

Caffeine: Toward Uniformed Representation and Acceleration for Deep Convolutional Neural Networks

Chen Zhang¹, Member, IEEE, Guangyu Sun, Member, IEEE, Zhenman Fang², Member, IEEE, Peipei Zhou, Peichen Pan, and Jason Cong, Fellow, IEEE

Abstract—With the recent advancement of multilayer convolutional neural networks (CNNs) and fully connected networks (FCNs), deep learning has achieved amazing success in many areas, especially in visual content understanding and classification. To improve the performance and energy efficiency of the computation-demanding CNN, the FPGA-based acceleration emerges as one of the most attractive alternatives. In this paper, we design and implement Caffeine, a hardware/software co-designed library to efficiently accelerate the entire CNN and FCN on FPGAs. First, we propose a uniformed convolutional matrix-multiplication representation for both computation-bound convolutional layers and communication-bound FCN layers. Based on this representation, we optimize the accelerator micro-architecture and maximize the underlying FPGA computing and bandwidth resource utilization based on a revised roofline model. Moreover, we design an automation flow to directly compile high-level network definitions to the final FPGA accelerator. As a case study, we integrate Caffeine into the industry-standard software deep learning framework Caffe. We evaluate Caffeine and its integration with Caffe by implementing VGG16 and AlexNet networks on multiple FPGA platforms. Caffeine achieves a peak performance of 1460 giga fixed point operations per second on a medium-sized Xilinx KU060 FPGA board; to our knowledge, this is the best published result. It achieves more than 100× speed-up on FCN layers over prior FPGA accelerators. An end-to-end evaluation with Caffe integration shows up to 29× and 150× performance and energy gains over Caffe on a 12-core Xeon

server, and 5.7× better energy efficiency over the GPU implementation. Performance projections for a system with a high-end FPGA (Virtex7 690t) show even higher gains.

Index Terms—Caffe, CNN FPGA engine, convolutional neural network (CNN), deep learning, hardware/software co-design.

I. INTRODUCTION

IN THE last few years, deep learning has achieved amazing success in many areas, especially in computer vision and speech recognition. Among various deep learning algorithms, convolutional neural networks (CNNs) has become the most popular for visual content understanding and classification, with significantly higher accuracy than traditional algorithms in various computer vision tasks, such as face recognition, image and video processing [1]–[3]. Now CNN is becoming one of the key algorithms in many modern applications, and is attracting enthusiastic interest from both the academic community [1], [3], [4] and industry heavyweights like Google, Facebook, and Baidu [5]–[7]. With the increasing image classification accuracy improvements, the size and complexity of the multilayer neural networks in CNN have grown significantly, as evidenced by the rapid evolution of real-life CNN models, such as AlexNet, ZFNet, GoogleLeNet, and VGG [8]–[11]. This puts overwhelming computing pressure on conventional general-purpose CPUs in light of the recent slowdown of Moore’s law. Therefore, various accelerators—based on GPUs, FPGAs, and even ASICs—have been proposed to improve the performance of CNN designs [12]–[15]. Due to its low power, high energy efficiency, and reprogrammability, the FPGA-based approach is now one of the most promising alternatives and has stimulated extensive interest [13], [16]–[29].

Most prior FPGA acceleration studies on CNN [13], [16]–[22], [26] mainly focus on the convolution layer in CNN, since it is computation-bound and is the most timing-consuming layer. However, this leads to three limitations. First, other unaccelerated layers in CNN cannot get that high energy efficiency from FPGAs. Second, there is significant intermediate data communication overhead between unaccelerated layers on a CPU and the accelerated convolution (CONV) layer on an FPGA through the PCIe connection, which diminishes the overall performance gains [30]. Third, after the FPGA acceleration of the CONV layer, other layers—especially the indispensable fully connected network (FCN) layer that is communication-bound—can become the new bottleneck in CNN. Based on our profiling (detailed in Section II-B), the FCN layer actually occupies more than 50% of the total execution time in CNN after the CONV layer is accelerated on an FPGA.

Manuscript received January 3, 2017; revised May 24, 2017 and September 18, 2017; accepted November 12, 2017. Date of publication October 18, 2018; date of current version October 16, 2019. This work was supported in part by the Center for Domain-Specific Computing Industrial Sponsors, including Fujitsu Labs, Huawei, Intel, Mentor Graphics, and NEC, in part by NSF China under Award 61572045, in part by UCLA/PKU Joint Research Institute, in part by the Chinese Scholarship Council, and in part by the AsiaInfo Inc. This paper was recommended by Associate Editor D. Chen. (*Corresponding author: Chen Zhang.*)

C. Zhang is with the Center for Energy-Efficient Computing and Applications, Peking University, Beijing 100871, China, and also with the Center for Domain-Specific Computing, University of California at Los Angeles, Los Angeles, CA 90095 USA (e-mail: chen.ceca@pku.edu.cn).

G. Sun is with the Center for Energy-Efficient Computing and Applications, Peking University, Beijing 100871, China.

Z. Fang and P. Zhou are with the Center for Domain-Specific Computing, University of California at Los Angeles, Los Angeles, CA 90095 USA.

P. Pan is with Falcon Computing Solutions Inc., Los Angeles, CA 95054 USA.

J. Cong is with the Center for Energy-Efficient Computing and Applications, Peking University, Beijing 100871, China, also with the Center for Domain-Specific Computing, University of California at Los Angeles, Los Angeles, CA 90095 USA, and also with Falcon Computing Solutions Inc., Los Angeles, CA 95054 USA.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TCAD.2017.2785257

To address the above limitations, two of the latest studies [23], [24] started implementing the entire CNN on an FPGA. The work [23] transforms a convolution layer into a regular matrix-multiplication (MM) in the FCN layer, and implements an MM-like accelerator for both layers. The other work [24] takes an opposite approach: it transforms a regular MM into a convolution, and implements a convolution accelerator for both CONV and FCN layers. While these two studies make a good start on accelerating the entire CNN on an FPGA, the straightforward transformation does not consider potential optimizations. They demonstrated a performance of approximately 1.2 giga fixed point operations per second (GOPS), leaving large room for improvement.

In this paper, we aim to address the following key challenges in efficient FPGA acceleration of the entire CNN. First, what is the right mathematical representation for a uniformed acceleration of the computation-bound CONV layer and the communication-bound FCN/deep neural network (DNN) layer?¹ Second, how do we design and implement an efficient and reusable FPGA engine for CNN that maximizes the underlying FPGA computing and bandwidth resource utilization, while still maintaining enough programmability for various layer configurations? Third, how do we provide software programmers an easy-to-use interface such that they can still write high-level network definitions while taking advantage of our Caffeine FPGA engine?

To find the right programming model and efficient implementation for CNN kernels, we first analyze the widely used *regular MM representation* in most CPU and GPU studies. These studies usually convert a convolution layer to a regular MM in the FCN layer, and leverage the well-optimized (with vectorization) CPU libraries like Intel MKL and GPU libraries like cuBLAS for a regular MM [12], [31]. However, the convolutional MM to regular MM transformation requires data duplication in CNN. According to this paper, this duplication results in up to 25 \times more data volume for the input feature maps (FMs) in the CONV layer, and thus diminishes the gains of FPGA acceleration considering that FPGA platforms have extremely limited bandwidth (about 10–20 GB/s [32]) compared to CPU/GPU platforms (up to 700 GB/s [33]). More importantly, according to this paper in Section III-C, the FPGA effective bandwidth is very sensitive to memory access burst lengths, which requires a more careful design for bandwidth-bound FCN layers on FPGAs.

To avoid the data duplication and improve the bandwidth utilization, we propose to use a *convolutional MM representation*. Instead of a straightforward mapping in prior work [24], we batch a group of input FMs in the FCN layer together into a single one in the new representation, which we call *input-major mapping*, so as to improve the data reuse of the weight kernels. Another alternative of this input-major mapping is achieved by reversing the input FM matrix and weight kernel matrix, which we call *weight-major mapping*, based on the observation that the latter matrix is much larger than the former one in the FCN layer. As a result, the weight-major mapping may have more data reuse, especially for the input FMs which are easier to be reused by each weight access than those in the input-major mapping considering the hardware

¹As analyzed in Section II-B, other layers in CNN are relatively simple and have marginal impact on the final performance and FPGA resource consumption. We do implement those layers in the same FPGA, but we will mainly discuss the CONV and FCN layers in this paper for simplicity. Note that the FCN layer is also a major component of DNNs that are widely used in speech recognition. For simplicity, we just use the term “FCN.”

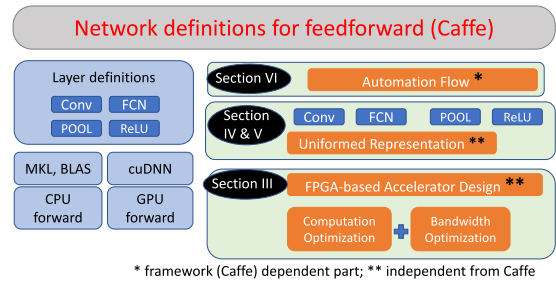


Fig. 1. Overview of Caffeine framework.

resource limitation. Considering the complex data reuse and memory burst access under the hardware resource limitation, it is quite challenging to identify which one is absolutely better between the input-major and weight-major convolutional mappings. For a quantitative comparison, we apply an accurate roofline-based model to guide their design space explorations under different neural network shapes and batch sizes.

Based on the above uniformed representation, we design and implement an efficient and reusable CNN/DNN FPGA accelerator engine called Caffeine.² First, Caffeine maximizes the FPGA computing capability by optimizing multilevel data parallelism within CNN, as well as fine-grained and coarse-grained pipeline parallelism. Second, Caffeine maximizes the underlying memory bandwidth utilization by combining both on-chip and off-chip data reorganizations for the convolutional MM representation. As a result, Caffeine can achieve high performance for both the computation-bound CONV layer and communication-bound FCN layer (more than 100 \times speed-up over prior work [24]). To improve the portability of Caffeine across different FPGA platforms, we design our FPGA accelerator in a systolic-like micro-architecture using high-level synthesis (HLS) so that it can be easily scaled up to a larger design [36]. In addition, Caffeine also supports various CNN layer configurations with different precision requirements (i.e., both floating-point and fixed-point operations). Finally, we further provide an automation flow for software programmers so that they can easily take advantage of our FPGA accelerator engine while still programming the highlevel CNN networks, just as they do for CPUs and GPUs. As a case study, we integrate Caffeine with the industry-standard Caffe deep learning framework [12]. We summarize our Caffeine work in Fig. 1.

In summary, this paper makes the following contributions.

- 1) We propose a uniformed mathematical representation (convolutional MM) for efficient FPGA acceleration of both CONV and FCN layers in CNN/DNN. In addition, we also propose a novel optimization framework based on the roofline model to find the optimal mapping of the uniformed representation to the specialized accelerator. Our optimization framework recommends weight-major mapping and input-major mapping according to platform constraints and NN configurations.
- 2) We customize an HW/SW co-designed efficient and reusable CNN/DNN engine called Caffeine, where the FPGA accelerator maximizes the utilization of computing and bandwidth resource. Caffeine achieves a peak performance of 1460 GOPS for the CONV layer and 346 GOPS for the FCN layer with 8-bit fixedpoint operations on a medium-sized FPGA board (KU060). The

²The name Caffeine comes from Caffe FPGA engine, but it is a generic library and not limited to the CAFFE. It can also be extended for other frameworks like Torch and TensorFlow [34], [35].




Fig. 2. Inference (also known as feedforward) phase in CNN.

performance and energy gains are even higher when projecting to a larger VC709 FPGA board.

- 3) We provide an automation flow for users to program CNN in highlevel network definitions, and the flow directly generates the final FPGA accelerator. We also provide the Caffe–Caffeine integration, which achieves $29\times$ and $150\times$ end-to-end performance and energy gains over a 12-core CPU and $5.7\times$ better energy efficiency over a GPU.

II. CNN OVERVIEW AND ANALYSIS

A. Algorithm of CNNs

As a typical supervised learning algorithm, there are two major phases in CNN: 1) a *training phase* and 2) an *inference (also known as feed-forward) phase*. Since many industry applications train CNN in the background and only perform inferences in a real-time scenario, we mainly focus on the inference phase in this paper. The aim of the CNN inference phase is to get a correct inference of classification for input images. Shown in Fig. 2, it is composed of multiple layers, where each image is fed to the first layer. Each layer receives a number of FMs from a previous layer and outputs a new set of FMs after filtering by certain kernels. The convolutional layer, activation layer, and pooling layer are for FM extraction, and the fully connected layers are for classification. A more detailed introduction is in our conference paper [15].

Convolutional (CONV) layers are the main components of CNN. The computation of a CONV layer is to extract feature information by adopting a filter on feature maps from a previous layer. It receives N feature maps as input and outputs M feature maps.

Pooling (POOL) layers are used to achieve spatial invariance by subsampling neighboring pixels, usually finding the maximum value in a neighborhood in each input feature map.

Activation (ReLU) layers are used to adopt an activation function (e.g., an ReLU function) on each pixel of FMs from previous layers to mimic the biological neuron’s activation [8].

Fully connected (FCN) layers are used to make final predictions. An FCN layer takes “features” in a form of vector from a prior feature extraction layer, multiplies a weight matrix, and outputs a new feature vector, whose computation pattern is a dense matrix-vector multiplication.

B. Analysis of Real-Life CNNs

State-of-the-art CNNs for large visual recognition tasks usually contain billions of neurons and show a trend to go deeper and larger. Table I lists some of the CNN models that have won the ImageNet Large-Scale Visual Recognition Challenge (ILSVRC) contest since 2012. These networks all contain millions of neurons, and hundreds of millions of parameters that include weights and intermediate FMs. Therefore, storing these parameters in DRAM is mandatory for those real-life CNNs. In this paper, we will mainly use the 16-layer VGG16 model [11].

TABLE I
RECENT CNN MODELS THAT WON THE ILSVRC CONTEST

Real-life CNNs	Year	Neurons	layers	Parameters
AlexNet [8]	2012	650,000	8	250 MB
ZFNet [9]	2013	78,000,000	8	200 MB
VGG [11]	2014	14,000,000	16	500 MB

TABLE II
COMPUTATION COMPLEXITY, STORAGE COMPLEXITY, AND EXECUTION TIME BREAKDOWN OF CNN LAYERS IN THE VGG16 MODEL

	CONV	POOL	ReLU	FCN
Comput. ops(10^7)	3E3(99.5%)	0.6(0%)	1.4(0%)	12.3(0.4%)
Storage (MB)	113(19.3%)	0(0%)	0(0%)	471.6(80.6%)
Time% in pure sw	96.3%	0.0%	0.0%	3.7%
after CONV acc	48.7%	0.0%	0.0%	51.2%

Table II shows two key points. *First*, the CONV and FCN layers present two extreme features. CONV layers are very computation-intensive: they contain 19.3% of total data but need 99.5% of computation. FCN layers are memory-intensive: they need 0.4% of arithmetic operations but use 80.6% of the total amount of data. These two layers also occupy most of the execution time (more than 99.9%). *Second*, when CONV is accelerated, the FCN layer becomes the new bottleneck, taking over 50% of computation time. Therefore, we need to accelerate the entire CNN on an FPGA and maximize the use of both FPGA’s computation and bandwidth efficiency. Since a straightforward acceleration of the POOL and ReLU layers is good enough due to their simplicity, in this paper, we will mainly focus on discussing how to accelerate both the CONV and FCN layers.

III. SPECIALIZED CONVOLUTION ACCELERATOR

There are several design challenges that obstacles an efficient convolution accelerator design on an FPGA platform. First, the organization of processing engines (PEs) and buffer banks should be carefully designed in order to process on-chip data efficiently. Second, loop tiling is mandatory to fit a small portion of data on-chip. Third, integration with high-level frameworks such as Caffe not only needs to guarantee optimal performance with customized optimizations, but also requires enough programmability of the specialized hardware. In the following sections, we start from the original convolution code in Fig. 3 and apply a combination of optimizations to achieve a high-performance specialized hardware accelerator design.

A. Convolution Accelerator Overview

The computation pattern of a convolution layer is shown in Fig. 3. Variables in red are all layer parameters, which are set in CNN training and usually differ among layers. Loop tiling and computation model on FPGA. FPGAs have limited BRAM and DSP resources. In order to support real-life CNNs with hundreds of mega bytes or even giga bytes of weights and feature maps, our CNN accelerator design puts all the data in DRAM and caches a part of weights, feature maps and layer definitions in on-chip buffers before they are fed to PEs. Fig. 3 shows a standard convolution layer’s computation procedure. We further apply loop tiling to fit a convolution layer to the FPGA. In CNN structure designs, variables R and C (for the “rows” and “columns”) range from tens to thousands; variables N and M (for the number of input and output feature maps) range from tens to hundreds; KEL (for convolution kernel size) ranges from one to ten. So we do not

```

for(r=0; r<ROW; r++){
  for(c=0; c<COL; c++){
    for(o=0; o<OUT; o++){
      for(i=0; i<IN; i++){
        for(k1=0; k1<K1; k1++){
          for(k2=0; k2<K2; k2++){
            output[o][r][c] +=
              weight[o][i][k1][k2] *
              input[i][S*r+k1][S*r+k1];
          }
        }
      }
    }
  }
}
    
```

Fig. 3. Pseudo code of a convolution layer.

```

External data Transfer Optimization: Section IV/C
for(rr=0; rr<R; rr++){
  for(cc=0; cc<C; cc++){
    for(oo=0; oo<M; oo++){
      for(ii=0; ii<N; ii++){
        On-chip data computation: Section IV/B
        Data_type cache_out [To] [OutSize] [OutSize];
        Data_type cache_inp [Ti] [InSize] [InSize];
        Data_type cache_weight [To] [Ti] [KerSize] [KerSize];
        for(k1=0; k1<K1; k1++){
          for(k2=0; k2<K2; k2++){
            for(r=0; r<Tr; r++){
              for(c=0; c<Tc; c++){
                for(o=0; o<To; o++){
                  for(i=0; i<Ti; i++){
                    output[o][r][c] +=
                      weight[o][i][k1][k2] *
                      input[ii][S*r+k1][S*r+k1];
                  }
                }
              }
            }
          }
        }
      }
    }
  }
}
    
```

Fig. 4. Pseudo code of a tiled convolution layer.

tile on loops “ $k1$ & $k2$ ” because of their small sizes. Other loops are tiled into “tile loops” and “point loops.” Point loops are for on-chip data’s computation, whose optimization is discussed in Section III-B. Tile loops are for bringing data tiles on-chip, whose optimizations are discussed in Section III-C. Fig. 4 shows a pseudo code of a tiled convolution layer.

1) *Software-Definable Parameters*: As described in Fig. 3, a convolution layer is featured by a set of parameters $\langle M, N, R, C, K1, K2, S \rangle$. In order to enable our accelerator’s programmability by software at run time without changing the FPGA bitstream, we set parameters $\langle M, N, R, C, K1, K2, S \rangle$ (which are variables in the blank rectangle in Fig. 4) to be software-definable parameters. In our specialized hardware design, we make them registers to control loop pipelines and could be reset by decoding accelerator-specific instructions during runtime.

2) *Hardware-Definable Parameters*: Except for software-definable parameters, the other parameters in Fig. 4 are hardware-definable parameters, which are labeled in the black rectangle “OutSize & InSize & KerSize” for buffer sizes, “ To & Ti ” for parallel PEs and “Data_type” for floating/fixed<bit-width> point operators. They are set before bit-stream synthesis. Larger values of “OutSize & InSize & KerSize” result in more BRAM utilization and larger values of “ To & Ti ” result in more parallel PEs. Whenever a user wants to switch to a new FPGA device, they can simply reset hardware-definable parameters to customize a new accelerator bitstream with our library.

B. Scalable Accelerator Architecture

Fig. 5 shows the computation structure after our optimization. They are described in the following paragraphs.

Multilevel Data Parallelism: We implement two levels of data parallelism as suggested in [13] for the sake of better hardware utilization and circuit simplicity: 1) parallelism in computing multiple output feature maps and 2) parallelism in processing multiple input feature maps for each output feature

```

1. for(k1=0; k1<K1; k1++){
2. for(k2=0; k2<K2; k2++){
3. for(r=0; r<Tr; r++){
4. for(c=0; c<Tc; c++){
  #pragma for LOOP PIPELINE(fine-grained pipeline parallelism)
5. for(o=0; o<To; o++){
  #pragma for LOOP UNROLL(multilevel parallelism)
6. for(i=0; i<Ti; i++){
  #pragma for LOOP UNROLL(multilevel parallelism)
7. cache_out[o][r][c] +=
   cache_weight[o][i][k1][k2] * cache_input[ii][S*r+k1][S*r+k1];
   }
   }
   }
 }
 }
 }
 }
 }
    
```

Fig. 5. Pseudo code of optimized on-chip computation.

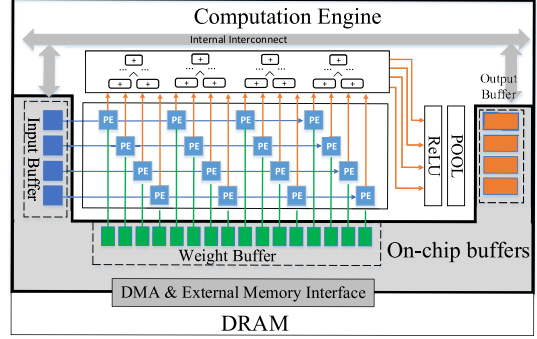


Fig. 6. Scalable accelerator architecture design.

map. Each PE is an arithmetic multiplication of input feature map pixels and corresponding weights.

1) *Combined Data Pipeline Streaming: Fined Grained Pipeline*: In order to fully utilize computation resource, our accelerator aims to achieve a pipeline initial interval (II) of 1; i.e., each PE is able to process a pair of input data on every cycle. We use a polyhedral-based optimization framework [37] to optimize the pipelining schedule by permuting the “ p ,” “ q ,” “ r ,” “ c ” loop levels to avoid loop carried dependence. Since pooling layers and ReLU layers are usually an optional layer following convolution layers, we also implement them in the instruction pipeline so that they can be processed immediately on convolution’s output. They can also be bypassed if there is no such layer following a convolution layer. They can also be configured through software-definable parameters. In coarse-grained pipelining, we use the double buffering technique to pre-fetch the next data tile for each PE so that the computation time can overlap with the data transfer overhead from the device DRAM to FPGA’s BRAM.

2) *Scalable Architecture*: The computations shown in Fig. 5 are a typical map and reduction pattern. We further use a systolic-like architecture to implement the above computations so that the hardware design could be scalable to a larger device with more parallel engines. Fig. 6 presents an overview of our scalable accelerator architecture, which is designed in portable HLS. Similar methods are also explored in work [25].

C. Accelerator Bandwidth Optimization

Since the FCN layer is bandwidth sensitive, we need to be careful about the accelerator bandwidth optimization. In order to have a sense of effective FPGA DRAM bandwidths under different memory access patterns, we test this on the latest Kintex Ultrascale KU060 FPGA as a representative with Xilinx SDAccel 2015.3 flow. Fig. 7(f) plots the effective DRAM bandwidth under different memory access burst lengths and bit-widths. We make two observations in efficient FPGA DRAM bandwidth utilization. First, the effective FPGA bandwidth (y-axis) goes up with the increase of burst

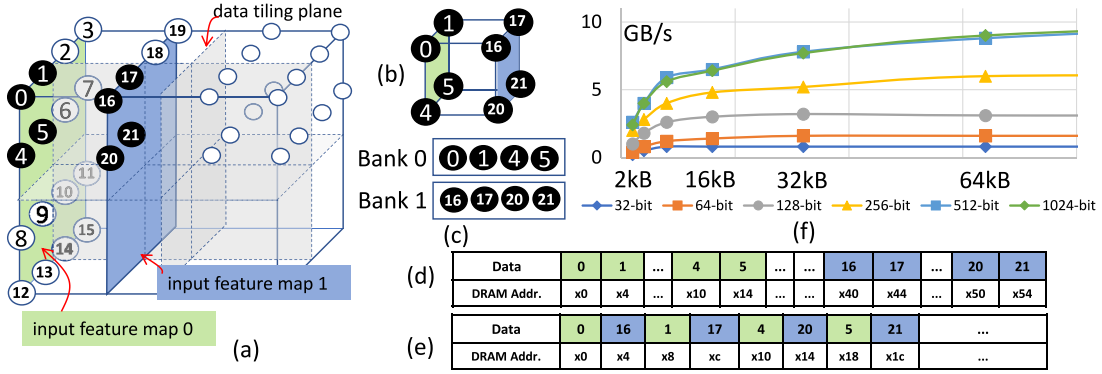


Fig. 7. Bandwidth optimization by DRAM layout reorganization. (a) Logical 3-D data layout. (b) Piece of data tile. (c) Physical data layout in on-chip buffer per BRAM bank 0 and 1. (d) Row-major data layout in DRAM space. (e) Proposed data layout in DRAM space. (f) Effective FPGA DRAM bandwidth with access length and bit-width.

Algorithm 1 DRAM Allocation and Data Organization

Input:

Parameters for feature map tensor shape, $[M, R, C]$
 Parameters for Input BRAM buffer, $[T_m, T_r, T_c]$

Output:

A linear list of tensor index in DRAM,
 List = $\{a_i \mid i \in [0, M \times R \times C]\}$
 1: **for** each $[i, j, k] \in [T_m, T_r, T_c]$ **do**
 2: **for** each $[i_t, j_t, k_t] \in [T_m, T_r, T_c]$ **do**
 3: $Tile_Size = T_m \times T_r \times T_c$
 4: $Tile_Addr = (i + j \times \frac{M}{T_m} + k \times \frac{M}{T_m} \times \frac{R}{T_r}) \times Tile_Size$
 5: $Point_Addr = i_t + j_t \times T_m + k_t \times T_m \times T_r$
 6: $Addr = Tile_Addr + Point_Addr$
 7: Append $Addr$ in the mappingList
 8: **end for**
 9: **end for**

length (x -axis) and finally flattens out above some burst length threshold, which is about 128 KB on 512-bit bitwidth in our experiment. Limited burst length will greatly degrade actual bandwidth performance, like 1 GB/s on 1 KB memory burst access. Second, longer interface bit-width can achieve higher peak bandwidth. The maximum effective bandwidth of 10 GB/s (about 83% of theoretical 12.8 GB/s) can be only reached at 512 bit-width and above, when the burst length is above 128 KB.

1) *Off-Chip Bandwidth Optimization Opportunity*: As analyzed earlier, the burst length and bit-width of DRAM interface are two dominating factors for FPGAs' effective bandwidth. However, the widely used data tiling technique usually results in a discontinuous DRAM access for the row-major data layout in DRAM. We illustrate this using an example in Fig. 7. Fig. 7(a) describes four input FMs in a logical 3-D representation, each with a size of 4×4 . Each dimension is tiled by 2 so that each tile has $2 \times 2 \times 2 = 8$ elements in total. The first tile of input FMs is shown in Fig. 7(b). Fig. 7(d) presents its corresponding data layout in DRAM in a row-major representation, which results in four discontinuous blocks. Therefore, it requires four DRAM accesses, each with a burst length of 2 floating points. This results in a pretty low memory bandwidth utilization and can greatly degrade the overall performance, especially for the bandwidth-intensive FCN layers.

2) *On-Chip Buffer Access Optimization Opportunity*: BRAM banks are usually organized for maximum parallel data access from massive parallel PEs. As illustrated in Fig. 7(c), elements (0, 1, 4, 5) from input FM 0 should be put in bank 0, while elements (16, 17, 20, 21) from input FM 1 should be put

in bank 1. However, such requirements would cause on-chip bank write conflicts using the original DRAM organization in Fig. 7(d). When loading continuous data blocks (0, 1) from DRAM to BRAM (similar for other pairs), they will be written to the same bank 0, which causes bank write conflicts and introduces additional overhead.

3) *Optimization*: To improve the effective memory bandwidth, we reorganize the DRAM layout as illustrated in Fig. 7(e). First, we move the data for an entire tile to a continuous space to improve the memory burst length and bit-width. Second, we interleave the data for different BRAM banks to reduce bank read/write conflicts. Algorithm 1 presents the method for transforming the cube indexes in Fig. 7(a) to indexes in linear DRAM space as shown in Fig. 7(e). Weight and output tensors use a similar method.

IV. UNIFORMED CONV AND FCN REPRESENTATION

A. Prior Representation on CPUs and GPUs

Prior CPU and GPU studies [12], [31] most often used the regular MM representation so as to leverage the well-optimized CPU libraries like Intel MKL and GPU libraries like cuBLAS. To achieve this uniformed acceleration, they convert a convolutional MM in the CONV layer to a regular MM in the FCN layer. However, such a transformation comes at the expense of data duplication, which diminishes the overall performance gains in bandwidth-limited FPGA platforms [23]. Fig. 9 illustrates the data duplication overhead by using MM for the CONV layer computation in AlexNet and VGG16 models. Compared to the original convolutional MM representation, the regular MM representation introduces $7.6 \times$ to $25 \times$ more data for the input FMs, and $1.35 \times$ to $4.8 \times$ more data for intermediate FMs and weights, which makes the CONV layer communication-bound.

B. New Representation Adapted for FPGAs

To avoid the data duplication overhead, we propose to use the convolutional MM representation, and transform the regular MM in the FCN layer to the convolutional MM in the CONV layer. Instead of a straightforward mapping as proposed in [24], we propose two optimized mappings to improve the data reuse and bandwidth utilization: 1) *input-major mapping* and 2) *weight-major mapping*.

1) *Straightforward Mapping*: For FCN shown in Fig. 8(a), an input vector with size N will do pairwise multiplication with a weight vector of size N and accumulate the results

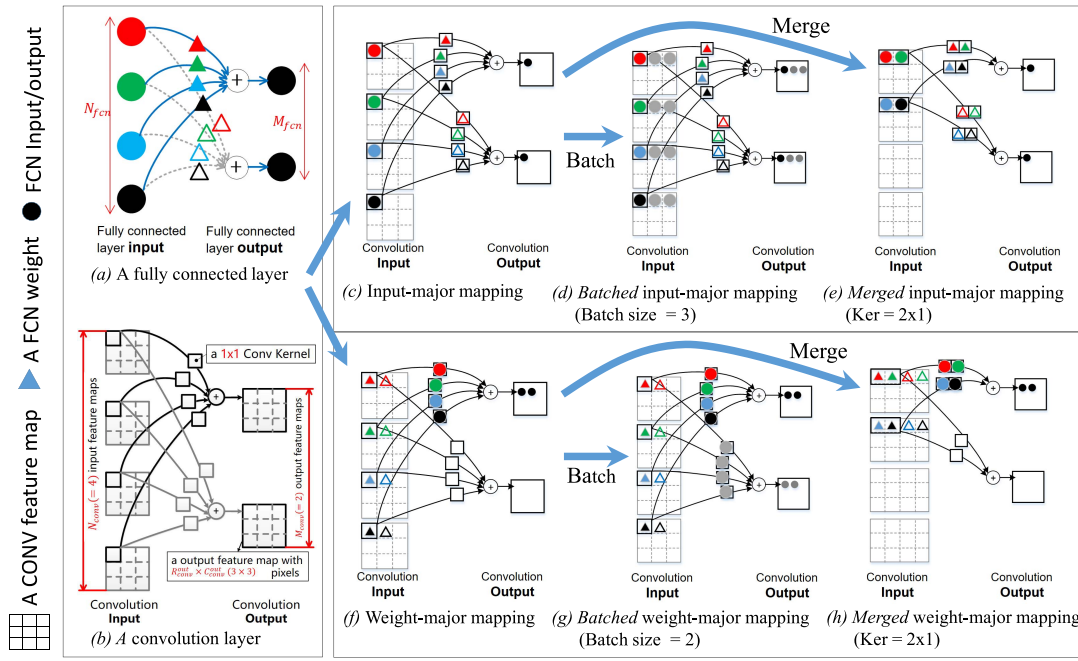


Fig. 8. Input-major and weight-major mapping from the FCN layer to the CONV layer. (a) Fully connected layer. (b) Convolution layer. (c) Input-major mapping. (d) Batched input-major mapping (batch size = 3). (e) Merged input-major mapping (Ker = 2×1). (f) Weight-major mapping. (g) Batched weight-major mapping (batch size = 2). (h) Merged weight-major mapping (Ker = 2×1) batch.




Fig. 9. Overhead of regular MM for CONV.




Fig. 10. Input-major mapping.

to get one output value. There are M weight vectors and M output values. For CONV shown in Fig. 8(b), similarly, N FMs will convolve with N weight kernels, and then element-wise addition is done for the convolution results to get one output FM. There are M sets of weight kernels, and we will get M output FMs.

In a straightforward mapping, each element in an input $1 \times N$ vector of FCN maps to one input FM sized as $R_i = 1, C_i = 1$ of CONV. And each element in an $1 \times N$ weight vector of FCN maps to one weight kernel of CONV sized as $K_1 = 1, K_2 = 1$. This can be viewed in Fig. 8(c) when batch size is 1. Prior work [24] first attempted to implement both CONV and FCN using a similar mapping, and demonstrated a performance of nearly 1.2 GOPS, leaving large room for improvement.

2) *Input-Major Mapping*: In real-life CNNs, multiple input images are processed in a batch to improve throughput. Therefore, in our *input-major mapping*, we can map a batch of elements from different input vectors in FCN to the same input FM in CONV. As a result, the data reuse of FCN weight kernels is improved when convolving the elements from different images in the batched input FMs. When batch size is batch, there are batch input vectors in FCN, and the reuse ratio of FCN weight kernels is batch. Note batch cannot be too large in the real-time inference phase.

To better illustrate the input-major mapping, we use Fig. 8(d) to show how we map FCN to CONV when batch = 3, $N = 4$ and $M = 2$. The four elements of the first input vector are mapped to the first element of each input FM, and the four elements of the second input vector are mapped to the second element of each input FM. Both the weight kernel size and stride size are 1×1 . While the weight kernels slide across the input FMs, they will generate batch elements in each output FM. In addition to the improved data reuse for weight kernels, this batching also improves the memory access burst length of FCN input and output FMs, which improves the bandwidth utilization, as explained in Section III-C.

Another way to improve the memory burst length is to increase the weight kernel size ker and batch ker elements within a single weight (or input) vector in FCN to the same weight kernel (or input FM) in CONV. Fig. 8(e) depicts an example where we change ker from 1×1 to 1×2 . Compared to Fig. 8(c), two weights are grouped in one weight kernel, and two input FMs are grouped into one input FM. Accordingly, stride size changes with ker to 1×2 .

Table V *column FCN-input* lists the parameters after input-major mapping from FCN to CONV. The number of input FMs decreases to (N/ker) , and the number of elements in one

input FM increases to $\text{batch} \times \text{ker}$. The number of elements in an output FM is batch .

3) *Weight-Major Mapping*: As another alternative to improve the data reuse and bandwidth utilization, we propose *weight-major mapping*, where input vectors of FCN map to weight kernels of CONV, and weight vectors of FCN map to input FMs of CONV. As shown in Fig. 8(f), every input vector of FCN in a batch transforms to one set of weight kernels. Weight vectors of FCN are aligned in input FMs in a way that weight elements at the same position of all weight vectors are grouped into the same input FM. Therefore, each FCN input can be reused M_{fcn} times (if it can be buffered on-chip) during the convolution, which greatly improves the data reuse. In addition, the memory burst length of FCN weights and FCN output FMs are greatly improved as well. Similarly, the batch size improves the data reuse of FCN weights and improves the memory burst length of FCN input FMs in weight-major mapping. In addition, it decides the number of FCN output FMs that are available to be processed simultaneously.

Similar to input-major mapping, we can increase the kernel size ker in FCN input FMs to increase the memory burst length, with an example of $\text{ker} = 2$ shown in Fig. 8(g). Table V column *FCN-weight* lists the parameters for weight-major mapping from FCN to CONV.

4) *Uniformed Representation*: Since FCN now maps to CONV, either using *input-major mapping* or *weight-major mapping*, we use a uniformed representation (column *uniformed*) for all cases in Table V. Considering the complex data reuse and memory burst access under different batch and kernel sizes, as well as the hardware resource constraints, it is quite challenging to identify whether input-major mapping or weight-major mapping is better. Therefore, we will conduct a quantitative design space exploration of concrete parameters in Section V.

V. DESIGN SPACE EXPLORATION

In this section, we discuss how to find the optimal solution of mapping a CNN/DNN onto our accelerator architecture. In Section V-A, we first use one concrete example to give readers a sense of the differences of the two mapping methods on their memory access features; and Section V-B gives formal formulations. In system performance, computation capability and memory access are two dominating factors to final achievable performance. We propose to use roofline models to accurately formulate the performance. In addition, as described in Fig. 7(f), DRAM's effective bandwidth is sensitive to access patterns. We further take DRAM bandwidth features in our formulations. In Sections V-B and V-C, we present our systematic methods of performance analysis and design space exploration.

A. Case Study

We use the real case of a fully connected layer from VGG16 model (FCN 1) in our case study. It has an input vector of 25088 and an output vector of 4096. We study the differences of two mapping methods to an accelerator with a hardware configuration of $\langle Tm, Tn, Tr \times Tc, \text{KernelSize} \rangle = \langle 32, 32, 4096, 3 \rangle$. In order to simplify the explanation, let us first discuss the mapping of Fig. 8(c) and (f) in this section. More complicated situations will be discussed in Section V-B with mathematical formulations.

Fig. 8(a) shows the original fully connected layer with “25088” inputs, “4096” outputs and “25088 × 4096” weights. According to the input-major mapping method described in

TABLE III
INPUT-MAJOR MAPPING (BATCH_SIZE = 1)

	input	weight	output
Original FCN layer	25088×1	25088×4096	4096×1
HW buffer name	input buffer	weight buffer	output buffer
HW buffer size	32×4096	$32 \times 32 \times 3 \times 3$	32×4096
Size of data tile	32×1	$32 \times 32 \times 1 \times 1$	32×1
Burst length	32	1024	32
# of memory access	784	100,352	128




Fig. 11. Weight-major mapping.

TABLE IV
WEIGHT-MAJOR MAPPING (BATCH_SIZE = 1)

	input	weight	output
Original FCN layer	25088×1	25088×4096	4096×1
HW buffer name	weight buffer	input buffer	output buffer
HW buffer size	$32 \times 32 \times 3 \times 3$	32×4096	32×4096
Size of data tile	$32 \times 1 \times 1$	32×4096	1×4096
Burst length	32	131072	4096
# of memory access	784	784	1

Section IV, the corresponding tiling method is 32×32 , which is shown as those bold connections in Fig. 8(a). Fig. 8(b) shows the input/weight/output accelerator buffers and the tiled FCN layer's mapping into corresponding buffers. So the total number of memory accesses (bursts) to the input vector is $(25088 \times 1) \div (32 \times 1) = 784$; the total number of memory accesses to the weights is $(25088 \times 4,096) \div (32 \times 32 \times 1 \times 1) = 100352$; the total number of memory accesses to the output vector is $(4096 \times 1) \div (32 \times 1) = 128$. Table III summarizes the total number of memory accesses and the burst length in each memory access. Similarly, Fig. 11 presents the weight-major mapping, where the tile size for the input buffer can be much larger (as discussed in Section IV). Table IV shows its corresponding data.

By comparing Tables III and IV, we can see that the weight-major mapping has significantly less numbers of memory accesses and longer burst lengths than the input-major mapping in this case study.

B. Analytical Comparison of Two Mapping Methods

In this section, we give a formulation of memory access patterns by considering workload size and platform constraints. We denote the hardware configuration of our accelerator as “number of (input, output, weight) buffer = $\langle Tn, Tm, Tn \times Tm \rangle$ ” and “size of each buffer = $\langle Tr \times Tc, Tr \times Tc, K \times K \rangle$,” which are exactly the following notations in Fig. 5.

Given the uniformed representation in Table V, the number of memory accesses can be correspondingly calculated as shown in Table VI. In this table, M, N, R, C , and K are following notations from Table V's column 2 (uniformed). When considering input-major mapping's and weight-mapping's concrete memory access behavior, we simply replace uniformed notations with the FCN-input or FCN-weight in Table V.

TABLE V
UNIFORMED REPRESENTATION PARAMETERS FOR CONV, FCN
INPUT-MAJOR MAPPING, AND FCN WEIGHT-MAJOR MAPPING

	Uniformed	Conv	FCN-Input	FCN-Weight
Input FM #	N	N_{conv}	$N_{fcn/ker}$	$N_{fcn/ker}$
Input FM size	$R_i \cdot C_i$	$R_{conv}^{in} \cdot C_{conv}^{in}$	batch \cdot ker	$M_{fcn} \cdot ker$
Output FM #	M	M_{conv}	batch	batch
Output FM size	$R_o \cdot C_o$	$R_{conv}^{out} \cdot C_{conv}^{out}$	batch	M_{fcn}
Kernel size	$K_1 \cdot K_2$	$K_1 \cdot K_2$	ker	ker
Stride	$S_1 \cdot S_2$	$S_1 \cdot S_2$	ker	ker

TABLE VI
NUMBER OF DRAM ACCESSES

	Uniformed	Input-major	Weight-major
Input	$\lceil \frac{N}{T_n} \rceil \lceil \frac{R_i \cdot C_i}{Tr \cdot Tc} \rceil$	$\lceil \frac{N_{fcn/ker}}{T_n} \rceil$	$\lceil \frac{N_{fcn/ker}}{T_n} \rceil$
Weight	$\lceil \frac{N}{T_n} \rceil \lceil \frac{M}{T_m} \rceil$	$\lceil \frac{N_{fcn/ker}}{T_n} \rceil \lceil \frac{M_{fcn}}{T_m} \rceil$	$\lceil \frac{N_{fcn/ker}}{T_n} \rceil \lceil \frac{batch}{T_m} \rceil$
Output	$\lceil \frac{M}{T_m} \rceil \lceil \frac{R_o \cdot C_o}{Tr \cdot Tc} \rceil$	$\lceil \frac{M_{fcn}}{T_m} \rceil \lceil \frac{batch}{Tr \cdot Tc} \rceil$	$\lceil \frac{batch}{T_m} \rceil \lceil \frac{M_{fcn}}{Tr \cdot Tc} \rceil$

The remaining part of Table VI summarizes input-major and weight-major mapping of memory access. Their major differences are in their “weight” and “output” DRAM access. For weight, input-major and weight-major mapping methods of DRAM accesses are $\lceil (N_{fcn/ker}/T_n) \rceil \lceil (M_{fcn}/T_m) \rceil$ and $\lceil (N_{fcn/ker}/T_n) \rceil \lceil (batch/T_m) \rceil$, respectively. The two formulations are almost the same except for “ M_{fcn} ” and “batch.” Real-life network configuration’s M_{fcn} is usually in a scale of thousands and T_m is in tens (it is constrained by DSP and BRAM resources, for example “ $(T_m, T_n) = (32, 32)$ ” uses 1024 multiplication and accumulation operators), while the tunable parameter batch is smaller or equal to T_m . So $\lceil (M_{fcn}/T_m) \rceil$ would be significantly larger than $\lceil (batch/T_m) \rceil$. For the output’s DRAM transfer, the considering denominator is “ $Tr \cdot Tc$ ” which is for FMs and usually very large. At our setting, $Tr \cdot Tc$ is $226 \times 30 = 6780$. With similar deductions, $\lceil (M_{fcn}/T_m) \rceil \lceil (batch/Tr \cdot Tc) \rceil$ would also be significantly larger than $\lceil (batch/T_m) \rceil \lceil (M_{fcn}/Tr \cdot Tc) \rceil$.

Thus, given an accelerator information $\langle T_m, T_n, Tr, Tc \rangle$ and FCN workload configuration $(N_{fcn}, M_{fcn}, ker, batch)$, we are able to calculate all DRAM traffic following formulations in Table VI. With the above formulations, we estimate the attainable performance by jointly considering both computation capability and bandwidth performance in the next section.

C. Revised Roofline Model for Caffeine

1) *Original Roofline Model*: The roofline model [38] is initially proposed in multicore systems to provide insight analysis of attainable performance by relating processors’ peak computation performance and the off-chip memory traffic. Equation 1 formulates the attainable throughput of an application on a specific hardware platform. Floating-point performance (GFLOPS) is used as the metric of throughput. The actual floating-point performance of an application kernel can be no higher than the minimum value of two terms. The first term describes the peak floating-point throughput provided by all available computation resources in the system, or computational roof. Operations per DRAM traffic, or the computation-to-communication (CTC) ratio, feature the DRAM traffic needed by a kernel in a specific system implementation. The second term bounds the maximum floating-point performance that the memory system can support for a given CTC ratio

$$\text{AttainablePerf.} = \min \left\{ \begin{array}{l} \text{Computational Roof} \\ \text{CTC Ratio} \times \text{BW.} \end{array} \right. \quad (1)$$

Previous work [13] uses the roofline model to optimize the FPGA accelerator design. However, the original roofline

model used in [13] ignores the fact that input/output/weight arrays have different data volumes in each tile. According to Fig. 7(f), different burst lengths and access patterns will result in different effective bandwidths. Thus, different designs have different final bandwidth rooflines, which makes the original roofline-based method’s prediction for bandwidth-intensive applications extremely inaccurate, like fully connected layers. As proposed in [13], the original total number of DRAM access in one layer’s computation is given by the following equation, where β denotes the corresponding size of input/output/weight data tile, and α denotes the number of times of corresponding data transfer for input/output/weight data

$$\text{DRAM_Access} = \sum_i^{\text{in, weight, out}} \alpha_i \times \beta_i. \quad (2)$$

In fact, (2) does not accurately model the total DRAM traffic. For example, as shown in Fig. 7(f), the effective bandwidth on 1 KB burst DRAM access is only 1 GB/s—10 \times lower than the maximum effective bandwidth of 10 GB/s. Therefore, the original roofline model becomes extremely inaccurate in bandwidth-sensitive workloads because it actually takes 10 \times longer time to make the data transfer than expected. Therefore, we would like to multiply a normalization factor of 10 \times on the original DRAM access number to approach the accurate effective DRAM traffic.

2) *Revised Roofline Model for Caffeine*: In general, we propose to normalize the DRAM traffic of input/output/weight accesses to the maximum effective bandwidth with a normalization factor γ

$$\text{DRAM_Access} = \sum_i^{\text{in, weight, out}} \gamma_i \times \alpha_i \times \beta_i \quad (3)$$

where γ is defined by $\gamma = \max_bandwidth/f(\beta)$. The f function is given by the curve of effective bandwidth with respect to the burst length, shown in Fig. 7(f).

Given a specific set of software-definable parameters for one layer $\langle N, R_i, C_i, M, R_o, C_o, K, S \rangle$ and a specific hardware definable parameter $\langle T_i, T_o, Tr, Tc \rangle$, as described in Section III-A, we can determine the x -axis and y -axis value in the roofline model by computing its computational performance and CTC ratio.

Similar to [13], the computational performance is given by

$$\begin{aligned} \text{Comput. Perf.} &= \frac{\text{total computation operations}}{\text{execution cycles}} \\ &= \frac{2 \cdot N \cdot M \cdot R_o \cdot C_o \cdot K_1 \cdot K_2}{\lceil N/T_i \rceil \cdot \lceil M/T_o \rceil \cdot R_o \cdot C_o \cdot K_1 \cdot K_2}. \end{aligned} \quad (4)$$

Our revised CTC ratio is given by

$$\begin{aligned} \text{CTC ratio} &= \frac{\text{total computation operations}}{\text{total DRAM traffic}} \\ &= \frac{2 \cdot N \cdot M \cdot R_o \cdot C_o \cdot K_1 \cdot K_2}{\gamma_{in} \cdot \alpha_{in} \cdot \beta_{in} + \gamma_{wght} \cdot \alpha_{wght} \cdot \beta_{wght} + \gamma_{out} \cdot \alpha_{out} \cdot \beta_{out}}. \end{aligned} \quad (5)$$

Given above revised roofline mode, we can have a design space exploration to find the optimal method mapping from the uniformed representation to our hardware accelerator.

D. Design Space Exploration

Since optimizing the CONV layer with the roofline model has been extensively discussed in [13], and there is a space

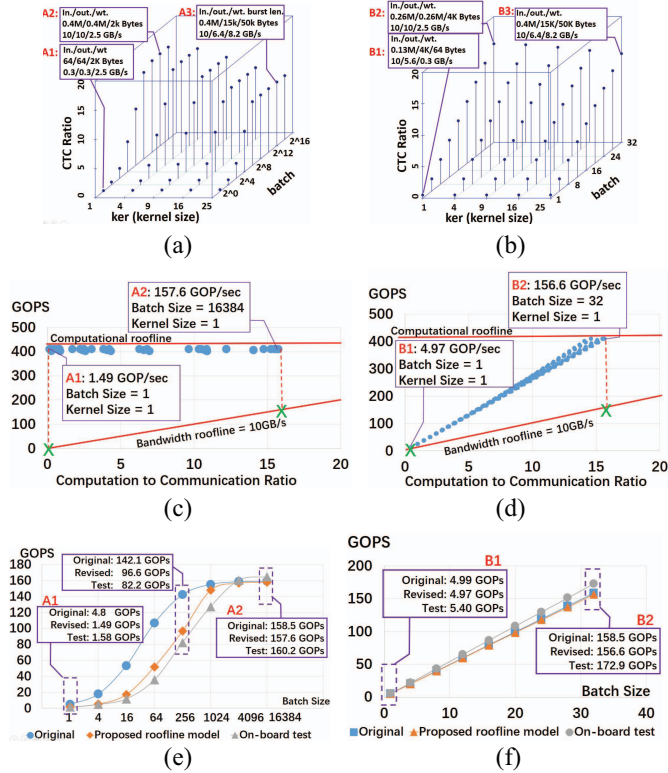


Fig. 12. (a) Input-major and (b) weight-major’s CTC ratio. (c) Input-major and (d) weight-major of revised roofline model. (e) Input-major and (f) weight-major comparison among original, revised roofline models and on-board tests of FCN.

constraint, we mainly focus on optimizing the mapping of the FCN layer to the uniformed representation using our revised roofline model. Specifically, it is a problem of choosing input-major/weight-major mapping methods and the optimal batch and ker parameters, given the FCN layer configuration and hardware configuration.

We use the VGG16 model’s FCN layer 1 as an example; it has an input of 25088 (N_{fcn}) neurons and output of 4096 (M_{fcn}) neurons, whose notations follow Table V. *Batch* and *ker* are tunable parameters for mapping FCN to the uniformed representation as described in Section IV. We use the hardware configuration from the Kintex Ultrascale KU060 platform and set hardware definable parameters as $\langle T_o, T_i, T_r, T_c, T_{K1}, T_{K2} \rangle = \langle 32, 32, 6272, 25 \rangle$. We choose our tile sizes based on the guidance of [13] to maximize the FPGA resource utilization. Users can configure their own tile sizes.

1) *FCN Input-Major Mapping*: Fig. 12(a) presents the design space of FCN input-major mapping in terms of CTC ratio under various batch (*batch*) and kernel (*ker*) sizes. First, given a fixed *ker*, the CTC ratio increases with *batch*, because *batch* FCN inputs reuse FCN weights, and memory burst length is increased by *batch* which results in higher effective DRAM bandwidth. The CTC ratio flattens out when *batch* is bigger than on-chip BRAM size. Second, given a fixed *batch*, the CTC ratio increases with *ker* when *batch* is small, because this increases memory burst length and thus benefits effective DRAM bandwidth. Finally, since the size of input FM is given by *batch* · *ker* in Table V, the maximum *batch* that could be cached in on-chip BRAM decreases when *ker* increases. Therefore, the CTC ratio decreases when *ker* increases on a large *batch*, because the output FM burst length (given by *batch* according to Table V) decreases. In the

input-major mapping, the maximum CTC ratio is achieved with a parameter $\langle \text{batch}, \text{ker} \rangle = \langle 16384, 1 \rangle$.

Fig. 12(c) presents input-major mapping’s attainable performance using our revised roofline model. Each point represents an implementation with its computation performance in GOPS and CTC ratio estimation, which are decided by parameters $\langle \text{batch}, \text{ker} \rangle$ according to our model. The red line (bandwidth roofline, slope = 10 GB/s) represents the max DRAM bandwidth that this FPGA platform supports. Any point located above this line indicates that this implementation requires higher bandwidth than what the platform can provide. Thus, it is bounded by platform bandwidth, and the attainable performance is then decided by the bandwidth roofline. From this figure, we can see that all implementations of FCN with input-major mapping are bounded by bandwidth. The highest attainable performance is achieved at the highest CTC ratio, where $\langle \text{batch}, \text{ker} \rangle = \langle 16384, 1 \rangle$, and this *batch* size is unreasonable in a real-time inference phase.

Fig. 12(e) presents the on-board test performance of input-major mapping and the comparison between performance estimations from original and revised roofline models. Our revised roofline model is much more accurate than the original one, and our estimated performance is very close to that of the on-board test.

2) *FCN Weight-Major Mapping*: Fig. 12(b) presents the design space of FCN weight-major mapping in terms of CTC ratio under various *batch* (*batch*) and *ker* (*ker*) sizes. As illustrated in Section IV-B3, *batch* represents the number of concurrent PEs processing different output FMs in weight-major mapping. Due to the FPGA resource constraints, we can only put 32 such PEs in the KU060 FPGA. Therefore, we set an up-limit of 32 to *batch* in weight-major mapping, which is pretty small. Given a fixed *ker*, the CTC ratio increases with *batch* since it increases the data reuse of FCN weights and the memory burst length of FCN inputs. The size of *ker* has marginal impact on weight-major mapping because it has pretty good bandwidth utilization, even for *ker* = 1.

Fig. 12(d) presents weight-major mapping’s attainable performance using our revised roofline model. Similar with input-major mapping, all implementations of weight-major mapping are bounded by bandwidth. In addition, small *batch* size also leads to lower computational performance due to less number of concurrent PEs in weight-major mapping. The highest attainable performance is achieved at the highest CTC ratio, where $\langle \text{batch}, \text{ker} \rangle = \langle 32, 1 \rangle$, which is reasonable in a real-time inference phase.

Fig. 12(e) presents the on-board test performance of weight-major mapping and the comparison between performance estimations from original and revised roofline models. Different than input-major mapping, weight-major mapping has very good data reuse as well as good effective bandwidth, as illustrated in Section IV. So the proposed roofline model is only slightly better than original model, and both models are close to the on-board test. In addition, weight-major mapping presents better performance than input-major mapping in cases of small *batch* sizes.

Due to the advantages of weight-major over input-major mapping in small *batch* sizes, in the remainder of this paper we will use weight-major mapping for the FCN layer with the best design point.

E. Design Space Exploration on Speech Applications

Previous sections are based on CNNs, which are mainly for computer vision tasks. However, in many other areas such as




Fig. 13. Design space exploration for hidden layers in [39]. (a) Input-major. (b) Weight-major.




Fig. 14. Design space exploration for bottleneck layers in [40]. (a) Input-major. (b) Weight-major.

speech and auto-encoder, fully connected neural network is also a major type of workload, such as networks presented on work [39]–[44].

Fig. 13 presents a design space of a hidden layer in work [39], which has a very typical shape like other FCN workloads.

Fig. 14 presents a design space of a bottleneck network, which is also frequently used in prior work [40]–[42]. Significantly less number of neurons is bottleneck layer’s major difference to typical networks. Compared to regular NN layers with 2048 or more neurons, bottleneck layers usually have much less neurons, such as 20 to 40 neurons in work [40]. This will greatly influences CTC ratios. As is presented in Fig. 14(b), solution “F4” has the highest performance in weight-major mapping method. On the same configuration (batch_size = 32, kernel_size = 16), weight-major mapping method wins input-major mapping. However, the highest input-major mapping solution “F3” achieves nearly 155 GPOS, which is almost the double of that of solution F4. Actually, input-major mapping wins when batch_size is larger than 128.

Under real service scenarios, there is a tradeoff between low latency and high-throughput when users use small networks such as bottleneck NN. When latency is more important, we recommend weight-major mapping which achieve higher performance on small batches. Otherwise we recommend, input-major mapping for throughput. Since the choice depends on real scenarios, we left the adventure for users.

VI. FROM HIGH-LEVEL NETWORK DESCRIPTION TO SPECIALIZED CNN ACCELERATOR

Programming hardware for nonexperts is usually very difficult. Therefore, we propose an automation flow to apply our proposed optimizations discussed in those sections to compile the high-level network descriptions directly into the FPGA-based specialized hardware accelerator. Our automation flow has two cooperating sides: 1) software automation, which provides a compiler to map the high-level network definitions to customized instructions for our specialized hardware and 2) hardware automation, which is responsible for generating a new FPGA accelerator bitstream.

A. Software Automation

With the proposed software-definable accelerator design, we implement an automated flow to bridge the neural network oriented high-level domain-specific language to our customized accelerator design. Fig. 15 presents the automation flow from Caffe standard inputs; these are defined in prototxt and caffemodel files in our hardware-optimized model, which includes all of the accelerator instructions (network definitions), DRAM space allocations and accelerator-specific weight data reorganizations. Overall, the key steps of our automation flow include the following.

- 1) *Network Parser (Network Model Parser and Compilation)*: We first parse the structure of CNN’s CONV/ReLU/POOL/FCN layers from Caffe’s network definition file, which is described in prototxt file, to a structured DAG-based data type to describe CNN’s data flow. In addition, we read in the original CNN layers’ weights and biases stored in Caffe’s caffemodel file. This is the only part of our automation flow that is specific to Caffe; all other parts can be reused in other frameworks.
- 2) *CNN Representation Transformation*: In the next step we transform FCN in the CNN DAG to a convolution MM format with roofline-based optimization techniques (as described in Sections IV and V). After the transformation, we generate accelerator-customized instructions to describe the whole CNN for the FPGA accelerator.
- 3) *Optimizer (Weights Transformation)*: In this step, we prepare the CNN layers’ weights and biases and transfer them into a format which is specifically optimized for our customized accelerator, as described in Section III-A. This transformation includes static FPGA DRAM space allocation, weights and biases reorganization, and floating-point to fixed-point format transformation when the accelerator is defined as fixed-point by the user.

The above transformed layer definitions and weights and biases data will be generated for a new CNN once and written into FPGA DRAM through the PCIe interface. It will be reused for all following input images, and there will be no further weights or instructions communication. For each input image, the FPGA accelerator will start from reading the first CNN layer instructions stored in FPGA DRAM and stop to reach CPU until the last layer instructions are finished.

B. Hardware Automation

In the analysis of CNN’s computation model in Section III-A, we discussed an application-specific hardware design with a series of computation and memory optimization techniques. Our hardware automation plan is to build an easy-to-use tool with such optimizations for users to customize the hardware design for their own FPGA devices.

The key required information is the number of DSP resources, on-chip storage capacity, and external memory bandwidth provided by the platform; these are the constraints to the performance of the accelerators. The output is a set of hardware-definable parameters which have been depicted in Fig. 4. With a highly structured hardware template, we use HLS to generate the customized RTL as well as device-specific bitstream with Xilinx’s SDAccel tool. The optimized microarchitecture proposed in Section III-B ensures its scalability to larger devices to overcome the difficulties in placement and routing.

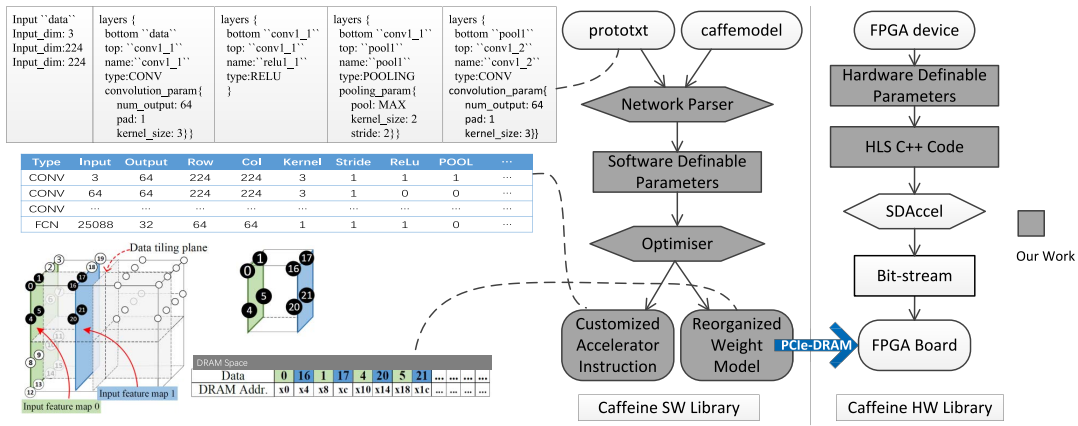


Fig. 15. Automation flow from high-level defined networks (Caffe) to hardware optimized accelerator.

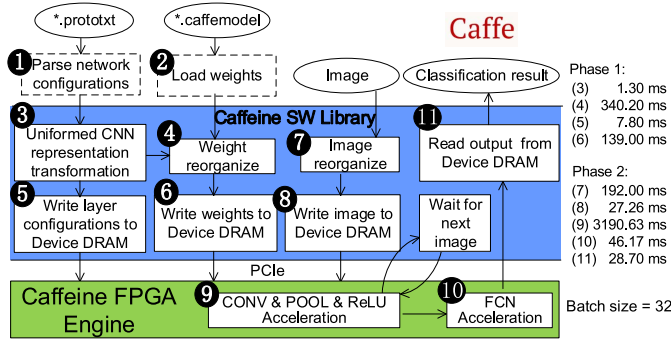


Fig. 16. Caffe-Caffeine integration.

C. Caffe-Caffeine Integration

As a case study, we integrate Caffeine with the industry-standard Caffe deep learning framework [12]. Note that Caffeine can also be integrated into other frameworks like Torch [34] and TensorFlow [35]. Fig. 16(left) presents an overview of Caffeine’s HW/SW library and its integration with Caffe. The integrated system accepts standard Caffe files for network configuration and weight values. As discussed earlier, the only part that is Caffe-specific is parsing the network configurations and loading weights (steps 1 and 2) into our Caffeine software library. Caffeine will take care of the rest.

There are two major execution phases in Caffeine. In *phase 1* (steps 3–6), it establishes the uniformed representation and automatically decides the optimal transformation, as illustrated in Section V, and then reorders weights for bandwidth optimization as illustrated in Section III-C. Finally, it initializes the FPGA device with weights and layer configurations. Phase 1 only needs to execute once unless users want to switch to a new CNN network. In *phase 2* (steps 7–11), Caffeine conducts the CNN acceleration: in batch mode, it will accumulate multiple CONV outputs and execute FCN once in a batch; in single mode, it will execute CONV and FCN once for each input image. A detailed execution time breakdown of Caffeine running the VGG16 network on a KU060 platform is shown in the right-hand part of Fig. 16 with a batch size of 32, where CONV layers dominate the entire execution again.

VII. CAFFEINE RESULTS

A. Experimental Setup

1) *CNN Models*: To demonstrate the software-definable features of Caffeine, we use two CNN models: 1) AlexNet [8]

and 2) VGG16 [11]. Users only need to write two configuration files for them.

2) *CPU and GPU Setup*: The baseline CPU we use is a two-socket server, each with a 6-core Intel CPU (E5-2609 @ 1.9 GHz). We use an NVIDIA GPU GTX1080 in our experiments. OpenBLAS and cuDNN 8.0 libraries are used for the CPU and GPU implementations [12]. In the following experiments, cuDNN is set to CUDNN_CONVOLUTION_FWD_ALGO_DIRECT mode, which the library is optimized on the original 6-loops shown in Fig. 3.

3) *FPGA Setup*: The main FPGA platform we use is the Xilinx KU3 board with a Kintex Ultrascale KU060 (20 nm) and a 8 GB DDR3 DRAM, where SDAccel 2015.3 is used to synthesize the bitstream. To demonstrate the portability of our hardware-definable architecture, we also extend our design to the VC709 (Virtex 690t, 28 nm) FPGA board. We create the IP design with Vivado HLS 2015.2 and use Vivado 2015.2 for synthesis.

B. Caffeine Results on Multiple FPGAs

To demonstrate the flexibility of Caffeine, we evaluate Caffeine using: 1) two FPGA platforms, KU060 and VC709 and 2) three data types, 32-bit floating-point, and 16-bit and 8-bit fixed-point; and 3) two network models, AlexNet and VGG16, as shown in Fig. 17.

First, Fig. 17(a) and (b) presents the VGG16 performance for 16-bit fixed-point on VC709 and KU060 platforms, respectively. VC709 can achieve higher peak performance (636 GOPS) and higher overall performance of all CONV+FCN layers (354 GOPS) than KU060’s peak 365 GOPS and overall 266 GOPS. Both figures show that most layers can achieve near-peak performance. Layer 1 is a special case because it only has three input FMs (three channels for RGB pictures). For both platforms, the FCN layer’s performance is quite similar (around 170 GOPS for overall performance of all FCN layers) because it is mainly bounded by bandwidth.

Second, Fig. 17(b)–(d) presents the differences between the 16-bit fixed-point, 8-bit fixed-point, and 32-bit floating-point on KU060. Both CONV and FCN layers show a drastic increase in performance from 32-bit floating-point to 16-bit fixed-point. For CONV layers, fixed-point saves computation resources and thus enables more parallelism. For FCN layers, fixed-point saves bandwidth because of its fewer bits. The

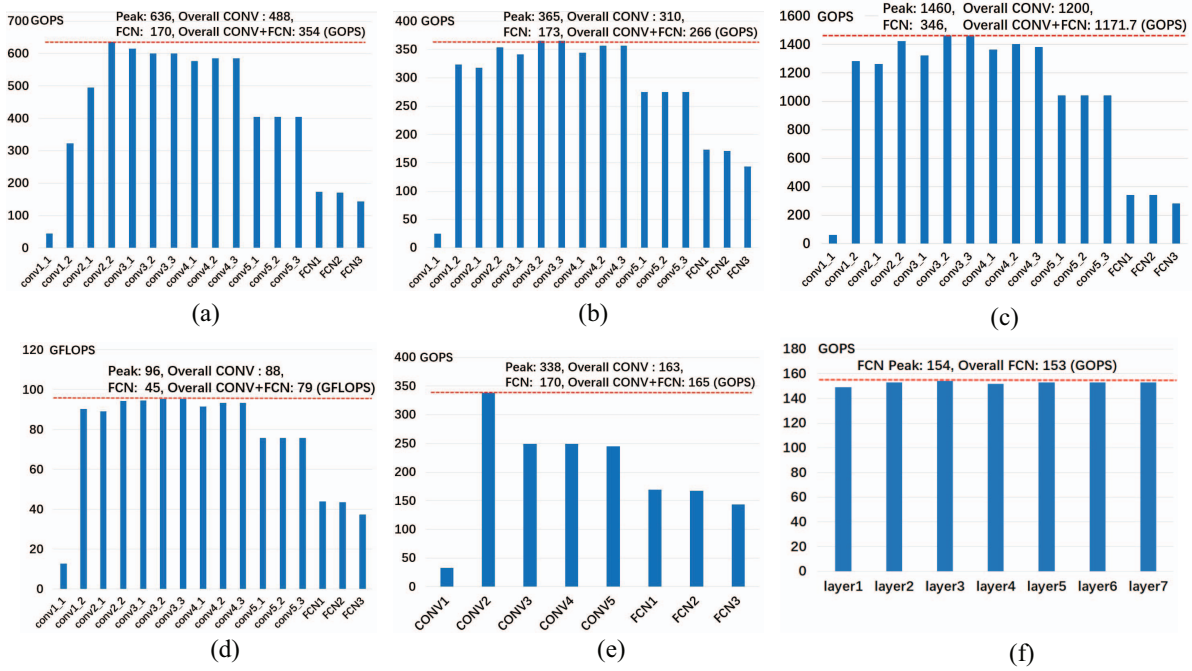


Fig. 17. Caffeine results on multiple FPGA boards for different CNN models and data types. (a) VC709 VGG 16-bit fixed-point. (b) KU VGG 16-bit fixed-point. (c) KU VGG 8-bit fixed-point. (d) KU VGG 32-bit floating-point. (e) KU AlexNet 16-bit fixed-point. (f) KU Speech FCN 16-bit fixed-point.

TABLE VII
COMPARISON WITH OTHER FPGA WORK

CNN models	Zhang[13]	Qiu[24]	Suda[23]	Ours	
	AlexNet	VGG			
Device	Virtex 480t	Zynq xc 7Z45	StratixV GSD8	Ultrascale KU060	Virtex 690t
Precision	float	fixed	fixed	fixed	fixed
	32 bit	16 bit	16 bit	16 bit	16 bit
DSP #	2240	780	1963	1058	2833
peak CONV Gops	83.8	254.8	-	365	636
overall CONV Gops	61.6	187.8	136.5	310	488
overall FCN Gops	-	1.2	-	173	170
CONV+FCN Gops	-	137	117.8	266	354

KU060 board with 8-bit operation can achieve as high as 1.46 TOPS peak performance for the CONV layer.

Third, Fig. 17(b) and (e) presents the KU060 platform’s performance on VGG16 and AlexNet. VGG16 has better performance since it has a more regular network shape which is more suitable for accelerators (better utilization after tiling).

Fourth, experimental results show that our results are quite near FPGA’s peak performance. For the KU060 FPGA case in Fig. 17(b), the theoretical peak performance with 1024 DSPs on a 16-bit fixed-point accelerator is “ $1024 \times 2 \times 0.2 \text{ GHz} = 409.6 \text{ GOPS}$,” while our attainable end-to-end test is 365 GOPS of peak performance. For KU060 FPGA with single-precision float in Fig. 17(d), theoretical peak performance is “100 GFLOPS,” while our evaluation peak performance is 96 GFLOPS.

Fifth, Fig. 17(f) shows experimental results on the fully connected network for speech [39]. With our approach, it achieves nearly 150 GOPS performance.

C. Comparison With Prior FPGA Work

We compare our accelerator design to three state-of-the-art studies in Table VII. We compare four terms of performance: 1) peak CONV layer performance; 2) overall performance of




Fig. 18. GPU versus FPGA performance.

all CONV layers; 3) overall performance of all FCN layers; and 4) overall performance of all CONV+FCN layers. This paper significantly outperforms all three prior studies in all terms of performance. Our FCN layer achieves more than 100× speed-up over previous work. In addition, very-low bit (binarized) network technique [27] is orthogonal to this paper.

D. End-to-End Comparison With CPUs and GPUs

We conduct an end-to-end comparison between Caffe–Caffeine integration with existing optimized CPU and GPU solutions [12] for VGG16 in Table VIII. For fair comparison, we use GOPS as the standard metric. With on-board (KU060) testing, our integration using 8-bit fixed-point operations demonstrates an end-to-end performance of 29× speed-up and 150× energy efficiency over 12-core CPU, and 5.7× and 2× energy efficiency over batch = 1 and batch = 16 cuDNN implementations respectively. Fig. 18 shows detailed layer-wise performance comparison between 8-bit fixed Ku060 FPGA implementation and GTX1080 GPU cuDNN. Our FPGA implementation has approximately similar performance to GPU when batch = 1. But batch = 16

TABLE VIII
END-TO-END COMPARISON WITH CPU/GPU PLATFORMS

Platforms	CPU		GPU		CPU+FPGA	
Device	E5-2609		GTX1080		KU060	VX 690t
Precision	float		float		fix16	fix8
Technology	22nm		16nm		20nm	20nm
Freq.(GHz)	1.9		2.1		0.2	0.15
Power(Watt)	150		180		25	26
Batch Size	1		16		1	1
Latency/img.(ms)	733.7		8.13		101.15	25.3
Speedup	1x		90x		7.3x	29x
J per image	110		1.46		4.23	2.5
Energy Efficiency	1x		75x		26x	43.5x

TABLE IX
FPGA RESOURCE UTILIZATION OF CAFFEINE

	DSP	BRAM	LUT	FF	Freq
VC fix-16	2833(78%)	1248(42%)	3E5(81%)	3E5(36%)	150MHz
KU fix-16	1058 (38%)	782(36%)	1E5(31%)	8E4(11%)	200MHz
KU fix-8	116(4%)	784(36%)	2E5(60%)	1.4E5(20%)	200MHz
KU float	1314(47%)	798(36%)	1.5E5(46%)	2E5(26%)	200MHz




Fig. 19. GPU implementation of input- and weight-major mappings. (a) Kernel size = 1*1. (b) Kernel size = 4*4.

GPU implementation has much higher performance (lower energy efficiency).

Finally, Table IX presents the FPGA resource utilization of the above implementations. SDAccel uses a partial reconfiguration to write bit-stream, and thus it has an up-limit of 60% of all available resources. We use about 50% of DSP resources on the KU060 board. We use 80% of DSP resources on the VC709 board. Note that for the 8-bit fixed-point implementation, it is more resource efficient and mainly uses the LUT resources. Caffeine on the KU060 board runs at a frequency of 200 MHz, and on VC709 it runs at a frequency of 150 MHz.

E. Input-Major/Weight-Major Mapping on GPUs

We further verify our idea on GPU implementations. We took one optimized implementation on the original 6-loops as shown in Fig. 3 from cuDNN library. We transform the VGG-16 FCN-2 layer to a convolutional layer using both input-major and weight-major mappings. Fig. 19 shows that for most of the cases under 1×1 and 4×4 kernel sizes, weight-major mapping outperforms input-major mapping.

F. Comparison With TPUs

Google's Tensor Processing Unit [44] cites work [13] and argues that their systolic micro-architecture design is more friendly for frequency tuning. In this paper, we also improve and use systolic design. However, TPU's performance on MLP for speech are greatly degraded because of strict bandwidth constraints. Our proposal of input-major/weight-major mapping in this paper can be helpful for TPU to optimize

the computation and communication ratio and thus improve overall performance.

VIII. CONCLUSION

In this paper, we proposed a uniformed convolutional MM representation to accelerate both the computation-bound convolutional layers and communication-bound fully connected layers of CNN/DNN on FPGAs. Based on the uniformed representation, we designed and implemented Caffeine, a HW/SW co-designed reusable library to efficiently accelerate the entire CNN/DNN on FPGAs. Finally, we also provide an automation flow to integrate Caffeine into the industry-standard software deep learning framework Caffe. We evaluated Caffeine and its integration with Caffe using both AlexNet and VGG networks on multiple FPGA platforms. Caffeine achieved up to 1460 GOPS on a KU060 board with 8-bit fixed-point operations, and more than $100 \times$ speed-up on fully connected layers over prior FPGA accelerators. Our Caffe integration achieved $29 \times$ and $150 \times$ performance and energy gains over a 12-core CPU, and $5.7 \times$ better energy efficiency over GPU on a medium-sized KU060 FPGA board.

REFERENCES

- [1] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf, "DeepFace: Closing the gap to human-level performance in face verification," in *Proc. CVPR*, 2014, pp. 1701–1708.
- [2] K. He, X. Zhang, S. Ren, and J. Sun, "Delving deep into rectifiers: Surpassing human-level performance on ImageNet classification," in *Proc. ICCV*, 2015, pp. 1026–1034.
- [3] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *Proc. CVPR*, 2014, pp. 580–587.
- [4] S. Ji, W. Xu, M. Yang, and K. Yu, "3D convolutional neural networks for human action recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 1, pp. 221–231, Jan. 2013.
- [5] A. Coates *et al.*, "Deep learning with COTS HPC systems," in *Proc. ICML*, 2013, pp. III-1337–III-1345.
- [6] Z. Zheng, W. Jiang, G. Wu, and E. Y. Chang, "SpeeDO: Parallelizing stochastic gradient descent for deep convolutional neural network," in *Proc. LearningSys*, 2015, pp. 1–6.
- [7] K. Yu, "Large-scale deep learning at Baidu," in *Proc. ACM CIKM*, 2013, pp. 2211–2212.
- [8] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Proc. NIPS*, 2012, pp. 1097–1105.
- [9] M. D. Zeiler and R. Fergus, "Visualizing and understanding convolutional networks," in *Proc. ECCV*, 2014, pp. 818–833.
- [10] C. Szegedy *et al.*, "Going deeper with convolutions," in *Proc. CVPR*, 2015, pp. 1–9.
- [11] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *Proc. ICLR*, 2015.
- [12] Y. Q. C. Jia. (2013). *An Open Source Convolutional Architecture for Fast Feature Embedding*. [Online]. Available: <http://caffe.berkeleyvision.org>
- [13] C. Zhang *et al.*, "Optimizing FPGA-based accelerator design for deep convolutional neural networks," in *Proc. ACM FPGA*, 2015, pp. 161–170.
- [14] T. Chen *et al.*, "DianNao: A small-footprint high-throughput accelerator for ubiquitous machine-learning," *ACM SIGPLAN Notices*, vol. 49, no. 4, pp. 269–284, 2014.
- [15] C. Zhang, Z. Fang, P. Zhou, P. Pan, and J. Cong, "Caffeine: Towards uniformed representation and acceleration for deep convolutional neural networks," in *Proc. IEEE/ACM Int. Conf. Comput.-Aided Design (ICCAD)*, 2016, pp. 1–8.
- [16] C. Farabet, C. Poulet, J. Y. Han, and Y. LeCun, "CNP: An FPGA-based processor for convolutional networks," in *Proc. IEEE FPL*, 2009, pp. 32–37.
- [17] S. Chakradhar *et al.*, "A dynamically configurable coprocessor for convolutional neural networks," *ACM SIGARCH Comput. Archit. News*, vol. 38, no. 3, pp. 247–257, 2010.


- [18] D. Aysegul *et al.*, "Accelerating deep neural networks on mobile processor with embedded programmable logic," in *Proc. IEEE NIPS*, 2013.
- [19] S. Cadambi, A. Majumdar, M. Becchi, S. Chakradhar, and H. P. Graf, "A programmable parallel accelerator for learning and classification," in *Proc. ACM PACT*, 2010, pp. 273–284.
- [20] M. Sankaradas *et al.*, "A massively parallel coprocessor for convolutional neural networks," in *Proc. IEEE ASAP*, 2009, pp. 53–60.
- [21] M. Peemen, A. A. A. Setio, B. Mesman, and H. Corporaal, "Memory-centric accelerator design for convolutional neural networks," in *Proc. IEEE ICCD*, 2013, pp. 13–19.
- [22] K. Ovtcharov *et al.*, "Accelerating deep convolutional neural networks using specialized hardware," *Microsoft Res. Whitepaper*, vol. 2, no. 11, 2015.
- [23] N. Suda *et al.*, "Throughput-optimized openCL-based FPGA accelerator for large-scale convolutional neural networks," in *Proc. ACM FPGA*, 2016, pp. 16–25.
- [24] J. Qiu *et al.*, "Going deeper with embedded FPGA platform for convolutional neural network," in *Proc. ACM FPGA*, 2016, pp. 26–35.
- [25] X. Wei *et al.*, "Automated systolic array architecture synthesis for high throughput CNN inference on FPGAs," in *Proc. ACM DAC*, 2017, pp. 1–6.
- [26] Y. Ma, Y. Cao, S. Vrudhula, and J.-S. Seo, "Optimizing loop operation and dataflow in FPGA acceleration of deep convolutional neural networks," in *Proc. ACM/SIGDA Int. Symp. Field Program. Gate Arrays*, 2017, pp. 45–54.
- [27] R. Zhao *et al.*, "Accelerating binarized convolutional neural networks with software-programmable FPGAs," in *Proc. ACM FPGA*, 2017, pp. 15–24.
- [28] J. Zhang and J. Li, "Improving the performance of openCL-based FPGA accelerator for convolutional neural network," in *Proc. FPGA*, 2017, pp. 25–34.
- [29] C. Zhang and V. Prasanna, "Frequency domain acceleration of convolutional neural networks on CPU-FPGA shared memory system," in *Proc. ACM/SIGDA Int. Symp. Field Program. Gate Arrays*, 2017, pp. 35–44.
- [30] Y.-K. Choi *et al.*, "A quantitative analysis on microarchitectures of modern CPU-FPGA platforms," in *Proc. DAC*, 2016, pp. 1–6.
- [31] J. Bergstra *et al.*, "Theano: A CPU and GPU math expression compiler," in *Proc. SciPy*, vol. 4, 2010, p. 3.
- [32] Vivado Design Suite, "Ultrascale architecture FPGAs memory interface solutions v7.0," Xilinx, San Jose, CA, USA, Rep., Apr. 2015.
- [33] S. Mittal, "A survey of techniques for managing and leveraging caches in GPUs," *J. Circuits Syst. Comput.*, vol. 23, no. 8, 2014, Art. no. 1430002.
- [34] *Torch7*. [Online]. Available: <http://torch.ch>
- [35] M. Abadi *et al.* (2016). *TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems*. [Online]. Available: <http://www.tensorflow.org>
- [36] K. Rupnow *et al.*, "High level synthesis of stereo matching: Productivity, performance, and software constraints," in *Proc. IEEE Int. Conf. Field Program. Technol. (FPT)*, 2011, pp. 1–8.
- [37] W. Zuo *et al.*, "Improving high level synthesis optimization opportunity through polyhedral transformations," in *Proc. ACM FPGA*, 2013, pp. 9–18.
- [38] S. Williams, A. Waterman, and D. Patterson, "Roofline: An insightful visual performance model for multicore architectures," *Commun. ACM*, vol. 52, no. 4, pp. 65–76, 2009.
- [39] Z.-J. Yan, Q. Huo, and J. Xu, "A scalable approach to using DNN-derived features in GMM-HMM based acoustic modeling for LVCSR," in *Proc. Interspeech*, 2013, pp. 104–108.
- [40] F. Grézl, M. Karafiát, S. Kontár, and J. Cernocky, "Probabilistic and bottle-neck features for LVCSR of meetings," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process. (ICASSP)*, vol. 4, 2007, pp. IV-757–IV-760.
- [41] J. Gehring, Y. Miao, F. Metzger, and A. Waibel, "Extracting deep bottleneck features using stacked auto-encoders," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process. (ICASSP)*, 2013, pp. 3377–3381.
- [42] D. Yu and M. L. Seltzer, "Improved bottleneck features using pretrained deep neural networks," in *Proc. 12th Annu. Conf. Int. Speech Commun. Assoc.*, 2011, pp. 237–240.
- [43] G. Hinton *et al.*, "Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups," *IEEE Signal Process. Mag.*, vol. 29, no. 6, pp. 82–97, Nov. 2012.
- [44] N. P. Jouppi *et al.*, "In-datacenter performance analysis of a tensor processing unit," in *Proc. ACM/IEEE 44th Annu. Int. Symp. Comput. Archit. (ISCA)*, 2017, pp. 1–12.



Dr. Zhang is a member of ACM.



Dr. Sun is currently serving as an Associate Editor of ACM JETC and TECS. He is a member of ACM and CCF.



Chen Zhang (M'03) received the B.S. degree in electronic engineering from the University of Electronic Science and Technology of China, Chengdu, China, in 2012 and the Ph.D. degree from Computer Science Department, Peking University, Beijing, China, in 2017.

He is currently an Associate Researcher (II) with Microsoft Research Asia, Beijing. His current research interests include high performance and energy-efficient computer architectures and systems in deep learning.


Guangyu Sun (M'07) received the B.S. and M.S. degrees from Tsinghua University, Beijing, China, in 2003 and 2006, respectively, and the Ph.D. degree in computer science from Pennsylvania State University, State College, PA, USA, in 2011.

He is an Associate Professor with the Center for Energy-Efficient Computing and Applications, Peking University, Beijing. His current research interests include computer architecture, electronic design automation, and acceleration system for modern applications.

Zhenman Fang (M'03) received the Ph.D. degree in computer science from Fudan University, Shanghai, China.

He recently joined Xilinx, San Jose, CA, USA, after a three-year Post-Doctoral Fellow with the University of California at Los Angeles, Los Angeles, CA. His current research interests include intersection of heterogeneous and energy-efficient computer architectures, big data workloads and systems, and system-level design automation.

Dr. Fang is a member of the ACM.



performance and energy

Peipei Zhou received the B.S. degree in electrical engineering from Chien-Shiung Wu Honor College Southeast University, Nanjing, China, in 2012, and the M.S. degree in electrical engineering from the University of California at Los Angeles, Los Angeles, CA, USA, in 2014, where she is currently pursuing the Ph.D. degree with Computer Science Department, under supervision of Prof. J. Cong.

Her current research interests include parallel/distributed architecture and programming model for computer architecture design.



Ph.D. Thesis Award from UIUC in 1996.

Peichen Pan received the Ph.D. degree in computer science from the University of Illinois at Urbana–Champaign (UIUC), Champaign, IL, USA, in 1995.

He is the Vice President of engineering with Falcon Computing Solutions Inc., Los Angeles, CA, USA. His current research interests include system-level and high-level synthesis, and FPGA acceleration of big-data applications, such as machine learning and genomic data processing.

Dr. Pan received the David J. Kuck Outstanding



Jason Cong (F'00) received the B.S. degree in computer science from Peking University, Beijing, China, in 1985, and the M.S. and Ph.D. degrees in computer science from the University of Illinois at Urbana–Champaign, Champaign, IL, USA, in 1987 and 1990, respectively.

He is currently a Chancellors Professor with the Computer Science Department and the Electrical Engineering Department, University of California at Los Angeles, Los Angeles, CA, USA.

Dr. Cong was elected as an ACM Fellow in 2008 and the National Academy of Engineering in 2017.