

Automatic Compiler Based FPGA Accelerator for CNN Training

Shreyas Kolala Venkataramaiah, Yufei Ma, Shihui Yin, Eriko Nurvithadhi*, Aravind Dasu†, Yu Cao, Jae-sun Seo

School of Electrical, Computer and Energy Engineering, Arizona State University, Tempe, AZ, USA

*Intel Labs, Intel Corporation, Hillsboro, OR, USA

†Programmable Solutions Group, Intel Corporation, San Jose, CA, USA

Email: skvenka5@asu.edu

Abstract—Training of convolutional neural networks (CNNs) on embedded platforms to support on-device learning is earning vital importance in recent days. Designing flexible training hardware is much more challenging than inference hardware, due to design complexity and large computation/memory requirement. In this work, we present an automatic compiler based FPGA accelerator with 16-bit fixed-point precision for complete CNN training, including Forward Pass (FP), Backward Pass (BP) and Weight Update (WU). We implemented an optimized RTL library to perform training-specific tasks, and developed an RTL compiler to automatically generate FPGA-synthesizable RTL based on user-defined constraints. We present a new cyclic weight storage/access scheme for on-chip BRAM and off-chip DRAM to efficiently implement non-transpose and transpose operations during FP and BP phases, respectively. Representative CNNs for CIFAR-10 dataset are implemented and trained on Intel Stratix 10 GX FPGA using proposed hardware architecture, demonstrating up to 479 GOPS performance.

Index Terms—Convolution neural networks, neural network training, back-propagation, hardware accelerator, FPGA

I. INTRODUCTION

CNNs have shown tremendous performance in many practical tasks including computer vision [1] and speech recognition [2]. Deep CNNs achieve high accuracy on large datasets, but an enormous amount of computation is required for training such networks. To support the high computation requirement, training tasks have been typically performed on datacenters with high-end GPUs. Nowadays, training on resource-constrained platforms is becoming more crucial for training networks with each user's private data. However, executing computation-/memory-intensive training tasks on hardware platforms with power and resource constraints become very challenging. This gives an opportunity to map these algorithms on FPGAs, which provide high configurability and power-efficiency compared to those of GPUs. They also provide a large volume of off-chip memory (DRAM) and shorter design time when compared to ASIC designs.

For CNN inference tasks, a number of FPGA accelerators have been proposed [3]–[7]. However, training deep neural networks on FPGA platform has not been investigated comprehensively. Compared to inference, the training phase

The authors would like to thank Intel Corporation for supporting and funding this research work. This work was also partially supported by NSF grant 1652866 and C-BRIC, one of six centers in JUMP, a SRC program sponsored by DARPA.

involves a much higher number of operations (>3X) with increased complexity [8]. The training phase also involves high intermediate data volume, necessitating high memory bandwidth and large storage. GPUs have been the de-facto for training tasks to meet immense computation requirements. However, GPUs' energy-efficiency is poor [9], and they are not well-suited for on-device learning with limited power budget.

To address this issue on the algorithm side, researchers have proposed low-precision training [10], [11], frequency domain training [12], and sparse weight update [13]. Techniques such as sparse weight update introduce irregular parallelism, making it more suitable for flexible FPGAs compared to GPUs [14]. FPGAs are well-suited for low-precision DNN algorithms as it provides large improvement in throughput and energy efficiency with low-precision arithmetic [15]. To that end, implementing configurable training hardware on FPGA becomes crucial to exploit these algorithmic advances.

On the hardware side, several prior FPGA works have implemented training of fully-connected neural networks [16]–[18]. A floating-point FPGA accelerator [19] reported training of small CNNs using an uniform computation structure with a fixed number of multiply-and-accumulate (MAC) units. F-CNN [20] presented a training framework where convolutions are done in FPGA and weight updates are performed in CPU. TrainWare [8] implemented dedicated hardware for weight update using a fixed $N_{kx} \times N_{ky}$ MAC array as the local gradients window is reused only $N_{kx} \times N_{ky}$ times during weight gradient computation. However, this is not suitable for FP/BP convolutions where there exists more kernel reuse. DeepTrain [21] presents an embedded platform for DNN training, but does not include back-propagation of pooling layers and DNN weight updates, which needs significant memory access. Overall, these works have not presented a compiler-based FPGA accelerator that supports all phases of training for various CNNs. Designing a standalone FPGA accelerator for CNN training involves managing limited memory resources to support batch operations and different CNN configurations.

In this work, we propose a flexible FPGA accelerator that performs stochastic gradient descent (SGD) based training of various CNNs. We extracted and designed training-specific operations and then developed a library based automatic RTL compiler to flexibly support training operations with different sizes of CNNs. The user provides the high-level CNN network

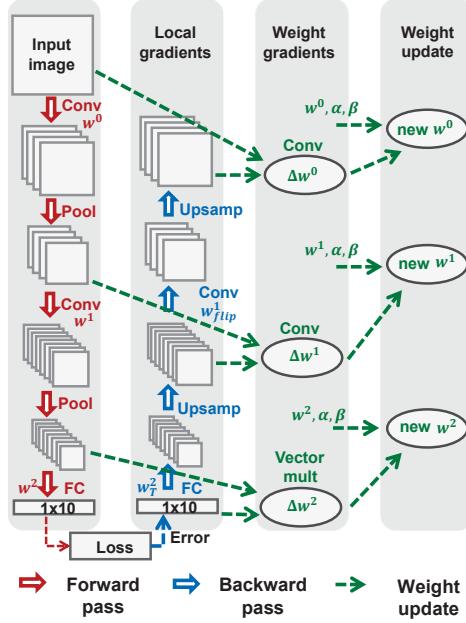


Fig. 1: SGD based CNN training dataflow illustrated for a simple 2C-2P-1FC model.

configurations along with the design variables to characterize FPGA hardware usage to the RTL compiler. The RTL compiler generates a FPGA compatible training accelerator based on the user's requirements. The key contributions of this work are:

- We present a comprehensive investigation of CNN training operations and challenges in FP, BP and WU stages.
- We developed a training-specific RTL module library and an RTL compiler to automatically implement CNN training accelerator with 16-bit fixed-point precision.
- A configurable FPGA hardware is presented for FP, BP and WU phases of the entire CNN training process using SGD with momentum.
- Our accelerator using Intel Stratix 10-GX FPGA is evaluated on training three different CNNs for CIFAR-10 dataset, achieving up to 479 GOPS of throughput.

II. CNN TRAINING ALGORITHM

Fig. 1 illustrates the dataflow of SGD based weight update for a simple 2C-2P-1FC CNN model. The CNN design variables and naming conventions are described in Table I. Output activation value $o_{x,y}^l$ is given by Eq. (1), where $w_{x,y}^l$ are kernel values and $a_{x,y}^{l-1}$ are activations from layer $l-1$.

$$o_{x,y}^l = \sum_{x'} \sum_{y'} w_{x,y}^l a_{(x+x'),(y+y')}^{l-1} \quad (1)$$

TABLE I: CNN design variables

	Kernel size width/height	Output feature map width/height/depth	Input feature map width/height/depth
Convolution dimensions	N_{kx}, N_{ky}	N_{ox}, N_{oy}, N_{of}	N_{ix}, N_{iy}, N_{if}
Loop unroll factors	P_{kx}, P_{ky}	P_{ox}, P_{oy}, P_{of}	P_{ix}, P_{iy}, P_{if}

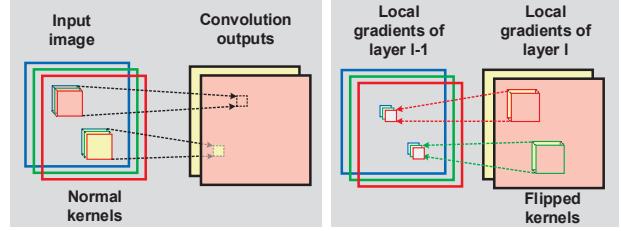


Fig. 2: Convolution operations and changes in kernels during FP and BP ($N_{of} = 2, N_{if} = 3$).

In supervised training, each input is associated with a label. After the completion of the FP, the performance of the network is estimated using a cost function. Eq. (2) shows a quadratic cost function of output layer L , where a_i is the obtained output value and y_i is the label. The derivative of the cost function with respect to output is also given in Eq. (2).

$$C = \frac{1}{2} \sum_i^L (a_i - y_i)^2, \quad \frac{\partial C}{\partial a_i^L} = (a_i - y_i) \quad (2)$$

Error values are back-propagated to all hidden layers and the required deviation of weight parameters to minimize the error is calculated. The derivative of the cost function with respect to weight parameters provides the required deviation for the weight parameters Δw to minimize the error. By applying the basic chain rule, weight deviation Δw can be obtained by convolving the derivative of the cost function with layer output activations, which we term as local gradients and feedforward activations. Local gradients of layer (l) can be obtained by convolving the gradients of the previous layer ($l-1$) with its own convolution kernel.

During these backward convolutions, the original kernel tensors are flipped. The differences of BP and FP convolutions are shown in Fig. 2. Fig. 2a shows FP convolutions of input image with three input channels ($N_{if} = 3$) and two sets of kernels to obtain two output feature maps ($N_{of} = 2$). During BP, convolutions are performed using local gradients of previous layer and FP kernels, where the number of input channels and convolution depth are interchanged. In Fig. 2b, it is shown that $N_{if} = 2$ and $N_{of} = 3$. Flipped kernels are used in BP convolutions to compute the local gradients.

$$\delta_{x,y}^l = \varphi_l(o_{x,y}^l) \sum_x \sum_{y'} \delta_{x',y'}^{l+1} w_{(x-x'),(y-y')}^{l+1} \quad (3)$$

$$\Delta w_n = \frac{\partial C}{\partial w_{x,y}^L} = \sum_x \sum_{y'} \delta_{x',y'}^l a_{(x+x'),(y+y')}^{l-1} \quad (4)$$

$$w_{i,j}^l(n) = -\alpha \Delta w_n + w_{i,j}^l(n-1) \quad (5)$$

$$w_{i,j}^l(n) = \beta \Delta w_{n-1} - \alpha \Delta w_n + w_{i,j}^l(n-1) \quad (6)$$

Local gradients of each layer l is computed using Eq. (3), where w is the flipped kernel. Eq. (4) is used for weight gradient computation, where l is local gradients of a layer and $\varphi_l(x)$ is activation gradients of layer l . The weight gradients of layer l is obtained by the convolution of local gradient layer l and

feedforward input activations of layer l , involving large kernel sizes. One feature map of feedforward activation is convolved with one feature map of local gradients to obtain one kernel gradient (intra-tile accumulation). Hence, this weight gradient convolution results in a 4D output. These weight gradients are averaged over a batch and new weights are computed using gradient descent algorithm given by Eq. (5), where α is learning rate, $w_{i,j}^l(n-1)$ is weights of previous batch and Δw_n is the average weight gradient. The weight update process can be accelerated by using past weight gradients as momentum. Eq. (6) shows the weight update in SGD with momentum, where β is a hyper-parameter.

The operations during BP are different to those of FP. In backward convolutions, the inputs are scaled by activation gradients, and convolutions are performed by applying 180-degree-rotated kernels. Similarly, fully-connected layers in BP also use transposed weight matrix to compute the local gradients. At the max-pooling node, the gradients propagate only through the selected maximum pixel location and all other pixels in the pooling window will be zero. Based on the pooling pixel index selected during FP, the gradients are upsampled and propagated back to the next layers.

During FP, we need to store not only the output activations, but also the activation gradients and max-pooling indices at all ReLU activations and max-pooling nodes. For ReLU, activation gradients are binary as the derivative of ReLU with respect to activations results in a step function. Our RTL library currently supports only ReLU activation function as it is less complex and widely used. During weight update of fully-connected layers, the weight gradients Δw are obtained by performing the outer product of the local gradient vector and the error vector. In convolution kernel updates, kernel gradient calculation involves convolution of input activations using local gradients as kernels, which are very large kernels. Each of these convolutions is considered as an FP convolution with $N_{if} = 1$ and results in N_{of} kernel gradients. To reuse FP convolution control logic, we employed an additional outer loop to iterate through the actual N_{if} local gradients.

Unlike CNN inference, CNN training usually requires higher precision. In this work, weights, activations, and local/weight gradients are represented with 16-bit fixed-point precision to ensure good training accuracy [10], [22]. Compared to floating-point precision, fixed-point precision training leads to more energy-efficient FPGA design, but requires more dedicated resolution/range assignment for different variables.

III. CNN TRAINING HARDWARE

A. RTL Compiler and Algorithm Mapping

To map various CNN algorithms with user defined hardware constraints onto FPGA, an RTL compiler for CNN training was developed. Fig. 3 shows the compiler tool flow from high-level CNN description to CNN training accelerator. According to the operations in each layer and FPGA design parameters (e.g. unroll and tiling factors), optimized handwritten Verilog modules are chosen from the RTL library to automatically generate a CNN training accelerator. The RTL library consists

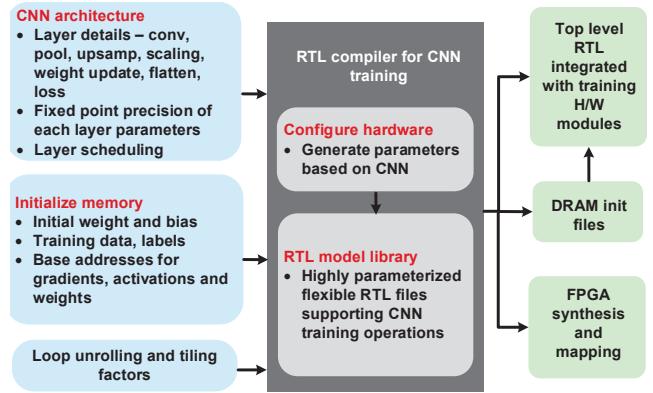


Fig. 3: Proposed RTL compiler automatically generates FPGA training accelerator from high-level CNN description.

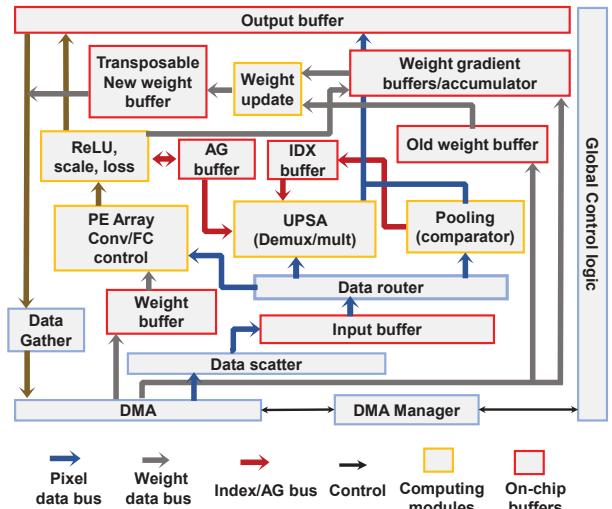


Fig. 4: Top-level block diagram of CNN training accelerator.

of Verilog modules that are specially designed to support training operations. Only the selected modules from the RTL library based on the training algorithm will be synthesized. Execution of training operations in one iteration of a batch can be scheduled sequentially similar to layer-by-layer execution of inference tasks. Each training image in a batch is processed sequentially. The scheduling of layer execution is done using the RTL compiler, and control logic parameters are generated.

B. Training Accelerator Architecture

Fig. 4 shows the top-level diagram and dataflow of the CNN training accelerator. The global control logic governs all modules to ensure proper CNN functionalities with layer-by-layer computation, and is controlled by the parameters generated by the RTL compiler. DRAM stores all the initial weight parameters, intermediate activations and computed weight/loss gradients using 16-bit fixed-point precision. DMA control generates the required DMA descriptors based on the layer type and tile sizes to read from and write to DRAM. A tile is a portion of data stored in on-chip buffers after/before

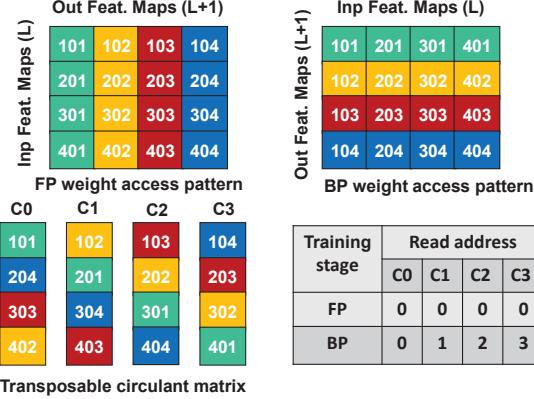


Fig. 5: Proposed transposable weight buffer stores weights in a circulant matrix, enabling both normal and transpose read.

reading/writing back to DRAM. Convolution, max-pooling and upsampling operations are considered as *key layers*, and ReLU, flatten, loss unit, and scaling unit are referred to as *affiliated layers*. Key layers read new data from DRAM and affiliated layers use outputs of key layers.

On-chip buffers store activation gradients and max-pooling indices. The pooling window size (e.g. 2x2) determines the bitwidth of max-pooling indices (e.g. 2-bit). After FP, loss is computed using outputs and labels. Our RTL library currently supports square hinge loss and euclidean loss functions, and this can be easily expanded to support other loss functions. Data scatter and data gather modules are used to convert the DRAM storage pattern to on-chip buffer storage pattern and vice versa. Data router reads the data from input buffers and routes it to the selected key layer according to the array sizes. Weight update unit and weight gradient buffers are used to compute new weights based on SGD with momentum.

C. MAC array

Fig. 6 shows the 2D systolic MAC array used for the training accelerator. MAC array size is determined by the RTL compiler based on the loop unroll factors P_{ox} , P_{oy} , P_{of} . In Fig. 6, each MAC row has a different set of weights but share the same input feature map data computing P_{of} output pixels. Each column shares the same weights, but different input data computing P_{ox} or P_{oy} output pixels in parallel. Data router reads the input data and routes it to MAC units considering pad and stride sizes of the layer. Weight router distributes weights or local gradients based on the training phase. Table in Fig. 6 summarizes how the MAC array is reused with different inputs/outputs for training phases of FP, BP and WU.

D. Transposable Weight Buffer

BP involves convolution of flipped kernels and the local gradients. Therefore, every convolution kernel is used twice in one iteration: 1) normal weights are applied during FP, and 2) rotated weights are used in BP (Fig. 2). To achieve this without duplicating kernel storage, the kernels are stored in special transposable buffers that we propose, where data

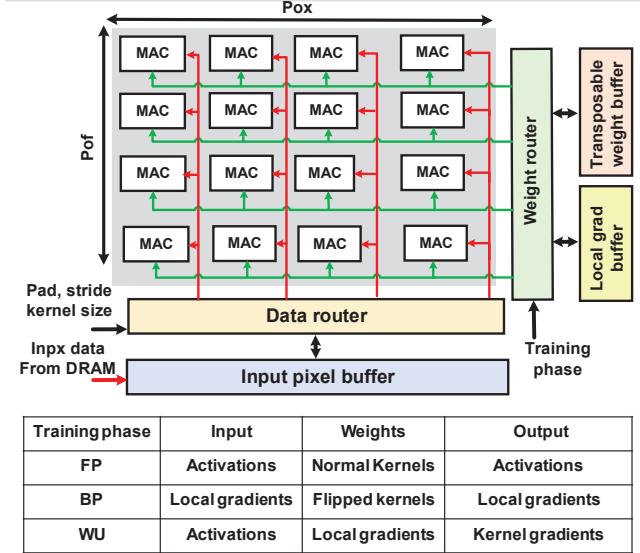


Fig. 6: Systolic MAC array is reused for training phases of FP, BP and WU, by feeding different activations/gradients/kernels.

can be read both in non-transpose and transpose modes. As shown in Fig. 5, the proposed transposable buffer stores the kernels in the form of a circulant matrix using column buffers. For 2D kernels, each $N_{kx} \times N_{ky}$ kernel is considered as one block and each row has P_{of} blocks of kernels, where P_{of} represents the number of output feature maps that can be computed in parallel. During backward convolution, not only the kernel is rotated by 180 degrees but also the input and output feature maps will be interchanged. In the proposed transposable buffer, every row of kernel blocks is circularly rotated and stored in the form of a circulant matrix in the single-port column buffers (Fig. 5). In the non-transpose mode, each column buffer shares the same read address, and in transpose mode, each column buffer obtains shifted addresses from the address translator unit. Address translator generates read/write addresses for column buffers for every transposable block. In each transposable block, the address vectors and the data are circularly shifted using shift registers.

E. Weight Update Unit

Weight gradients are calculated by convolving the feedforward activations with the local gradients. Convolution control logic is configurable to support tile-by-tile computation, intra-tile accumulation and large kernel sizes needed for weight gradient computation. Fig. 7 shows the dataflow after the computation of weight gradients. For every new training image in a batch, newly computed weight gradients are accumulated with old weight gradients. This accumulation is done tile-by-tile and repeated for the entire batch of images while the accumulated gradients are stored in DRAM. At the end of the batch, as the weight gradients get accumulated, old weights and past weight gradients are also read from DRAM, and new weights are computed following Eq. (6).

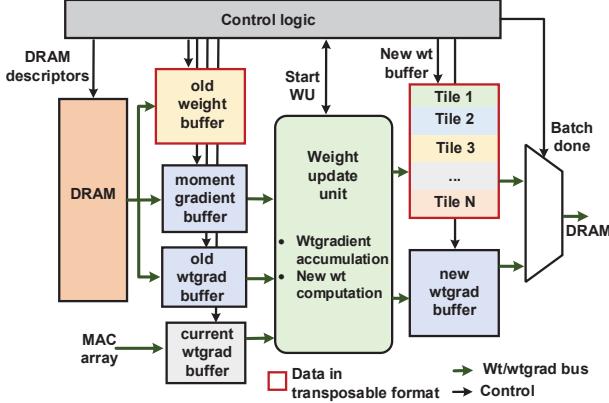


Fig. 7: Block diagram of weight update unit.

Weights are initially stored in transposable format in DRAM as aforementioned. The entire transposable weights of layer l are read from DRAM to the old weight buffer. New weights are computed tile-by-tile and written back in transposable format to the new weight buffer. After completing the last tile's computation, the new weights are written back to DRAM. Control logic translates the address for transposable read/write operations, generates DRAM descriptors according to tile count and generates addresses to read newly computed weight gradients. Fully-connected weight update follows the same dataflow, but gradients are computed by outer product of local gradient vector and activation vector. 16-bit fixed-point precision is used for all weights and gradient computation.

F. Efficient MAC Usage in Weight Update layers

During FP and BP, the MAC array is designed to compute convolutions for $P_{ox} \times P_{oy} \times P_{of}$ pixels in parallel. Regarding convolutions required for weight updates, however, the output feature map size N_{ox} , N_{oy} is less as the outputs are kernel gradients. This results in inefficient usage of MAC units, since most of them will be idle. It also consumes more output buffer storage in order to store $P_{ox} \times P_{oy} \times P_{of}$ block of output data. To overcome this, MAC load balance unit was designed to utilize the idle MAC units.

The MAC load balance unit employs additional input buffers to feed the data to the MAC units in parallel. If buffer usage is critical, this optimization can be disabled by the RTL compiler. Fig. 8 shows the operation of MAC load balancing unit, when $P_{ox}=8$, $P_{oy}=8$, $P_{of}=16$ and kernel size is $N_{ox}=3$, $N_{oy}=3$, $N_{of}=16$. In this example, four kernel gradients are computed in parallel, reducing the latency by 4X without additional MAC units. The output buffer is also efficiently used.

G. Upsampling and Scaling module

During BP, the local gradient at the max-pooling node is propagated to convolution layers only through the maximum pixel position selected in FP. The gradients of unselected pixels are zero, as they do not contribute to the error. If the max-pooling unit receives the input from ReLU node, then the upsampled gradients should also be scaled by the feedforward activation gradients to compute the gradients of ReLU node.

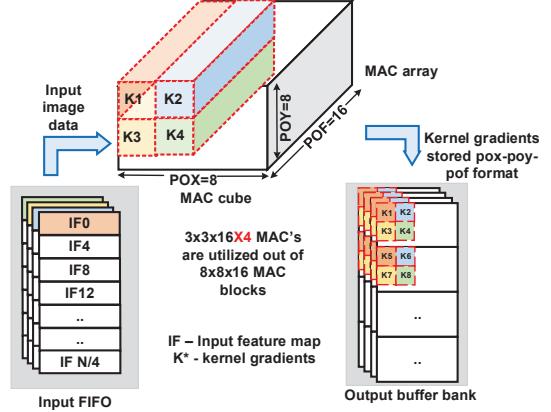


Fig. 8: Operation of MAC load balancing unit during convolution weight gradient computation.

During FP, max-pooling indices are stored tile-by-tile inside the on-chip index buffers. Each layer has its own index and activation gradient buffers. The local gradients computed in the previous iteration is read from DRAM and stored in input buffers. Data router unit rearranges the data of index, input and activation gradient buffers and sends it to the upsampling unit. Each upsampling unit consists of a demultiplexer and a multiplier unit. The gradient is conveyed as the demultiplexer input and the index serves as the select signal. For pooling window size of k , each processing block produces $k \times k$ pixel data corresponding to k rows of the output feature map. After each operation, k rows of activation gradients are read and the demultiplexer outputs are scaled.

IV. RESULTS

A. Experimental Setup

The FPGA accelerator generated by the compiler was synthesized using Intel Quartus 17.1 at 240MHz frequency. We used Stratix 10 GX FPGA as the target hardware, which includes 240 Mbits of BRAM, 5,760 DSP blocks, and 93K ALMs. The development kit [23] is equipped with 4Gb DDR3 DRAM with 16.9Gb/s bandwidth. Weights, weight/local gradients, and activations use 16-bit fixed point precision. We trained representative CNNs for CIFAR-10 dataset. '1X' CNN has the structure of 16C3-16C3-P-32C3-32C3-P-64C3-64C3-P-FC. 2X and 4X CNN models exhibit 2X and 4X more input/output feature maps for each layer, and could achieve higher accuracy. Unroll factor of 8 was used for output image x and y dimensions. For output feature maps, 16, 32, 64 was used as unroll factors for 1X, 2X and 4X CNNs, resulting in 8x8x16 (1,024), 8x8x32 (2,048), 8x8x64 (4,096) MAC arrays, respectively. Batch size (BS) of up to 40 and learning rate of 0.002 was used for training. Latency was measured using simulation of the synthesized accelerator. DRAM modules and Intel IPs were used in the testbench adhering to DRAM protocols. We also developed a custom fixed-point precision training model using PyTorch [24] to verify the functionality of the FPGA design with the same precision.

TABLE II: Evaluation of CNN training accelerator on Stratix 10 FPGA , using 16-bit fixed point precision. CIFAR10-1X refers to network structure of 16C3-16C3-P-32C3-32C3-P-64C3-64C3-P-FC, and 2X/4X designs refer to accordingly wider CNNs.

CNN network	Resource			Power (W)					Latency per epoch (s)			Throughput GOPs
	DSP	ALM	BRAM	DSP	RAM	Logic	clock	Pstatic	BS-10	BS-20	BS-40	
CIFAR-10 1X	1699 (30%)	20.8K (19%)	10.6(4.4%)	0.58	5.7	2.4	1.68	10.28	18.19	18.07	18.01	163
CIFAR-10 2X	3363 (58%)	415K (44%)	22.8(9.5%)	1.05	11.2	6.6	2.97	11	41.7	41.30	41	282
CIFAR-10 4X	5760(100%)	720K(76.2%)	54.5(22.4%)	3.48	14.6	11	4.95	16.47	98.2	96.87	96.18	479

TABLE III: Performance comparison with GPU.

Device	Throughput (GOPS)		Efficiency (GOPS/W)	
	Titan XP	FPGA	Titan XP	FPGA
Batch size	1	40	1/40	1
CIFAR-10 1X	45.67	551.87	163	0.50
CIFAR-10 2X	128.84	1337.98	282	1.30
CIFAR-10 4X	331.41	2353.79	479	2.91

B. Results and Analysis

Table II shows the comparison of CNN training performance and resource utilization for three different CNNs for CIFAR-10 dataset. The FPGA accelerator was generated from the RTL compiler using high-level description of training parameters and design variables. FPGA power numbers are obtained after routing stage from Quartus power analyzer and Intel Early Power Estimator tools using the data toggling activity from functional simulation at the junction temperature of 65°C. Tiling of activations and weight gradients greatly reduces the on chip buffer usage. BRAM utilization is low because of the tiling and size of the intermediate activations and number of parameters. Training of each image in a batch is done sequentially, larger batch sizes results in less number of weight updates in one epoch resulting in improvement in latency.

Performance comparison of our accelerator implementation on Stratix 10 FPGA and Titan XP GPU is shown in Table III. Our performance remains the same for different batch sizes as the images in a batch are processed sequentially one after the other. Our implementation shows better energy efficiency for smaller batch sizes. For batch size of 40, the 4X model shows less energy-efficiency than GPU, due to limited DRAM bandwidth (30X less than Titan XP). Stable and reliable training can also be achieved with smaller batch sizes as it provides more up-to-date gradient calculations [25].

To flexibly support arbitrary sizes of CNNs, all intermediate outputs are stored in DRAM. Fig. 9 shows the latency breakdown during different stages of training. Weight update layers will have large DRAM access latency due to access of

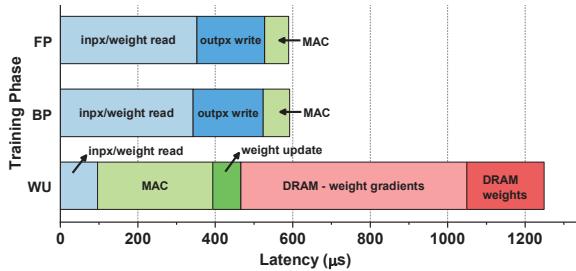


Fig. 9: Latency breakdown of CIFAR-10 4X CNN for FP, BP and WU for the last iteration of a batch.

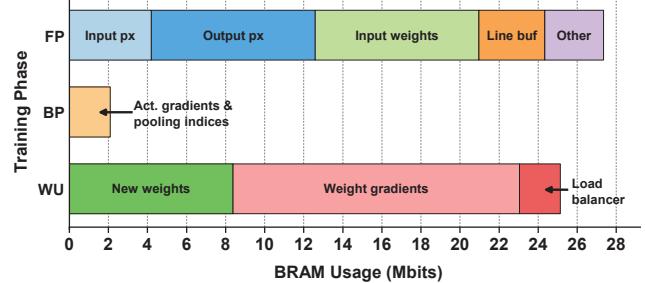


Fig. 10: Buffer usage breakdown of CIFAR-10 4X CNN.

past weight gradients, weights and storing back the updated values. 51% percent of the overall latency in one iteration of a batch is consumed in weight update layers. By sacrificing the flexibility of the hardware, this latency could be significantly reduced by using on-chip buffers for weight/gradient storage.

Old weight gradients are read from DRAM tile-by-tile during computation of current weight gradients. Double buffering scheme is employed to hide the memory access latency [3], which reduced the latency of weight update layers by 11%. The logic latency in weight update layers is reduced by 4X, using the load balancing technique for MAC arrays. Logic in weight update layers refer to convolution operations to generate weight gradients and weight update is referred to computation of new weights. Tile sizes are carefully chosen to efficiently map compute-/memory-bounded layers. All buffers can be controlled by tile sizes apart from weight buffers, where the entire weights are read from transposable DRAM.

Fig. 10 shows the breakdown of buffer utilization for three different phases of training. The weight buffer size is decided by the largest layer weights. Double buffering technique is used for all other buffers, thereby hiding DRAM latency. The 1X design achieves 73% CIFAR-10 accuracy at 50 epochs with learning rate of 0.002 and batch size of 40 (similar to baseline with floating-point precision). Higher accuracy will be achievable with addition of integer batch normalization and adaptive fixed point features [22] to our RTL module library.

V. CONCLUSION

In this paper, we presented an automatic RTL compiler based end-to-end CNN training accelerator. CNN training operations are implemented by optimized and parameterized custom Verilog modules, and the accelerator is flexible to support various FPGA design parameters. The training performance is evaluated on Intel Stratix-10 GX FPGA for three different CNNs for CIFAR-10 dataset. The proposed training accelerator achieves throughput of up to 479 GOPS at 240MHz for CNNs with 2M parameters.

REFERENCES

- [1] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 7132–7141, 2018.
- [2] Y. Zhang, M. Pezeshki, P. Brakel, S. Zhang, C. Laurent, Y. Bengio, and A. Courville, "Towards end-to-end speech recognition with deep convolutional neural networks," in *INTERSPEECH*, 2016.
- [3] C. Zhang, P. Li, G. Sun, Y. Guan, B. Xiao, and J. Cong, "Optimizing FPGA-based accelerator design for deep convolutional neural networks," in *Proceedings of the ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*, pp. 161–170, 2015.
- [4] Y. Ma, Y. Cao, S. Vrudhula, and J. Seo, "An automatic RTL compiler for high-throughput FPGA implementation of diverse deep convolutional neural networks," in *Proceedings of the International Conference on Field Programmable Logic and Applications (FPL)*, pp. 1–8, 2017.
- [5] J. Zhang and J. Li, "Improving the performance of OpenCL-based FPGA accelerator for convolutional neural network," in *Proceedings of the ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA)*, pp. 25–34, 2017.
- [6] H. Zeng, R. Chen, C. Zhang, and V. Prasanna, "A framework for generating high throughput CNN implementations on FPGAs," in *Proceedings of the ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA)*, pp. 117–126, 2018.
- [7] Y. Yang, Q. Huang, B. Wu, T. Zhang, L. Ma, G. Gambardella, M. Blott, L. Lavagno, K. Vissers, J. Wawrzynek, and K. Keutzer, "Synetgy: Algorithm-hardware co-design for ConvNet accelerators on embedded FPGAs," in *Proceedings of the ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA)*, pp. 23–32, 2019.
- [8] S. Choi, J. Sim, M. Kang, and L.-S. Kim, "TrainWare: A memory optimized weight update architecture for on-device convolutional neural network training," in *Proceedings of the International Symposium on Low Power Electronics and Design (ISLPED)*, 2018.
- [9] N. P. Jouppi, C. Young, N. Patil, D. Patterson, G. Agrawal, R. Bajwa, S. Bates, S. Bhatia, N. Boden, A. Borchers, *et al.*, "In-datacenter performance analysis of a tensor processing unit," in *Proceedings of the ACM/IEEE Annual International Symposium on Computer Architecture (ISCA)*, pp. 1–12, 2017.
- [10] S. Gupta, A. Agrawal, K. Gopalakrishnan, and P. Narayanan, "Deep learning with limited numerical precision," in *Proceedings of the International Conference on Machine Learning (ICML)*, pp. 1737–1746, 2015.
- [11] U. Köster, T. Webb, X. Wang, M. Nassar, A. K. Bansal, W. Constable, O. Elibol, S. Gray, S. Hall, L. Hornof, A. Khosrowshahi, C. Kloss, R. J. Pai, and N. Rao, "Flexpoint: An adaptive numerical format for efficient training of deep neural networks," in *Advances in Neural Information Processing Systems*, pp. 1742–1752, 2017.
- [12] J. H. Ko, B. Mudassar, T. Na, and S. Mukhopadhyay, "Design of an energy-efficient accelerator for training of convolutional neural networks using frequency-domain computation," in *Proceedings of the ACM/EDAC/IEEE Design Automation Conference (DAC)*, pp. 1–6, 2017.
- [13] X. Sun, X. Ren, S. Ma, and H. Wang, "meProp: sparsified back propagation for accelerated deep learning with reduced overfitting," in *Proceedings of the International Conference on Machine Learning (ICML)*, pp. 3299–3308, 2017.
- [14] E. Nurmukhamed, G. Venkatesh, J. Sim, D. Marr, R. Huang, J. Ong Gee Hock, Y. T. Liew, K. Srivatsan, D. Moss, S. Subhaschandra, *et al.*, "Can FPGAs beat GPUs in accelerating next-generation deep neural networks?," in *Proceedings of ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA)*, pp. 5–14, 2017.
- [15] E. Wang, J. J. Davis, R. Zhao, H.-C. Ng, X. Niu, W. Luk, P. Y. Cheung, and G. A. Constantinides, "Deep neural network approximation for custom hardware: Where we've been, where we're going," *arXiv preprint arXiv:1901.06955*, 2019.
- [16] Q. Liu, J. Liu, R. Sang, J. Li, T. Zhang, and Q. Zhang, "Fast neural network training on FPGA using quasi-newton optimization method," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 26, no. 8, pp. 1575–1579, 2018.
- [17] A. Gomperts, A. Ukil, and F. Zurfluh, "Development and implementation of parameterized FPGA-based general purpose neural networks for online applications," *IEEE Transactions on Industrial Informatics*, vol. 7, no. 1, pp. 78–89, 2011.
- [18] G. Rafael, C. Ricardo, C. Joaquín, C. Angel, and W. M. Maeda, "FPGA implementation of a pipelined on-line backpropagation," *Journal of VLSI Signal Processing*, vol. 40, no. 2, pp. 189–213, 2005.
- [19] Z. Liu, Y. Dou, J. Jiang, Q. Wang, and P. Chow, "An FPGA-based processor for training convolutional neural networks," in *Proceedings of the International Conference on Field Programmable Technology (ICFPT)*, pp. 207–210, 2017.
- [20] W. Zhao, H. Fu, W. Luk, T. Yu, S. Wang, B. Feng, Y. Ma, and G. Yang, "F-CNN: An FPGA-based framework for training convolutional neural networks," in *Proceedings of the IEEE International Conference on Application-specific Systems, Architectures and Processors (ASAP)*, pp. 107–114, 2016.
- [21] D. Kim, T. Na, S. Yalamanchili, and S. Mukhopadhyay, "Deeptrain: A programmable embedded platform for training deep neural networks," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 37, no. 11, pp. 2360–2370, 2018.
- [22] X. Chen, X. Hu, H. Zhou, and N. Xu, "FxpNet: training a deep convolutional neural network in fixed-point representation," in *Proceedings of the IEEE International Joint Conference on Neural Networks (IJCNN)*, pp. 2494–2501, 2017.
- [23] "Intel Stratix 10 GX Development Kit." https://www.intel.com/content/www/us/en/programmable/products/boards_and_kits/dev-kits/altera-kit-s10-fpga.html.
- [24] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, "Automatic differentiation in PyTorch," in *NIPS 2017 Autodiff Workshop*, 2017.
- [25] D. Masters and C. Luschi, "Revisiting small batch training for deep neural networks," *arXiv preprint arXiv:1804.07612*, 2018.