An Efficient Hardware Design for Accelerating Sparse CNNs With NAS-Based Models

Yun Liang^D, Senior Member, IEEE, Liqiang Lu^D, Yicheng Jin^D, Jiaming Xie, Ruirui Huang, Member, IEEE, Jiansong Zhang, Member, IEEE, and Wei Lin

Abstract—Deep convolutional neural networks (CNNs) have achieved remarkable performance at the cost of huge computation. As the CNN models become more complex and deeper, compressing CNNs to sparse by pruning the redundant connection in the networks has emerged as an attractive approach to reduce the amount of computation and memory requirement. On the other hand, FPGAs have been demonstrated to be an effective hardware platform to accelerate CNN inference. However, most existing FPGA accelerators focus on dense CNN models, which are inefficient when executing sparse models as most of the arithmetic operations involve addition and multiplication with zero operands. In this work, we propose an accelerator with software-hardware co-design for sparse CNNs on FPGAs. To efficiently deal with the irregular connections in the sparse convolutional layers, we propose a weight-oriented dataflow that exploits element-matrix multiplication as the key operation. Each weight is processed individually, which yields low decoding overhead. Then, we design an FPGA accelerator that features a tile look-up table (TLUT) and a channel multiplexer (CMUX). The TLUT is designed to match the index between sparse weights and input pixels. Using TLUT, the runtime decoding overhead is mitigated by using an efficient indexing operation. Moreover, we propose a weight layout to enable efficient on-chip memory access without conflicts. To cooperate with the weight layout, a CMUX is inserted to locate the address. Finally, we build a neural architecture search (NAS) engine that leverages the reconfigurability of FPGAs to generate an efficient CNN model and choose the optimal hardware design parameters. The experiments demonstrate that our accelerator can achieve 223.4-309.0 GOP/s for the modern CNNs on Xilinx ZCU102, which provides a 2.4×-12.9× speedup over previous dense CNN accelerators on FPGAs. Our FPGA-aware NAS approach shows 2x speedup over MobileNetV2 with 1.5% accuracy loss.

Manuscript received October 30, 2020; revised January 18, 2021; accepted February 23, 2021. Date of publication March 17, 2021; date of current version February 21, 2022. This work was supported in part by the Beijing Natural Science Foundation under Grant JQ19014; in part by the Beijing Academy of Artificial Intelligence (BAAI); and in part by the Key-Area Research and Development Program of Guangdong Province under Grant 2019B010155002. This article was recommended by Associate Editor W. Zhang. (Corresponding author: Yun Liang.)

Yun Liang is with the Center for Energy-efficient Computing and Applications, Peking University, Beijing 100871 China, and also with Peng Cheng Laboratory, Shenzhen 518066, China (e-mail: ericlyun@pku.edu.cn).

Liqiang Lu, Yicheng Jin, and Jiaming Xie are with the Center for Energyefficient Computing and Applications, Peking University, Beijing 100871 liqianglu@pku.edu.cn; China (e-mail: vicheng.jin@pku.edu.cn; jmxie@pku.edu.cn).

Ruirui Huang, Jiansong Zhang, and Wei Lin are with Alibaba Cloud intelligence, Alibaba Group, Hangzhou 330110, China (e-mail: ruirui.huang@ alibaba-inc.com; muduan.zjs@alibaba-inc.com; weilin.lw@alibaba-inc.com).

Digital Object Identifier 10.1109/TCAD.2021.3066563

Index Terms—Accelerator, convolutional neural network (CNN), FPGA, neural architecture search (NAS), sparse.

I. INTRODUCTION

NSPIRED by the biological nervous system, deep learning has recently achieved remarkable accuracy improvement. Convolutional neural networks (CNNs), the most commonly used model in deep learning, have been adopted in various domains, including image and speech recognition [1]–[4]. The significant accuracy improvement of CNNs comes at the cost of huge computational complexity as it requires a comprehensive assessment of all the regions across the feature maps. Toward such overwhelming computation pressure, FPGAs have emerged as a promising solution due to their high performance, energy efficiency, and programability [5]–[8].

In a typical CNN model, each neuron is regarded as a node in the network while the weight represents connections between nodes in two adjacent layers. Pruning the connections in the deep neural networks has been proved as an effective solution to compress the overall computation and memory requirements of these models while maintaining high accuracy. In general, compression techniques can be divided into two categories: 1) unstructured compression and 2) structured compression. The unstructured compression techniques prune the weights with irregularity in a fine-grain manner of pixels [9]–[11]. For example, Han *et al.* [9] have shown that there is significant redundancy (up to 90%) for certain DNNs, which can be pruned without sacrificing accuracy. The structured compression aims at pruning the networks with a certain shape in the weight [12]–[15]. However, the structured pruning often leads to a lower compression rate as it shows a strict mathematical formalization.

In this article, we mainly focused on accelerating CNNs with unstructured compression on FPGAs. Our approach can also be applied to structure compression. Though pruning techniques theoretically reduce the number of operations in the convolution algorithm and potentially provide the opportunity for faster inference process, existing accelerators on FPGAs for dense models are not suitable for sparse CNN models. Most of these works optimize their dataflows based on loop operations, such as loop interchange and loop unrolling [16]–[20]. The dense accelerator can result in high hardware inefficiency since most multiplication operations

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License.

For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

Weight: N×M×R×S Output feature maps: N×H×W



Fig. 1. Typical convolutional layer.



Fig. 2. Invalid computation caused by redundant connections in sparse CNNs. (a) Spatial dimension. (b) Input channel dimension. (c) Output channel dimension.

involve zero operands [5], [6], [16], [21]–[25]. The implementation of sparse DNNs has been studied in recent years on FPGAs [26]. These accelerators mainly focus on the fully connected (FC) layers, which use matrix-vector multiplication operations and are used for recurrent neural networks (RNNs) and LSTMs. However, the major operators of the modern CNN's computation are convolution operations. For example, the convolution operations occupy 90% of the total computation in GoogLeNet. Although the spatial convolution can be mapped to matrix–vector multiplications, this will increase the local memory requirement since the pixels in the input feature maps have to be copied multiple times when being flattened to a vector.

The challenges to design an efficient FPGA accelerator can be summarized as follows.

- The convolutional layers involve complex connections between input feature maps and output feature maps for sparse CNNs. Clearly, each output pixel is connected with part of the input pixels through the sliding kernels. The connection becomes irregular when the network becomes sparse. It is difficult to design a dataflow to deal with the irregularity but can leverage the high parallelism of FPGA and maintain FPGA efficiency.
- 2) The sparse weights are encoded in sparse format, which requires extra coordinate computation to locate the weights. However, the distribution of the sparse weights (nonzeros) is irregular, which leads to inefficient memory access and low on-chip bandwidth utilization.

3) A CNN model generally consists of different types of convolutional layers. Given a specific architecture design, the performance can be different when the layer parameter changes. Therefore, it is challenging to design a hardware-friendly CNN model that can maximize the performance.

To address the first challenge, we propose a weight-oriented dataflow where each PE performs element-matrix multiplication instead of spatial convolution. Here, the element refers to the sparse weight and the matrix refers to the input tile. In this dataflow, the sparse weights are processed separately. By doing this, we successfully avoid the design issues related to sparsity, such as irregular connections, load imbalance, etc. For the second challenge, we propose a weight layout, which can enable efficient on-chip memory access of the weights. In this layout, the weights processed in parallel are stored continuously, and the results are accumulated from different BRAM banks to avoid access conflicts. Moreover, we design an efficient accelerator for sparse CNNs that features a tile look-up table (TLUT) and a channel multiplexer (CMUX). TLUT can reduce the overhead of runtime index matching and CMUX helps to locate the output address easily when updating the results. Finally, we build a neural architecture search (NAS) engine based on analytical models that are used to predict the latency and resource utilization. For a specific deep learning task, we use the engine to explore the design space and identify the optimal CNN model architecture with hardware design parameters.

A preliminary version of this article was reported in [27]. In [27], we propose an architecture design for accelerating sparse CNNs on FPGAs. In this article, we extend previous work with software–hardware co-design to further improvement the performance. In particular, we propose a FPGA-aware NAS framework to search for the optimal hardware design parameters and network architectures simultaneously. We perform architecture search on ImageNet and draw comparisons with several state-of-the-art hand-crafted and autodesigned models.

In conclusion, this work makes the following contributions.

- 1) We propose a dataflow with element-matrix multiplication as the key operation, where the element and the matrix refer to the sparse weight and input tile, respectively.
- 2) We propose an architecture design for the dataflow with a set of optimization techniques. We use a look-up table to match the sparse weight with the corresponding input pixels. We also design the weight layout and compression format, which can enable efficient on-chip memory access.
- We develop an analytical model to estimate the latency. This model considers different types of operators in modern CNNs, e.g., pointwise convolution and depthwise convolution.
- 4) We develop a NAS engine to automatically generate CNN model that can match our hardware design. This engine searches both hardware design parameters and possible CNN models under resource constraints, and outputs the CNN model with high accuracy and low latency.

The experiments demonstrate that our accelerator can achieve 309.0, 223.4, 291.4 and 257.4 GOP/s for VGG, Alexnet, Resnet-152, and GoogLeNet on Xilinx ZCU102, respectively. Our accelerator achieves a $2.4 \times -12.9 \times$ speedup over the previous dense CNN FPGA accelerators. Compared to the TitanX GPU platform, our accelerator shows $7.56 \times$ energy efficiency. Our FPGA-aware NAS approach shows $2 \times$ speedup over MobileNetV2 with 1.5% accuracy loss.

II. BACKGROUND

A. Sparse CNN Model

CNNs are a class of deep, feedforward artificial neural networks, which are composed of a series of layers, including convolutional layers, pooling layers, and fully connected layers (FC layer). The convolutional layer is the most important layer in which the kernels extract features from the input feature map. Fig. 1 shows the typical convolution operation. The convolution operation uses a small $R \times S$ kernel to slide through the input feature map. The pixels inside the sliding window conduct a multiply-and-add operation with the weights in the kernel to compute a pixel value in the $H \times W$ output feature map. There are usually many input feature maps (also known as input channels) and output feature maps (output channels) in a single convolutional layer, and the numbers of input feature maps and output feature maps are M and N as shown in Fig. 1, respectively. Note that the convolution results in the different input channels are accumulated to obtain the output channel results.

CNNs usually have a large number of weights, which could introduce the problem of overfitting. The weights pruning techniques [9], [28] have been proven to be an effective method to reduce the computation and memory size while maintaining the overall model accuracy. For example, deep compression [9], [28] can reduce the number of weights in AlexNet [29] and VGG-16 [3] by $9 \times$ and $13 \times$, respectively. These are known as unstructured pruning techniques. There are other pruning techniques that prune the weights with structured patterns [15], [30], [31]. The advantage of structured pruning techniques is that they are hardware friendly. However, they often yield a low compression rate due to the strict mathematical formalization. The sparse CNN accelerators we propose can be used for both structured and unstructured pruning techniques.

B. Neural Architecture Search

NAS aims at automating NN architecture design, in analogy to deep learning automating feature engineering. Generally, a NAS program consists of search space, cost functions, and search algorithm. First, the search space generates a concrete NN architecture by combining different types of convolutional layers. Then, the NN architecture is evaluated by the cost functions, which consider the accuracy, network size, and the execution latency on a target platform. As for the algorithm, the earliest NAS algorithms train it from scratch on the whole data set using a controller RNN [32], [33]. However, these approaches mainly focused on the model accuracy without the consideration of the execution latency. Besides, they

TABLE I Analysis of Recent Sparse CNN Dataflow

dataflow	type	inner computation	Decoding overhead
SCNN[37]	pixel-oriented	Cartesian product	high
CambriconX[38]	kernel-oriented	vector dot product	medium
Cnvlutin[39]	pixel-oriented	vector multiplication & reduction	medium
SparTen[40]	pixel-oriented	vector dot product with inner join	high
Bit-Pragmatic[41]	bit-oriented	vector dot product with shifters	high
Ours	weight-oriented	element-matrix multiplication	low

are prohibitively computation intensive and are limited to small data sets and cell-level search spaces. Recently, there are NAS works introducing the hardware latency in the cost function [34]–[36]. However, these approaches only target a fixed hardware architecture, such as GPU platforms and mobile phones. The search algorithm does not take the hardware reconfigurability into account, which cannot be applied to FPGA platform.

Different from previous NAS approaches, we take full advantage of the flexibility of FPGA design where we incorporate hardware parameters into our NAS framework. More concretely, we build a resource model to estimate the FPGA resource utilization with architectural parameters and a latency model with both convolution parameters and architectural parameters. In this manner, FPGA architecture design can be taken into account in the search space. Section VI will provide the details of our hardware-aware NAS approach.

III. FPGA DATAFLOW DESIGN

A. Dataflows for Sparse CNNs on ASIC

There have been prior efforts on designing dataflows for sparse CNNs on ASIC platforms. However, these dataflows will be inefficient for FPGA platforms due to the distinct architectures. In Table I, we classify prior ASIC designs based on the inner computation of the dataflow. SCNN architecture [37] applies the pixel-oriented dataflow where the innermost computation is a Cartesian product. Using the Cartesian product, this dataflow multiplies input pixels with weights and returns multiple partial sums. This method requires significant coordinates computation to locate the sparse weights. Besides, the partial sums are connected with different output pixels, which bring great challenges for pipelining on FPGAs due to complex data dependency. Cambricon-X [38] design applies direct and step indexing technique to select the input pixels by detecting the nonzeros. Cambricon-X performs the vector dot product across channels by gathering the weights into a vector, which needs to dynamically select the input vector. This dataflow only performs parallel computation in channel dimensions, which also leads to poor parallelism on FPGAs. Cvnlutin design [39] leverages the sparsity in the input feature maps by using zero-skip computation. However, this dataflow requires runtime control to identify the nonzeros in the input pixels. Besides, the vector multiplication results are dynamically reduced to the output via an adder tree, resulting in high



Fig. 3. Weight-oriented dataflow. (a) Pseudo code of the dataflow. (b) Inner computation of the dataflow.

decoding overhead. SparTen [40] shares the similar inner computation to Cvnlutin. The difference is that SparTen applies an inner join scheme to gather the partial sums where the nonzeros are represented using a bit-mask. Since the input sparsity depends on the results of the previous layer, such bit-mask requires to dynamically encode the nonzeros into bit-mask representation. Bit-Pragmatic [41] focuses on the sparsity in bit level. Based-on the bit-serial unit of Stripes [42], Bit-Pragmatic performs parallelized bit-serial multiplications, and gathers the partial sums via a reduction tree. However, the zero bit of input pixels cannot be determined offline, which increases the logical overhead to detect nonzero bits in the input vector.

B. Dataflows for Dense CNNs on FPGAs

There have been dense CNNs dataflows on FPGAs [6], [16], [21], [43], [44]. However, these dataflows will lead to invalid multiplications caused by the redundant connections between weights and input/ouput channels for sparse CNNs. As shown in Fig. 2, the invalid multiplications can be from spatial kernel, input channel, and output channel dimensions, respectively. The input feature maps share the same index with the weight in the spatial kernel dimension and in the input channel dimension. In other words, the input pixel whose index matches the weight is needed when convolving the input with the kernel. Besides, different kernels are connected to different output feature maps, and the zero weight will not contribute to the corresponding output feature map.

C. Our Weight-Oriented Dataflow

We propose to transform the convolution computation to element-matrix multiplication by processing each weight as a single element. We compress the sparse weights into two arrays: 1) *SPw array*, where the nonzero weights in the same input channel are compressed into a vector and 2) *NZ array*, which records the number of nonzero weights in each input channel. One input channel is processed at a time. Fig. 3(a) shows the pseudocode of our dataflow, which consists of three steps. In step 1, we gather the necessary input pixels into an input tile according the position (h, w, m). A $T_H \times T_W$ tile in the output feature map is connected with $T_H \times T_W$ pixels in the input feature map through a specific weight. Given a specific kernel size and the sliding stride, a $T_H \times T_W$ tile corresponds to a $T_{H'} \times T_{W'}$ tile in the input feature map as follows:

$$T_{H'} = R + \text{stride} \times (T_H - 1), \quad T_{W'} = S + \text{stride} \times (T_W - 1)$$
(1)

where the kernel size is $R \times S$. Then, the input tile slides with a vertical stride T_H and a horizontal stride T_W as shown in Fig. 3(a). Step 2 is the inner computation of our dataflow where T_N weights are multiplied with the input tile in parallel.

Fig. 3(b) presents the details of the inner computation in the weight-oriented dataflow. Based on the position of the weight, we select a tile of input pixels that are connected with the weight. More clearly, given an output tile, each weight corresponds to a certain subinput tile determined by the position of the weight in the kernel. For example, the value "1" in the top-left corner of the sparse weight multiplies with all the 4×4 top-left tiles of input feature maps. The weights are from different output channels. Finally, the multiplication results will be accumulated the output pixels according to the index (n, h, w) in *Step 3*.

Our dataflow and its element-matrix multiplication inner computation have the following advantages. First, our dataflow processes the sparse weights one by one separately. By doing this, we can effectively exploit the sparsity and meanwhile, reduce the sparsity decoding overhead. Second, our dataflow provides sufficient parallelism on FPGAs. More clearly, the output pixels in the spatial kernel and output channel dimensions are computed in parallel. Third, our dataflow has low data dependency overhead. The results from step 2 in Fig. 3 are accumulated to different output pixels, which have no read-and-write conflicts.

IV. ARCHITECTURE OPTIMIZATION

In Section III, we transform the convolution operation to element–matrix multiplication. However, implementing this dataflow on FPGA arises two challenges. The first challenge



Fig. 4. Architecture overview.



Fig. 5. TLUT to locate the subinput tile.

is to select the necessary pixels for a specific weight. A single weight is connected to only part of the pixels in the input feature maps, and the weight in the different positions of the kernel is connected to different input pixels, as shown in step 1 of Fig. 3. Second, to ensure multiple results can be accumulated to the output buffer in parallel in step 2, a dedicated data layout is required under the hardware constraints of FPGA memory structure (e.g., dual-port BRAM). Furthermore, the PEs should be pipelined to increase the throughput.

A. Architecture Overview

As shown in Fig. 4, the input buffer contains four rows of feature maps. The output buffer size is set to store all pixels in one row of feature maps. Our FPGA accelerator consists of T_N PEs with each PE having $T_H \times T_W$ multipliers. Each PE is connected with a TLUT to match the weight and the required input pixels (Section IV-B). In Section IV-C, we propose a novel weight layout where the parallel weights are stored continuously. Besides, the layout can ensure the results from the PE array are accumulated to different output banks without data access conflicts in the pipeline. To cooperate with the layout, in Section IV-E, we propose a CMUX to locate the channel address. The CMUX receives the weight index in the sparse format and outputs which bank the results should be accumulated to. Since the weight distribution across output channels might be unbalanced, we analyze the load balancing problem in Section IV-F.

B. TLUT Module

As aforementioned, the weights represent the connections between the input feature map and the output feature map. However, when the weight is sparse, the connection loses its structured topology. To bridge the gap between irregular connections to input pixels and the regular PE array, we insert a TLUT between the input tile and PEs. Fig. 5 depicts how the weight and the input pixels are paired. When the kernel is sliding in the input tile, the weight in a $R \times S$ kernel is connected to a set of input pixels in the input tile. These pixels are batched together into a new tile. For example, in Fig. 2, the position of the weight with value "1" is (0, 0), which corresponds the top-left tile, and we can directly fetch the pixels from the TLUT module, which has been prefetched when the start point (h, w, m) is determined.

There are $R \times S$ subtiles in total with $R \times S$ positions in the kernel. These tiles are stored in local registers in the TLUT module. As the PE array processes multiple weights in parallel, each PE has its own TLUT module. As shown in Fig. 4, the input tile is reused by duplicating the pixel into multiple TLUT modules. The weight is reused by multiple pixels in the input tile. The TLUT module replaces runtime index matching with a simple array indexing operation by introducing additional local registers. This helps to save the logic resources significantly since the runtime index matching requires a large number of multiplexers.

C. PE Design and Weight Layout

The PE receives the decoded weight and the selected tile from TLUT. We initiate a PE array with each PE conducting an element–matrix multiplication operation. In step 2 of our dataflow, we compute multiple output pixels from different output channels in parallel. There are T_N homogeneous PEs process multiple weights and input tiles in parallel. Furthermore, we apply pipelining technique to our PE design. Pipelining allows multiple operations in step 2 to process concurrently to increase throughput, and the pipeling efficiency is determined by the iteration interval (II). According to Fig. 3, the iteration interval is bounded by the weight access bandwidth and output access bandwidth.

To enable simultaneous update of multiple output channels, the output buffer is partitioned to $T_N \times T_H \times T_W$ banks where each bank *i* in the channel dimension stores the weights from the $n = (T_N \times x + i)$ output channel as shown in Fig. 8. Traditionally, the weights are sorted in the ascending order of channels. If more than one weight need to be read from the same bank, this will lead to a long read latency. To address this problem, we rearrange the weight layout according to its remainder Re by dividing the output channel n with T_N (Re = $n \mod T_N$), as shown in Fig. 6, so that the results from the PE array are accumulated to the output buffers. For example, in Fig. 6, four weights are processed in parallel. In our weight layout, the results from the PE array need to be accumulated to the output channel (0, 5, 2, 7) in iteration i+1, whose remainder is (0, 1, 2, 3). In this manner, multiple write operations refer to different banks, resulting in an improved iteration interval.

D. Sparse Format

To cooperate with the weight layout, we propose an interleaved output channel coordinate list (IOCOO) format to



Fig. 6. Weight layout in the output channel dimension. *II* is the loop iteration interval of pipeline. (a) A straightforward data layout by sorting the weights according to the channel number in ascending order. (b) Proposed data layout by sorting the weights according to the channel number in cyclic order.

store the sparse weights. More clearly, weights in one input channel are stored in a vector. Each element is 5-tuple (n, r, s, value, valid), which represents the indices and the value of the weight. These tuples are stored with different bit width as follows. Using this format, the compressed weight can be directly sent to PEs without decoding overhead, leading to a high PE bandwidth utilization.

tuples	r	s	n	valid	value
bit width	4	4	10	1	16

Fig. 7 shows the compression rate of a few well-known formats, such as CSR, CSC, and COO, and recently proposed formats, including C²SR of MatRaptor [45] and CISS of Tensaurus [46]. For CSR, CSC, and C²SR formats, we flatten the weight into a $(1024 \times 3 \times 3) \times 1024$ matrix. The compression rate is given by

compression rate =
$$\frac{\text{the compressed data size}}{\text{dense data size}}$$

In Fig. 7, CSR has a lower compression rate because of the matrix shape. The COO format shows a higher memory requirement as indices are stored individually. The compression rate of our approach is similar to CSC because of the same number of pointers. In the CISS format, extra information is required to store pointers of higher dimensions as the weight is a 4-D-tensor. C^2SR shares a similar idea to our design where each row is assigned a fixed channel in a cyclic manner to avoid memory conflict. Though our format requires a little higher memory, less logic resource is needed for decoding. For example, using CSR format, the spatial coordinate of the weight needs to be calculated according to the row pointer.

We also analyze the available bandwidth of different formats, as shown in Fig. 7. The bandwidth is calculated with the assumption that the PE number is 8, and the frequency is 1 GHz. Traditional formats, such as CSR, CSC, and COO, only have a single entry for compressed weights resulting low bandwidth. The bandwidth of CISS and IOCOO is similar since both of these two formats partition the weight according to PE numbers. C²SR has higher bandwidth. This is because each PE in MatRaptor [45] is responsible for multiple weights



Fig. 7. Compression rate of different format. We assume the filter size is $1024 \times 1024 \times 3 \times 3$.



Fig. 8. CMUX to locate the output channel address.

in one iteration. To enable enough bandwidth, C^2SR partitions the weight matrix into more pieces leading to higher bandwidth.

E. CMUX Module

In the PE array, each PE generates a tile of results that belong to a distinct output channel. The address that the results need to be accumulated to is determined by the index in the format. A CMUX is inserted between the PE array and the output buffer to locate the address as shown in Fig. 8. The CMUX consists of T_N input wires, which represent the number of banks in the output channel dimension. The CMUX module will first compute the output address according to the remainder, e.g., the first input wire means its remainder is 1. Then, CMUX will output which bank the results need to be accumulated.

F. Load Balancing Analysis

In our architecture, PEs strictly process T_N weights with different remainders together. However, the weights with different remainders cannot be evenly distributed. So we align the weights with unnecessary data among all the remainders so that the number of weights across different remainders is equal. There is a *valid* signal in IOCOO format to indicate whether the weight is valid. As a result, the latency is always bounded by the remainder with the maximum nonzeros. The computation efficiency can be computed as follows:

$$Compute_{eff} = \frac{\# \text{ of valid}}{\# \text{ of valid} + \# \text{ of unnecessary}}.$$
 (2)

In the example of Fig. 9, the parallelism factor T_N is set to 4, 6, and 8 with a fixed number of nonzeros 28. For example, when $T_N = 8$, the computation efficiency is [28/(28 + 12)] = 70%. In the experimental section, we will analyze the computation efficiency using real networks.



Fig. 9. Unnecessary computation under proposed weight layout.

V. IMPLEMENTATION DETAILS

A. Memory System

The on-chip memory of FPGAs is not large enough to hold all the channels of feature maps. Besides, there exist data reuse opportunities both horizontally and vertically since there is overlapping when the kernel slides across the input feature maps. Line buffer design is widely used in previous accelerators and can effectively reuse the input data [6], [21], [24], [47]. Hence, we apply line buffer design to load and calculate feature maps. We implement line buffer design using loop tiling techniques where the required data in tiled loops are stored in the on-chip BRAM. Different tiling strategies can lead to different data reuse opportunities. In our design, we choose to tile the loop in the channel dimension with factor K_M and K_N , as shown in Fig. 3(a). Because when the kernel sliding across feature maps, the relationship between the data of different channels is irrelevant or independent. Assuming the sliding stride of convolution kernel is one, each input line buffer contains $K_M \times W$ elements, and each output line buffer contains $K_N \times W$ elements. To sustain sufficient on-chip bandwidth for PE computation, we partition each buffer according to the parallelization degrees. For example, each output line buffer is partitioned with factor T_W in width dimension and factor T_N in channel dimension.

The modules in the PE are also pipelined to increase throughput. Clearly, there are two input tiles working in a ping-pong manner to overlap the latency of TLUT and the latency of the PE array. The latency of generating TLUT can be regarded as a constant. The latency of the PE array depends on the loop count of L_K and the pipeline depth, as shown in Fig. 3. In general, the latency of loop L_K is much larger than the latency of TLUT; therefore, our PE design can achieve high efficiency.

B. Implementation of Other Layers

In general, the modern CNN networks contain different kernel sizes. For example, Resnet has 1×1 and 3×3 kernels in the residual block, and GoogLeNet has 1×1 , 3×3 , and 5×5 kernels in the inception module. Since each weight is processed independently in our dataflow, our architecture can flexibly handle different kernel sizes. To unify the structure of TLUT, we transform all the kernels to the 3×3 kernel. Fig. 10 shows an example that transforms the 5×5 kernel to the 3×3 kernel. The 5×5 kernel is padded to 6×6 kernel with zeros then split into four 3×3 kernels. Apart from the convolutional layers, there are other layers in CNN models. In our architecture, we implement two widely used layers: 1) pooling layer



Fig. 10. Uniform design different kernel sizes.

TABLE II Hardware Design Parameters

Parameters	Description					
Memory design parameters						
T_H	Tiling factor of output image height					
K_M Tiling factor of input channel						
K_N Tiling factor of output channel						
	Computation design parameters					
T_H	Parallelization factor of output image height					
T_W	Parallelization factor of output image width					
T_N	Parallelization factor of output channel					

and 2) rectified linear unit (ReLU) layer. The pooling layer outputs the maximum values in subregions of input feature maps. ReLU layer sets any input value less than zero to zero. These two layers are implemented by introducing comparison operators when writing the results to off-chip memory.

To accelerate an end-to-end CNN model, our design also supports ReLU layer, pooling layer, and FC layer. FC layers connect all the neurons in the previous layer to every single neuron in the weight matrix. We treat FC layer as a convolution layer with 1×1 kernel. Max pooling layers are widely used in CNNs, which output the maximum values in subregions of input feature maps. ReLU layers set any input value less than zero to zero. As shown in Fig. 4, pooling and ReLU logic are set before storing the output line buffer to the off-chip memory. ReLU is a pixelwise operation, which is implemented by introducing comparison operators for each pixel. Pooling layer is implemented by gating the nonmaximum value in the pooling region when storing the output pixels.

VI. FPGA-AWARE NAS

In previous sections, we provide an efficient architecture design for sparse CNN acceleration. Our hardware implementation involves several design parameters as shown in Table II. These parameters will affect both FPGA resource utilization and performance. More importantly, the execution latency of a CNN model is determined by both hardware parameters and convolution parameters, which makes it hard to find the optimal hardware design and an FPGA-friendly CNN model. Here, we propose a NAS framework that searches both architecture parameters and CNN model parameters. Specifically, we first formulate a resource model as the search constraint, and a latency model that is incorporated into the loss function. Then, we perform architecture search on a sparse supernet, which minimizes the total latency meanwhile maintains the accuracy.

A. Resource Model

We use the number of memory banks to estimate the BRAM usage, which is mainly used for input buffer, output buffer, and encoded weight buffer. In our design, the PE array generates multiple $T_H \times T_W$ output tiles in T_N different output channels. Considering double buffer design, the number of output buffer banks is $2T_H \times T_W \times T_N$. In the line buffer design, the input line buffers are rotated to reuse the overlapped area during the kernel sliding. According to (1), the number of input buffer banks is $(T_H + T_{H'}) \times T_{W'}$. As for the compressed weight in IOCOO format, five tuples are represented with different bitwidth. The value and output channel number are stored individually. While the rest tuples are packed together and stored as a 9-bit element. Each part is partitioned with factor T_N . Therefore, the bank number of sparse weight is $3 \times T_N$. In summary, the total number of banks can be written as follows:

$$Banks^{1} = 2T_{H} \times T_{W} \times T_{N} + (T_{H} + T_{H'}) \times T_{W'} + 3 \times T_{N}.$$
(3)

Most DSP resource is consumed to perform multiplications between the input element and the sparse weight. Here, we assume a single DSP can be implemented as one multiplier.² In the CMUX module, each output tile needs to locate the output channel number, which takes three DSPs to calculate the address (3 comes from Vivado HLS). In this manner, the DSP utilization can be estimated as follows:

$$DSPs = T_H \times T_W \times T_N + 3 \times T_N.$$
(4)

Modeling the LUT consumption on FPGA is more complex. For simplicity, we only model the LUT consumption for TLUT modules and CMUX modules. In our dataflow, there are T_N TLUTs working in parallel. The LUT consumption depends on the input tile size $T_H \times T_W$. Besides, there are T_N CMUX modules, which is a crossbar with T_N input wires. In summary, the LUT consumption is formulated as

$$LUTs = T_N \times (\alpha \log(T_H \times T_W) + \beta \log(T_N))$$
(5)

where $\alpha \log(T_H \times T_W)$ is the LUT consumption to store a single input tile, and $\beta \log(T_N)$ is the LUT consumption for one CMUX module. α and β can be obtained on different platforms in advance. We first get the LUT consumption for a set of (T_H, T_W) based on Vivado High level Synthesis Tool. Then, we get α and β based on linear regression.

B. Latency Model

The overall latency is either bounded by the PE computation time or the data transfer time in the line buffer design. First, we model the latency of loading/storing T_H rows of feature maps to/from the line buffer. Assuming each data are stored in 16 bits fixed point and the kernel sliding stride is 1, the data transfer time can be formulated as follows:

$$T_{\rm tra} = \frac{T_H \times W \times \max(K_M, K_N) \times 16 \text{bits}}{\text{Bandwidth}}$$
(6)

²In Xilinx FPGA, one BRAM can store 1024 words of 18bits. The required BRAMs is determined by the data size in one bank.

where K_M , K_N are tiling factors in the channel dimension in Table II.

To simplify the estimation of the computation time, we use eff_{ave} to represent the average computation efficiency as discussed in Section IV-F. We define the computation time as the time to generate $T_H \times K_N$ output elements

$$T_{\rm com} = \left(\frac{\rm sparsity}{\rm eff_{ave}} \times \left\lceil \frac{K_N}{T_N} \right\rceil \times K_M \times \left\lceil \frac{W_{\rm out}}{T_W} \right\rceil \times II + P_{\rm depth} \right) \\ \times \frac{1}{\rm Freq}$$
(7)

where *Freq* is the operating frequency of FPGAs. *II* denotes the iteration interval of the pipeline. In our implementation, the loop L_K in Fig. 3(a) is perfectly pipelined, so II = 1. P_{depth} is the pipeline depth, which can be ignored when the loop trip count is large enough.

Involving sparsity can result in low computation latency. Therefore, we also consider the initial time to load the first $T_{H'}$ rows of input feature map and the sparse weight

$$T_{\text{init}} = \frac{K_M \times K_N \times T_{H'} + 3 \times R \times S \times K_M \times K_N \times \text{sparsity}}{\text{Bandwidth}/16\text{bits}}.$$
(8)

Putting it all together, the total latency is determined by the initial time and maximum between data transfer time and computation time.

$$T_{\text{total}} = \left\lceil \frac{M}{K_M} \right\rceil \times \left\lceil \frac{N}{K_N} \right\rceil \times \left(\left\lceil \frac{H}{T_H} \right\rceil \times \max(T_{\text{tra}}, T_{\text{com}}) + T_{\text{init}} \right).$$
(9)

C. Search Algorithm

Our NAS algorithm models the FPGA architecture. As aforementioned, our architecture is parameterized with parallelization factor T_H , T_W , T_N and buffering factor K_M , K_N , which affect the resource utilization and execution latency. Therefore, our problem setting specifies each layer o with three groups of parameters, including model weights w, and architecture parameters p and hardware design parameters h. More specifically, when deriving child networks from the supernet, each layer is sampled from a multinomial distribution parameterized by p over a predefined candidate set { o^i }. The hardware design parameters are also taken into account in twofolds. First, the resource model ensures that the resource requirement of the search result can be met for a given FPGA device. Second, the latency model is integrated into the loss function for software–hardware co-design.

We use the same search space and supernet architecture as shown in Tables III and IV. We adopt the mobilenet inverted residual block as the basic building block of the supernet, whose kernel size can be chosen from $\{3, 5, 7\}$ and expansion ratio from $\{3, 6\}$. Furthermore, to permit flexibility in network depth, a special skip layer is added to the candidate set if the input and output of a layer are of the same size.

Fig. 11 shows our search algorithm, which comprises three stages, i.e., warm up, searching, and retraining. In the warm-up stage, we adopt ADMM pruning to obtain a sparse supernet [10]. The supernet is trained normally for several epochs

 TABLE III

 CONFIGURATIONS OF CANDIDATE OPERATIONS

Operation Type	Expansion Ratio	Kernel Size
3×3 _MBConv3	3	3
3×3_MBConv6	6	3
5×5 _MBConv3	3	5
5×5_MBConv6	6	5
7×7_MBConv3	3	7
7×7_MBConv6	6	7
skip layer	-	-

TABLE IVSUPERNET SPECIFICATION. HERE MIXOP DENOTES THE MIXEDOPERATION, WHICH CAN BE CHOSEN FROM SEVEN CANDIDATEOPERATIONS, 3×3 _CONV DENOTES A NORMAL 3×3 CONVOLUTION,
AND 3×3 _MBCONV1 REPRESENTS A 3×3 MOBILENET INVERTED
RESIDUAL BLOCK OF EXPANSION RATIO 1

Input Size	Operator	Output	s	n
224×224x3	3×3_Conv	32	2	1
112×112×32	3×3 _MBConv1	16	1	1
112×112×16	MixOp	24	2	4
56×56×24	MixOp	40	2	4
$28 \times 28 \times 40$	MixOp	80	2	4
$14 \times 14 \times 80$	MixOp	96	1	4
14×14×96	MixOp	192	2	4
7×7×192	MixOp	320	1	1
7×7×320	GAP	-	-	1
320	FC	1000	-	1



Fig. 11. FPGA-aware NAS flow.

before we introduce the ADMM regularizer to promote sparsity. After convergence, we zero out the least significant connections according to the magnitude. A binary mask is then applied over each pruned parameter to prevent backpropagation from tampering with the weights of removed connections. In the second stage, we perform NAS over the sparse supernet. Our loss function integrates the above-mentioned resource constraints and latency model, which is given by

$$\begin{array}{ll} \underset{\{w_i\},\{p_i\},\{h_i\}}{\text{minimize}} & \text{Loss}_{\text{CE}} + \lambda \mathbb{E} \big[\text{latency} \big] \\ \text{subject to} & \text{LUTs}(h_i) \leq \text{LUT}_{\text{max}} \\ & \text{DSPs}(h_i) \leq \text{DSP}_{\text{max}} \\ & h_i \in \mathcal{H}, \quad i = 1, \dots, N \end{array}$$
(10)

where λ is a scaling factor, Loss_{CE} is the accuracy loss, and LUT_{max} and DSP_{max} denote, respectively, the limitation of LUT and DSP resources. The weight decay term is omitted here for simplicity. The second term of the loss function stands



Fig. 12. Automatic hardware generation for ASIC and FPGA.

for the expected latency of the whole network, which can be calculated as follows:

$$\mathbb{E}\left[\text{latency}\right] = \sum_{i} \sum_{j} p_{i}^{j} \times T_{\text{total}}\left(\sigma_{i}^{j}, h_{i}\right).$$
(11)

Here, σ_i^j and p_i^j are the *j*th candidate in the *i*th layer and its assigned possibility. The estimated latency T_{total} is derived by assigning a specific candidate and hardware configuration as shown in (9).

To simplify the problem, we optimize hardware parameter h in a separate subproblem:

minimize
$$T_{\text{total}}(o, h)$$

subject to $\text{LUTs}(h) \leq \text{LUT}_{\text{max}}$
 $\text{DSPs}(h) \leq \text{DSP}_{\text{max}}$
 $h \in \mathcal{H}.$ (12)

Since the search space \mathcal{H} is a finite discrete in our case, a simple parameter sweep will find the globally optimal solution h^* . In our case, the latency overhead of exhaustive search is negligible compared to neural network training. Then, the original loss function can be rewritten as

$$\underset{\{w_i\},\{p_i\}}{\text{minimize Loss}_{CE}} + \lambda \sum_{i} \sum_{j} p_i^j \times T_{\text{total}}^* \left(o_i^j \right)$$
(13)

where $T_{\text{total}}^*(o_i^j)$ is a shorthand for $T_{\text{total}}(o_i^j, h_i^*)$.

In this manner, the simplified objective resembles a regular hardware-aware NAS problem that does not involve hardware parameters, and we can solve the remaining part according to the technique in [35]. After the search algorithm has sufficiently converged, we obtain the optimal subnetwork by only keeping the most promising candidate at each layer. Since we have pruned the supernet in the warm-up stage, the compact network retains its sparsity. Finally, the best architecture is retrained from scratch for the final evaluation.

Note that our framework is able to support more complex cases where \mathcal{H} can be a large discrete space or even a continuous one, and where a more sophisticated performance model is adopted. For example, given a learned performance model T_{total} parameterized by θ , we could backpropagate gradients through it and optimize hardware parameters h with gradient descent by fixing θ .

D. Automatic Hardware Generation

In this section, we integrated our architecture design with our NAS engine into an automatic flow for hardware generation. The intermediate passes are summarized in Fig. 12. The inputs consist of a specific task and budget constraints.

 TABLE V

 Performance Comparison With Previous Implementation

	[21]	[16]	[48]	[49]	[43]	Ours	Ours	Ours	Ours
CNN type	VGG	VGG	VGG	VGG	Alexnet	VGG	Alexnet	Resnet-152	GoogLeNet
Daviaa	Arria-10	Arria-10	Xilinx	Zuna 70706	Zynq	Zynq	Zynq	Zynq	Zynq
Device	GX1150	GX1150	Virtex-7	Zynq ZC700	ZC706	ZCU102	ZCU102	ZCU102	ZCU102
Frequency (MHz)	150	370	200	100	-	200	200	200	200
Precision	16bit fixed	32bit float	16bit fixed	-	-	16bit fixed	16bit fixed	16bit fixed	16bit fixed
DSP Utilization	1518 (100%)	1320 (87%)	3200 (89%)	-	-	1144 (45%)	1144 (45%)	1144 (45%)	1144 (45%)
Logic Utilization	161K (38%)	182K (43%)	237K (55%)	-	-	552K (92%)	552K (92%)	552K (92%)	552K (92%)
BRAM	1900 (70%)	1250 (46%)	1244 (85%)	-	-	912 (48%)	912 (48%)	912 (48%)	912 (48%)
Performance (GOP/s)	$64.5 (eff.)^1$	128.5 (eff.)	281.6 ²	297	71.2	309.0	223.4	291.4	257.4
DSP efficiency (GOP/s/DSP)	0.042	0.097	0.088	-	-	0.27	0.19	0.25	0.23
Logic efficiency (GOP/s/logic(K))	0.40	0.71	1.12	-	-	0.56	0.41	0.53	0.47
Power (W)	45.0	41.7	-	9.6	9.6	23.6	23.6	23.6	23.6

¹ eff. is the effective performance on sparse networks.

² 281.6 is estimated according to the number of frame per second in the paper.

TABLE VI TRAINING CONFIGURATIONS

Stage		Epoch	LR	Sched	bs
	pretrain	40	0.05	cosine	256
warmup	prune	60	0.01	step	256
	retrain	30	0.01	cosine	256
search		60	0.025	cosine	256
retrain		180	0.05	cosine	256

For ASIC design, the budget involves energy and area constraints under certain technology. For FPGAs, the budget is the on-chip resources. Then, our flow explores the space of the hardware–software co-design using the proposed NAS engine. The co-design information includes the target execution latency and the accuracy of a specific task. The co-design tuner helps to find the optimal architecture parameters and model parameters in the search space. Finally, our automatic flow generates a CNN model and hardware implementation written in high-level template. The template can be synthesized using Xilinx Vivado HLS to get the FPGA bitstream, or using Mentor Catapult HLS tool [50] to generate Verilog RTL.

VII. EXPERIMENT

In this section, we first introduce the experiments setting. In Section VII-B, we show the performance of our accelerator for the state-of-the-art CNNs and compare it with previous dense CNN FPGA accelerators. Then, we measure the resource utilization and analyze the utilization breakdown. In Section VII-D, we examine the hardware efficiency of different configurations using four state-of-the-art networks. Finally, we evaluate our NAS approach on FPGAs and compare to other NAS work.

A. Experimental Setup

We evaluate our design on Xilinx ZCU102 platform. ZCU102 consists of an UltraScale FPGA, quad ARM Cortex-A53 processors, 500-MB DDR3. Our FPGA implementation is operated at 200-MHz frequency on this platform. To measure the runtime power, we plugged in a power meter in the FPGA platform. In this work, we first use Xilinx Vivado HLS (v2017.4) tool chain to transform C code into RTL implementation. Then, we employ Xilinx SDSoC (v2017.4) to compile the source code into bitstream. We apply [9] and [28] methods to train the CNN model with sparsity using the Caffe framework [51]. Specifically, we set the expected sparsity of the network by setting the value that is less than a threshold to zero, followed by retraining the network to regain any lost accuracy. In our experiment, we use the state-of-the-art CNNs, including Alexnet, VGG-16, Resnet-152, and GoogLeNet. We achieve 10.8%, 11.7%, 23.5%, and 34.2% sparsity³ of Alexnet, VGG-16, Resnet-152, and GoogLeNet without accuracy loss, respectively.

We perform NAS on the full ImageNet. The validation set contains 50 000 images randomly sampled from the original training set. As for ADMM pruning, we set rho⁴ to 1e-3, 1e-2, and 1e-1, respectively in multirho training. Both the ADMM pruning and the NAS framework are implemented in PyTorch [52]. We use $\lambda = 0.1$ in our experiments, and summarize other training configurations in Table VI. Note that we spend more GPU hours (130 epochs versus 40) in the warm-up stage because of pruning. All training stages are run on 4 NVIDIA V100 GPUs.

B. Performance Analysis

In this section, we show the performance of our accelerator using modern CNNs. We set the accelerator configuration as $\langle T_H, T_W, T_N \rangle = \langle 8, 8, 16 \rangle$, which involves 1024 multipliers. In this configuration, the peak available performance can be computed as

peak performance = # of multipliers \times frequency \times 2.

Here, 2 means multiplication and addition operations. The peak performance of our design is 409.6 GOP/s.

We also compare our design with previous FPGA accelerators in Table V. References [16] and [21] are dense CNN accelerators, and [43], [48], and [49] are sparse CNN accelerators. The performance in Table V represents the effective performance. For the dense CNN accelerators, the effective performance is computed by multiplying the performance

³Sparsity is defined as the percentage of nonzeros.

⁴Rho is the scale factor of the penalty term in ADMM pruning.



Fig. 13. Computation efficiency. (a) Alexnet. (b) VGG-16. (c) GoogLeNet. (d) Resnet.

of dense CNNs with sparsity. According to Table V, our implementation achieves 223.4 GOP/s effective performance on sparse Alexnet, which shows $2.4 \times$ speedup compared with [43].⁵ You and Wu [49] showed similar performance to our design, but it applies low bit precision, which requires less resources. The performance on VGG network is 309.0 GOP/s, which is $3.6 \times -4.8 \times$ higher than [16] and [21]. Niu *et al.* [48] showed higher performance because they pruned the network in the frequency domain, which results in elementwise multiplication pattern. This computation pattern shows less complexity compared with the convolution operator. For Resnet-152 and GoogLeNet, our design achieves 291.4 and 257.4 GOP/s, respectively.

To make a fair comparison across different platforms, we also present the DSP-efficiency and logic-efficiency on each platform. On average, our design exhibits 0.24 GOP/s/DSP DSP-efficiency, which shows $2.5 \times -5.7 \times$ improvement compared with prior works [16], [21], [48]. On the other hand, our design shows lower logic efficiency. This mainly comes from the TLUT module and CMUX module for the data decoding. Most previous designs target dense CNNs that exhibit regular data access. Therefore, they have higher logic efficiency.

The speedup of our design is because the weight-oriented dataflow can effectively eliminate the useless multiplications. In addition, the dataflow maintains a high utilization of onchip resources. Previous implementations cannot efficiently exploit the zeros in the computation, which results in a waste of on-chip resources. On the other hand, previous dense CNN accelerators are highly optimized and DSPs are fully utilized to conduct multiplications. In our implementation, only half of DSPs is utilized, and the performance is bounded by the logic resource.

The inefficiency of our implementation mainly comes from three aspects. First, there exist some invalid weights in our weight layout, which leads to imbalanced workload among PEs. Section VII-D presents the details of the load imbalance problem. Second, the feature map size in the CNN layers cannot divide T_H and T_W evenly. We choose 8×8 as the output tile size. Taking the last convolutional layer of VGG as an example, the feature map size is 14×14 leading to a 12.5%waste of computation. Third, as mentioned in Section IV-C, we apply pipeline technique in PEs. When the workload is small

TABLE VII RESOURCE UTILIZATION AND LATENCY

	BRAM	DSP ⁶	FF	LUT		
Buffers	609	0	0	0		
TLUT	0	0	48894	59150		
PEs	0	288	1088	32		
CMUX	0	24	394	43976		
Others	8	52	18408	12758		
Total	617	364	68784	132344		
Predict(error)	556(9.9%)	312(14.3%)	-	103158(9.6%)		
Available	1824	2520	548160	274780		
Actual latency	191011 cycles					
Predicted latency	182523 cycles (4.4% error)					

after pruning, the latency of PE can be bounded by the depth of pipeline. In our implementation, the pipeline depth is eight cycles. Compared with the VGG network, Resnet-152s and GoogLeNet consist of many convolutional layers with 1×1 kernels, leading to low performance. The speedup of VGG-16 is higher than that of Alexnet, because VGG-16 is a structured and regular network. The kernel size of all layers is 3×3 ; however, Alexnet contains many layers in different types.

C. Resource Utilization Characteristics

Table VII shows the resource utilization breakdown with the configuration ($T_H = T_W = 6$, $T_N = 8$). In Table VII, we also present the prediction accuracy of our resource model and latency model. The cycle number of latency is tested on the convolutional layer with 112×112 feature map size and $64 \times 64 \times 3 \times 3$ kernel size. The inaccuracy sources of DSP utilization mainly come from the padding operation and FIFOs in the line buffer design. A few extra BRAMs are used for FIFOs and pipeline buffers. The actual latency is obtained using Xilinx HLS simulation tools. The prediction error of latency model results from the padding operation, which only occupied a few cycles.

Fig. 14 shows the resource utilization of different configurations obtained from Xilinx Vivado tool (v2017.4). In Fig. 14, the LUT utilization increases as the parallelism factor T_N increases because of CMUX. When T_N is large, the utilization of BRAMs is mainly determined by the parallelization degree of feature maps (T_H , T_W). When T_N is small, BRAM utilization is similar. This indicates that the consumption of BRAM is determined by the input and output data size, instead of the partition factors. For example, when $T_N = 4$ in (3), each bank

⁵This article [43] only reported the performance and the platform.

Fig. 14. Resource utilization (HaWbNc means $T_H = a, T_W = b, T_N = c$).

H6W6N4

-16W6N8

16W6N16

H8W8N4

H8W8N8

I8W8N16

DSP

LUT

is constructed using two BRAMs. While each bank only needs one BRAM when $T_N = 8$.

D. Computation Efficiency

In our implementation, the PE array processes P_N weights with different remainders at the same time. However, the remainder distribution is irregular, which can result in load imbalance problem as discussed in Section IV-F. Fig. 13 shows the efficiency across different layers with different parallelism factors. We find that the format efficiency increases when the modular factor T_N becomes larger. Because a large T_N will bridge the gap between the maximum number of valid values and the average number of valid values among different remainders. We find that the efficiency increases as the network goes deeper. This because the number of channels increases as the network goes deeper, which makes the total number of nonzeros larger. A large number of nonzeros can compensate for the gap between the maximal and minimal number of remainders. Also, we observe that the computation efficiency of GoogLeNet and Resnet-152 is much higher. This is because the sparsity of these two networks is relatively large, which leads to a large number of nonzeros. Besides, the computation efficiency of some layers in GoogLeNet is low. Because the channel number cannot be divided evenly by T_N . For example, the output channel number of inception_4b layer in GoogLeNet is 24, which is not a multiple of $T_N = 16$. In conclusion, our dataflow can maintain high computation efficiency for different configurations and networks.

E. Comparison With ASIC Dataflows

In this section, we build several cycle-accurate models for the comparison with ASIC sparse tensor accelerators. Cambricon-X [38] and SCNN [37] are CNN-specific ASIC accelerators. We only model the behavior of the PE array of these designs. To make a fair comparison, we scale the number of multipliers to 1024. Tensaurus [46] and ExTensor [53] are general sparse tensor accelerators. MatRaptor [45] focuses on the sparse matrix–matrix multiplication (GEMM). When comparing with them, we transform the convolutional layer to GEMM using the image-to-column (im2col) operation. The input image is regarded as a dense matrix.



Fig. 15. Comparison with other ASIC dataflows.

Fig. 15 gives the comparison results. We also draw the line of theoretical speed up calculated by (1/sparsity). When the sparsity is higher than 70%, Cambricon-X and our design show near-ideal speedup. However, SCNN shows a lower speedup because of a large amount of fanout of the PE array when most weights are nonzero. On the other hand, our design outperforms Cambricon-X when the sparsity decreases. This is because Cambricon-X requires index comparison to conduct sparse vector dot product, which is inefficient for low sparsity. Our weight-oriented dataflow minimizes the indexing overhead by introducing a static TLUT module.

The speedup of general sparse tensor accelerators is always lower than the CNN-specific accelerators. Such limitation mainly comes from the im2col operation where the input images are duplicated multiple times and flattened into a matrix. On the other hand, CNN accelerators naturally leverage the convolution properties to avoid data rearrange overhead. To be specific, Tensaurus unifies the sparse computation as an operation between a scalar and a fiber, where each PE performs the multiplication between one weight and one input vector. MatRaptor applies the rowwise matrix multiplication dataflow where one weight is multiplied with the corresponding row of the input matrix. Though these two dataflows share a similar idea to our design that each PE is responsible for the multiplication between one weight and multiple inputs, their speedup is limited by the sparsity. PEs in Tensaurus and MatRaptor are parallelized with inputs from different rows, which require a synchronization operation between different PEs when accumulating partial sums. Such synchronization overhead is small when the sparsity is extremely low. However, in sparse CNN models, the sparsity is usually around 0.1-0.2, which can lead to high synchronization overhead. ExTensor is parallelized using multiple dot product, which requires to compare the index between two vectors. When the sparsity is high, most indices can be matched. Therefore, ExTensor outperforms Tensaurus and MatRaptor when the sparsity is higher than 0.4.

F. Scalability and Comparison With GPU

We also test our design on the ZC706 platform to demonstrate the scalability. Our implementation is operated at 166-MHz frequency on this platform. ZC706 has 900 DSPs, 1090 BRAMs, and 305K logic cells. We set the configuration parameter as ($T_H = T_W = 8$, $T_N = 8$) and achieve 134.2 GOP/s on Resnet-152, which means our design can

100

80

60 40 20

0

H4W4N4

H4W4N8

H4W4N16

BRAM

Resource Utilization (%)

Device	TitanX ¹	TitanX ²	ZC706	ZCU102
Technology	28 nm	28 nm	28 nm	16 nm
Frequency (MHz)	1075	1075	166	200
Precision	32bits float	32bits float	16bits fixed	16bits fixed
conv average (GOP/s)	212(eff.)	119	134	291
Power (W)	130	134	9.4	23.6
Energy efficiency (GOP/s/W)	1.63	0.88	12.66	12.33

TABLE VIII COMPARISON WITH GPU PLATFORM USING RESNET-152

¹ The sparse network is considered as the dense network and accelerated using CuDNN.

² The sparse network is accelerated using CuSparse.

TABLE IX LAYER TYPES FOR LATENCY PROFILING. S-CONV: SPATIAL CONVOLUTION. DW-CONV: DEPTHWISE CONVOLUTION

Layer ID	Filter	Туре	Channel	Resolution
C1	1×1	S-Conv	256×256	112×112
C2	1×1	S-Conv	256×256	56×56
C3	5×5	DW-Conv	32×32	14×14
C4	7×7	DW-Conv	3×64	14×14
C5	3×3	S-Conv	32×32	112×112
C6	3×3	S-Conv	64×64	112×112
C7	3×3	DW-Conv	128×128	112×112
C8	3×3	S-Conv	256×256	14×14

be scaled to different platforms. Besides, we conduct a comparison between GPU and FPGA platforms, as shown in Table VIII. We measure the performance of dense Resnet-152 using the latest CuDNN on NVIDIA TitanX platform. To make a fair comparison, we also apply CuSparse library to accelerate sparse Resnet-152. The sparse version shows a lower performance because of the memory uncoalesing problem. In conclusion, our design shows $1.37 \times$ speedup and $7.56 \times$ energy efficiency compared with the TitanX platform.

G. Latency Profiling

We first profile the latency of different layer types, which help to find the FPGA-friendly neural network architecture. In Table IX, we list eight representative layer types, including spatial convolution and depthwise convolution. Fig. 16 shows the detailed profiling results involving initial time, compute time, and data transfer time. According to Fig. 16, the compute time is the bottleneck for most layers (C4-C8). However, the transfer time of the layer that shows less data reuse is higher than the compute time (C1 and C2). These layers are often spatial convolution with 1×1 filter and depthwise convolution. Besides, we also observe that the initial time can also affect the performance when the feature map size is small (C3 and C4). For example, for the C3 layer, the feature map only contains three input channels, which is less computation intensive. This phenomenon can also be found in C4 layer whose channel number is 32.

H. NAS Results

As shown in Table X, compared to other compact models that are with similar accuracy, our method consistently improves upon inference latency. Specifically, when confronted with MobileNetV2-1.0 [54], a manually designed



Fig. 16. Latency profiling with different layer parameters.



Fig. 17. Accuracy versus latency of sampled subnetworks. (a) Subnetworks sampled from the pruned supernet. (b) Subnetworks sampled from the unpruned supernet.

architecture that targets no specific platform, our model achieves $1.49 \times$ speedup. Note that ProxylessNAS and FBNet are NAS networks, but target a fixed hardware, e.g., CPU, GPU, and mobile phones, which did not search architecture parameters. While maintaining accuracy on par with FBNet-B [34] and ProxylessNAS-mobile [35], our architecture is $2.36 \times$ and $3.65 \times$ faster, respectively. Nevertheless, our method takes longer time than [35] in the warmup stage, because we conduct ADMM pruning on the supernet additionally. Yet, since we can perform architecture search multiple times on the same supernet, the pruning cost will only occur once.

In our experiments, we choose a conservative sparsity of 30% to avoid significant accuracy loss. Mobilenet-like compact networks on ImageNet have been observed to be less tolerant to network pruning, which can be attributed to less redundancy in parameters and a more challenging task. Fig. 17 shows the accuracy versus latency of 300 subnetworks randomly sampled from the dense and sparse supernet, respectively, with the highlighted point denoting the resultant architecture. This graph indicates: 1) the ranking of the child networks still correlates with their actual performance since the distribution characteristics are basically retained after pruning and 2) the search algorithm succeeds in finding out an architecture near the Pareto-front with acceptable latency. We also provide the resultant network architecture and hardware parameters for ZCU102 in Table XI.

To evaluate the scalability of our NAS method, we conduct the search algorithm on ZC706 FPGA and ZCU102 FPGA, which show different hardware resources. Table XII reports the results on two devices with different resource constraints and hardware search space. Based on the size of the available

 TABLE X

 Performance Comparison With Other Handcrafted Models (MobileNetV2 and ShuffleNetV2) and NAS Techniques (ProxylessNAS and FBNet). MobileNetV2-1.0 and ShuffleNetV2-1.0 Are the Original Network

Model	Search method	HW	Search space	#Params (M)	#OPs (M)	FPGA latency (ms)	Top-1 (%)
MobileNetV2-0.75	-	-	-	2.7	290	34.471	67.9
MobileNetV2-1.0	-	-	-	3.4	600	40.791	72.0
MobileNetV2-1.4	-	-	-	6.9	770	54.486	74.7
ProxylessNAS-CPU	gradient	X	cell-level	4.4	954	61.889	75.3
ProxylessNAS-GPU	gradient	X	cell-level	7.1	930	109.421	75.1
ProxylessNAS-mobile	gradient	X	cell-level	4.1	640	99.770	74.6
ShuffleNetV2-0.5	-	-	-	1.4	82	10.867	60.3
ShuffleNetV2-1.0	-	-	-	2.3	292	16.786	69.1
ShuffleNetV2-1.5	-	-	-	3.5	598	22.019	72.6
ShuffleNetV2-2.0	-	-	-	7.4	1182	28.019	74.9
FBNet-A	gradient	X	cell-level	4.3	498	56.328	73.0
FBNet-B	gradient	X	cell-level	4.2	590	64.532	74.1
FBNet-C	gradient	X	cell-level	5.0	750	87.766	74.9
Ours	gradient	1	cell-level	3.9	642	27.3	73.3

TABLE XI Searched Network Architecture

Input	Operator	Ksize	Output	Expansion		
$224^2 \times 3$	Conv	3×3	32	-		
$112^2 \times 32$	MBConv	3×3	16	1		
$112^{2} \times 16$	MBConv	5×5	24	3		
$56^2 \times 24$	MBConv	3×3	24	6		
$56^2 \times 24$	MBConv	7×7	40	3		
$28^2 \times 40$	MBConv	3×3	40	3		
$28^2 \times 40$	MBConv	5×5	40	3		
$28^2 \times 40$	MBConv	3×3	80	6		
$14^2 \times 80$	MBConv	3×3	80	6		
$14^2 \times 80$	MBConv	7×7	80	3		
$14^2 \times 80$	MBConv	5×5	80	3		
$14^2 \times 80$	MBConv	3×3	96	6		
$14^2 \times 96$	MBConv	3×3	96	3		
$14^2 \times 96$	MBConv	7×7	96	6		
$14^2 \times 96$	MBConv	3×3	192	3		
$7^2 \times 192$	MBConv	3×3	192	6		
$7^2 \times 192$	MBConv	3×3	192	3		
$7^2 \times 192$	MBConv	5×5	192	3		
$7^2 \times 192$	MBConv	3×3	320	6		
$7^2 \times 320$	Conv	1×1	1280	-		
1280	FC	-	1000	-		
Searched hardware parameters: $T_W = T_H = 3, T_N = 22$						

resource, we set the hardware parameter search range, which differs in the output channel number that computed in parallel. The search algorithm returns models with almost the same top-1 accuracy under these two settings. As we can see, our method can adapt to various hardware platforms and reflect hardware constraints in the search process.

VIII. RELATED WORK

Architecture for Dense CNNs on FPGAs: Prior efforts to accelerate CNNs have shown substantial successes on FPGAs. Ma et al. [21] made an in-depth analysis of loop optimization techniques in spatial convolution, which includes loop tiling, loop unrolling, and loop interchange. Zhang and Li [16] focused on reducing the on-chip memory bandwidth requirement. Wei et al. [23] implemented CNN on an FPGA using a systolic array architecture, which can achieve high clock frequency under high resource utilization. Zhang et al. [55] proposed AccDNN tool, which included high-quality RTL network layer IPs, a fine-grained layer-based pipeline architecture, and an automatic design space exploration tool. Besides, there are some works that implement fast algorithms to further accelerate CNNs [47], [56]–[59].

Architecture for Sparse CNNs on ASICs: Recently, some works explore the dataflow and architecture to accelerate sparse CNNs on ASICs. Han et al. [60] proposed EIE CNN accelerator, which exploits sparsity both in input feature maps and filters but only focused on the FC layer. The FC layer is computed using matrix multiplication, in EIE design, the matrix is stored in CSC format and multiple columns are computed in parallel. Parashar et al. [37]proposed SCNN accelerator with a dataflow named PT-IS-CP (planartiled input-stationary Cartesian-product). Zhang et al. [38] presented Cambricon-X accelerator which applies step indexing techniques. In Cambricon-X design, the nonzeros in the same row are divided into multiple segments with the same size in subsequent addresses. The row that contains nonzeros less than the size will be aligned to the size. In recent years, some ASIC accelerators apply hardware-software design that prunes the weight with structured pattern [15], [31].

NAS: Early NAS algorithms [33] are inefficient in terms of search time and hardware-friendliness during inference. There are two trends related to our work in the subsequent works that tackle these issues. One-shot NAS [35] constructs a supernet and defines candidate architectures as its subgraphs. Rather than training from scratch each time, the weights of a sampled architectures are generated by the pretrained supernet. DARTS [61] relaxes the discrete search space into a concrete distribution by assigning a real-valued weight to each candidate path. Instead of optimizing all paths jointly, ProxylessNAS [35] samples a few paths in each training step to reduce GPU memory consumption. Device-aware multiobjective NAS [34], [35] explicitly incorporates resource efficiency into the objective function, either device related (such as latency and energy consumption) or device agnostic (such as FLOPs and model size). These works often adopt a compact search space that is inspired by hand-crafted networks. Depending on the nature of the target hardware, the efficiency is either measured through runtime measuring [36],

TABLE XII NAS ON DIFFERENT DEVICE

Device	Resource constraints				Hardware parameter search range			Latency	Top 1 acc $(\%)$
	DSP	Logic	BRAM	Frequency	P_N range	P_H range	P_W range	Latency	$\left 10p-1 \operatorname{acc} (n) \right $
ZC706	900	350K	1090×18 Kb	166M	[4,16]	[3,8]	[3,8]	33.7	73.0
ZCU102	2520	600K	1824×18 Kb	200M	[4,32]	[3,8]	[3,8]	27.3	73.3

a premeasured lookup table [34], [35]. In parallel to our work, there is a recent trend of incorporating NAS into SW-HW codesign frameworks. This line of works fuses NN architecture parameters and hardware implementation parameters into a single search space, thereby optimizing them simultaneously via stochastic coordinate descent [62] or gradient-based methods [63]. Our method is to some extent similar to the one mentioned in [63], in that both adopt a mobilenet-like search space and a gradient-based approach to NAS. Nevertheless, to the best of our knowledge, this is the first work in this area targeting sparsity in both NN and accelerator design.

IX. CONCLUSION

In this work, we proposed an FPGA accelerator for sparse CNNs. We first propose a weight-oriented dataflow that exploits element-matrix multiplication. Based on this dataflow, we design an FPGA architecture mainly composed of a TLUT and a CMUX. Besides, we propose a weight layout where the weights calculated in parallel are stored continuously. To cooperate with the weight layout, CMUX is inserted to locate the address which can ensure no data access conflict. Finally, we develop an FPGA-aware NAS approach to find the hardware-friendly network structure. The experiments demonstrate that our accelerator can achieve 223.4-309.0 GOP/s for the modern CNNs on Xilinx ZCU102, which provides a $2.4 \times -12.9 \times$ speedup over previous dense CNN FPGA accelerators. Our FPGA-aware NAS approach shows $2 \times$ speedup over MobileNetV2 with 1.5% accuracy loss.

REFERENCES

- K. He, X. Zhang, S. Ren, and J. Sun, "Delving deep into rectifiers: Surpassing human-level performance on imagenet classification," in *Proc. ICCV*, 2015, pp. 1026–1034.
- [2] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *Proc. CVPR*, 2014, pp. 580–587.
- [3] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," 2014. [Online]. Available: arXiv:1409.1556.
- [4] J. Redmon, S. K. Divvala, R. B. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proc. CVPR*, 2016, pp. 779–788.
- [5] N. Suda et al., "Throughput-optimized OpenCL-based FPGA accelerator for large-scale convolutional neural networks," in Proc. FPGA, 2016, pp. 16–25.
- [6] J. Qiu et al., "Going deeper with embedded FPGA platform for convolutional neural network," in Proc. FPGA, 2016, pp. 26–35.
- [7] X. Zhang, J. Zou, X. Ming, K. He, and J. Sun, "Efficient and accurate approximations of nonlinear convolutional networks," 2014. [Online]. Available: http://arxiv.org/abs/1411.4229.
- [8] Y.-H. Chen, J. Emer, and V. Sze, "Eyeriss: A spatial architecture for energy-efficient dataflow for convolutional neural networks," ACM SIGARCH Comput. Archit. News, vol. 44, no. 3, pp. 367–379, 2016.
- [9] S. Han, J. Pool, J. Tran, and W. J. Dally, "Learning both weights and connections for efficient neural network," in *Proc. NIPS*, 2015, pp. 1135–1143.

- [10] T. Zhang *et al.*, "A systematic DNN weight pruning framework using alternating direction method of multipliers," 2018. [Online]. Available: http://arxiv.org/abs/1804.03294
- [11] J. Yu, A. Lukefahr, D. J. Palframan, G. S. Dasika, R. Das, and S. A. Mahlke, "Scalpel: Customizing DNN pruning to the underlying hardware parallelism," ACM SIGARCH Comput. Archit. News, vol. 45, no. 2, pp. 548–560, 2017.
- [12] J. Luo, J. Wu, and W. Lin, "ThiNet: A filter level pruning method for deep neural network compression," 2017. [Online]. Available: arXiv:1707.06342.
- [13] S. Lin, R. Ji, Y. Li, Y. Wu, F. Huang, and B. Zhang, "Accelerating convolutional networks via global & dynamic filter pruning," in *Proc. IJCAI*, 2018, pp. 2425–2432.
- [14] W. Wen, C. Wu, Y. Wang, Y. Chen, and H. Li, "Learning structured sparsity in deep neural networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 2074–2082.
- [15] C. Ding *et al.*, "CirCNN: Accelerating and compressing deep neural networks using block-circulant weight matrices," in *Proc. MICRO*, 2017, pp. 395–408.
- [16] J. Zhang and J. Li, "Improving the performance of OpenCL-based FPGA accelerator for convolutional neural network," in *Proc. FPGA*, 2017, pp. 25–34.
- [17] Q. Xiao, L. Lu, J. Xie, and Y. Liang, "FCNNLib: An efficient and flexible convolution algorithm library on FPGAs," in *Proc. 57th ACM/IEEE Design Autom. Conf. (DAC)*. 2020, pp. 1–6.
- [18] X. Wei, Y. Liang, X. Li, C. H. Yu, P. Zhang, and J. Cong, "TGPA: Tilegrained pipeline architecture for low latency CNN inference," in *Proc. Int. Conf. Comput.-Aided Design (ICCAD)*, 2018, pp. 1–8.
- [19] L. Jia, L. Lu, X. Wei, and Y. Liang, "Generating systolic array accelerators with reusable blocks," *IEEE Micro*, vol. 40, no. 4, pp. 85–92, Jul./Aug. 2020.
- [20] Q. Xiao and Y. Liang, "Zac: Towards automatic optimization and deployment of quantized deep neural networks on embedded devices," in *Proc. IEEE/ACM Int. Conf. Comput.-Aided Design (ICCAD)*. 2019, pp. 1–6.
- [21] Y. Ma, Y. Cao, S. B. K. Vrudhula, and J.-S. Seo, "Optimizing loop operation and dataflow in FPGA acceleration of deep convolutional neural networks," in *Proc. FPGA*, 2017, pp. 45–54.
- [22] Y. Guan *et al.*, "FP-DNN: An automated framework for mapping deep neural networks onto FPGAs with RTL-HLS hybrid templates," in *Proc. FCCM*, 2017, pp. 152–159.
- [23] X. Wei *et al.*, "Automated systolic array architecture synthesis for high throughput CNN inference on FPGAs," in *Proc. DAC*, 2017, pp. 1–6.
- [24] Q. Xiao, Y. Liang, L. Lu, S. Yan, and Y.-W. Tai, "Exploring heterogeneous algorithms for accelerating deep convolutional neural networks on FPGAs," in *Proc. DAC*, 2017, pp. 1–6.
- [25] Y. Liang, L. Lu, and J. Xie, "OMNI: A framework for integrating hardware and software optimizations for sparse CNNs," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, early access, Sep. 14, 2020, doi: 10.1109/TCAD.2020.3023903.
- [26] S. Han *et al.*, "ESE: Efficient speech recognition engine with sparse LSTM on FPGA," in *Proc. FPGA*, 2017, pp. 75–84.
- [27] L. Lu, J. Xie, R. Huang, J. Zhang, W. Lin, and Y. Liang, "An efficient hardware accelerator for sparse convolutional neural networks on FPGAs," in *Proc. IEEE 27th Annu. Int. Symp. Field-Programmable Custom Comput. Mach. (FCCM)*. 2019, pp. 17–25.
- [28] S. Han, H. Mao, and W. J. Dally, "Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding," 2015. [Online]. Available: arXiv:1510.00149
- [29] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Proc. NIPS*, 2012, pp. 1106–1114.
- [30] C. Deng, S. Liao, Y. Xie, K. K. Parhi, X. Qian, and B. Yuan, "PERMDNN: Efficient compressed DNN architecture with permuted diagonal matrices," in *Proc. MICRO*, 2018, pp. 89–202.
- [31] Z. Xuda et al., "Cambricon-S: Addressing irregularity in sparse neural networks through a cooperative software-hardware approach," in *Proc. MICRO*, 2018, pp. 15–28.

- [32] B. Zoph and Q. V. Le, "Neural architecture search with reinforcement learning," 2016. [Online]. Available: arXiv:1611.01578.
- [33] B. Zoph, V. Vasudevan, J. Shlens, and Q. V. Le, "Learning transferable architectures for scalable image recognition," *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 8697–8710.
- [34] B. Wu *et al.*, "FBNet: Hardware-aware efficient ConvNet design via differentiable neural architecture search," in *Proc. CVPR*, 2018, pp. 10734–10742.
- [35] H. Cai, L. Zhu, and S. Han, "ProxylessNAS: Direct neural architecture search on target task and hardware," 2018. [Online]. Available: arXiv:1812.00332.
- [36] M. Tan *et al.*, "MnasNet: Platform-aware neural architecture search for mobile," in *Proc. CVPR*, 2018, pp. 2820–2828.
- [37] A. Parashar et al., "SCNN: An accelerator for compressed-sparse convolutional neural networks," in Proc. ISCA, 2017, pp. 27–40.
- [38] S. Zhang et al., "Cambricon-X: An accelerator for sparse neural networks," in Proc. MICRO, 2016, pp. 1–12.
- [39] J. Albericio, P. Judd, T. Hetherington, T. Aamodt, N. E. Jerger, and A. Moshovos, "Cnvlutin: Ineffectual-neuron-free deep neural network computing," in *Proc. 43rd Int. Symp. Comput. Archit.*, 2016, pp. 1–13.
- [40] A. Gondimalla, N. Chesnut, M. Thottethodi, and T. N. Vijaykumar, "SparTen: A sparse tensor accelerator for convolutional neural networks," in *Proc. 52nd Annu. IEEE/ACM Int. Symp. Microarchit.*, 2019, pp. 151–165.
- [41] J. Albericio et al., "Bit-pragmatic deep neural network computing," in Proc. 50th Annu. IEEE/ACM Int. Symp. Microarchit., 2017, pp. 382–394.
- [42] P. Judd, J. Albericio, T. Hetherington, T. M. Aamodt, and A. Moshovos, "Stripes: Bit-serial deep neural network computing," in *Proc. 49th Annu. IEEE/ACM Int. Symp. Microarchit. (MICRO).* 2016, pp. 1–12.
- [43] S. Li, W. Wen, Y. Wang, S. Han, Y. Chen, and H. Li, "An FPGA design framework for CNN sparsification and acceleration," in *Proc. FCCM*, 2017, p. 28.
- [44] J. H. Ko, B. A. Mudassar, T. Na, and S. Mukhopadhyay, "Design of an energy-efficient accelerator for training of convolutional neural networks using frequency-domain computation," in *Proc. DAC*, 2017, p. 59.
- [45] N. Srivastava, H. Jin, J. Liu, D. Albonesi, and Z. Zhang, "MatRaptor: A sparse-sparse matrix multiplication accelerator based on row-wise product," in *Proc. 53rd Annu. IEEE/ACM Int. Symp. Microarchit.(MICRO)*. 2020, pp. 766–780.
- [46] N. K. Srivastava, H. Jin, S. Smith, H. Rong, D. H. Albonesi, and Z. Zhang, "Tensaurus: A versatile accelerator for mixed sparse-dense tensor computations," in *Proc. HPCA*, 2020, pp. 689–702.
- [47] L. Lu, Y. Liang, Q. Xiao, and S. Yan, "Evaluating fast algorithms for convolutional neural networks on FPGAs," in *Proc. FCCM*, 2017, pp. 101–108.
- [48] Y. Niuet al., "SPEC2: Spectral sparse cnn accelerator on FPGAs," in Proc. IEEE 26th Int. Conf. High Perform. Comput. Data Anal. (HiPC), 2019, pp. 195–204.
- [49] W. You and C. Wu, "A reconfigurable accelerator for sparse convolutional neural networks," in *Proc. ACM/SIGDA Int. Symp. Field-Programmable Gate Arrays (FPGA)*, 2019, p. 119.
- [50] Mentor. (2018). Catapult High-Level Synthesis. [Online]. Available: https://www.mentor.com/hls-lp/
- [51] Y. Jia et al., "Caffe: Convolutional architecture for fast feature embedding," 2014. [Online]. Available: arXiv:1408.5093.
- [52] A. Paszke *et al.*, "PyTorch: An imperative style, high-performance deep learning library," 2019. [Online]. Available: arXiv:1912.01703.
- [53] K. Hegde et al., "ExTensor: An accelerator for sparse tensor algebra," in Proc. Int. Symp. Microarchit., 2019, pp. 319–333.
- [54] M. Sandler, A. G. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "MobileNetv2: Inverted residuals and linear bottlenecks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 4510–4520.
- [55] X. Zhang et al., "AccDNN: An IP-based DNN generator for FPGAs," in Proc. FCCM, 2018, p. 210.
- [56] U. Aydonat, S. O'Connell, D. Capalija, A. C. Ling, and G. R. Chiu, "An OpenCLTM deep learning accelerator on arria 10," in *Proc. FPGA*, 2017, pp. 55–64.
- [57] C. Zhang and V. K. Prasanna, "Frequency domain acceleration of convolutional neural networks on CPU-FPGA shared memory system," in *Proc. FPGA*, 2017, pp. 35–44.
- [58] L. Lu and Y. Liang, "SpWA: An efficient sparse winograd convolutional neural networks accelerator on FPGAs," in *Proc. DAC*, 2018, pp. 1–6.
- [59] Y. Liang, L. Lu, Q. Xiao, and S. Yan, "Evaluating fast algorithms for convolutional neural networks on FPGAs," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 39, no. 4, pp. 857–870, Apr. 2020.

- [60] S. Han et al., "EIE: Efficient inference engine on compressed deep neural network," in Proc. ISCA, 2016, pp. 243–254.
- [61] H. Liu, K. Simonyan, and Y. Yang, "DARTS: Differentiable architecture search," 2018. [Online]. Available: arXiv:1806.09055.
- [62] C. Hao et al., "FPGA/DNN co-design: An efficient design methodology for IoT intelligence on the edge," Proc. 56th ACM/IEEE Design Autom. Conf. (DAC), 2019, p. 206.
- [63] Y. Li et al., "EDD: Efficient differentiable DNN architecture and implementation co-search for embedded AI solutions," in Proc. 57th ACM/IEEE Design Autom. Conf. (DAC), 2020, pp. 1–6.



Yun (Eric) Liang (Senior Member, IEEE) received the Ph.D. degree in computer science from the National University of Singapore, Singapore, in 2010.

He is an Associate Professor (with tenure) with the Center for Energy-efficient Computing and Applications, School of EECS, Peking University, China. He has authored over 90 scientific publications in premier international journals and conferences in related domains. His research interests include computer architecture, compiler, electronic

design automation, and embedded system.

Dr. Liang's research has been recognized by best paper awards at FCCM 2011 and ICCAD 2017 and best paper nominations at PPoPP 2019, DAC 2017, ASPDAC 2016, DAC 2012, FPT 2011, and CODES+ISSS 2008. He serves as an Associate Editor for ACM Transactions in Embedded Computing Systems, ACM Transactions on Reconfigurable Technology and Systems, and Embedded System Letters. He also serves in the program committees in the premier conferences in the related domain, including MICRO, DAC, HPCA, FPGA, ICCAD, FCCM, and ICS.



Liqiang Lu received the B.S. degree from the Institute of Microelectronics, Peking University, Beijing, China, in 2017, where he is currently pursuing the Ph.D. degree with the School of EECS.

His research focuses on algorithm-level and architecture-level optimizations of FPGA for machine learning applications.



Yicheng Jin received the B.S. degree in computer science from Yuanpei College, Peking University, Beijing, China, in 2020, where he is currently pursuing the undergraduate degree.

He holds interest in machine learning systems and software–hardware co-design for machine learning applications.

Jiaming Xie, photograph and biography not available at the time of publication.



Ruirui (Raymond) Huang (Member, IEEE) received the Ph.D. degree in computer architecture research with a focus on highly reliable, secure, and available computing design from Cornell University, Ithaca, NY, USA, in 2013.

He is the Director of Cloud Architecture and is responsible for the overall cloud architecture design of the cloud platform and products in Alibaba Cloud, Hangzhou, China. Before joining Alibaba Cloud, he was an Architect for the Xeon server SoC in Intel, Mountain View, CA, USA. He has published papers

in ASPLOS, ISCA, HPCA, DAC, FSE, and FPGA.



Jiansong Zhang (Member, IEEE) received the bachelor's and master's degrees from Tsinghua University, Beijing, China, in 2004 and 2006, respectively, and the Ph.D. degree from the Hong Kong University of Science and Technology, Hong Kong, in 2014.

He is a Research Scientist with Alibaba Damo Academy who works on heterogeneous computing for machine learning and big data processing. Before joining Alibaba Group, Hangzhou, China, he was a Researcher working on heterogeneous computing

for wireless communication and networking, IoT systems, and datacenter acceleration with Microsoft Research Asia. He has published 10 papers in Sigcomm, NSDI, Mobicom, Ubicomp, Mobisys, HotNets, HotChips, FCCM, Infocom, and MLSys. His research interests are building innovative software and hardware systems for various applications.

Wei Lin, photograph and biography not available at the time of publication.