Ax-BxP: Approximate Blocked Computation for Precision-reconfigurable Deep Neural Network Acceleration

REENA ELANGO VAN, Purdue University

SHUBHAM JAIN, IBM T. J. Watson Research Center

ANAND RAGHUNATHAN, Purdue University

Precision scaling has emerged as a popular technique to optimize the compute and storage requirements of Deep Neural Networks (DNNs). Efforts toward creating ultra-low-precision (sub-8-bit) DNNs for efficient inference suggest that the minimum precision required to achieve a given network-level accuracy varies considerably across networks, and even across layers within a network. This translates to a need to support variable precision computation in DNN hardware. Previous proposals for precision-reconfigurable hardware, such as bit-serial architectures, incur high overheads, significantly diminishing the benefits of lower precision. We propose Ax-BxP, a method for approximate blocked computation wherein each multiply-accumulate operation is performed block-wise (a block is a group of bits), facilitating re-configurability at the granularity of blocks. Further, approximations are introduced by only performing a subset of the required block-wise computations to realize precision re-configurability with high efficiency. We design a DNN accelerator that embodies approximate blocked computation and propose a method to determine a suitable approximation configuration for any given DNN. For the AlexNet, ResNet50, and MobileNetV2 DNNs, Ax-BxP achieves $1.1\times$ – $1.74\times$ and $1.02\times$ – $2\times$ improvement in system energy and performance, respectively, over an 8-bit fixed-point (Fxp8) baseline, with minimal loss ($<1\%$ on average) in classification accuracy. Further, by varying the approximation configurations at a finer granularity across layers and data-structures within a DNN, we achieve $1.12\times$ – $2.23\times$ and $1.14\times$ – $2.34\times$ improvement in system energy and performance, respectively.

CCS Concepts: • **Computer systems organization** → **Reconfigurable computing**; • **Computing methodologies** → **Neural networks**; • **Hardware** → **Neural systems**; **Reconfigurable logic applications**;

Additional Key Words and Phrases: Approximate computing, precision-reconfigurable DNN acceleration

ACM Reference format:

Reena Elangovan, Shubham Jain, and Anand Raghunathan. 2022. Ax-BxP: Approximate Blocked Computation for Precision-reconfigurable Deep Neural Network Acceleration. *ACM Trans. Des. Autom. Electron. Syst.* 27, 3, Article 28 (January 2022), 20 pages.

<https://doi.org/10.1145/3492733>

This work was supported by C-BRIC, one of six centers in JUMP, a Semiconductor Research Corporation (SRC) program, sponsored by DARPA.

Authors' addresses: R. Elangovan and A. Raghunathan, School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN, USA; emails: {elangovr, raghunathan}@purdue.edu; S. Jain, IBM T. J. Watson Research Center, Yorktown Heights, NY, USA; email: shubham.jain35@ibm.com.

This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

1084-4309/2022/01-ART28 \$15.00

<https://doi.org/10.1145/3492733>

1 INTRODUCTION

Deep Neural Networks (DNNs) have become very popular in recent years due to their ability to achieve state-of-the-art performance in a variety of cognitive tasks such as image classification, speech recognition, and natural language processing [1–3]. The remarkable algorithmic performance of DNNs comes with extremely high computation and storage requirements. While these challenges span both training and inference, we focus on the latter scenario where the high computation requirements of DNN inference limit their adoption in energy- and cost-constrained devices [4].

The use of low precision has emerged as a popular technique for realizing DNN inference efficiently in hardware [5, 6]. Lowering the precision or bit-width favorably impacts all components of energy consumption including processing elements, memory, and interconnect. State-of-the-art commercial DNN hardware widely supports 8-bit precision for DNN inference, and recent research continues to explore inference with even lower precision [7–11].

Recent efforts [7, 12, 13] suggest that realizing ultra-low-precision (sub-8-bit) DNN inference with minimal accuracy degradation is quite challenging if the precision for all data-structures is scaled uniformly. Therefore, the use of variable precision (across DNNs, across layers within a DNN, and across different data structures) has gained considerable interest. For instance, HAQ [7] shows that MobileNet and ResNets require precision varying from 3 to 8 bits across network layers to match the accuracy of a full-precision network.

To support variable precision, one option is to utilize conventional fixed-precision hardware that is provisioned for the highest precision and gate the unused portions of logic to save power when computations do not use the full precision supported. However, this approach does not fully utilize the potential of aggressive precision-scaling, since the unused hardware lowers utilization. Alternatively, variable-precision DNN accelerators with bit-serial [14–16] or bit-fused [17] fabrics have been designed to support re-configurability. However, this re-configurability comes at a high cost, as the bit-serial arithmetic circuits incur significant energy and latency overheads with respect to their fixed-precision counterparts of equivalent bit-width, due to multi-cycle operation and control logic [18]. Figure 1 quantifies the energy overhead (at iso-area) incurred while performing 8-bit MAC computations using digit-serial MAC units of 2-bits and 4-bits (synthesized to 15 nm with Synopsys Design Compiler). As shown, the increased flexibility provided by digit-serial hardware is accompanied by a $2\times$ – $6\times$ higher energy cost when 8-bit operations must be performed. This limits the energy benefits that can be realized from variable-precision DNNs, where bit-width varies from 2–8 bits across layers and across networks [7, 12, 13].

To design DNN hardware that caters to variable precision requirements with minimal overheads, we leverage the intrinsic tolerance of DNNs to approximate computation [19, 20], which is the basis for reduced precision itself. Specifically, we propose Ax-BxP, an approximate computation method to execute DNN inference in which weights and activations are composed of fixed-length blocks (groups of bits), and computations are performed block-wise. Approximations are introduced by only performing a subset of the block-wise computations to enable efficient re-configurability. We present a methodology to choose the best approximation configuration for each layer in a DNN. We also propose architectural enhancements to realize Ax-BxP in a standard systolic array-based DNN inference accelerator with simple and low-overhead design modifications. We show that (i) Ax-BxP with varying approximation configurations across DNNs achieves $1.1\times$ – $1.74\times$ and $1.02\times$ – $2\times$ improvement in system-level energy and performance, respectively, and (ii) Ax-BxP with varying approximation configurations across layers within DNNs obtains $1.12\times$ – $2.23\times$ and $1.14\times$ – $2.34\times$ improvement in system-level energy and performance, respectively, in both cases with small loss in classification accuracy ($<1\%$ on average) with respect to an 8-bit **fixed-point (FxP)** baseline.

Fig. 1. Computation cost vs. re-configurability tradeoffs.

(a) Blocked Fixed-Point

(b) Blocked Computation

Fig. 2. BxP: Overview.

2 PRELIMINARIES

2.1 Blocked Fixed Point Representation

The FxP format is widely used for efficient realization of low-precision DNNs. **Blocked fixed-point (BxP)** format is an adaptation of the FxP format wherein values are partitioned into fixed-length blocks. In particular, an $(N * K)$ bit signed FxP number X_{FxP} can be represented as a BxP number X_{BxP} composed of N blocks of K bits each, as shown in Figure 2(a). The blocks of X_{BxP} are arranged in the decreasing order of significance (place value) by default, where the significance of the i th block X_i is 2^{i*K} . We assume a sign-magnitude representation wherein the most significant bit within the most significant block (X_{N-1}) is the sign bit. The magnitude of X_{BxP} , denoted $|X_{BxP}|$, can be derived from its blocks, as shown in the figure.

2.2 Blocked Multiplication

Figure 2(b) demonstrates a blocked¹ multiplication between two BxP format numbers X_{BxP} and Y_{BxP} . As shown, each block of X_{BxP} is multiplied with each block of Y_{BxP} to generate N^2 partial products (P). Subsequently, the partial products are shifted and accumulated using $N^2 - 1$ additions. Equation (1) expresses an exact blocked multiplication operation, where P_{ij} [$P_{ij} = X_i * Y_j$] represents the partial product of the i th block of X_{BxP} (X_i) and j th block of Y_{BxP} (Y_j).

$$X_{BxP} * Y_{BxP} = \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} P_{ij} * 2^{(i+j)*K} \quad (1)$$

3 APPROXIMATE BLOCKED COMPUTATION

In this section, we discuss Ax-BxP, the proposed approximate blocked computation method for designing efficient precision-reconfigurable DNNs. We first detail the key approximation concepts and introduce the Ax-BxP format for representing MAC operands. Subsequently, we present a systematic methodology for designing Ax-BxP DNNs with minimal impact on application-level accuracy. Finally, we demonstrate the integration of Ax-BxP into a standard systolic array-based DNN accelerator using simple hardware enhancements.

¹Blocked multiplication is also known in the literature as digit serial multiplication (e.g., Reference [21]) and bit-level composable multiplication (e.g., Reference [17]).

Fig. 3. Approximate blocked computation: Overview.

3.1 Approximation Method

The main idea in Ax-BxP is to perform blocked computation by performing only a subset of the required block-wise computations. Figure 3(a) illustrates the concept, where the multiplication of operands (X_{BxP} and Y_{BxP}) is performed by computing and accumulating only $L = |\Pi_{Ax}|$ out of the total of N^2 partial product terms. Formally, we characterize Ax-BxP multiplication using a set $\Pi_{Ax} \subseteq \Pi_{exact}$, wherein the final output (out-approx) is given by summing the partial products (P_s) in Π_{Ax} . Ax-BxP involves two key design choices: (i) L (the size of set Π_{Ax}) and (ii) the choice of partial products to include within the set Π_{Ax} . These design choices affect both the computational errors introduced and the energy benefits, and therefore need to be explored judiciously to produce the best possible energy-accuracy tradeoffs.

Figure 3(b) presents the energy-accuracy tradeoff provided by Ax-BxP for the AlexNet DNN on the CIFAR10 dataset, across various choices of L with fixed values of N and K (note that $K = 2, N = 3$, with $L = 9$ corresponds to exact blocked computation). As L decreases, the accuracy decreases minimally, whereas the energy benefits increase drastically. The computational accuracy reported is for the best choice of Π_{Ax} among all choices of Π_{Ax} , identified using the methodology described in Section 3.3. To estimate these energy benefits, we synthesized the exact and approximate RTL designs (described in Section 3.4) to the 15 nm technology node using Synopsys Design Compiler. The results suggest a favorable energy-accuracy tradeoff, which arises due to the typical data distribution seen in DNNs [11], wherein a majority (>90%) of the operations can be computed accurately at low-precision, but computing the remaining operations at higher precision is critical to preserving accuracy. We also evaluated the energy benefit across various values of N by fixing the values of K and L ($K = 4, L = N$). As shown in Figure 3(c), the energy of approximate block-wise computations increase linearly with N , as we require $O(N)$ block-wise partial products to be evaluated and accumulated. In contrast, the energy of exact blocked computation increases quadratically, as it requires $O(N^2)$ block-wise partial products.

For a given L , when fewer than N blocks of X_{BxP} or Y_{BxP} (or both) are used to construct Π_{Ax} , we can achieve memory footprint savings in addition to computation savings by storing only the required blocks. To leverage this, we introduce a new Ax-BxP tensor format for storing the Ax-BxP operands in the following subsection. Section 3.3 describes the *significance-* and *value-based* methods that we use for choosing the required blocks of the operands.

3.2 Ax-BxP Tensor

We introduce a new format for storing Ax-BxP operand blocks and their indices. The Ax-BxP format uses the following fields: (i) total number of blocks (N), (ii) block-size (K), (iii) the set $\mathcal{I}$ containing indices of the operand blocks chosen for Ax-BxP computation, (iv) number of chosen

Fig. 4. Ax-BxP Tensor: Memory layout.

blocks ($\tilde{N} := |I| \leq N$), and (v) data blocks (*data*), which are arranged in decreasing order of significance. The size of the *data* field is $\tilde{N} * K$ bits. Note that, during exact blocked computation, $\tilde{N} = N$ and I is not required.

An Ax-BxP Tensor (typically, the weights or activations of an entire layer) is composed of scalar elements in the Ax-BxP Format, where the elements share common values of the parameters (viz., $N, K, \tilde{N}, I$). Ax-BxP Tensor is presented as a C++ template class in Figure 4, along with an illustrative memory layout of the different fields. Since the space required to store the parameters is amortized across an entire tensor, the associated memory footprint is negligible. Furthermore, when $\tilde{N} < N$ the size of the *data* field is reduced, resulting in savings in memory footprint and memory traffic in addition to computation.

3.3 Design Methodology

In this subsection, we present the key design considerations involved in approximating DNNs using Ax-BxP and the methodology we use to select the Ax-BxP format for each layer in the network. We first characterize the Ax-BxP design space. Subsequently, we provide pruning techniques to reduce the complexity of exploring this design space and algorithms for the systematic design of Ax-BxP DNNs.

3.3.1 Design Space. For a given **bit-width (BW)**, where $BW = N * K$, an Ax-BxP MAC operation (characterized by the set Π_{Ax}) can be designed in numerous ways. We define Ω as a set enumerating all possible ways of constructing Π_{Ax} . Equation (2) expresses the size of Ω ($|\Omega|$), which is determined by free variables L and N . As shown, for a given BW, we are free to choose N (i.e., number of blocks) to be an integer from 1 to BW. Subsequently, we can select the approximation-level by determining L , i.e., number of partial products to be used during MAC operations. The value of L can be 1 to N^2 , where $L = N^2$ represents an exact blocked computation. Last, there are $\binom{N^2}{L}$ ways of selecting L out of N^2 partial products.

$$|\Omega| = \sum_{N=1}^{BW} \sum_{L=1}^{N^2} \binom{N^2}{L} \quad (2)$$

3.3.2 Design Space Constraints. To reduce the search space $|\Omega|$, which is exponential in N , we put forth the following arguments to bound $|\Omega|$ by constraining BW, K, N , and L :

- **Bitwidth (BW):** Since a bit-width of 8 for both activations and weights is sufficient to preserve accuracy during inference [22], we constraint $BW \leq 8$.
- **Bits in a block (K):** We also bound K such that $1 < K \leq 4$. By setting $K > 1$, we avoid the latency and energy overheads associated with bit-serial ($K = 1$) implementations [17]. Moreover, we introduce an upper-bound on K ($K \leq 4$) to avoid $N = 1$ (i.e., an FxP implementation).

Fig. 5. Design search space for $L = 4, N = 4$: Illustration.

- *Number of Blocks (N):* We set $N = \lceil BW/K \rceil$. Therefore, for $BW \leq 8$ and $1 < K \leq 4$, the allowed values of N are 2, 3, and 4.
- *Size of set Π_{Ax} (L):* Last, we constraint $L \leq N$ based on the energy-accuracy tradeoffs discussed in Section 3.1 and shown in Figure 3. We found that $L \leq N$ provides ample design choices, wherein we can obtain significant energy benefits with minimal impact on computational accuracy. Apart from a reduction in design space, bounding $L \leq N$ also helps in minimizing the design complexity of both the control logic and Ax-BxP PEs at the systolic-array level, mitigating the associated reconfiguration overheads (discussed further in Section 3.4).

Equation (3) expresses the size of reduced design space (Ω_c) obtained after constraining the variables BW, K, N , and L .

$$|\Omega_c| = \sum_{N=2,3,4} \sum_{L=1}^N \binom{N^2}{L} \quad (3)$$

Using numerical methods, we evaluate $|\Omega_c|$ to be 2,655 for a single DNN layer. For a DNN with n layers, $|\Omega_c| = 2,655^n$. This space is too large to be exhaustively explored for practical DNNs, especially, since re-training is needed to alleviate the accuracy degradation caused by approximation. We therefore propose techniques to prune the search space by eliminating sub-optimal configurations.

3.3.3 Design Space Pruning. We prune the search space by restricting the contents of the set Π_{Ax} for a given L . Figure 5 illustrates the possible choices of Π_{Ax} , wherein W_{BxP} and A_{BxP} are the Ax-BxP weight and activation tensors, respectively, of a DNN layer. Further, $\tilde{N}_W$ ($\tilde{N}_A$) and I_W (I_A) represent number of blocks and block indices, respectively, of W_{BxP} (A_{BxP}) used for computing the MAC operation ($W_{BxP} * A_{BxP}$). As shown, there are a variety of ways of choosing 4 (L) out of 16 (N^2) partial products. The selected partial products (shown in green) cast a shape on the 2D array that represents all possible N^2 partial products. The shape could be scattered, irregular, or regular. In our design exploration, we restrict ourselves to regular shapes, which substantially reduces the search space complexity. Formally, restricting to regular shapes constraints L as shown in Equation (4):

$$L = \tilde{N}_W * \tilde{N}_A. \quad (4)$$

Based on previous studies that show activations to be more sensitive to precision-scaling than weights during DNN inference operations [10, 22], we further prune the search space such that $\tilde{N}_A \geq \tilde{N}_W$. In other words, we never select configurations (e.g., the right-most configuration in Figure 5), wherein a weight operand (W_{BxP}) has more blocks than an activation operand (A_{BxP}). Equation (5) shows the size of the design search space (Ω_{c+p}) obtained after pruning, wherein

Fig. 6. Significance of the non-zero activation blocks in a layer of quantized AlexNet.

Fig. 7. Operand blocks selection heuristics: Error analysis.

$\binom{N}{\tilde{N}_W}$ and $\binom{N}{\tilde{N}_A}$ are the number of possible ways of selecting $\tilde{N}_W$ and $\tilde{N}_A$ blocks, respectively, out of N blocks. It is worth mentioning that the use of regular configurations enables an efficient systolic array implementation so we can accrue benefits from Ax-BxP due to reduced computation, memory footprint, and memory traffic.

$$|\Omega_{c+p}| = \sum_{N=2,3,4} \sum_{\tilde{N}_A=1}^N \sum_{\tilde{N}_W=1}^{\tilde{N}_A} \binom{N}{\tilde{N}_A} \binom{N}{\tilde{N}_W} \quad (5)$$

3.3.4 Design Heuristics. Next, we present the two heuristics, viz., *static-idx* and *dynamic-idx*, to select $\tilde{N}_W$ ($\tilde{N}_A$) blocks of W_{BxP} (A_{BxP}).

In the *static-idx* heuristic, the operand blocks are chosen in a *significance-aware* manner where the blocks of higher significance are always chosen over the blocks of lower significance. For a given $\tilde{N}$, we first find the index of the most-significant non-zero block of the operand tensor and choose the next $\tilde{N}$ consecutive blocks in the decreasing order of significance. Since the blocks of data in the *data* field of the Ax-BxP tensor are arranged in the decreasing order of significance by default, we require only the start index ($\mathcal{I}[N-1]$) or the end index ($\mathcal{I}[0]$) to determine the indices of all the blocks. Recall that $\mathcal{I}$ is common to all the scalar elements of an operand tensor.

Note that when $\mathcal{I}$ is chosen using the *static-idx* heuristic, the small-valued scalar elements in an operand tensor cannot be represented with sufficient resolution. For instance, consider the activation histogram (generated using the ImageNet dataset) of a quantized AlexNet layer shown in Figure 6, where the activation values are represented using 2 blocks of 2 bits each, (i.e., $N = 2, K = 2$). In the histogram, both blocks may be non-zero for large values that reside in bins $[2, 4]$, whereas only the least significant block is non-zero for the values in bins $[-3, 1]$. Therefore, when $\tilde{N} = 1$, the *static-idx* heuristic chooses the most significant block for each scalar element, and the small values in bins $[-3, 1]$ are approximated to zero. To avoid this loss of information in the smaller-value scalar elements, we introduce the *dynamic-idx* heuristic, where the set $\mathcal{I}$ is chosen specifically for each scalar element in a tensor.

In the *dynamic-idx* heuristic, $\mathcal{I}$ (i.e., the index of the most significant non-zero block) is determined for each scalar value of a tensor. Subsequently, each scalar is represented using $\tilde{N}$ consecutive blocks starting from its most-significant non-zero block. Note that the dynamic heuristic is based on both significance and value of the blocks. Figure 7 quantifies the advantage of the *dynamic-idx* heuristic over the *static-idx* heuristic by showing the error distribution in the activations (obtained using the ImageNet dataset) for an Alexnet layer. As shown, *dynamic-idx* can achieve noticeably lower error in comparison to *static-idx*. However, memory savings are less for the *dynamic-idx* heuristic in comparison to the *static-idx* heuristic, as we need to store $\mathcal{I}$ for each scalar element. For a given N and $\tilde{N}$, the overhead of appending $\mathcal{I}$ to each data element is $\lceil \log_2((N - \tilde{N}) + 1) \rceil$ bits. For example, if $K = 2, N = 4, \tilde{N} = 2$, we have 2 additional index bits for

Table 1. Design Space for Approximate Blocked Computation

Block-size	Ax-BxP configuration = $\{(K, \tilde{N}_W, \tilde{N}_A)\}$
$K = 2$	$\{(2, 1, 4), (2, 1, 3), (2, 2, 2), (2, 1, 2), (2, 1, 1)\}$
$K = 3$	$\{(3, 1, 3), (3, 1, 2), (3, 1, 1)\}$
$K = 4$	$\{(4, 1, 2), (4, 1, 1)\}$

every 4 ($K \cdot \tilde{N}$) compute bits. Therefore, the overall memory footprint decreases by $\sim 25\%$ (from 8 bits to 6 bits) compared to 8-bit fixed point. However, with *static-idx*, the same configuration would require only 4 bits, resulting in a $\sim 50\%$ memory savings.

Since L and the set Π_{Ax} can be derived for a given K , $\tilde{N}_W$, and $\tilde{N}_A$, the design space Ω can now be restricted to the set of all 3-tuples $\{K, \tilde{N}_W, \tilde{N}_A\}$ that satisfy all the constraints discussed thus far. The Ax-BxP configurations in Ω are listed against the block-size K in Table 1. Furthermore, we define two modes of Ax-BxP, *Static* and *Dynamic*, where the operand blocks are chosen using the *static-idx* heuristic and the *dynamic-idx* heuristic, respectively.

3.3.5 Designing DNNs using Ax-BxP. We now present a systematic methodology to design DNNs using Ax-BxP. Algorithm 1 describes the pseudo code that we utilize to identify best Ax-BxP configuration for each data-structure of each DNN layer. It takes a pre-trained DNN model (DNN_{Fxp}), a training dataset (Tr_{data}), a target block size (K_{tgt}), and a limit on allowed accuracy degradation (γ) as inputs and produces an Ax-BxP DNN (DNN_{AxBxP}) as output. We first utilize Tr_{data} to evaluate the baseline network accuracy (line 4) and construct data-value histograms ($DsHist_{list}$) of each data-structure within the network (line 5). Next, we identify the best Ax-BxP configuration for a DNN layer using the histograms of the associated weight (W_{Fxp}) and activation (A_{Fxp}) pair (lines 6–12). As detailed in Algorithm 1, to obtain best Ax-BxP configuration, we first form a pruned search space (Ω_{c+p}) (line 7), and subsequently, explore the choices within Ω_{c+p} to find the best Ax-BxP configuration, which is represented by N_W , $\tilde{N}_W$, N_A , and $\tilde{N}_A$ (line 8). Next, data-structures are converted to Ax-BxP tensor using the Convert-To-AxBxP function (lines 9–10) and inserted into DNN_{AxBxP} network (line 11). Once all data-structures are converted to Ax-BxP tensor, we re-train DNN_{AxBxP} until the network accuracy is within the desired degradation limit ($< \gamma$) or the maximum allowed trained epochs (maxEpoch) is exceeded.

We determine the best Ax-BxP configuration layer-by-layer for a given DNN as shown in Algorithm 1 (lines 6–12). We start with a pre-trained DNN_{Fxp} , where exact computations are performed in each layer. Once we find the best Ax-BxP configuration for a layer, we convert the operands of that layer to Ax-BxP format and proceed to the next layer. Algorithm 2 describes the methodology to choose the best Ax-BxP configuration for a given DNN layer, block-size K_{tgt} and target accuracy degradation γ . We set the operands to each of the Ax-BxP formats in Ω_{c+p} (lines 4–5), and evaluate the resulting DNN_{AxBxP} (line 6). We perform the evaluation on a subset of the training dataset to speed up the search. In our experiments with the ImageNet dataset, we perform re-training for 1 epoch on 120,000 images followed by testing on 5,000 images to determine accuracy. Subsequently, we choose the first encountered Ax-BxP configuration that provides the desired accuracy (line 7).

We note that after the design space pruning described in Section 3.3.3, the design space for finding the best approximation configuration is drastically reduced to 2, 3, and 5 choices per layer for $K = 4, 3$, and 2, respectively. Since we perform a greedy search on a layer-by-layer basis, the size of the pruned design space for a DNN with N layers is at most $5N$. For C choices, a DNN with N layers and evaluation time of T seconds per choice, the time taken to find the best Ax-BxP configuration is $C * N * T$.

ALGORITHM 1: Designing Ax-BxP DNN

```

1: Input:{ $DNN_{Fxp}$ : Pre-trained FxP DNN,  $Tr_{data}$ : Training dataset,  $\gamma$ : Max Accuracy Loss,  $K_{tgt}$ : Target
   block size}
2: OUTPUT: Approximate Blocked DNN
3:  $DNN_{AxBxP} = DNN_{Fxp}$  /* initialize to the given FxP DNN */
4:  $Acc_{Fxp} = \text{computeAccuracy}(DNN_{Fxp})$ 
5:  $DsHist_{list} = \text{getDist}(DNN_{Fxp}, Tr_{data}) \forall \text{ datastructures}$ 
6: for each [Weight ( $W_{Fxp}$ ), Activations ( $F_{xp}$ )] pair  $\in DsHist_{list}$ 
7:    $\Omega_{c+p} = \text{formPrunedSearchSpace}(W_{Fxp}, A_{Fxp}, K_{tgt})$  /* Form the pruned search space of Ax-BxP
   configurations */
8:    $(N_W, \tilde{N}_W, N_A, \tilde{N}_A) = \text{getBestConfig}(\Omega_{c+p}, DNN_{AxBxP}, K_{tgt}, \gamma)$  /* Search for the best Ax-BxP
   configuration */
9:    $W_{AxBxP} = \text{Convert-To-AxBxP}(W_{Fxp}, N_W, \tilde{N}_W, K_{tgt})$ 
10:   $A_{AxBxP} = \text{Convert-To-AxBxP}(A_{Fxp}, N_A, \tilde{N}_A, K_{tgt})$ 
11:  insert-AxBxP-Tensors( $DNN_{AxBxP}, W_{AxBxP}, A_{AxBxP}$ )
12: end for
13: numEpochs = 0
14: while  $\gamma < (Acc_{Fxp} - \text{computeAccuracy}(DNN_{AxBxP}))$  and numEpochs < maxEpoch /* Re-train to
   recover accuracy */
15:   Ax-BxP-Aware-Training( $DNN_{AxBxP}, Tr_{data}$ )
16:   numEpochs++
17: return  $DNN_{AxBxP}$ 

```

ALGORITHM 2: getBestConfig

```

1: Input: { $\Omega_{c+p}$ : Pruned Design Space,  $DNN_{AxBxP}, K_{tgt}$ : Target block size,  $\gamma$ : Max Accuracy Loss}
2: OUTPUT: Best Ax-BxP Config
3: while  $\gamma < (Acc_{Fxp} - Acc_{AxBxP})$ 
4:   Ax-BxP-config =  $\Omega_{c+p}.\text{pop}()$ 
5:   Convert-and-insert-AxBxP-Tensors( $DNN_{AxBxP}, W_{AxBxP}, A_{AxBxP}, \text{Ax-BxP-config}$ )
6:    $Acc_{AxBxP} = \text{evaluate}(DNN_{AxBxP})$ 
7: return Ax-BxP-config

```

ALGORITHM 3: Convert-To-AxBxP

```

1: Input:{ $X_{Fxp}$ : FxP tensor, ( $K, N_X, \tilde{N}_X$ ): Ax-BxP configuration}
2: OUTPUT: Ax-BxP tensor
3: For each scalar x in  $X_{Fxp}$ 
4:    $Block_{list} = \text{Get-Significance-Sorted-Blocks}(x, K, N_X)$ 
5:    $I_x = \text{get-idx-first-Non-Zero-block}(Block_{list})$ 
6:    $X_{AxFxp} = \text{pick-insert-Blocks-in-Range}(I_x, I_x - \tilde{N}_X)$ 
7: return  $X_{AxBxP}$ 

```

Algorithm 3 outlines the pseudo code for converting FxP tensors to Ax-BxP tensors using the *dynamic-idx* heuristic. It takes an FxP tensor (X_{Fxp}) and Ax-BxP configuration ($K, N_X, \tilde{N}_X$) as inputs and produces an Ax-BxP tensor (X_{AxBxP}) as output. A key function of this algorithm is to determine indices (I_x) of the chosen blocks. To achieve this objective, we first convert fixed point scalars to blocked fixed-point scalars (line 2) and subsequently pick $\tilde{N}_X$ contiguous blocks starting

Fig. 8. Ax-BxP accelerator architecture.

from the first non-zero block. Next, the chosen blocks and indices are inserted into the Ax-BxP tensor (line 4). After all scalars have been converted the X_{Ax-BxP} tensor is returned (line 5).

3.4 Ax-BxP DNN Accelerator

Figure 8 shows the proposed Ax-BxP DNN accelerator, which is a conventional two-dimensional systolic-array-based DNN accelerator with enhancements such as Ax-BxP PEs, control logic, and peripheral units (ToAx-BxP) that support Ax-BxP computation. We design the Ax-BxP DNN accelerator for a fixed K (although K is a parameter in the RTL, it is specified at synthesis time). While the proposed approach can be applied with any DNN dataflow, for illustration, we focus on the output stationary dataflow. The control logic partitions the Ax-BxP operands into blocks and determines the corresponding shift amounts. These operand blocks and shift amounts are used to perform the Ax-BxP MAC operations in the Ax-BxP PEs. Finally, the output activations computed by the Ax-BxP PEs are converted into the Ax-BxP format by the ToAx-BxP logic using the methodology described in Algorithm 3 (see Section 3.3.5). We now discuss the design of the Ax-BxP PE and the Control logic in detail.

3.4.1 Ax-BxP Processing Element (Ax-BxP PE). The Ax-BxP PEs contain $N = \lceil 8/K \rceil$ signed multipliers of bitwidth $K + 1$ and N shifters to support the Ax-BxP computations. The partial products generated by the multipliers are shifted by the shift amounts determined by the control logic and are accumulated at high-precision in the 32-bit accumulator to generate the output activations. It is worth noting that as L decreases, the throughput achieved by Ax-BxP PEs increases, since a fixed number (N) of multiplications are performed in each cycle.

For a given K and a given mode of Ax-BxP, Ax-BxP PEs support all the Ax-BxP configurations in Ω_{c+p} with block size K by allowing different shift amounts. The different shift amounts are realized in the shifters using multiplexers. The number of multiplexers increases with the number of unique shift amounts to be supported, resulting in increases in energy and area overheads. It is straightforward to show that larger number of unique shift amounts must be supported by the Ax-BxP PE during dynamic Ax-BxP compared to static Ax-BxP. Therefore, the energy and area of Ax-BxP PEs are comparatively larger in the case of dynamic Ax-BxP.

3.4.2 Control. The control logic in the Ax-BxP DNN accelerator partitions each of the Ax-BxP operands A_{Ax-BxP} and W_{Ax-BxP} into signed blocks of $K + 1$ bits and determines the shift amounts corresponding to each of these blocks. The shift amounts are determined based on the indices of

the operand blocks, where the shift amount corresponding to the i th block is computed as $i * K$. The indices of the operand blocks of A_{AxBxP} (W_{AxBxP}) are determined from the parameters $\tilde{N}_A$ and $\mathcal{I}_A$ ($\tilde{N}_W$ and $\mathcal{I}_W$). Note that these parameters are obtained from the AxBxP format of the operands.

For a given layer of the DNN, in the static mode, the parameters $\tilde{N}_A$ and $\mathcal{I}_A$ ($\tilde{N}_W$ and $\mathcal{I}_W$) are fixed for all the scalar elements of A_{AxBxP} (W_{AxBxP}) and therefore, are broadcast to the control blocks at the start of a layer's computations. As discussed in Section 3.3.4, $\mathcal{I}_A$ ($\mathcal{I}_W$) is the index of the most significant block of A_{AxBxP} (W_{AxBxP}), from which the indices of $\tilde{N}_A$ ($\tilde{N}_W$) consecutive blocks are derived. In the dynamic mode, the parameter $\tilde{N}_A$ ($\tilde{N}_W$) is fixed while $\mathcal{I}_A$ ($\mathcal{I}_W$) varies within the operand tensor A_{AxBxP} (W_{AxBxP}). Therefore, the parameter $\tilde{N}_A$ ($\tilde{N}_W$) is broadcast to the control blocks at the start of a layer's computations. The parameters $\mathcal{I}_A$ and $\mathcal{I}_W$ are obtained from each scalar element of A_{AxBxP} and W_{AxBxP} , respectively, since the indices of the chosen operand blocks are stored individually for every scalar element.

It is worth noting that the design complexity of both the Ax-BxP PEs and the control logic depends on the number of unique shift amounts to be supported, which in turn depends on L . Therefore, constraining L to be $\leq N$ minimizes the design complexity while still achieving good classification accuracy (as shown in Section 5).

4 EXPERIMENTAL METHODOLOGY

In this section, we present the experimental methodology used to evaluate Ax-BxP.

Accuracy Evaluation: We evaluate Ax-BxP using three state-of-the-art image recognition DNNs for the ImageNet dataset, viz., ResNet50, MobileNetV2, and AlexNet. The ImageNet dataset has 1.28M training images and 50,000 test images. We use the entire test dataset for the accuracy evaluation. We perform up to five epochs of re-training for all the Ax-BxP configurations considered. We use an Intel Core i7 system with NVIDIA GeForce 2080 (Turing) graphics card for the simulations.

System Energy and Performance Evaluation: We design the Ax-BxP DNN accelerator by expanding the conventional systolic array accelerator modeled in ScaleSim [23] to include enhancements such as Control logic, ToAxBxP logic, and the AxBxP PEs, all synthesized to the 15 nm technology node using Synopsys Design Compiler. We consider an output-stationary systolic array of size 32×32 and on-chip scratch-pad memory of 2 MB, operating at 1 GHz. The on-chip memory is modeled using CACTI [24]. We design our baseline FxP8 accelerator, also synthesized to the 15 nm technology node using Synopsys Design Compiler, as a conventional systolic array with FxP8 PEs that can implement 8-bit MAC operations. The system-level energy and performance benefits of the Ax-BxP accelerator are evaluated against the FxP8 baseline. To evaluate the benefits over a mixed-precision baseline, we adopt the HAQ [7, 25] precision configuration and design a power-gated FxP8 baseline with power-gated FxP8 PEs, i.e., we design the PEs for a maximum precision of 8 bits and power-gate the unused portions during lower-precision computations. Additionally, we evaluate the benefits of the proposed Ax-BxP in the Bit-Fusion accelerator [17].

5 RESULTS

In this section, we demonstrate the energy and performance benefits of Ax-BxP at the **processing element (PE)** level and system level using the proposed Ax-BxP DNN accelerator. We also demonstrate the benefits of Ax-BxP in the Bit-Fusion [17] accelerator.

5.1 PE-level Energy and Area Benefits

Figure 9 shows the energy and area benefits of the Ax-BxP PE for $K = 2, 3, 4$ when compared to an **8-bit fixed-point (FxP8)** PE. On average, the energy and area benefits of the Ax-BxP PE

Fig. 9. Computation energy and area benefits over an FxP8 baseline.

Fig. 10. Computation energy breakdown (left) and area breakdown (right) for Ax-BxP and the FxP8 baseline.

in dynamic mode are $1.69\times$ and $1.12\times$, respectively. In static mode, the average energy and area benefits are $1.87\times$ and $1.25\times$, respectively.

Recall from section 3.4 that for a given K , the Ax-BxP PEs can support any Ax-BxP configuration by allowing shifts by different amounts. The energy and area overheads of Ax-BxP PEs increase with an increase in the number of unique shift amounts to be supported. The number of unique shift amounts to be supported is proportional to $N = \lceil 8/K \rceil$. Therefore the energy benefits decrease as K decreases for both static and dynamic modes of operation. Furthermore, since greater numbers of shift amounts need to be supported for the dynamic mode, the energy benefits during dynamic Ax-BxP are lower than static Ax-BxP for all K . Figure 10 (left) shows the energy breakdown of multipliers, shifters and adders, and the accumulator in Ax-BxP PEs for all K in both static and dynamic modes of Ax-BxP. For a given value of K , we observe that the overhead due to re-configurability (i.e., shifters and adders) is greater in dynamic mode vs. static mode. Figure 10 (right) shows the area breakdown of the Ax-BxP PEs and the FxP8 PE. Similar to energy, we find that for a given K , the re-configurability overhead is greater in dynamic vs. static mode. The energy and area overheads of re-configurability are larger for smaller K because of the greater number of shift amounts to be supported.

5.2 System Benefits and Accuracy of Dynamic Ax-BxP with Network-level Precision Re-configuration

Figure 11 shows the system-level benefits of dynamic Ax-BxP inference compared to the FxP8 baseline at iso-area. In this experiment, for a given network, we maintain the Ax-BxP configurations uniform across its layers. We find that the best Ax-BxP configuration in Ω_{c+p} for a given K , i.e., the configuration that provides maximum energy benefits with small accuracy loss ($\sim 1\%$), varies across the networks considered.

Table 2 shows the best Ax-BxP configuration for a each K and each network considered. The proposed Ax-BxP DNN accelerator efficiently supports varying configurations across networks and achieves system-level energy reduction of $1.19\times$ – $1.78\times$, $1.01\times$ – $1.72\times$, and $1.1\times$ – $2.26\times$ for

Fig. 11. System-level benefits with dynamic Ax-BxP.

Table 2. Best Dynamic Ax-BxP Configurations

Block-size	AlexNet	ResNet50	MobileNetV2
$K = 4$	(4, 1, 1)	(4, 1, 2)	(4, 1, 2)
$K = 3$	(3, 1, 1)	(3, 1, 2)	(3, 1, 2)
$K = 2$	(2, 1, 2)	(2, 1, 2)	(2, 2, 2)

$K = 4, 3$, and 2 , respectively, compared to the FxP8 baseline at iso-area. Figure 11 shows the breakdown of the normalized system-energy into off-chip memory access energy, on-chip Buffer access energy and the Systolic-Array energy. We achieve a significant reduction in each of these components in our proposed Ax-BxP accelerator across DNNs and across the approximation configurations considered.

The $1.04\times$ – $3.41\times$ reduction in the systolic-array energy stems primarily from the Ax-BxP PE benefits discussed in Section 5.1. Additionally, for a given K , when L decreases, the throughput of the Ax-BxP DNN accelerator increases, as discussed in Section 3.4. Therefore, the overall inference cycles reduces, resulting in further reduction in the systolic array energy. However, the Imagenet classification accuracy decreases with smaller L . Furthermore, by selecting to store a subset of the operand blocks, we achieve $1.01\times$ – $1.6\times$ and $1.03\times$ – $2.17\times$ reduction in off-chip and on-chip memory access energy, respectively. For a given K , the memory-access energy decreases as L decreases. This is because of the reduction in memory footprint and the number of accesses despite the overhead of storing the $\lceil \log_2((N - \tilde{N}) + 1) \rceil$ bits of operand block indices during dynamic Ax-BxP.

The performance benefits during dynamic Ax-BxP compared to FxP8 at iso-area are also shown in Figure 11. We obtain $1.13\times$ – $2.02\times$, $1.01\times$ – $2.08\times$, and $1.02\times$ – $2.59\times$ performance benefits for $K = 4, 3$, and 2 , respectively, for the configurations listed in Table 2. For a given K , smaller L results in increased throughput, which in turn results in increased performance.

5.3 System-level Energy Benefits and Accuracy with Intra-network Precision Reconfiguration

The benefits of Ax-BxP can be further highlighted in the context of variable-precision DNNs that have different layer-wise precision requirements. The uniform Ax-BxP configuration of $(2, 1, 1)$ provides the maximum system-level energy benefits. However, it suffers from significant accuracy degradation. To improve the classification accuracy and to maximize the system-benefits, we vary

Fig. 12. System-level benefits and ImageNet accuracy for mixed-precision networks.

the precision in a coarse-grained manner across the layers of these DNNs with $K = 2$. We use the Ax-BxP configurations (2, 1, 2) and (2, 1, 1) for ResNet50 and AlexNet and (2, 2, 2) and (2, 1, 2) for MobileNetV2. The layer-wise precision configurations considered are shown in Figure 12. Following previous studies [10, 22, 26], we do not approximate the first and last layer computations, as they have been shown to impact the classification accuracy significantly. We achieve $1.12\times$ – $2.23\times$ reduction in system energy and $1.13\times$ – $2.34\times$ improvement in system performance compared to the FxP8 baseline with small loss in classification accuracy.

Next, we compare Ax-BxP against the HAQ [7, 25] mixed-precision configuration implemented on a conventional systolic array accelerator with power-gated PEs. We design the PEs to support the worst-case precision of 8-bits and power-gate the unused portions during sub-8-bit computations. Figure 13 shows the normalized system-energy breakdown, performance, and ImageNet accuracy of Ax-BxP and HAQ implementations. We observe that Ax-BxP achieves $1.04\times$ – $1.6\times$ system-energy reduction and $1.1\times$ – $2.34\times$ performance improvement compared to HAQ. The memory access energy (off-chip + on-chip) of Ax-BxP is $0.95\times$ and $1.06\times$ of the HAQ memory-access energy for ResNet50 and MobileNetV2, respectively. The small overhead in case of MobileNetV2 is caused by storing the Ax-BxP operand block indices. Despite this overhead, Ax-BxP achieves superior system benefits compared to HAQ by substantially reducing the systolic-array energy by $1.32\times$ – $2.95\times$. This is the result of lower overall inference cycles for Ax-BxP, achieved through superior systolic-array utilization compared to the power-gated FxP8 implementation of HAQ. As shown in Figure 13, HAQ does not provide any performance improvement compared to FxP8, because the power-gated PEs cannot increase the throughput of the systolic-array.

5.4 Benefits of Ax-BxP in the Bit-fusion Accelerator

To demonstrate the broad applicability of Ax-BxP, we implement it on top of the Bit-fusion accelerator architecture [17]. The energy and performance benefits of dynamic Ax-BxP in the Bit-fusion accelerator compared to exact FxP8 computations are shown in Figure 14. We have considered the Ax-BxP configurations with $K = 2$, since the bit-bricks in Bit-fusion PE are designed for a block-size of 2. By performing approximations using Ax-BxP, we could achieve energy benefits of upto $3.3\times$ and performance benefits of up to $4.6\times$ in the Bit-fusion accelerator. The Bit-fusion PEs

Fig. 13. Comparison to HAQ precision configuration implemented in a power-gated FxP8 systolic array.

Fig. 14. Ax-BxP Benefits in Bit-fusion.

achieve a comparatively higher throughput for a given L . The maximum increase in throughput is $16\times$ when $L = 1$, resulting in the $4.6\times$ benefits for the configuration $(2, 1, 1)$.

We further evaluate the benefits of Ax-BxP over HAQ when both are implemented on the Bit-fusion accelerator. As shown in Figure 15, Ax-BxP achieves $1.01\times$ – $1.31\times$ reduction in system energy and $1.09\times$ – $1.75\times$ improvement in system performance compared to HAQ. Similar to the discussion in Section 5.3, the small overhead in the memory access energy (on-chip + off-chip) of Ax-BxP compared to HAQ in MobileNetV2 ($1.06\times$) can be attributed to the overhead in storing the Ax-BxP operand block-indices. However, Ax-BxP achieves $1.25\times$ to $1.98\times$ reduction in systolic-array energy, which results in an overall reduction in system energy compared to HAQ. Further, Ax-BxP achieves higher throughput than HAQ as a result of superior utilization of the on-chip bit-bricks (2-bit signed multipliers) of the Bit-fusion accelerator. For instance, in layers 42 to 52 of MobileNetV2, HAQ [7] multiplies 6-bit weights with 6-bit activations, requiring 9 bit-bricks, whereas Ax-BxP (Figure 12) performs these computations using 2 blocks of weights and 2 blocks of activations, requiring only 4 bit-bricks. By dynamically choosing the weight and activation blocks, Ax-BxP minimizes the accuracy degradation (shown in Figure 12).

5.5 System Benefits and Accuracy of Static Ax-BxP

Figure 16 shows the system benefits of static Ax-BxP compared to the FxP8 baseline at iso-area. The energy benefits during static Ax-BxP are greater than the dynamic Ax-BxP across networks and across configurations. This is because the Ax-BxP PEs are simpler in terms of the number of shift amounts to be supported during the static mode, compared to the dynamic mode. Furthermore, the memory footprint of the operands are lower during the static mode compared to the dynamic

Fig. 15. Comparison of AxBxP and HAQ when implemented in the Bit-fusion Accelerator.

Fig. 16. System-level benefits with static Ax-BxP.

mode, since the cost of storing and fetching $\mathcal{I}_W$ and $\mathcal{I}_A$ are amortized across the tensors W and A , respectively. As a result, the memory access energy is lower and the compute energy is lower during static Ax-BxP. However, for equal re-training effort (five epochs), the ImageNet accuracy degradation with static Ax-BxP is significantly higher than dynamic Ax-BxP across configurations and networks. The performance benefits with static Ax-BxP are also shown in Figure 16. The performance benefits during static Ax-BxP are greater than dynamic Ax-BxP across networks and Ax-BxP configurations. This is because, in static mode, the Ax-BxP PEs are smaller, which is exploited in the iso-area design to achieve higher throughput.

6 RELATED WORK

The high computation and storage demands posed by DNNs have motivated several efforts to focus on precision scaling. Many of the early efforts [22, 26–29] on precision scaling are effective in small networks, but they suffer significant accuracy degradation in large networks.

More recent efforts [8–11, 30, 31] have developed advanced quantization techniques and training methodologies that work well for a wide range of networks and effectively reduce the bit-widths of data-structures to below 8 bits. Notably, PACT [10] has demonstrated successful inference using only 2-bit precision for weights and activations, except in the first and last layers, which are evaluated at 8-bit precision. Other efforts such as BQ [9] and WRPN [8] also achieve good inference accuracy using 2-bit weights and activations by using techniques such as Balanced Quantization and Model scaling, respectively. Deep Compression [30] employs a combination of pruning, quantization, and Huffman coding to reduce the model size. Bi-Scaled DNN [11] and Compensated DNN [31] leverage the value statistics in DNNs and design number-formats and error compensation schemes that effectively reduce the overall bit-width. Although these efforts enable DNN inference with ultra-low (below 8-bit) precision, a common precision is not optimal across networks

or even across layers within a network. For example, it is a common practice to retain the first and last layers at higher precision while quantizing intermediate layers to very low precision to preserve accuracy. This trend is carried further by works like HAQ [7], which demonstrate that the minimum precision requirement varies within a network, across different layers.

Since varying precision requirements are inherent across DNNs, recent efforts [14, 15, 17] have focused on the design of precision re-configurable hardware. BISMO [14] proposes a parallelized bit-serial architecture that offers maximum flexibility in terms of the bit-widths it can support. Stripes [15] is a similar work that uses bit-serial hardware design to support variable precision computation. Albeit offering maximum flexibility, the performance of bit-serial hardware is limited by high latency and energy caused by its serial nature. However, fixed-precision hardware needs to be designed to support maximum precision and hence is over-designed for applications with low precision requirement. Instead of performing computations serially at the granularity of bits, Bit-fusion [17] explores serial computation at the granularity of a group of bits and demonstrates superior energy benefits compared to state-of-the-art bit-serial and fixed-precision accelerators. However, *none of these efforts explore the use of approximations to improve the efficiency of variable-precision hardware*. Exploiting the resilience of DNNs to approximations, we propose approximate blocked computation as a next step that is complementary to previous efforts. Computations are performed block-wise, where blocks are a group of bits of fixed length. Our approximation methodology reduces the number of block-wise computations that need to be performed while maintaining the inference accuracy.

The design of approximate multipliers has been extensively explored in the literature. These efforts can be broadly classified into three categories—efforts that focus on the design of general-purpose approximate circuits, efforts that approximate partial-product accumulation, and efforts that approximate partial product generation. Efforts such as References [32–35] focus on general-purpose approximate circuit design using voltage over-scaling and logic simplification techniques. The logic simplification approach used in these methods eliminates transistors or gates to trade off accuracy for efficiency. A more structured way to approximate multiplier circuits is to systematically minimize the computation and accumulation of partial products.

Energy reduction during partial product accumulation can be achieved by approximate adders [36–38] or by approximate accumulation techniques [39]. While these efforts focus on minimizing the energy consumption of accumulation, we note that the multipliers are the primary sources of energy consumption during MAC operations. Several previous works [40–45] have explored the approximation of partial product generation.

In Reference [40], the authors propose a 2×2 under-designed multiplier block and build arbitrarily large power-efficient in-accurate multipliers. The inaccurate 4:2 counter in Reference [41] can effectively reduce the partial product stages of the Wallace multiplier. In Reference [42], the authors substitute multiplication with additions and shift operations by representing the integer operands as logarithms with an error correction factor. The computation-sharing multiplier proposed in References [43, 44] specifically targets computation re-use in vector-scalar products. Reference [45] proposes an approximate multiplier that performs approximation on only the multiplication of lower-order bits. These efforts achieve only computational energy benefits. In contrast, since our proposed Ax-BxP method minimizes the number of operand blocks used in computation, we achieve savings in terms of memory footprint and memory traffic in addition to computational energy savings. Other efforts, such as References [46–51] have taken a similar approach.

Operand bit-width truncation to minimize partial product generation is explored by References [46–48]. However, these efforts exhibit poor performance during small bit-width computations. In Reference [49], the authors extract an m -bit segment from an n -bit operand and perform an m -bit ($m < n$) bit multiplication, achieving significant energy benefits. However the segments must

be at least $n/2$ bits long, thus limiting the energy savings. An improvement over Reference [49] is proposed by Reference [50] that reduces the segment size beyond $n/2$ while minimizing error, enabling dynamic range multiplication. However, this approach involves complex circuitry, such as leading one-bit detectors, barrel shifters, and so on, which introduces considerable delay, area, and energy overheads that decrease the approximation benefits. The partial product perforation method proposed by Reference [51] aims at generating fewer partial products by dropping a few bits of one operand during multiplication. Since this approach reduces the bit-width (precision) of just one operand, it does not fully utilize the benefits of precision scaling. Moreover, it requires complex error correction methods that further limit the benefits of approximation method.

Additionally, none of the approximation efforts discussed thus far have focused on variable precision computations, which we have explored in our work. Reference [52] proposes dynamic range **floating-point (FP)** format to support varying precision. However, the area and power cost of supporting FP computations is much higher than **fixed-point (Fxp)** computations. In Reference [31], the authors propose error-compensation techniques for reduced-precision Fxp multiplication. A novel number format to represent dual-precision Fxp numbers is proposed in Reference [11]. Our proposed Ax-Bxp format supports a wide range of precision requirements. It enables efficient re-configurability at the block granularity while minimizing the approximation errors.

7 CONCLUSION

Efforts to design ultra-low precision DNNs suggest that the minimum bit-width requirement varies across and within DNNs. Optimally supporting such varying precision configurations in DNN accelerators is therefore a key challenge. We address this challenge algorithmically and in hardware using our proposed **Approximate Blocked Computation method, Ax-Bxp**. We demonstrate the effectiveness of Ax-Bxp in achieving good classification accuracy on the ImageNet dataset with state-of-the-art DNNs such as AlexNet, ResNet50, and MobileNetV2. Ax-Bxp provides up to $1.74\times$ and $2\times$ benefits in system energy and performance, respectively, while varying configurations across networks. Further, with varying configurations at the layer level, Ax-Bxp achieves up to $2.23\times$ and $2.34\times$ improvements in system energy and performance, respectively.

REFERENCES

- [1] J. X. Chen. 2016. The evolution of computing: AlphaGo. *Comput. Sci. Eng.* 18, 4 (2016), 4–7.
- [2] Dario Amodei, T. Brown et al. 2020. Language models are few-shot learners. arXiv:arXiv:2005.14165.
- [3] Mingxing Tan and Quoc Le. 2019. EfficientNet: Rethinking model scaling for convolutional neural networks. In *Proceedings of the 36th International Conference on Machine Learning*. 6105–6114.
- [4] S. Venkataramani, K. Roy, and A. Raghunathan. 2016. Efficient embedded learning for IoT devices. In *Proceedings of the Asia and South Pacific Design Automation Conference (ASP-DAC)*. 308–311.
- [5] V. Sze, Y. Chen, T. Yang, and J. S. Emer. 2017. Efficient processing of deep neural networks: A tutorial and survey. *Proc. IEEE* 105, 12 (Dec. 2017), 2295–2329.
- [6] Doe Hyun, N. P. Jouppi et al. 2017. In-datacenter performance analysis of a tensor processing unit. *SIGARCH Comput. Archit. News* 45, 2 (2017). DOI: <https://doi.org/10.1145/3140659.3080246>
- [7] Kuan Wang, Zhijian Liu, Yujun Lin, Ji Lin, and Song Han. 2019. HAQ: Hardware-Aware automated quantization with mixed precision. In *Proceedings of the Conference on Computer Vision and Pattern Recognition*.
- [8] Asit Mishra, Eriko Nurvitadhi, Jeffrey J. Cook, and Debbie Marr. 2017. WRPN: Wide reduced-precision networks. arXiv:arXiv:1709.01134.
- [9] Shuchang Zhou, Yuzhi Wang, He Wen, Qinyao He, and Yuheng Zou. 2017. Balanced quantization: An effective and efficient approach to quantized neural networks. *J. Comput. Sci. Technol.* 32, 4 (2017).
- [10] Jungwook Choi, Zhuo Wang, Swagath Venkataramani, Pierce I-Jen Chuang, Vijayalakshmi Srinivasan, and Kailash Gopalakrishnan. 2018. Pact: Parameterized clipping activation for quantized neural networks. arXiv:arXiv:1805.06085.

- [11] S. Jain, S. Venkataramani, V. Srinivasan, J. Choi, K. Gopalakrishnan, and L. Chang. 2019. BiScaled-DNN: Quantizing long-tailed datastructures with two scale factors for deep neural networks. In *Proceedings of the 56th ACM/IEEE Design Automation Conference (DAC)*. 1–6.
- [12] Patrick Judd, Jorge Albericio, Tayler Hetherington, Tor M. Aamodt, Natalie Enright Jerger, and Andreas Moshovos. 2016. Proteus: Exploiting numerical precision variability in deep neural networks. In *Proceedings of the International Conference on Supercomputing*.
- [13] Soheil Hashemi, Nicholas Anthony, Hokchhay Tann, R. Iris Bahar, and Sherief Reda. 2017. Understanding the impact of precision quantization on the accuracy and energy of neural networks. In *Proceedings of the Design, Automation, and Test in Europe Conference*.
- [14] Yaman Umuroglu, Lahiru Rasnayake, and Magnus Sjalander. 2018. Bismo: A scalable bit-serial matrix multiplication overlay for reconfigurable computing. In *Proceedings of the 28th International Conference on Field Programmable Logic and Applications (FPL)*. 3007–3077. DOI: [10.1109/FPL.2018.00059](https://doi.org/10.1109/FPL.2018.00059)
- [15] P. Judd, J. Albericio, T. Hetherington, T. M. Aamodt, and A. Moshovos. 2016. Stripes: Bit-serial deep neural network computing. In *Proceedings of the IEEE/ACM International Symposium on Microarchitecture*. IEEE.
- [16] Sayeh Sharify, Alberto Delmas Lascorz, Kevin Siu, Patrick Judd, and Andreas Moshovos. 2018. Loom: Exploiting weight and activation precisions to accelerate convolutional neural networks. In *Proceedings of the Design Automation Conference*. ACM.
- [17] Hardik Sharma, Jongse Park, Naveen Suda, Liangzhen Lai, Benson Chau, Joon Kyung Kim, Vikas Chandra, and Hadi Esmaeilzadeh. 2018. Bit fusion: Bit-level dynamically composable architecture for accelerating deep neural networks. In *Proceedings of the Annual International Symposium on Computer Architecture*.
- [18] V. Camus, L. Mei, C. Enz, and M. Verhelst. 2019. Review and benchmarking of precision-scalable multiply-accumulate unit architectures for embedded neural-network processing. *IEEE J. Emerg. Select. Topics Circ. Syst.* 9, 4 (2019), 697–711. <http://dx.doi.org/10.1109/JETCAS.2019.2950386>.
- [19] Swagath Venkataramani, Ashish Ranjan, Kaushik Roy, and Anand Raghunathan. 2014. AxNN: Energy-efficient neuromorphic systems using approximate computing. In *Proceedings of the International Symposium on Low Power Electronics and Design*. 2732.
- [20] Vinay K. Chippa, Swagath Venkataramani, Srimat T. Chakradhar, Kaushik Roy, and Anand Raghunathan. 2013. Approximate computing: An integrated hardware approach. In *Proceedings of the Asilomar Conference on Signals, Systems and Computers*. 111–117. <http://dx.doi.org/10.1109/ACSSC.2013.6810241>.
- [21] R. I. Hartley and K. K. Parhi. 1995. *Digit-serial Computation*. Springer, Boston, MA. DOI: <https://doi.org/10.1007/978-1-4615-2327-7>
- [22] S. Zhou et al. 2016. DoReFa-Net: Training low bitwidth convolutional neural networks with low bitwidth gradients. *arXiv preprint arXiv:1606.06160*.
- [23] Ananda Samajdar, Yuhao Zhu, Paul Whatmough, Matthew Mattina, and Tushar Krishna. 2019. SCALE-Sim: Systolic CNN accelerator. arXiv:arXiv:1811.02883.
- [24] Naveen Muralimanohar, Rajeev Balasubramanian, and Norman P. Jouppi. 2009. *CACTI 6.0: A Tool to Model Large Caches*. Technical Report. HP Laboratories.
- [25] Kuan Wang, Zhijian Liu, Yujun Lin, Ji Lin, and Song Han. 2019. Retrieved from <https://github.com/mit-han-lab/haq>.
- [26] Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. 2017. Quantized neural networks: Training neural networks with low precision weights and activations. *J. Mach. Learn. Res.* 18, 1 (2017).
- [27] Fengfu Li, Bo Zhang, and Bin Liu. 2016. Ternary weight networks. arXiv:arXiv:1605.04711.
- [28] Matthieu Courbariaux, Itay Hubara, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. 2016. BinaryNet: Training deep neural networks with weights and activations constrained to +1 or -1. arXiv:arXiv:1602.02830.
- [29] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. 2016. XNOR-Net: ImageNet classification using binary convolutional neural networks. In *Computer Vision – ECCV*. Springer.
- [30] Song Han, Huizi Mao, and William J. Dally. 2016. Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. In *Proceedings of the 4th International Conference on Learning Representations (ICLR)*. <http://arxiv.org/abs/1510.00149>
- [31] S. Jain, S. Venkataramani, V. Srinivasan, J. Choi, P. Chuang, and L. Chang. 2018. Compensated-DNN: Energy efficient low-precision deep neural networks by compensating quantization errors. In *Proceedings of the 55th ACM/ESDA/IEEE Design Automation Conference (DAC)*. 1–6. <http://dx.doi.org/10.1109/DAC.2018.8465893>
- [32] Y. Liu, T. Zhang, and K. K. Parhi. 2010. Computation error analysis in digital signal processing systems with overscaled supply voltage. *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.* 18, 4 (2010), 517–526.
- [33] S. Venkataramani, A. Sabne, V. Kozhikkottu, K. Roy, and A. Raghunathan. 2012. SALSA: Systematic logic synthesis of approximate circuits. In *Proceedings of the Design Automation Conference 2012*. 796–801.
- [34] Z. Vasicek and L. Sekanina. 2015. Evolutionary approach to approximate digital circuits design. *IEEE Trans. Evolut. Comput.* 19, 3 (2015), 432–444.

- [35] R. Venkatesan, A. Agarwal, K. Roy, and A. Raghunathan. 2011. MACACO: Modeling and analysis of circuits for approximate computing. In *Proceedings of the IEEE/ACM International Conference on Computer-aided Design (ICCAD)*. 667–673.
- [36] Ajay K. Verma, Philip Brisk, and Paolo Ienne. 2008. Variable latency speculative addition: A new paradigm for arithmetic circuit design. In *Proceedings of the Conference on Design, Automation and Test in Europe*. 12501255.
- [37] Ning Zhu, Wang Ling Goh, WeiJa Zhang, Kiat Seng Yeo, and Zhi Hui Kong. 2010. Design of low-power high-speed truncation-error-tolerant adder and its application in digital signal processing. *IEEE Trans. Very Large Scale Integ. (VLSI) Syst.* 18, 8 (2010).
- [38] V. Gupta, D. Mohapatra, S. P. Park, A. Raghunathan, and K. Roy. 2011. IMPACT: IMPrecise adders for low-power approximate computing. In *Proceedings of the IEEE/ACM International Symposium on Low Power Electronics and Design*. 409–414.
- [39] I. Qiqieh, R. Shafik, G. Tarawneh, D. Sokolov, and A. Yakovlev. 2017. Energy-efficient approximate multiplier design using bit significance-driven logic compression. In *Proceedings of the Design, Automation Test in Europe Conference Exhibition (DATE)*. 7–12.
- [40] P. Kulkarni, P. Gupta, and M. Ercegovac. 2011. Trading accuracy for power with an underdesigned multiplier architecture. In *Proceedings of the 24th International Conference on VLSI Design*. 346–351.
- [41] C. Lin and I. Lin. 2013. High accuracy approximate multiplier with error correction. In *Proceedings of the IEEE 31st International Conference on Computer Design (ICCD)*. 33–38.
- [42] Patricio Bulić Uroš Lotrič. 2012. Applicability of approximate multipliers in hardware neural networks. *Neurocomputing* 96 (2012), 57–65.
- [43] Syed Shakib Sarwar, Swagath Venkataramani, Aayush Ankit, Anand Raghunathan, and Kaushik Roy. 2018. Energy-efficient neural computing with approximate multipliers. *Journal on Emerging Technologies in Computing Systems*. DOI: <https://doi.org/10.1145/3097264>
- [44] Jongsun Park, Hunsoo Choo, K. Muhammad, SeungHoon Choi, Yonghee Im, and Kaushik Roy. 2000. Non-adaptive and adaptive filter implementation based on sharing multiplication. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*.
- [45] Khaing Yin Kyaw, Wang Ling Goh, and Kiat Seng Yeo. 2010. Low-power high-speed multiplier for error-tolerant application. In *Proceedings of the IEEE International Conference of Electron Devices and Solid-State Circuits (EDSSC)*. 1–4.
- [46] M. J. Schulte and E. E. Swartzlander. 1993. Truncated multiplication with correction constant [for DSP]. In *Proceedings of the IEEE Workshop on VLSI Signal Processing*. 388–396.
- [47] S. S. Kidambi, F. El-Guibaly, and A. Antoniou. 1996. Area-efficient multipliers for digital signal processing applications. *IEEE Trans. Circ. Syst. II: Analog Dig. Sig. Process.* 43, 2 (1996), 90–95.
- [48] Jer Min Jou, Shiann Rong Kuang, and Ren Der Chen. 1999. Design of low-error fixed-width multipliers for DSP applications. *IEEE Trans. Circ. Syst. II: Analog Dig. Sig. Process.* 46, 6 (1999), 836–842.
- [49] S. Narayanamoorthy, H. A. Moghaddam, Z. Liu, T. Park, and N. S. Kim. 2015. Energy-efficient approximate multiplication for digital signal processing and classification applications. *IEEE Trans. Very Large Scale Integ. (VLSI) Syst.* 23, 6 (2015), 1180–1184.
- [50] S. Hashemi, R. I. Bahar, and S. Reda. 2015. DRUM: A dynamic range unbiased multiplier for approximate applications. In *Proceedings of the IEEE/ACM International Conference on Computer-aided Design (ICCAD)*. 418–425.
- [51] G. Zervakis, K. Tsoumanis, S. Xydis, D. Soudris, and K. Pekmetzi. 2016. Design-efficient approximate multiplication circuits through partial product perforation. *IEEE Trans. Very Large Scale Integ. (VLSI) Syst.* 24, 10 (2016), 3105–3117.
- [52] J. Y. F. Tong, D. Nagle, and R. A. Rutenbar. 2000. Reducing power by optimizing the necessary precision/range of floating-point arithmetic. *IEEE Trans. Very Large Scale Integ. (VLSI) Syst.* 8, 3 (2000), 273–286.

Received February 2021; revised October 2021; accepted October 2021