EF-Train: Enable Efficient On-device CNN Training on FPGA through Data Reshaping for Online Adaptation or Personalization

YUE TANG, XINYI ZHANG, PEIPEI ZHOU, and JINGTONG HU, University of Pittsburgh, USA

Conventionally, DNN models are trained once in the cloud and deployed in edge devices such as cars, robots, or unmanned aerial vehicles (UAVs) for real-time inference. However, there are many cases that require the models to adapt to new environments, domains, or users. In order to realize such domain adaption or personalization, the models on devices need to be continuously trained on the device. In this work, we design EF-Train, an efficient DNN training accelerator with a unified channel-level parallelism-based convolution kernel that can achieve end-to-end training on resource-limited low-power edge-level FPGAs. It is challenging to implement on-device training on resource-limited FPGAs due to the low efficiency caused by different memory access patterns among forward and backward propagation and weight update. Therefore, we developed a data reshaping approach with intra-tile continuous memory allocation and weight reuse. An analytical model is established to automatically schedule computation and memory resources to achieve high energy efficiency on edge FPGAs. The experimental results show that our design achieves 46.99 GFLOPS and 6.09 GFLOPS/W in terms of throughput and energy efficiency, respectively.

CCS Concepts: \bullet Hardware \rightarrow Integrated circuits; Reconfigurable logic and FPGAs; Hardware accelerators;

Additional Key Words and Phrases: On-device training, edge FPGAs, data reshaping

ACM Reference format:

Yue Tang, Xinyi Zhang, Peipei Zhou, and Jingtong Hu. 2022. EF-Train: Enable Efficient On-device CNN Training on FPGA through Data Reshaping for Online Adaptation or Personalization. *ACM Trans. Des. Autom. Electron. Syst.* 27, 5, Article 49 (June 2022), 36 pages.

https://doi.org/10.1145/3505633

1 INTRODUCTION

Deep Neural Networks (DNNs) have been widely used in edge devices such as cars, robotics [1], and **unmanned aerial vehicles (UAVs)** [2] to accomplish various tasks, including autonomous driving, object detection, and so forth. FPGAs are promising platforms with higher computational density, communication bandwidth, and energy efficiency and can be configured based on different tasks. Nowadays, FPGAs have been widely used in various edge device domains. For example, edge-scale FPGAs are commonly utilized in object detection tasks with high frames per

This work is supported by the NSF CNS-2122320 Towards Unsupervised Learning on Resource Constrained Edge Devices with Novel Statistical Contrastive Learning Scheme.

Authors' address: Y. Tang, X. Zhang, P. Zhou, and J. Hu, Department of Electrical and Computer Engineering Swanson School of Engineering University of Pittsburgh 1208 Benedum Hall, 3700 O'Hara St, Pittsburgh, PA, 15261; emails: {yut51, xinyizhang, peipei.zhou, jthu}@pitt.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2022 Association for Computing Machinery.

1084-4309/2022/06-ART49 \$15.00

https://doi.org/10.1145/3505633

49:2 Y. Tang et al.

second and low power consumption [3]. The Corazon-AI built on Xilinx Zynq is a perfect fit for various computer-vision-based applications including video surveillance, advanced driver-assistance systems (ADASs), medical robotics, industrial automation, and augmented reality [4]. Combined with reconfigurability, FPGAs have been adopted in several autonomous platforms such as pony.ai [5] and ZF ProAI [6]. In medical applications, an FPGA-based low-latency multi-layer perception (MLP) processor for real-time cancer detection has been developed, since FPGA-based design can directly interface with sensors, display devices, and reduce data movement delays [7]. Burger et al. applied an FPGA-based embedded device to monitor users' electrocardiograms (ECGs) in a pervasive Internet of Things (IoT) system [8]. FPGAs have also performed well in other areas such as agricultural robots [9], UAVs [10], and so forth.

In traditional FPGA-based edge device applications, DNNs are pre-trained in the cloud before being deployed in FPGAs, which is not efficient for domain adaption. When the environments, tasks, or users change, data needs to be collected from the edge FPGAs and transmitted to the cloud. Then, the cloud retrains a new model, transmitting the model back to the edge devices. The whole process is inefficient and time-consuming. Therefore, it is often desirable for edge FPGAs to continuously and locally learn from new data. Such on-device learning can directly improve model accuracy and adapt to new environments. Currently, several algorithms have been proposed to enable edge devices to achieve domain adaption locally. For example, a MobileDA framework [11] has been developed to allow a novel teacher network trained in the server to distill the knowledge for a student network running in the edge device, and the algorithm was employed on an embedded GPU and NVIDIA Jetson TX2. A transductive transfer learning model HDCNN [12] has been proposed to allow adaptation without requiring collecting large volumes of labeled training data in the target domain, and the algorithm was tested on 1080 Ti GPU. To implement these complex and fantastic software-level algorithms on FPGA-based edge devices, an FPGA-based training accelerator is indispensable. However, traditional FPGA-based edge device applications lack such hardware-level designs for training operations, which prevents FPGA-based devices from applying these algorithms directly.

Furthermore, directly training **Convolutional Neural Network (CNN)** models on local FPGAs can facilitate personalization. For example, in some medical applications such as home monitoring [13], long-term ECG monitoring [14], and so forth, the distinction of different users' physical conditions will impact data distribution, so models need to be fine-tuned based on specific users. The system in [8] utilized cloud services to log a user's condition over time and continuously improve the system's performance. It would be more effective if models could be directly updated on the FPGA device. Besides, learning at the edge can provide better privacy since users do not need to upload data into the central cloud [15].

However, it has been challenging to implement on-device training on FPGAs. Previous works mainly focused on implementing CNN inference on FPGAs. For example, Zhang et al. [16] exploited various optimization techniques including loop unrolling, loop tiling, and loop transformation on the FPGA accelerator and proposed a roofline model to quantitatively analyze its computing throughput and required memory bandwidth. Various designs [17, 18] have been proposed to map well-trained neural networks on FPGAs for inference with high throughput and low latency. Compared with CNN inference, it is more complex to efficiently implement CNN training on FPGAs in terms of the following aspects. First, the inference process only includes forward propagation (FP), whereas the training process includes FP, backward propagation (BP), and weight update (WU), which leads to a 3X computation operation count with more types of operations [19]. Second, the large volume of activation data in FP needs to be used in BP and WU, and the loss data generated in BP is also required in WU. Such data dependency across multiple layers makes it difficult for on-board memory management and data reusing in dynamic

random access memory (DRAM) in an end-to-end training system [20]. Third, since FP, BP, and WU have different memory access patterns, simply using the memory optimizations of FP in the whole training process leads to low memory access efficiency in BP and WU. Because of the above-mentioned challenges, CNN training on FPGAs has not been comprehensively investigated.

Recently, several FPGA-based architectures have been designed to accelerate training on large-scale FPGAs. F-CNN [21] first performed FP and BP on a Maxeler MPC-X dataflow FPGA node but WU on CPU. Designs such as [22, 23] aimed to further reduce the CNN training latency and improve throughput. However, these works mainly focused on cloud-level devices with abundant resources. A straightforward training implementation on edge FPGAs is still challenging.

To tackle the challenges in implementing on-device training on edge-level FPGAs, we propose EF-Train, a new efficient FPGA-based training accelerator with a unified channel-level parallelismbased convolution kernel to handle the computation complexity. The unified kernel means that it processes convolution operations of FP, BP, and WU in the training phase utilizing the same computation resources on the FPGA chip. The channel-level parallelism means that the kernel allocates these computation resources to process multiple channels of feature maps in parallel. We also propose a data reshaping approach to solve the communication bottleneck in realistic endto-end training processes. The overview of the design framework is shown in Figure 1. The data reshaping approach is a compile-time optimization that achieves intra-tile and inter-tile memory access continuity and weight reuse in mini-batch training. The proposed design can be implemented on resource-limited FPGAs without sacrificing precision. Neural networks can be trained on both small batches and large batches. Since training and inference are conducted separately in realistic applications, while FPGAs are configurable to implement different designs on the same hardware platform for different applications, our design can be applied to those well-developed FPGA-based inference devices. In a relatively long life cycle of the inference phase, the original design can guarantee high throughput and low latency. If users or the environments change, the device can be switched to implement our design immediately to learn from local data for online adaptation or personalization rather than transmitting data to the cloud centers and waiting for the cloud centers to transmit the well-trained model back to the device. Our main contributions are as follows:

- We propose EF-Train, an efficient FPGA-based CNN training accelerator with a unified convolution kernel to process FP, BP, and WU with full precision. The accelerator exploits channel-level parallelism to achieve high computation utilization for both small and large batch sizes. Our accelerator supports end-to-end CNN training with convolutional (Conv) layers, fully connected (FC) layers, batch normalization (BN) layers, rectified linear unit (ReLU) layers, and pooling layers (Section 3).
- We propose a data reshaping approach to solve the off-chip communication bottleneck. The features and weights are stored in off-chip memory with intra-tile continuous memory allocation to remove discontinuous memory accesses within a tile. We also reduce inter-tile discontinuous memory accesses by scheduling loop orders between tiles. We further exploit weight reuse among multiple images in a mini-batch to improve communication efficiency when the batch size is larger than one (Section 4).
- We build a performance and resource model for the proposed accelerator. Based on the model, a computation and memory resource scheduling tool is established to determine design parameters for different FPGA devices and different neural networks (Section 5).
- We deploy the training process on PYNQ-Z1 and ZCU102 for various CNNs on both Cifar-10
 and ImageNet datasets. Experimental results show that our design can achieve 46.99 GFLOPS
 and 6.09 GFLOPS/W in terms of throughput and energy efficiency, respectively (Section 6).

49:4 Y. Tang et al.

Challenges for On-device Training on Edge-level FPGAs Computation Complexity: Computation Rettlement

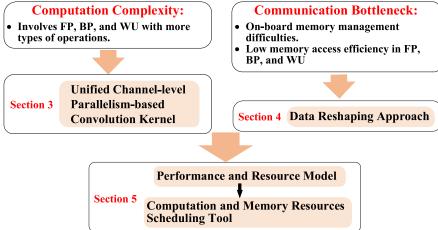


Fig. 1. Overview of our design framework.

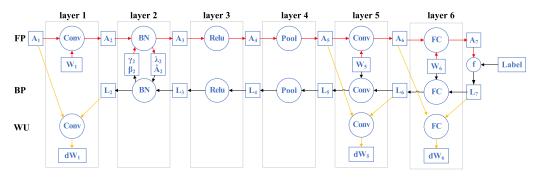


Fig. 2. CNN training process.

2 BACKGROUND AND MOTIVATIONS

2.1 CNN Training

The training process of a five-layer CNN is shown in Figure 2, including the FP process (red arrows), the BP process (black arrows), and the WU process (yellow arrows). The network includes two Conv layers, one FC layer, one BN layer, one ReLU layer, and one pooling layer, which are practical and can make up most neural networks in real-world scenarios.

In the FP process, activation is propagated layer by layer. In a Conv layer such as layer 1, the input activation A_1 conducts **multiply-accumulate (MAC)** operations with the weights W_1 . A BN layer is always followed by a Conv layer. In layer 2, the inputs of the BN layer include the input activation A_2 and learnable parameters γ_2 and β_2 . The immediate outputs include λ_2 , \hat{A}_2 , and output activation A_3 . A_3 then goes through the ReLU and pooling layers. Finally, the FC layer provides classification results for the input image.

In the BP process, the loss will be calculated and propagated back to the first layer. The loss of the last layer is calculated by the loss function f. This article adopts the most commonly used cross-entropy loss function. The **stochastic gradient descent (SGD)** is applied in CNN training. In a Conv layer, layer 5 for example, the loss L_6 needs to be padded first to ensure the size of the

convolution results L_5 is the same as the size of A_5 . The tensors for weights W_5 are transposed on dimensions for output channels and input channels. The original kernel tensors need to be flipped. Then, L_5 is calculated by the convolution operation between the transposed weights and L_6 . In the maximum pooling layer, layer 4, A_4 is compared with A_5 to determine which element on L_4 should obtain the value from the corresponding position on L_5 . If layer 4 is an average pooling layer, the values for each element of each patch in L_5 will be directly accumulated and propagated to L_4 . In the ReLU layer, layer 3, an element of L_3 will return zero if the value in the same position of A_3 is less than zero. Otherwise, the value of the corresponding position in L_4 will be propagated back. In the BN layer, layer 2, γ_2 and β_2 are updated according to the value of λ_2 , \hat{A}_2 , and L_3 . Then the loss is propagated back to L_2 .

The gradients of weights in layer i is calculated after the loss of layer i+1 is propagated. In layer 5, dW_5 is calculated by conducting MAC operations for A_5 and L_6 . The gradients are accumulated inside a batch. In a mini-batch, after the above-mentioned operations are conducted for all images, W_5 will be updated by subtracting $dW_5 \times learning \ rate$.

2.2 Related Works

DRAM Access Issues for Current FPGA-based Inference Accelerators: Currently, FPGAs have been widely adopted in edge domains thanks to the well-developed FPGA-based inference accelerators [16, 17]. Among the inference accelerators, many works [16, 24, 25] mainly focused on selecting optimal design parameters to improve the acceleration performance for individual Conv layers. Optimizing techniques such as loop tiling and loop unrolling are adopted by these works. Although the proposed algorithms achieved higher performance efficiency on a given Conv layer, the proposed designs only presented isolated accelerators without completing end-to-end inference where all layers of a neural network are tested continuously. In an end-to-end system, the layers' intermediate results are usually transferred between on-chip buffer and off-chip DRAM due to the limited on-chip storage size, so the impact of off-chip memory accesses should be considered in realistic scenarios. For most edge-level FPGAs, direct memory access (DMA) is a commonly used effective data swapping way for continuous address data access. However, in current FPGA-based DNN deployments, when the on-chip memory cannot hold all the features and weights of a Conv layer, data need to be fetched and processed in tiles based on the computation pattern. Such tiling schemes break the continuity of data addresses in DRAM and thus reduce the DMA transmission efficiency. The detailed analysis will be further discussed in Section 4. This discontinuity can degrade the DMA transferring speed from about 8GB/s to around 1GB/s [26]. The optimal algorithms proposed in the above-mentioned accelerators are based on the assumption that data are well pre-allocated between adjacent layers, so tiles can be loaded from and stored back to the DRAM continuously. However, in actual end-to-end systems, such allocation overhead is extremely large compared to the acceleration time.

Solutions for the DRAM Access Issues in the Inference Phase: Issues related to DRAM memory access have been addressed in recent works. For example, ROMANet [27] proposed a design space exploration (DSE) by searching for the appropriate data partitioning and scheduling for each layer of a network to reduce the number of memory accesses. DRMap [28] proposed a generic DRAM mapping policy and a DSE to reduce the DRAM access latency and energy. These two works were implemented on Tensor Processing Units (TPUs). [29] defined a multi-bank on-chip memory management (MOMM) problem to minimize the DRAM access overhead in the processing of CNNs on a neural processing unit (NPU) with a multi-bank on-chip memory. However, since FPGAs have different hardware architectures with TPUs or NPUs, their optimizing algorithms cannot be directly applied to FPGA-based designs. For example, in the TPU-based

49:6 Y. Tang et al.

designs [27, 28], the architecture of the on-chip accelerator is already fixed, with fixed MAC arrays and fixed on-chip buffers for input features, output features, and weights, respectively. In FPGA-based designs, only the total number of **digital signal processors (DSPs)** and on-chip memory sizes are given, and the allocation and connection of MAC arrays and individual buffers are configured by the designer. Therefore, optimizations on FPGA-based designs should not only reduce DRAM access latency or frequency based on the off-chip DRAM access policy but also be comparable to support the on-chip acceleration designs.

A few FPGA-based works reorganized the DRAM layout to relieve the memory access discontinuity and validated the approaches on realistic end-to-end tests. For example, [26] compared three different layout schemes of input features in the inference phase and finally found that the channel-major scheme where the input features are fetched and stored along the input channel direction first could improve access continuity and reduce data duplication. Caffeine [30] combined both on-chip and off-chip data reorganizations for the convolutional matrix-multiplication representation to maximize the underlying memory bandwidth utilization. FlexCNNe [31] further optimized data layout optimizations on the concatenation layers.

However, all these works [26]–[31] are based on the computation and memory access pattern in the inference phase, which only has FP. The training phase involves FP, BP, and WU, where their data access pattern for output features, input features, and weights are different. Therefore, the above-mentioned approaches cannot be directly applied in CNN training, and a new optimized design considering FP, BP, and WU together is required.

FPGA-based Training Accelerators: As mentioned in Section 1, CNN training on FPGAs has not been comprehensively investigated. The training process is much more complicated than the inference process, so it is sub-optimal to directly adopt the frameworks of inference accelerators for training.

Due to the computation complexity and communication bottleneck, currently, only a few works aim to achieve efficient FPGA-based training. With FPGA clusters, FPDeep explored layer-level parallelism for training a CNN model in a fine-grained pipeline [32], which has superior scalability to a large number of FPGAs. However, such larger clusters are not suitable to be adopted on edge-level applications. For training on a single FPGA, an automatic compiler for a training accelerator on Stratix 10 was developed in the precision of 16-bit fixed point [22]. DarkFPGA adopted batch-level parallelism using 8-bit integers for training a VGG-like network on the Maxeler MAX5 platform [23]. It achieved high throughput when the batch size was large. A sparse CNN training accelerator was designed on VCU1525. The accelerator was implemented on a pre-trained CNN model with 85% of parameters pruned [33]. However, these existing works mainly focused on cloud-level devices with abundant computation and memory resources.

Besides, even with cloud-level resources, reduced precision and pruning approaches have also been utilized to decrease computation intensity and communication bottleneck. Although quantization adopted in prior training accelerators [23, 34] led to remarkable benefits in terms of resource usage and power consumption, these works have not provided any evidence that such quantization techniques can maintain high accuracy on a large dataset (e.g., ImageNet) with dense neural networks. Currently, training with full precision is still preferred in most realistic applications, and its high computation and memory overhead should be faced directly. However, none of the abovementioned state-of-the-art training accelerators targeted resource-limited edge FPGAs with full precision, which is more challenging to implement in end-to-end CNN training but is more practical in real-world scenarios. Therefore, an optimized design is necessary to implement on-device training on resource-limited FPGAs without sacrificing precision.

Implementation of BN Layers: Apart from the computation-intensive Conv layer, the BN layer is also a key component and is essential for the training process. In inference, a BN layer can

be folded into the adjacent CONV layer, since it just performs a simple linear transformation [35]. However, the batch normalization process in the training phase is much more complex. It needs to calculate the expected value and variance of the data in the whole mini-batch, which involves lots of on-chip and off-chip data transmission. Lu et al. [35] optimized the computation flow of BN layers during FP and BP and implemented BN layers in their CNN training accelerator. Unlike [35], which adopts the 8-bit fixed point in Conv layers and FP16 in BN layers, our work supports BN layers with full precision, which brings more challenges for computation and transmission requirements.

2.3 Motivations of the Proposed Design

To implement on-device training on resource-limited FPGAs, we need to solve the computation complexity and communication bottleneck illustrated in Figure 1.

In the training phase, the FP, BP, and WU processes are conducted iteratively and need to be completed on the same accelerator. Edge FPGAs have limited computational resources. Using separate kernels for FP, BP, and WU leads to resource underutilization and low energy efficiency. Therefore, to efficiently process the complex computation for FP, BP, and WU, we need to design a training accelerator that can handle the three processes in a unified convolution kernel and can achieve a high parallelism degree considering the flexibility of DNN architectures. For a Conv layer, there are three levels of parallelism that are adopted in FPGA-based accelerators: batch-level parallelism, feature-map-level parallelism, and channel-level parallelism. Figure 3(a) illustrates the process of batch-level parallelism, where Tb nominates the number of **output feature maps (OFMs)** of images in a mini-batch that are processed in parallel. Figure 3(b) shows the process of feature-maplevel parallelism, where $Tf \times Tf$ features of OFMs are processed in parallel. Figure 3(c) shows the process of channel-level parallelism, where *Tm* nominates the number of output channels of OFMs that are processed in parallel, and *Tn* nominates the number of input channels of **input feature** maps (IFMs) that are processed in parallel. The degree of parallelism depends on the amount of utilized computation units on the hardware. Table 1 shows the comparisons of the three levels of parallelism. Considering a Conv layer with B images in a mini-batch, it is assumed that the number of input channels is N, the number of output channels is M, the size of an OFM is $R \times C$, and the size of a weight kernel is $K \times K$. Tmops = $B \times M \times N \times R \times C \times K \times K$ MAC operations are required to process such a layer. For the batch-level parallelism, it takes $\lceil \frac{B}{Th} \rceil \times M \times N \times R \times C \times K \times K$ cycles to complete the Conv layer. Such parallelism can achieve high throughput when the batch size is large, and the size of the feature map and the number of channels have little impact on the performance. For example, in previous works, DarkFPGA [23] built its accelerator with batch-level parallelism and achieved high throughput when the batch size was 128. However, when the batch size is small or even one (online learning), most computation units will remain idle. For example, when B < Tb, completing the Conv layer costs Tmops cycles, and $\frac{Tb-B}{Tb}$ of computation resources remain idle. It leads to a low parallelism degree, and such under-utilization of resources makes the performance sub-optimal. For the feature-map-level parallelism, which has been adopted by works like [22], it takes $B \times M \times N \times \lceil \frac{R}{Tf} \rceil \times \lceil \frac{C}{Tf} \rceil \times K \times K$ cycles to finish a Conv layer. The batch size and the number of channels have little impact on such parallelism. The parallelism will benefit from layers with large feature map size but has under-utilization for layers with small feature map size. For example, when R < Tf and C < Tf, completing the Conv layer costs Tmops cycles, and $\frac{Tf-R}{Tf} \times \frac{Tf-C}{Tf}$ of computation resources will remain idle. However, in CNN training, the size of a feature map may vary from large size (like 224 × 224 for the input image of the ImageNet) to 1 × 1 for the FC layer. The feature-map-level parallelism will be inefficient to process the layers with a small feature map. For channel-level parallelism, it takes $B \times \lceil \frac{M}{Tm} \rceil \times \lceil \frac{N}{Tn} \rceil \times R \times C \times K \times K$ cycles to

49:8 Y. Tang et al.

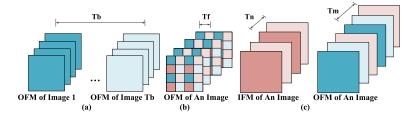


Fig. 3. Three levels of parallelism. (a) Batch-level parallelism. (b) Feature-map-level parallelism. (c) Channel-level parallelism.

Parallelism	Batch-level Parallelism	Feature-map-level Parallelism	Channel-level Parallelism
Large Batch Size	advantaged	little impact	little impact
Small Batch Size	disadvantaged	little impact	little impact
Large Feature Map Size	little impact	advantaged	little impact
Small Feature Map Size	little impact	disadvantaged	little impact
Large Channel Number	little impact	little impact	advantaged
Small Channel Number	little impact	little impact	disadvantaged

Table 1. Comparisons of the Three Levels of Parallelism

complete the Conv layer. It acquires a high parallelism degree with a large channel number, and the batch size and feature map size have little impact on it. When the channel number is small, for example, when N < Tn, completing the Conv layer costs $B \times \lceil \frac{M}{Tm} \rceil \times N \times R \times C \times K \times K$ cycles, and $\frac{Tn-N}{Tn}$ of computation resources will remain idle. However, for most neural networks, only the first layer has a small input channel number (N=3). For other layers, the channel size (e.g., 32, 64, etc.) is usually larger than the maximum degree of parallelism that an edge FPGA can achieve. Therefore, channel-level parallelism is widely adopted by FPGA-based inference accelerators [16, 24]. Generally speaking, the channel-level parallelism can achieve a constantly high degree of parallelism across multiple layers, so it is adopted in the proposed design as shown in Figure 1. The proposed accelerator with a channel-level parallelism-based convolution kernel to process FP, BP, and WU will be introduced in detail in Section 3.

Furthermore, the communication bottleneck is also challenging for edge-level FPGAs in end-to-end training. As illustrated in Figure 2, the activation data in FP needs to be used in BP and WU, and the loss data generated in BP is also required in WU. The length and heterogeneity of the data dependency paths in different layers make external memory accesses inevitable [20]. Previous training accelerators attempted to avoid external memory accesses. For example, FPDeep [32] scaled CNN computations to larger clusters so that only on-chip memory is needed for the CONV layers. However, such larger clusters cannot be used on edge devices. [36] implemented LeNet-10 on an FPGA and stored the inputs and outputs of one layer on the chip. Such design can only support small networks, but for many larger networks (e.g., Vgg-16, AlexNet, etc.), the on-chip memory of an edge FPGA is not big enough to hold weights or features in every Conv layer. Therefore, several works [23, 33, 34] applied quantization or pruning to reduce off-chip memory access. However, unlike inference, where compressed networks cause little decrease in accuracy [37], these training works have not proved that their compression techniques can maintain high accuracy on large datasets with dense networks. To guarantee accuracy, it is necessary to implement CNN training

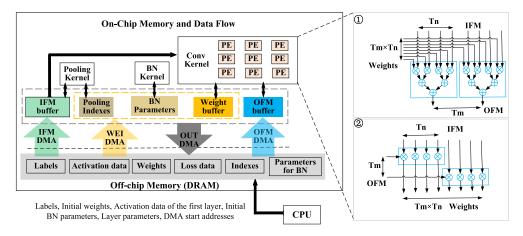


Fig. 4. Accelerator architecture.

with full precision. Our goal is to design a general accelerator supporting end-to-end training with both dense and small networks without sacrificing precision, so it is necessary to appropriately manage external memory access and allocate on-chip buffers. As mentioned in Section 2.2, the tiling schemes involved in on-chip accelerator design break the continuity of data addresses in DRAM and thus reduce the DMA transmission efficiency. Therefore, it is necessary to improve the address continuity of data to improve the efficiency of data swapping considering the complex data patterns in FP, BP, and WU altogether. To solve this communication issue, as shown in Figure 1, a data reshaping approach is proposed and will be introduced in detail in Section 4.

3 FPGA-BASED CNN TRAINING ACCELERATOR

In this section, we propose an FPGA-based accelerator exploiting channel-level parallelism to deal with the training process. A unified convolution kernel is designed to process FP, BP, and WU with full precision.

3.1 The Architecture of the Training Accelerator

The proposed training accelerator is shown in Figure 4. We implement our accelerator on an end-to-end training system. At first, the CPU transmits labels, initial weights, the activation data of the first layer, layer parameters, initial parameters for BN layers, and the DMA start addresses for each layer to the off-chip DRAM. The layer parameters include computation type (e.g., Conv, ReLU, BN, or pooling) and the shape information. The DMA start addresses are calculated off-line according to the off-chip memory layout based on our data reshaping approach mentioned in Section 4. Our accelerator executes computation-intensive kernels based on data dependencies within a CNN model, while the entropy loss function is computed on the off-chip ARM core.

As illustrated in Figure 4, the off-chip memory (DRAM) stores data for activation, loss, weights, labels, indexes for pooling, and parameters for BN. Data are transmitted through the DMA AXI-stream bus to on-chip memory for computation. There are four DMA stream channels: IFM DMA, OFM DMA, WEI DMA, and OUT DMA. These four channels are independent and can work in parallel. On the FPGA chip, a unified Conv kernel is designed to process FP, BP, and WU with the same computation resources, i.e., the DSPs. The Conv kernel is composed of multiple **processing elements (PEs)** to implement MAC operations. The adder tree structure is adopted for the proposed kernel since it is flexible to support different computation patterns for FP, BP, and WU.

49:10 Y. Tang et al.

Table 2. Definitions of Symbols

Notation	Description
i,j	Index of a Conv layer
В	Batch size
N^i	Number of the input channels of the <i>i</i> th Conv layer
M^i	Number of the output channels of the <i>i</i> th Conv layer
R^i	Number of the rows of the OFMs for the <i>i</i> th Conv layer
C^{i}	Number of the columns of the OFMs for the <i>i</i> th Conv layer
K^{i}	Size of the weights kernel for the <i>i</i> th Conv layer
S^i	Stride for the <i>i</i> th Conv layer
$A_i[b, m, r, c]$	Activation for the <i>i</i> th Conv layer in FP
$L_i[b, n, r, c]$	Loss for the <i>i</i> th Conv layer in BP
$W_i[n, m, kr, kc]$	Weights for the <i>i</i> th Conv layer
$dW_i[n, m, kr, kc]$	Weight gradients for the <i>i</i> th Conv layer in WU
Tm	Number of the output channels in a tile of output features
	in each Conv layer
Tn	Number of the input channels in a tile of input features in
	each Conv layer
Tr^i	Number of the rows in a tile of output features in the <i>i</i> th Conv layer
Tc^i	Number of the columns in a tile of output features in the <i>i</i> th Conv layer
M^i _on	Number of the output channels of the weights stored on-chip in
	the <i>i</i> th Conv layer
R ^j _in	Number of the rows of the IFMs for the <i>j</i> th Conv layer
C^j _in	Number of the columns of the IFMs for the <i>j</i> th Conv layer
Tr^{j} _in	Number of the rows in a tile of input features in the <i>j</i> th Conv layer
Tc ^j _in	Number of the columns in a tile of input features in the <i>j</i> th Conv layer

The connection of each multiplier and adder for FP and BP is shown in ① of Figure 4, while the connection of each multiplier and adder for WU is shown in ② of Figure 4. The Pooling Kernel focuses on the pooling operation. The BN Kernel achieves batch normalization and updates BN parameters during FP and BP. ReLU is always followed by a Conv or BN layer. The accelerator compares the output features with 0 when storing output features back to the DRAM in Conv or BN layers, so ReLU does not need a unique functional unit. Five types of on-chip **block RAMs** (**BRAMs**) are used to buffer IFMs, weights or weights gradients, OFMs, pooling indexes, and BN parameters. We adopt double-buffer designs so that data transmission and computation can be conducted in parallel.

3.2 The Forward and Backward Propagation of a Convolutional Layer

Our accelerator adopts channel-level parallelism, loop unrolling, and loop tiling. The symbols are defined in Table 2. In channel-level parallelism, Tn and Tm are determined by available computation resources (i.e., DSPs) on the FPGA chip and are fixed for all Conv layers. The degree of parallelism is determined by $Tn \times Tm$.

Our accelerator achieves SGD in CNN training. The forward and backward propagation processes of a Conv layer processing the bth image in a mini-batch can be formulated in Equations (1) and (2), where W_i' is the transposed and flipped tensor of W_i . As illustrated in ① of Figure 4, in FP and BP, the Conv Kernel conducts MAC operations for weights and input features from activation

or loss. The IFM buffer stores a tile of activation or loss transmitted via the IFM DMA channel, and the Weight buffer stores weights transmitted via the WEI DMA channel. The OFM buffer stores a tile of MAC outcomes. Computation results are transmitted to the DRAM via the OUT DMA channel. If a Conv layer is followed by a ReLU layer, for FP, the data in the OFM buffer will be compared with 0 before entering into the OUT DMA channel. For BP, the activation of the previous layer will be transmitted via the OFM DMA channel, and the accelerator decides which value should be propagated according to Equation (3).

$$A_{i+1}[b, m, r, c] = \sum_{n=1}^{N^i} \sum_{kr=1}^{K^i} \sum_{kc=1}^{K^i} A_i[b, n, S^i \times r + kr, S^i \times c + kc] \times W_i[m, n, kr, kc]$$
 (1)

$$L_{i}[b, n, r, c] = \sum_{m=1}^{M^{i}} \sum_{kr'=1}^{K^{i}} \sum_{kc'=1}^{K^{i}} L_{i+1}[b, m, S^{i} \times r + kr', S^{i} \times c + kc'] \times W'_{i}[n, m, kr', kc']$$
(2)

$$L_{i}[b, m, r, c] = \begin{cases} L_{i+1}[b, m, r, c], & A_{i}[b, m, r, c] > 0, \\ 0, & others \end{cases}$$
 (3)

3.3 The Weight Update of a Convolutional Layer

The gradients of weights can be calculated in Equation (4). The generated hardware implementation of the PE architecture processing WU operations in the Conv kernel is shown in ② of Figure 4. During the WU, the Conv Kernel conducts MAC operations for the activation data transmitted via the IFM DMA channel and the loss data transmitted via the OFM DMA channel. The gradients are stored in the Weight buffer. Once the Conv Kernel completes the computation for the last image in a mini-batch, the original weights are transmitted via the WEI DMA channel. Then, weights are updated by deducting the product of the gradients and learning rate. New weights are sent back to DRAM via the OUT DMA channel.

$$dW_{i}[m, n, kr, kc] = \sum_{b=1}^{B} \sum_{r=1}^{R^{i}} \sum_{c=1}^{C^{i}} L_{i+1}[b, m, r, c] \times A_{i}[b, n, S^{i} \times r + kr, S^{i} \times c + kc]$$
(4)

3.4 The Forward and Backward Propagation of a Pooling Layer

In the FP process of a pooling layer, the activation is transmitted via the IFM DMA channel and stored in the IFM buffer. In the maximum pooling, the Pooling Kernel compares adjacent pixels, transfers the results back to DRAM via the OUT DMA channel, and records the index for the maximal pixel into the Pooling Indexes buffer. The index of a pixel is a 2-bit integer. For average pooling, the kernel just calculates the average value of a patch of features. In the BP process of maximum pooling, the indexes are loaded back via the WEI DMA channel, and loss from the previous layer is loaded via the IFM DMA channel. The Pooling Kernel compares the indexes and stores the propagated value into the IFM buffer. The BP process of the maximum pooling is formulated in Equation (5). For average pooling, the loss values of a patch are directly accumulated. After a tile of data is processed, the calculated loss is sent back via the OUT DMA channel.

$$L_{i}[b, m, S^{i} \times r + kr, S^{i} \times c + kc] = \begin{cases} L_{i+1}[b, m, r, c], A_{i+1}[b, m, r, c] = A_{i}[b, m, S^{i} \times r + kr, S^{i} \times c + kc], \\ 0, others \end{cases}$$
(5)

3.5 The Forward Propagation of a BN Layer

Our BN kernel is based on the computation flow in [35]. However, unlike the prior work, which utilizes half-precision, we adopt full precision, which brings more computation and transmission challenges. The BN parameters in Figure 4 include learnable parameters $\gamma_i[m]$ and $\beta_i[m]$ and immediate parameters $\lambda_i[m]$ and $\hat{A}_i[b, m, r, c]$, where m is the index of the channel. The $\gamma_i[m]$ and $\beta_i[m]$

49:12 Y. Tang et al.

are used to generate the immediate parameters and the output activation $A_{i+1}[b,m,r,c]$ during FP. During BP, the immediate parameters and the loss propagated from the next layer $L_{i+1}[b,m,r,c]$ are used to update the learn parameters and propagate the loss $L_i[b,m,r,c]$ back. Since the size of $\gamma_i[m]$, $\beta_i[m]$, and $\lambda_i[m]$ is M (the number of the output channels), the on-chip BRAMs are large enough to hold these data in a BN layer. Therefore, we use the BN Parameters buffer to store these parameters as well as the expected value and variance. The $\hat{A}_i[b,m,r,c]$ is transmitted to DRAM together with $A_{i+1}[b,m,r,c]$.

In FP, the BN Kernel first loads $\gamma_i[m]$ and $\beta_i[m]$ from DRAM to the BN Parameters buffer via the WEI DMA channel. Then it loads $A_i[b,m,r,c]$ via the IFM DMA channel and calculates the expected value $E(X)_i[m]$ and variance $V(X)_i[m]$ according to Equations (6) to (8). To avoid disarranging the DRAM data layout for adjacent Conv layers, we load data tile by tile using the same data format as that in Conv layers. The expected value and variance are calculated after the entire data of a mini-batch is accessed. Then input activation is loaded from the beginning to calculate the immediate parameters according to Equations (9) and (10), where ϵ is a constant parameter. $\lambda_i[m]$ is stored in the BN Parameters buffer, while $\hat{A}_i[b,m,r,c]$ is transmitted to DRAM via the OUT channel in parallel with activation loading. Finally, the output activation is calculated according to Equation (11). The BN operation completes after $A_{i+1}[b,m,r,c]$, $\gamma_i[m]$, $\beta_i[m]$, and $\lambda_i[m]$ are stored to DRAM.

$$E(X)_{i}[m] = \frac{1}{B \times R^{i} \times C^{i}} \sum_{b=1}^{B} \sum_{r=1}^{R^{i}} \sum_{c=1}^{C^{i}} A_{i}[b, m, r, c]$$
(6)

$$E(X^{2})_{i}[m] = \frac{1}{B \times R^{i} \times C^{i}} \sum_{b=1}^{B} \sum_{r=1}^{R^{i}} \sum_{c=1}^{C^{i}} A_{i}^{2}[b, m, r, c]$$
(7)

$$V(X)_{i}[m] = E(X^{2})_{i}[m] - (E(X)_{i}[m])^{2}$$
(8)

$$\lambda_i[m] = \frac{1}{\sqrt{V(X)_i[m] + \epsilon}} \tag{9}$$

$$\hat{A}_{i}[b, m, r, c] = (A_{i}[b, m, r, c] - E(X)_{i}[m]) \times \lambda_{i}[m]$$
(10)

$$A_{i+1}[b, m, r, c] = \hat{A}_i[b, m, r, c] \times \gamma_i[m] + \beta_i[m]$$
(11)

3.6 The Backward Propagation of a BN Layer

In BP, $\hat{A}_i[b, m, r, c]$, $\lambda_i[m]$ and $L_{i+1}[b, m, r, c]$ are used to update the learnable parameters $\gamma_i[m]$ and $\beta_i[m]$, and $L_i[b, m, r, c]$ is propagated back. $\lambda_i[m]$, $\gamma_i[m]$, and $\beta_i[m]$ are first uploaded via the WEI channel and stored in the BN Parameters buffer. Then, the BN Kernel loads $\hat{A}_i[b, m, r, c]$ and $L_{i+1}[b, m, r, c]$ via the IFM and OFM channel, respectively, to calculate the gradients for $\gamma_i[m]$ and $\beta_i[m]$ according to Equations (12) and (13). The learnable parameters are updated by deducting the gradients, while $L_i[b, m, r, c]$ is calculated according to Equation (14).

$$d\gamma_i[m] = \sum_{h=1}^B \sum_{r=1}^{R^i} \sum_{c=1}^{C^i} L_{i+1}[b, m, r, c] \times \hat{A}_i[b, m, r, c]$$
 (12)

$$d\beta_i[m] = \sum_{b=1}^B \sum_{r=1}^{R^i} \sum_{c=1}^{C^i} L_{i+1}[b, m, r, c]$$
(13)

Off-chip data transmission	Off-chip data transmission
1. for (row=0; row <r; (col="0;" (ti="0;" (to="0;" 2.="" 3.="" 4.="" col+="Tc)" col<c;="" for="" row+="Tr)" th="" ti+="Tn)" ti<n;="" to+="Tm)" to<m;="" {="" {<=""><th>1. for (to=0; to<m; 2.<="" th="" to+="Tm)" {=""></m;></th></r;>	1. for (to=0; to <m; 2.<="" th="" to+="Tm)" {=""></m;>
On-chip computation	On-chip computation
$ \begin{array}{ll} \textbf{5. for}(i=0; i< K; i++) \{\\ 6. $	5. for(trr=0; trr <tr; 10.<="" 6.="" 7.="" 8.="" 9.="" for(i="0;" for(j="0;" for(tcc="0;" for(too="0;" i++)="" i<k;="" j++)="" j<k;="" tcc++)="" tcc<tc;="" td="" too++)="" too<tm;="" trr++)="" {=""></tr;>
(a)	(b)

Fig. 5. Pseudo-code of a tiled convolution layer. (a) Pseudo-code for FP and BP. (b) Pseudo-code for WU.

$$L_{i}[b,m,r,c] = \gamma_{i}[m] \times \lambda_{i}[m] \times \left(L_{i+1}[b,m,r,c] - \frac{d\beta_{i}[m]}{B \times R^{i} \times C^{i}} - \hat{A}_{i}[b,m,r,c] \times \frac{d\gamma_{i}[m]}{B \times R^{i} \times C^{i}}\right)$$
(14)

4 DATA RESHAPING APPROACH

In this section, we propose a data reshaping approach to solve the communication bottleneck between the on-chip buffer and off-chip memory in realistic end-to-end training processes. We first analyze the discontinuous memory access for the isolate accelerator with the unified channel-level parallelism-based Conv kernel proposed in Section 3. Then we introduce our data reshaping approach, which involves three aspects. We first achieve intra-tile continuous memory allocation by reorganizing the DRAM layouts for input features, output features, and weights. Then, we re-schedule the loop order to achieve inter-tile continuous memory allocation. These two parts are optimized together. Finally, considering that the training process involves convolution operations among a mini-batch, we propose and apply a weight reuse strategy based on the proposed data layout.

4.1 Analysis on Discontinuous Memory Access

The pseudo-code of a tiled convolution layer is shown in Figure 5. The pseudo-code in Figure 5(a) is applied in FP and BP, following previous FPGA-based inference works [16, 30], while the pseudo-code in Figure 5(b) is applied in WU, based on the accelerator design proposed in Section 3. As discussed in Section 2, the continuity of data significantly influences the DMA transmission efficiency. In this section, we analyze the data discontinuity when features are placed with the BCHW pattern and the BHWC pattern, where B represents batch, C represents channel, H represents height (row), and W represents width (column).

Features Are Placed in the BCHW Pattern: Figure 6(a) shows the data layout of $M^i \times R^i \times C^i$ output features stored in DRAM for the ith layer. The output features are placed with the BCHW pattern commonly used in CNN accelerating CPU and GPU platforms [38, 39]. While OpenVINO [39] is primarily for CPUs, it would also work for CPU, GPU, and FPGA platforms. In this layout, a cube represents an element of the features, and the indexes represent the orders of the elements stored in DRAM. In FPGA-based DNN deployments, data are fetched and processed in tiles. As shown in Figure 6, the size of a tile is $Tm \times Tr^i \times Tc^i$ for output features.

The output features of layer i are also the input features for its next layer j. As shown in Figure 7(a), the size of the input features in layer j is $N^j \times R^j in \times C^j in$. For input features, the size of a tile is $Tn \times Tr^j in \times Tc^j$ in.

49:14 Y. Tang et al.

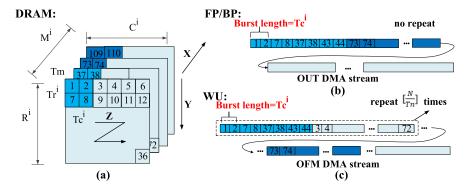


Fig. 6. Data layout of output features before reshaping. (a) Data stored in DRAM. (b) Data transmitted via the OUT DMA channel in FP/BP. (c) Data transmitted via the OFM DMA channel in WU.

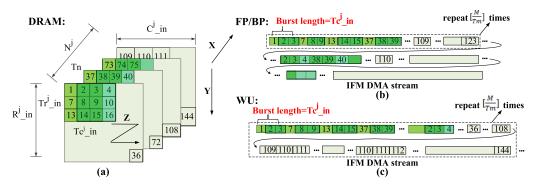


Fig. 7. Data layout of input features before reshaping. (a) Data stored in DRAM. (b) Data transmitted via the IFM DMA channel in FP/BP. (c) Data transmitted via the IFM DMA channel in WU.

In each DMA stream, the AXI-stream bus allows a pipeline data stream when the data addresses are continuous. Burst length represents the number of data with continuous addresses in the data stream. When a discontinuity happens, the DMA needs to be restarted. Therefore, our goal is to avoid discontinuity, i.e., elongate the burst length for different transmission patterns.

During FP, the Conv Kernel conducts MAC operations with weights and input features, and then it generates output features. The output features are transmitted to the DRAM via the OUT DMA channel, which is shown in Figure 6(b). For the next layer, the input features are fetched from DRAM to the FPGA chip via the IFM DMA stream, which is shown in Figure 7(b). As illustrated in Figure 5(a), in a Conv layer, the OFM buffer is reused to store and accumulate the immediate convolution results between each tile of input features and each tile of weights. The first tile of output features is generated by accumulating the convolution results when the input feature tiles move from the first input channel to the last channel. It corresponds to the movement in the **X** direction in Figure 7. Then, the next tiles of output features are generated in the **X** direction in Figure 6, so the data access pattern of input features (the dashed box in Figure 7(b)) repeats $\lceil \frac{M}{Tm} \rceil$ times. After the output feature tiles move from the first output channel to the last output channel, they begin to move in the **Z** direction, and the input feature tiles follow the **Z** direction as well. From Figures 6(b) and 7(b), the address of data is discontinuous both inside and outside of a tile. The burst length of the output features in the *i*th layer is Tc^i , and the burst length of the input features in the *j*th layer is Tc^i .

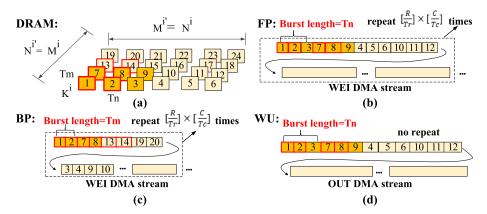


Fig. 8. Data layout of weights before reshaping. (a) Weights stored in DRAM. (b) Weights transmitted via the WEI DMA channel in FP. (c) Weights transmitted via the WEI DMA channel in BP. (d) Weights transmitted via the OUT DMA channel in WU.

In WU, the data access pattern inside a tile is the same as that in FP/BP. However, the intertile data access pattern is different from that for FP/BP. This is because, in WU, the Conv Kernel conducts MAC operations for input features (the activation data transmitted via the IFM DMA channel) and output features (the loss data transmitted via the OFM DMA channel) to calculate weight gradients. Therefore, as shown in Figure 5(b), the WEI buffer is reused to store and accumulate the immediate convolution result between each tile of input features and each tile of output features. The first tile of weight gradients is generated when the input and output feature tiles move from the first row and the first column to the last row and the last column. It corresponds to the movement in the Z direction in Figures 6 and 7. Then, the next tile of weight gradients is generated along the input channel direction. The tiles of input features move along the X direction, while the pattern of output feature tiles (the dashed box in Figure 6(c)) repeats $\lceil \frac{N}{T_n} \rceil$ times. After the gradients of weights are calculated from the first input channel to the last input channel, the next tiles are generated along the output channel direction. Thus, the pattern of input feature tiles (the dashed box in Figure 7(c)) repeats $\lceil \frac{M}{Tm} \rceil$ times, while the output feature tiles move along the **X** direction. As shown in Figure 7(c), the burst length for input features is Tc^{j} _in. As shown in Figure 6(c), the burst length for output features is Tc^{i} .

In CNN training, the data layout of weights is also more complex compared to the inference process. As illustrated in Figure 8(a), $M^i \times N^i \times K^i \times K^i$ weights of layer i are stored in DRAM. In FP, weights are fetched in the input channel first and then the output channel when the output features are generated along the \mathbf{X} direction in Figure 6. Then, the output features are generated along the \mathbf{Z} direction, while the weight access pattern (the dashed box in Figure 8(a)) repeats $\lceil \frac{R}{T_r} \rceil \times \lceil \frac{C}{T_c} \rceil$ times. WU shares the same intra-tile weight access pattern with FP, but it does not need to repeat during inter-tile data access. The burst lengths for FP and WU are both Tn. In BP, each $K^i \times K^i$ kernel needs to be flipped. Such reallocation can be processed on the FPGA chip. However, since the numbers of input channels and output channels are interchanged, the memory access pattern of a tile is also changed. The weight kernels are transposed between the input channel dimension and the output channel dimension. In Figure 8, the yellow cubes represent a tile of weights in FP and WU, and the cubes with the red box represent a tile of weights in BP. As illustrated in Figure 8(c), in BP, the number of output channels becomes $M^{i'} = N^i$, the number of input channels becomes $M^{i'} = M^i$, and the burst length is Tm.

49:16 Y. Tang et al.

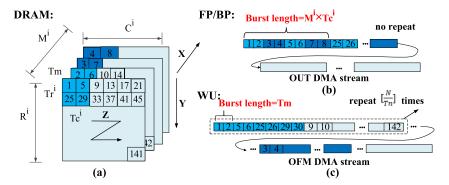


Fig. 9. Data layout of output features with the BHWC memory allocation and feature reuse. (a) Data stored in DRAM. (b) Data transmitted via the OUT DMA channel in FP/BP. (c) Data transmitted via the OFM DMA channel in WU.

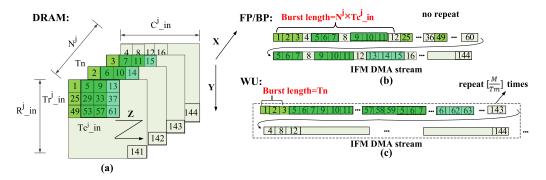


Fig. 10. Data layout of input features with the BHWC memory allocation and feature reuse. (a) Data stored in DRAM. (b) Data transmitted via the IFM DMA channel in FP/BP. (c) Data transmitted via the IFM DMA channel in WU.

Features Are Placed in the BHWC Pattern: As can be seen from Figures 6, 7, and 8, the tiled data breaks the data continuity of memory access in FP, BP, and WU for the isolate accelerator. In FPGA-based inference works, the BHWC pattern is also commonly used in end-to-end designs to optimize memory access [26, 30]. Figures 9(a) and 10(a) show the data layout of features placed in the BHWC pattern following previous inference-based works. According to the loop order in Figure 5(a), in FP and BP, tiles move in the channel dimension first and then move in the **Z** direction. Therefore, it is effective to fetch $\lceil \frac{N}{Tn} \rceil$ tiles of input features to the on-chip memory and reuse the data after $\lceil \frac{M}{Tm} \rceil$ tiles of output features are calculated. With such optimizations, the data discontinuity of features is alleviated in FP and BP. As shown in Figures 9(b) and 10(b), the burst length for output features is $M^i \times Tc^i$, and the burst length for input features is $N^j \times Tc^j$ _in. Besides, the FPGA accelerator does not need to repeatedly load the input feature tiles from the DRAM.

However, in WU, the Conv Kernel conducts MAC operations for input features and output features to calculate weight gradients, so input and output feature tiles should move in the \mathbf{Z} direction first to calculate the weight gradients of $Tm \times Tn$ weight kernels and then move in the channel direction, which is illustrated in Figure 5(b). Therefore, features cannot be continuously fetched to the on-chip buffer and reused as that in the inference phase unless the on-chip memory is large enough to hold all features of each layer. When the on-chip memory cannot hold all features of

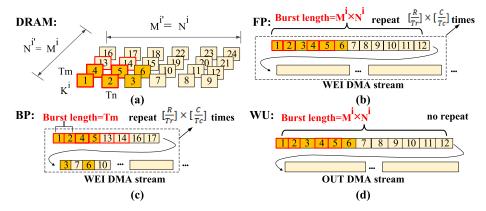


Fig. 11. Data layout of weights placed tile by tile based on inference-based data flow. (a) Weights stored in DRAM. (b) Weights transmitted via the WEI DMA channel in FP. (c) Weights transmitted via the WEI DMA channel in BP. (d) Weights transmitted via the OUT DMA channel in WU.

a Conv layer in resource-limited FPGAs, the burst length for output features is Tm, and the burst length for input features is Tn. The data layouts are shown in Figures 9(c) and 10(c).

In CNN inference, weights will not change in the whole process, and they are loaded in the same pattern for different layers. Therefore, in the inference phase, weights can be pre-allocated tile by tile to ensure continuous memory access. The pre-allocated data layout is illustrated in Figure 11(a). As shown in Figures 11(b) and 11(d), the burst length is $M^i \times N^i$ in FP and WU. However, as illustrated in Figure 11(c), the weight kernels are transposed between the input channel dimension and the output channel dimension, and the tiling scheme in BP breaks the memory access continuity. Since weights are updated after one iteration of FP, BP, and WU, it is impossible to pre-allocate them before each iteration. Therefore, data discontinuity is inevitable in BP. As shown in Figure 11, the burst length is Tm.

4.2 Optimizing Discontinuous Memory Access

To optimize the discontinuous memory access, our data reshaping approach includes the following steps. First, we achieve intra-tile continuous memory allocation for both features and weights by reorganizing the DRAM layouts for output features, input features, and weights, which are shown in Figures 12, 13, and 14, respectively. Then we schedule the loop order based on the pseudo-code in Figure 15 to achieve inter-tile continuous memory allocation. Finally, weights are reused among a mini-batch based on the proposed data layouts. In this section, we reorganize the data layouts and schedule the loop order together to achieve both intra-tile and inter-tile continuous memory address allocation.

Intra-tile Continuous Memory Allocation: Inspired by previous inference works [26, 30], employing a channel-last data layout can improve data continuity for the channel-level parallelism-based accelerator. However, as explained in Section 4.1, simply changing the data layout cannot optimize the memory access continuity in FP, BP, and WU together. The memory access patterns in the three processes need to be considered together. In CNN inference, the selection of Tm and Tn is flexible. However, to ensure data continuity of weight kernels in both FP and BP, we fix Tm = Tn in our training accelerator so that weights can be loaded tile by tile in both FP and BP.

Figure 12(a) shows the data layout of the output features in DRAM after data reshaping. The first Tm channels of OFMs are placed in the row-column-channel pattern. The next Tm channels of OFMs are followed with the same pattern. When applying loop tiling, we assign the tiling

49:18 Y. Tang et al.

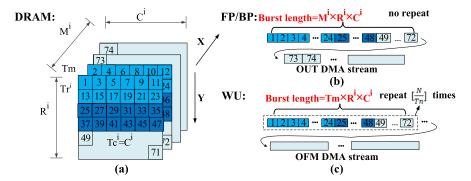


Fig. 12. Data layout of output features after reshaping. (a) Data stored in DRAM. (b) Data transmitted via the OUT DMA channel in FP/BP. (c) Data transmitted via the OFM DMA channel in WU.

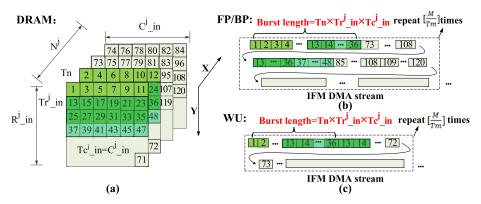


Fig. 13. Data layout of input features after reshaping. (a) Data stored in DRAM. (b) Data transmitted via the IFM DMA channel in FP/BP. (c) Data transmitted via the IFM DMA channel in WU.



Fig. 14. Data layout of weights after reshaping. (a) Weights stored in DRAM. (b) Weights transmitted via the WEI/OUT DMA channel in FP/WU. (c) Weights transmitted via the WEI DMA channel in BP.

parameter $Tc^i = C^i$ so that data are continuous inside a tile for both FP and WU. From Figures 12(b) and 12(c), the burst lengths of output features during FP, BP, and WU are larger than the size of a tile.

The selection of Tm = Tn, $Tc^i = C^i$, and $Tc^i_i = C^i_i$ also guarantees that features of different layers share similar data layouts and tiling schemes no matter whether they serve as output features or input features of a Conv layer in FP/BP/WU. Therefore, the intra-tile continuity of input features is also guaranteed. The data layout of the input features in DRAM for the jth

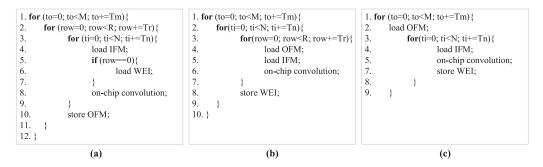


Fig. 15. Pseudo-code of loop order scheduling between tiles. (a) Loop order for FP and BP. (b) Loop order for WU. (c) Loop order for WU when $R^i \leq Tr^i$.

layer is shown in Figure 13(a). According to Figures 13(b) and 13(c), the burst length equals the size of a tile.

After selecting Tm = Tn, weights can be placed and fetched tile by tile during FP, BP, and WU. The data layout is illustrated in Figure 14. Before data reshaping, weights need to be repeatedly transmitted between the FPGA chip and the DRAM in FP and BP, which is inefficient especially for mini-batch training. Therefore, weight reuse based on our unique data layout is necessary, which will be introduced in detail in Section 4.3. After reshaping, the burst length for FP and WU is $M^i \times N^i$, while the burst length for BP is $Tm \times Tn$; these are shown in Figures 14(b) and 14(c), respectively.

Inter-tile Loop Order Optimization: The proposed data reshaping approach also achieves inter-tile data continuity by rescheduling the loop order in Figure 5. The loop order of 1, 2, and 3 in Figure 5(a) does not have data dependency. Based on our data layout, we move loop 3 to the outermost loop so that the output features share similar memory access patterns in FP/BP and WU. The loop order of off-chip data transmission in FP/BP is shown in Figure 15(a). As shown in Figures 12 and 13, in FP/BP, tiles of input features are fetched in the **X** direction first to generate the first output features tile. Then the tiles of output features are generated and stored in the **Y** direction first, so the input feature tiles' movement follows the **Y** direction as well. Then the output feature tiles are generated and stored in the **X** direction, and the access pattern of input features repeats $\lceil \frac{M}{Tm} \rceil$ times. The burst length of output features in the OUT DMA channel is $M^i \times R^i \times C^i$.

The loop order in Figure 15(b) is adopted in WU. From Figures 12 and 13, tiles of both input features and output features are fetched and stored in the Y direction first to calculate weight gradients for the first tile. Then weights are updated along the input channel dimension, so the input feature tiles move in the X direction, while the output feature access pattern (the dashed box in Figure 12(c)) repeats $\lceil \frac{N}{T_n} \rceil$ times. After that, weights are updated along the output channel dimension, so the output feature tiles move in the X direction, while the input feature access pattern (the dashed box in Figure 13(c)) repeats $\lceil \frac{N}{T_n} \rceil$ times. The burst length of output features in the OFM DMA channel is $Tm \times R^i \times C^i$. When the IFM buffer and the OFM buffer are large enough to hold the $Tn \times R^i = Tn$ in input and $Tm^i \times R^i \times C^i$ output features, i.e., $R^i \leq Tr^i$, the output features do not need to be repeatedly loaded. The loop order can be optimized as shown in Figure 15(c).

4.3 Weight Reuse in Mini-batch Training

Based on the above-mentioned optimization, we further reduce DRAM data access by reusing weights in mini-batch training. Different from inference, training involves processing a batch of

49:20 Y. Tang et al.

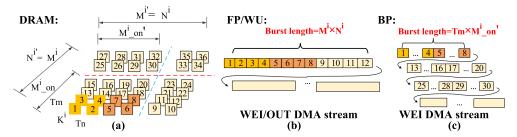


Fig. 16. Data layout of weights after weight reuse. (a) Weights stored in DRAM. (b) Weights transmitted via the WEI/OUT DMA channel in FP/WU. (c) Weights transmitted via the WEI DMA channel in BP.

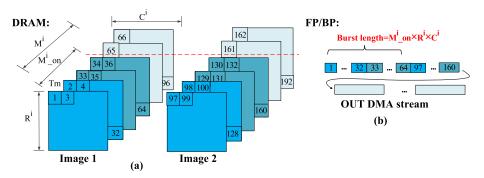


Fig. 17. Data layout of output features in mini-batch training after weight reuse. (a) Output features of two images in a batch. (b) Output features transmitted the via the OUT DMA channel in FP/BP.

data at once, so data reuse is necessary to decrease the transmission times of weights between onchip buffer and off-chip memory. On FPGAs, a BRAM bank size is large enough to store multiple tiles of weights. Therefore, we propose a weight reuse strategy based on our data layout. Thanks to our loop order shown in Figure 15, we can load weights only when the accelerator processes the output feature tile lying in the first row. As illustrated in Figure 16, when the accelerator processes a tile of features in the first row of the first image in a batch, M^i _on $\times N^i \times K^i \times K^i$ weights are loaded and stored in the WEI double buffers, where M^i _on is the multiple of Tm depending on the on-chip BRAM resources. After the first M^i _on channels of OFMs in the image are processed, the first M^i _on channels of OFMs of the next image will be processed, so weights do not need to be uploaded again. The next M^i _on channels of the first image will be processed after the first M^i _on channels of all images in the batch are processed. Therefore, weights do not need to be transmitted back and forth. After the above-mentioned steps, the burst length is $M^i \times N^i$ for FP/WU and $Tm \times M^i$ _on' for BP; this is illustrated in Figures 16(b) and 16(c), respectively.

In mini-batch training, weight reuse will not affect the burst length of output features in WU and input features. For output features in FP and BP, after the first M^i _on channels of OFMs of the first image are transmitted to DRAM, the next image of the batch will be processed before other channels of the prior image. Therefore, as shown in Figure 17, the burst length is M^i _on $\times R^i \times C^i$.

5 PERFORMANCE AND RESOURCE MODEL

In this section, we establish an analytic model to calculate the latency and resources for our design. Unlike previous works [16, 24], which only focused on the performance of a bare accelerator running on separate Conv layers, our model considers the discontinuity of off-chip memory

access in a realistic end-to-end training process. Based on the model, we build a scheduling tool to determine design parameters for given FPGA devices and given network models.

5.1 Performance Model

For a $[M^i, N^i, R^i, C^i, K^i, K^i, S^i]$ Conv layer i, we assume the parameters of a tile are $[Tm, Tn, Tr^i, Tc^i]$. Tm and Tn are fixed for all layers since they are determined by the number of DSPs, while Tr^i and Tc^i are adjustable according to different layer parameters. In our design, Tm = Tn, and $Tc^i = C^i$. The computation latency of a tile of features in FP, BP, and WU can be represented as $t^i_{COMP} = Tr^i \times Tc^i \times K^i \times K^i$ clock cycles.

The continuity of memory address significantly impacts the off-chip communication efficiency. To consider the memory access discontinuity, we assume the start time of the DMA stream is t_{start} . When discontinuity happens, DMA restarts. We have tested the start time on both the PYNQ-Z1 and the ZCU102 board, and $t_{start}\approx 400$ cycles under a 100MHz clock. We determine the data width parameters p to model the off-chip/on-chip communication bandwidth. For 32-bit floating point, if the DMA stream width is 128 bits, p=4. Since the burst length of input features equals the size of a tile, discontinuity happens every time a tile of input features is fetched. The latency of loading a tile of input features is formulated as $t_{IFM}^i = t_{start} + \lceil \frac{Tn}{p} \rceil \times ((Tr^i-1) \times S^i + K^i) \times ((Tc^i-1) \times S^i + K^i)$ clock cycles. The weight loading latency can be represented as $t_{WEI}^i = \lceil \frac{Tm \times Tn}{p} \rceil \times K^i \times K^i$ clock cycles, and the latency of storing a tile of output features is formulated as $t_{OUT}^i = \lceil \frac{Tm}{p} \rceil \times Tr^i \times Tc^i$ clock cycles. t_{start} is added to t_{WEI}^i and t_{OUT}^i only when the discontinuity happens, which will be discussed in detail as follows. We define $t_{LOAD}^i = \max\{t_{IFM}^i, t_{WEI}^i\}, t_{PROD1}^i = \max\{t_{IFM}^i, t_{COMP}^i\}, t_{PROD2}^i = \max\{t_{LOAD}^i, t_{COMP}^i\}, \text{ and } t_{STORE}^i = \max\{t_{COMP}^i, t_{OUT}^i\}.$ We assume in layer i that $M^i_on \times Tn \times K^i \times K^i$ weights are stored on-chip. If the batch size is B,

We assume in layer i that $M^i _on \times Tn \times K^i \times K^i$ weights are stored on-chip. If the batch size is B, weights will be loaded only during the iteration when the proposed accelerator processes the first image in the batch. For other iterations, the latency of processing $M^i _on$ channels of an image in FP can be formulated as follows:

$$Lat1^{i} = \left[\frac{N^{i}}{Tn} - 1\right] \times t_{PROD1}^{i} + t_{IFM}^{i} + t_{COMP}^{i}$$

$$\tag{15}$$

$$Lat2^{i} = \left[\frac{N^{i}}{Tn} - 1\right] \times t_{PROD1}^{i} + t_{IFM}^{i} + t_{STORE}^{i}$$
(16)

$$Lat3^{i} = \left(\left\lceil \frac{M^{i}_on}{Tm} \right\rceil \times \left\lceil \frac{R^{i}}{Tr^{i}} \right\rceil - 1 \right) \times Lat2^{i} + Lat1^{i} + t_{OUT}^{i} + t_{start}^{i}.$$
 (17)

Weights need to be loaded when our accelerator processes the first image in the mini-batch. In FP, t_{start} can be neglected in weight transmission since the burst length equals the size of weights, which means the addresses are continuous during the whole Conv layer. Therefore, the latency of the proposed accelerator processing M^i _on channels of the first image can be formulated as follows:

$$Latb1^{i} = \left[\frac{N^{i}}{Tn} - 1\right] \times t_{PROD2}^{i} + t_{LOAD}^{i} + t_{COMP}^{i}$$

$$\tag{18}$$

$$Latb2^{i} = \left[\frac{N^{i}}{Tn} - 1\right] \times t_{PROD2}^{i} + t_{LOAD}^{i} + t_{STORE}^{i}$$
(19)

$$Latb3^{i} = \left\lceil \frac{M^{i}_on}{Tm} \right\rceil \times \left\lceil \frac{R^{i}}{Tr^{i}} - 1 \right\rceil \times Lat2^{i} + \left\lceil \frac{M^{i}_on}{Tm} - 1 \right\rceil \times Latb2^{i} + Latb1^{i} + t_{OUT}^{i} + t_{start}^{i}. \quad (20)$$

49:22 Y. Tang et al.

The latency of our accelerator processing the whole Conv layer in FP is formulated in Equation (21):

$$Lat^{i} = \left[\frac{M^{i}}{M^{i} on}\right] \times ((B-1) \times Lat3^{i} + Latb3^{i}). \tag{21}$$

In BP, the situation is similar to that in FP, except that the addresses of weights are discontinuous after M^i _on channels. The accelerator loads M^i _on \times $Tn \times K^i \times K^i$ weights together when processing the first tile of the first image, so it costs $t^i_{WEI} = \lceil \frac{M^i _on \times Tn}{p} \rceil \times K^i \times K^i + t^i_{start}$ clock cycles. $Lat1^i$, $Lat2^i$, $Lat3^i$, $Lat2^i$, $Latb1^i$, and $Latb2^i$ remain unchanged, while $Latb3^i = (\lceil \frac{M^i _on}{Tm} \rceil \times \lceil \frac{R^i}{Tr^i} \rceil - 1) \times Lat2^i + Latb1^i + t^i_{OUT} + t^i_{start}$.

 $(\lceil \frac{M^i_on}{Tm} \rceil \times \lceil \frac{R^i}{Tr^i} \rceil - 1) \times Lat2^i + Latb1^i + t^i_{OUT} + t^i_{start}.$ In WU, loss features are loaded from the off-chip memory to the OFM buffer. Transmitting a tile of loss features costs $t^i_{OFM} = t^i_{start} + Tr^i \times Tc^i \times \lceil \frac{Tm}{p} \rceil$ clock cycles. Weights are updated after all the gradients of the batch are accumulated, so transmitting the updated weights costs the same time as loading weights, which means $t^i_{OUT} = t^i_{WEI}$. As with FP, t_{start} can be neglected when calculating t^i_{WEI} . We define $t^i_{LOAD} = \max\{t^i_{IFM}, t^i_{OFM}\}, t^i_{PROD1} = \max\{t^i_{LOAD}, t^i_{COMP}\}, t^i_{PROD2} = \max\{t^i_{IFM}, t^i_{COMP}\}, \text{ and } t^i_{STORE} = \max\{t^i_{COMP}, t^i_{OUT}\}.$ The latency of WU of the ith Conv layer is formulated as follows:

$$Lat1^{i} = \left[\frac{R^{i}}{Tr^{i}} - 1\right] \times t_{PROD1}^{i} + t_{LOAD}^{i} + t_{COMP}^{i}$$
 (22)

$$Latb1^{i} = \left[\frac{R^{i}}{Tr^{i}} - 1\right] \times t_{PROD1}^{i} + t_{LOAD}^{i} + t_{STORE}^{i}$$
(23)

$$Lat^{i} = \left(\left((B-1) \times \left\lceil \frac{M^{i}_on}{Tm} \right\rceil \times \left\lceil \frac{N^{i}}{Tn} \right\rceil + 1 \right) \times Lat1^{i} + \left(\left\lceil \frac{M^{i}_on}{Tm} \right\rceil \times \left\lceil \frac{N^{i}}{Tn} \right\rceil - 1 \right) \times Latb1^{i} + t^{i}_{OUT} \right) \times \left\lceil \frac{M^{i}}{M^{i}_on} \right\rceil. \quad (24)$$

As illustrated in Figure 15(c), when $R^i \leq Tr^i$, the output features do not need to be repeatedly loaded. Under this circumstance, the latency of WU is formulated as follows:

$$Lat1^{i} = \left[\frac{N^{i}}{Tn^{i}} - 1\right] \times t_{PROD2}^{i} + t_{LOAD}^{i} + t_{COMP}^{i}$$

$$\tag{25}$$

$$Latb1^{i} = \left[\frac{N^{i}}{Tn^{i}} - 1\right] \times \left(t_{PROD2}^{i} + t_{OUT}^{i}\right) + t_{LOAD}^{i} + t_{COMP}^{i} + t_{OUT}^{i}$$
 (26)

$$Lat^{i} = \left[\frac{M^{i}}{M^{i}_on}\right] \times \left[\frac{M^{i}_on}{Tm}\right] \times \left((B-1) \times Lat1^{i} + Lat1b^{i}\right). \tag{27}$$

5.2 Resource Model

$$D_{Conv} = q \times Tm \times Tn < total DSPs number$$
 (28)

$$B_{IFM} = \max_{i} B_{IFM}^{i} = \max_{i} \left\{ Tn \times \left[\frac{((Tr^{i} - 1) \times S^{i} + K^{i}) \times ((Tc^{i} - 1) \times S^{i} + K^{i}) \times BITs}{Size \ of \ a \ BRAM \ Bank} \right] \right\} \quad (29)$$

$$B_{OFM} = \max_{i} B_{OFM}^{i} = \max_{i} \left\{ Tm \times \left[\frac{Tr^{i} \times Tc^{i} \times BITs}{Size \ of \ a \ BRAM \ Bank} \right] \right\}$$
(30)

$$B_{WEI} = \max_{i} B_{WEI}^{i} = \max_{i} \left\{ Tm \times Tn \times \left[\frac{K^{i} \times K^{i} \times \left\lceil \frac{N^{i}}{2 \times Tn} \right\rceil \times \left\lceil \frac{M^{i} - on}{Tm} \right\rceil \times BITs}{Size \ of \ a \ BRAM \ Bank} \right] \right\}$$
(31)

$$B_{Conv} = 2 \times (B_{IFM} + B_{OFM} + B_{WEI}) < total BRAMs number$$
 (32)

It should be noted that in realistic end-to-end system design, the boundary of D_{Conv} and B_{Conv} should be slightly smaller than the total DSP and BRAM numbers. This is because except for the MAC operations, several operations also take up a small fraction of on-chip resources. For example, some non-Conv layers (e.g., maximum pooling, average pooling, ReLU, etc.), which are inevitable in practical end-to-end training processes, need extra DSPs to make comparisons and extra BRAMs to buffer the indexes. Besides, some neural networks have irregular weight kernel shapes for different Conv layers. Adding an extra buffer to fetch a tile of weights from the on-chip Weight buffer to the Conv Kernel can relieve the routing congestion in realistic FPGA implementation. Besides, since FP, BP, and WU have different loop orders, extra DSPs are utilized to calculate BRAM addresses under different layer parameters. This address calculation is much more complex than that in inference. Therefore, in practical design, the estimated boundary of the on-chip resources should be set lower than the available resources empirically. The details will be further explained in Section 5.3.

5.3 Computation and Memory Resources Scheduling Tool

Based on the above-mentioned model, we build a computation and memory resources scheduling tool for different devices and different networks. Algorithm 1 shows the framework of our scheduling tool. As mentioned in Section 5.2, D_{Conv} and B_{Conv} are lower than the total DSP and BRAM numbers in realistic FPGA implementation. Therefore, it is wise to set a boundary for D_{Conv} and B_{Conv} that is lower than the available on-chip resources. According to the experimental results in Section 6, assigning 80% of DSPs and 75% BRAMs to the estimated boundary for D_{Conv} and B_{Conv} should be enough. Then we determine Tm and Tn according to the DSP number. Then we choose the optimal Tr^i , Tc^i , and M^i _on for each layer according to Equations (15) to (27). Specifically, in steps 3 and 4, we find the lower bound for B_{IFM} and B_{OFM} by assuming that the buffers can only hold one row for the largest feature maps. Then, from step 5 to step 12, we try to assign resources for Weight buffers so that they can hold as many weights for each layer as possible. After we determine B_{WEI} and M^i _on for each layer, we re-assign IFM and OFM buffers under the constraints shown in Equations (29), (30), and (32) and find the optimal Tr^i and Tc^i for each layer. After Tm, Tn, and $[Tr^i, Tc^i, M^i$ _on]_{1≤i} are determined, we can calculate the DMA start addresses for each layer offline based on the data reshaping approach in Section 4.

6 EXPERIMENTS

The proposed work is evaluated on edge-level FPGAs PYNQ-Z1 and ZCU102 with working frequency at 100MHz. The accelerator is designed with Vivado HLS, which generates IP core from C language. The obtained IP cores are connected, synthesized, and implemented in Vivado (v2019.1).

49:24 Y. Tang et al.

ALGORITHM 1: Computation and Memory Resources Scheduling

Input:

CNN layers parameters $[M^i, N^i, R^i, C^i, K^i, K^i, S^i]_{1 \le i \le n}$, batch size B, data type parameter q, DMA stream width, total DSPs number, total BRAMs number;

Output:

```
Tm, Tn, [Tr^{i}, Tc^{i}, M^{i}\_on]_{1 \le i \le n}, B_{IFM}, B_{OFM}, B_{WEI};
1: Estimate the boundary for D_{Conv} and B_{Conv};
2: Assign Tm, Tn according to Equation (28), while Tm = Tn;
3: Find k = \arg\max_{i} \{R^i \times C^i\};
4: Determine the lower bound for B_{IFM} and B_{OFM}, i.e., inf B_{IFM} = B_{IFM}^k, inf B_{OFM} = B_{OFM}^k
   when Tc^k = C^k, Tr^k = 1;
5: for i = 1; i \le n; i + + do
      Calculate B_{WEI}^{i} based on (31) when M^{i}_on = M^{i}, and initialize l = 1;
      if 2 \times (\inf B_{IFM} + \inf B_{OFM} + B_{WEI}^i) \ge estimated B_{Conv} boundary then
7:
8:
         Find the minimal M^i_on satisfying \frac{M^i}{I} \le M^i_on, M^i_on mod Tm = 0, and go to Step 7;
9:
10:
11: end for
12: Calculate B_{WEI} and M^{i}_on for each layer based on Equation (31);
13: for i = 1; i \le n; i + + do
      Set Tc^i = C^i, and select all Tr_m^i satisfying Equations (29), (30), and (32), where 1 \le Tr_m^i \le R^i;
14:
      Determine Tr^i = \arg \min Lat_m^i based on Equations (15) to (27);
17: Calculate B_{IFM} and B_{OFM} based on Equations (29) and (30);
```

The Vivado Project Summary reports resource utilization and power after implementation. Finally, we employ Xilinx SDK to program SoC on PYNQ-Z1 and ZCU102 to achieve end-to-end CNN training.

6.1 Effectiveness of the Data Reshaping Approach

In this section, we need to validate the effectiveness of the proposed data reshaping approach ([Tm, Tn] = [16, 16]). We test the Conv layers of the AlexNet on ZCU-102. We select the batch size B as 4 and the DMA stream width as 128 bits. We adopt the results using the BCHW data layout and the results using the BHWC data layout as baselines ($[Tm_{Base}, Tn_{Base}] = [32, 8]$). The BCHW pattern does not involve any optimization. For the BHWC pattern, N/Tn tiles of the input features and M/Tm tiles of the output features are buffered in the on-chip BRAM for data reuse based on the loop order in the inference phase. Weights are pre-allocated tile by tile based on the data flow in inference. The comparisons are shown in Tables 3, 4, and 5.

As mentioned in Section 2.3, our goal is to design a general accelerator supporting end-to-end training with both dense and small networks without sacrificing precision, so it is necessary to appropriately manage external memory access and allocate on-chip buffers. When applying loop tiling, the tiling schemes involved in the accelerator design break the continuity of data addresses in DRAM and thus reduce the transmission efficiency between on-chip buffer and off-chip DRAM. Table 3 shows the experimental results of our baseline, which is a bare accelerator with the unified channel-level parallelism-based convolution kernel. It does not involve any optimizations related to the off-chip DRAM access policy. As illustrated in Section 4, the burst

AlexNet	Process	$[Tr_{Base}^{i}, Tc_{Base}^{i}]$	Acceleration (Cycles)	Reallocation (Cycles)	Total (Cycles)
			(Cycles)	(Cycles)	(Cycles)
Conv 1	FP	[11, 11]	6,732,837	151,846,336	158,579,173
	BP	N/A	N/A	N/A	N/A
	WU	[11, 11]	4,496,029	152,110,235	156,606,264
Conv 2	FP	[27, 27]	7,105,292	69,743,160	76,848,452
	BP	[27, 27]	7,066,705	68,271,764	75,338,469
	WU	[27, 27]	9,258,823	57,303,397	66,562,220
Conv 3	FP	[13, 13]	2,410,532	101,062,954	103,473,486
	BP	[13, 13]	2,401,320	98,646,892	101,048,212
	WU	[13, 13]	4,448,898	83,566,193	88,015,091
Conv 4	FP	[13, 13]	3,596,425	150,012,382	153,608,807
	BP	[13, 13]	3,596,400	149,621,995	153,218,395
	WU	[13, 13]	6,669,238	126,214,297	132,883,535
Conv 5	FP	[13, 13]	2,401,212	102,632,162	105,033,374
	BP	[13, 13]	2,410,637	99,408,011	101,818,648
	WU	[13, 13]	4,448,751	84,518,969	88,967,720
Total			67,043,099	1,494,958,747	1,562,001,846

Table 3. Experimental Results of the Baseline with the BCHW Data Layout

Table 4. Experimental Results of the Baseline with the BHWC Data Layout and Data Reuse

AlexNet	Process	$[Tr_{Base}^{i}, Tc_{Base}^{i}]$	Acceleration (Cycles)	Reallocation (Cycles)	Total (Cycles)
Conv 1	FP	[11, 11]	8,094,251	N/A	8,094,251
	BP	N/A	N/A	N/A	N/A
	WU	[11, 11]	4,495,794	161,048,775	165,544,569
Conv 2	FP	[27, 27]	7,383,996	N/A	7,383,996
	BP	[27, 27]	7,382,504	68,200,715	75,583,219
	WU	[27, 27]	7,848,249	N/A	7,848,249
Conv 3	FP	[13, 13]	2,531,247	N/A	2,531,247
	BP	[13, 13]	2,529,216	100,372,954	102,902,170
	WU	[13, 13]	3,345,845	N/A	3,345,845
Conv 4	FP	[13, 13]	3,745,972	N/A	3,745,972
	BP	[13, 13]	3,745,922	148,657,460	152,403,382
	WU	[13, 13]	4,999,576	N/A	4,999,576
Conv 5	FP	[13, 13]	2,529,173	N/A	2,529,173
	BP	[13, 13]	2,531,318	100,586,051	103,117,369
	WU	[13, 13]	3,364,408	N/A	3,364,408
Total			64,527,471	578,865,955	643,393,426

length before data reshaping is much smaller than the size of a tile. To ensure that the accelerator conducts MAC operations with correct features and weight matrices in realistic end-to-end training, data should be reallocated before being transmitted from DRAM to the on-chip accelerator. Therefore, our baseline includes the on-chip acceleration time and off-chip reallocation time. After

49:26 Y. Tang et al.

AlexNet	Process	$[Tr^i, Tc^i]$	Without Weight Reuse (Cycles)	After Weight Reuse (Cycles)
Conv 1	FP	[2, 55]	11,498,545	11,419,835
	BP	N/A	N/A	N/A
	WU	[2, 55]	9,598,744	9,299,086
Conv 2	FP	[27, 27]	7,283,187	7,312,794
	BP	[27, 27]	7,128,663	7,146,578
	WU	[27, 27]	7,910,148	7,430,533
Conv 3	FP	[13, 13]	2,491,672	2,510,310
	BP	[13, 13]	2,461,694	2,671,392
	WU	[13, 13]	3,402,418	2,706,696
Conv 4	FP	[13, 13]	3,689,930	3,708,934
	BP	[13, 13]	3,688,961	3,972,757
	WU	[13, 13]	5,053,485	4,014,651
Conv 5	FP	[13, 13]	2,462,778	2,475,263
	BP	[13, 13]	2,490,897	2,686,910
	WU	[13, 13]	3,373,373	2,677,726
Total			72,534,495	70,033,465

Table 5. Experimental Results Validating Data Reshaping Approach

applying data reshaping, data can be fetched from DRAM to the accelerator directly without extra reallocation.

For Conv 1, the number of input channels is only three, $Tn_{Base}=8$, and Tn=16. Therefore, 5/8 computation resources remain idle for the baseline, while 13/16 computation resources for our proposed design remain idle. That's why the acceleration time for the baseline is shorter than the latency in our proposed design. However, features should be reallocated before entering the next layer (for FP) or after being generated from the prior layer (for WU). As shown in Table 3, the reallocation time is much longer than the acceleration time. For Conv 2 to Conv 5, $Tr^i \geq R^i$ and $Tc^i \geq C^i$, so features do not need to be reallocated between adjacent layers, but weights still need to be reallocated before entering the Conv layer (for FP and BP) or updated from the Conv layer (for WU). To sum up, the total acceleration time for the baseline is close to that for our proposed design under the same degree of parallelism $(Tm_{Base} \times Tn_{Base} = Tm \times Tn)$ and tile size boundary of features ($\max_i Tr^i_{Base} \times \max_i Tc^i_{Base} = \max_i Tr^i \times \max_i Tc^i$), but the extra reallocation time in realistic end-to-end training is even longer than the acceleration time. Therefore, accelerating without considering the actual data layout in DRAM between adjacent layers is inefficient in realistic end-to-end training.

The baseline in Table 4 uses the BHWC data layout and applies data reuse to alleviate the discontinuous memory access. As illustrated in Figures 9 to 11, features and weights are continuous in a long burst length in FP, so data are not reallocated during the Conv layers. Such an approach is efficient in the inference phase. In BP, although the memory access pattern of features is the same as that in FP, the weight transmission patterns are quite different. As shown in Figure 11(c), the burst length is much less than the size of a tile, so weights should be reallocated in each Conv layer. The extra reallocation time is much longer than the acceleration time. In WU, the on-chip buffer can hold all the features for Conv 2 to Conv 5 layers, so it is practical to load all the features to the FPGA chip without extra reallocation. However, in the Conv 1 layer, the on-chip memory cannot hold all the features. Even though the input features can be pre-allocated before entering into the

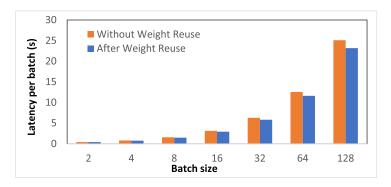


Fig. 18. Experimental results of the data reshaping approach without weight reuse and after weight reuse.

neural network since they serve as the inputs for the whole process, the output features that are calculated in BP cannot be allocated ahead of time. Therefore, the Conv 1 layer also requires extra reallocation time in WU, which is quite inefficient.

As mentioned in Section 4, we optimize the DRAM access incrementally. We first achieve intratile continuous memory allocation by reorganizing the DRAM layouts for input features, output features, and weights. Then we re-schedule the loop order to achieve inter-tile continuous memory allocation. These two parts are combined to improve memory access continuity together when the batch size is one. Considering the training process involves convolution operations among a minibatch, we further propose and apply a weight reuse strategy based on the proposed data layout. Table 5 shows the experimental results of the data reshaping approach without weight reuse and after weight reuse. The batch size is also four, and the latency for FP/BP without and after weight reuse is nearly the same. This is because, in FP and BP, input features and weights are transmitted together. As can be seen from Section 5, when $\lceil \frac{Tn}{p} \rceil \times ((Tr^i - 1) \times S^i + K^i) \times ((Tc^i - 1) \times S^i + K^i) > 1$ $\lceil \frac{Tm \times Tn}{p} \rceil \times K^i \times K^i$, $t^i_{IFM} > t^i_{WEI}$. Therefore, reusing the weights may not reduce the latency as a whole. However, in WU, the latency with weight reuse is apparently less than that without reuse. This is because the transmission of weights happens during storing the results in WU, which cannot be totally covered by t_{COMP}^i (e.g., in the last iteration of the loop in line 2 from Figure 15(b) and the loop in line 3 from Figure 15(c)). As a whole, reusing weights can reduce the latency of the training phase of the whole network. Figure 18 shows the latency without and with weight reuse when the batch size ranges from 2 to 128. It shows that when the batch size increases, applying the reuse strategy has more apparent advantages than only achieving intra-tile and inter-tile data access continuity.

6.2 Accuracy of the Performance Model

After displaying the effectiveness of the data reshaping approach, we use the AlexNet to validate the accuracy of the performance model. Our scheduling tool first determines optimal tiling parameters, which are shown in Table 6. Then the latency is estimated by our model and tested on-board separately. As shown in Table 6, the estimated results are close to the tested results. The results verify the accuracy of the performance model.

6.3 CNN Training Performance

In this section, we conduct end-to-end evaluations on different neural networks. We first compare our design with the automatic compiler-based FPGA accelerator [22]. It adopted a combination of channel-level parallelism and feature-map-level parallelism with the unrolling factors for

49:28 Y. Tang et al.

AlexNet	Process	$[Tr^i, Tc^i, M^i_on]$	Our Model (Cycles)	On-board (Cycles)	Deviation
Conv 1	FP	[2, 55, 96]	11,504,640	11,419,835	0.74%
	BP	N/A	N/A	N/A	N/A
	WU	[2, 55, 96]	9,043,384	9,299,086	2.75%
Conv 2	FP	[27, 27, 112]	7,309,808	7,312,794	0.04%
	BP	[27, 27, 48]	7,126,784	7,146,578	0.28%
	WU	[27, 27, 112]	7,423,616	7,430,533	0.09%
Conv 3	FP	[13, 13, 112]	2,478,272	2,510,310	1.28%
	BP	[13, 13, 112]	2,566,987	2,671,392	3.91%
	WU	[13, 13, 112]	2,682,240	2,706,696	0.90%
Conv 4	FP	[13, 13, 112]	3,646,400	3,708,934	1.69%
	BP	[13, 13, 112]	3,861,220	3,972,757	2.81%
	WU	[13, 13, 112]	3,960,960	4,014,651	1.34%
Conv 5	FP	[13, 13, 112]	2,432,368	2,475,263	1.73%
	BP	[13, 13, 112]	2,618,372	2,686,910	2.55%
	WU	[13, 13, 112]	2,640,640	2,677,726	1.38%
Total			69,295,691	70,033,465	1.05%

Table 6. Experimental Results Validating the Performance Model

columns, rows, and output channels. It initially stored weights tile by tile in a transposable format in DRAM and read the entire weights of a Conv layer from DRAM to their on-chip buffer. The baseline implemented a "1X" CNN on the CIFAR-10 dataset with the structure as Conv 1 ($[M^i, N^i, R^i, C^i, K^i, S^i] = [16, 3, 32, 32, 3, 1]$) - Conv 2 ([16, 16, 32, 32, 3, 1]) - Pooling - Conv 3 ([32, 16, 16, 16, 3, 1]) - Conv 4 ([32, 32, 16, 16, 3, 1]) - Pooling - Conv 5 ([64, 32, 8, 8, 3, 1]) - Conv 6 ([64, 64, 8, 8, 3, 1]) - Pooling - FC ([10, 1024, 1, 1, 1, 1]), using 16-bit fixed-point precision. We test the same network on both PYNQ-Z1 and ZCU102 boards. The DMA stream bandwidth is 128 bits for ZCU102 and 32 bits for PYNQ-Z1. Our design focuses on implementing on-device FPGAs without sacrificing precision, so 32-bit floating point is adopted. The Vivado utilization report provides the utilization of BRAMs and DSPs, and the power report provides the total on-chip power. We measure the latency of training the whole batch with the batch size of 128. Then we calculate the latency per image and the throughput.

Table 7 shows the comparison results between the baseline [22] and our design in terms of resource utilization, throughput, energy efficiency, and so forth. The Stratix 10 GX adopted in the baseline is an advanced FPGA board developed by Intel. It is unfair to compare the throughput directly for different devices. However, energy efficiency is an important metric to judge the performance of edge devices, and thus we use energy efficiency as the metric for different designs on different FPGAs. We nominate the throughput and efficiency by multiplying the bit width of the data type. Although using the fixed-point data type is much more DSP efficient and power efficient than adopting floating point under the same bit width, our nominal efficiency still can outperform that of the baseline. The reason is that the baseline has more data transmission latency, especially for WU, where accessing weight gradients, weights, and storing back the updated values leads to DRAM access latency. Fifty-one percent of the overall latency in one iteration of a batch is consumed in WU [22]. Figure 19 shows the latency breakdown of our design. The total latency for each training process is calculated by summarizing the latency for each Conv layer, and the latency for MAC is the theoretical computation latency calculated by accumulating t_{COMP}^i for each Conv layer based on the performance model. Since the "1X" CNN is a relatively small network,

	Baseline [22]	Ours	Ours
Platform	Stratix 10 GX	PYNQ-Z1	ZCU102
Frequency (MHz)	240	100	100
DSP Utilization	1,699 (30%)	212 (96.4%)	1,315 (52.2%)
D_{Conv} ($D_{Conv}/Used DSPs$)		180 (84.9%)	1,280 (97.3%)
BRAM Utilization	10.6 (4.4%)	123 (87.9%)	324 (35.5%)
$B_{Conv}(B_{Conv}/Used\ BRAMs)$		108 (87.8%)	288 (88.9%)
Power (W)	20.6	1.85 (11.14X)	6.89 (2.99X)
Data Type	Fixed 16	FP 32	FP 32
Batch Size	40	128	128
Latency/Image (ms)	0.36	14.32	2.08
Throughput	163 GOPS	4.08 GFLOPS	28.15 GFLOPS
Nominal Throughput	2608	130.56	900.8
$(GOPS \times precision)$	2000	130.30	900.0
Energy Efficiency	7.90 GOPS/W	2.21 GFLOPS/W	4.09 GFLOPS/W
Nominal Effciency			
$(GOPS \times precision/W)$	126.4	70.72	130.88 (1.04X)

Table 7. Experimental Results on the "1X" CNN

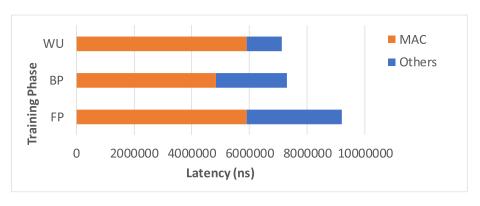


Fig. 19. Latency breakdown of CIFAR-10 "1X" CNN for FP, BP, and WU when the batch size is 128.

the number of loops is also small. According to the performance model, although double buffers are adopted, the computation and data transmission are conducted sequentially in the first and last iteration of the loop, while they are in parallel for the middle iterations. Therefore, when the number of loops is small, the proposed design also includes much data transmission latency for FP, BP, and WU. However, other optimizations like loop order scheduling and weight reuse in a mini-batch reduce the number of off-chip memory access. Therefore, our computation latency is still much more than 50% of the total latency in FP, BP, or WU, which takes up a larger proportion compared with the baseline.

Besides, the baseline stored the entire weights of a Conv layer from DRAM to the on-chip buffer. Their design cannot support denser networks where the on-chip buffer cannot hold the entire weights of each Conv layer. However, our design does not have such restrictions and can support many larger networks.

49:30 Y. Tang et al.

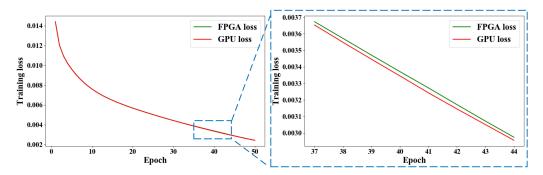


Fig. 20. The loss curves during the training phase.

In Table 7, the D_{Conv} and B_{Conv} are the DSP and BRAM numbers for the Conv layer estimated by our resource model. The percentage shows the ratio between the estimated resources and actual resources used in the whole end-to-end training. As mentioned in Sections 5.2 and 5.3, D_{Conv} and B_{Conv} are lower than the DSP and BRAM numbers used in the realistic end-to-end training process. For "1X" CNN, the extra on-chip resources mainly function for maximum pooling layers. Besides, a few DSPs are utilized to calculate BRAM addresses of features and weights.

To validate the correctness of our design, we also implement the whole training phase of the "1X" CNN on ZCU102 and compare the training result with that on GPU. We load the Cifar-10 dataset from the **secure digital (SD)** card to the DRAM and run 50 epochs. The batch size is also 128, and the learning rate is 0.008. We use the V-100 GPU from AWS to validate the training process. The loss curves are shown in Figure 20. Since we adopt full precision and have not changed the training algorithm, the training result should be nearly the same as that on GPU. As can be seen in Figure 20, the two curves are really close to each other. We also test the trained model on the test dataset. The test accuracy is 65.22% running on GPU and 64.82% running on FPGA.

Most state-of-the-art works [22, 23] mainly implemented their design on the Cifar-10 dataset, whose input image is really small ($3 \times 32 \times 32$) compared to real-world on-device learning scenarios. To verify that our accelerator with the data reshaping approach can support larger networks with larger feature sizes, we test our design on AlexNet and Vgg-16 for ImageNet whose input image parameters are $3 \times 227 \times 227$ and $3 \times 224 \times 224$, respectively.

The data reshaping approach enables our accelerator to support end-to-end training with the following situations: (1) when the feature map size of a layer increases and the on-chip memory is not big enough to hold all the feature maps of the layer and (2) when the number of channels increases and the weight buffer cannot hold all the weights of a layer. For AlexNet, its convolution kernel size ranges from 11 \times 11 to 1 \times 1 and feature map size ranges from 227 \times 227 to 1 \times 1, which covers the above-mentioned situations. Besides, the stride of the first Conv layer of AlexNet is 4. Implementing Conv layers with different stride sizes is more complex than only verifying the design on CNNs where the stride remains one. Therefore, AlexNet is ideal to verify that our design can support DNNs with a larger feature map size and larger weight density and can deal with different Conv layers shapes. Figure 21(a) shows the throughput and latency of a batch for training the AlexNet model with batch size ranging from 2 to 128. When the batch size is 128, the throughput reaches 34.52 GFLOPS. Because of weight reuse, the weight transmission bottleneck is ameliorated when the batch size increases, so the throughput in larger batch size is slightly higher than that for small batch size. However, unlike batch-level parallelism-based designs [23] where the performance varies a lot under different batch sizes, the performance of our channel-level parallelism-based design is less affected by the batch size. As shown in Figure 21, the throughput when the batch size is two is still above 32 GFLOPS.

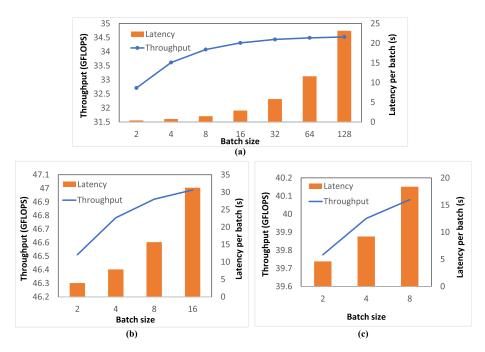


Fig. 21. Experimental results training different CNNs. (a) Throughput and latency of AlexNet. (b) Throughput and latency of Vgg-16 without BN layers. (c) Throughput and latency of Vgg-16 with BN layers.

We also test our design on Vgg-16, which has denser parameters, and the performance is shown in Figure 21(b). Due to the DRAM memory size limitation of ZCU102, the maximum batch size is 16. As illustrated in Figure 21, our design achieves higher throughput on Vgg-16 compared with AlexNet. This is because, for channel-level parallelism, the number of input channels of the first Conv layer is only three, which is smaller than Tn, so the computation resources are not fully utilized in this layer. This effect is also mentioned in Section 6.1. However, such underutilization only happens in the first Conv layer and is alleviated when the neural network becomes deeper. Hence, in the deeper network, Vgg-16, we achieve higher throughput.

To verify that our design can support the BN layer, which is a key component of typical CNN architectures, we also test the proposed design on Vgg-16 with BN layers. The performance is shown in Figure 21(c). Apart from the loss and activation, the immediate BN parameters also need to be stored in DRAM. Due to the memory size limitation, the maximum batch size is eight. Unlike computation-intensive Conv layers, BN layers involve lots of data transmission processes. Some complex operations like extracting a root also cost extra computation resources and reduce the timing performance. Therefore, the overall throughput is a little less than that for Vgg-16 without BN layers.

Table 8 also shows the resource utilization and energy efficiency of the FPGA for these networks. With the same estimated DSPs and BRAMs for Conv layers (D_{Conv} and B_{Conv}), AlexNet requires more BRAMs than Vgg-16. This is because, compared to Vgg-16, AlexNet has a less regular weight kernel shape (ranging from 11 × 11 to 1 × 1), so we add an extra buffer to fetch a tile of weights from the on-chip Weight buffer before the Conv Kernel conducting MAC operations. Such optimization can release routing congestion caused by complex BRAM address calculation and allocation in FP, BP, and WU processes. Apart from the extra buffer, a small fraction of DSPs and BRAMs function for non-Conv layers. The accelerator also utilizes a few DSPs to calculate BRAM addresses.

49:32 Y. Tang et al.

Network	AlexNet	Vgg-16 without BN	Vgg-16 with BN
DSP Utilization	1,513 (60.0%)	1,508 (59.8%)	1,680 (66.7%)
D_{Conv} ($D_{Conv}/Used$ DSPs)	1,280 (84.6%)	1,280 (84.9%)	1,280 (76.2%)
BRAM Utilization	857 (94.0%)	787 (86.3%)	812 (89.0%)
B_{Conv} ($B_{Conv}/Used$ BRAMs)	672 (78.4%)	672 (85.4%)	672 (82.8%)
Power (W)	7.736	7.712	8.208
Batch Size	128	16	8
Throughput (GFLOPS)	34.52	46.99	40.08
Efficiency (GFLOPS/W)	4.46	6.09	4.88

Table 8. Experimental Results on AlexNet and Vgg-16

Therefore, as mentioned in Section 5.3, the estimated boundary of D_{Conv} and B_{Conv} in realistic end-to-end system design should be slightly smaller than the total DSP and BRAM numbers. From our experimental results, assigning 80% of DSPs and 75% of BRAMs should be enough.

As for the Vgg-16 with BN layers, extra computation resources are utilized to do complex operations such as division, root extraction, and so forth. Therefore, Vgg-16 with BN layers costs more DSP resources compared with Vgg-16 without BN layers. Additional BRAMs are also utilized to buffer BN parameters for a batch.

In our end-to-end training validation, we utilize 1,508 DSPs for the Vgg-16 model. The theoretical peak performance with 1,508 DSPs on the 32-bit floating-point accelerator is $\frac{1508}{5} \times 2 \times 0.1$ GHz= 60.3 GFLOPS, while our attainable end-to-end test is 46.99 GFLOPS including pooling and ReLU operations.

6.4 Comparison with State-of-the-art Works

Comparisons of the best performance between our design and other state-of-the-art FPGA-based training accelerators are shown in Table 9. In the table, "N/A" means that the metric is not provided, and "≈" means that the value is obtained by approximate estimation. Since the platforms, the neural networks for training, and the data type are different, it is extremely difficult to fairly compare between different training accelerators. However, our design still shows desirable performance even under such circumstances.

To better illustrate the uniqueness of the proposed design, we also compare our work with the accelerators that also adopted 32-bit floating point. The comparisons are shown in Tables 10 and 11. The design in [36] was tested on LeNet-10, which is a really small network with the structure as Conv 1 ($[M^i, N^i, R^i, C^i, K^i, S^i] = [32, 3, 32, 32, 31]$) - Pooling - Conv 2 ([32, 32, 16, 16, 3, 1]) -Pooling - Conv 3 ([64, 32, 8, 8, 3, 1]) - Pooling - FC ([64, 1024, 1, 1, 1, 1]) - FC ([10, 64, 1, 1, 1, 1]). As explained in Section 6.3, the underutilization of computation resources in the first Conv layer reduces the overall throughput. Therefore, the performance of the proposed design on this small network cannot be as superior as that in deeper networks like Vgg-16. However, our design is a general architecture that can support both small networks and larger networks, while the accelerator in [36] only targeted small networks. It first achieved feature-map-level parallelism in a uniform computation engine and then unrolled channel-level parallelism factors to improve utilization of computation resources. For the memory access issues, the input and output features of each layer are all stored on the FPGA chip, which restricts the work from extending to support the networks where the on-chip BRAMs are not large enough to hold entire features of a Conv layer. However, larger networks like AlexNet and Vgg are commonly applied in practical applications. Unlike [36], our work not only enables on-device training on larger CNN models but also achieves

Accelerator	Liu et al. 2017 [36]	DarkFPGA 2020 [23]	Venkataramanaiah et al. 2020 [40]	FeCaffe 2020 [41]	Ours
Platform	ZU19EG	XCVU9P	Stratix 10 MX	Stratix 10	ZCU102
Technology	16nm	16nm	14nm	14nm	16nm
DSP Util.	1,500	4,202	1,040	1,796	1,508
Freq. (MHz)	200	200	185	253	100
Power (W)	14.24	13.5	≈20	N/A	7.712
Network	LeNet-10	Vgg-like	ResNet-20	AlexNet	Vgg-16
Dataset	CIFAR-10	CIFAR-10	CIFAR-10	ImageNet	ImageNet
Data Type	FP 32	Fixed 8	FP 16	FP 32	FP 32
Throughput	86.12	1417	≈180	≈24	46.99
	GFLOPS	GOPS	GFLOPS	GFLOPS	GFLOPS
Energy Effi.	6.05	104.96	≈9	N/A	6.09
	GFLOPS/W	GOPS/W	GFLOPS/W		GFLOPS/W
Nominal	2,755.84	11,336	≈2,880	≈768	1,503.68
Thro. (GOPS \times					
precision)					
Nominal	193.6	839.68	≈144	N/A	194.88
Effi. (GOPS $ imes$					
precision/W)					

Table 9. Comparison of Different FPGA-based Training Accelerators

Table 10. Experimental Results on LeNet-10 Compared with Liu et al. [36]

	Liu et al. [36]	Ours
Platform	ZU19EG	ZCU102
Frequency (MHz)	200	100
DSP Utilization	1,699 (76.2%)	1,315 (52.2%)
BRAM Utilization	174 (17.7%)	340 (37.3%)
Power (W)	14.24	7.14
Throughput	86.12 GFLOPS	15.47 GFLOPS
Energy Efficiency	6.05 GFLOPS/W	2.17 GFLOPS/W

higher throughput when the network becomes deeper. Besides, the number of operations of LeNet-10 reported in [36] is 74.43 MFLOPs. However, according to *total number of training operations* = $2 \times (3 \times \sum_{i=1}^{n} M^i \times N^i \times R^i \times C^i \times K^i \times K^i - M^1 \times N^1 \times R^1 \times C^1 \times K^1 \times K^1)$, the actual number of operations that we obtain is only 25.17 MFLOPs. In this formula, 2× is due to the FP 32 MAC involving 2 operations, and 3× is due to the fact that each layer needs to conduct FP, BP, and WU except the first layer, which only needs to conduct FP and WU.

FeCaffe [41] introduced a Caffe framework with OpenCL, which can integrate FPGA to perform CNN network training. It only provided DSP utilization and throughput, which are shown in Table 11. Compared to the FeCaffe framework, our design utilized fewer computation resources but achieved higher throughput implementing AlexNet.

49:34 Y. Tang et al.

	FeCaffe [41]	Ours
Platform	Stratix 10	ZCU102
Frequency (MHz)	253	100
DSP Utilization	1,796 (31.2%)	1,513 (60.0%)
BRAM Utilization	N/A	857 (94.0%)
Power (W)	N/A	7.736
Throughput	≈24 GFLOPS	34.52 GFLOPS
Energy Efficiency	N/A	4.46 GFLOPS/W

Table 11. Experimental Results on AlexNet Compared with FeCaffe [41]

The work in [40] also adopted both feature-map-level parallelism and channel-level parallelism, similar to its preliminary work in [22]. The best nominal energy efficiency reaches 144 (GOPS \times precision), which is lower than our best nominal energy efficiency, which is 194.88 (GOPS \times precision). As for the memory access issues, the accelerator in [40] targeted devices equipped with high bandwidth memory (HBM2). Compared with DMA, the HBM2 is superiorly advanced with 16 pseudo channels providing a high number of I/O data pins. However, HBM2 is a new high-speed memory technology and is only integrated into a few modern FPGAs like Stratix 10 MX. Most FPGA-based edge devices still rely on DMA to communicate between the FPGA chip and off-chip DRAM. Besides, [22] and [40] only tested their designs on the Cifar-10 dataset, where the input image size is only 32×32 , which is really small so that their on-chip BRAMs can easily hold P entire feature maps, where P is the unrolling factors in the channel dimension. However, our design can support both small and large feature map sizes.

DarkFPGA [23] placed DRAM data layout in the **channel-height-width-batch (CHWB)** pattern based on its batch-level parallelism-based design. It achieves higher nominal energy efficiency because the 8-bit fixed points can improve the energy efficiency and DSP efficiency out of proportion. The previous study has shown that if the data precision is no more than 8-bit, two MACs can be calculated on one Xilinx DSP48, reducing the DSP usage by half [42]. However, for 32-bit floating point, one MAC operation takes up five DSPs in the Xilinx FPGA board. Besides, XCVU9P is an extremely high-end cloud-level FPGA that has superior efficiency than commonly used edge FPGAs. However, as mentioned in Section 2.3, the batch-level parallelism adopted by DarkFPGA only achieved high throughput when the batch size is large. From their experiments, when the batch size is below 16, its throughput is below 100 GOPS, which is around 800 GOPS×precision after nominating, while our nominal throughput is stably above 1,000 GOPS×precision among different batch sizes. Besides, as with [22, 40], DarkFPGA also implemented their design on the Cifar-10 dataset with a relatively small feature map size.

7 CONCLUSION

In this article, we design EF-train, an efficient DNN training accelerator enabling edge FPGAs to continuously learn on the device, which makes it possible for current FPGA-based edge-level applications to achieve domain adaption and personalization. We propose an FPGA-based CNN training accelerator with a unified convolution kernel to process FP, BP, and WU with full precision and a data reshaping approach to ensure continuous memory access during end-to-end training processes. We implement end-to-end CNN training effectively for low-power edge devices with restricted resources. The experimental results show that our design achieves 46.99 GFLOPS and 6.09 GFLOPS/W in terms of throughput and energy efficiency, respectively.

REFERENCES

- [1] Scott R. Granter, Andrew H. Beck, and David J. Papke Jr. 2017. AlphaGo, deep learning, and the future of the human microscopist. *Archives of Pathology & Laboratory Medicine* 141 (2017), 619–621.
- [2] Tengchan Zeng, Omid Semiari, Mohammad Mozaffari, Mingzhe Chen, Walid Saad, and Mehdi Bennis. 2020. Federated learning in the sky: Joint power allocation and scheduling with UAV swarms. In 2020 IEEE International Conference on Communications (ICC'20). IEEE, 1–6.
- [3] Cong Hao, Xiaofan Zhang, Yuhong Li, Sitao Huang, Jinjun Xiong, Kyle Rupnow, Wen-mei Hwu, and Deming Chen. 2019. FPGA/DNN co-design: An efficient design methodology for 1ot intelligence on the edge. In 2019 56th ACM/IEEE Design Automation Conference (DAC'19). IEEE, 1–6.
- [4] Xilinx. ([n. d.]). Corazon AI. http://www.xilinx.com/products/boards-and-kits/1-1bua5s3.html.
- [5] Xilinx. ([n. d.]). Pony.ai Sensor Fusion using multiple Xilinx devices. https://www.xilinx.com/applications/automotive/automated-driving.html.
- [6] Xilinx. ([n. d.]). ZF ProAI Gen 3 using Xilinx Zynq UltraScale+ MPSoC. https://www.xilinx.com/applications/automotive/automated-driving.html.
- [7] Ahmed Sanaullah, Chen Yang, Yuri Alexeev, Kazutomo Yoshii, and Martin C. Herbordt. 2018. Real-time data analysis for medical diagnosis using FPGA-accelerated neural networks. *BMC Bioinformatics* 19, 18 (2018), 19–31.
- [8] Alwyn Burger, Chao Qian, Gregor Schiele, and Domenik Helms. 2020. An embedded CNN implementation for ondevice ECG analysis. In 2020 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops'20). IEEE, 1–6.
- [9] Corey Lammie, Alex Olsen, Tony Carrick, and Mostafa Rahimi Azghadi. 2019. Low-power and high-speed deep FPGA inference engines for weed classification at the edge. IEEE Access 7 (2019), 51171–51184.
- [10] Xiaofan Zhang, Cong Hao, Haoming Lu, Jiachen Li, Yuhong Li, Yuchen Fan, Kyle Rupnow, Jinjun Xiong, Thomas Huang, Honghui Shi, et al. 2019. Skynet: A champion model for DAC-SDC on low power object detection. arXiv preprint arXiv:1906.10327 (2019).
- [11] Jianfei Yang, Han Zou, Shuxin Cao, Zhenghua Chen, and Lihua Xie. 2020. Mobileda: Toward edge-domain adaptation. IEEE Internet of Things Journal 7, 8 (2020), 6909–6918.
- [12] Md Abdullah Al Hafiz Khan, Nirmalya Roy, and Archan Misra. 2018. Scaling human activity recognition via deep learning-based domain adaptation. In 2018 IEEE International Conference on Pervasive Computing and Communications (PerCom'18). IEEE, 1–9.
- [13] Fábio Mendonça, Sheikh Shanawaz Mostafa, Fernando Morgado-Dias, and Antonio G. Ravelo-García. 2021. A method based on cardiopulmonary coupling analysis for sleep quality assessment with FPGA implementation. Artificial Intelligence in Medicine 112 (2021), 102019.
- [14] Amrita Rana and Kyung Ki Kim. 2020. Comparison of artificial neural networks for low-power ECG-classification system. *Journal of Sensor Science and Technology* 29, 1 (2020), 19–26.
- [15] Mengwei Xu, Feng Qian, Qiaozhu Mei, Kang Huang, and Xuanzhe Liu. 2018. Deeptype: On-device deep learning for input personalization service with minimal privacy concern. Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies 2 (2018), 1–26.
- [16] Chen Zhang, Peng Li, Guangyu Sun, Yijin Guan, Bingjun Xiao, and Jason Cong. 2015. Optimizing FPGA-based accelerator design for deep convolutional neural networks. In Proceedings of the 2015 ACM/SIGDA International Symposium on Field-programmable Gate Arrays. 161–170.
- [17] Xiaofan Zhang, Junsong Wang, Chao Zhu, Yonghua Lin, Jinjun Xiong, Wen-mei Hwu, and Deming Chen. 2018. DNNbuilder: An automated tool for building high-performance DNN hardware accelerators for FPGAs. In 2018 IEEE/ACM International Conference on Computer-Aided Design (ICCAD'18). IEEE, 1–8.
- [18] Xuechao Wei, Yun Liang, and Jason Cong. 2019. Overcoming data transfer bottlenecks in FPGA-based DNN accelerators via layer conscious memory management. In 2019 56th ACM/IEEE Design Automation Conference (DAC'19). IEEE, 1–6.
- [19] Seungkyu Choi, Jaehyeong Sim, Myeonggu Kang, and Lee-Sup Kim. 2018. 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. 1–6.
- [20] Yudong Tao, Rui Ma, Mei-Ling Shyu, and Shu-Ching Chen. 2020. Challenges in energy-efficient deep neural network training with FPGA. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. 400–401.
- [21] Wenlai Zhao, Haohuan Fu, Wayne Luk, Teng Yu, Shaojun Wang, Bo Feng, Yuchun Ma, and Guangwen Yang. 2016.
 F-CNN: An FPGA-based framework for training convolutional neural networks. In 2016 IEEE 27th International Conference on Application-specific Systems, Architectures and Processors (ASAP'16). IEEE, 107–114.
- [22] Shreyas Kolala Venkataramanaiah, Yufei Ma, Shihui Yin, Eriko Nurvithadhi, Aravind Dasu, Yu Cao, and Jae-sun Seo. 2019. Automatic compiler based FPGA accelerator for CNN training. In 2019 29th International Conference on Field Programmable Logic and Applications (FPL'19). IEEE, 166–172.

49:36 Y. Tang et al.

[23] Cheng Luo, Man-Kit Sit, Hongxiang Fan, Shuanglong Liu, Wayne Luk, and Ce Guo. 2020. Towards efficient deep neural network training by FPGA-based batch-level parallelism. *Journal of Semiconductors* 41, 2 (2020), 022403.

- [24] Weiwen Jiang, Edwin H.-M. Sha, Xinyi Zhang, Lei Yang, Qingfeng Zhuge, Yiyu Shi, and Jingtong Hu. 2019. Achieving super-linear speedup across multi-FPGA for real-time DNN inference. ACM Transactions on Embedded Computing Systems (TECS) 18, 5s (2019), 1–231.
- [25] Sheng-Chun Kao, Geonhwa Jeong, and Tushar Krishna. 2020. Confuciux: Autonomous hardware resource assignment for DNN accelerators using reinforcement learning. In 2020 53rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO'20). IEEE, 622–636.
- [26] Yijin Guan, Hao Liang, Ningyi Xu, Wenqiang Wang, Shaoshuai Shi, Xi Chen, Guangyu Sun, Wei Zhang, and Jason Cong. 2017. FP-DNN: An automated framework for mapping deep neural networks onto FPGAs with RTL-HLS hybrid templates. In 2017 IEEE 25th Annual International Symposium on Field-programmable Custom Computing Machines (FCCM'17). IEEE, 152–159.
- [27] Rachmad Vidya Wicaksana Putra, Muhammad Abdullah Hanif, and Muhammad Shafique. 2021. ROMANet: Fine-grained reuse-driven off-chip memory access management and data organization for deep neural network accelerators. IEEE Transactions on Very Large Scale Integration (VLSI) Systems 29, 4 (2021), 702-715.
- [28] Rachmad Vidya Wicaksana Putra, Muhammad Abdullah Hanif, and Muhammad Shafique. 2020. DRMap: A generic DRAM data mapping policy for energy-efficient processing of convolutional neural networks. In 2020 57th ACM/IEEE Design Automation Conference (DAC'20). IEEE, 1–6.
- [29] Duseok Kang, Donghyun Kang, and Soonhoi Ha. 2021. Multi-bank on-chip memory management techniques for CNN accelerators. *IEEE Transactions on Computers* (2021), 1–1. DOI: 10.1109/TC.2021.3076987
- [30] Chen Zhang, Guangyu Sun, Zhenman Fang, Peipei Zhou, Peichen Pan, and Jason Cong. 2018. Caffeine: Toward uniformed representation and acceleration for deep convolutional neural networks. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems 38, 11 (2018), 2072–2085.
- [31] Atefeh Sohrabizadeh, Jie Wang, and Jason Cong. 2020. End-to-end optimization of deep learning applications. In *Proceedings of the 2020 ACM/SIGDA International Symposium on Field-programmable Gate Arrays.* 133–139.
- [32] Tianqi Wang, Tong Geng, Ang Li, Xi Jin, and Martin Herbordt. 2020. FPDeep: Scalable acceleration of CNN training on deeply-pipelined FPGA clusters. *IEEE Transactions on Computers* 69, 8 (2020), 1143–1158.
- [33] Hiroki Nakahara, Youki Sada, Masayuki Shimoda, Kouki Sayama, Akira Jinguji, and Shimpei Sato. 2019. FPGA-based training accelerator utilizing sparseness of convolutional neural network. In 2019 29th International Conference on Field Programmable Logic and Applications (FPL'19). IEEE, 180–186.
- [34] Sean Fox, Julian Faraone, David Boland, Kees Vissers, and Philip H. W. Leong. 2019. Training deep neural networks in low-precision with high accuracy using FPGAs. In 2019 International Conference on Field-programmable Technology (ICFPT 19). IEEE, 1–9.
- [35] Jinming Lu, Jun Lin, and Zhongfeng Wang. 2020. A reconfigurable DNN training accelerator on FPGA. In 2020 IEEE Workshop on Signal Processing Systems (SiPS'20). IEEE, 1–6.
- [36] Zhiqiang Liu, Yong Dou, Jingfei Jiang, Qiang Wang, and Paul Chow. 2017. An FPGA-based processor for training convolutional neural networks. In 2017 International Conference on Field Programmable Technology (ICFPT'17). IEEE, 207–210.
- [37] Kaiyuan Guo, Lingzhi Sui, Jiantao Qiu, Jincheng Yu, Junbin Wang, Song Yao, Song Han, Yu Wang, and Huazhong Yang. 2017. Angel-eye: A complete design flow for mapping CNN onto embedded FPGA. IEEE Transactions on Computeraided Design of Integrated Circuits and Systems 37, 1 (2017), 35–47.
- [38] Sharan Chetlur, Cliff Woolley, Philippe Vandermersch, Jonathan Cohen, John Tran, Bryan Catanzaro, and Evan Shelhamer. 2014. CUDNN: Efficient primitives for deep learning. arXiv preprint arXiv:1410.0759 (2014).
- [39] OpenVINO. ([n. d.]). Optimization Guide. https://docs.openvino.ai/2020.2/_docs_optimization_guide_dldt_optimization_guide.html.
- [40] Shreyas K. Venkataramanaiah, Han-Sok Suh, Shihui Yin, Eriko Nurvitadhi, Aravind Dasu, Yu Cao, and Jae-sun Seo. 2020. FPGA-based low-batch training accelerator for modern CNNs featuring high bandwidth memory. In Proceedings of the 39th International Conference on Computer-aided Design. 1–8.
- [41] Ke He, Bo Liu, Yu Zhang, Andrew Ling, and Dian Gu. 2020. FeCaffe: FPGA-enabled Caffe with OpenCL for deep learning training and inference on Intel Stratix 10. In *Proceedings of the 2020 ACM/SIGDA International Symposium on Field-programmable Gate Arrays*. 314–314.
- [42] Yuhong Li, Cong Hao, Xiaofan Zhang, Xinheng Liu, Yao Chen, Jinjun Xiong, Wen-mei Hwu, and Deming Chen. 2020.
 AEdd: Efficient differentiable DNN architecture and implementation co-search for embedded ai solutions. In 2020 57th ACM/IEEE Design Automation Conference (DAC'20). IEEE, 1–6.

Received June 2021; revised December 2021; accepted December 2021