

# LayerPipe: Accelerating Deep Neural Network Training by Intra-Layer and Inter-Layer Gradient Pipelining and Multiprocessor Scheduling

Nanda K. Unnikrishnan and Keshab K. Parhi

Dept. Electrical and Computer Engineering, University of Minnesota

**Abstract**—The time required for training the neural networks increases with size, complexity, and depth. Training model parameters by back-propagation inherently creates feedback loops. These loops hinder efficient pipelining and scheduling of the tasks within the layer and between consecutive layers. Prior approaches, such as PipeDream, have exploited the use of *delayed gradient* to achieve inter-layer pipelining. However, these approaches treat the entire backpropagation as a single task; this leads to an increase in computation time and processor underutilization. This paper presents novel optimization approaches where the gradient computations with respect to the weights and the activation functions are considered independently; therefore, these can be computed in parallel. This is referred to as *intra-layer* optimization. Additionally, the gradient computation with respect to the activation function is further divided into two parts and distributed to two consecutive layers. This leads to balanced scheduling where the computation time of each layer is the same. This is referred to as *inter-layer* optimization. The proposed system, referred to as *LayerPipe*, reduces the number of clock cycles required for training while maximizing processor utilization with minimal inter-processor communication overhead. *LayerPipe* achieves an average speedup of 25% and upwards of 80% with 7 to 9 processors with less communication overhead when compared to PipeDream.

## I. INTRODUCTION

Deep neural networks (DNNs) are omnipresent in our daily lives due to their ability to solve a wide range of complex real-world problems. DNNs form the backbone for many tasks such as image recognition, language translation, autonomous driving, and recommendation systems [1]–[6]. The DNN models are trained once and repeatedly used for inference. However, because of the large computations associated with training, the computation time of the training is orders of magnitude larger than that of the inference. The massive surge in data center workloads that involve deep learning has led to new devices such as Google’s TPU [7], [8], NVidia’s Tesla [9], or Xilinx Alveo [10] in addition to custom accelerators [11]–[15].

Prior works have optimized training using parallel computations such as data parallelism [16] which replicates the DNN model across processors, splits the data into multiple smaller batches, and distributes the computational workload of the model across different processors. However, data parallelization suffers from significant inter-processor communication overhead. In addition, with larger and more complex models with model parameters requiring storage in the range of GigaBytes [17], it becomes challenging to store the model as a whole in a single processor. To overcome the issues with data parallelization, the use of *delayed gradients* [18] have been exploited to achieve pipelining in [19], [20] as a viable alternative. However, while these techniques effectively improve training times, they do not achieve their maximum potential due to pipeline imbalance issues and processor underutilization. This has resulted in pipeline parallelization techniques resorting back to data parallelization [20]

This research was supported in part by the National Science Foundation under grant number CCF-1954749.

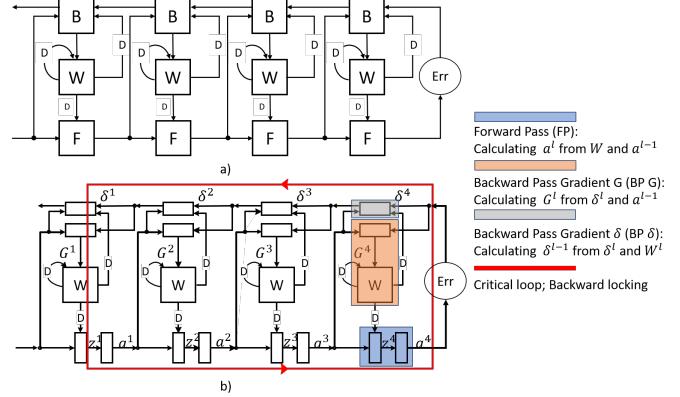


Fig. 1. A sample four-layer network showing the individual operation in the forward and backward passes. a) Conventional abstraction into 3 operations: Forward, Backward and Weight update. b) Detailed view of different operations in forward and backward. The critical loop of the network is highlighted.

for pipeline balancing. In this work, we propose novel *intra-layer* and *inter-layer* optimization techniques to achieve maximum processor utilization with minimal inter-processor communication overhead.

Fig. 1(a) shows a sample four-layer network highlighting fundamental operations of layers in the forward (F) and backward (B) passes during training. These layers can be either convolutional layers or fully-connected (FC) layers. During training, the system computes the forward pass to generate the outputs, calculates the error using ground truth, and then propagates the error backwards through the network. Fig. 1(b) highlights the critical loop during training which is a severe bottleneck in the computation time of DNN training. Unlike [19], [20] that treat the backward computations as a single computation block as shown in Fig. 1(a), we split the backward pass into its fundamental operations, which involve the calculation of gradients of the error with respect to weight ( $G$ ) and with respect to the activation function ( $\delta$ ) as shown in Fig. 1(b). These two gradients can be computed in parallel using different processors. This is referred to as *intra-layer* parallelism. Additionally, the computation of  $\delta$  is further divided into two parts which are distributed to two consecutive layers. This is referred to as *inter-layer* optimization.

This paper makes three key contributions listed below.

- 1) A formal derivation of pipeline models in PipeDream [20] using concepts such as pipelining, retiming, and *delayed gradients*.
- 2) A novel intra-layer optimization technique where the gradients of the error with respect to the weights and activation functions can be computed in parallel, reducing computation time and increasing processor utilization efficiency.
- 3) A novel fine-grain inter-layer pipelining by dividing the gra-

dient computation of the error with respect to the activation function into two parts and distributing the two parts to two consecutive layers. This leads to a balanced inter-layer pipeline and reduction of total system latency with negligible inter-processor communication overhead.

The rest of this paper is organized as follows: Section II presents fundamental concepts of backpropagation, data, and pipeline parallelism. Section III describes the LayerPipe framework for efficient pipeline parallelism using proposed intra-layer and inter-layer optimizations. Finally, section IV evaluates and compares LayerPipe against PipeDream, a state-of-the-art pipelined accelerator architecture for training DNNs [20].

## II. BACKGROUND AND PRIOR WORK

To optimize training time, we examine the operations within the backpropagation algorithm.

### A. Backpropagation

We deconstruct the backpropagation algorithm into its primary operations to better understand how it can be implemented and optimized. The training loop for any supervised learning problem has two parts: a forward pass or *inference* and a backward pass for training. Fig. 1(b) illustrates the data-flow graph of these operations on a sample four-layer neural network. The lower half of the data-flow graph shows the forward pass computations, while the upper half shows the backward pass computations. As shown in Fig. 1 there exist multiple nested feedback loops in the network. This is the main reason why system-level techniques such as pipelining are not straightforward, as delays cannot be introduced into a feedback loop system without affecting the output. The forward and backward operations are summarized by Eqs. (1) to (5). In the forward pass, the weights  $\mathbf{W}^l$  and the activation output,  $\mathbf{a}^{l-1}$ , from the preceding layer are used to compute  $\mathbf{z}^l$  and  $\mathbf{a}^l$ . The activation function  $f()$  refers to non-linear activation functions such as ReLU, sigmoid or tanh.  $\mathbf{W}^l$  represent the weights or filters, also referred to as model parameters, associated with the layer  $l$ ,  $\mathbf{z}^l$  is the output of the convolution operation in the convolutional layer or linear operation in the FC layer. The output  $\mathbf{a}^l$  represents the activation output at layer  $l$ . The backward pass consists of two operations:  $\mathbf{G}$  calculation, and  $\delta$  calculation.  $\mathbf{G}^l$  is the gradient of the error function,  $E$ , with respect to the weights  $\mathbf{W}^l$  at layer  $l$ , and it is computed from  $\delta^l$  and  $\mathbf{a}^{l-1}$ .  $\delta^l$  is the gradient of the error function with respect to the activations of the layer  $l$  and is backpropagated to the previous layer. The notation  $\odot$  represents the Hadamard product of matrices.

$$\mathbf{z}^{(l)} = \mathbf{W}^{(l)} \mathbf{a}^{(l-1)} \quad (1)$$

$$\mathbf{a}^{(l)} = f(\mathbf{z}^{(l)}) \quad (2)$$

$$\delta^{(l-1)} = \frac{\partial E}{\partial \mathbf{a}^{(l-1)}} \odot f'(\mathbf{z}^{(l-1)}) \quad (3)$$

$$\text{where } \frac{\partial E}{\partial \mathbf{a}^{(l-1)}} = (\mathbf{W}^{(l)T} \delta^{(l)}) \quad (4)$$

$$\mathbf{G}^{(l)} = \left( \frac{\partial E}{\partial \mathbf{W}^{(l)}} \right) = \delta^{(l)} \mathbf{a}^{(l-1)T} \quad (5)$$

Training neural networks are characterized by the presence of several large feedback loops in the design. For example, literature has identified three fundamental *lockings* for the backpropagation algorithm [21]. First is forward locking, where a layer cannot be executed unless all the previous layers in the directed forward graph are executed. Second, update locking, a layer cannot be updated until

all dependent operations have been executed in the forward pass. Last, backward locking, a layer cannot be updated before all the dependent operations are executed in forward and backward passes. Among the three, the backward locking problem has been identified as a severe bottleneck for training and therefore has been the focus of several efforts [19]–[25].

### B. Prior work on parallelism

Given that training times have been increasing with network size and complexity, several methods have been proposed to parallelize training. These techniques can be broadly classified as intra-batch parallelism, and inter-batch parallelism [20]. Intra-batch parallelism is the most common form of parallelism currently deployed to accelerate training [16]. In intra-batch parallelism, a single iteration of training is split across available processors. Intra-batch techniques can be further classified into two categories: *data parallelism* and *model parallelism* [16], [26]. The first approach, data parallelism, distributes the workload by replicating the model and splitting the input mini-batch among different processors. Each model trains independently, and the gradients are updated several times before the weights. In the second approach, model parallelism, a large model is split into multiple processors [1]. Furthermore, model parallelism can split the computations along different dimensions such as channel (C), width (W), or height (H). Lastly, a subcategory of model parallelism is *layer parallelism* [21], [22] where the computation is split along the depth of the architecture. Model parallel approaches limit the size of the network to be stored in the processor; however, they incur large communication overheads for transferring the intermediate results, a limiting factor in systems that are communication-bounded.

Recently, pipeline parallelism has emerged as a popular technique for speeding up backpropagation [20]. Inter-batch pipeline parallelism processes multiple mini-batches of the data in parallel over multiple iterations of the backpropagation algorithm. Intra-batch pipeline parallelism [19], however, splits a mini-batch into smaller micro-batches to achieve the same goal. Inter-batch pipeline parallelism is advantageous as it avoids low processor utilization by frequent synchronization. However, significant drawbacks of pipeline parallel designs include large communication overheads and difficulty balancing the workload between processors due to coarse-grained pipelining. PipeDream, for example, uses dynamic programming to find the optimal workload partition on a per-layer basis. Further works have attempted to address the workload balancing issues with fine-grain pipelining and allocation [23], [24]. However, while relying on precedence graphs to derive the schedule, existing approaches do not fully exploit the precedence constraints to derive partitions that minimize communication overhead.

## III. LAYERPIPE: THEORY AND ALGORITHM

The limitation of existing pipelined parallel algorithms in DNNs is that they do not tackle the problem from the fundamental level. This may lead to inefficient algorithms and inaccuracies as it does not account for all critical factors. However, recent improvements to tackle this problem, such as weight stashing and activation recomputation, are often rooted in rudimentary intuition and are not based on any formal theory of such systems. The following section identifies the need to formally derive these pipelined models and set up a systematic approach. From Fig. 1, we can observe the parallels with traditional signal processing architectures, allowing us to exploit the existing well-established architecture optimizations [27].

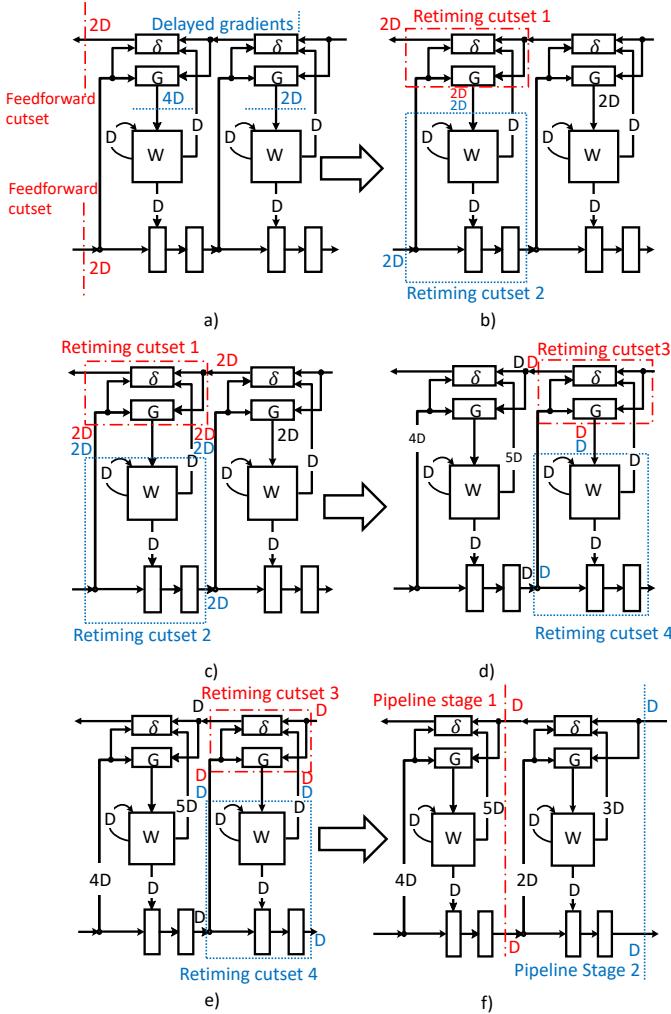


Fig. 2. Step-by-step derivation (a– f) of pipeline stages in a 2-layer design. a) Insertion of delays at feedforward cutsets and adding delayed gradients. b) Defining retiming cutsets 1 and 2 for the first stage. c) Pushing the delays through the retiming cutsets. d) Defining retiming cutsets 3 and 4 for stage 2. e) Pushing the delays through the retiming cutsets. f) Final pipelined DFG with two pipelined stages.

#### A. Pipelining and retiming the backpropagation algorithm using delayed gradients: intra-layer optimization

Prior work on this topic has primarily designed the system with an intuition of how the variables should behave when pipeline stages are added. For example, weight stashing [20] or activation stashing is the concept where the weights or activations are stored locally within a processor during the forward pass till the corresponding data reaches the same processor in the backward pass. In some instances, if the forward pass is computationally cheap, the activations are recomputed [19]. Using the architecture optimization techniques from DSP systems [27], we define two processes that must be followed to derive pipelined models: i) Locations in a system where delays can be legally inserted; ii) Necessary conditions for moving delays between edges on the data-flow graph (DFG). For the derivation, we use a delay element to represent *inter-iteration dependency* [28]. Here the delay element does not refer to a physical delay but a conceptual delay. The delay holds the data for an iteration or the transition or pipeline stage that transfers data to the adjacent processor.

In the first process, there are only two legal locations for delay

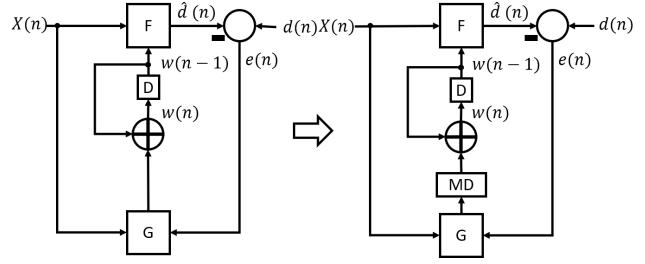


Fig. 3. Delay insertion in DLMS algorithm [18].

placement within the system. Delays may be inserted through *pipelining* on all feedforward cutsets; a cutset is a line that divides a graph into two subgraphs. A feedforward cutset is one where all the edges along the cutset are in the same direction. In the case of neural networks, the only feedforward cutsets are at the inputs and outputs of the network as highlighted in Fig. 2. Delays cannot be inserted in non-feedforward cutsets barring some exceptions. One such exception is the case of slowly varying weights in neural networks analogous to the delayed least mean squares (DLMS) algorithm [18]. The weight update step may use an older version of the gradients in gradient descent algorithms like the DLMS algorithm. It is assumed that due to the small step size, the change in the weight with each iteration is gradual. Thus using an older or *stale* version of the weights or gradients will not significantly affect the convergence characteristics of the learning process. This idea forms the basis of PipeDream. We point out its equivalence with DLMS algorithm that has been widely analyzed and applied in adaptive filter applications. The weight update equation for a DLMS based approach is given by:

$$\mathbf{W}(n) = \mathbf{W}(n-1) - \eta \mathbf{G}(n-M) \quad (6)$$

where  $\mathbf{W}(n)$  is the weight parameter as sample  $n$ ,  $\mathbf{G}$  is the gradient of the error with respect to the weight parameter at sample  $n$ ,  $M$  corresponds to the degree of staleness,  $\eta$  is the step size or learning rate, and  $n$  is the current iteration. The  $F$  block computes the forward filter and the  $G$  block computes the gradient using the input  $X(n)$  and the error  $e(n)$ . The error is the difference between the output and the ground truth (the desired value  $d(n)$ ). This leads to the second legal location for placement of delays, between the gradient calculation and weight update states using *delayed gradients*. With the addition of  $M$  delay elements within the feedback loop in Eq. (6), the system can use these elements for pipelining. As there is no synchronization step in every block, DLMS does not suffer from any overheads for the weight update operation. This process of using delayed gradients is shown in Fig. 3, where MD refers to the  $M$  delays in Eq. (6).

The second process uses *retiming* [29] to move delays in the DFG to the desired location. One delay can be inserted in each outward edge of a feedback cutset during retiming if one delay is removed from each inward edge and vice-versa. Feedback cutsets are cutsets with at least one edge in both directions. At a node level, this could be the operation's outputs and inputs. Thus we can use the two processes defined above to illustrate a step-by-step process for deriving a pipelined model as shown in Fig. 2. Two intermediate layers of the network are chosen as a representative set for the entire network. The aim is to eventually insert a pipeline stage after the forward pass of the layer and before the backward pass of the layer.

The first step in the process is to insert the required number of delays in the network as shown in Fig. 2(a). At each feedforward

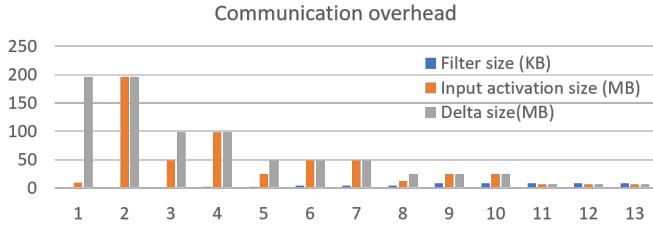


Fig. 4. Layer-wise communication overhead for VGG16.

cutset, the total number of delays added via pipelining is the number of pipeline stages required, two in this example. At the location of delayed gradients, the number of delays to be added is twice the number of pipeline stages (or layers) after the current layer, i.e., four delays and two delays for layers 1 and 2, respectively, in this example.

The second step uses retiming to push the required number of delays to the first pipeline stage, as shown in Figs. 2(b) and (c). We define two retiming cutsets, 1 and 2, to retime the delays highlighted in red and blue. Retiming cutset 1 pushes two delays from each outgoing edge into each incoming edge. Similarly, retiming cutset 2 pushes two delays from each incoming edge into each outgoing edge.

Finally, in the third step, we again use retiming to move the required number of delays to the second pipeline stage, as shown in Figs. 2(d) and (e). We define two retiming cutsets, 3 and 4, used to retime the delays highlighted in red and blue. Retiming cutsets 3 and 4 are identical to cutsets 1 and 2 except that these shift one delay instead of two. Using this, we can generalize retiming in a layer using cutsets like 1 and 2, where each cutset shifts  $N$  delays, where  $N$  is the number of pipeline stages after the layer.

The final pipelined DFG is shown in Fig. 2(f) with its two pipelined stages and associated delay elements. The consequence of retiming was the insertion of delays between the  $W$  and  $\delta$  operation. Similarly, delay elements are inserted between  $a$  and the backward pass calculations. These delays correspond to the weight stashing technique used by PipeDream and the activation stashing/recomputation as shown in PipeDream/Gpipe. Thus retiming, pipelining, and delayed gradients can be used to derive any pipelined model formally. Additionally, this novel method can precisely determine how many cycles the intermediate values should be stored for multiple iterations, a result that prior works had arrived at with intuition.

#### B. Pipeline balancing with parallelism: inter-layer optimization

With a formal derivation of the pipelined model, the next task is to balance the workload across multiple operations. One of the critical challenges for pipelined models is that the network architecture is not uniform, and the computation requirements across layers can vary significantly. Thus any distribution of layers across multiple processors will inevitably be imbalanced. The drawback of prior approaches is that they often lump the entire backward pass into a single operation [19], [20]. Furthermore, even when fine-grain pipelining is applied [23], [24], these approaches do not exploit the characteristics of the data-flow graph (DFG) to find optimal distribution strategies.

An examination of the data-flow graph shows that the cost to move an operation to the adjacent processor is not the same for all backpropagation operations. For example, Fig. 4 shows the communication overheads associated with moving variables between different processors in the convolutional layers of VGG16. It is observed that moving  $\delta$  or  $a$  between processors leads to significant communication overhead in the order of MBs. However, communicating the filters

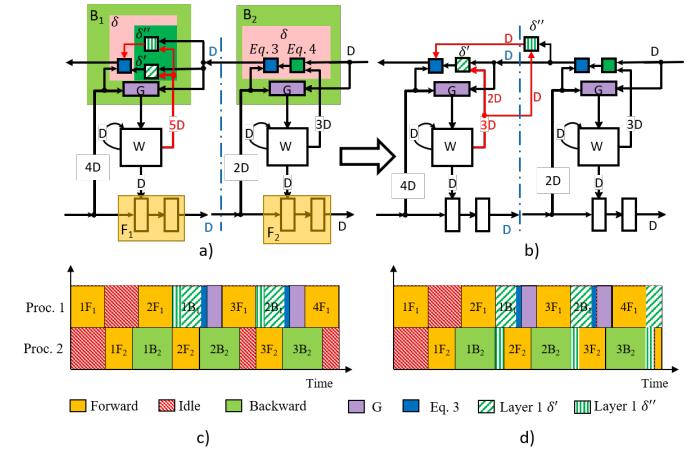


Fig. 5. Block diagram of operation parallelism and its impact on scheduling. a) Original DFG with each operation mapped to its corresponding layer processor. b) Proposed DFG where  $\delta''$  computation is moved to the adjacent processor. c) Imbalanced processor workloads in the original DFG. d) Balanced processor workloads due to the proposed DFG. The notation  $nF_p$  indicates the forward computations for minibatch  $n$  that are assigned to Processor  $p$ . Similarly,  $nB_p$  follows the same notation for the backward computations.  $nB_p$  can be further subdivided into  $nB_pG$ ,  $nB_p\delta'$ , and  $nB_p\delta''$ .

is insignificant as this transfer is in the order of KBs. In pipelined designs, at the pipeline boundaries, variables  $a$  and  $\delta$  have to be transferred between the processors, leading to a mandatory overhead. Any load balancing attempt will require the existing mandatory overhead, in addition to the communication required to transfer its inputs and outputs to the layers. Analyzing the backward pass in Fig. 1, we can observe the input dependencies of all the operations.  $G^l$  computation at layer  $l$  requires  $a^{l-1}$  and  $\delta^l$ , and  $\delta$  as shown in Eq. (5). Specifically Eq. (4) requires  $\delta^l$  and  $W^l$ . Note that  $\delta^l$  is computed in layer  $l+1$ . Using the example of a 2-layer network mapped to 2 processors, we explore the feasibility of shifting the computation of the backward pass of the central processor to its adjacent processors. Shifting the  $G$  computation of processor  $k$  to processor  $k+1$  will incur significant overhead as  $a_k$ , activation output from processor  $k$ , will have to be broadcast to both processors  $k$  and  $k+1$ . Similarly, shifting to processor  $k-1$  is infeasible due to the overhead of  $\delta_{k+1}$ ,  $\delta$  output from  $(k+1)$ th processor. Shifting the computation of  $\delta_k$  to processor  $k+1$  will place the computation in the same processor as its input  $a_{k+1}$ . Also, communication of  $G$  between the processors is insignificant compared to the mandatory overheads. Furthermore, the computation result is not consumed within processor  $k$  and can be directly forwarded to the destination. Thus  $\delta$  is a prime candidate to distribute with its adjacent processor at  $k+1$ .

Fig. 5 summarizes the steps required to divide  $\delta$  and transfer it to the adjacent processor. In the original DFG from Fig. 1,  $\delta$  computations are split into 3-parts.  $\delta'$  is the portion of the computation that remains within layer  $l$ .  $\delta''$  is the portion of the computation that is shifted to layer  $l+1$ . In essence, layer  $l+1$  borrows as much computation from layer  $l$  as needed to balance the computations of consecutive processors. As  $\delta$  is parallelizable in the input and output channel dimensions, it would be simple to parallelize the operation and split it between processors. This technique of dividing the  $\delta$  and merging it with the adjacent layer enables *inter-layer* optimization. This approach is used to derive a load-balanced DFG as shown in Fig. 2(b). In this example, the total completion time of processor 1 is much longer than processor 2, which leads to underutilization and

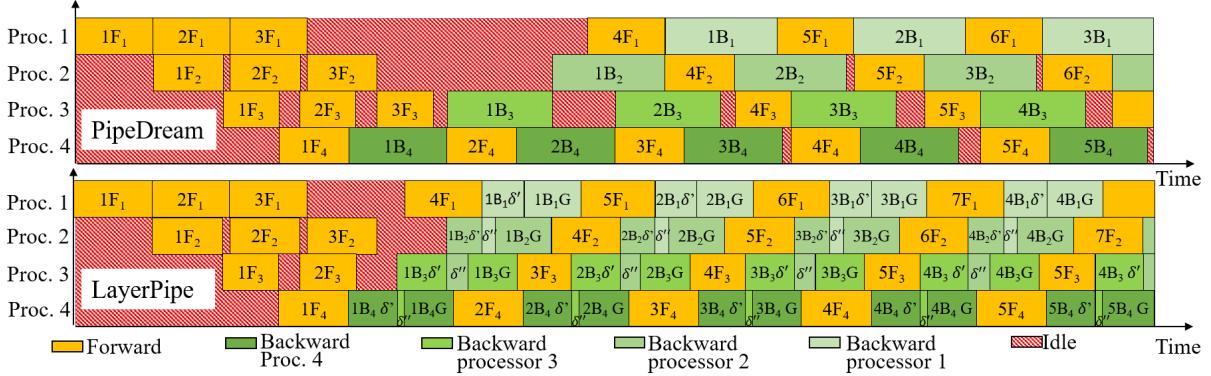


Fig. 6. Contrast of scheduling approaches for PipeDream and LayerPipe.

idle time for processors. In Fig. 5(b), we use retiming once again to place the delays such that  $\delta'$  and  $\delta''$  are computed in adjacent processors. The current depiction highlights simple layer connections; however, there can be multiple ways these brain-inspired neural networks can be connected [30], including divergent or convergent paths. In examples such as ResNet or U-net [31] with convergent paths for  $\delta$ , an additional summation step is required before the result can be used; thus,  $\delta''$  would only be moved to the processor that computes this step.

Figs. 5(c) and (d) represent the comparison between a layer-based coarse and fine-grained layer parallelism. In Fig. 5(c), the difference in computation times between the processors is approximately a third of the computation time of  $\delta$ . Therefore,  $\delta$  is partitioned such that  $\delta'$  computes two-thirds of the output channels and  $\delta''$  computes one-third of the output channels. The  $W$  parameters required for the operation are transferred to Processor 2, and  $\delta''$  is computed in Processor 2. Thus we observe that Fig. 5(d) has a perfectly balanced pipeline, leading to an increase in throughput and reduction in latency.

### C. Scheduling and partitioning algorithms

The heart of the problem is designing an algorithm to process a data-flow graph (DFG) and generate a schedule in a processor-constrained environment. The algorithm is designed to augment or improve upon existing parallelization techniques like PipeDream and GPipe by providing a more theoretical basis for partitioning the network. The algorithm serves two purposes. First, coarse-grained partitioning based on layers and inter-layer pipelining maximize throughput and minimize communication overhead. Second, a fine-grained schedule based on precedence graph and critical path minimizes latency.

The proposed layer partitioning scheme can be described by Algorithm 1. The proposed algorithm leverages the techniques in the MARS algorithm [32] to schedule feedback loops. The first step in the process is to evaluate and find all layers ( $L$ ) and all the critical loops in the DFG (line 1). The critical loop, along with its path, can be found in  $O(de_d)$  time using the minimum cycle mean (MCM) algorithm [33], where  $d$  is the number of delays in the DFG and  $e_d$  are all edges between the delays. In the second step, we run a profiler over the network and calculate the computation times, communication overhead, and memory overhead of each layer in the network. We then classify computations as movable or immovable based on the constraints on communication and memory (line 4). If a computation's communication overhead and overhead memory fall below a threshold, we classify that computation as movable;

### Algorithm 1 Partitioning algorithm for balanced pipeline generation.

**Input:** DFG of DNN, #processors ( $N_p$ ), processors

**Output:** Critical loops  $C_l$ , processor allocation  $P_a$

```

1: //Step 1: Find all  $C_{loops}$  and DNN layers  $L$  in the DFG
2:  $C_l$ ,  $L$  = find_critical_loops(DFG)
3: //Step 2 starts here
4: //Profile: For each DNN layer  $l$  in DFG find layer compute time  $t_c$ , fixed computation
5:  $t_{fix}$ , and flexible computation time  $t_{flex}$ . Store in  $T_c$ ,  $T_{fix}$ , and  $T_{flex}$ 
6:  $T_{tot}$  = for  $i$  in  $T_c[i]$  do sum( $T_l[i]$ )
7: //Find maximum processor time  $T_p$ 
8:  $T_p = \frac{T_{tot}}{N_p}$ 
9: //Step 3 starts here
10: //Initialize flag, processor index  $p_{idx}$  to 0 and processor idle time  $T_{idle}$  to  $T_p$ 
11: flag = 0
12: while flag = 0 do
13:    $p_{idx} = 0$ ;  $T_{idle} = T_p$ 
14:   for each  $l$  in reversed( $L$ ) do
15:     if  $T_{flex}[l] < T_{idle}$  then
16:       allocate  $T_{flex}[l]$  to processors[ $p_{idx}$ ] and update  $P_a$ 
17:        $T_{idle} = T_{idle} - T_{flex}[l]$ 
18:     else
19:       //Partition  $T_{flex}[l]$  ( $\delta$ ) to  $\delta'$  and  $\delta''$  with operational parallelism (OP)
20:        $\delta', \delta'' = OP(T_{idle}, T_{flex}[l])$ 
21:       allocate  $\delta''$  to processors[ $p_{idx}$ ] and update  $P_a$ 
22:        $p_{idx} = p_{idx} + 1$ 
23:       allocate  $\delta'$  to processors[ $p_{idx}$ ] and update  $P_a$ 
24:        $T_{idle} = T_p - \delta'$ 
25:     end if
26:     if  $T_{fix}[l] > T_{idle}$  then
27:        $p_{idx} = p_{idx} + 1$ 
28:     end if
29:     allocate  $T_{fix}[l]$  to processors[ $p_{idx}$ ] and update  $P_a$ 
30:      $T_{idle} = T_{idle} - T_{fix}[l]$ 
31:   end for
32:   if  $p_{idx} > N_p$  then
33:     //Relax the max processor time constraint and flags stays 0
34:      $T_p = \alpha \times T_p$ 
35:   else
36:     flag = 1
37:   end if
38: end while
39: return  $L_c$ ,  $P_a$ 

```

otherwise, it is immovable. We return time taken for all movable computations as  $T_{flex}$  and immovable computations as  $T_{fix}$ . For simplicity, we assume only movement from layer  $l$  to layer  $l + 1$  is allowed. Taking the total computation time of the network ( $T_{tot}$ ), we can determine the target workload of each processor or the maximum processor time ( $T_p$ ) for the required number of processors ( $N_p$ ) in lines 6 and 8 of Algorithm 1.

In the third step of Algorithm 1 (line 9), we iterate through all the layers of the network in reverse and try to allocate it to processors as follows. As the flexible portion,  $T_{flex}[l]$ , can only be

TABLE I  
NETWORK ARCHITECTURE FOR SAMPLE 4 CONVOLUTION LAYER DESIGN.

Layer	Filter size	Input channels	Output channels	Padding	Stride
1	5	3	32	2	2
2	5	32	64	2	2
3	3	64	128	1	2
4	3	128	128	1	1

TABLE II  
SUMMARY OF COMPUTATION TIMES AND COMMUNICATION OVERHEADS FOR THE SAMPLE 4-LAYER NETWORK.

Layer	Computation time (cycles)			
	1	2	3	4
FP	$1.20 \times 10^6$	$5.02 \times 10^6$	$1.81 \times 10^6$	$3.63 \times 10^6$
BP_G	$2.16 \times 10^6$	$5.63 \times 10^6$	$2.11 \times 10^6$	$3.92 \times 10^6$
BP $_{\delta}$	$4.01 \times 10^7$	$2.01 \times 10^7$	$7.23 \times 10^6$	$3.63 \times 10^6$
Total	$4.35 \times 10^7$	$3.07 \times 10^7$	$1.12 \times 10^7$	$1.12 \times 10^7$
Communication overhead				
FP Overhead	12.69MB	6.57MB	3.29MB	3.52MB
BP Overhead	4.59MB	12.25MB	6.13MB	3.06MB
Additional Overhead	0.07KB	0.78KB	0.56KB	1.13KB

allocated to layer  $l + 1$ , this is allocated first. If the computation can be accommodated entirely in the current processor  $p_{idx}$ , i.e.,  $T_{flex}[l] < T_{idle}$  we allocate this computation to the current processor,  $processors[p_{idx}]$  (line 16). If this is not possible, we use operational parallelism to partition  $\delta$  ( $T_{flex}[l]$ ) into  $\delta'$  and  $\delta''$  such that  $\delta''$  fits within the remaining workload available in the processor. We then move to the next processor and assign the remaining computation  $\delta'$ . We then try to allocate the fixed portion of the computation in a similar manner. If the computation can be accommodated entirely in the current processor  $p_{idx}$ , i.e.,  $T_{flex}[l] < T_{idle}$ , we allocate this computation to the current processor with index  $p_{idx}$ . If this is not possible, we move to the next processor and assign the computation (lines 26–30). This process is repeated until all layers of the network have been assigned. If more processors are needed than those available, we relax the target workload requirements of each processor by a factor of  $\alpha$  (line 34) to allow for longer compute times and restart at step 3. This process is repeated until the number of processors assigned matches the number of processors targeted.

Fig. 6 summarizes the differences between a traditional pipeline parallel schedule like PipeDream and a balanced, fine-grained pipeline schedule like LayerPipe. In the first stage, we derived the allocations of each processor, including the suggested operational parallelism from Algorithm 1. Using the information from the critical loops, we derive a schedule that prioritizes computations along the critical loop. Note that the critical loops generally run through the  $\delta$  computation (Fig. 1) and the  $G$  computation does not appear in the critical loop except in the first layer. Furthermore, as the only dependence of  $\delta^{l-1}$  in layer  $l$  from layer  $l + 1$  is  $\delta^l$  we need not wait for the  $G$  computation from layer  $l + 1$  to complete before starting layer  $l$ . Therefore, we can derive a schedule that prioritizes the  $\delta$  computations over  $G$  computations allowing for fine-grained pipelining that reduces the overall latency of the system. When operational parallelism is active, the two partitions  $\delta'$  and  $\delta''$  are independent; therefore, these can be computed in parallel, further reducing the system latency.

#### IV. EXPERIMENTAL EVALUATION

##### A. Methodology

In order to test the effectiveness of the proposed system, we benchmark the performance of LayerPipe against the standard pipeline parallelism algorithm, PipeDream [20]. The system simulated consists of multiple Processors that can communicate intermediate results

TABLE III  
COMPARISON OF COMPUTATION TIMES IN CYCLES OF PIPELINE PARALLELISM ALGORITHMS FOR THE SAMPLE FOUR-LAYER NETWORK.

Processors	1	2	3
PipeDream	$4.35 \times 10^7$	$3.07 \times 10^7$	$2.23 \times 10^7$
LayerPipe			
Assigned computation	$3.22 \times 10^7$	$2.09 \times 10^7$	$2.23 \times 10^7$
Borrowed computation	$0.00 \times 10^0$	$1.13 \times 10^7$	$9.86 \times 10^6$
Total	$3.22 \times 10^7$	$3.22 \times 10^7$	$3.22 \times 10^7$

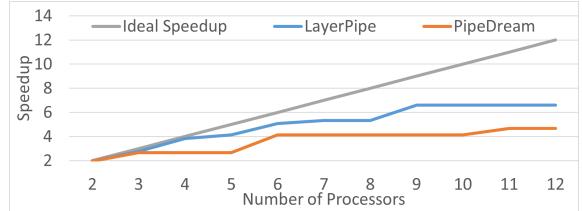


Fig. 7. Comparison of speedups between LayerPipe and PipeDream with different number of processors on the convolutional layers of VGG16.

among themselves without any external memory. The algorithm is evaluated by varying the number of processors while balancing the pipelines. Each processor consists of a single systolic array within a TPU or neural processing unit inside a GPU. For this experiment, it is assumed the systolic array is operated in a weight-stationary data-flow [34]; however, the same techniques are applicable for other data-flows. In order to model the systolic array, we developed a python simulator based on the SCALESim [35] library to estimate the computation times and communication requirements of the systolic array. The new simulator was validated against SCALESim, a cycle-accurate systolic array simulator verified against RTL simulations.

To account for communication overheads, the simulator keeps track of two kinds of overheads. First, mandatory overheads, such as communication of activations and  $\delta$  between processors, are required irrespective of the algorithm for any pipeline parallel design. Second, the additional overheads introduced by the LayerPipe algorithm are calculated. These overheads account for additional inputs that processors must communicate to support the LayerPipe algorithm, such as layer weights.

The simulator tests whether the proposed algorithm is hardware agnostic by varying the systolic array size and batch size. The results are then averaged to obtain the final network performance results.

TABLE IV  
COMPARISON BETWEEN THE THEORETICAL SPEEDUP OF LAYERPIPE VERSUS PIPEDREAM FOR THE CONVOLUTION LAYER OF VGG16. THE RESULTS ARE AVERAGED ACROSS SYSTOLIC ARRAY AND BATCH SIZES.

Processor	LayerPipe	PipeDream	Improvement	Communication Overhead
2	1.93	1.80	7.2%	0KB
3	2.75	2.32	18.4%	0KB
4	3.59	2.73	31.7%	0KB
5	4.35	3.10	40.2%	0.26KB
6	5.01	3.10	61.8%	0.01KB
7	5.55	3.10	79.0%	0.01KB
8	5.69	3.10	83.5%	0.01KB
9	6.24	3.47	80.0%	2.73KB
10	6.69	4.37	52.9%	2.61KB
11	6.88	6.13	12.4%	1.49KB
12	6.88	6.13	12.4%	0KB

##### B. Sample four-layer network

The proposed algorithm is first evaluated on a simple network consisting of four convolutional layers of different sizes. This provides a detailed demonstration of how the workload is scheduled between processors leading to a balanced schedule. The parameters of the

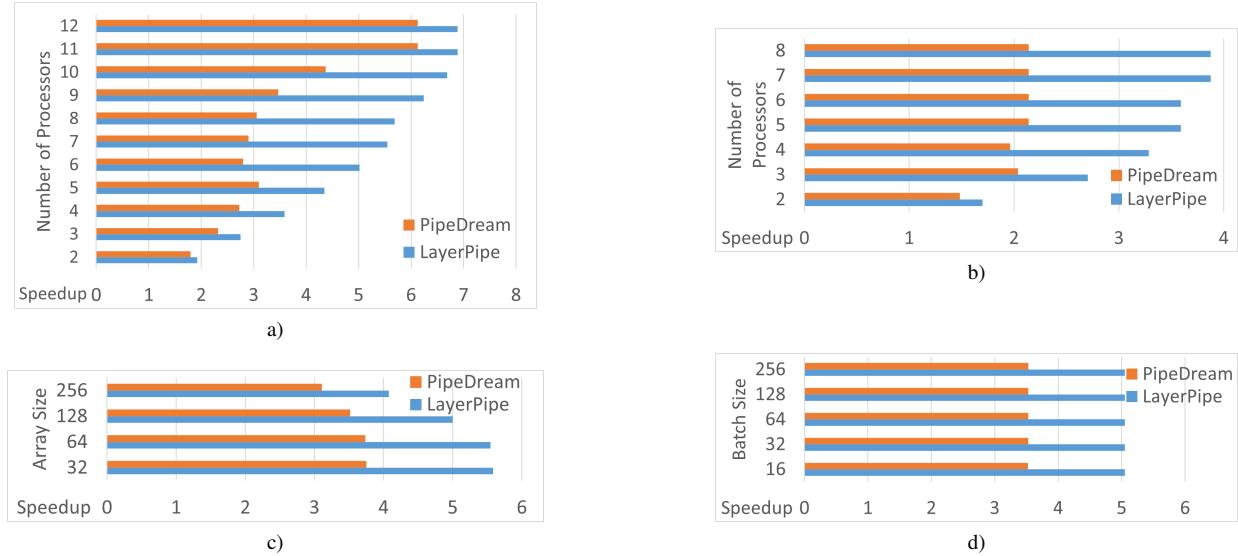


Fig. 8. Performance comparison between LayerPipe and PipeDream for different number of processors on the convolutional layers of a) VGG16 b) ResNet50. The results were averaged over various batch sizes, systolic array sizes, and normalized to a single processor's performance. The individual distributions for VGG16 is shown for c) array size and d) batch size.

architecture of the sample network in Fig. 1 are described in Table I. Here the filter size, input channels, output channels, padding, and stride represent the parameters of the convolution layer. The input to the system is a  $224 \times 224$  image with 3 channels. The system consists of three systolic arrays, each with 32 processing elements, and the minibatch size is 32. The computation times for each of the layers and the communication overheads are listed in Table II. FP refers to the forward pass, BP\_G and BP $_{\delta}$  refer to the calculation of  $G$  and  $\delta$  in the backward pass. FP and BP overheads indicate how much additional communication is required if a pipeline stage is added immediately after and before the layer, respectively. These overheads are mandatory for any pipeline parallel design and are unavoidable. The additional overhead is the additional communication overhead required to support the LayerPipe algorithm. As shown in Table II, the additional communication overhead of LayerPipe is insignificant.

Table III compares LayerPipe, the proposed operation scheduling, against PipeDream, a traditional pipeline parallel design for the sample network in a three-processor system. In PipeDream, layers 1 and 2 are mapped to their own processor, while layers 3 and 4 are assigned to the third processor. This assignment leads to severely imbalanced pipelines, as seen in the computation times of each processor in Table III. LayerPipe balances the workload across the processors by borrowing computation time from the previous layer. In this sample network, the computation available to borrow is sufficient to balance the pipelines. This results in a 26% improvement in the system throughput ( $3.22 \times 10^7$  versus  $4.35 \times 10^7$ ) with an additional communication overhead of 2.54KB.

### C. Convolutional neural networks

Extending the sample four-layer network analysis, we perform a detailed comparison between LayerPipe and PipeDream for VGG16 and ResNet50. The tests were performed by sweeping the systolic array size from  $32 \times 32$  to  $256 \times 256$  and the minibatch size from 16 to 256 in powers of 2. Table IV compares the speedup and communication overhead between LayerPipe and PipeDream. LayerPipe achieves on average 43% improvement over PipeDream with a maximum increase of 2.73KB in the communication overhead.

Additionally, it achieves greater than 80% improvement with 8 processors. Note, after 11 processors, the ideal time for a single processor is less than the fixed computation time that cannot be borrowed from a layer. Therefore, this level provides the maximum speedup achievable, and further pipelining is no longer advantageous in both PipeDream and LayerPipe, as shown by the saturated curve in Fig. 7.

Fig. 8 summarizes the performance of LayerPipe versus PipeDream by varying the batch and systolic array sizes and testing it on VGG16 and ResNet50. Figs. 8(a) and (b) show the speedup comparison between LayerPipe and PipeDream by varying the number of processors from 2 to 12 while averaging the results across batch and array sizes. LayerPipe consistently outperforms PipeDream on all benchmarks. Fig. 8(c) summarizes the performance of the two methods across different systolic array sizes for the processors. It is seen that using processors with smaller arrays improves the performance of both PipeDream and LayerPipe, but LayerPipe performance improvement is far more significant. A similar analysis in Fig. 8(d) for batch size shows that the performance remains constant for the range of batch sizes tested. This indicates that the pipeline is very susceptible to the systolic array's size but independent of the batch size.

## V. CONCLUSION

This paper presented LayerPipe, a novel approach to achieve intra-layer and inter-layer optimization to generate balanced schedules for the pipelined design. LayerPipe achieves an average speedup of 25% and upwards of 80% with 7 to 9 processors with an insignificant communication overhead compared to PipeDream. The use of delayed gradient may lead to some performance degradation. Since the computations in PipeDream are functionally equivalent to the LayerPipe, any performance degradation in LayerPipe will be the same as in PipeDream. It has been shown that the use of *relaxed look-ahead* can overcome any degradation in delayed LMS [36]. Future work will design accelerators that incorporate relaxed look-ahead, can adapt to multi-GPU clusters, incorporate complex hierarchical communication models for overhead computations, and will address a detailed analysis of branches for complex DNN topologies.

## REFERENCES

- [1] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Advances in Neural Information Processing Systems*, 2012, pp. 1097–1105.
- [2] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [3] C. Szegedy *et al.*, "Going deeper with convolutions," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 1–9.
- [4] M. Johnson *et al.*, "Google's multilingual neural machine translation system: Enabling zero-shot translation," *Transactions of the Association for Computational Linguistics*, vol. 5, pp. 339–351, 2017.
- [5] P. Covington, J. Adams, and E. Sargin, "Deep neural networks for YouTube recommendations," in *Proceedings of the ACM conference on recommender systems*, 2016, pp. 191–198.
- [6] A. Radford *et al.*, "Language models are unsupervised multitask learners," *OpenAI blog*, vol. 1, no. 8, p. 9, 2019.
- [7] N. P. Jouppi *et al.*, "In-datacenter performance analysis of a tensor processing unit," in *Proceedings of the ACM/IEEE International Symposium on Computer Architecture (ISCA)*, Jun. 2017, pp. 1–12.
- [8] N. Jouppi, C. Young, N. Patil, and D. Patterson, "Motivation for and evaluation of the first tensor processing unit," *IEEE Micro*, vol. 38, no. 3, pp. 10–19, 2018.
- [9] NVIDIA, "NVIDIA DGX-1 With Tesla V100 System Architecture white paper," Tech. Rep., 2017.
- [10] Xilinx, "Accelerating DNNs with Xilinx Alveo accelerator cards," Tech. Rep., 2018.
- [11] J. Li *et al.*, "TNPU: An efficient accelerator architecture for training convolutional neural networks," in *Proceedings of the Asia and South Pacific Design Automation Conference (ASP-DAC)*, 2019, pp. 450–455.
- [12] M. Rhu *et al.*, "Compressing DMA engine: Leveraging activation sparsity for training deep neural networks," in *Proceedings of the IEEE International Symposium on High Performance Computer Architecture (HPCA)*, Feb. 2018, pp. 78–91.
- [13] J. Liu *et al.*, "Processing-in-memory for energy-efficient neural network training: A heterogeneous approach," in *Proceedings of the IEEE/ACM International Symposium on Microarchitecture (MICRO)*, Oct. 2018, pp. 655–668.
- [14] J. Zhu, J. Jiang, X. Chen, and C. Tsui, "SparseNN: An energy-efficient neural network accelerator exploiting input and output sparsity," in *Proceedings of the Design, Automation Test in Europe Conference Exhibition (DATE)*, Mar. 2018, pp. 241–244.
- [15] N. Unnikrishnan and K. K. Parhi, "A gradient-interleaved scheduler for energy-efficient backpropagation for training neural networks," in *Proceedings of the International Symposium on Circuits and Systems (ISCAS)*, 2020, pp. 1–5.
- [16] T. Ben-Nun and T. Hoefer, "Demystifying parallel and distributed deep learning: An in-depth concurrency analysis," *ACM Comput. Surv.*, vol. 52, no. 4, Aug. 2019.
- [17] E. Real, A. Aggarwal, Y. Huang, and Q. V. Le, "Regularized evolution for image classifier architecture search," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, no. 01, pp. 4780–4789, Jul. 2019.
- [18] G. Long, F. Ling, and J. G. Proakis, "The LMS algorithm with delayed coefficient adaptation," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 37, no. 9, pp. 1397–1405, 1989.
- [19] Y. Huang *et al.*, "GPipe: Efficient training of giant neural networks using pipeline parallelism," in *Advances in Neural Information Processing Systems*, 2019, pp. 103–112.
- [20] D. Narayanan *et al.*, "PipeDream: Generalized pipeline parallelism for DNN training," in *Proceedings of the ACM Symposium on Operating Systems Principles (SOSP)*, 2019, pp. 1–15.
- [21] M. Jaderberg *et al.*, "Decoupled neural interfaces using synthetic gradients," in *Proceedings of the International Conference on Machine Learning*, 2017, pp. 1627–1635.
- [22] Z. Huo, B. Gu, and H. Huang, "Training neural networks using features replay," in *Advances in Neural Information Processing Systems*, 2018, pp. 6659–6668.
- [23] L. Zhao *et al.*, "BaPipe: Exploration of balanced pipeline parallelism for DNN training," *arXiv preprint arXiv:2012.12544*, 2020.
- [24] T. Wang *et al.*, "FPDeep: Scalable acceleration of CNN training on deeply-pipelined FPGA clusters," *IEEE Transactions on Computers*, vol. 69, no. 8, pp. 1143–1158, 2020.
- [25] Y. Li *et al.*, "Pipe-SGD: A decentralized pipelined SGD framework for distributed deep net training," in *Advances in Neural Information Processing Systems*, S. Bengio *et al.*, Eds., 2018, pp. 8045–8056.
- [26] R. Mayer and H.-A. Jacobsen, "Scalable deep learning on distributed infrastructures: Challenges, techniques, and tools," *ACM Computing Surveys*, vol. 53, no. 1, Feb. 2020.
- [27] K. K. Parhi, *VLSI Digital Signal Processing Systems: Design and Implementation*. Hoboken, NJ, USA: Wiley, 1999.
- [28] K. K. Parhi and D. G. Messerschmitt, "Static rate-optimal scheduling of iterative data-flow programs via optimum unfolding," *IEEE Transactions on Computers*, vol. 40, no. 2, pp. 178–195, Feb. 1991.
- [29] C. E. Leiserson and J. B. Saxe, "Retiming synchronous circuitry," *Algorithmica*, vol. 6, no. 1–6, pp. 5–35, 1991.
- [30] K. K. Parhi and N. K. Unnikrishnan, "Brain-inspired computing: Models and architectures," *IEEE Open Journal of Circuits and Systems*, vol. 1, pp. 185–204, 2020.
- [31] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Proceedings of the Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2015, pp. 234–241.
- [32] C.-Y. Wang and K. K. Parhi, "High-level dsp synthesis using concurrent transformations, scheduling, and allocation," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 14, no. 3, pp. 274–295, 1995.
- [33] K. Ito and K. K. Parhi, "Determining the minimum iteration period of an algorithm," *Journal of VLSI signal processing systems for signal, image and video technology*, vol. 11, no. 3, pp. 229–244, Dec. 1995.
- [34] V. Sze, Y. Chen, T. Yang, and J. S. Emer, "Efficient processing of deep neural networks: A tutorial and survey," *Proceedings of the IEEE*, vol. 105, no. 12, pp. 2295–2329, Dec 2017.
- [35] A. Samajdar *et al.*, "SCALE-Sim: Systolic CNN Accelerator Simulator," *arXiv e-prints*, p. arXiv:1811.02883, Oct. 2018.
- [36] N. R. Shanbhag and K. K. Parhi, "Relaxed look-ahead pipelined LMS adaptive filters and their application to ADPCM coder," *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, vol. 40, no. 12, pp. 753–766, 1993.