs. For the analysis in Fig. 10, we consider the Cifar-100 dataset and a  $64 \times 64$  MAC array size. In addition, only the configurations that result in up to 10% accuracy loss are depicted. Similar results are obtained for Cifar-10 and the rest of the  $N$  values. As Fig. 10 shows, there is not a dominant solution, but the optimal design choice depends on the target accuracy threshold and the network's topology.

The Pareto front spans across a wide range of approximate multipliers and configurations. As a general observation, it is better to apply our control variate approach with the recursive approximate multiplier if the accuracy loss constraint is very tight. Under relaxed accuracy constraints, the perforated multiplier should be preferred since it delivers mainly the highest power savings. The truncated multiplier is usually in the middle delivering a considerable power reduction for limited accuracy loss. However, in the case of VGG16, there is not any Pareto-optimal approximate design that uses the truncated multiplier. Finally, it is important to emphasize that our control variate technique is applied to different approximate multipliers allowing to obtain design points in the Pareto front that would be impossible to reach with just a single approximate multiplier. Hence, the versatility of our approach leads to more fine-grained traversal of the accuracy-power design space.

### C. MAC<sup>+</sup> OVERHEAD DISCUSSION

As described in Section IV, the proposed control variate technique requires an additional column of MAC<sup>+</sup> units in order to add  $V$  to the final accumulated value, and thus mitigate the accuracy loss due to the approximate multiplications. Although in the analysis in Fig. 7-9 we have evaluated the area and power consumption of the entire approximate MAC array, i.e., including the MAC<sup>+</sup> units, it is useful to disaggregate the area and power consumption of the extra MAC<sup>+</sup> modules from the total area and power consumption, in order to analyze their impact on the overall architecture. Table 5 shows the area and power of the MAC<sup>+</sup> units expressed as percentages over the total area and power consumption of the approximate systolic MAC array. As shown in Table 5, the impact of the MAC<sup>+</sup> units is negligible. The maximum area (power) overhead has been found to be just 1.38% (1.52%) for the truncated (perforated) multiplier and array size of  $16 \times 16$ . As expected, the area overhead increases as the applied approximation ( $m$ ) increases. This is mainly explained by the fact that as  $m$  increases, the area gains in the MAC\* units increase significantly while the MAC<sup>+</sup> are hardly affected. Hence,



**TABLE 5.** Evaluation of the area and power overheads of MAC<sup>+</sup>.

<i>Approximate Perforated Multiplier in MAC*</i>					<i>Approximate Recursive Multiplier in MAC*</i>					<i>Approximate Truncated Multiplier in MAC*</i>				
Percentage of Total Area (%)					Percentage of Total Area (%)					Percentage of Total Area (%)				
<i>m</i>	16 × 16	32 × 32	48 × 48	64 × 64	<i>m</i>	16 × 16	32 × 32	48 × 48	64 × 64	<i>m</i>	16 × 16	32 × 32	48 × 48	64 × 64
1	1.07	0.55	0.38	0.28	2	0.95	0.48	0.33	0.25	5	1.15	0.59	0.40	0.30
2	1.18	0.61	0.41	0.31	3	1.02	0.53	0.36	0.27	6	1.30	0.66	0.45	0.33
3	1.36	0.71	0.47	0.36	4	1.15	0.59	0.41	0.31	7	1.38	0.71	0.48	0.36
Percentage of Total Power (%)					Percentage of Total Power (%)					Percentage of Total Power (%)				
<i>m</i>	16 × 16	32 × 32	48 × 48	64 × 64	<i>m</i>	16 × 16	32 × 32	48 × 48	64 × 64	<i>m</i>	16 × 16	32 × 32	48 × 48	64 × 64
1	1.22	0.63	0.43	0.32	2	0.99	0.50	0.34	0.25	5	1.06	0.54	0.37	0.28
2	1.32	0.68	0.46	0.35	3	1.06	0.54	0.36	0.27	6	1.19	0.60	0.39	0.29
3	1.52	0.80	0.53	0.40	4	1.16	0.58	0.39	0.30	7	1.25	0.64	0.43	0.32

considering the significantly higher number of MAC<sup>\*</sup>, the area overhead increases with  $m$ , being however 1.52% at maximum. Moreover, the area overhead decreases as  $N$  increases. This is mainly explained by the fact that the total area of the MAC<sup>+</sup> units scales linearly with respect to  $N$  while the total area scales quadratically with  $N$ . The same conclusion can be drawn about the impact of the extra MAC<sup>+</sup> modules on the power consumption of the systolic array, and the its dependency on  $m$  and  $N$ . Concluding, Fig. 7, Fig. 8, and Table 5 clearly demonstrate the scalability of our approach.

## VI. RELATED WORKS

There has been great interest around approximate computing for neural network inference. References [8] and [28] employed approximate multipliers to different convolution layers and [7] proposed a compact and energy-efficient multiplier-less artificial neuron. Similarly, [29] used reconfigurable constant coefficient multipliers that require only adders and shifters. The authors in [30], [31], [32] replace multiplication with approximate logarithmic multipliers. However, [7], [8], [32] are based on retraining to recover accuracy loss caused by the usage of approximation while [30], [31] consider energy-inefficient 32-bit inference. The authors in [33] proposed a framework for approximation-aware retraining. Still, the execution time remains very high for complex datasets and networks. [34] utilized approximate MAC by splitting the addition and limiting the carry propagation. Nevertheless, [8], [34] are evaluated on the LeNet network, a very shallow architecture which cannot provide the amount of operations recent DNNs do. Therefore, both of these methods can be deemed inapplicable in modern scenarios which require deeper network architectures. The authors in [35] approximate bit-decomposition DNN architectures by limiting the maximum value of the partial sums. Though, [35] targets in-memory inference. The work in [36] enhances the approximate multipliers library of [37] and shows that, in simple DNNs, even without retraining, considerable energy gains for a small accuracy loss can be attained. Still, in more complex DNNs the energy gains are not maintained [36]. The authors in [10] propose a non-uniform architecture that utilized

approximate multipliers from [37]. Their work tunes the weights accordingly and avoids retraining. However, [10] requires a heterogeneous design and generates a different approximate accelerator per DNN. Similarly, inspired by big.LITTLE computing, Spantidi et al. [13] implemented heterogeneous DNN accelerators that consist of 8-bit NPUs in conjunction with lower bit-width NPUs to enhance overall throughput while reducing energy consumption during NN inference. In [3], the authors use low-precision MAC units with larger NPUs to reduce power density. In [11], approximate multipliers with reconfigurable accuracy at run-time are generated. Similar to [10], they also apply layer-wise approximation, however with limited power reduction. In [9], [12] approximate reconfigurable multipliers are used, and a mapping strategy is employed to set the approximation level per weight. Reference [9] applies layer-wise approximation while [12] introduced a more fine-grained filter-wise approximation. However, [9] and [12] significantly increase the size of the DNN to store the required configurations. The work in [38] allows the usage of reduced-complexity multipliers based on Canonic Sign Digit approximation to represent the filter weights of CNNs that have already been trained. Nevertheless, [38] requires DNN-specific approximations. In [39] the authors present an interleaving method that employs approximate multipliers to minimize energy consumption of MAC-oriented signal processing algorithms with minimal performance loss. Towards improving the multiplication performance for CNN inference, [40] proposes an architecture comprising a preprocessing precision controller and approximate multiplier designs of varying precision using the static and dynamic segment methods. Both [39] and [40] apply their proposed work on 16-bit inference, while recent CNN accelerators are mostly targeting 8-bit precision [1]. Voltage scaling can help reduce power consumption and manage temperature; however, it significantly decreases the throughput of DNNs [2], [41] while voltage over-scaling [42] can lead to unpredictable timing errors, which may severely impact accuracy [43].

## VII. CONCLUSION

In this work, we introduced control variate approximation to increase the accuracy of approximate DNN accelerators



without requiring any DNN retraining. Our mathematical analysis demonstrates that our technique mitigates the error induced by the approximate multipliers, by effectively nullifying the mean convolution error and reducing its variance. As our extensive experimentation over three diverse approximate multipliers, six DNNs, and four MAC array sizes demonstrates, our control variate approximation enables using aggressive approximate multipliers to design approximate DNN accelerators that boost the power savings with limited accuracy loss.

## REFERENCES

- [1] N. P. Jouppi et al., "In-datacenter performance analysis of a tensor processing unit," in *Proc. ACM/IEEE 44th Annu. Int. Symp. Comput. Archit. (ISCA)*, Jun. 2017, pp. 1–12.
- [2] H. Amrouch, G. Zervakis, S. Salamin, H. Kattan, I. Anagnostopoulos, and J. Henkel, "NPU thermal management," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 39, no. 11, pp. 3842–3855, Nov. 2020.
- [3] G. Zervakis, I. Anagnostopoulos, S. Salamin, O. Spantidi, I. Roman-Ballesteros, J. Henkel, and H. Amrouch, "Thermal-aware design for approximate DNN accelerators," *IEEE Trans. Comput.*, vol. 71, no. 10, pp. 2687–2697, Oct. 2022.
- [4] G. Armeniakos, G. Zervakis, D. Soudris, and J. Henkel, "Hardware approximate techniques for deep neural network accelerators: A survey," *ACM Comput. Surv.*, vol. 55, no. 4, pp. 1–36, Nov. 2022.
- [5] M. Alioto, V. De, and A. Marongiu, "Energy-quality scalable integrated circuits and systems: Continuing energy scaling in the twilight of Moore's law," *IEEE J. Emerg. Sel. Topics Circuits Syst.*, vol. 8, no. 4, pp. 653–678, Dec. 2018.
- [6] H. Jiang, F. J. H. Santiago, H. Mo, L. Liu, and J. Han, "Approximate arithmetic circuits: A survey, characterization, and recent applications," *Proc. IEEE*, vol. 108, no. 12, pp. 2108–2135, Dec. 2020.
- [7] S. S. Sarwar, S. Venkataramani, A. Ankit, A. Raghunathan, and K. Roy, "Energy-efficient neural computing with approximate multipliers," *ACM J. Emerg. Technol. Comput. Syst.*, vol. 14, no. 2, pp. 1–23, Apr. 2018.
- [8] V. Mrazek, S. S. Sarwar, L. Sekanina, Z. Vasicek, and K. Roy, "Design of power-efficient approximate multipliers for approximate artificial neural networks," in *Proc. IEEE/ACM Int. Conf. Comput.-Aided Design (ICCAD)*, Nov. 2016, pp. 1–7.
- [9] Z.-G. Tasoulas, G. Zervakis, I. Anagnostopoulos, H. Amrouch, and J. Henkel, "Weight-oriented approximation for energy-efficient neural network inference accelerators," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 67, no. 12, pp. 4670–4683, Dec. 2020.
- [10] V. Mrazek, Z. Vasicek, L. Sekanina, M. A. Hanif, and M. Shafique, "ALWANN: Automatic layer-wise approximation of deep neural network accelerators without retraining," in *Proc. IEEE/ACM Int. Conf. Comput.-Aided Design (ICCAD)*, Nov. 2019, pp. 1–8.
- [11] G. Zervakis, H. Amrouch, and J. Henkel, "Design automation of approximate circuits with runtime reconfigurable accuracy," *IEEE Access*, vol. 8, pp. 53522–53538, 2020.
- [12] O. Spantidi, G. Zervakis, I. Anagnostopoulos, H. Amrouch, and J. Henkel, "Positive/Negative approximate multipliers for DNN accelerators," in *Proc. IEEE/ACM Int. Conf. Comput. Aided Design (ICCAD)*, Nov. 2021, pp. 1–9.
- [13] O. Spantidi, G. Zervakis, S. Alsalam, I. Roman-Ballesteros, J. Henkel, H. Amrouch, and I. Anagnostopoulos, "Targeting DNN inference via efficient utilization of heterogeneous precision DNN accelerators," *IEEE Trans. Emerg. Topics Comput.*, vol. 11, no. 1, pp. 112–125, Jan. 2023.
- [14] A. Raha and V. Raghunathan, "Towards full-system energy-accuracy tradeoffs: A case study of an approximate smart camera system," in *Proc. 54th ACM/EDAC/IEEE Design Autom. Conf. (DAC)*, Jun. 2017, pp. 1–6.
- [15] M. de la Guia Solaz, W. Han, and R. Conway, "A flexible low power DSP with a programmable truncated multiplier," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 59, no. 11, pp. 2555–2568, Nov. 2012.
- [16] F. Frustaci, S. Perri, P. Corsonello, and M. Alioto, "Approximate multipliers with dynamic transduction for energy reduction via graceful quality degradation," *IEEE Trans. Circuits Syst. II, Exp. Briefs*, vol. 67, no. 12, pp. 3427–3431, Dec. 2020.
- [17] V. Mrazek, Z. Vasicek, L. Sekanina, H. Jiang, and J. Han, "Scalable construction of approximate multipliers with formally guaranteed worst case error," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 26, no. 11, pp. 2572–2576, Nov. 2018.
- [18] H. Pei, X. Yi, H. Zhou, and Y. He, "Design of ultra-low power consumption approximate 4–2 compressors based on the compensation characteristic," *IEEE Trans. Circuits Syst. II, Exp. Briefs*, vol. 68, no. 1, pp. 461–465, Jan. 2021.
- [19] A. G. M. Strollo, E. Napoli, D. De Caro, N. Petra, and G. D. Meo, "Comparison and extension of approximate 4–2 compressors for low-power approximate multipliers," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 67, no. 9, pp. 3021–3034, Sep. 2020.
- [20] G. Zervakis, K. Tsoumanis, S. Xydis, D. Soudris, and K. Pekmestzi, "Design-efficient approximate multiplication circuits through partial product perforation," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 24, no. 10, pp. 3105–3117, Oct. 2016.
- [21] P. Kulkarni, P. Gupta, and M. Ercegovac, "Trading accuracy for power with an underdesigned multiplier architecture," in *Proc. 24th International Conf. VLSI Design*, Jan. 2011, pp. 346–351.
- [22] H. Waris, C. Wang, C. Xu, and W. Liu, "AxRMs: Approximate recursive multipliers using high-performance building blocks," *IEEE Trans. Emerg. Topics Comput.*, vol. 10, no. 2, pp. 1229–1235, Apr. 2022.
- [23] G. Zervakis, O. Spantidi, I. Anagnostopoulos, H. Amrouch, and J. Henkel, "Control variate approximation for DNN accelerators," in *Proc. 58th ACM/IEEE Design Autom. Conf. (DAC)*, Dec. 2021, pp. 481–486.
- [24] C. Li, W. Luo, S. S. Sapatnekar, and J. Hu, "Joint precision optimization and high level synthesis for approximate computing," in *Proc. 52nd ACM/EDAC/IEEE Design Autom. Conf. (DAC)*, Jun. 2015, pp. 1–6.
- [25] S. Venkataramani, K. Roy, and A. Raghunathan, "Substitute-and-simplify: A unified design paradigm for approximate and quality configurable circuits," in *Proc. Design, Autom. Test Eur. Conf. Exhib. (DATE)*, 2013, pp. 1367–1372.
- [26] F. Vaverka, V. Mrazek, Z. Vasicek, and L. Sekanina, "TFApprox: Towards a fast emulation of DNN approximate hardware accelerators on GPU," in *Proc. Design, Autom. Test Eur. Conf. Exhib. (DATE)*, Mar. 2020, pp. 294–297.
- [27] M. S. Ansari, V. Mrazek, B. F. Cockburn, L. Sekanina, Z. Vasicek, and J. Han, "Improving the accuracy and hardware efficiency of neural networks using approximate multipliers," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 28, no. 2, pp. 317–328, Feb. 2020.
- [28] Z. Vasicek, V. Mrazek, and L. Sekanina, "Automated circuit approximation method driven by data distribution," in *Proc. Design, Autom. Test Eur. Conf. Exhib. (DATE)*, Mar. 2019, pp. 96–101.
- [29] J. Faraone, M. Kumm, M. Hardieck, P. Zipf, X. Liu, D. Boland, and P. H. W. Leong, "AddNet: Deep neural networks using FPGA-optimized multipliers," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 28, no. 1, pp. 115–128, Jan. 2020.
- [30] M. S. Kim, A. A. D. Barrio, L. T. Oliveira, R. Hermida, and N. Bagherzadeh, "Efficient Mitchell's approximate log multipliers for convolutional neural networks," *IEEE Trans. Comput.*, vol. 68, no. 5, pp. 660–675, May 2019.
- [31] H. Saadat, H. Bokhari, and S. Parameswaran, "Minimally biased multipliers for approximate integer and floating-point multiplication," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 37, no. 11, pp. 2623–2635, Nov. 2018.
- [32] R. Pilipovic, P. Bulic, and U. Lotric, "A two-stage operand trimming approximate logarithmic multiplier," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 68, no. 6, pp. 2535–2545, Jun. 2021.
- [33] C. De la Parra, A. Guntoro, and A. Kumar, "Efficient accuracy recovery in approximate neural networks by systematic error modelling," in *Proc. 26th Asia South Pacific Design Autom. Conf. (ASP-DAC)*, Jan. 2021, pp. 365–371.
- [34] M. A. Hanif, F. Khalid, and M. Shafique, "CANN: Curable approximations for high-performance deep neural network accelerators," in *Proc. 56th ACM/IEEE Design Autom. Conf. (DAC)*, Jun. 2019, pp. 1–6.
- [35] T. Soliman, C. De La Parra, A. Guntoro, and N. Wehn, "Adaptable approximation based on bit decomposition for deep neural network accelerators," in *Proc. IEEE 3rd Int. Conf. Artif. Intell. Circuits Syst. (AICAS)*, Jun. 2021, pp. 1–4.
- [36] V. Mrazek, L. Sekanina, and Z. Vasicek, "Libraries of approximate circuits: Automated design and application in CNN accelerators," *IEEE J. Emerg. Sel. Topics Circuits Syst.*, vol. 10, no. 4, pp. 406–418, Dec. 2020.

- [37] V. Mrazek, R. Hrbacek, Z. Vasicek, and L. Sekanina, "EvoApprox8b: Library of approximate adders and multipliers for circuit design and benchmarking of approximation methods," in *Proc. Design, Autom. Test Eur. Conf. Exhib. (DATE)*, Mar. 2017, pp. 258–261.
- [38] M. Riaz, R. Hafiz, S. A. Khaliq, M. Faisal, H. T. Iqbal, M. Ali, and M. Shafique, "CAXCNN: Towards the use of canonic sign digit based approximation for hardware-friendly convolutional neural networks," *IEEE Access*, vol. 8, pp. 127014–127021, 2020.
- [39] G. Park, J. Kung, and Y. Lee, "Design and analysis of approximate compressors for balanced error accumulation in MAC operator," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 68, no. 7, pp. 2950–2961, Jul. 2021.
- [40] I. Hammad, L. Li, K. El-Sankary, and W. M. Snelgrove, "CNN inference using a preprocessing precision controller and approximate multipliers with various precisions," *IEEE Access*, vol. 9, pp. 7220–7232, 2021.
- [41] J. Song, Y. Cho, J.-S. Park, J.-W. Jang, S. Lee, J.-H. Song, J.-G. Lee, and I. Kang, "7.1 an 11.5TOPS/W 1024-MAC butterfly structure dual-core sparsity-aware neural processing unit in 8nm flagship mobile SoC," in *Proc. IEEE Int. Solid- State Circuits Conf. (ISSCC)*, Feb. 2019, pp. 130–132.
- [42] G. Paim, G. Zervakis, G. Pahwa, Y. S. Chauhan, E. A. C. da Costa, S. Bampi, J. Henkel, and H. Amrouch, "On the resiliency of NCFET circuits against voltage over-scaling," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 68, no. 4, pp. 1481–1492, Apr. 2021.
- [43] S. Salamin, G. Zervakis, O. Spantidi, I. Anagnostopoulos, J. Henkel, and H. Amrouch, "Reliability-aware quantization for anti-aging NPU," in *Proc. Design, Autom. Test Eur. Conf. Exhib. (DATE)*, Feb. 2021, pp. 1460–1465.



research projects as a member of the Institute of Communication and Computer Systems (ICCS), Athens, Greece. His main research interests include low-power design, accelerator microarchitectures, approximate computing, machine learning, and flexible electronics. He has received one Best Paper Nomination at DATE 2022. He serves as a reviewer for many IEEE and ACM Transactions journals and a member of the technical program committee of several major design conferences.

**GEORGIOS ZERVAKIS** (Member, IEEE) received the Diploma and Ph.D. degrees from the School of Electrical and Computer Engineering (ECE), National Technical University of Athens (NTUA), Greece, in 2012 and 2018, respectively. He is an Assistant Professor with the University of Patras. Before, that he was a Research Group Leader with the Chair for Embedded Systems (CES), Karlsruhe Institute of Technology (KIT), from 2019 to 2022. From 2015 to 2019, he worked in many EU-funded



Professor with the Computer Science, Electronics, Modeling and Systems Department, University of Calabria, Rende, Italy. He has authored more than 70 articles in the field of VLSI design. His research interests include low-power and high-performance VLSI circuits, design techniques for emerging technologies, reconfigurable architectures, and embedded systems. He is a member of the editorial board of *Microelectronics Journal*.

**FABIO FRUSTACI** (Senior Member, IEEE) received the M.S. and Ph.D. degrees in electronic engineering from the Mediterranean University of Reggio Calabria, Reggio Calabria, Italy, in 2003 and 2007, respectively. In 2006, he was a Visiting Scholar with the ECE Department, University of Rochester, Rochester, NY, USA. From 2011 to 2013, he was a Visiting Researcher with the EECS Department, University of Michigan, Ann Arbor, MI, USA. He is an Associate



Embedded Systems Software Laboratory. Her research interests include neural network optimization for embedded systems and approximate computing.

**OURANIA SPANTIDI** received the Ph.D. degree from the School of Electrical, Computer, and Biomedical Engineering, Southern Illinois University at Carbondale. She is an Assistant Professor with the Department of Computer Science, Eastern Michigan University, and the Director of the Embedded AI Systems Laboratory. She was with the School of Electrical, Computer, and Biomedical Engineering, Southern Illinois University at Carbondale, where she was a member of the



embedded many-core architectures. His research interests include machine learning, heterogeneous hardware accelerators, and hardware/software co-design.

**IRAKLIS ANAGNOSTOPOULOS** (Member, IEEE) received the Ph.D. degree from the Microprocessors and Digital Systems Laboratory, National Technical University of Athens. He is an Associate Professor with the School of Electrical, Computer and Biomedical Engineering, Southern Illinois University at Carbondale. He is the Director of the Embedded Systems Software Laboratory, which works on run-time resource management of modern and heterogeneous



plinary research areas (including over 115 journals) across the computing stack, starting from semiconductor physics to circuit design all the way up to computer architecture. His research in HW security and reliability have been funded by German Research Foundation (DFG), Advantest Corporation, and U.S. Office of Naval Research (ONR). His research interests include in-memory computing, reliability of advanced technology, and emerging technologies. He holds 10x HiPEAC Paper Awards.

**HUSSAM AMROUCH** (Member, IEEE) received the Ph.D. degree (summa cum laude) from KIT, in 2015. He is currently a Professor, heading the Chair of AI Processor Design, Technical University of Munich (TUM), Germany. He is also heading the Brain-Inspired Computing, Munich Institute of Robotics and Machine Intelligence (MIRMI); and the Head of the Semiconductor Test and Reliability, University of Stuttgart, Germany. He has more than 270 publications in multidisciplinary



1500 "Dependable Embedded Systems." He is a site coordinator of the DFG TR89 collaborative research center "Invasive Computing." He is the Chairperson of the IEEE Computer Society, Germany Chapter. For two terms, he served as the Editor-in-Chief for the *ACM Transactions on Embedded Computing Systems*. He is currently the Editor-in-Chief of *IEEE Design and Test Magazine*. He is/has been an Associate Editor of major ACM and IEEE journals.

**JÖRG HENKEL** (Fellow, IEEE) is with Karlsruhe Institute of Technology. He was a Research Staff Member with NEC Laboratories, Princeton, NJ, USA. He has received six best paper awards from ICCAD, ESWeek, and DATE. He has led several conferences as the General Chair, including ICCAD, ESWeek; and serves as the Steering Committee Chair/member for leading conferences and journals for embedded and cyber-physical systems. He coordinates the DFG Program SPP

...