

ASCEND: A Scalable and Energy-Efficient Deep Neural Network Accelerator with Photonic Interconnects

Yuan Li, *Student Member, IEEE*, Ke Wang, *Student Member, IEEE*, Hao Zheng, *Student Member, IEEE*, Ahmed Louri, *Fellow, IEEE*, Avinash Karanth, *Senior Member, IEEE*

Abstract—The complexity and size of recent deep neural network (DNN) models have increased significantly in pursuit of high inference accuracy. Chiplet-based accelerator is considered a viable scaling approach to provide substantial computation capability and on-chip memory for efficient process of such DNN models. However, communication using metallic interconnects in prior chiplet-based accelerators poses a major challenge to system performance, energy efficiency, and scalability. Photonic interconnects can adequately support communication across chiplets due to features such as distance-independent latency, high bandwidth density, and high energy efficiency. Furthermore, the salient ease of broadcast property makes photonic interconnects suitable for DNN inference which often incurs prevalent broadcast communication. In this paper, we propose a scalable chiplet-based DNN accelerator with photonic interconnects named ASCEND. ASCEND introduces (1) a novel photonic network that supports seamless intra- and inter- chiplet broadcast communication, and flexible mapping of diverse convolution layers, and (2) a tailored dataflow that exploits the ease of broadcast property and maximizes parallelism by simultaneously processing computations with shared input data. Simulation results using multiple DNN models show that ASCEND achieves 71% and 67% reduction in execution time and energy consumption, respectively, as compared to other state-of-the-art chiplet-based DNN accelerators with metallic or photonic interconnects.

Index Terms—Deep neural network, Photonic interconnect, Chiplet, Accelerator, Dataflow

I. INTRODUCTION

RECENT deep neural network (DNN) models have significantly increased in complexity and size with the goal of improving inference accuracy [1]–[8]. As a result, the underlying computing systems must scale up in computation capability and on-chip memory for efficient process of such DNN models [4]. Chiplet-based accelerator [9]–[12] is considered a viable scaling approach as the scaling of a monolithic chip slows down due to concerns related to power density, yield, and fabrication cost [11], [13]. However, communication across chiplets using metallic interconnects in prior chiplet-based accelerators [12] poses a major challenge to system performance, energy efficiency, and scalability. This is because the long-distance

communication across chiplets accentuates latency and latency discrepancy, inevitably leading to performance degradation and difficulty in data movement orchestration. Besides, the energy consumption of communication across chiplets is higher than within a monolithic chip [12], [14].

Disruptive technologies such as photonic interconnects can potentially overcome the fundamental limitations of metallic interconnects [15]–[18]. Data can propagate through waveguide within one hop regardless of the distance between source and destination, maintaining low and uniform communication latency in a chiplet-based accelerator. Communication bandwidth can be increased through techniques such as wavelength-division multiplexing (WDM) and space-division multiplexing (SDM) [19]. Photonic interconnects have also shown advantage in energy efficiency for long-distance communication as often seen in chiplet-based accelerators [15], [17]. Despite the above superior features of photonic interconnects, the salient ease of broadcast property [15], [16] makes photonic interconnects especially suitable for DNN inference which often incurs prevalent broadcast communication [20]–[23].

Prior photonic networks [24]–[35] often target inter-processor communication typically observed in CPUs or GPUs, and support uniform bandwidth provision at relatively high cost. Besides, the ease of broadcast property of photonic interconnects is not fully exploited. Some prior photonic networks [27], [34] only employ broadcast to facilitate cache coherence protocol. Other photonic networks [29], [31], [33], though constructed by single-write-multiple-read (SWMR) channels, disable the broadcast capability. As a result, a novel photonic network which is tailored to DNN inference and efficiently supports broadcast communication is necessary.

Employing photonic interconnects in chiplet-based accelerators also alters the primary target of dataflow optimization. Prior dataflow optimizations for accelerators with metallic interconnects [12], [36]–[41] often prioritize exploiting locality over broadcast communication. For example, some dataflow optimizations [12], [42] exploit locality of weights at the cost of only being able to broadcast input features. By contrast, [40] exploits locality of input features at the cost of only being able to broadcast weights. Due to the distance-independent latency feature and ease of broadcast property of photonic interconnects, a tailored dataflow that enables broadcast of both types of input data (weights and input features) is beneficial, when a photonic network is implemented in a chiplet-based accelerator.

In this paper, we propose a chiplet-based DNN accelerator

This work was supported in part by National Science Foundation grants CCF-1702980, CCF-1812495, CCF-1901165, CCF-1953980, CCF-1513606, CCF-1703013, and CCF-1901192.

Yuan Li, Ke Wang, Hao Zheng, and Ahmed Louri are with the Department of Electrical and Computer Engineering, George Washington University, Washington, DC 20052 USA. Email: {liyuan5859, cory, haozheng, louri}@gwu.edu.

Avinash Karanth is with the School of Electrical Engineering and Computer Science, Ohio University, Athens, OH 45701 USA. Email: karanth@ohio.edu




Fig. 1. Computations in a convolution layer.

with photonic interconnects named ASCEND. ASCEND includes (1) a novel photonic network that facilitates massive broadcast communication, and (2) a tailored dataflow that exploits the ease of broadcast property to improve parallelism. The contributions of this paper include:

A Novel Photonic Network: We construct a unit 2D processing element (PE) array by selectively grouping local PEs and corresponding PEs across different chiplets in columns and rows, respectively. A waveguide facilitates the broadcast communication from the global buffer (GLB) to this PE array through WDM while a second waveguide reuses the wavelengths for unicast communication from each individual PE to the GLB. A chiplet-based accelerator is constructed by aggregating multiple such PE arrays and connecting them to the GLB through SDM. The resulting photonic network supports (1) seamless one-hop intra- and inter-chiplet broadcast communication, and (2) flexible mapping of diverse convolution layers at the granularity of a unit 2D PE array.

A Tailored Dataflow: We introduce a broadcast-based output-stationary dataflow that exploits the broadcast communication capability of the proposed photonic network and facilitates high parallelism. Specifically, this dataflow enforces intra-chiplet broadcast of input features and inter-chiplet broadcast of weights by spatially mapping computations with shared input features and weights to columns and rows of PEs in the unit 2D PE arrays, respectively. Furthermore, the output-stationary nature of this dataflow minimizes the unicast communication of writing back intermediate data from PEs to the GLB.

Evaluation and Design Space Exploration: We compare ASCEND with other state-of-the-art chiplet-based accelerators with metallic or photonic interconnects using multiple DNN models. Simulation results show that ASCEND achieves up to 71% and 67% reduction in execution time and energy consumption, respectively. We further perform design space exploration by varying multiple factors such as the size of the unit 2D PE array and the capacity of the GLB.

Algorithm 1: Nested loop representation

```

1 for c ← [0:C] do
2   for k ← [0:K] do
3     for h ← [0:H] do
4       for w ← [0:W] do
5         for r ← [0:R] do
6           for s ← [0:S] do
7             O[k, h-r+1, w-s+1] += I[h, w, c] × W[k, r, s, c]

```




Fig. 2. Multiplications with shared weights or input features.

II. BACKGROUND AND MOTIVATION

A. Communications in DNN

The computations involved in a typical convolution layer can be presented as a 6-dimension nested loop over weight kernels, input feature maps (ifmaps), and output feature maps (ofmaps), assuming the batch size of ifmaps to be 1. As illustrated in Fig. 1 and Algorithm 1, the dimensions include the number of weight kernels (k), the number of input channels (c), the height (r) and width (s) of weight kernels, and the height (h) and width (w) of ifmaps. The height (e) and width (f) of ofmaps are not independent and can be derived from the above 6 dimensions. In the single-batch case, $e = h - r + 1$ and $f = w - s + 1$ (assume stride of 1). As shown in Algorithm 1, there are two types of read-only input data: weights $W[k, r, s, c]$ and input features $I[h, w, c]$. Meanwhile, the read-and-write intermediate computation results, known as partial sums (psums), are accumulated to obtain the final output features $O[k, h-r+1, w-s+1]$.

Unlike the dynamic communication patterns often observed in generic applications in CPUs and GPUs, the communication patterns incurred in DNN inference are predetermined by factors such as the dimension values (C, K, H, W, R, S) in the nested loop, the parameters of the underlying computing hardware, and the utilized dataflow. Since each psum $I[h, w, c] \times W[k, r, s, c]$ is unique and only involved in accumulation once, we focus on identifying the broadcast communication incurred during the separate transmission of two input data types: weights $W[k, r, s, c]$ and input features $I[h, w, c]$. Fig. 2 lists multiplications involved in a convolution layer along $k, h/w$, and r/s dimensions. Please note that the c dimension is not shown in Fig. 2, as there is no data sharing and broadcast communication along this dimension. We utilize a symbol "*" to represent the value in c dimension. We observe that multiplications along the k and r/s dimensions share the same input feature $I[h, w, c]$, while multiplications along the h/w dimension share the same weight $W[k, r, s, c]$, indicating the possible broadcast communication for both types of input data. However, prior dataflow optimizations [12], [40], [42] are not developed to fully exploit the broadcast communication. For example, [42] and [12] spatially distribute multiplications along the r/s and k dimensions (1 and 2 in Fig. 2), respectively, to exploit the locality of weights at the cost of only being able to broadcast input features. By contrast, [40] spatially distributes multiplications along the h/w dimension (3 in Fig. 2) to exploit the locality of output features at the cost of only being able to broadcast weights. Given that the




Fig. 3. A WDM photonic link connecting two transmitters and receivers.




Fig. 4. Optical tunable splitter that works at (a) off-resonance state and (b) a transient state with a split ratio of $\alpha/(1-\alpha)$.

transmission of both weights and input features can incur broadcast communication, a tailored dataflow that spatially distributes multiplications along the k and h/w dimensions (4) in Fig. 2) and enables simultaneous broadcast of weights and input features is beneficial when photonic interconnects are employed. Fully-connected layers can also be mathematically framed using the nested loop representation in Algorithm 1, by restricting $H=R$ and $W=S$. DNN models also include other layer types such as pooling and normalization. Our work focuses on the convolution and fully-connected layers as they dominate the computation and memory communication [43], [44].

B. Photonic Interconnects

We present a photonic link that connects two sets of transmitter and receiver by multiplexing two wavelengths in Fig. 3. The light of wavelengths λ_0 and λ_1 is generated by an off-chip laser source and coupled into a waveguide using an optical coupler [45]. Two micro-ring resonators (MRRs) [46], MRR0 and MRR1, work as modulators to modulate the incoming electrical signals on wavelengths λ_0 and λ_1 , respectively. Another two MRRs, MRR2 and MRR3, work as filters to select modulated wavelengths and forward them to the corresponding photodetectors [16]. Each set of modulator and filter MRRs can only work on a specific wavelength (e.g., MRR0 and MRR2 work only on wavelength λ_0). The electrical signals generated from photodetectors are then amplified through transimpedance amplifiers (TIAs) and forwarded to comparators to retrieve the initial data being transmitted. All MRRs that function as either modulators or filters, are tuned by separate resistive heaters with specific thermal tuning modules to mitigate thermal and process variations [16], [29]. The example in Fig. 3 only shows the multiplexing of two wavelengths, prior work has shown as many as 64 wavelengths multiplexed in a waveguide with each wavelength operating at 10 Gbps data rate [25], [47]–[49].

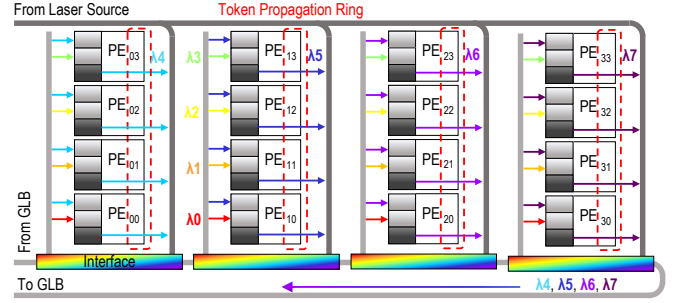


Fig. 5. Architecture and wavelength allocation of a 4×4 unit 2D PE Array.

In addition to the common components shown in Fig. 3, ASCEND includes a special component named tunable splitter [50] to facilitate broadcast communication. Different from modulators and filters that work at either on- or off- resonance, a tunable splitter works at a transient state between the on- and off- resonance. As shown in Fig. 4, the regions outside and inside a tunable splitter ring are with n-type and p-type dope, respectively, to form a PIN diode structure. When no bias voltage is applied to the PIN diode as shown in Fig. 4 (a), the tunable splitter is at off-resonance and light from the input port is directly forwarded to the through port. When applying a proper bias voltage to the PIN diode as shown in Fig. 4 (b), the tunable splitter works at the transient state to guide α fraction of light to the drop port while forwarding the remaining $(1-\alpha)$ fraction of light to the through port. The split ratio is defined as $\alpha/(1-\alpha)$. [50] reports that different split ratios in the range of $[0.4, 1.8]$ can be obtained by tuning the bias voltage in the range of $[0, 5V]$. The applied bias voltage is tuned by a digital-to-analog converter (DAC). In the case that a split ratio beyond the range of $[0.4, 1.8]$ is needed, multiple tunable splitters must be cascaded [51].

Many prior photonic networks [14], [24], [25], [30], [31], [33] for chiplet-based architectures are developed for inter-processor communication typically observed in CPUs and GPUs. The resulting uniform bandwidth provision approach leads to excessive energy and area overhead. For example, the number of MRRs in photonic crossbars in [30], [31], [33] scales quadratically with the number of chiplets. Furthermore, though constructed by SWMR channels which are naturally suitable for broadcast communication, the broadcast capability of the above photonic crossbars are disabled due to power and other concerns. Unlike prior photonic networks for chiplet-based architectures, ASCEND photonic network is tailored for DNN inference and facilitates massive broadcast communication and high parallelism.

III. ASCEND ARCHITECTURE

A. Unit 2D Processing Element Array

Recall Fig. 2 where we spatially distribute multiplications along both k and h/w dimensions to achieve simultaneous broadcast of input features and weights, respectively. As a result, PEs in a chiplet-based accelerator are grouped into a unit 2D array to accommodate the above multiplications. The purpose of constructing a unit 2D PE array is to (1) explore

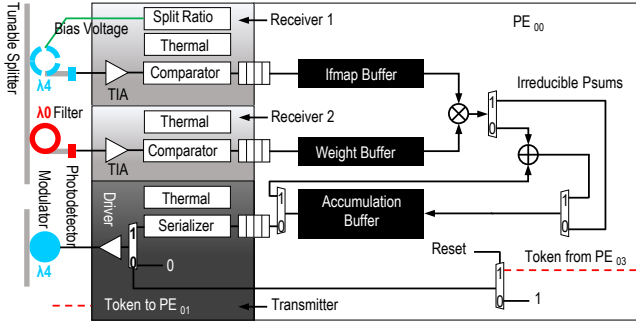


Fig. 6. ASCEND PE architecture, taken PE00 in Fig. 5 as an example.

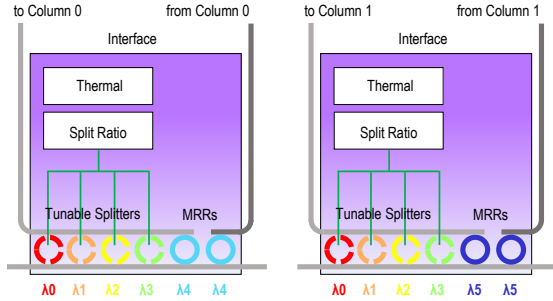


Fig. 7. The interfaces attached to Column0 and Column1 in Fig. 5.

the optimal organization of PEs with high energy-efficiency, and (2) construct large-scale chiplet-based accelerators in a scalable manner by aggregating one or multiple unit 2D PE arrays.

1) *Unit 2D PE Array Architecture*: Fig. 5 illustrates the architecture and wavelength allocation of a 4×4 unit 2D PE array. The architectural details of PE00 in Fig. 5 are presented in Fig. 6. Each PE includes a multiply-accumulate (MAC) unit and register buffers to store weights, input features and intermediate psums. There are one transmitter for PE-to-GLB unicast communication and two receivers for GLB-to-PE broadcast communication. Please note that one receiver is connected to a tunable splitter for per-column broadcast communication as the wavelength (λ_4 in Fig. 6) is shared by all PEs in the same column and only a fraction of light is guided to the photodetector. By contrast, the other receiver is connected to a filter for per-row broadcast communication as the wavelength (λ_0 in Fig. 6) is dedicated for communication to a PE (PE00 in Fig. 6). Since PEs in a column utilize the same wavelength for PE-to-GLB unicast communication, a token-based approach is implemented for arbitration. As shown in Fig. 5 and Fig. 6, a token is propagated among PEs in a column through a token propagation ring. Interfaces attached to different columns are very similar as shown in Fig. 7. A set of tunable splitters are responsible for guiding an appropriate fraction of light of wavelengths ($\lambda_0, \lambda_1, \lambda_2, \lambda_3$ in Fig. 7) to the corresponding column while forwarding the remaining fraction of light to downstream columns. The split ratio shared by a set of tunable splitters is determined according to the position of the corresponding column. For example, the split ratio values for the interfaces associated

with Column0 and Column1 are $1/3$ and $1/2$, as there are three and two downstream columns, respectively. Within each interface, there are also two MRRs that keep working at on-resonance state to filter and merge the wavelength (λ_4 for Column0) for per-column broadcast communication and PE-to-GLB unicast communication.

2) *Wavelength Allocation*: Four wavelengths $\lambda_0, \lambda_1, \lambda_2, \lambda_3$ are utilized to broadcast weights from the GLB to each row of PEs. For example, wavelength λ_0 is utilized to broadcast weights from the GLB to PE00, PE10, PE20, and PE30 in the first row of the unit 2D PE array. Additional four wavelengths $\lambda_4, \lambda_5, \lambda_6, \lambda_7$ are utilized to broadcast input features from the GLB to each column of PEs. For example, wavelength λ_4 is utilized to broadcast input features from the GLB to PE00, PE01, PE02, and PE03. The wavelengths for per-column broadcast communication are also reused for PE-to-GLB unicast communication (e.g., wavelength λ_4 is reused for unicast communication from PEs in the first column to the GLB). Please note that multiple independent waveguides can be implemented using SDM to increase the bandwidth provision for PE-to-GLB communication. All eight wavelengths involved in Fig. 4, $\lambda_0, \lambda_1, \lambda_2, \lambda_3, \lambda_4, \lambda_5, \lambda_6, \lambda_7$, are multiplexed in a waveguide using WDM.

3) *Network Power Consumption of Unit 2D PE Array*: The network power consumption of a unit 2D PE array is directly affected by its size (number of PEs involved) and shape (ratio of array height and width). As array size increases, the overall power consumption of modulators and associated heaters decreases as each transmitter can broadcast to an increasing volume of receivers. However, the laser power consumption increases drastically due to insertion loss increase when more PEs are attached to each broadcast channel. We explore the impact of array size on network power consumption and observe that optimal power consumption is obtained at 16×16 array size. For simplicity, we continue using the 4×4 unit 2D PE array to explain the proposed ASCEND architecture. Similarly, the network power consumption of a unit 2D PE array is also affected by the shape of the array, given a fixed number of PEs involved. Non-square array shapes (e.g., 2×8 and 8×2) inevitably lead to insertion loss imbalance between per-column and per-row broadcast channels. As we assume that each wavelength is generated with similar power from the off-chip laser source, a fraction of power of wavelengths utilized in broadcast channels with low insertion loss will be wasted.

B. ASCEND Network

1) *Network Overview*: Fig. 8 presents an ASCEND architecture with eight accelerator chiplets and eight PEs per accelerator chiplet. This chiplet-based accelerator is constructed by aggregating four 4×4 unit 2D PE arrays. Given the per-column broadcast communication support of a unit 2D PE array discussed before, we allocate a column of PEs in a unit 2D PE array to a single accelerator chiplet (e.g., PE0, PE1, PE2, and PE3 in Chiplet0 are from the same column of a unit 2D PE array). Therefore, the per-column broadcast communication in a unit 2D PE array is equivalent to intra-chiplet broadcast

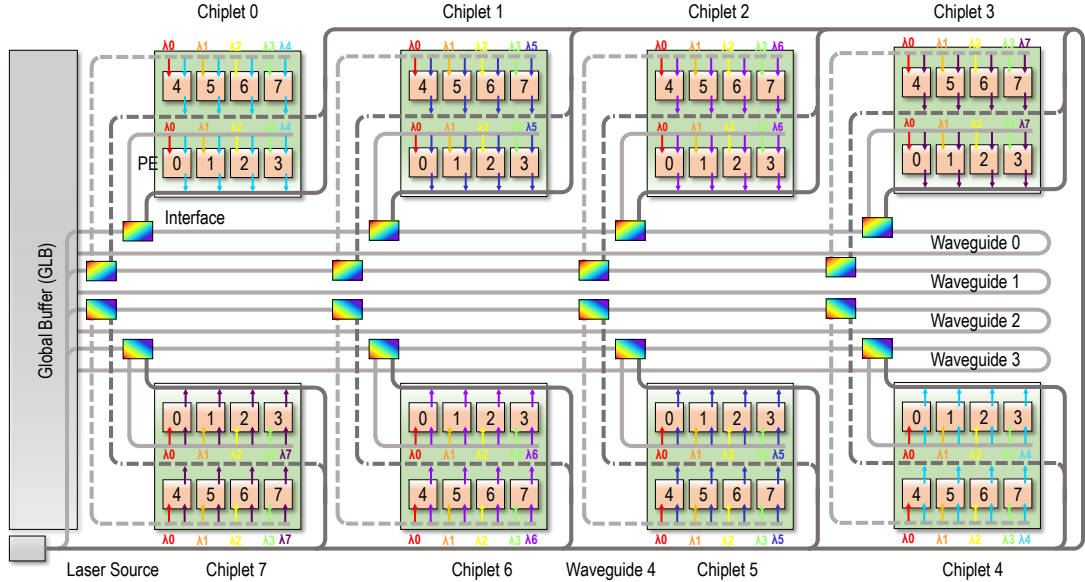


Fig. 8. ASCEND architecture with eight accelerator chiplets and eight PEs per accelerator chiplet, constructed by four 4×4 unit 2D PE arrays. The four unit 2D PE arrays are connected to the GLB through four separate waveguides (Waveguide0, Waveguide1, Waveguide2, and Waveguide3) using SDM. Another waveguide (Waveguide4) is utilized to support simultaneous GLB-to-PE broadcast and PE-to-GLB unicast communication.

communication in the constructed chiplet-based accelerator. Similarly, we allocate a row of PEs in a unit 2D PE array to the same position of different accelerator chiplets (e.g., PE0 in Chiplet0, PE0 in Chiplet1, PE0 in Chiplet2, and PE0 in Chiplet3 are from the same row of a unit 2D PE array), making the per-row broadcast communication equivalent to inter-chiplet broadcast communication. In Fig. 8, each row of sixteen PEs across four accelerator chiplets belong to a unit 2D PE array. PEs within the same accelerator chiplet but in different unit 2D PE arrays are separately connected by waveguides presented by solid and dashed lines. The four involved unit 2D PE arrays in Fig. 8 are connected to the GLB die with four separate waveguides using SDM. For example, the unit 2D PE array including PE0, PE1, PE2, and PE3 in Chiplet0 to Chiplet3 is connected to the GLB with Waveguide0. We observe wavelength reuse between unit 2D PE arrays as separate waveguides are utilized. In the ASCEND architecture shown in Fig. 8, wavelengths λ_0 , λ_1 , λ_2 , and λ_3 are reused for inter-chiplet broadcast communication while wavelengths λ_4 , λ_5 , λ_6 , and λ_7 are reused for intra-chiplet broadcast communication. Waveguide4 is used to deliver light to each PE for PE-to-GLB unicast communication.

2) *Inter-chiplet Broadcast Communication*: The inter-chiplet broadcast function in ASCEND broadcasts the same weight to PEs in the same position of different accelerator chiplets. The broadcast communication from the GLB to PE0 in all eight accelerator chiplets is done by modulating wavelength λ_0 on both Waveguide0 and Waveguide3. Similarly, the broadcast communication from the GLB to PE1 in all accelerator chiplets is done by modulating wavelength λ_1 on both Waveguide0 and Waveguide3, while the broadcast communication from the GLB to PE4 in all accelerator chiplets is done by modulating wavelength λ_0 on both Waveguide1 and Waveguide2. During inter-chiplet broadcast communi-

Algorithm 2: ASCEND dataflow

```

1 // Package level
2 for e1 ← [0:E1] do
3   for f1 ← [0:F1] do
4     parallel_for e2 ← [0:E2] do
5       parallel_for f2 ← [0:F2] do
6         parallel_for k1 ← [0:K1] do
7 // Chiplet level
8   for k2 ← [0:K2] do
9     parallel_for k3 ← [0:K3] do
10    parallel_for e3 ← [0:E3] do
11    parallel_for f3 ← [0:F3] do
12 // PE level
13   for c ← [0:C] do
14    for r ← [0:R] do
15    for s ← [0:S] do
16      k=k3+K3×(k2+K2×k1)
17      e=e3+E3×(e2+E2×e1)
18      f=f3+F3×(f2+F2×f1)
19      O[k,e,f]+=I[r+e-1,s+f-1,c]×W[k,r,s,c]

```

cation, the tunable splitters in the interfaces along a waveguide are tuned to appropriate split ratios to guide a fraction of laser power in λ_0 to λ_3 to the local accelerator chiplet while forwarding the remaining fraction of laser power to downstream accelerator chiplets. For example, the tunable splitters in the interfaces attached to Chiplet0 are all tuned to have a split ratio of 1/3 because there are in total 3 downstream chiplets along either Waveguide0 or Waveguide1. The laser power at the drop port of a tunable splitter is collected and forwarded to the PE with a filter working on the same wavelength, which means this particular PE is a destination of inter-chiplet broadcast communication.

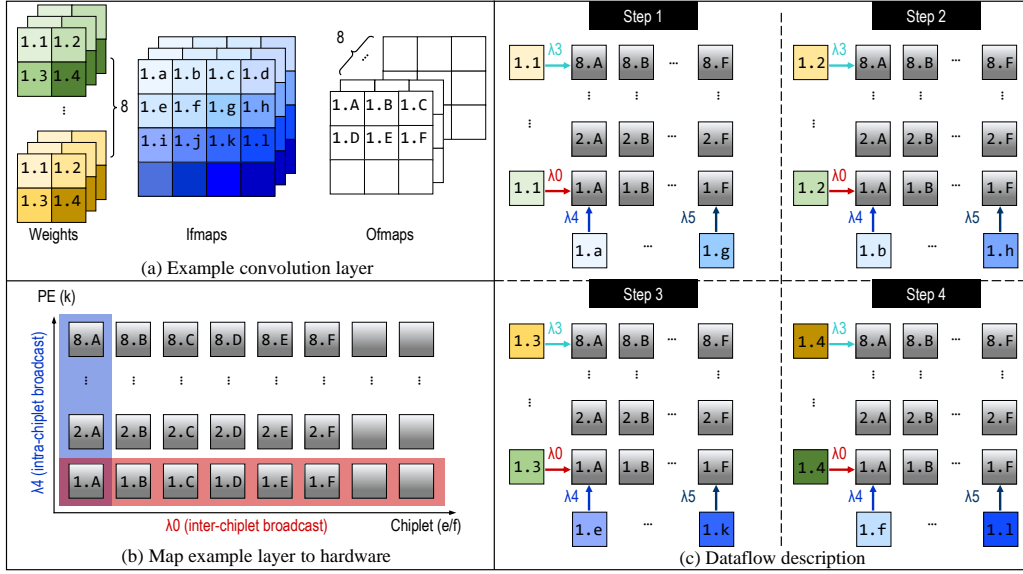


Fig. 9. Processing a convolution layer $[r, s, e, f, c, k]=[2, 2, 4, 4, 3, 8]$ on the ASCEND architecture as shown in Fig. 8 that supports intra- and inter-chiplet broadcast communication. ASCEND dataflow processes output features on the same e/f dimension on different accelerator chiplets while processing output features with different k dimension values on different PEs in the same accelerator chiplet.

3) *Intra-chiplet Broadcast Communication*: The intra-chiplet broadcast function in ASCEND broadcasts the same input feature to PEs in the same accelerator chiplet. The broadcast communication from the GLB to all PEs in Chiplet0 is done by modulating wavelength λ_4 on both Waveguide0 and Waveguide1. Similarly, the broadcast communication from the GLB to all PEs in Chiplet1 is done by modulating wavelength λ_5 on both Waveguide0 and Waveguide1, while the broadcast communication from the GLB to all PEs in Chiplet4 is done by modulating wavelength λ_4 on both Waveguide2 and Waveguide3. During intra-chiplet broadcast communication, the MRR filters in the interfaces along a waveguide work at on-resonance state and completely guide wavelengths for intra-chiplet broadcast communication to the drop port. The laser power is then collected and propagated through local PEs. The tunable splitter attached to the receiver of a specific PE is utilized to guide an appropriate fraction of laser power to the corresponding photodetector while forwarding the remaining fraction of laser power to downstream PEs. For example, the tunable splitter attached to PE0 of Chiplet0 is tuned to a split ratio of 1/3 as there are three downstream PEs (PE1, PE2, and PE3).

4) *PE-to-GLB Unicast Communication*: The intra- and inter-chiplet broadcast functions in ASCEND only address the transmission of input data: weights and input features. The intermediate psums and final output features are transmitted to the GLB through PE-to-GLB unicast function. This function reuses the wavelengths originally allocated for intra-chiplet broadcast communication. For example, wavelength λ_4 is allocated for both intra-chiplet broadcast communication and PE-to-GLB unicast communication in Chiplet0. The wavelength conflict of these two functions is resolved by implementing separate waveguides. As local PEs share the same wavelength for PE-to-GLB unicast communication, a token-based approach is employed. PE that possesses the single-bit token can transmit

its intermediate psums or output features back to the GLB. Once the transmission is complete, the single-bit token is released and propagated to the next local PE through an electrical token propagation ring. The token is originally possessed by the first local PE after reset (active low reset signal in Fig. 6). Because of the uniform computation operations across all PEs, a single-bit electrical token propagation ring is sufficient compared to more sophisticated token arbitration waveguide approach [34]. The bandwidth for PE-to-GLB unicast communication is smaller than the bandwidth for GLB-to-PE broadcast communication. This bottleneck is alleviated by adopting an output-stationary-based dataflow as discussed in the following section. The bandwidth for PE-to-GLB unicast communication can also be expanded by implementing multiple waveguides using SDM.

C. ASCEND Dataflow

ASCEND dataflow, as shown in Algorithm 2 and Fig. 9, is optimized based on three unique features of the proposed photonic network. First, ASCEND supports intra- and inter-chiplet broadcast communication by leveraging the ease of broadcast property of photonic interconnects. Second, by using an output-stationary dataflow, we minimize data exchange between PEs. Third, output-stationary dataflow prioritizes reducing psum movement, which significantly reduces the bandwidth demand for PE-to-GLB unicast communication. Moreover, by multiplexing different wavelengths, we can increase the number of psums sent simultaneously back to GLB from different chiplets.

Consider the convolution layer shown in Fig. 9 (a) as an example. We represent different weight kernels (output channel k dimension) with different colors, and label a weight in a specific kernel using $X:Y$ terminology where X and Y represent the input channel in the c dimension and the position of this weight, respectively. Input features are represented using the

TABLE I
NETWORK PARAMETERS

Simba	Chiplet level	20 Gbps / PE read / write bandwidth	Electrical mesh
	Package level	320 Gbps / chiplet read / write bandwidth	Electrical mesh
POPSTAR	Chiplet level	20 Gbps / PE read / write bandwidth	Electrical mesh
	Package level	310 Gbps / chiplet read bandwidth 100 Gbps / chiplet write bandwidth 10 wavelengths, 10 Gbps / wavelength	Photonic crossbar
ASCEND	Chiplet level	20 Gbps / PE read bandwidth 10 Gbps / PE write bandwidth (shared)	
	Package level	340 Gbps / chiplet read bandwidth 20 Gbps / chiplet write bandwidth 32 wavelengths, 10 Gbps / wavelength	

same terminology as for weights. For output features, X in the $X:Y$ terminology represents the output channel in the k dimension. Fig. 9 (b) describes how the example convolution layer is mapped to the ASCEND architecture shown in Fig. 8 to fully exploit the broadcast capability of the ASCEND photonic network. We map two rows of output features in an ofmap to different chiplets ($E2=2$, $F2=3$ in the dataflow shown in Algorithm 2) while filling the rest PEs in each chiplet with corresponding output features in other ofmaps ($K3=8$ in the dataflow shown in Algorithm 2). As we allocate output features at the same ofmap to different accelerator chiplets, the inter-chiplet broadcast capability of ASCEND photonic network can be leveraged to transmit weights from the GLB to PEs. Meanwhile, as we allocate output features from different ofmaps to PEs within a chiplet, the intra-chiplet broadcast capability of ASCEND photonic network can be leveraged to transmit input features. By doing so, both types of input data are transmitted to PEs through broadcast communication.

Fig. 9 (c) describes the detailed computation and communication operations involved in one iteration of the c loop in Algorithm 2. Since $R=S=2$, the operations are done in four steps. We focus on computation and communication operations related to two PEs responsible for operations related to output features $1.A$ and $1.F$. Operations related to other PEs can be easily inferred. In `Step1`, weight labeled 1.1 and in green color is transmitted to $1.A$ and $1.F$ using inter-chiplet broadcast wavelength λ_0 . Meanwhile, input features labeled $1.a$ and $1.g$ are transmitted to $1.A$ and $1.F$ using intra-chiplet broadcast wavelengths λ_4 and λ_5 , respectively. $1.A$ and $1.F$ perform MAC operations when corresponding weights and input features are delivered, and generate $1.1 \times 1.a$ and $1.1 \times 1.g$, respectively. There are similar operations in the following steps. The psums generated at `Step4` are stored in the local accumulation buffers for the next iteration of the c loop (Line13 in Algorithm 2). Once the entire c loop is completed, the final output features are obtained and transmitted to the GLB.

D. Flexible Mapping of Convolution Layers

Consider a layer $[r, s, e, f, c, k]=[2, 2, 2, 2, 3, 16]$, the number of output features on an ofmap is $e \times f=2 \times 2=4$ while the number of output channels is $k=16$. When mapping this convolution layer to the ASCEND architecture shown in Fig. 8, we observe that only four accelerator chiplets are utilized. Meanwhile, the computations along k dimension have to be iteratively performed while there are idle accelerator chiplets in the system. To resolve this issue, we virtually construct a 16×4 PE array instead of a 8×8 PE array in the k and e/f dimensions by simultaneously broadcasting the same input feature in Waveguide0 to Waveguide3 in Fig. 8. This approach exploits Line6 of the ASCEND dataflow shown in Algorithm 2.

Consider another convolution layer with parameters $[r, s, e, f, c, k]=[2, 2, 4, 4, 3, 4]$, the number of output features on an ofmap is $e \times f=4 \times 4=16$ while the number of output channels $k=4$. This represents an opposite situation as compared to the above example. When mapping this convolution layer to the ASCEND architecture shown in Fig. 8, we observe that only 4 PEs in each accelerator chiplet are utilized. Meanwhile, the computations along e/f dimension have to be iteratively performed while there are idle PEs in all accelerator chiplets in the system. To resolve this issue, we virtually construct a 4×16 PE array in the k and e/f dimensions by simultaneously broadcasting the same weight in Waveguide0 to Waveguide3 in Fig. 8. This approach exploits Line10 and Line11 of the ASCEND dataflow shown in Algorithm 2.

IV. EVALUATION METHODOLOGY

A. Simulation Platform

In order to evaluate ASCEND and other chiplet-based DNN accelerators [12], [33], we extend the open-source MAESTRO simulator [52] to support the non-uniform distribution of latency and bandwidth between PEs. The execution time includes both computation time and communication time. The extended simulator tracks the number of arithmetic operations and the number of accesses to each in-package memory hierarchy (e.g. GLB, local register buffer, etc.) to calculate the computation time and in-package communication time, respectively. The calculation takes the hierarchical network architecture into account and ensures that communication does not exceed the bandwidth limit of the corresponding link. The delay for tuning the optical tunable splitters is set to 500 ps [50]. The off-package communication time is obtained from the DRAMSim2 simulator [53]. We assume that the communication time is maximally overlapped by the computation time.

B. Power Model

We evaluate the power consumption of computations using *Synopsys Design Compiler*. The power consumption values of accessing in-package memory hierarchies and off-package DRAM are obtained using CACTI 6.0 [62] and DRAMSim2, respectively. The power consumption of the in-package metallic-based interconnects is obtained using DSENT [55], while the

TABLE II
STANDARD PHOTONIC PARAMETERS

Component	Value	Component	Value
Laser source	5 dB [47]	Ring drop	1 dB [54]
Coupler	1 dB [47]	Ring through	0.02 dB [55]
Splitter	0.2 dB [49]	Photodetector	0.1 dB [47]
Waveguide	1 dB/cm [47]	Waveguide-to-receiver	0.5 dB [56]
Waveguide bend	1 dB [56]	Receiver sensitivity	-20 dBm [47]
Waveguide crossover	0.05 dB [56]	Ring heating	2 mW [57]

TABLE III
AGGRESSIVE PHOTONIC PARAMETERS

Component	Value	Component	Value
Laser source	5 dB [47]	Ring drop	0.7 dB [55]
Coupler	1 dB [47]	Ring through	0.01 dB [58]
Splitter	0.2 dB [49]	Photodetector	0.1 dB [47]
Waveguide	1 dB/cm [47]	Waveguide-to-receiver	0.5 dB [56]
Waveguide bend	0.01 dB [59]	Receiver sensitivity	-26 dBm [60]
Waveguide crossover	0.05 dB [56]	Ring heating	320 μ W [61]

power consumption of photonic interconnects is derived from Equation (1):

$$P_{total} = P_{laser} + P_{TX} + P_{RX} + P_{thermal} \quad (1)$$

The overall power consumption P_{total} includes three parts: laser power P_{laser} , power consumption of transmitting circuitry P_{TX} , and power consumption of receiving circuitry P_{RX} . We calculate P_{TX} and P_{RX} using the same parameters as in [61], [63]. Please note that the power consumption for ring heating is not included in both P_{TX} and P_{RX} . The values for P_{TX} and P_{RX} are 0.9 mW and 0.6 mW, respectively.

The laser power P_{laser} includes four parts: photodetector sensitivity P_{rs} , insertion loss C_{loss} , extinction ratio power penalty $P_{extinction}$, and system margin M_{system} , as shown in Equation (2):

$$P_{laser} = P_{rs} + C_{loss} + P_{extinction} + M_{system} \quad (2)$$

Table II and Table III list standard and aggressive photonic parameters, respectively, from which the photodetector sensitivity P_{rs} and insertion loss C_{loss} can be obtained or derived. We adopt the standard photonic parameters in Table II for energy consumption estimation, unless otherwise stated. $P_{extinction}$ represents the power penalty caused by extinction ratio which is assumed to be 2 dB [64]. System margin M_{system} is assumed to be 4 dB [65]. The purpose of the system margin is to allocate a certain amount of power to additional sources of power penalty that may develop during the system lifetime.

C. Chiplet-based DNN Accelerators for Comparison

ASCEND is compared with Simba [12] and POPSTAR [33]. Simba is the state-of-the-art chiplet-based DNN accelerator with only metallic interconnects. To the best of our knowledge, there are no chiplet-based DNN accelerators with photonic interconnects. Hence, we select a chiplet-based architecture POPSTAR originally designed for general applications, and replace the general CPU chiplets with accelerator chiplets in Simba to create another baseline for fair comparison. We

assume the system includes 32 chiplets and 32 PEs per chiplet for ASCEND. PE clock frequency is set to 1 GHz similar to [12]. ASCEND adopts 16×16 unit 2D PE array unless otherwise stated. To maintain the same computation capability, the MAC vector width of each PE is 32 in ASCEND. The local buffer size of a PE in ASCEND is 4 kB (128 B per unit MAC vector width) while the local buffer size of a PE in Simba and POPSTAR is 43 kB [12]. The GLB size in ASCEND is 2 MB (64 B per unit MAC vector width), which is the same as in Simba and POPSTAR [12]. The network parameters of ASCEND and two other baselines are listed in Table I. We attempt to keep the bandwidth values at both chiplet and package levels comparable across ASCEND and two baselines. For example, we keep the bandwidth values at chiplet level the same in ASCEND and two baselines by adjusting the clock frequency of the electrical mesh networks in Simba and POPSTAR. However, some bandwidth values cannot be tuned to be exactly the same due to specific features of different network architectures.

D. Benchmarks

We choose four DNN models, VGG-16 [5], ResNet-50 [1], DenseNet-201 [2], and EfficientNet-B7 [8] as the evaluation benchmarks. ResNet-50 includes more variations of weight kernel size and computation intensity, while VGG-16 includes more communication-intensive fully-connected layers that can test network performance in extreme scenarios. There are 21 and 12 different convolution or fully-connected layers in ResNet-50 and VGG-16, respectively. We will test all 33 layers in a layer-by-layer manner, as each layer exhibits different parameters which have implications on performance and energy consumption of our design and the other two baselines. Please note that we have removed redundant layers with the same configuration parameters. For example, `res2a_branch1` in ResNet-50 has been removed because it has the same configuration parameters as `res2[a-c]_branch2c`. Additionally, we accumulate the execution time and energy consumption values of all layers to obtain an implication of the overall execution time and energy consumption in a complete inference pass. Please note that only the convolution and fully-connected layers are taken into account during the accumulation process.

V. EXPERIMENT RESULTS

A. Execution Time and Energy Consumption

Fig. 10 shows the execution time comparison of ASCEND, Simba and POPSTAR in 33 different ResNet-50 and VGG-16 layers. The execution time values are normalized to the execution time of Simba. As compared to Simba, ASCEND achieves execution time reduction in the range of 21% (L1: `conv1`) to 75% (L21: `fc-1000`). The difference in reduction of execution time comes from (1) the average number of hops of inter-chiplet communication in Simba, (2) the ofmap e/f dimension and output channel k dimension that determine the utilization rate of PEs in ASCEND, and (3) the input feature reuse distance that largely determines the intra-chiplet broadcast efficiency in ASCEND. As compared to POPSTAR, ASCEND achieves execution time reduction in the range of 7% (L8: `res3[a-d]_branch2b`) to 55% (L6: `res3a_branch1`).




Fig. 10. Per-layer execution time comparison of Simba, POPSTAR and ASCEND. All values are normalized to Simba.




Fig. 11. Per-layer energy consumption comparison of Simba, POPSTAR and ASCEND. All values are normalized to Simba.

This indicates the effectiveness of the architecture and dataflow co-design. ASCEND performs better than POPSTAR because (1) ASCEND exploits the ease of broadcast feature better than POPSTAR through package-level data partition, and (2) ASCEND allocates higher bandwidth for communication between the GLB and accelerator chiplets. On average, ASCEND performs 52% and 29% faster than Simba and POPSTAR, respectively.

Fig. 11 shows the energy consumption comparison of ASCEND, Simba and POPSTAR in 33 different ResNet-50 and VGG-16 layers using standard photonic parameters listed in Table II. The energy consumption values are normalized to the energy consumption of Simba. As compared to Simba, ASCEND achieves energy consumption saving in the range of 25% (L1: conv1) to 72% (L20: res5[b-c]_branch2a). This mainly comes from the low energy consumption of inter-chiplet communication in ASCEND. As compared to POPSTAR, ASCEND achieves energy consumption saving in the range of 7% (L1: conv1) to 56% (L33: fc-1000). The energy savings observed in different layers are similar because ASCEND photonic inter-chiplet network requires fewer MRRs than the photonic crossbar in POPSTAR. On average, we observe that ASCEND achieves 57% and 46% energy saving as compared to Simba and POPSTAR, respectively.

In addition to the layer-by-layer estimation, we compare the execution time of a complete inference in four different DNN models in Simba, POPSTAR and ASCEND. The results are shown in Fig. 12 (a). We observe that ASCEND achieves execution time reduction in the range of 47% (DenseNet) and 71% (ResNet-50) as compared to Simba. We adopt both standard photonic parameters listed in Table II and aggressive photonic parameters listed in Table III for energy consumption estimation of POPSTAR and ASCEND and present the results in Fig. 12 (b). When using the standard photonic parameters, ASCEND achieves energy reduction in the range of 37% (DenseNet) and 67% (ResNet-50) as compared to Simba. When more aggressive photonic parameters are utilized, more energy reduction in the range of 47% (DenseNet) and 74% (ResNet-50) is achieved by ASCEND as compared to Simba.




Fig. 12. Normalized (a) execution time and (b) energy consumption of one complete inference in different DNN models.




Fig. 13. (a) Execution time and (b) energy consumption breakdown of one complete ResNet-50 inference when comparing ASCEND with Simba and POPSTAR. (c) Network energy consumption breakdown of ASCEND in one complete ResNet-50 inference.

B. Analysis on a ResNet-50 Inference Pass

We present the detailed analysis of execution time and energy consumption (using standard photonic parameters listed in Table II) of a complete ResNet-50 inference. Please note that the values only include the convolution layers and fully-connected layers. We make several observations from the execution time diagram in Fig. 13 (a). First, the numbers of cycles for computation are the same in Simba and POPSTAR, as these two baselines have the same chiplet architecture and dataflow. Second, the numbers of cycles for communication in Simba and POPSTAR are higher than the number of cycles for computation, taking 73% and 66% of overall execution time. Third, the number of cycles for communication in ASCEND is very small due to direct connection between the GLB and each PE, and fully leveraging the broadcast capability of photonic interconnects. Fig. 13 (b) illustrates the energy consumption




Fig. 14. (a) Execution time and (b) energy consumption comparison when applying weight-stationary [12], output-stationary [40], and ASCEND dataflows to the ASCEND architecture. All values are normalized to weight-stationary dataflow.

breakdown of Simba, POPSTAR and ASCEND when processing a complete ResNet-50 inference. The energy reduction as compared to Simba and POPSTAR mainly comes from (1) lower energy consumption of communication network and (2) fewer accesses to the memory hierarchy. When breaking down the network energy consumption of ASCEND as shown in Fig. 13 (c), we observe that the energy consumption values for thermal heating, laser, transmitters and receivers are 1.3 mJ, 1.3 mJ, 0.5 mJ, and 5.5 mJ, respectively. The significant difference of energy consumption values of P_{TX} and P_{RX} illustrates that our design successfully leverages the broadcast capability of photonic interconnects. The ASCEND throughput and energy consumption are 5649 frames per second and 21.7 mJ when running ResNet-50 model and assuming batch size of one.

C. Impact of ASCEND Dataflow

Fig. 14 (a) shows the execution time comparison of weight-stationary dataflow in [12], output-stationary dataflow in [40] with partition along k dimension at the package level, and ASCEND dataflow. All three dataflows are implemented on the ASCEND architecture for fair comparison. The weight-stationary dataflow does not fully exploit the two-level broadcast capability of ASCEND photonic network. Partitions along k dimension at package level and along c dimension at chiplet level prevent full utilization of inter-chiplet weight broadcast and intra-chiplet input feature broadcast, respectively. Further, the weight-stationary dataflow incurs inter-PE communication which yields high latency overhead in ASCEND photonic network. The average execution time reduction of ASCEND dataflow over weight-stationary dataflow is 65%. The output-stationary dataflow [40] is originally designed for single-chip DNN accelerators. We extend it by partitioning along k dimension at the package level. The output-stationary dataflow [40] maps output features to PEs one ofmap at a time. Due to mismatch between ofmap size and system scale, full broadcast capability of ASCEND photonic network can not be often achieved. The average execution time reduction of ASCEND dataflow over the output-stationary dataflow is 23%.

Fig. 14 (b) shows the energy consumption comparison of all three dataflows implemented on the ASCEND architecture using the standard photonic parameters listed in Table II. The average energy saving of ASCEND dataflow over the weight-stationary dataflow [12] is 84%. The excessive energy consumed by the weight-stationary dataflow mainly comes from (1) excessive

GLB access due to prevalent inter-PE communication, and (2) high fraction of unicast communication that leads to high modulation energy. The average energy saving of ASCEND dataflow over the output-stationary dataflow [40] is 39%.

D. Area Estimation

We estimate the area of ASCEND PE (excluding the transmitter and receivers) using Synopsys Design Compiler and a 28 nm technology library. The area of PE excluding the transmitter and receivers is 0.72 mm^2 . We assume that the area for a transmitter or a receiver is 0.0096 mm^2 per wavelength [66]. Hence, the area overhead of the peripheral circuitry (E/O and O/E) of an ASCEND PE is about 3.9%. The area of an accelerator chiplet in ASCEND is 24.07 mm^2 . When assuming $5 \mu\text{m}$ MRR radius [67], the overall area of MRRs is 0.01 mm^2 . When further assuming 4 electrical wires (for data transmission and thermal tuning) per MRR and $36 \mu\text{m}$ micro-bump pitch size [68], the overall area of micro-bumps is 0.68 mm^2 . As most MRRs and micro-bumps can be implemented underneath the accelerator chiplet, we assume that they do not incur extra area overhead.

VI. CONCLUSIONS

In this paper, we propose a chiplet-based DNN accelerator with photonic interconnects named ASCEND. The salient features of ASCEND include (1) a novel photonic network that supports seamless intra- and inter- chiplet broadcast communication and flexible mapping of diverse convolution layers, and (2) a tailored dataflow that exploits the ease of broadcast property of photonic interconnects to maximize parallelism in DNN inference. The combined benefits of the above two features provide high-performance and energy-efficient communication support for scalable chiplet-based DNN accelerators. Simulation studies using multiple DNN models show that ASCEND achieves significant reduction in execution time and energy consumption, and exhibits better scalability, as compared to other state-of-the-art chiplet-based accelerators.

REFERENCES

- [1] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770-778, 2016.
- [2] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely Connected Convolutional Networks," in *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4700-4708, 2017.
- [3] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," in *Advances in Neural Information Processing Systems*, pp. 1097-1105, 2012.
- [4] R. Mayer and H. A. Jacobsen, "Scalable Deep Learning on Distributed Infrastructures: Challenges, Techniques, and Tools," *ACM Computing Survey*, vol. 53, no. 1, pp. 1-37, 2020.
- [5] K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition," *arXiv preprint arXiv: 1409.1556*, pp. 1-14, 2014.
- [6] V. Sze, Y. Chen, T. Yang, and J. S. Emer, "Efficient Processing of Deep Neural Networks: A Tutorial and Survey," *Proc. of the IEEE*, vol. 105, no. 12, pp. 2295-2329, 2017.
- [7] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going Deeper with Convolutions," in *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1-9, 2015.

- [8] M. Tan and Q. Le, "EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks," in *International Conference on Machine Learning (ICML)*, pp. 6105-6114, 2019.
- [9] G. Ascia, V. Catania, A. Mineo, S. Monteleone, M. Palesi, and D. Patti, "Improving Inference Latency and Energy of DNNs through Wireless Enabled Multi-Chip-Module-based Architectures and Model Parameters Compression," in *Proc. of the IEEE/ACM International Symposium on Networks-on-Chip (NOCS)*, pp. 1-6, 2020.
- [10] R. Hwang, T. Kim, Y. Kwon, and M. Rhu, "Centaur: A Chiplet-based, Hybrid Sparse-Dense Accelerator for Personalized Recommendations," in *Proc. of the ACM/IEEE International Symposium on Computer Architecture (ISCA)*, pp. 968-981, 2020.
- [11] A. Kannan, N. E. Jerger, and G. H. Loh, "Enabling Interposer-based Disintegration of Multi-core Processors," in *Proc. of the IEEE/ACM International Symposium on Microarchitecture (MICRO)*, pp. 546-558, 2015.
- [12] Y. S. Shao, J. Clemons, R. Venkatesan, B. Zimmer, M. Fojtik, N. Jiang, B. Keller, A. Klinefelter, N. Pinckney, P. Raina, S. G. Tell, Y. Zhang, W. J. Dally, J. Emer, C. T. Gray, B. Khailany, and S. W. Keckler, "Simba: Scaling Deep-Learning Inference with Multi-Chip-Module-based Architecture," in *Proc. of the IEEE/ACM International Symposium on Microarchitecture (MICRO)*, pp. 14-27, 2019.
- [13] X. Hu, D. Stow, and Y. Xie, "Die Stacking is Happening," *IEEE Micro*, vol. 38, no. 1, pp. 22-28, 2018.
- [14] P. Fotouhi, S. Werner, J. Lowe-Power, and S. J. B. Yoo, "Enabling Scalable Chiplet-based Uniform Memory Architectures with Silicon Photonics," in *Proc. of the International Symposium on Memory Systems (MEMSYS)*, pp. 222-234, 2019.
- [15] D. A. B. Miller, "Rationale and Challenges for Optical Interconnects to Electronic Chips," *Proc. of the IEEE*, vol. 88, no. 6, pp. 728-749, 2000.
- [16] D. A. B. Miller, "Device Requirements for Optical Interconnects to Silicon Chips," *Proc. of the IEEE*, vol. 97, no. 7, pp. 1166-1185, 2009.
- [17] A. Shacham, K. Bergman, and L. P. Carloni, "On the Design of a Photonic Network-on-Chip," in *Proc. of the IEEE/ACM International Symposium on Networks-on-Chip (NOCS)*, pp. 53-64, 2007.
- [18] R. Soref, "The Past, Present, and Future of Silicon Photonics," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 12, no. 6, pp. 1678-1687, 2006.
- [19] K. Bergman, L. P. Carloni, A. Biberman, J. Chan, and G. Hendry, *Photonic Network-on-Chip Design*. Springer, 2014.
- [20] H. Kwon, A. Samajdar, and T. Krishna, "MAERI: Enabling Flexible Dataflow Mapping over DNN Accelerators via Reconfigurable Interconnects," in *Proc. of the ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, pp. 461-475, 2018.
- [21] Y. Li, A. Louri, and A. Karanth, "Scaling Deep-Learning Inference with Chiplet-based Architecture and Photonic Interconnects," in *Proc. of the ACM/IEEE Design Automation Conference (DAC)*, pp. 931-936, 2021.
- [22] Y. Li, A. Louri, and A. Karanth, "SPRINT: A High-Performance, Energy-Efficient, and Scalable Chiplet-based Accelerator with Photonic Interconnects for CNN Inference," *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, vol. 33, no. 10, pp. 2332-2345, 2022.
- [23] Y. Li, A. Louri, and A. Karanth, "SPACX: Silicon Photonics-based Scalable Chiplet Accelerator for DNN Inference," in *Proc. of the IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, pp. 831-845, 2022.
- [24] Y. Demir, Y. Pan, S. Song, N. Hardavellas, J. Kim, and G. Memik, "Galaxy: A High-Performance Energy-Efficient Multi-Chip Architecture Using Photonic Interconnects," in *Proc. of the ACM International Conference on Supercomputing (ICS)*, pp. 303-312, 2014.
- [25] P. Grani, R. Proietti, V. Akella, and S. J. B. Yoo, "Design and Evaluation of AWGR-based Photonic NoC Architectures for 2.5D Integrated High-Performance Computing Systems," in *Proc. of the IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, pp. 289-300, 2017.
- [26] Y. Kao and H. J. Chao, "BLOCON: A Bufferless Photonic Clos Network-on-Chip Architecture," in *Proc. of the IEEE/ACM International Symposium on Networks-on-Chip (NOCS)*, pp. 81-88, 2011.
- [27] N. Kirman, M. Kirman, R. K. Dokania, J. F. Martinez, A. B. Apsel, M. A. Watkins, and D. H. Albonesi, "Leveraging Optical Technology in Future Bus-based Chip Multiprocessors," in *Proc. of the IEEE/ACM International Symposium on Microarchitecture (MICRO)*, pp. 492-503, 2006.
- [28] C. Li, M. Browning, P. V. Gratz, and S. Palermo, "LumiNOC: A Power-Efficient, High-Performance, Photonic Network-on-Chip," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)*, vol. 33, no. 6, pp. 826-838, 2014.
- [29] A. Narayan, Y. Thonnart, P. Vivet, and A. K. Coskun, "PROWAVES: Proactive Runtime Wavelength Selection for Energy-Efficient Photonic NoCs," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)*, vol. 40, no. 10, pp. 2156-2169, 2020.
- [30] A. Narayana, Y. Thonnart, P. Vivet, A. Joshi, and A. K. Coskun, "System-Level Evaluation of Chip-Scale Silicon Photonic Networks for Emerging Data-Intensive Applications," in *Proc. of the Design Automation and Test in Europe Conference (DATE)*, pp. 1444-1449, 2020.
- [31] A. Narayan, Y. Thonnart, P. Vivet, C. F. Tortolero, and A. K. Coskun, "WAVES: Wavelength Selection for Power-Efficient 2.5D-Integrated Photonic NoCs," in *Proc. of the Design Automation and Test in Europe Conference (DATE)*, pp. 516-521, 2019.
- [32] Y. Pan, P. Kumar, J. Kim, G. Memik, Y. Zhang, and A. Choudhary, "Firefly: Illuminating Future Network-on-Chip with Nanophotonics," in *Proc. of the ACM/IEEE International Symposium on Computer Architecture (ISCA)*, pp. 429-440, 2009.
- [33] Y. Thonnart, S. Bernabe, J. Charbonnier, C. Bernard, D. Coriat, C. Fuguet, P. Tissier, B. Charbonnier, S. Malhouitre, D. Saint-Patric, M. Assous, A. Narayan, A. Coskun, D. Dutoit, and P. Vivet, "POPSTAR: A Robust Modular Optical NoC Architecture for Chiplet-based 3D Integrated Systems," in *Proc. of the Design, Automation and Test in Europe Conference (DATE)*, pp. 1456-1461, 2020.
- [34] D. Vantrease, R. Schreiber, M. Monchiero, M. McLaren, N. P. Jouppi, M. Fiorentino, A. Davis, N. Binkert, R. G. Beausoleil, and J. H. Ahn, "Corona: System Implications of Emerging Nanophotonic Technology," in *Proc. of the ACM/IEEE International Symposium on Computer Architecture (ISCA)*, pp. 153-164, 2008.
- [35] A. K. Ziabari, J. L. Abellan, R. Ubal, C. Chen, A. Joshi, and D. Kaeli, "Leveraging Silicon-Photonic NoC for Designing Scalable GPUs," in *Proc. of the ACM International Conference on Supercomputing (ICS)*, pp. 273-282, 2015.
- [36] L. Cavigelli, D. Gschwend, C. Mayer, S. Willi, B. Muheim, and L. Benini, "Origami: A Convolutional Network Accelerator," in *Proc. of the ACM Great Lakes Symposium on VLSI (GLSVLSI)*, pp. 199-204, 2015.
- [37] S. Chakradhar, M. Sankaradas, V. Jakkula, and S. Cadambi, "A Dynamically Configurable Coprocessor for Convolutional Neural Networks," in *Proc. of the ACM/IEEE International Symposium on Computer Architecture (ISCA)*, pp. 247-257, 2010.
- [38] T. Chen, Z. Du, N. Sun, J. Wang, C. Wu, Y. Chen, and O. Temam, "DianNao: A Small-Footprint High-Throughput Accelerator for Ubiquitous Machine-Learning," in *Proc. of the ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, pp. 269-284, 2014.
- [39] Y. Chen, J. Emer, and V. Sze, "Eyeriss: A Spatial Architecture for Energy-Efficient Dataflow for Convolutional Neural Networks," in *Proc. of the ACM/IEEE International Symposium on Computer Architecture (ISCA)*, pp. 367-379, 2016.
- [40] Z. Du, R. Fasthuber, T. Chen, P. Jenne, L. Li, T. Luo, X. Feng, Y. Chen, and O. Temam, "ShiDianNao: Shifting Vision Processing Closer to the Sensor," in *Proc. of the ACM/IEEE International Symposium on Computer Architecture (ISCA)*, pp. 92-104, 2015.
- [41] H. J. Yoo, S. Park, K. Bong, D. Shin, J. Lee, and S. Choi, "A 1.93 TOPS/W Scalable Deep Learning/Inference Processor with Tetra-Parallel MIMD Architecture for Big Data Applications," in *Proc. of the IEEE International Solid-State Circuits Conference (ISSCC)*, pp. 80-81, 2015.
- [42] M. Sankaradas, V. Jakkula, S. Cadambi, S. Chakradhar, I. Durdanovic, E. Cosatto, and H. P. Graf, "A Massively Parallel Coprocessor for Convolutional Neural Networks," in *IEEE International Conference on Application-Specific Systems, Architectures and Processors*, pp. 53-60, 2009.
- [43] M. Gao, J. Pu, X. Yang, M. Horowitz, and C. Kozyrakis, "TETRIS: Scalable and Efficient Neural Network Acceleration with 3D Memory," in *Proc. of the ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, pp. 751-764, 2017.
- [44] X. Yang, M. Gao, Q. Liu, J. Setter, J. Pu, A. Nayak, S. Bell, K. Cao, H. Ha, P. Raina, C. Kozyrakis, and M. Horowitz, "Interstellar: Using Halide's Scheduling Language to Analyze DNN Accelerators," in *Proc. of the ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, pp. 369-383, 2020.
- [45] M. Riccardo, L. Cosimo, C. Lee, G. Kamil, and M. Paolo, "Coupling Strategies for Silicon Photonics Integrated Chips," *Photonics Research*, vol. 7, pp. 201-239, 2019.
- [46] W. Bogaerts, P. De Heyn, T. Van Vaerenbergh, K. De Vos, S. Kumar Selvaraja, T. Claes, P. Dumon, P. Bienstman, D. Van Thourhout, and R. Baets, "Silicon Microring Resonators," *Laser Photonics Reviews*, vol. 6, no. 1, pp. 47-73, 2012.

- [47] R. Morris, A. Karanth, and A. Louri, "Dynamic Reconfiguration of 3D Photonic Networks-on-Chip for Maximizing Performance and Improving Fault Tolerance," in *Proc. of the IEEE/ACM International Symposium on Microarchitecture (MICRO)*, pp. 282-293, 2012.
- [48] S. Van Winkle, A. Karanth, R. Bunescu, and A. Louri, "Extending the Power-Efficiency and Performance of Photonic Interconnects for Heterogeneous Multicores with Machine Learning," in *Proc. of the IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, pp. 480-491, 2018.
- [49] S. Werner, J. Navaridas, and M. Lujan, "Designing Low-Power, Low-Latency Networks-on-Chip by Optimally Combining Electrical and Optical Links," in *Proc. of the IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, pp. 265-276, 2017.
- [50] E. Peter, A. Thomas, A. Dhawan, and S. R. Sarangi, "Active Microring based Tunable Optical Power Splitters," *Optics Communications*, vol. 359, pp. 311-315, 2016.
- [51] J. Bashir, E. Peter, and S. R. Sarangi, "A Survey of On-Chip Optical Interconnects," *ACM Computing Survey*, vol. 51, no. 6, pp. 1-34, 2019.
- [52] H. Kwon, P. Chatarasi, V. Sarkar, T. Krishna, M. Pellauer, and A. Parashar, "MAESTRO: A Data-Centric Approach to Understand Reuse, Performance, and Hardware Cost of DNN Mappings," *IEEE Micro*, vol. 40, no. 3, pp. 20-29, 2020.
- [53] P. Rosenfeld, E. Cooper-Balis, and B. Jacob, "DRAMSim2: A Cycle Accurate Memory System Simulator," *IEEE Computer Architecture Letters (CAL)*, vol. 10, no. 1, pp. 16-19, 2011.
- [54] H. Jayatileka, M. Caverley, N. A. F. Jaeger, S. Shekhar, and L. Chrostowski, "Crosstalk Limitations of Microring-Resonator based WDM Demultiplexers on SOI," in *IEEE Optical Interconnects Conference (OI)*, pp. 48-49, 2015.
- [55] C. Sun, C. O. Chen, G. Kurian, L. Wei, J. Miller, A. Agarwal, L. S. Peh, and V. Stojanovic, "DSENT - A Tool Connecting Emerging Photonics with Electronics for Opto-Electronic Networks-on-Chip Modeling," in *Proc. of the IEEE/ACM International Symposium on Networks-on-Chip (NOCS)*, pp. 201-210, 2012.
- [56] R. Morris and A. Karanth, "Power-Efficient and High-Performance Multi-Level Hybrid Nanophotonic Interconnect for Multicores," in *Proc. of the IEEE/ACM International Symposium on Networks-on-Chip (NOCS)*, pp. 207-214, 2010.
- [57] G. Li, X. Zheng, J. Yao, H. Thacker, I. Shubin, Y. Luo, K. Raj, J. E. Cunningham, and A. V. Krishnamoorthy, "25 Gb/s 1V-Driving CMOS Ring Modulator with Integrated Thermal Tuning," *Optics Express*, vol. 19, no. 21, pp. 20435-20443, 2011.
- [58] S. Pasricha and S. Bahirat, "OPAL: A Multi-Layer Hybrid Photonic NoC for 3D ICs," in *Proc. of the Asia and South Pacific Design Automation Conference (ASP-DAC)*, pp. 345-350, 2011.
- [59] M. Bahadori, M. Nikdast, Q. Cheng, and K. Bergman, "Universal Design of Waveguide Bends in Silicon-on-Insulator Photonics Platform," *Journal of Lightwave Technology*, vol. 37, no. 13, pp. 3044-3054, 2019.
- [60] A. Biberman, K. Preston, G. Hendry, N. Sherwood-Droz, J. Chan, J. S. Levy, M. Lipson, and K. Bergman, "Photonic Network-on-Chip Architectures Using Multilayer Deposited Silicon Materials for High-Performance Chip Multiprocessors," *ACM Journal on Emerging Technologies in Computing Systems (JETC)*, vol. 7, no. 2, pp. 1-25, 2011.
- [61] A. Joshi, C. Batten, Y. Kwon, S. Beamer, I. Shamim, K. Asanovic, and V. Stojanovic, "Silicon-Photonic Clos Networks for Global On-Chip Communication," in *Proc. of the IEEE/ACM International Symposium on Networks-on-Chip (NOCS)*, pp. 124-133, 2009.
- [62] N. Muralimanohar, R. Balasubramonian, and N. P. Jouppi, "CACTI 6.0: A Tool to Model Large Caches," *HP Laboratories*, vol. 27, pp. 1-24, 2009.
- [63] R. Polster, Y. Thonnart, G. Waltener, J. Gonzalez, and E. Cassan, "Efficiency Optimization of Silicon Photonic Links in 65-nm CMOS and 28-nm FDSOI Technology Nodes," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems (TVLSI)*, vol. 24, no. 12, pp. 3450-3459, 2016.
- [64] C. DeCusatis, *Handbook of Fiber Optic Data Communication: A Practical Guide to Optical Networking*. Academic Press, 2013.
- [65] A. V. Krishnamoorthy, R. Ho, X. Zheng, H. Schwetman, J. Lexau, P. Koka, G. Li, I. Shubin, and J. E. Cunningham, "Computer Systems based on Silicon Photonic Interconnects," *Proc. of the IEEE*, vol. 97, no. 7, pp. 1337-1361, 2009.
- [66] Y. Thonnart, M. Zid, J. L. Gonzalez-Jimenez, G. Waltener, R. Polster, O. Dubray, F. Lepin, S. Bernabe, S. Menezes, G. Pares, O. Castany, L. Boutafa, P. Grosse, B. Charbonnier, and C. Baudot, "A 10Gb/s Si-Photonic Transceiver with 150 μ W 120 μ s-Lock-Time Digitally Supervised Analog Microring Wavelength Stabilization for 1Tb/s/mm² Die-to-Die Optical Networks," in *Proc. of IEEE International Solid-State Circuits Conference (ISSCC)*, pp. 350-352, 2018.
- [67] G. Li, X. Zheng, H. Thacker, J. Yao, Y. Luo, I. Shubin, K. Raj, J. E. Cunningham, and A. V. Krishnamoorthy, "40 Gb/s Thermally Tunable CMOS Ring Modulator," in *International Conference on Group IV Photonics (GFP)*, pp. 1-3, 2012.
- [68] H. Zheng, K. Wang, and A. Louri, "A Versatile and Flexible Chiplet-based System Design for Heterogeneous Manycore Architectures," in *Proc. of the ACM/IEEE Design Automation Conference (DAC)*, pp. 1-6, 2020.



Yuan Li received the BS degree in physics from the University of Science and Technology of China in 2010, and the MS degree in microelectronics from the University of Newcastle upon Tyne in 2011. He is currently working toward the PhD degree in computer engineering at the George Washington University. His research interests include machine learning architectures, accelerator-rich heterogeneous systems, and emerging interconnect and memory technologies. He is a student member of the IEEE.



Ke Wang received the B.S. degree in Electrical Engineering from Peking University in 2013, and the M.S. degree in Electrical Engineering from Worcester Polytechnic Institute in 2015. He is currently working toward the Ph.D. degree in Computer Engineering in the School of Engineering and Applied Science at the George Washington University. His research work focuses on high-performance, energy-efficient, fault-tolerant, and secure interconnection networks.



Hao Zheng received the B.S. degree in electrical engineering from Beijing Jiaotong University, Beijing, China, and the M.S. degree in electrical engineering from George Washington University, Washington, DC, USA, where he is currently pursuing the Ph.D. degree in computer engineering. His research interests are in the areas of computer architecture and parallel computing, with emphasis on interconnection networks, machine learning techniques for efficient computing, and energy-efficient manycore architecture designs.



Ahmed Louri is the David and Marilyn Karlgaard Endowed Chair Professor of Electrical and Computer Engineering at the George Washington University, which he joined in August 2015. He is also the director of the High Performance Computing Architectures and Technologies Laboratory. Dr. Louri received the Ph.D. degree in Computer Engineering from the University of Southern California, Los Angeles, California in 1988. From 1988 to 2015, he was a professor of Electrical and Computer Engineering at the University of Arizona, and during that time, he

served six years (2000 to 2006) as the Chair of the Computer Engineering Program. From 2010 to 2013, Dr. Louri served as a program director in the National Science Foundation's (NSF) Directorate for Computer and Information Science and Engineering. He directed the core computer architecture program and was on the management team of several cross-cutting programs. Dr. Louri conducts research in the broad area of computer architecture and parallel computing, with emphasis on interconnection networks, optical interconnects for scalable parallel computing systems, reconfigurable computing systems, and power-efficient and reliable Network-on-Chips (NoCs) for multicore architectures. Recently he has been concentrating on: energy-efficient, reliable, and high-performance many-core architectures; accelerator-rich reconfigurable heterogeneous architectures; machine learning techniques for efficient computing, memory, and interconnect systems; emerging interconnect technologies (photonic, wireless, RF, hybrid) for NoCs; future parallel computing models and architectures (including convolutional neural networks, deep neural networks, and approximate computing); and cloud-computing and data centers. He is the recipient of the 2020 IEEE Computer Society Edward J. McCluskey Technical Achievement Award, "for pioneering contributions to the solution of on-chip and off-chip communication problems for parallel computing and manycore architectures." Dr. Louri is a Fellow of the IEEE, and he is currently the Editor-in-Chief of the IEEE Transactions on Computers. More information can be found at <https://hpcat.seas.gwu.edu/Director.html>.



Avinash Karanth received the BE degree in electronics and communications in February 2000 from the Manipal Institute of Technology, Mangalore University, and the MS and PhD degrees in the Electrical and Computer Engineering Department from The University of Arizona in May 2003 and August 2006, respectively. Presently, he is the Joseph Jachinowski Professor in the School of Electrical Engineering and Computer Science at Ohio University in Athens, Ohio. Dr. Karanth directs the Technologies for Emerging Computer Architecture Lab (TEAL) at

Ohio University. His research interests include computer architecture, optical interconnects, Network-on-Chips (NoCs) and emerging technologies such as nanophotonics, 3D, and wireless interconnects. He is the recipient of the NSF CAREER Award in 2011, the Presidential Research Scholar Award in 2017, the Best Paper Award at the ICCD 2013 conference and his papers have been nominated for Best Paper at the IEEE Symposium on Network-on-Chips (NoCs) in May 2010 and the IEEE Asia & South Pacific Design Automation Conference (ASP-DAC) in January 2009. He is a senior member of the IEEE and member of ACM.