

DLAU: A Scalable Deep Learning Accelerator Unit on FPGA

Chao Wang, *Member, IEEE*, Lei Gong, Qi Yu, Xi Li, *Member, IEEE* Yuan Xie, *Fellow, IEEE*
and Xuehai Zhou, *Member, IEEE*

Abstract—As the emerging field of machine learning, deep learning shows excellent ability in solving complex learning problems. However, the size of the networks becomes increasingly large scale due to the demands of the practical applications, which poses significant challenge to construct a high performance implementations of deep learning neural networks. In order to improve the performance as well to maintain the low power cost, in this paper we design DLAU, which is a scalable accelerator architecture for large-scale deep learning networks using FPGA as the hardware prototype. The DLAU accelerator employs three pipelined processing units to improve the throughput and utilizes tile techniques to explore locality for deep learning applications. Experimental results on the state-of-the-art Xilinx FPGA board demonstrate that the DLAU accelerator is able to achieve up to 36.1x speedup comparing to the Intel Core2 processors, with the power consumption at 234mW.

Index Terms—FPGA; Deep Learning; neural network; hardware accelerator.

I. INTRODUCTION

IN the past few years, machine learning has become pervasive in various research fields and commercial applications, and achieved satisfactory products. The emergence of deep learning speeded up the development of machine learning and artificial intelligence. Consequently, deep learning has become a research hot spot in research organizations [1]. In general, deep learning uses a multi-layer neural network model to extract high-level features which are a combination of low-level abstractions to find the distributed data features, in order to solve complex problems in machine learning. Currently the most widely used neural models of deep learning are Deep Neural Networks (DNNs) [2] and Convolution Neural Networks (CNNs) [3], which have been proved to have excellent capability in solving picture recognition, voice recognition and other complex machine learning tasks.

However, with the increasing accuracy requirements and complexity for the practical applications, the size of the neural networks becomes explosively large scale, such as the Baidu Brain with 100 Billion neuronal connections, and the Google cat-recognizing system with 1 Billion neuronal connections. The explosive volume of data makes the data centers quite

C. Wang, L.Gong, Q. Yu, X.Li and X.Zhou are with University of Science and Technology of China, Hefei, 230027, Anhui, China. E-mail: {cswang,llxx,xhzhou}@ustc.edu.cn, {SA514002,yuiq}@mail.ustc.edu.cn).

Y. Xie is with University of California at Santa Barbara, 93106, United States, E-mail:yuanxie@ece.ucsb.edu.

Manuscript received January 10, 2016.

Copyright (c) 2015 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending an email to pubs-permissions@ieee.org.

power consuming. In particular, the electricity consumption of data centers in U.S. are projected to increase to roughly 140 billion kilowatt-hours annually by 2020 [4]. Therefore, it poses significant challenges to implement high performance deep learning networks with low power cost, especially for large-scale deep learning neural network models. So far, the state-of-the-art means for accelerating deep learning algorithms are Field-Programmable Gate Array (FPGA), Application Specific Integrated Circuit (ASIC), and Graphic Processing Unit (GPU). Compared with GPU acceleration, hardware accelerators like FPGA and ASIC can achieve at least moderate performance with lower power consumption. However, both FPGA and ASIC have relatively limited computing resources, memory, and I/O bandwidths, therefore it is challenging to develop complex and massive deep neural networks using hardware accelerators. For ASIC, it has a longer development cycle and the flexibility is not satisfying. Chen et al presents a ubiquitous machine-learning hardware accelerator called DianNao [6], which initiated the field of deep learning processor. It opens a new paradigm to machine learning hardware accelerators focusing on neural networks. But DianNao is not implemented using reconfigurable hardware like FPGA, therefore it cannot adapt to different application demands. Currently around FPGA acceleration researches, Ly and Chow [5] designed FPGA based solutions to accelerate the Restricted Boltzmann Machine (RBM). They created dedicated hardware processing cores which are optimized for the RBM algorithm. Similarly Kim et al [7] also developed a FPGA based accelerator for the restricted Boltzmann machine. They use multiple RBM processing modules in parallel, with each module responsible for a relatively small number of nodes. Other similar works also present FPGA based neural network accelerators [9]. Qi et al. present a FPGA based accelerator [8], but it cannot accommodate changing network size and network topologies. To sum up, these studies focus on implementing a particular deep learning algorithm efficiently, but how to increase the size of the neural networks with scalable and flexible hardware architecture has not been properly solved.

To tackle these problems, we present a scalable deep learning accelerator unit named DLAU to speed up the kernel computational parts of deep learning algorithms. In particular, we utilize the tile techniques, FIFO buffers, and pipelines to minimize memory transfer operations, and reuse the computing units to implement the large-size neural networks. This approach distinguishes itself from previous literatures with following contributions:

1. In order to explore the locality of the deep learning

TABLE I
PROFILING OF HOT SPOTS OF DNN

Algorithms	Matrix Multiplication	Activation	Vector
Feedforward	98.60%	1.40%	
RBM	98.20%	1.48%	0.30%
BP	99.10%	0.42%	0.48%

application, we employ tile techniques to partition the large scale input data. The DLAU architecture can be configured to operate different sizes of tile data to leverage the trade-offs between speedup and hardware costs. Consequently the FPGA based accelerator is more scalable to accommodate different machine learning applications.

2. The DLAU accelerator is composed of three fully pipelined processing units, including TMMU, PSAU, and AFAU. Different network topologies such as CNN, DNN, or even emerging neural networks can be composed from these basic modules. Consequently the scalability of FPGA based accelerator is higher than ASIC based accelerator.

II. TILE TECHNIQUES AND HOT SPOT PROFILING

Restricted Boltzmann Machines (RBMs) have been widely used to efficiently train each layer of a deep network. Normally a deep neural network is composed of one input layer, several hidden layers and one classifier layer. The units in adjacent layers are all-to-all weighted connected. The prediction process contains feedforward computation from given input neurons to the output neurons with the current network configurations. Training process includes pre-training which locally tune the connection weights between the units in adjacent layers, and global training which globally tune the connection weights with Back Propagation process.

The large-scale deep neural networks include iterative computations which have few conditional branch operations, therefore they are suitable for parallel optimization in hardware. In this paper we first explore the hot spot using the profiler. Results in Fig. 1 illustrates the percentage of running time including Matrix Multiplication (MM), Activation, and Vector operations. For the representative three key operations: feed forward, Restricted Boltzmann Machine (RBM), and back propagation (BP), matrix multiplication play a significant role of the overall execution. In particular, it takes 98.6%, 98.2%, and 99.1% of the feed forward, RBM, and BP operations. In comparison, the activation function only takes 1.40%, 1.48%, and 0.42% of the three operations. Experimental results on profiling demonstrate that the design and implementation of MM accelerators is able to improve the overall speedup of the system significantly.

However, considerable memory bandwidth and computing resources are needed to support the parallel processing, consequently it poses a significant challenge to FPGA implementations compared with GPU and CPU optimization measures. In order to tackle the problem, in this paper we employ tile techniques to partition the massive input data set into tiled subsets. Each designed hardware accelerator is able to buffer the tiled subset of data for processing. In order to support the

large-scale neural networks, the accelerator architecture are reused. Moreover, the data access for each tiled subset can run in parallel to the computation of the hardware accelerators.

Algorithm 1 Pseudocode Code of the Tiled Inputs

Require:

N_i : the number of the input neurons
 N_o : the number of the output neurons
 Tile_Size: the tile size of the input data
 batchsize: the batch size of the input data
for $n = 0; n < batchsize; n++$ **do**
 for $k = 0; k < N_i; k+ = Tile_Size$ **do**
 for $j = 0; j < N_o; j++$ **do**
 $y[n][j] = 0;$
 for $i = k; i < k + Tile_Size \&\& i < N_i; i++$ **do**
 $y[n][j] += w[i][j] * x[n][i]$
 if $i == N_i - 1$ **then**
 $y[n][j] = f(y[n][j]);$
 end if
 end for
 end for
 end for
end for

In particular, for each iteration, output neurons are reused as the input neurons in next iteration. To generate the output neurons for each iteration, we need to multiply the input neurons by each column in weights matrix. As illustrated in Algorithm 1, the input data are partitioned into tiles and then multiplied by the corresponding weights. Thereafter the calculated part sum are accumulated to get the result. Besides the input/output neurons, we also divided the weight matrix into tiles corresponding to the tile size. As a consequence, the hardware cost of the accelerator only depends on the tile size, which saves significant number of hardware resources. The tiled technique is able to solve the problem by implementing large networks with limited hardware. Moreover, the pipelined hardware implementation is another advantage of FPGA technology compared to GPU architecture, which uses massive parallel SIMD architectures to improve the overall performance and throughput. According to the profiling results depicted in Table I, during the prediction process and the training process in deep learning algorithms, the common but important computational parts are matrix multiplication and activation functions, consequently in this paper we implement the specialized accelerator to speed up the matrix multiplication and activation functions.

III. DLAU ARCHITECTURE AND EXECUTION MODEL

Fig. 1 describes the DLAU system architecture which contains an embedded processor, a DDR3 memory controller, a DMA module, and the DLAU accelerator. The embedded processor is responsible for providing programming interface to the users and communicating with DLAU via JTAG-UART. In particular it transfers the input data and the weight matrix to internal BRAM blocks, activates the DLAU accelerator, and returns the results to the user after execution. The DLAU is integrated as a standalone unit which is flexible and adaptive

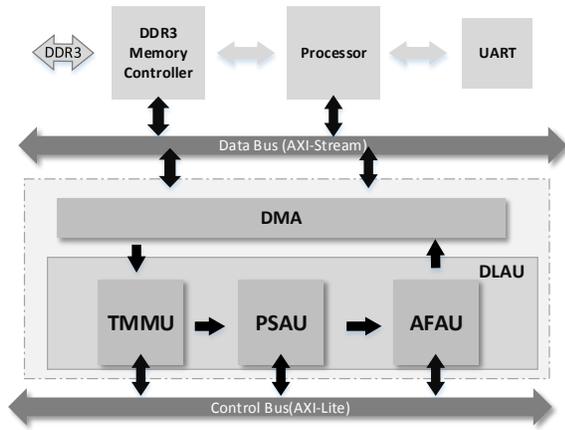


Fig. 1. DLAU Accelerator Architecture.

to accommodate different applications with configurations. The DLAU consists of 3 processing units organized in a pipeline manner: Tiled Matrix Multiplication Unit (TMMU), Part Sum Accumulation Unit (PSAU), and Activation Function Acceleration Unit (AFAU). For execution, DLAU reads the tiled data from the memory by DMA, computes with all the three processing units in turn, and then writes the results back to the memory.

In particular, the DLAU accelerator architecture has following key features:

FIFO Buffer: Each processing unit in DLAU has an input buffer and an output buffer to receive or send the data in FIFO. These buffers are employed to prevent the data loss caused by the inconsistent throughput between each processing unit.

Tiled Techniques: Different machine learning applications may require specific neural network sizes. The tile technique is employed to divide the large volume of data into small tiles that can be cached on chip, therefore the accelerator can be adopted to different neural network size. Consequently the FPGA based accelerator is more scalable to accommodate different machine learning applications.

Pipeline Accelerator: We use stream-like data passing mechanism (e.g. AXI-Stream for demonstration) to transfer data between the adjacent processing units, therefore TMMU, PSAU, and AFAU can compute in streaming-like manner. Of these three computational modules, TMMU is the primary computational unit, which reads the total weights and tiled nodes data through DMA, performs the calculations, and then transfers the intermediate Part Sum results to PSAU. PSAU collects Part Sums and performs accumulation. When the accumulation is completed, results will be passed to AFAU. AFAU performs the activation function using piecewise linear interpolation methods. In the rest of this section, we will detail the implementation of these three processing units respectively.

A. TMMU architecture

Tiled Matrix Multiplication Unit (TMMU) is in charge of multiplication and accumulation operations. TMMU is specially designed to exploit the data locality of the weights and is responsible for calculating the Part Sums. TMMU employs

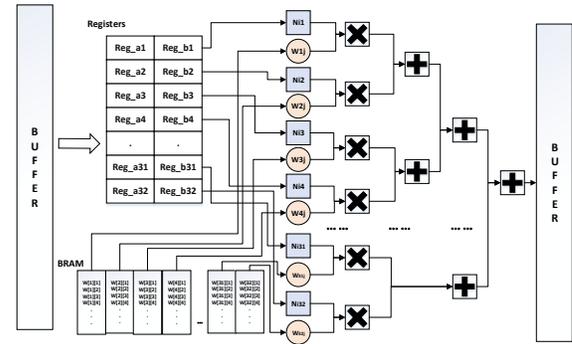


Fig. 2. TMMU Schematic Diagram.

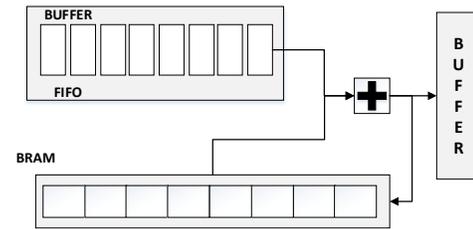


Fig. 3. PSAU Schematic Diagram

an input FIFO buffer which receives the data transferred from DMA and an output FIFO buffer to send Part Sums to PSAU. Fig. 2 illustrates the TMMU schematic diagram, in which we set tile size=32 as an example. TMMU firstly reads the weight matrix data from input buffer into different BRAMs in 32 by the row number of the weight matrix ($n=i\%32$ where n refers to the number of BRAM, and i is the row number of weight matrix). Then, TMMU begins to buffer the tiled node data. In the first time, TMMU reads the tiled 32 values to registers Reg_a and starts execution. In parallel to the computation at every cycle, TMMU reads the next node from input buffer and saves to the registers Reg_b. Consequently the registers Reg_a and Reg_b can be used alternately.

For the calculation, we use pipelined binary adder tree structure to optimize the performance. As depicted in Fig. 2, the weight data and the node data are saved in BRAMs and registers. The pipeline takes advantage of time-sharing the coarse-grained accelerators. As a consequence, this implementation enables the TMMU unit to produce a Part Sum result every clock cycle.

B. PSAU architecture

Part Sum Accumulation Unit (PSAU) is responsible for the accumulation operation. Fig. 3 presents the PSAU architecture, which accumulates the part sum produced by TMMU. If the Part Sum is the final result, PSAU will write the value to output buffer and send results to AFAU in a pipeline manner. PSAU can accumulate one Part Sum every clock cycle, therefore the throughput of PSAU accumulation matches the generation of the Part Sum in TMMU.

C. AFAU architecture

Finally, Activation Function Acceleration Unit (AFAU) implements the activation function using piecewise linear interpolation ($y=ai*x+bi, x \in [x_i, x_{i+1})$). This method has been widely applied to implement activation functions with negligible accuracy loss when the interval between x_i and x_{i+1} is insignificant. Eq. (1) shows the implementation of sigmoid function. For $x > 8$ and $x \leq -8$, the results are sufficiently close to the bounds of 1 and 0, respectively. For the cases in $-8 < x \leq 0$ and $0 < x \leq 8$, different functions are configured. In total we divide the sigmoid function into four segments.

$$f(x) = \begin{cases} 0 & \text{if } x \leq -8 \\ 1 + a[\lfloor \frac{-x}{k} \rfloor]x - b[\lfloor \frac{-x}{k} \rfloor] & \text{if } -8 < x \leq 0 \\ a[\lfloor \frac{x}{k} \rfloor]x + b[\lfloor \frac{x}{k} \rfloor] & \text{if } 0 < x \leq 8 \\ 1 & \text{if } x > 8 \end{cases} \quad (1)$$

Similar to PSAU, AFAU also has both input buffer and output buffer to maintain the throughput with other processing units. In particular, we use two separate BRAMs to store the values of a and b . The computation of AFAU is pipelined to operate sigmoid function every clock cycle. As a consequence, all the three processing units are fully pipelined to ensure the peak throughput of the DLAU accelerator architecture.

IV. EXPERIMENTS AND DATA ANALYSIS

In order to evaluate the performance and cost of the DLAU accelerator, we have implemented the hardware prototype on the Xilinx Zynq Zedboard development board, which equips ARM Cortex-A9 processors clocked at 667MHz and programmable fabrics. For benchmarks, we use the Mnist data set to train the $784 \times M \times N \times 10$ Deep Neural Networks in Matlab, and use $M \times N$ layers weights and nodes value for the input data of DLAU. For comparison, we use Intel Core2 processor clocked at 2.3GHz as the baseline.

In the experiment we use Tile size=32 considering the hardware resources integrated in the Zedboard development board. The DLAU computes 32 hardware neurons with 32 weights every cycle. The clock of DLAU is 200MHz (one cycle takes 5ns). Three network sizes— 64×64 , 128×128 , and 256×256 are tested.

A. Speedup Analysis

We present the speedup of DLAU and some other similar implementations of the deep learning algorithms in Table II. Experimental results demonstrate that the DLAU is able to achieve up to 36.1x speedup at 256×256 network size. In comparison, Ly&Chows work [5] and Kim et.als work [7] present the work only on Restricted Boltzmann Machine algorithms, while the DLAU is much more scalable and flexible. DianNao [6] reaches up to 117.87x speedup due to its high working frequency at 0.98GHz. Moreover, as DianNao is hardwired instead of implemented on a FPGA platform, therefore it cannot efficiently adapt to different neural network sizes.

Fig. 4 illustrates the speedup of DLAU at different network sizes— 64×64 , 128×128 , and 256×256 respectively. Experimental results demonstrate a reasonable ascendant speedup

TABLE II
COMPARISONS BETWEEN SIMILAR APPROACHES

Work	Network	Clock	Speedup	Baseline
Ly&Chow [5]	256×256	100MHz	$32 \times$	2.8GHz P4
Kim et.al [7]	256×256	200MHz	$25 \times$	2.4GHz Core2
DianNao [6]	General	0.98GHz	$117.87 \times$	2GHz SIMD
Zhang et.al [3]	256×256	100MHz	$17.42 \times$	2.2GHz Xeon
DLAU	256×256	200MHz	$36.1 \times$	2.3GHz Core2

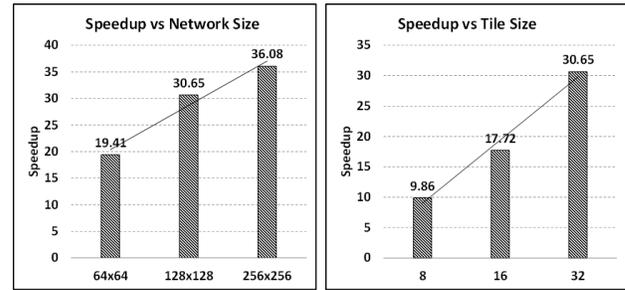


Fig. 4. Speedup at Different Network Sizes and Tile Sizes.

TABLE III
RESOURCE UTILIZATION OF DLAU AT 32×32 TILE SIZE

Component	BRAMs	DSPs	FFs	LUTs
TMMU	32	158	25356	32461
PSAU	1	2	754	632
AFAU	2	7	2216	3291
Total	35	167	28326	36384
Available	280	220	106400	53200
Utilization	12.5%	75.9%	26.6%	68.4%

with the growth of neural networks sizes. In particular, the speedup increases from 19.2x in 64×64 network size to 36.1x at the 256×256 network size. The right part of Fig. 4 illustrates how the tile size has an impact on the performance of the DLAU. It can be acknowledged that bigger tile size means more number of neurons to be computed concurrently. At the network size of 128×128 , the speedup is 9.2x when the tile size is 8. When the tile size increases to 32, the speedup reaches 30.5x. Experimental results demonstrate that the DLAU framework is configurable and scalable with different tile sizes. The speedup can be leveraged with hardware cost to achieve satisfying trade-offs.

B. Resource utilization and Power

Table III summarizes the resource utilization of DLAU in 32×32 tile size including the BRAM resources, DSPs, FFs, and LUTs. TMMU is much more complex than the rest two hardware modules therefore it consumes most hardware resources. Taking the limited number of hardware logic resources provided by Xilinx XC7Z020 FPGA chip, the overall utilization is reasonable. The DLAU utilizes 167 DSP blocks due to the use of the Floating-point addition and the Floating-point multiplication operations.

Table IV compares the resource utilization of DLAU with other two FPGA based literatures. Experimental results depict

TABLE IV
RESOURCE COMPARISONS BETWEEN SIMILAR APPROACHES

Implementation	FPGA	BRAMs	DSPs	FFs	LUTs
Ly&Chow [5]	XC2VP70	257	N/A	30403	29885
Kim et.al [7]	N/A	589824	18	11790	7662
DLAU	XC7Z020	35	167	28326	36384

TABLE V
POWER CONSUMPTION OF THE UNITS

Component	Power	Component	Power
Accelerator-TMMU	189mW	Processor	1307mW
Accelerator-PSAU	5mW	DDR Controller	177mW
Accelerator-AFAU	25mW	Peripherals	26mW
Accelerator-DMA	15mW	Clocks	70mW
Accelerator-Total	234mW	System Total	1814mW

that our DLAU accelerator occupies similar number of FFs and LUTs to Ly&Chow’s work [5], while it only consumes 35/257=13.6% on the BRAMs. Comparing to the Kim et.al’s work [7], the BRAM utilization of DLAU is insignificant. This is due to the tile techniques so that large scale neural networks can be divided into small tiles, therefore the scalability and flexibility of the architecture is significantly improved.

In order to evaluate the power consumption of accelerator, we use Xilinx Vivado tool set to achieve power cost of each processing unit in DLAU and the DMA module. The results in Table IV-B depict that the total power of DLAU is only 234mW, which is much lower than that of DianNao (485mW). The results demonstrate that the DLAU is quite energy efficient as well as highly scalable compared to other accelerating techniques. To compare the energy and power between FPGA based accelerator and GPU based accelerators, we also implement a prototype using the state-of-the-art NVIDIA Tesla K40c as the baseline. K40c has 2880 stream cores working at peak frequency 875MHz, and the Max Memory Bandwidth is 288 (GB/sec). In comparison, we only employ 1 DLAU on the FPGA board working at 100MHz. In order to evaluate the speedup of the accelerators in a real deep learning applications, we use DNN to model 3 benchmarks, including Caltech101, Cifar-10, and MNIST, respectively. Fig. 5 illustrates the comparison between FPGA based GPU+cuBLAS implementations. It reveals that the power consumption of GPU based accelerator is 364 times higher than FPGA based accelerators. Regarding the total energy consumption, the FPGA based accelerator is 10x more energy efficient than GPU, and 4.2x than GPU+cuBLAS optimizations.

Finally Fig. 6 illustrates the floor plan of the FPGA chip. The left corner depicts the ARM processor which is hard-wired in the FPGA chip. Other modules, including different components of the DLAU accelerator, the DMA, and memory interconnect, are presented in different colors. Regarding the programming logic devices, TMMU takes most of the areas as it utilizes a significant number of LUTs and FFs.

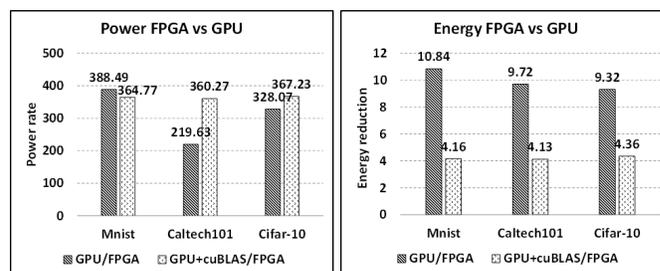


Fig. 5. Power and Energy Comparison between FPGA and GPU

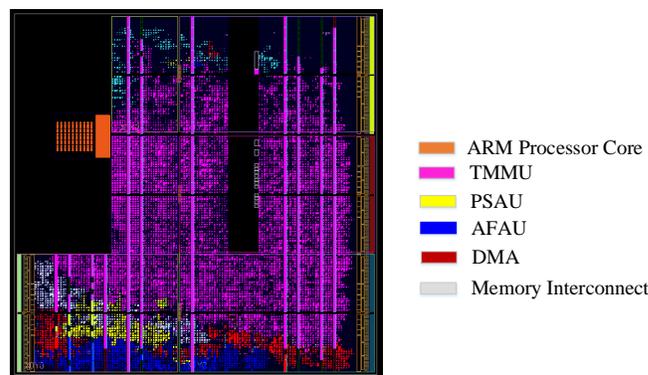


Fig. 6. Floorplan of the FPGA Chip

V. CONCLUSION AND FUTURE WORK

In this article we have presented DLAU, which is a scalable and flexible deep learning accelerator based on FPGA. The DLAU includes three pipelined processing units, which can be reused for large scale neural networks. DLAU uses tile techniques to partition the input node data into smaller sets and compute repeatedly by time-sharing the arithmetic logic. Experimental results on Xilinx FPGA prototype show that DLAU can achieve 36.1x speedup with reasonable hardware cost and low power utilization.

The results are promising but there are still some future directions, including optimization of the weight matrix and memory access. Also the trade-off analysis between FPGA and GPU accelerators is another promising direction for large scale neural networks accelerations.

REFERENCES

- [1] LeCun, Y., Y. Bengio, and G. Hinton, Deep learning. *Nature*, 2015. 521: p. 436-444.
- [2] Hauswald, J., et al. DjiNN and Tonic: DNN as a service and its implications for future warehouse scale computers. in *ISCA* 2015.
- [3] Zhang, C., et al. Optimizing FPGA-based Accelerator Design for Deep Convolutional Neural Networks. in *FPGA* 2015.
- [4] Thibodeau, P. Data centers are the new polluters. 2014 [cited 2015].
- [5] Ly, D.L. and P. Chow, A high-performance FPGA architecture for restricted boltzmann machines, in *FPGA* 2009.
- [6] Chen, T., et al., DianNao: a small-footprint high-throughput accelerator for ubiquitous machine-learning, in *ASPLOS* 2014. p. 269-284.
- [7] Kim, S.K., et al. A highly scalable restricted boltzmann machine FPGA implementation. in *FPL* 2009.
- [8] Qi Yu, et al. A Deep Learning Prediction Process Accelerator Based FPGA. *CCGRID* 2015: 1159-1162
- [9] Jiantao Qiu, et al. Going Deeper with Embedded FPGA Platform for Convolutional Neural Network. *FPGA* 2016: 26-35