

unzipFPGA: Enhancing FPGA-based CNN Engines with On-the-Fly Weights Generation

Stylianos I. Venieris^{†*}, Javier Fernandez-Marques^{‡*}, Nicholas D. Lane^{†§}

[†]Samsung AI Center, Cambridge, UK, [‡]University of Oxford, [§]University of Cambridge

Email: {s.venieris, nic.lane}@samsung.com, javier.fernandezmarques@cs.ox.ac.uk

**Indicates equal contribution.*

Abstract—Single computation engines have become a popular design choice for FPGA-based convolutional neural networks (CNNs) enabling the deployment of diverse models without fabric reconfiguration. This flexibility, however, often comes with significantly reduced performance on memory-bound layers and resource underutilisation due to suboptimal mapping of certain layers on the engine’s fixed configuration. In this work, we investigate the implications in terms of CNN engine design for a class of models that introduce a pre-convolution stage to decompress the weights at run time. We refer to these approaches as *on-the-fly*. To minimise the negative impact of limited bandwidth on memory-bound layers, we present a novel hardware component that enables the on-chip on-the-fly generation of weights. We further introduce an input selective processing element (PE) design that balances the load between PEs on suboptimally mapped layers. Finally, we present *unzipFPGA*, a framework to train on-the-fly models and traverse the design space to select the highest performing CNN engine configuration. Quantitative evaluation shows that *unzipFPGA* yields an average speedup of $2.14\times$ and 71% over optimised status-quo and pruned CNN engines under constrained bandwidth and up to $3.69\times$ higher performance density over the state-of-the-art FPGA-based CNN accelerators.

I. INTRODUCTION

The emergence of convolutional neural networks (CNNs) as a core component in a variety of AI systems has led to the design of numerous FPGA-based accelerators. Currently, the accelerator landscape spans from CNN-specific processors [1]–[3] to highly custom streaming architectures [4]–[7]. One of the most widely adopted paradigms is the single computation engine [8]–[17], where a powerful processing engine is time-shared to execute the CNN layers sequentially. This approach enables the reuse of the accelerator resources across various CNNs and minimises the need for fabric reconfiguration upon deployment.

Despite its flexibility, performance is often bounded by: 1) layers with low computation-to-communication ratio that become memory-bound [13]–[15] and 2) suboptimal mapping of diverse layers on the fixed engine configuration that leads to underutilised processing elements (PEs) [15], [18], [19]. These two factors set a hard limit to the actual sustained performance that this family of accelerators can reach, indicating an emerging need for novel solutions to minimise their impact.

At the same time, a growing body of work, orthogonal to other techniques such as pruning or quantisation, focuses on compressing CNNs through lossy non-structural methods [20]–[25]. Under this new paradigm, the weights of a model are deployed in a compact form and are “inflated” at run time. Given that several CNN layers are constrained by the limited off-chip memory bandwidth [13]–[15], [18], storing the compressed weights on-chip and reconstructing them on-the-fly can play a key role in alleviating the memory-boundedness

and enabling better utilisation of the computational resources. Nevertheless, the novel dataflow and execution scheme of such models brings up a new challenge regarding their optimised mapping. Existing accelerators have been designed for conventional deep models, adopting either a streaming or layer-by-layer execution [26]. Hence, despite the significant potential of on-the-fly models, their different execution paradigm renders conventional architectures futile in serving them.

In this work, we present *unzipFPGA*, a novel CNN inference system that alleviates the withstanding limitations of single computation engines. To minimise the impact of memory-bound layers, *unzipFPGA* first introduces a mechanism for the on-the-fly generation of weights, reducing transfers from the off-chip memory. Moreover, we propose an input selective PE design that counteracts the underutilisation of the computational resources due to the suboptimal mapping of diverse layers. Overall, this work makes the following key contributions:

- A novel CNN hardware architecture for on-the-fly generation of CNN weights. We introduce a weights generator with custom memory organisation and a highly optimised datapath that enables sustaining high processing rates. The configurable parameters are exposed to the system-level design space exploration to yield the highest performing allocation of resources between the weights generator and the core CNN engine for the given CNN-FPGA pair.
- A new CNN engine design comprising an array of input selective PEs. A subset of PEs is enhanced with efficient switches that enable other PEs to perform work stealing. In this manner, the proposed design applies workload balancing, improving up to 20% the performance on layers that previously severely underutilised the instantiated PEs.
- A framework for deriving and mapping on-the-fly models from pretrained CNNs on a given FPGA, yielding $2.48\times$ average performance density gain over various state-of-the-art FPGA designs. Compared to a state-of-the-art pruning framework [27], *unzipFPGA* delivers up to 90% and 69% compression and throughput gain under the same accuracy.

II. BACKGROUND & RELATED WORK

A. Designing Lightweight Convolutional Neural Networks

The plethora of existing techniques to modify CNNs for faster inference can be categorised into: pruning [28]–[30], which removes redundant parameters; quantisation [31]–[35], which results in low-precision compact models; or, sparsification [36]–[38], which leverages compressed data formats. In addition, a number of frameworks combine several of these techniques. Most notably: Deep Compression [39] which, given

an over-parametrised model, applies pruning, quantisation and Huffman encoding; RedCNN [40] which prunes channels based on an activation overlap metric; and, more recently, APQ [41], which designs a CNN that meets given compute, memory and latency constraints through a joint optimisation formulation.

B. On-the-Fly Convolutional Neural Networks

Orthogonal to these methods, various works have explored ways of factorising the filters in CNNs in order to produce compressed model representations. Common to these techniques is the need for a *decompressing* stage that generates the filters *on-the-fly* during inference. A selection of such techniques include: [20] that uses an auxiliary NN to generate each layer’s weights in the main network given an embedding of the weights. In [21], weight filters are constructed as a dense combination of a set of Fourier Bessels bases that are generated deterministically at run time. Another technique exploiting deterministic bases was presented in [22]–[24], where bases are formed from orthogonal variable spreading factor (OVSF) binary codes. It enables the construction of model weights by learning a linear combination of OVSF bases during training. Finally, in [25] the filters of each layer are sampled from a single learnable filter, inducing weight sharing across filters. This sampling process is deterministic and enables faster processing via an associated integral image-based implementation.

In this work, we make use of OVSF codes to compress filters in a CNN, directly reducing the overheads due to off-chip memory accesses to retrieve weights, based on three aspects: 1) these codes are *binary* and thus can be efficiently stored on-chip [42]; 2) their theoretical properties are well studied by the wireless community [43], [44]; 3) they offer good compression-accuracy trade-off in various AI tasks [23], [24]. Nonetheless, [45] is the sole existing FPGA-based OVSF design, presenting a direct implementation specific for communication systems. To effectively use OVSF with CNNs, the underlying design needs to be tailored and optimised for the CNN dataflow.

C. On-the-Fly OVSF Models

In this section, we give an overview on how OVSF codes are generated and their use for constructing CNN weights.

1) *Constructing OVSF bases*: OVSF codes can be recursively constructed as an n -step expansion of an $L \times L$ Hadamard matrix, H_n , with $L = 2^n$, $n \in \mathbb{N}$. The resulting H_n is comprised of L mutually orthogonal binary codes.

The length, L , of the OVSF codes depends on the shape of the tensor that is generated, but also on the granularity at which this process happens. In this work, we construct $N_{\text{in}} \times K \times K$ filters by concatenating N_{in} $K \times K$ matrices, requiring the generation of $L=K^2$ OVSF codes. Alternative approaches exist [24], each with different compute and memory overheads.

2) *Generating CNN filters with OVSF bases*: A $N_{\text{in}} \times K \times K$ weight filter, f_i , in a convolutional layer can be represented as a truncated linear combination of OVSF codes. The codes are first reshaped to match the shape of f_i . Mathematically, this linear combination can be represented as $f_i = \sum_{j=0}^{\lfloor \rho \cdot l \rfloor} \alpha_i^j B_i^j$, where $\rho \in [0, 1]$ is the ratio of codes to use in order to generate filter f_i , l the total number of OVSF codes of length $L = K^2 N_{\text{in}}$, B_i^j the

j -th OVSF code and $\alpha_i^j \in \mathbb{R}$ its associated scalar. When $\rho=1$, the number of scalars, $|\alpha_i|$, is equal to the number of weights in f_i . These scalars are learnt via standard backpropagation. A compressed representation of f_i is obtained when $\rho < 1$. Upon deployment, the filters are first generated and then the main inference computation proceeds as normal.

By design, OVSF codes cannot directly be used to generate 3×3 filters, which are popular in CNN models. In Section V, we present a technique to overcome this limitation.

D. Challenges of FPGA-based CNN Inference Engines

Until now, a wide array of FPGA-based CNN accelerators have been proposed. Existing designs span from custom streaming architectures [4], [5] and accelerators for quantised [6], [7], [46]–[49] and sparse CNNs [50]–[56], up to instruction-based processors [1]–[3]. One of the most well adopted paradigms are the single computation engines [8]–[16], due to their balanced trade-off of programmability and performance. Currently, despite the progress in processing unit design, further gain in the attainable performance of such engines is hindered by two main factors: i) memory-bound layers that are dominated by the communication with the external memory [13]–[15], [57]. While embedded platforms provide limited bandwidth [58]–[60], *e.g.* less than 4.5 GB/s for Ultra96 and ZC706, sustaining peak bandwidth even on larger devices, such as ZCU104, is nontrivial [59]. This is aggravated as multiple applications are collocated on a single device [60]–[62]; and ii) underutilised PEs due to the mismatch of diverse layer shapes [15]–[19].

Memory-centric Designs. The memory bandwidth problem faced by CNN engines has been studied in previous work. EIE [63] uses the Deep Compression method [39] to compress FC layers’ weights. However, as FC layers have been mostly abandoned in modern CNNs, its applicability is limited. Angel-Eye [64] compresses all layers through precision quantisation. Cambricon-X [50] transfers only the non-zero weights, while Cambricon-S [51] and Scalpel [65] apply coarse weight pruning, but with significant accuracy drop. CircCNN [66] uses block-circulant matrices for weights, but requires complex FFT hardware. [67] uses permuted diagonal matrices for sparse weights, but only targets FC layers. [68] exploits large batch sizes to increase weights reuse and, thus, is not suitable for latency-critical applications that cannot tolerate batching [18].

Focusing on activations, [69] fuses adjacent layers to cache intermediate activations, while Eyeriss [70], [71] and [52] employ encoding schemes to minimise their bandwidth footprint. Other solutions have either relied on large devices [72] and multiple FPGAs [73] to fit all weights on-chip, or utilised highly customised designs to exploit multi-precision cascades [10] or fine-grained pruning [54] at the cost of notable accuracy drop.

Tackling PE Underutilisation. So far, a limited number of designs have focused on ii). [74] groups CONV layers based on the compatibility of their shapes. [4] maps each layer to a dedicated compute stage, which can be used only for shallower networks, but does not scale to the deeper models of today. Furthermore, a limited number of works rely on FFT-based designs with flexible dataflow [14] and costly ASIC solutions [75]–[77] with highly flexible PE interconnect.

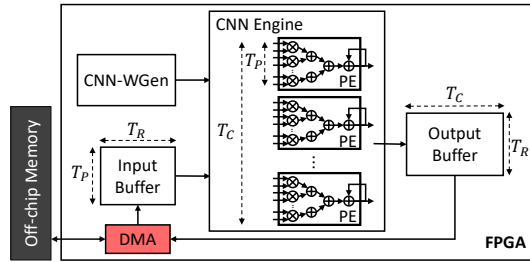


Fig. 1: Overview of unzipFPGA's architecture.

In contrast to these works, we propose an approach that is independent of the CNN engine by not requiring any modification to the engine architecture itself. unzipFPGA can benefit any existing single computation engine by augmenting it with its hardware weights generator and enhancing its PE array with lightweight switches, without affecting the PE's internal processing units. As such, unzipFPGA is orthogonal and complementary to quantisation [64], activations' encoding [52], [70], fusion [69] and zero-skipping PEs [51], [53], [78]–[80].

III. UNZIPFPGA

A. Overview

To remedy the limitations of existing FPGA-based CNN engines, unzipFPGA employs a weights generation scheme that minimises the external memory bandwidth requirements. The proposed system introduces a hardware weights generator based on OVFS bases that enables the on-the-fly, on-chip generation of weights, significantly reducing the off-chip memory accesses. To alleviate the underutilisation of PEs due to layer shape mismatch, unzipFPGA introduces input selective PEs, enabling seamless load balancing with minimal hardware modifications.

The overall flow of the proposed framework is as follows. First, given a CNN, unzipFPGA derives an OVFS variant, by replacing the CONV weights with a trainable linear combination of OVFS bases and selecting each layer's compression ratio ρ (Sec. II-C). Next, the OVFS model is trained using the supplied training set (Sec. V). The OVFS model is passed to the design space exploration (DSE) phase to explore different resource allocations between the weights generator (CNN-WGen) and the CNN engine (Fig. 1) (Sec. IV). Upon completion, the DSE yields the highest performing configuration of unzipFPGA's architecture (Sec. III-B) tailored to the given CNN-device pair and the system is deployed on the target FPGA platform.

B. Architecture

Fig. 1 shows the proposed architecture consisting of three components: the CNN hardware weights generator (CNN-WGen), the core CNN engine and the I/O subsystem. From an operational perspective, the layers of the target CNN are scheduled sequentially, with the computation of the current layer and the I/O communication of adjacent layers overlapped in a pipelined manner. A key component of the proposed architecture is the CNN-WGen module, which relaxes the off-chip memory bandwidth requirements by generating the weights at run time with minimal access to the external memory.

C. CNN Engine

To execute various layer shapes and types, the core CNN engine comprises a parametrised array of PEs for the execution

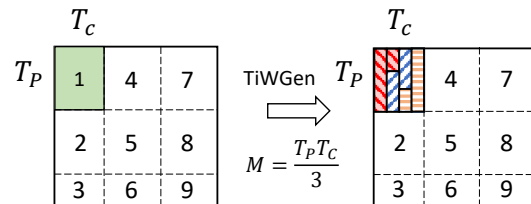


Fig. 2: An example of the TiWGen technique. With a tile size of M , each tile is generated in $\lceil T_P T_C / M \rceil \cdot \rho K^2$ cycles.

of block general matrix multiply (GEMM). By interpreting convolutions as matrix multiplication, both CONV and FC layers are supported. In this manner, a CONV layer with $N_{in} H \times W$ input activations, N_{out} output channels, $K \times K$ filters, p padding and S stride involves the multiplication between an $R \times P$ activations matrix and a $P \times C$ weights matrix in order to produce a $R \times C$ output matrix, with $R = \lceil \frac{H+2p-K}{S} + 1 \rceil \lceil \frac{W+2p-K}{S} + 1 \rceil$, $P = N_{in} K^2$ and $C = N_{out}$. The engine is parametrised with respect to the tile size $\langle T_R, T_P, T_C \rangle$ of each matrix dimension $\langle R, P, C \rangle$, the number of PEs and the MAC operators within each PE.

To produce a $T_R \times T_C$ output tile, $\lceil \frac{P}{T_P} \rceil$ tiles from the activations and weights matrices are processed and accumulated sequentially. unzipFPGA's mapping strategy (Fig. 1) ties T_P to the MAC units per PE to exploit the parallelism within each T_P -wide dot product, and T_C to PEs to parallelise the dot products at each output column. Overall, the rows of the $T_R \times T_P$ activations tile are processed in a pipelined manner to maximise throughput. This is equivalent to an output stationary dataflow [70], [81], which minimises the memory accesses for the output activations by caching partial sums on-chip. Nonetheless, unzipFPGA is adaptable to other dataflows.

D. CNN Hardware Weights Generator

To attain high performance, the adopted weights generation algorithm is mapped to hardware via two key techniques: a *tiled weights generation* (TiWGen) method and a hardware weights generator (CNN-WGen).

1) *Tiled Weights Generation*: To be able to target layers of various dimensions, we introduce a tiling method for the weights generation process, denoted by TiWGen. As shown in Fig. 2, TiWGen divides each $T_P \times T_C$ weights tile into subtiles of size M , with M being uniform across the CNN's layers. Tiling on top of the weights generation method makes the data flow of diverse layers identical to each other. With this approach, the value of M becomes independent of the CNN model and is solely bound by the on-chip computational resources allocated to the weights generator. Thus, M provides a tunable trade-off between weights generation speed and resource consumption.

Alg. 1 describes the overall TiWGen technique. Each $T_P \times T_C$ tile of the weights matrix is processed sequentially (line 1). This is done by partitioning each tile into $\lceil \frac{T_P T_C}{M} \rceil$ subtiles (line 2). After all basis vectors of the current subtile have been processed (lines 4-9), the associated part of the output tile is updated and the algorithm proceeds to the next subtile. When all subtiles of a tile have been generated, the weights matrix is updated (line 12) and the algorithm continues to the next iteration until all weights tiles have been formed.

Algorithm 1: Generation of a layer's weights using TiWGen

Input: Layer's weights matrix shape: $P \times C = N_{in}K^2 \times N_{out}$
 Row and column tile sizes T_P and T_C
 α values with $\alpha \in \mathbb{R}^{N_{in}N_{out} \lceil \rho K^2 \rceil}$

Output: Weights matrix W

```

1 for  $t \leftarrow 1$  to  $\lceil \frac{P}{T_P} \rceil \cdot \lceil \frac{C}{T_C} \rceil$  do // tiles loop - # PIPELINE
2   for  $i \leftarrow 1$  to  $\lceil \frac{T_P T_C}{M} \rceil$  do // subtiles loop - # PIPELINE
3     subtile $_i^t \leftarrow 0$ 
4     for  $j \leftarrow 1$  to  $\rho K^2$  do // basis vectors loop - # PIPELINE
5       for  $k \leftarrow 1$  to  $M$  do // # UNROLL
6          $incr_k \leftarrow vec_j(k) \cdot \alpha_k$  // Multiplier array
7         subtile $_i^t(k) \leftarrow subtile_i^t(k) + incr_k$  // Adder array
8       end
9     end
10    tile $^t \leftarrow$  UpdateTile(tile $^t$ , subtile $_i^t$ )
11  end
12   $W \leftarrow$  UpdateMatrix( $W$ , tile $^t$ )
13 end
  
```

2) *Microarchitecture*: Key enabler to the proposed design is the CNN-wGen module, which is responsible for generating the weights in an orderly manner and feeding them to the CNN engine. Its main components comprise (Fig. 3): i) a compute datapath formed by two vector units (multiplier and adder arrays), ii) the Alpha buffer storing the α values, and iii) the OVSF generator that is responsible for outputting the M -sized basis vector subtiles as dictated by the TiWGen scheme.

CNN-wGen's strategy for mapping the TiWGen method to hardware is annotated in Alg. 1. CNN-wGen implements TiWGen by pipelining the three outer loops over tiles, subtiles and basis vectors, and unrolling the inner loop that processes the M -sized subtile. To unroll the inner loop, CNN-wGen employs two M -wide vector units that perform M -parallel multiplications and additions, respectively. In this manner, tuning M can balance the parallelism-resource usage trade-off.

Compute Datapath. With reference to Fig. 3, the vector arithmetic units must have a fixed number of inputs that meets the resource budget of the target FPGA and namely the available DSPs. For the i -th subtile (line 3 in Alg. 1), ρK^2 vectors of size M are produced by the OVSF generator and the associated α values fetched from the Alpha buffer, and are fed to the multiplier array in a pipelined manner. All M elements are processed in parallel by the M -wide vector units, leading to the unrolling of the inner loop on line 5 of Alg. 1. The adder array processes the outputs of the multiplier array by accumulating the ρK^2 intermediate results. Finally, when the processing proceeds to the next subtile (*i.e.* next iteration of the loop on line 2), the control unit (CU in Fig. 3) resets the accumulators' state. Overall, the design-time configurable parameter M determines the size of the vector units and controls in this way the performance-resource trade-off of CNN-wGen. The tuning of M is exposed to the DSE (Sec. IV) to determine how many resources are allocated to CNN-wGen.

Alpha Buffer. Following TiWGen, each subtile contains weights from N_f distinct $K \times K$ filters (Eq. (1)). To sustain the throughput of CNN-wGen, an equal number of α s has to be fetched in parallel from the Alpha buffer. This is accomplished by designing the buffer with appropriate memory organisation and addressing. Each layer contains $N_{in}N_{out} \lceil \rho_l K_l^2 \rceil$ distinct α

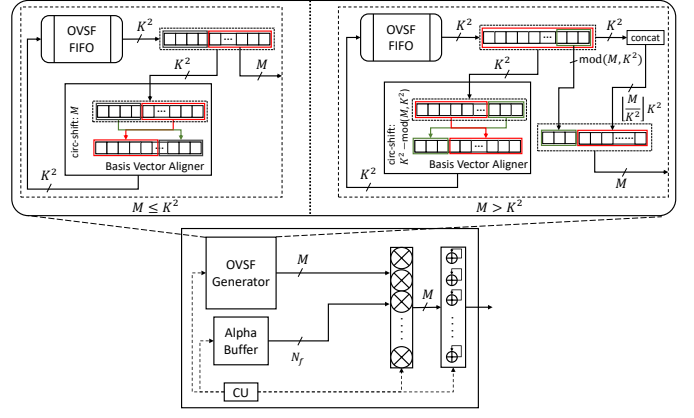


Fig. 3: Microarchitecture of the CNN-wGen module.

values. The RAM blocks are organised with $N_P^{\text{Alpha}} = N_f$ ports and a depth of D^{Alpha} (Eq. (2)) to accommodate N_L layers.

$$N_f = \left\lceil \frac{\min(T_P, M)}{K_{\max}^2} \right\rceil \left\lfloor \frac{M}{T_P} \right\rfloor + \text{mod}(M, T_P) \left\lceil \frac{M}{K_{\max}^2} \right\rceil \quad (1)$$

$$D^{\text{Alpha}} = \sum_{l=1}^{N_L} \underbrace{\left\lceil \frac{N_{in}^l N_{out}^l \lceil \rho_l K_l^2 \rceil}{N_P^{\text{Alpha}}} \right\rceil}_{\text{no. of } \alpha \text{ values}} \quad (\text{Buffer depth}) \quad (2)$$

where N_L is the number of layers, $N_{\{in, out\}}^l$ the l -th layer's number of input/output channels and ρ_l the compression ratio.

OVSF Generator. Based on TiWGen, the basis vectors are processed in a blocked manner with a tile size of M . This approach leads to two pipelined loops over the $\lceil \frac{T_P T_C}{M} \rceil$ subtiles (line 2) and the ρK^2 basis vectors (line 4) and the unrolled loop of processing the M -element subtile with the vector units (line 5). To produce the i -th subtile, the OVSF generator has to feed the compute datapath with ρK^2 basis vectors that are tiled as dictated by TiWGen's parameter M .

In order not to slow down the operation of CNN-wGen, the OVSF generator has to match the rate of the vector units by providing M bits/cycle. A conventional approach involves laying out the tiled vectors in an on-chip buffer through static data marshalling so that the ports match the required rate (*i.e.* M ports) with a depth equal to the number of reads/tile (*i.e.* #basis vectors \times #subtiles). However, such a design would replicate the basis vectors either in the same address (*e.g.* when $M > K^2$) or in multiple addresses (*e.g.* storing rotated versions as required by different subtiles). This leads to inefficient utilisation of on-chip memory due to significant replication.

Another approach that would avoid the basis vector replication involves the instantiation of a K^2 -deep OVSF memory with each location storing one K^2 -bit vector. Such a design requires significantly lower amount of storage and provides an access rate of 1 vector/cycle by reading from the appropriate address. However, to obtain the M -bit subtile from the K^2 -bit vector, complex multiplexer selection circuitry has to be instantiated. This approach can affect the attainable clock frequency or add latency cycles and, thus, degrade CNN-wGen's performance.

To alleviate these limitations when mapping TiWGen's tiling scheme, a custom OVSF generator was developed (Fig. 3 - top). By utilising a FIFO for the OVSF vectors together with a basis vector aligner, the OVSF generator introduces a rate-matching mechanism that sustains the processing

rate of the compute datapath while efficiently utilising the on-chip memory. The generator performs a different operation depending on the values of M and K_i^2 of layer i (Fig. 3).

Initially, the OVFS FIFO stores the $(K_i^2 K_i^2)$ -bit basis vectors. The current vector is read from the FIFO into the top register. If $M \leq K_i^2$, the M least significant bits (LSBs) are outputted to the compute datapath. At the same time, the basis vector is processed by the *basis vector aligner*, which performs an M -bit left circular shift and writes the rotated vector to the OVFS FIFO. If $M > K_i^2$, the basis vector is self-concatenated $\lfloor \frac{M}{K_i^2} \rfloor$ times and written to the output's LS part. Simultaneously, the $\text{mod}(M, K_i^2)$ LSBs of the basis vector are written to the output's MSBs and the constructed vector is passed to the compute datapath. In this case, the aligner performs a left circ-shift of $K_i^2 - \text{mod}(M, K_i^2)$ and writes the result to the FIFO.

With this approach, when the basis vectors are read again out of the OVFS FIFO after ρK_i^2 cycles (*i.e.* in the next iteration of the loop on line 2), they are correctly aligned to directly match TiWGen's tiling pattern. For instance, after the generation of the red-striped subtile in Fig. 2, the FIFO-read basis vectors will be correctly aligned in order to generate the blue-striped subtile without the need for costly selection logic or redundant storage. For CNNs with multiple filter sizes, the basis vector aligner is instantiated with as many circ-shift options. As the distinct filter sizes are known *a priori*, only the required shifting logic is inserted, avoiding expensive generic multiplexers, and the appropriate per-layer bit-shift is selected at run time.

Overall, the proposed design offers a twofold gain. First, it alleviates the redundant replicated storage of basis vectors and avoids the hardware cost of partitioning multiplexers that would require excessive LUTs usage. Second, it provides the necessary bandwidth to the compute datapath while efficiently utilising the on-chip memory through the OVFS FIFO and the resource-efficient aligner design. As the values of K_i^2 for each layer and M are known at design time and after the DSE phase respectively, the OVFS generator is statically instantiated at compile time with the appropriate design.

E. Input Selective PEs for Reducing Underutilisation

When processing compute-bound layers, existing single computation engines are often bounded by the underutilisation of the computational resources [11], [15], [18]. This is due to the fixed CNN engine configuration that does not match all layer dimensions. For instance, mapping a layer with $N_{\text{out}}=64$ (*i.e.* $C=64$) on an engine with 128 PEs (*i.e.* $T_C=128$) would leave 50% of the PEs idle, halving the attainable performance. Fig. 4 shows unzipFPGA's approach of alleviating this by introducing input selective PEs. The initial PE is augmented with registers and switches. However, not all PEs have the same components; only the PEs that remain underutilised even for a single layer are further equipped with a compact switch that selects the inputs to the dot-product unit. In this manner, the PE can utilise either i) the default weight written by the weights generator in the weights buffer or ii) in the absence of this weight, the weight passed from the neighbouring PE. In the second case, the weights are propagated along the PE array

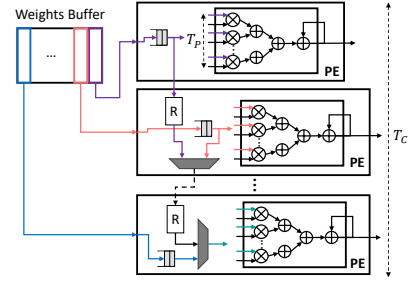


Fig. 4: unzipFPGA's input selective PE array for CNNs.

so that a different weight is used by each augmented PE on each cycle. Moreover, the Input Buffer (Fig. 1) is reorganised accordingly to provide parallel access to multiple rows.

Effectively, this design works as a load-balancing mechanism that partially unrolls the T_R dimension and thus distributes the work more evenly among the PEs. By restricting connectivity only to neighbouring units and enhancing only the underutilised PEs, the additional circuitry is low-overhead and delivers up to 20% higher performance on compute-bound layers.

IV. DESIGN SPACE EXPLORATION

To estimate the performance and resource usage of different architectural parameters, an analytical modelling framework was developed. All models have been verified empirically.

Performance Model. The workload of a CNN with N_L layers is represented as a sequence of $W_i = \langle R_i, P_i, C_i \rangle$ *workload tuples* with $i \in \{1, \dots, N_L\}$. Given a design point $\sigma = \langle M, T_R, T_P, T_C \rangle$, the CNN-WGen's runtime for generating the i -th layer's weights to compute an output tile is given by

$$t_{\text{CNN-WGen}}^i(\sigma, W_i) = \lfloor \rho \cdot l \rfloor \cdot \left\lceil \frac{T_P \cdot T_C}{M} \right\rceil \cdot \left\lceil \frac{P_i}{T_P} \right\rceil \quad (3)$$

where ρ and l are the OVFS ratio and basis length respectively, and with one factor for each of the pipelined loops in Algorithm 1. With the α values transferred upfront and the OVFS method generating all weights on-chip, the off-chip memory transfers involve only the input and output activations:

$$t_{\text{mem in}}^i(\sigma, W_i) = \frac{T_R \cdot P \cdot WL}{bw_{\text{in}}}, \quad t_{\text{mem out}}^i(\sigma, W_i) = \frac{T_R \cdot T_C \cdot WL}{bw_{\text{out}}} \quad (4)$$

where WL is the adopted wordlength, and $bw_{\{\text{in}, \text{out}\}}$ are the memory bandwidths for transferring inputs/outputs.

With the T_C and T_P dimensions unrolled, the computation of an output tile requires the pipelined processing of $\frac{P_i}{T_P}$ tiles for each of the T_R rows. Hence, the CNN engine's runtime for each output tile is estimated as $t_{\text{eng}}^i(\sigma, W_i) = T_R \left\lceil \frac{P_i}{T_P} \right\rceil$. By applying the input selective PEs, the runtime is refined as

$$t_{\text{eng}^*}^i(\sigma, W_i) = \left(T_C - C_i + \left\lceil \frac{T_R \cdot C_i - (T_C - C_i) \cdot (C_i + 1)}{T_C} \right\rceil \right) \cdot \left\lceil \frac{P_i}{T_P} \right\rceil \quad (5)$$

where the T_R dimension is partially unrolled by processing different rows of T_R through the underutilised PEs.

Overall, the accelerator forms a pipeline of three coarse stages: the concurrent input transfer and weights generation, the CNN engine processing and the output transfer. In this context, the initiation interval of the architecture is given by

$$II^i(\sigma, W_i) = \max \left(\max \left(t_{\text{mem in}}^i, t_{\text{CNN-WGen}}^i \right), t_{\text{eng}^*}^i, t_{\text{mem out}}^i \right) \quad (6)$$

and $t_{\text{total}}^i(\sigma, W_i) = II^i(\sigma, W_i) \left\lceil \frac{R_i}{T_R} \right\rceil \left\lceil \frac{C_i}{T_C} \right\rceil$ yields the total runtime for the i -th layer. Thus, given a CNN's workload tuple

$W = \langle W_i \mid \forall i \in \{1, \dots, N_L\} \rangle$, the throughput in inferences per sec (inf/s) is estimated as $T(\sigma, W) = 1 / \sum_{i=1}^{N_L} t_{\text{total}}^i(\sigma, W_i)$.

Resource Consumption Model. From a resource perspective, the main design constraints are the DSPs and on-chip RAM blocks of the target FPGA. Assuming that all MAC operators are mapped to DSPs, the values of $\langle M, T_P, T_C \rangle$ are constrained as $D_{\text{MAC}} \times (M + T_P T_C) \leq D_{\text{fpga}}$, with D_{fpga} the available DSPs and D_{MAC} the DSPs/MAC. In our case, 16-bit fixed-point precision is used, where $D_{\text{MAC}}=1$ on the evaluated Xilinx FPGAs.

In terms of on-chip RAM, the accelerator has the I/O and Alpha buffers with wordlength WL and the binary OVFSF FIFO, with a total capacity requirement as given by Eq. (7).

$$\left(2(T_R T_P + T_R T_C) + D^{\text{Alpha}} N_P^{\text{Alpha}} \right) WL + K_{\text{max}}^2 K_{\text{max}}^2 \leq C_{\text{fpga}} \quad (7)$$

where the factor of 2 accounts for double-buffering and C_{fpga} is the on-chip RAM capacity of the target device.

To further estimate the consumption of LUTs, we used a set of place-and-route measurements and fitted linear regression models as a function of unzipFPGA’s tunable parameters. Overall, we formally capture the resource consumption of a design point σ by means of vector $\mathbf{rsc}(\sigma)$ that holds the utilised amount of DSPs, BRAMs and LUTs. Similarly, we denote the resource vector of the target platform by $\mathbf{rsc}_{\text{Avail.}}$.

Configuration Optimisation. To yield the highest performing design for the given CNN-FPGA pair, we cast DSE as a formal optimisation problem: $\max_{\sigma} T(\sigma, W)$ s.t. $\mathbf{rsc}(\sigma) \leq \mathbf{rsc}_{\text{Avail.}}$. Given a CNN-FPGA pair, we perform exhaustive search, exploring different resource allocations between CNN-WGen and the CNN engine. All designs that violate the resource constraints are pruned as infeasible to accelerate the exploration.

V. LIGHTWEIGHT OVFSF MODELS

To efficiently derive lightweight OVFSF models from a CNN, we introduce two techniques: 1) a regression scheme that enables the use of readily available pretrained models, reducing the costly training on large datasets; and 2) two mechanisms to overcome the limitation of OVFSF codes not being able to directly represent 3×3 filters, with marginal accuracy drop.

Convolutional layers in OVFSF models learn weighting coefficients of OVFSF codes that are linearly combined to generate the filters. To exploit the availability of pretrained models, we obtain the scalars, α_i , from a pretrained filter \hat{f}_i by framing the minimisation problem as a regression:

$$\alpha_i^* = \underset{\alpha_i}{\operatorname{argmin}} \left\| f_i - \hat{f}_i \right\|_2^2 = \underset{\alpha_i}{\operatorname{argmin}} \left\| \sum_{j=0}^{\lfloor \rho-1 \rfloor} \alpha_i^j B_i^j - \hat{f}_i \right\|_2^2 \quad (8)$$

When ratio $\rho=1$ and shape of filter f_i matches the OVFSF condition of having 2^n with $n \in \mathbb{N}$ parameters, then this minimisation guarantees f_i being indistinguishable from \hat{f}_i . When $\rho < 1$, we follow an iterative approach to greedily discard OVFSF bases, based on their scalars’ magnitude, until the specified ρ is reached. The regression stage comprises a 2-layer MLP network and 100 iterations per filter.

We evaluated two approaches to overcome the limitation of generating 3×3 filters using OVFSF codes: first, by learning a 3×3 crop from a 4×4 filter (which can be represented using OVFSF codes); secondly, by means of an auxiliary layer that

TABLE I: Accuracy and number of parameters for ResNet34 models on ImageNet following different compression schemes. Performance measured on the Zynq 7045 platform at different memory bandwidths.

Model Arch.	Compression Method	Params (millions)	Accuracy (%)	Performance (inf/sec) (1×, 2×, 4×)
ResNet34	-	21.8	73.3	(8.6, 16.8, 28.7)
ResNet34	Tay82	17.4	72.7	(10.7, 21.0, 35.6)
ResNet34	Tay72	15.1	71.9	(13.3, 25.8, 44.0)
ResNet34	Tay56	9.4	67.8	(18.3, 36.3, 63.8)
ResNet34	Tay45	6.3	63.1	(21.8, 43.4, 79.8)
ResNet34	OVFSF50	17.2	72.8	(18.1, 21.8, 31.1)
ResNet34	OVFSF25	7.2	71.5	(18.4, 27.3, 33.5)
ResNet34	Tay82+OVFSF50	13.2	71.1	(18.6, 30.0, 37.3)
ResNet34	Tay82+OVFSF25	6.7	70.6	(18.8, 31.0, 38.9)
ResNet34	Tay72+OVFSF50	11.9	70.3	(18.8, 32.0, 40.2)
ResNet34	Tay72+OVFSF25	4.9	68.9	(18.9, 33.3, 42.0)

performs adaptive average pooling transforming a 4×4 filter into the desired 3×3 shape. We empirically found that taking a 3×3 crop worked best for all models except ResNet50.

VI. EVALUATION

A. Experimental Setup

We target two FPGA platforms with different resource characteristics: ZC706 mounting the mid-tier Z7045 and ZCU104 with the larger ZU7EV, with a clock rate of 150 and 200 MHz respectively. Our hardware designs were synthesised and placed-and-routed with Xilinx Vivado HLS and Vivado Design Suite (v2019.2) and run on both boards. The corresponding Arm CPU was used to set up the off-chip memory transactions, launch execution and measure the end-to-end performance of each design. In the evaluation, 16-bit fixed-point precision was used, with unzipFPGA also supporting other quantisation schemes. The available off-chip memory bandwidth was controlled by using a different number of memory ports and amount of word packing, spanning from 1.1 GB/s (1×) to 13.4 GB/s (12×).

Benchmarks. We evaluate on CNNs of varying depth, workload and memory footprint. In particular, we target the widely used family of residual networks [82]. Concretely, we use ResNet18, ResNet34 and ResNet50 on the ImageNet dataset. We also evaluate unzipFPGA on SqueezeNet1.1 [83], a highly optimised network for resource-constrained devices.

Training Scheme. We have developed unzipFPGA’s training component on top of *PyTorch* (1.5) [84]. To derive the OVFSF models, we modified the official *PyTorch*-based ResNet by replacing all 3×3 convolutional layers within residual blocks with their OVFSF counterparts. In all our experiments, we employed pretrained ImageNet models from *torchvision* (0.6.0). After a regression stage that transforms standard models into OVFSF ones, the models were fine-tuned for 30 epochs using an Adam optimiser [85] and learning rate decay every 10 epochs. For each given model, we trained two OVFSF variants following different distributions of ratios ρ for layers in each of the four residual blocks. First, OVFSF50 with ratios=[1.0, 0.5, 0.5, 0.5]; and OVFSF25 with ratios=[1.0, 0.4, 0.25, 0.125]. We follow the same procedure and ratios for SqueezeNet’s *Fire* modules.

Baselines. We introduce two highly optimised single computation engines executing: a) the vanilla CNN and b) pruned variants. For b), we use a state-of-the-art method [27] which applies

TABLE II: Accuracy and number of parameters for ResNet18 models on ImageNet following different compression schemes. Performance measured on the Zynq 7045 platform at different memory bandwidths.

Model Arch.	Compression Method	Params (millions)	Accuracy (%)	Performance (inf/sec) (1×, 2×, 4×)
ResNet18	-	11.7	69.8	(12.0, 23.5, 40.1)
ResNet18	Tay88	9.1	68.8	(14.3, 28.0, 46.4)
ResNet18	Tay82	7.9	67.3	(14.3, 27.8, 45.4)
ResNet18	Tay72	6.0	64.8	(18.2, 35.3, 57.6)
ResNet18	Tay56	3.7	58.3	(23.8, 47.3, 82.2)
ResNet18	OVSF50	9.1	69.2	(19.4, 33.8, 49.9)
ResNet18	OVSF25	4.1	67.3	(19.4, 34.8, 51.0)
ResNet18	Tay82+OVSF50	6.3	66.2	(24.5, 43.2, 57.9)
ResNet18	Tay82+OVSF25	2.8	64.4	(24.5, 43.6, 59.7)

channel pruning based on the first-order Taylor approximation contribution of each filter to the model’s loss. This process is carried out iteratively until a target compression ratio is reached. We refer to a pruned model that keeps 82% of the filters as Tay82 and follow the same naming scheme for other ratios. The baseline architecture comprises the same design as unzipFPGA with the weights transferred from the off-chip memory into an additional $T_P \times T_C$ buffer, if they do not fit on-chip. Both a) and b) are parametrised with tile sizes $\langle T_R, T_P, T_C \rangle$ and roofline modelling [19] is used to obtain the highest throughput configuration for the target CNN-FPGA pair.

B. Performance Comparison

This section assesses the actual performance gains of unzipFPGA compared to our optimised baselines. Tables I and II show the achieved validation set accuracy and actual performance of each design as measured on ZC706 under varying bandwidth budget. Across bandwidths ($1 \times / 2 \times / 4 \times$ where $4 \times$ is the 4.5 GB/s peak measured bandwidth on ZC706), unzipFPGA’s OVSF50 and OVSF25 designs outperform the faithful baseline by $2.1 \times / 1.3 \times / 1.1 \times$ and $2.1 \times / 1.6 \times / 1.2 \times$ respectively for ResNet34, and by $1.6 \times / 1.6 \times / 1.24 \times$ and $1.4 \times / 1.5 \times / 1.3 \times$ respectively for ResNet18. As bandwidth availability increases, the baseline becomes less memory-bound and the performance gap closes. Table III shows the comparison of unzipFPGA with the faithful baseline for SqueezeNet on ZU7EV with peak measured bandwidth of 13.4 GB/s ($12 \times$). Both OVSF50 and OVSF25 designs yield increasing throughput gains as the bandwidth becomes more restricted, with OVSF25 sustaining over 57% speedup for up to $4 \times$ bandwidth. Under $1 \times$ bandwidth, OVSF25 offers minimal additional gains. This is because, below a compression ratio, even though the memory needs are further reduced, activations begin to dominate I/O, and hence further weights reduction does not provide significant benefits. Activations compression techniques [52], [70] can be orthogonally combined to obtain further gains.

Compared to the pruned baselines, unzipFPGA’s OVSF models are more resilient at high compression ratios while resulting in similar accuracy at lower compression ratios. In terms of throughput, unzipFPGA delivers faster processing at more constrained bandwidths. Concretely, ResNet34-OVSF50 is 80% faster than Tay82 at $1 \times$ bandwidth, with less than 1 percentage point (pp) accuracy drop. Despite being almost

TABLE III: Comparing unzipFPGA with faithful baseline on SqueezeNet on ImageNet. Performance measured on the UltraScale+ ZU7EV platform at different memory bandwidths.

Model Arch.	Compression Method	Params (millions)	Accuracy (%)	Performance (inf/sec) (1×, 2×, 4×, 12×)
SqueezeNet	-	1.24	58.2	(54.7, 108.9, 217.8, 515.6)
SqueezeNet	OVSF50	1.07	57.6	(97.4, 189.7, 339.1, 594.1)
SqueezeNet	OVSF25	0.86	57.1	(97.4, 189.7, 342.6, 600.5)

identical in terms of model size and accuracy, Tay82’s approach, which prioritises the pruning of layers with the least accuracy impact, leads to the pruning of mostly compute-bound layers when targeting ResNet34. On the other hand, ResNet34-OVSF50 compresses more effectively memory-bound layers, leading to significantly higher throughput at low bandwidths. A similar pattern is observed for ResNet18. At higher compression ratios, ResNet34-OVSF25 yields 3.7 pp higher accuracy than Tay56, despite using 25% fewer parameters.

To explore the benefits of combining unzipFPGA’s OVSF execution scheme with channel pruning, we derive, train and map on unzipFPGA Tay-OVSF models. This approach yields competitive lightweight models that are not attainable through pruning alone. For instance, ResNet18 with Tay82+OVSF25 is 25% smaller than ResNet18-Tay56 and achieves 6.1 pp higher accuracy, while achieving 34.6% and 23.5% higher throughput over ResNet18-Tay72 with less than 0.5 pp accuracy drop.

Comparison with existing FPGA designs. Here, we evaluate unzipFPGA with diverse existing FPGA designs. In all cases, we use the OVSF50 variant with less than 1-pp accuracy drop. As shown in Table IV, we compare with Snowflake [86] for ResNet18, the SOTA sparse CNN FPGA design [54] for ResNet34, and the light-CNN-optimised Light-OPU [2] and [74] that addresses PE underutilisation for SqueezeNet. On Z7045, unzipFPGA achieves $2.33 \times$ and $1.12 \times$ higher throughput than [86] and [54], respectively. For SqueezeNet, our design delivers $1.83 \times$ and $1.67 \times$ higher performance density in inf/s/DSP and inf/s/Logic, respectively, than Light-OPU and $1.40 \times$ in inf/s/DSP and $1.14 \times - 1.27 \times$ in inf/s/Logic over [74] with the same or 36% less on-chip memory.

ResNet50 Results. The original ResNet50 reaches 76.15% accuracy with 25.56M parameters. Our ResNet50-OVSF50 improves accuracy to 76.23% with 22.83M parameters. Table V presents the comparison for ResNet50. On Z7045, unzipFPGA outperforms Snowflake by $1.59 \times$ in inf/s. Compared with designs on larger devices, our design achieves higher performance density (GOp/s/DSP) by $3.71 \times$, $3.69 \times$, $1.29 \times$ and $1.76 \times - 6.16 \times$ over ResNetAccel, xDNN, DNNVM and ALAMO, respectively, and consistently higher GOp/s/Logic. FTDL reaches higher GOp/s/DSP and $1.47 \times$ lower GOp/s/Logic, but targets a platform with $2.33 \times$ larger on-chip memory and $2 \times$ higher bandwidth, which substantially reduce the off-chip memory accesses and the associated latency.

Resource Usage. We select unzipFPGA and baseline designs with up to 1-pp accuracy drop and compare their post place-and-route resource usage on Z7045, reported in [DSPs, BRAM, LUTs] tuples for $4 \times$ bandwidth. For ResNet34, the faithful baseline consumes [99%,83%,77%], Tay82 [99%,79%,77%], OVSF50 [100%,81%,78%] and Tay82+OVSF50 [100%,87%,81%].

TABLE IV: Comparison with existing FPGA work on ResNet18 (4.03 Gops), ResNet34 (7.40 Gops) and SqueezeNet (0.78 Gops).

	ResNet18 [86]	unzipFPGA: ResNet18*	Sparse ResNet34 [54] using Deep Compression	unzipFPGA: ResNet34*	SqueezeNet [2]	SqueezeNet [74]	unzipFPGA: SqueezeNet*	
FPGA	Z7045	Z7045	Z7045	Z7045	K325T	V485T	V690T	ZU7EV
Clock (MHz)	250	150	166	150	200	170	170	200
Precision	16b fixed	16b fixed	16b fixed	16b fixed	8b fixed	16b fixed	16b fixed	16b fixed
DSPs [†]	900	900	900	900	840	2800	3600	1728
Logic Capacity	218.6 kLUTs	218.6 kLUTs	218.6 kLUTs	218.6 kLUTs	203.8 kLUTs	303.6 kLUTs	433.2 kLUTs	230.0 kLUTs
Block RAM	2.40 MB	2.40 MB	2.40 MB	2.40 MB	1.95 MB	4.52 MB	6.46 MB	4.75 MB
DSP Util. [†]	28.4%	100%	56.8%	100%	83.8%	80%	80%	100%
inf/s	21.38	49.90	27.84	31.1	420.90	913.40	1173.00	792.20
inf/s/DSP [†]	0.0237	0.0576	0.0309	0.0369	0.2505	0.3260	0.3258	0.4584
inf/s/Logic	0.0978	0.2282	0.1273	0.1422	2.0652	3.0085	2.7077	3.444

* using OVFSF50, ** batch size = 1, † 18×18, 19×18 and 25×18 DSP configurations, inf/s/DSP is adjusted based on precision for fair comparison (multiplied by 0.5 for 8b).

TABLE V: Comparison with existing FPGA work on ResNet50 (8.41 Gops).

	Snowflake [1]	unzipFPGA: ResNet50*	ResNetAccel [87]	xDNN [88]	DNNVM [12]	ALAMO [15]	FTDL [89]	unzipFPGA: ResNet50*	
FPGA	Z7045	Z7045	Arria 10 GX1150	VU9P	ZU9	Arria10 GX1150	Stratix10 GX2800	VU125	ZU7EV
Clock (MHz)	250	150	150	500	500	240	300	650	200
Precision	16b fixed	16b fixed	16b fixed	8b fixed	8b fixed	16b fixed	16b fixed	16b fixed	16b fixed
DSPs [†]	900	900	3036	6840	2520	3036	11,520	1200	1728
Logic Capacity	218.6 kLUTs	218.6 kLUTs	427.2 kALMs	1182.0 kLUTs	274.0 kLUTs	427.2 kALMs	933.0 kALMs	716.0 kLUTs	230.0 kLUTs
Block RAM	2.40 MB	2.40 MB	6.60 MB	9.48 MB	4.01 MB	6.60 MB	28.62 MB	11.075 MB	4.75 MB
DSP Util. [†]	28.4%	100%	56.8%	100%	83.8%	80%	80%	100%	100%
inf/s	17.7	28.18	33.93	153.57	80.95	71.38	77.55	151.22	71.71
inf/s/DSP [†]	0.0196	0.0313	0.0111	0.0112	0.0321	0.0235	0.0067	0.1260	0.0415
inf/s/Logic	0.0809	0.1289	0.0794	0.0649	0.1477	0.1671	0.0831	0.2112	0.3117

* using OVFSF50, ** batch size = 1, † 18×18, 19×18 and 25×18 DSP configurations, inf/s/DSP is adjusted based on precision for fair comparison (multiplied by 0.5 for 8b).

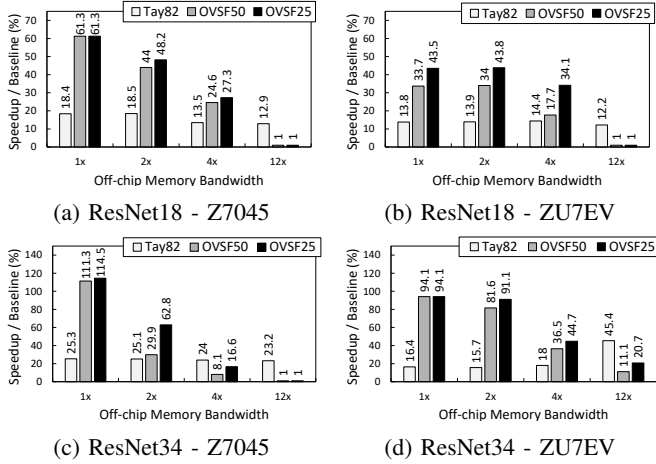


Fig. 5: Speedup over optimised ResNet18 (a-b) and ResNet34 (c-d) baselines when varying the available off-chip memory bandwidth.

For ResNet18, the faithful baseline [78%,99%,70%], Tay88 [78%,99%,69%], OVFSF50 [100%,87%,75%] and Tay82+OVFSF50 [100%,83%,80%]. For ResNet50, OVFSF50 on ZU7EV consumes [100%,87%,78%]. Finally, the input selective PE mechanism adds a minimal LUTs overhead of less than 7%.

C. Sensitivity to Off-Chip Memory Bandwidth

Fig. 5 shows the impact of varying off-chip memory bandwidth over performance on the two target platforms. The figure compares the speedup of unzipFPGA and the Tay82 baseline over the vanilla baseline when varying the external memory bandwidth from 1× to 12×. The bandwidth’s impact is most prominent on the larger ZU7EV, where the performance gains are sustained higher across 1×-4×. In the case of the mid-tier Z7045, we observe a sharper drop in the speedup as the bandwidth increases. This is due to the more limited computational resources of Z7045, which makes most CNN layers compute-bound. In contrast, the abundance of computational resources on ZU7EV makes the CNN layers

more memory-bound. For instance, at 4× bandwidth (4.5 GB/s), the vanilla ResNet18 baseline yields DSP utilisation of 71% on the compute-bound Z7045 and 53% on the more memory-bound ZU7EV. In this case, unzipFPGA significantly improves both cases by mapping ResNet18-OVFSF25 with 89% and 71% DSP utilisation. As a result, unzipFPGA sustains its gains across a wider range of bandwidths and outperforms Tay82, until the bandwidth-abundant case (12×) where computational resources become the critical factor. In this case, Tay82’s lower number of operations due to pruning leads to higher performance.

Across the designs, the input selective PEs contribute an additional speedup of up to 20%, with varying gains depending on the CNN-FPGA pair and the available bandwidth. For ResNet34-OVFSF25 on ZU7EV, disabling this mechanism leads to 0/13.9/3.3/5.9% lower throughput for the four bandwidths, with a similar pattern observed for the rest of the CNNs. Our input selective PEs effectively improve the performance of suboptimally mapped layers in compute-bound settings, whereas no gain is obtained for the most bandwidth-constrained case (1×) where the designs are severely memory-bound, limiting further improvements through higher PE utilisation.

VII. CONCLUSION

This paper presents unzipFPGA, a CNN inference system that mitigates the limitations of existing FPGA-based CNN engines. By generating weights on-the-fly and selectively balancing the PE work, unzipFPGA outperforms both status-quo and pruned CNN engines for the same bandwidth, while consistently achieving large performance density improvements over state-of-the-art CNN accelerators. The presented system opens up several future paths. Concretely, we envision: the automated tuning of OVFSF ratios in a per-layer basis depending on their compute/memory boundedness; and enabling the concurrent deployment of multiple CNNs with lower bandwidths, where unzipFPGA offers larger speedups over existing designs.

REFERENCES

- [1] V. Gokhale, A. Zaidy, A. X. M. Chang, and E. Culurciello, "Snowflake: An Efficient Hardware Accelerator for Convolutional Neural Networks," in *2017 IEEE International Symposium on Circuits and Systems (ISCAS)*, 2017, pp. 1–4.
- [2] Y. Yu, T. Zhao, K. Wang, and L. He, "Light-OPU: An FPGA-Based Overlay Processor for Lightweight Convolutional Neural Networks," in *The 2020 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA)*, 2020, p. 122–132.
- [3] Y. Yu, T. Zhao, M. Wang, K. Wang, and L. He, "Uni-OPU: An FPGA-Based Uniform Accelerator for Convolutional and Transposed Convolutional Networks," *IEEE Transactions on Very Large Scale Integration Systems (TVLSI)*, vol. 28, no. 7, pp. 1545–1556, 2020.
- [4] Huimin Li, Xitian Fan, Li Jiao, Wei Cao, Xuegong Zhou, and Lingli Wang, "A High Performance FPGA-based Accelerator for Large-Scale Convolutional Neural Networks," in *2016 26th International Conference on Field Programmable Logic and Applications (FPL)*, 2016, pp. 1–9.
- [5] S. I. Venieris and C. Bouganis, "fpgaConvNet: Mapping Regular and Irregular Convolutional Neural Networks on FPGAs," *IEEE Transactions on Neural Networks and Learning Systems (TNNLS)*, vol. 30, no. 2, pp. 326–342, 2019.
- [6] Y. Umuroglu, N. J. Fraser, G. Gambardella, M. Blott, P. Leong, M. Jahre, and K. Vissers, "FINN: A Framework for Fast, Scalable Binarized Neural Network Inference," in *Proceedings of the 2017 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA)*, 2017, p. 65–74.
- [7] E. Wang, J. J. Davis, P. Y. K. Cheung, and G. Constantinides, "LUTNet: Learning FPGA Configurations for Highly Efficient Neural Network Inference," *IEEE Transactions on Computers (TOC)*, 2020.
- [8] Y. Guan, H. Liang, N. Xu, W. Wang, S. Shi, X. Chen, G. Sun, W. Zhang, and J. Cong, "FP-DNN: An Automated Framework for Mapping Deep Neural Networks onto FPGAs with RTL-HLS Hybrid Templates," in *2017 IEEE 25th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*, 2017, pp. 152–159.
- [9] M. S. Abdelfattah, D. Han, A. Bitar, R. DiCecco, S. O'Connell, N. Shanker, J. Chu, I. Prins, J. Fender, A. C. Ling, and G. R. Chiu, "DLA: Compiler and FPGA Overlay for Neural Network Inference Acceleration," in *2018 28th International Conference on Field Programmable Logic and Applications (FPL)*, 2018, pp. 411–4117.
- [10] A. Kouris, S. I. Venieris, and C. Bouganis, "A Throughput-Latency Co-Optimised Cascade of Convolutional Neural Network Classifiers," in *2020 Design, Automation Test in Europe Conference Exhibition (DATE)*, 2020, pp. 1656–1661.
- [11] C. Zhang, G. Sun, Z. Fang, P. Zhou, P. Pan, and J. Cong, "Caffeine: Toward Uniformed Representation and Acceleration for Deep Convolutional Neural Networks," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)*, vol. 38, no. 11, pp. 2072–2085, 2019.
- [12] Y. Xing, S. Liang, L. Sui, X. Jia, J. Qiu, X. Liu, Y. Wang, Y. Shan, and Y. Wang, "DNNVM: End-to-End Compiler Leveraging Heterogeneous Optimizations on FPGA-based CNN Accelerators," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)*, 2019.
- [13] A. Kouris, S. I. Venieris, and C. Bouganis, "CascadeCNN: Pushing the Performance Limits of Quantisation in Convolutional Neural Networks," in *2018 28th International Conference on Field Programmable Logic and Applications (FPL)*, 2018, pp. 155–1557.
- [14] Y. Niu, R. Kannan, A. Srivastava, and V. Prasanna, "Reuse Kernels or Activations? A Flexible Dataflow for Low-Latency Spectral CNN Acceleration," in *The 2020 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA)*, 2020, p. 266–276.
- [15] Y. Ma, Y. Cao, S. Vrudhula, and J. Seo, "Automatic Compilation of Diverse CNNs Onto High-Performance FPGA Accelerators," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)*, vol. 39, no. 2, pp. 424–437, 2020.
- [16] M. S. Abdelfattah, Ł. Dudziak, T. Chau, R. Lee, H. Kim, and N. D. Lane, "Best of Both Worlds: AutoML Codesign of a CNN and its Hardware Accelerator," in *Design Automation Conference (DAC)*, 2020.
- [17] S. C. Kao and T. Krishna, "GAMMA: Automating the HW Mapping of DNN Models on Accelerators via Genetic Algorithm," in *2020 IEEE/ACM International Conference On Computer Aided Design (ICCAD)*, 2020, pp. 1–9.
- [18] S. I. Venieris and C. Bouganis, "Latency-Driven Design for FPGA-based Convolutional Neural Networks," in *2017 27th International Conference on Field Programmable Logic and Applications (FPL)*, 2017, pp. 1–8.
- [19] C. Zhang, P. Li, G. Sun, Y. Guan, B. Xiao, and J. Cong, "Optimizing FPGA-Based Accelerator Design for Deep Convolutional Neural Networks," in *Proceedings of the 2015 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA)*, 2015, p. 161–170.
- [20] D. Ha, A. Dai, and Q. V. Le, "HyperNetworks," in *International Conference on Learning Representations (ICLR)*, 2017.
- [21] Q. Qiu, X. Cheng, Robert Calderbank, and G. Sapiro, "DCFNet: Deep Neural Network with Decomposed Convolutional Filters," in *Proceedings of the 35th International Conference on Machine Learning (ICML)*, 2018, pp. 4198–4207.
- [22] J. Fernández-Marqués, W.-S. T. Vincent, S. Bhattachara, and N. D. Lane, "BinaryCmd: Keyword Spotting with deterministic binary basis," in *Conference on Machine Learning and Systems (MLSys)*, 2018.
- [23] V. W.-S. Tseng, S. Bhattacharya, J. F. Marqués, M. Alizadeh, C. Tong, and N. D. Lane, "Deterministic Binary Filters for Convolutional Neural Networks," in *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence (IJCAI)*, 2018, pp. 2739–2747.
- [24] J. Fernández-Marqués, V. W.-S. Tseng, S. Bhattachara, and N. D. Lane, "On-the-Fly Deterministic Binary Filters for Memory Efficient Keyword Spotting Applications on Embedded Devices," in *Proceedings of the 2nd International Workshop on Embedded and Mobile Deep Learning (EMDL)*. ACM, 2018, p. 13–18.
- [25] Y. Yang, J. Yu, N. Jojic, J. Huan, and T. S. Huang, "FSNet: Compression of Deep Convolutional Neural Networks by Filter Summary," in *International Conference on Learning Representations (ICLR)*, 2020.
- [26] S. I. Venieris, A. Kouris, and C.-S. Bouganis, "Toolflows for Mapping Convolutional Neural Networks on FPGAs: A Survey and Future Directions," *ACM Comput. Surv. (CSUR)*, vol. 51, no. 3, 2018.
- [27] P. Molchanov, A. Mallya, S. Tyree, I. Frosio, and J. Kautz, "Importance Estimation for Neural Network Pruning," in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [28] J.-H. Luo, J. Wu, and W. Lin, "ThiNet: A Filter Level Pruning Method for Deep Neural Network Compression," in *IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [29] D. Blalock, J. J. G. Ortiz, J. Frankle, and J. Gutttag, "What is the State of Neural Network Pruning?" in *Conference on Machine Learning and Systems (MLSys)*, 2020.
- [30] Y. He, G. Kang, X. Dong, Y. Fu, and Y. Yang, "Soft Filter Pruning for Accelerating Deep Convolutional Neural Networks," in *International Joint Conference on Artificial Intelligence (IJCAI)*, 2018.
- [31] R. Krishnamoorthi, "Quantizing Deep Convolutional Networks for Efficient Inference: A Whitepaper," 2018.
- [32] B. Jacob, S. Kligys, B. Chen, M. Zhu, M. Tang, A. Howard, H. Adam, and D. Kalenichenko, "Quantization and Training of Neural Networks for Efficient Integer-Arithmetic-Only Inference," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [33] Z. Dong, Z. Yao, A. Gholami, M. Mahoney, and K. Keutzer, "HAWQ: Hessian Aware Quantization of Neural Networks with Mixed-Precision," in *IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [34] J. Fernandez-Marques, P. N. Whatmough, A. Mundy, and M. Mattina, "Searching for Winograd-aware Quantized Networks," in *Conference on Machine Learning and Systems (MLSys)*, 2020.
- [35] M. Alizadeh, J. Fernández-Marqués, N. D. Lane, and Y. Gal, "A Systematic Study of Binary Neural Networks' Optimisation," in *International Conference on Learning Representations (ICLR)*, 2019.
- [36] W. Wen, C. Wu, Y. Wang, Y. Chen, and H. Li, "Learning Structured Sparsity in Deep Neural Networks," in *Proceedings of the 30th International Conference on Neural Information Processing Systems (NeurIPS)*, 2016.
- [37] S. Bhattacharya and N. D. Lane, "Sparsification and Separation of Deep Learning Layers for Constrained Resource Inference on Wearables," in *Proceedings of the 14th ACM Conference on Embedded Network Sensor Systems (SenSys)*. ACM, 2016, p. 176–189.
- [38] T. Gale, M. Zaharia, C. Young, and E. Elsen, "Sparse GPU Kernels for Deep Learning," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, 2020.
- [39] S. Han, H. Mao, and W. J. Dally, "Deep Compression: Compressing Deep Neural Networks with Pruning, Trained Quantization and Huffman Coding," in *International Conference on Learning Representations (ICLR)*, 2015.

- [40] Y. Wang, C. Xu, C. Xu, and D. Tao, "Beyond Filters: Compact Feature Map for Portable Deep Model," in *Proceedings of the 34th International Conference on Machine Learning (ICML)*, 2017, pp. 3703–3711.
- [41] T. Wang, K. Wang, H. Cai, J. Lin, Z. Liu, and S. Han, "APQ: Joint Search for Network Architecture, Pruning and Quantization Policy," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [42] S. Kim, M. Kim, C. Shin, J. Lee, and Y. Kim, "Efficient implementation of OVFS code generator for UMTS systems," in *2009 IEEE Pacific Rim Conference on Communications, Computers and Signal Processing*, 2009, pp. 483–486.
- [43] B. D. Andreev *et al.*, "Orthogonal Code Generator for 3G Wireless Transceivers," in *Proceedings of the 13th ACM Great Lakes Symposium on VLSI (GLSVLSI)*, 2003, pp. 229–232.
- [44] T. Rintakoski, M. Kuulusa, and J. Nurmi, "Hardware unit for OVFS/Walsh/Hadamard code generation [3G mobile communication applications]," in *2004 International Symposium on System-on-Chip (ISSOC)*, 2004, pp. 143–145.
- [45] G. Purohit, V. K. Chaubey, K. S. Raju, and P. V. Reddy, "FPGA based implementation and testing of OVFS code," in *2013 International Conference on Advanced Electronic Systems (ICAES)*, 2013, pp. 88–92.
- [46] A. Prost-Boucle, A. Bourge, and F. Pétrot, "High-Efficiency Convolutional Ternary Neural Networks with Custom Adder Trees and Weight Compression," *ACM Trans. Reconfigurable Technol. Syst. (TRETS)*, vol. 11, no. 3, 2018.
- [47] M. Blott, T. B. Preußner, N. J. Fraser, G. Gambardella, K. O'Brien, Y. Umuroglu, M. Leeser, and K. Vissers, "FINN-R: An End-to-End Deep-Learning Framework for Fast Exploration of Quantized Neural Networks," *ACM Trans. Reconfigurable Technol. Syst. (TRETS)*, vol. 11, no. 3, 2018.
- [48] Y. Zhao, X. Gao, X. Guo, J. Liu, E. Wang, R. Mullins, P. Y. K. Cheung, G. Constantinides, and C. Xu, "Automatic Generation of Multi-Precision Multi-Arithmetic CNN Accelerators for FPGAs," in *2019 International Conference on Field-Programmable Technology (ICFPT)*, 2019, pp. 45–53.
- [49] Y. Umuroglu, Y. Akhauri, N. J. Fraser, and M. Blott, "LogicNets: Co-Designed Neural Networks and Circuits for Extreme-Throughput Applications," in *2020 30th International Conference on Field-Programmable Logic and Applications (FPL)*, 2020, pp. 291–297.
- [50] S. Zhang, Z. Du, L. Zhang, H. Lan, S. Liu, L. Li, Q. Guo, T. Chen, and Y. Chen, "Cambricon-X: An Accelerator for Sparse Neural Networks," in *2016 49th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, 2016, pp. 1–12.
- [51] X. Zhou, Z. Du, Q. Guo, S. Liu, C. Liu, C. Wang, X. Zhou, L. Li, T. Chen, and Y. Chen, "Cambricon-S: Addressing Irregularity in Sparse Neural Networks through A Cooperative Software/Hardware Approach," in *2018 51st Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, 2018, pp. 15–28.
- [52] A. Parashar, M. Rhu, A. Mukkara, A. Pugliesi, R. Venkatesan, B. Khailany, J. Emer, S. W. Keckler, and W. J. Dally, "SCNN: An Accelerator for Compressed-Sparse Convolutional Neural Networks," in *2017 ACM/IEEE 44th Annual International Symposium on Computer Architecture (ISCA)*, 2017, pp. 27–40.
- [53] A. Gondimalla, N. Chesnut, M. Thottethodi, and T. Vijaykumar, "SparTen: A Sparse Tensor Accelerator for Convolutional Neural Networks," in *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, 2019, pp. 151–165.
- [54] L. Lu, J. Xie, R. Huang, J. Zhang, W. Lin, and Y. Liang, "An Efficient Hardware Accelerator for Sparse Convolutional Neural Networks on FPGAs," in *IEEE 27th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*, 2019, pp. 17–25.
- [55] K. Hegde, H. Asghari-Moghaddam, M. Pellauer, N. Crago, A. Jaleel, E. Solomonik, J. Emer, and C. W. Fletcher, "Extensor: An accelerator for sparse tensor algebra," in *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, 2019, p. 319–333.
- [56] N. Srivastava, H. Jin, S. Smith, H. Rong, D. Albonesi, and Z. Zhang, "Tensaurus: A Versatile Accelerator for Mixed Sparse-Dense Tensor Computations," in *2020 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, 2020, pp. 689–702.
- [57] K. Siu, D. M. Stuart, M. Mahmoud, and A. Moshovos, "Memory Requirements for Convolutional Neural Network Hardware Accelerators," in *2018 IEEE International Symposium on Workload Characterization (IISWC)*, 2018, pp. 111–121.
- [58] G. Csordas, M. Asiaticci, and P. Jenne, "In Search of Lost Bandwidth: Extensive Reordering of DRAM Accesses on FPGA," in *2019 International Conference on Field-Programmable Technology (ICFPT)*, 2019, pp. 188–196.
- [59] K. Manev, A. Vaishnav, and D. Koch, "Unexpected Diversity: Quantitative Memory Analysis for Zynq UltraScale+ Systems," in *2019 International Conference on Field-Programmable Technology (ICFPT)*, 2019, pp. 179–187.
- [60] S. I. Venieris and C. S. Bouganis, "f-CNN²: A Toolflow for Mapping Multiple Convolutional Neural Networks on FPGAs," in *2018 28th International Conference on Field Programmable Logic and Applications (FPL)*, 2018, pp. 1–8.
- [61] L. Yang, Z. Yan, M. Li, H. Kwon, W. Jiang, L. Lai, Y. Shi, T. Krishna, and V. Chandra, "Co-Exploration of Neural Architectures and Heterogeneous ASIC Accelerator Designs Targeting Multiple Tasks," in *Proceedings of the 57th ACM/EDAC/IEEE Design Automation Conference (DAC)*, 2020.
- [62] E. Baek, D. Kwon, and J. Kim, "A Multi-Neural Network Acceleration Architecture," in *2020 ACM/IEEE 47th Annual International Symposium on Computer Architecture (ISCA)*, 2020, pp. 940–953.
- [63] S. Han, X. Liu, H. Mao, J. Pu, A. Pedram, M. A. Horowitz, and W. J. Dally, "EIE: Efficient Inference Engine on Compressed Deep Neural Network," in *2016 ACM/IEEE 43rd Annual International Symposium on Computer Architecture (ISCA)*, 2016, pp. 243–254.
- [64] K. Guo, L. Sui, J. Qiu, J. Yu, J. Wang, S. Yao, S. Han, Y. Wang, and H. Yang, "Angel-Eye: A Complete Design Flow for Mapping CNN Onto Embedded FPGA," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)*, vol. 37, no. 1, pp. 35–47, 2018.
- [65] J. Yu, A. Lukefahr, D. Palframan, G. Dasika, R. Das, and S. Mahlke, "Scalpel: Customizing dnn pruning to the underlying hardware parallelism," in *Proceedings of the 44th Annual International Symposium on Computer Architecture (ISCA)*, 2017, p. 548–560.
- [66] C. Ding, S. Liao, Y. Wang, Z. Li, N. Liu, Y. Zhuo, C. Wang, X. Qian, Y. Bai, G. Yuan, X. Ma, Y. Zhang, J. Tang, Q. Qiu, X. Lin, and B. Yuan, "CirCNN: Accelerating and Compressing Deep Neural Networks Using Block-Circulant Weight Matrices," in *2017 50th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, 2017, pp. 395–408.
- [67] C. Deng, S. Liao, Y. Xie, K. K. Parhi, X. Qian, and B. Yuan, "PermdNN: Efficient Compressed DNN Architecture with Permuted Diagonal Matrices," in *Proceedings of the 51st Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, 2018, p. 189–202.
- [68] Y. Shen, M. Ferdman, and P. Milder, "Escher: A CNN Accelerator with Flexible Buffering to Minimize Off-Chip Transfer," in *2017 IEEE 25th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*, 2017, pp. 93–100.
- [69] M. Alwani, H. Chen, M. Ferdman, and P. Milder, "Fused-layer cnn accelerators," in *2016 49th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, 2016, pp. 1–12.
- [70] Y. Chen, T. Krishna, J. S. Emer, and V. Sze, "Eyeriss: An Energy-Efficient Reconfigurable Accelerator for Deep Convolutional Neural Networks," *IEEE Journal of Solid-State Circuits (JSSC)*, vol. 52, no. 1, pp. 127–138, 2017.
- [71] A. Montgomerie-Corcoran and C. Savvas-Bouganis, "DEF: Differential Encoding of Featuremaps for Low Power Convolutional Neural Network Accelerators," in *Proceedings of the 26th Asia and South Pacific Design Automation Conference (ASP-DAC)*, 2021.
- [72] U. Aydonat, S. O'Connell, D. Capalija, A. C. Ling, and G. R. Chiu, "An OpenCL™ Deep Learning Accelerator on Arria 10," in *Proceedings of the 2017 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA)*, 2017, p. 55–64.
- [73] J. Fowers, K. Ovtcharov, M. Papamichael, T. Massengill, M. Liu, D. Lo, S. Alkalay, M. Haselman, L. Adams, M. Ghandi, S. Heil, P. Patel, A. Sapek, G. Weisz, L. Woods, S. Lanka, S. K. Reinhardt, A. M. Caulfield, E. S. Chung, and D. Burger, "A Configurable Cloud-Scale DNN Processor for Real-Time AI," in *2018 ACM/IEEE 45th Annual International Symposium on Computer Architecture (ISCA)*, 2018, pp. 1–14.
- [74] Y. Shen, M. Ferdman, and P. Milder, "Maximizing CNN Accelerator Efficiency Through Resource Partitioning," in *Proceedings of the 44th Annual International Symposium on Computer Architecture (ISCA)*, 2017.
- [75] W. Lu, G. Yan, J. Li, S. Gong, Y. Han, and X. Li, "FlexFlow: A Flexible Dataflow Accelerator Architecture for Convolutional Neural Networks," in *2017 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, 2017, pp. 553–564.

- [76] F. Tu, S. Yin, P. Ouyang, S. Tang, L. Liu, and S. Wei, "Deep Convolutional Neural Network Architecture With Reconfigurable Computation Patterns," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems (TVLSI)*, vol. 25, no. 8, pp. 2220–2233, 2017.
- [77] H. Kwon, A. Samajdar, and T. Krishna, "MAERI: Enabling Flexible Dataflow Mapping over DNN Accelerators via Reconfigurable Interconnects," in *Proceedings of the Twenty-Third International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2018, p. 461–475.
- [78] J. Albericio, P. Judd, T. Hetherington, T. Aamodt, N. E. Jerger, and A. Moshovos, "Cnvlutin: Ineffectual-Neuron-Free Deep Neural Network Computing," in *Proceedings of the 43rd International Symposium on Computer Architecture (ISCA)*, 2016, pp. 1–13.
- [79] J. Albericio, A. Delmás, P. Judd, S. Sharify, G. O’Leary, R. Genov, and A. Moshovos, "Bit-Pragmatic Deep Neural Network Computing," in *Proceedings of the 50th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, 2017, pp. 382–394.
- [80] A. D. Lascorz, S. Sharify, I. Edo, D. M. Stuart, O. M. Awad, P. Judd, M. Mahmoud, M. Nikolic, K. Siu, Z. Poulos, and A. Moshovos, "ShapeShifter: Enabling Fine-Grain Data Width Adaptation in Deep Learning," in *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, 2019, p. 28–41.
- [81] 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 *2015 ACM/IEEE 42nd Annual International Symposium on Computer Architecture (ISCA)*, 2015, pp. 92–104.
- [82] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.
- [83] F. N. Iandola, S. Han, M. W. Moskewicz, K. Ashraf, W. J. Dally, and K. Keutzer, "SqueezeNet: AlexNet-level accuracy with 50x fewer parameters and < 0.5 MB model size," *arXiv preprint arXiv:1602.07360*, 2016.
- [84] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "PyTorch: An Imperative Style, High-Performance Deep Learning Library," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2019, pp. 8026–8037.
- [85] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," in *International Conference on Learning Representations (ICLR)*, 2015.
- [86] A. X. M. Chang, A. Zaidy, V. Gokhale, and E. Culurciello, "Compiling Deep Learning Models for Custom Hardware Accelerators," *arXiv preprint arXiv:1708.00117*, 2017.
- [87] Y. Ma, M. Kim, Y. Cao, S. Vrudhula, and J. Seo, "End-to-End Scalable FPGA Accelerator for Deep Residual Networks," in *IEEE International Symposium on Circuits and Systems (ISCAS)*, 2017, pp. 1–4.
- [88] Xilinx, "Adaptive Machine Learning Acceleration," <https://www.xilinx.com/products/acceleration-solutions/xilinx-machine-learning-suite.html>, 2020, [Retrieved: April 3, 2021].
- [89] R. Shi, Y. Ding, X. Wei, H. Li, H. Liu, H. K. . H. So, and C. Ding, "FTDL: A Tailored FPGA-Overlay for Deep Learning with High Scalability," in *57th ACM/IEEE Design Automation Conference (DAC)*, 2020, pp. 1–6.