

Towards Systems Education for Artificial Intelligence: A Course Practice in Intelligent Computing Architectures*

Jianlei Yang[†]
SCSE, BDBC, Beihang University
Beijing, P. R. China
jianlei@buaa.edu.cn

Xiaopeng Gao
SCSE, Beihang University
Beijing, P. R. China
gxp@buaa.edu.cn

Weisheng Zhao
SME, BDBC, Beihang University
Beijing, P. R. China
weisheng.zhao@buaa.edu.cn

ABSTRACT

With the rapid development of artificial intelligence (AI) community, education in AI is receiving more and more attentions. There have been many AI related courses in the respects of algorithms and applications, while not many courses in system level are seriously taken into considerations. In order to bridge the gap between AI and computing systems, we are trying to explore how to conduct AI education from the perspective of computing systems. In this paper, a course practice in intelligent computing architectures are provided to demonstrate the system education in AI era. The motivation for this course practice is first introduced as well as the learning orientations. The main goal of this course aims to teach students for designing AI accelerators on FPGA platforms. The elaborated course contents include lecture notes and related technical materials. Especially several practical labs and projects are detailed illustrated. Finally, some teaching experiences and effects are discussed as well as some potential improvements in the future.

CCS CONCEPTS

• **Computer systems organization** → **Neural networks**; • **Hardware** → **Application specific integrated circuits**; **Hardware accelerators**.

KEYWORDS

Artificial Intelligence, System Education, Neural Network Accelerators, Intelligent Computing Architectures

ACM Reference Format:

Jianlei Yang, Xiaopeng Gao, and Weisheng Zhao. 2020. Towards Systems Education for Artificial Intelligence: A Course Practice in Intelligent Computing Architectures. In *ACM Great Lakes Symposium on VLSI 2020 (GLSVLSI '20)*, September 7–9, 2020, Virtual Event, China. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3386263.3406935>

*A Microelectronic Systems Education Paper

[†]Corresponding author is Jianlei Yang. Email: jianlei@buaa.edu.cn

This work is supported in part by the National Natural Science Foundation of China (61602022), State Key Laboratory of Software Development Environment (SKLSDE-2018ZX-07), CCF-Tencent IACR20180101 and the 111 Talent Program B16001.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ACM GLSVLSI '20, September 7–9, 2020, Virtual Event, China

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7944-1/20/09...\$15.00

<https://doi.org/10.1145/3386263.3406935>

1 INTRODUCTION

In recent years, artificial intelligence (AI) has made great progress due to the various deep learning (DL) algorithms and applications. Actually in scientific research community, AI's research has been extended from algorithm/application level to system level. More and more attentions have been taken into system design problems in AI computing. From the point of view of ACM Turning Lectures [1–3], the computer architecture's evolution is driven by AI's revolution. What's more, many interdisciplinary research topics will be inspired by AI and computer architectures/systems, such as the recent rising Conference on Machine Learning and Systems (MLSys) [4].

Although AI technologies have been well developed in the scope of scientific research and industrial applications, more attentions should be taken into education scope in order to satisfy the urgent requirements of industrial technologies developments. Nowadays the curriculum system in most of colleges or universities is relative old and lacks of enough novelty compared with the rapid development in technical stacks and social applications. If most of the students are trained with lacking of enough system capabilities, they cannot meet the full-stack talents requirement of the AI industry. Hence, the personnel training for AI community cannot lack of computing system hierarchies. Meanwhile, the education in computer system has to consider the AI-inspired architectures and systems. In summary, the gap between education and industry has to be bridged by incorporating the computing architecture/systems and AI algorithms/applications.

There have been many courses in AI algorithms and applications as well as the traditional computer architecture and systems. And the related resourceful course experiments have a very good training effect. However, it lacks of sufficient and timely exploration in the teaching AI-inspired systems and architectures. If AI education only focuses on the higher level topics, the students will lack of the vision and capability to understand the system performance issues on AI computing. Therefore, with the goal of cultivating talents in the full stack of AI and starting from cultivating students' computer system capabilities, we will introduce the demand of computing system in the AI era into the curriculum education. The main goals of this course exploration are listed as following:

- *Bridge the education gap between AI and computing architecture/system*: broaden the students' research vision and towards full-stack development skills in AI-related techniques.
- *Introduce emerging research topics on AI-inspired architecture/system into college courses*: motivate students to learn and study computer architecture/system instead of only focusing on the algorithms/applications in AI era.

Table 1: Related courses in ML architecture and system.

Affiliation	Course Name
UIUC	Machine Learning in Silicon ¹
Beihang Univ.	Intelligent Computing Architectures
MIT	Hardware Architecture for Deep Learning ²
Stanford	Hardware Accelerators for Machine Learning ³
Univ. of CAS	Intelligent Computing Systems ⁴
UC Berkeley	AI-Sys ⁵
U. Washington	System for ML ⁶
Georgia Tech.	Hardware Acceleration for Machine Learning ⁷

¹ <https://courses.grainger.illinois.edu/ece598ns/fa2017/>² <https://www.eecs.mit.edu/academics-admissions/academic-information/subject-updates-spring-2019/6s0826888>³ <https://cs217.stanford.edu/>⁴ <http://novel.ict.ac.cn/aics/>⁵ <https://ucbrise.github.io/cs294-ai-sys-sp19/>⁶ <https://dlsys.cs.washington.edu/>⁷ http://tusharkrishna.ece.gatech.edu/teaching/hml_s19/

2 BACKGROUND AND RELATED COURSES

The Turing Lecture [1] points out that domain-specified computing architecture is an important trend in the future customization computing. Especially in recent years, how to design high-performance and energy-efficient domain-specified architectures have become as very popular research directions due to the rapid growth of AI applications on the urgent requirements for computing [5].

As shown in Figure 1, different kind of computing platforms are compared in terms of adaptivity (AD), performance (PE), power efficiency (PO), programmability (PR) and scalability (SC) [6, 7]. General computing platforms, such as CPU or GPU, have good advantages in adaptivity, scalability and programmability, but their performance and power efficiency are greatly limited for specialized computing purpose. Application specified platforms, such as NPUs, could achieve very high performance, and power efficiency, but are very poor in adaptivity and scalability. Reconfigurable computing platforms, such as FPGAs, have a good balance in performance, efficiency and adaptivity, but do not have obvious advantages in each dimension while their advantages are convenient for rapid design iteration. In addition, there are some emerging NVM-based neuromorphic computing platforms, which have good advantages in scalability, performance and efficiency, but also have obvious limitations in adaptivity and programmability due to lack of available development tools or methodologies. As a result, it is difficult to find a computing platform that is excellent in all dimensions for domain customized architecture and chip design. In this course, the idea of domain-specified architecture design is introduced and some developing methodologies are provided to train students for design skills improvements. In terms of platform consideration, FPGA is adopted in this course for rapid development iteration.

Recently, some relevant courses have appeared in many universities as shown in Table 1. Most of these courses are focusing on AI-inspired architecture and system education, which have received great attention and good effects. Prof. Naresh R. Shanbhag in UIUC

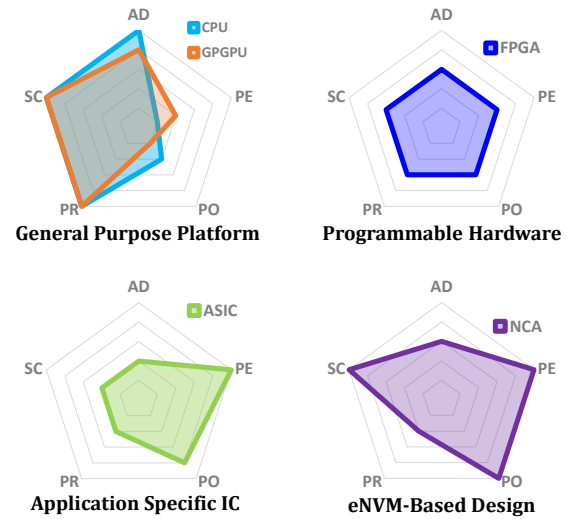


Figure 1: Comparison among different design platforms [6, 7], where AD means Adaptivity, PE means Performance, PO means Power Efficiency, PR means Programmability, SC means Scalability.

offers a course named *Machine Learning in Silicon* since Fall 2017. This course aims to teach students how to design machine learning algorithms in chips directly, which mainly focuses on circuit and system implementations [8]. Prof. Vivienne Sze and Joel Emer in MIT offer a course named *Hardware Architecture for Deep Learning* since Fall 2017. This course is more likely a tutorial to present their researches in DL architectures [9]. Prof. Ardavan Pedram and Kunle Olukotun in Stanford Univ. offer a course named *Hardware Accelerators for Machine Learning* since Fall 2018. This course introduces many latest hardware accelerators for machine learning applications. Dr. Yunji Chen in Univ. of CAS offer a course named *Intelligent Computing Systems* since Fall 2018. This course mainly promote their research achievements of DianNao architectures as well as their Cambricon chips and ecosystems [10]. They have published a relevant teaching textbook and provided resourceful experimental materials. This course has been promoted to several universities in China, which has a great influence in education community. Since 2019, there have been several courses focused on ML system level. Prof. Ion Stoica and Joseph E. Gonzalez in UC Berkeley offered a course named *AI-Sys*. It describes the latest trends in systems designs to better support the next generation of AI applications, and applications of AI to optimize the architecture and the performance of systems. Prof. Tianqi Chen in Univ. of Washington offered a course named *System for ML* since Spring 2019. This course is designed to fill the gap in how to build and optimize these deep learning specified systems. Prof. Tushar Krishna in Georgia Tech. offered a course named *Hardware Acceleration for Machine Learning* since Spring 2019. This course present various development resources that can enable researchers and practitioners to quickly get started on DNN accelerator design, and highlight important benchmarking metrics and design considerations that should be used




Figure 2: Course prerequisites and following courses.

for evaluating the rapidly growing number of DNN hardware platforms being proposed in academia and industry. In summary, these courses listed above provided a very useful early exploration by introducing circuit/chip/architecture/system levels considerations on how to build and optimize deep learning computing systems.

Our course, named *Intelligent Computing Architectures*, mainly focused on the architectures explorations for deep learning tasks since Fall 2017. Most attentions are paid to present recent advances in architectures towards the goal of enabling efficient processing of DNNs. This course requires the students to have a wide scope of knowledge and strong practical capabilities which is very challenging to finish the experiments and projects. But the students' system developing capabilities will be greatly improved once they have gone through the intensive training in this course.

3 COURSE DESIGN AND IMPLEMENTATION

3.1 Course Prerequisites

This course presents how to design domain-specific hardware architectures, which requires students to learn and practice both hardware and software knowledge. The prerequisites include three parts: computer architecture and system, hardware design, and algorithm related background. For the system part, students are first required to learn computer architectures to understand how to build instructions set and optimize hardware implementations. Also some background of computer compiler are required for mapping ML algorithms as instructions. For the hardware design part, students are required to learn digital circuit & system, Verilog HDL language, as well as FPGA development techniques for hardware implementations. For the algorithms related background, students are required to know some ML introductions and know well about programming techniques, including C, C++ and Python programming. With these backgrounds, students in this course could learn how to design ML architectures and implement them into FPGA platform with the support of dataflow instructions compilation. Based on these achievements, students could try to learn how to design and optimize ML systems with the architectural level considerations.

3.2 Teaching Objectives

In this course, some emerging research are introduced into classroom so that it involves many basic knowledge and skills, especially the software-hardware co-design methodologies. This course requires students to have very strong learning capabilities and operational capabilities. Based on several typical applications in the field of intelligent computing, this course introduces the design and implementation methodologies of mapping software/algorithms to hardware architectures. From this course, students could understand the working principles of intelligent computing in the point of view of computer architectures. With the course experimental training, students could obtain the skills of designing and implementing domain-specified hardware for intelligent algorithms. This course requires students to learn and know how to independently design and prototype several typical intelligent computing accelerators. Their engineering opinions of heterogeneous system design could be built for extending further research vision. The main teaching objectives of this course could be listed as following:

- (1) Know some typical intelligent computing algorithms, including machine learning, data mining, etc. Know some basic problems and techniques in hardware acceleration area.
- (2) Master some design and optimization methodologies of digital systems. Know how to map intelligent algorithms into hardware architectures.
- (3) Familiar with development tools, such as Verilog, Modelsim, Vivado, etc. Familiar with deep learning frameworks, such as Caffe, PyTorch, TensorFlow, etc.
- (4) Know how to access and utilize technical documentations.
- (5) Master embedded system design and optimization techniques.
- (6) Obtain modelling, design and analysis skills of intelligent computing systems, as well as the innovation capabilities in solving practical problems through acquired knowledge and skills.

3.3 Teaching Contents and Practices

As listed in Table 2, this course consists of 32 teaching hours and mainly includes three parts: basic preparation, accelerator design, further explorations. In the basic preparation part (6 hours), intelligent computing architectures are first introduced as well as the mainstream computing models and methods. It motivates students' interests to learn this course and know some research and technical backgrounds in related communities. Meanwhile, some DNNs are described in the class and students evaluate some DNN models to understand their computational dataflow.

In the accelerator design part (20 hours), system architectures and dataflow mappings are detailed presented. The basic FPGA design flows are introduced for development in Zynq platforms. Some resourceful Zynq materials, including documentations, design examples, application notes, etc., are provided to students for learning and developing. Then, some popular architectures of DNN accelerators are discussed, which are arisen from latest research papers or chips in recent 6 years, such as DianNao[11], Eeyriss[12], TPU[13], etc. Furthermore, a simple DNN architecture is provided as a reference to our labs and projects. Several typical operations are defined as instructions or pseudo-instructions to support our architectures for running whole DNN networks. Meanwhile, dataflow compiling or mapping techniques are introduced. With the predefined ISA,

DNN dataflow is translated as instructions which could run in the designed architectures. Some optimization techniques could be involved in this process to improve the accelerator's performance or energy efficiency.

In the further explorations part (6 hours), ASIC implementation flow is introduced and accelerators' applications are discussed. Beyond the implementations on FPGA platforms, some ASIC design flows could broaden the students' research or technical visions. Also some further perspectives of DNN accelerators' applications could motivate students to learn and explore potential possibilities in this emerging research direction.

Besides the teaching contents in the class, several labs and projects are provided as shown in Table 3. Lab 1-4 help students to understand the DNN algorithms and frameworks as well as the FPGA developing skills. Based on these four labs, students are required to build a DNN accelerator step by step among project 1 to project 3. Aiming to finish these tasks, students are usually required to spend much more times than in class. Lab 1 is provided to make students familiar with the frequently used DNN algorithms. Lab 2 aims to evaluate DNNs with deep learning frameworks, which is important to balance the model accuracy and performance/efficiency in

architectural explorations. Lab 3 provides a chance to learn Zynq FPGA developing skills while many design examples are provided as a reference. Lab 4 requires students to design a multiply-and-accumulation (MAC) module with Verilog, which should be validated in PL part of Zynq FPGA. It should be not difficult for students who have learnt Verilog and digital systems design courses before. These four labs aims to make students familiar with DNN fundamentals and hardware design platforms.

These labs and projects are provided progressively towards a practical DNN accelerator design. The main target of this course aims to design DNNs accelerator in FPGA which is covered by the remaining three projects. In project 1, the designed MAC module in Lab 4 is integrated with ARM processor as main controller in PS part. Vivado SDK is utilized for software developing while the PS-PL communication mechanisms should be well exploited for software-hardware co-design. Project 2 aims to design a LeNet on Zynq FPGA while all computing and data buffering tasks could be finished within on-chip resources. Since the LeNet model size is relatively small, such as less than 1Mb, the model weights and intermediate data could be handled just with on-chip block RAMs. Project 3 is our final project, which targets to design a VGGNet on Zynq FPGA. The VGGNet model size is very large so that we have to store most of weights and intermediate data in off-chip DDR memory.

Table 2: Teaching contents and class hours in this course, while total teaching hours are 32 hours in 16 weeks (2 hours for each week, total 16 weeks).

Teaching Contents	Schedule
Introduction to intelligent computing architectures: the motivations to learn this course, the teaching scope and required learning capabilities.	2 Hours Week 1
Mainstream computing models and methods: DNNs and graph computing methods, especially targeted for hardware design issues.	4 Hours Week 2-3
Domain-specified architecture design methodologies: design principles and development flow from algorithms to hardware, targeted on FPGA or ASIC implementations.	4 Hours Week 4-5
Compiling or mapping methodologies: system modeling, functionalities partitioning, dataflow mapping and scheduling, performance and efficiency optimization techniques.	6 Hours Week 6-8
DNN accelerators design: algorithms evaluations, dataflow architectures, hardware design with Verilog/HLS implementations targeted on Xilinx Zynq FPGA platforms.	10 Hours Week 9-13
ASIC implementation flow: a brief introduction including behavior description, logic synthesis, physical implantation, system verification, timing analysis and optimization, etc.	4 Hours Week 14-15
Application perspectives of hardware accelerators: an outlook of practical applications in big data analytic, DL, CV, robotics, etc.	2 Hours Week 16

Table 3: Labs and projects design in this course.

No.	Lab or Project Details (<i>Week Schedule</i>)
Lab 1	Deep neural networks: learn and write C++/Python code for some deep neural networks to understand the learning principles and network structures. <i>Week 3</i>
Lab 2	Deep learning frameworks: run deep neural networks with Caffe/TensorFlow/PyTorch frameworks, evaluate the model accuracy and performance. <i>Week 4</i>
Lab 3	Zynq FPGA development: learn Zynq FPGA development flows, including Vivado design suite, IP usage, AXI bus protocol, PS and PL co-design, etc. <i>Week 5-6</i>
Lab 4	MAC module design with Verilog: implement Verilog modules for matrix multiplication and accumulation, validate on PL part of Zynq FPGA. <i>Week 7-8</i>
Prj 1	MAC design on Zynq FPGA: implement MAC modules on PL part and perform validation with the controlling of PS part in Zynq FPGA. <i>Week 9-10</i>
Prj 2	LeNet design on Zynq FPGA: implement the data path and controller to run LeNet on Zynq FPGA, only on-chip BRAMs are exploited, where the PS part is in charge of data input/output and dataflow controlling, the PL part is in charge of computations. <i>Week 11-13</i>
Prj 3	VGGNet design on Zynq FPGA: VGGNet implementation which is similar to LeNet, where the off-chip DRAM is also utilized for buffering the intermediate data as well as the on-chip BRAMs. <i>Week 14-16</i>

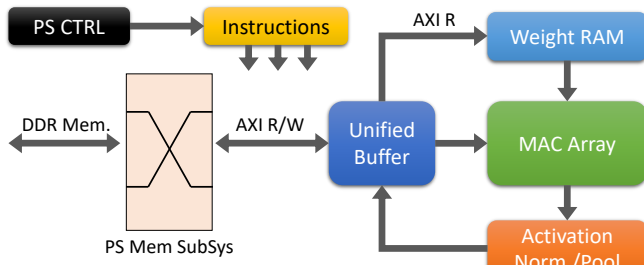


Figure 3: A typical architecture reference for DNN accelerator design, where three factors are mainly considered: data transfer bandwidth, MAC array scale, on-chip buffer size.

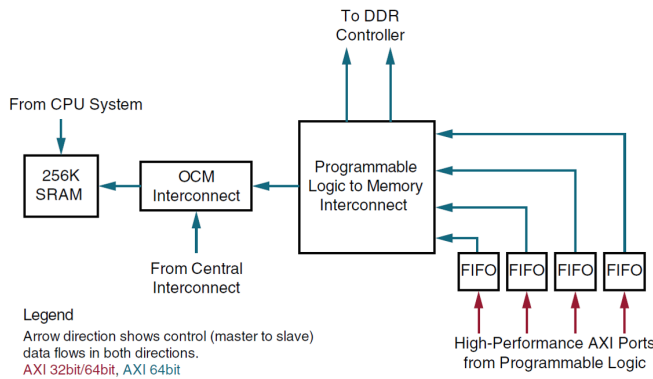


Figure 4: Data access from PL part to DDR via the PL Interface to PS Memory Subsystem [14].

A typical architecture is provided as shown in Figure 3, which is a simple reference for students to design DNN accelerators. The original data including network models and images are pre-stored in on-board SD card. The PS part in Zynq is in charge of system controlling, including data input/output, data caching and computing, etc. The memory system includes CPU cache, OCM, block RAMs, off-chip DRAM, etc. Data access between PL and PS is maintained by a memory interconnect block as shown in Figure 4. Especially when DRAM is utilized for data accessing by PL part, it requires to use the PL interface to PS memory sub-system. Once system booting is finished, DNN models are loaded into DRAM and partially stored in unified buffer (in BRAMs) while weights are stored in weight RAMs. MAC array is the core computation part to finish multiplication and accumulation. After computation of MAC array, a post-processing part is utilized to compute activation, normalization, pooling, etc.

DNN computation dataflow is compiled as instructions to control typical operations, including computing, data caching, data transfer, etc. Several typical instructions includes:

- Read_Host_Memory: read data from PS to unified buffer.
- Read_Wights: read weights from DRAM to weight FIFO.
- MatrixMultiply/CONV: convolutional computation.
- Activate: activation or pooling computation.
- Write_Host_Memory: write results from unified buffer to DRAM.




Figure 5: Finite state machine description for a typical computing dataflow.

Table 4: Resources utilization comparison between Google TPU and our referred NaïveTPU design.

Resources	Google TPU	Our NaïveTPU
Matrix Multiply Unit	256×256	32×32
Accumulators RAM	$4K \times 256 \times 32b$	$4K \times 32 \times 32b$
Unified Buffer	$96K \times 256 \times 8b$	$16K \times 32 \times 8b$

The DNN computation dataflow is processed by a row-by-row manner while the finite state machine (FSM) is described as Figure 5. It starts from data preparation, i.e., data have been pre-loaded into unified buffer or weight RAM. If the required data have been ready for each row convolutions, they are sent to MAC array for CONV computing. In our course, a typical systolic array is provided for CONV computing. As shown in Table 4, our utilized NaïveTPU is a reduced version of Google TPU. The MAC array scale is reduced according to the available logic resources in our utilized Zynq chip. Unified buffer, FIFOs, and RFs are also determined by the available storage resources in our utilized Zynq chip.

For course labs and projects, we provide a Zynq FPGA board for each student as shown in Figure 6. The utilized board is named AX7020, with a Zynq chip of XC7Z0202CLG400I, including a dual Cortex-A9 ARM core (PS), 4.9Mb BRAM, and on-board 8GBit DDR3 SDRAM. For running LeNet, on-chip BRAM is enough for data caching. But for VGGNet, off-chip DRAM is required to store most of the model weights and intermediate data. Students are required to understand how to access DDR from PL via the interface to PS memory sub-system. Some important demos or reference designs are provided to help students to finished their project, including: SD card access API, DDR memory access API, AXI controller, BRAM usage, UART uasge, etc. All of these examples are included in the AX7020 suite with comprehensive code examples and documentations.

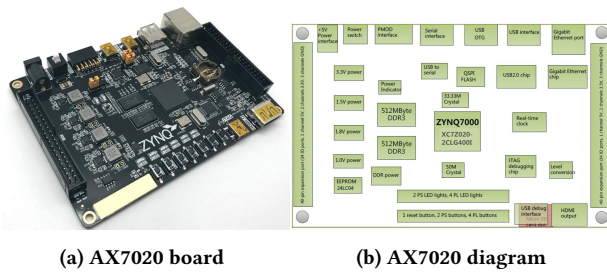


Figure 6: The utilized AX7020 Zynq FPGA board [15] in this course.

4 TEACHING PRACTICE AND DISCUSSION

In this section, we mainly discuss some key observations we have acquired from teaching practice in each Fall semester since 2017 to illustrate the efficiency and feasibility of our course in motivating students to study AI-inspired architectures.

This course requires students to spend many time in reading documentations and system developing. These efforts are necessary to well understand the architectures and FPGA developing flows. From our observations, most of the students are very interested in learning these knowledge and willing to spend many efforts to finish our labs/projects. All students could finish the lab 1-4 to project 1 well since these tasks are fundamental requirements. About 80% students could finish the LeNet design project and demo their functions to us. The remaining 20% students usually have some difficulties in timing optimization and software-hardware co-design stage. For VGGNet design, about 60% students could demo their full functionalities. Some of them even could achieve a relatively high performance when evaluated on ImageNet dataset. The remaining 40% students usually could not handle the complicated FSMs since there are many data transfer and data reuse strategies for VGGNet computations.

From our observations, the most excellent students (about 10%) in this course could design a very robust DNN accelerator in Zynq FPGA while their performance and efficiency could be very close to the reported results in some early research papers [16]. Meanwhile, a survey from the students in this course of first two years shows that the top excellent students (about 10%) choose deep learning accelerators as their research directions in the postgraduate study. It indicates that most of the students in this course are very interested in our teaching even though it is very challenging and has many difficulties. These students have learnt lots of knowledge in DNN architectures and achieved strong system developing skills in this course. What's more, their inspired interests and confidence enable them contribute to this research field in the future.

5 LIMITATIONS AND FUTURE WORK

Course Workloads. From our course practice in previous three years, we have noticed that students have a very heavy workloads both in leaning new knowledge and mastering new developing skills. We plan to suggest students to further improve their FPGA developing skills and practice more DNN algorithms in their pre-acquired courses. Hence, we could introduce more architectural details in DNN accelerators. And furthermore, students could try

to design more complicated architectures and optimize their performance and efficiency. These achievements could provide a stronger background for their future research.

System Capabilities. This course mostly focuses on architectural design issues while lacking of enough system level implementations. Aiming to enable students to be qualified with full-stack techniques, we would like to further offer a new course from Fall 2020, named *Intelligence Computing Systems*, to provide more system level techniques. This new course will mostly focus on compilers of mapping DNN algorithms into specified architectures, as well as the related optimization methodologies.

6 CONCLUSIONS

In this paper, we present a new course focused on intelligent computing architectures. Based on recent researches, this course introduces lots of domain-specified architectures and provides resourceful experimental materials. These training tasks are very challenging for students but they are motivated to study and practice deeply. We believe that such practice could be helpful to cultivate students towards full-stack system developing capabilities in artificial intelligence era.

REFERENCES

- [1] John L Hennessy and David A Patterson. A new golden age for computer architecture. *Communications of the ACM*, 62(2):48–60, 2019.
- [2] Geoffrey Hinton. The deep learning revolution. Turing Lecture at FCRC 2019, June 2019.
- [3] Yann LeCun. The deep learning revolution: The sequel. Turing Lecture at FCRC 2019, June 2019.
- [4] The Conference on Machine Learning and Systems (MLSys). <https://mlsys.org/>.
- [5] Jason Cong, Vivek Sarkar, Glenn Reinman, and Alex Bui. Customizable domain-specific computing. *IEEE Design & Test of Computers*, 28(2):6–15, 2010.
- [6] Divya Mahajan, Jongse Park, Emmanuel Amaro, Hardik Sharma, Amir Yazdanbakhsh, Joon Kyung Kim, and Hadi Esmaeilzadeh. Tabla: A unified template-based framework for accelerating statistical machine learning. In *IEEE International Symposium on High Performance Computer Architecture (HPCA)*, pages 14–26, 2016.
- [7] Jongse Park, Hardik Sharma, Divya Mahajan, Joon Kyung Kim, Preston Olds, and Hadi Esmaeilzadeh. Scale-out acceleration for machine learning. In *Proceedings of IEEE/ACM International Symposium on Microarchitecture (MICRO)*, pages 367–381, 2017.
- [8] Mingyu Kang, Prakash Srivastava, Vikram Adve, Nam Sung Kim, and Naresh R Shanbhag. An energy-efficient programmable mixed-signal accelerator for machine learning algorithms. *IEEE Micro*, 39(5):64–72, 2019.
- [9] Vivienne Sze, Yu-Hsin Chen, Tien-Ju Yang, and Joel S Emer. Efficient processing of deep neural networks: A tutorial and survey. *Proceedings of the IEEE*, 105(12):2295–2329, 2017.
- [10] Cambricon Tech. Inc. <http://www.cambricon.com/>.
- [11] Tianshi Chen, Zidong Du, Ninghui Sun, Jia Wang, Chengyong Wu, Yunji Chen, and Olivier Temam. Dianao: A small-footprint high-throughput accelerator for ubiquitous machine-learning. *ACM SIGARCH Computer Architecture News*, 42(1):269–284, 2014.
- [12] Yu-Hsin Chen, Joel Emer, and Vivienne Sze. Eyeriss: A spatial architecture for energy-efficient dataflow for convolutional neural networks. *ACM SIGARCH Computer Architecture News*, 44(3):367–379, 2016.
- [13] Norman P Jouppi, Cliff Young, Nishant Patil, David Patterson, Gaurav Agrawal, Raminder Bajwa, Sarah Bates, Suresh Bhatia, Nan Boden, Al Borchers, et al. In-datacenter performance analysis of a tensor processing unit. In *Proceedings of ISCA*, pages 1–12, 2017.
- [14] Xilinx Inc. DS190: Zynq-7000 SoC First Generation Architecture. https://www.xilinx.com/support/documentation/data_sheets/ds190-Zynq-7000-Overview.pdf, July 2, 2018.
- [15] ALINX Inc. ALINX AX7020 Xilinx Zynq FPGA. <http://www.alinx.com/index.php/default/content/97.html>, 2020.
- [16] Jiantao Qiu, Jie Wang, Song Yao, Kaiyuan Guo, Boxun Li, Erjin Zhou, Jincheng Yu, Tianqi Tang, Ningyi Xu, Sen Song, et al. Going deeper with embedded FPGA platform for convolutional neural network. In *Proceedings of ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*, pages 26–35, 2016.