

RESEARCH

Open Access



Deep learning accelerators: a case study with MAESTRO

Hamidreza Bolhasani*  and Somayyeh Jafarali Jassbi

*Correspondence:
hamidreza.bolhasani@srbiau.
ac.ir
Department of Computer
Engineering, Science
and Research Branch, Islamic
Azad University, Tehran, Iran

Abstract


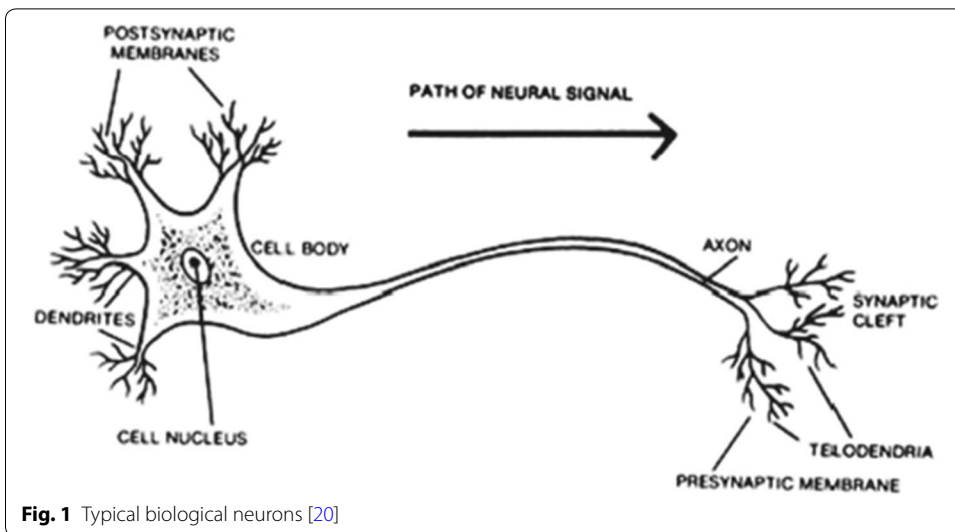
In recent years, deep learning has become one of the most important topics in computer sciences. Deep learning is a growing trend in the edge of technology and its applications are now seen in many aspects of our life such as object detection, speech recognition, natural language processing, etc. Currently, almost all major sciences and technologies are benefiting from the advantages of deep learning such as high accuracy, speed and flexibility. Therefore, any efforts in improving performance of related techniques is valuable. Deep learning accelerators are considered as hardware architecture, which are designed and optimized for increasing speed, efficiency and accuracy of computers that are running deep learning algorithms. In this paper, after reviewing some backgrounds on deep learning, a well-known accelerator architecture named MAERI (Multiply-Accumulate Engine with Reconfigurable interconnects) is investigated. Performance of a deep learning task is measured and compared in two different data flow strategies: NLR (No Local Reuse) and NVDLA (NVIDIA Deep Learning Accelerator), using an open source tool called MAESTRO (Modeling Accelerator Efficiency via Spatio-Temporal Resource Occupancy). Measured performance indicators of novel optimized architecture, NVDLA shows higher L1 and L2 computation reuse, and lower total runtime (cycles) in comparison to the other one.

Keywords: Deep learning, Convolutional neural networks, Deep neural networks, Hardware accelerator, Deep learning accelerator

Introduction

The main idea of neural networks (NN) is based on biological neural system structure, which consists of several connected elements named neurons [1]. In biological systems, neurons get signals from dendrites and pass them to the next neurons via axon as shown in Fig. 1.

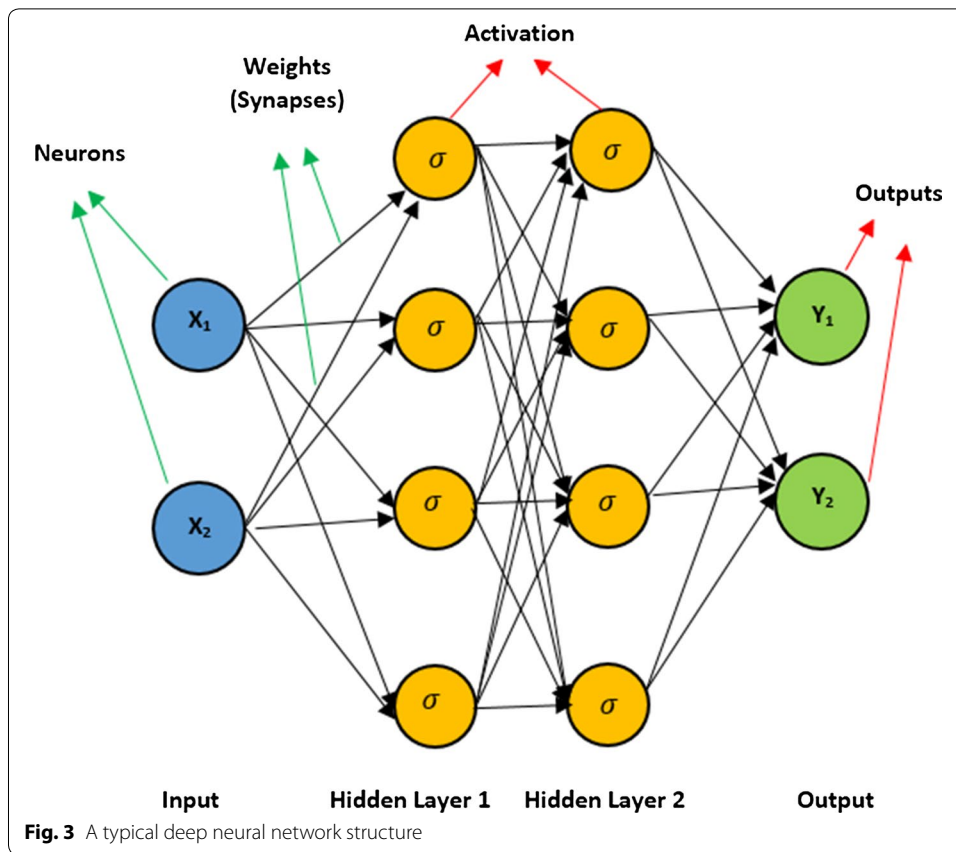
Neural networks are made up of artificial neurons for handling brain tasks like learning, recognition and optimization. In this structure, the nodes are neurons, links can be considered as synapses and biases as activation thresholds [2]. Each layer extracts some information related to the features and forwards them with a weight to the next layer. Output is the sum of all these information gains multiplied by their related weights. Figure 2 represents a simple artificial neural network structure.



Deep neural networks are complex artificial neural networks with more than two layers. Nowadays, these networks are widely used for several scientific and industrial purposes such as visual object detection, segmentation, image classification, speech recognition, natural language processing, genomics, drug discovery, and many other areas [3].

Deep learning is a new subset of machine learning including algorithms that are used for learning concepts in different levels, utilizing artificial neural networks [4].

As Fig. 3 shows, if each neuron and its weight are represented by X_i and W_{ij} respectively, the output result (Y_j) would be:



$$Y_j = \sum_{i=1}^n \sigma(W_{ij}X_i) \tag{1}$$

where σ is the activation function. A popular function that is used for activation in deep neural networks is ReLU (Rectified Linear Unit) function, which is defined in Eq. (2).

$$[\sigma(z)]_j = \max\{z_j \text{ and } 0\} \tag{2}$$

Leaky ReLU, tanhh and Sigmoid functions are some other activation functions with less frequent usage [5].

$$\sigma(z) = \frac{1}{1 + e^{-z}} \tag{3}$$

As shown in Fig. 4, each layer of a deep neural network’s role is to extract some features and send them to the next layer with its corresponding weight. For example, in the first layer, color properties (green, red blue) are gained; in the next layer, edge of objects are determined and so on.

Convolutional neural networks are a type of deep neural networks that is mostly used for recognition, mining and synthesis applications like face detection, handwriting recognition and natural language processing [6]. Since parallel computations is an unavoidable part of CNNs, several efforts and research works have been done for designing




Fig. 4 Deep learning setup for object detection [21]

an optimized hardware for it. As a result, many application-specific integrated circuits (ASICs) as hardware accelerators have been introduced and evaluated in the recent decade [7]. In the next section, some of the most successful and impressive works related to CNN accelerators are introduced.

Related works

Tianshi et al. [8] proposed DianNao as a hardware accelerator for large-scale convolutional neural networks (CNNs) and deep neural networks (DNNs). The main focus of the suggested model is on the memory structure to be optimized for big neural network computations. The experimental results showed speedup in computation and reduction of overhead in performance and energy. This research also demonstrated that the accelerator can be implemented in very small area in order of 3 mm^2 and 485 mW power.

Zidong et al. [9] suggested ShiDianNao as a CNN accelerator for image processing close to a CMOS or CCD sensor. The performance and energy of this architecture is compared to CPU, GPU and DianNao, which has been discussed in previous work [8]. Utilizing SRAM instead of DRAM made it 60 times more energy efficient than DianNao. It is also $50\times$, $30\times$ and $1.87\times$ faster than a mainstream CPU, GPU and DianNao, with just 65 nm usage area and 320 mW power.

Wenyan et al. [6] offered a flexible dataflow accelerator for convolutional neural networks called FlexFlow. Working on different types of parallelism is the substantial contribution of this model. Results of the tests showed $2\text{--}10\times$ performance speedup and $2.5\text{--}10\times$ power efficiency in comparison with three investigated baseline architectures.

Eyriss is a spatial architecture for energy efficient data flows for CNNs which presented by Yu-Hsin et al. [10]. This hardware model is based on a dataflow named row stationary (RS). This dataflow minimizes energy consumption by reusing computation of filter weights. The proposed RS dataflow is investigated on AlexNet CNN configuration, which proved energy efficiency improvement.

Morph is a flexible accelerator for 3D CNN-based video processing that offered by Katrik et al. [7]. Since the previous work and proposed architectures didn't specifically focus on video processing, this model can be considered as a novelty in this area. Comparison of energy consumption in this architecture with previous idea, Eyriss [10] showed a high level of reduction that means energy saving. The main reason of this improvement is effective data reuse which reduces the access to higher level buffers and high cost off-chip memory.

Michael et al. [11] described Buffets that is an efficient and composable accelerator and independent of any particular design. Through this research, explicit decoupled data orchestration (EDDO) is introduced which allows evaluation of energy efficiency

in accelerators. Result of this work showed that with a smaller usage area, higher energy efficiency and lower control overhead is acquired.

Deep learning applications

Deep learning has a wide range of applications in recognition, classification and prediction, and since it tends to work like the human brain and consequently does the human jobs in a more accurate and low cost manner, its usage is dramatically increasing. More than 100 papers published from 2015 to 2020, helped categorize the main applications as below:

- Computer vision
- Translation
- Smart cars
- Robotics
- Health monitoring
- Disease prediction
- Medical image analysis
- Drug discovery
- Biomedicine
- Bioinformatics
- Smart clothing
- Personal health advisors
- Pixel restoration for photos
- Sound restoration in videos
- Describing photos
- Handwriting recognition
- Predicting natural disasters
- Cyber physical security systems [12]
- Intelligent transportation systems [13]
- Computed tomography image reconstruction [14]

Method

As mentioned previously, artificial intelligence and deep learning applications are growing drastically, but they have high complexity computation, energy consumption, costs and memory bandwidth. All these reasons were major motivations for developing deep learning accelerators (DLA) [15]. A DLA is a hardware architecture that is specially designed and optimized for deep learning purposes. Recent DLA architectures (e.g. OpenCL) have mainly focused on maximizing computation reuse and minimizing memory bandwidth, which led to higher speed and performance [16].

Generally, most of the accelerators support just fixed data flow and are not reconfigurable, but for doing huge deployments, they need to be programmable. Hyoukjun et al. [15] proposed a novel architecture named MAERI (Multiply-Accumulate Engine with Reconfigurable Interconnects), which is reconfigurable and employs ART (Augmented Reduction Tree) which showed 8~459% better utilization for different data flows over

a strict network-on-chip (NoC) fabric. Figure 5 shows the overall structure of MAERI DLA.

In another research, Hyoukjun et al. offered a framework called “MAESTRO” (Modeling Accelerator Efficiency via Spatio-Temporal Resource Occupancy) for predicting energy performance and efficiency in DLAs [17]. MAESTRO is an open-source tool that is capable of computing many NoC parameters for a proposed accelerator and related data flow such as maximum performance (roofline throughput), compute runtime, total runtime, NoC analysis, L1 to L2 NoC bandwidth, L2 to L1 bandwidth analysis, buffer analysis, L1 and L2 computation reuse, L1 and L2 weight reuse, L1 and L2 input reuse and so on. The topology, tool flow and relationship between each of its blocks of this framework are presented in Fig. 6.

Results and discussion

In this paper, we used MAESTRO to investigate buffer, NoC, and performance parameters of a DLA in comparison to a classical architecture for a specific deep learning data flow. For running MAESTRO and getting the related analysis, some parameters should be configured, as follows:

- LayerFile: Including the information related to the layers of neural network.
- DataFlow File: Information related to data flow.
- Vector Width: Width of the vectors.
- NoC Band width: Bandwidth of NoC.
- Multicast Supported: This logical indicator (True/False) is for defining that the NoC supports multicast or not.
- NumAverageHopsinNoC: Average number of hops in the NoC.
- NumPEs: Number of processing elements.

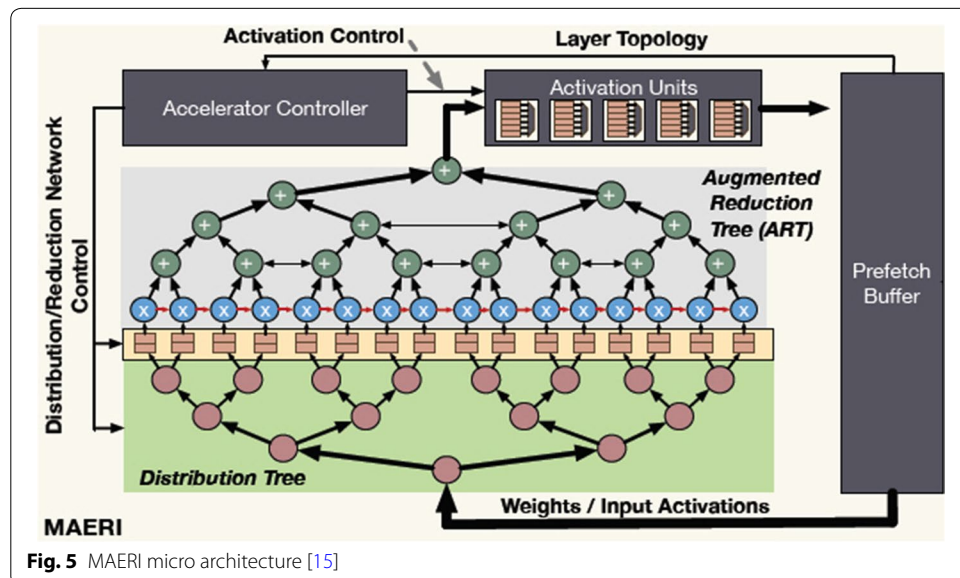


Fig. 5 MAERI micro architecture [15]

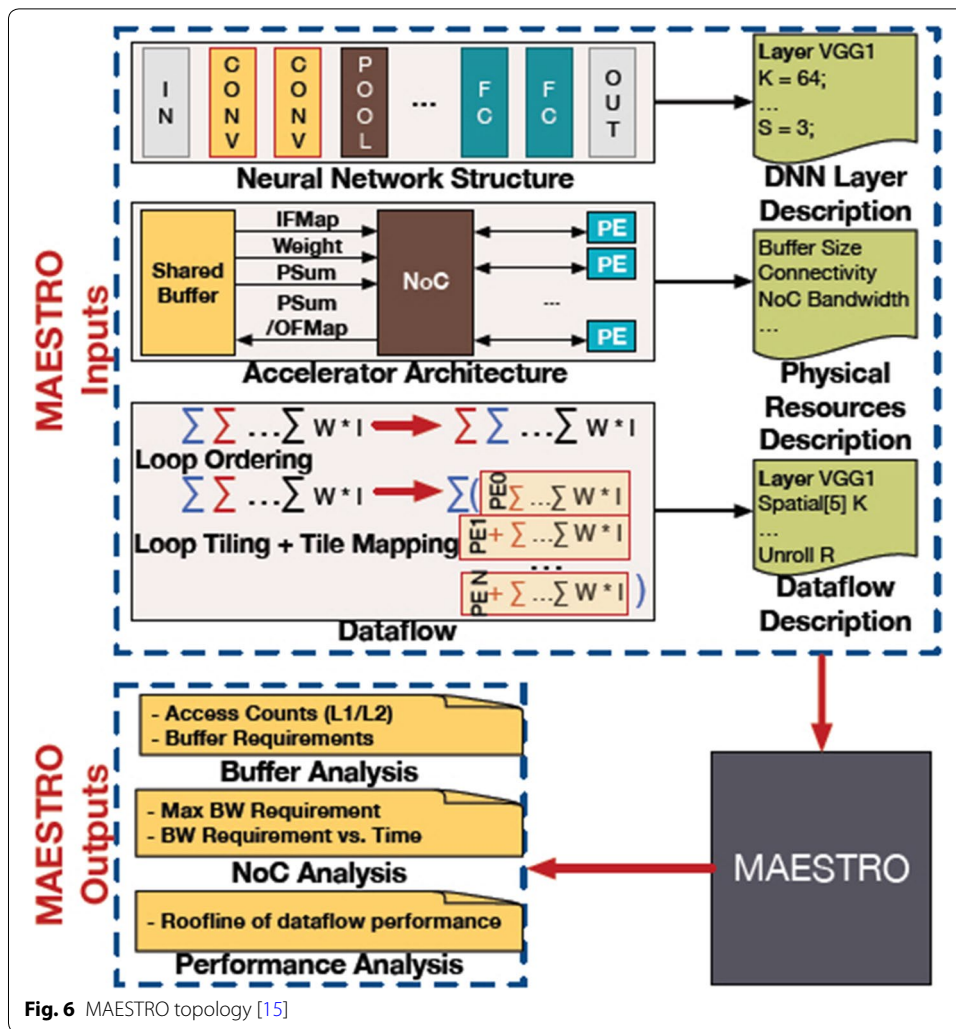


Fig. 6 MAESTRO topology [15]

For the simulation of this paper, we configured the mentioned parameters as presented in Table 1.

As presented in Table 1, we have selected Vgg16_conv11 as LayerFile, which is a convolutional neural network that has proposed by K. Simonyan and A. Zisserman. This deep convolutional network model was offered for image recognition with 92.7% accuracy on ImageNet dataset [18].

Two different data flow strategies are investigated and compared in this study: NLR and NVDLA. NLR stands for “No Local Reuse” which expresses its specific strategy and NVDLA is a novel DLA designed by NVIDIA Co [19].

Other parameters such as vector width, NoC bandwidth, multicast support capability, average numbers of hops and numbers of processing elements in NoC have been selected based on a real hardware condition.

Simulation results demonstrated that NVDLA has better performance, runtime, higher computation reuse and lower memory bandwidth in comparison to NLR as presented in Table 2 and Figs. 7, 8, and 9.

Table 1 MAESTRO configuration parameters

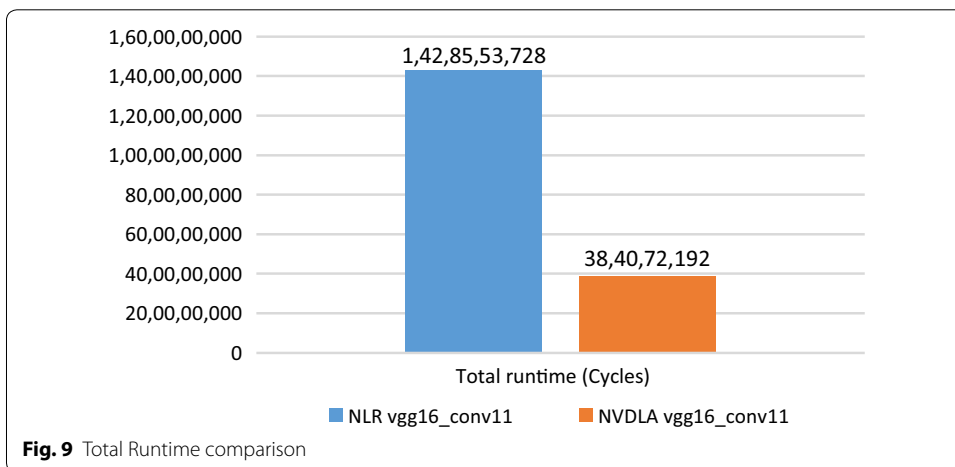
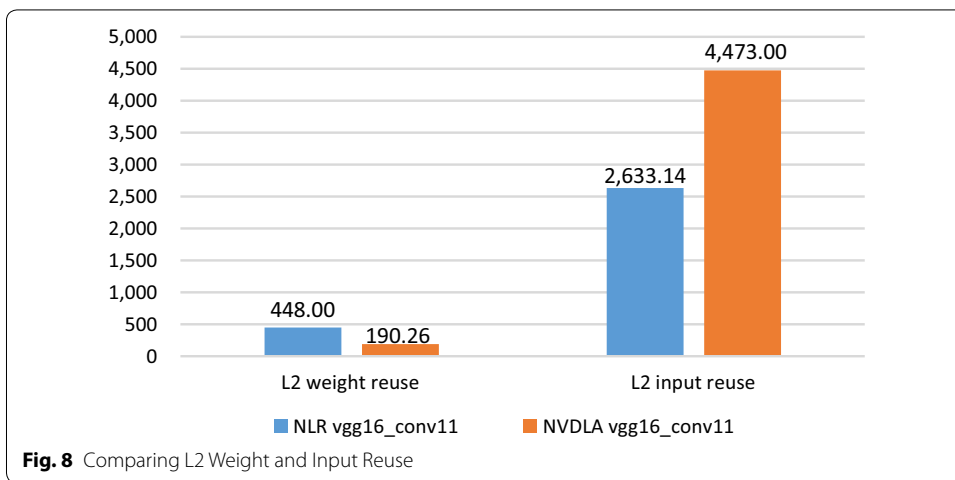

| No. | Input parameter | Value |
|-----|---------------------|------------------|
| 1 | LayerFile | Vgg16_conv11 |
| 2 | dataFlowFile | NLR.m NVDLA.m |
| 3 | vectorWidth | 64 |
| 4 | NoCBandwidth | 128 |
| 5 | multicastSupported | True(1) |
| 6 | numAverageHopsinNoC | 4 |
| 7 | numPEs | 32 |

Table 2 Simulation Results For NLR and NVDLA

| Data Flow | NLR | NVDLA |
|---|---------------|-------------|
| Buffer analysis | | |
| L1 Buffer Requirement (Byte) | 18.00 | 66.00 |
| L2 Buffer Requirement (KB) | 1.12 | 4.12 |
| L1RdSum | 7,225,344 | 451,584 |
| L1WrSum | 7,225,344 | 451,584 |
| L2RdSum | 462,422,016 | 28,901,376 |
| L2WrSum | 462,422,016 | 28,901,376 |
| L1 weight reuse | 1 | 16 |
| L1 input reuse | 4 | 16 |
| L2 weight reuse | 448 | 190.26 |
| L2 input reuse | 2633 | 4473 |
| NoC analysis | | |
| L1 to L2 NoC BW | 128 | 32 |
| L2 to L1 NoC BW | 160 | 1024 |
| Performance analysis | | |
| L1 to L2 Sum | 56 | 32 |
| L1 to L2 Delay | 4.43 | 4.25 |
| L2 to L1 Delay | 0 | 0 |
| Roofline Throughput (GFLOPS with 1 GHZ clock) | 896 | 128 |
| Compute Runtime | 169 | 421 |
| Total Runtime (cycles) | 1,428,553,728 | 384,072,192 |

Conclusion

Artificial intelligence, machine learning and deep learning are growing trends affecting our lives in almost all aspects of human's life. These technologies make our life easier by assigning routine tasks of human resources to the machines that are much more accurate and fast. Therefore, any effort for optimizing performance, speed, and accuracy of these technologies is valuable. In this research, we focused on performance improvements of the hardware that are used for deep learning purposes named deep learning accelerators. Investigating recent researches conducted on these hardware accelerators shows that they can optimize costs, energy consumption, run time about 8–459% based on MAERI's investigation by minimizing memory bandwidth and maximizing computation reuse. Utilizing an open source tool named MAESTRO, we compared buffer, NoC



and performance parameters of NLR and NVDLA data flows. Results showed higher computation reuse for both L1 and L2 of the NVDLA data flow which is designed and

optimized for deep learning purposes and studied as deep learning accelerator in this study. The results showed that the customized hardware accelerator for deep learning (NVDLA) had much shorter total runtime in comparison with NLR.

Abbreviations

MAERI: Multiply-accumulate engine with reconfigurable interconnects; NLR: No local reuse; NVDLA: NVIDIA deep learning accelerator; MAESTRO: Modeling accelerator efficiency via spatio-temporal resource occupancy; ReLU: Rectified linear unit; DLA: Deep learning accelerator; NN: Neural network; CNN: Convolutional neural network; DNN: Deep neural network; RS: Row stationary; ASIC: Application-specific integrated circuits; ART: Augmented reduction tree; NoC: Network on chip; L1RdSum: L1 read sum; L1WrSUM: L1 write sum; L2RdSum: L2 read sum; L2WrSUM: L2 write sum.

Acknowledgements

Not applicable.

Authors' contributions

Investigating deep learning accelerators functionality. Analyzing a deep learning accelerator's architecture. Performance measurement of NVIDIA deep learning accelerator as a case study. Higher computation reuse and lower total runtime for the studied deep learning accelerator in comparison with non-optimized architecture.

Funding

Not applicable.

Availability of data and materials

Available.

Competing interests

Evaluating a deep learning accelerator's performance.

Received: 13 May 2020 Accepted: 1 November 2020

Published online: 12 November 2020

References

- Jürgen S. Deep learning in neural networks: an overview. *Neural Netw.* 2015;61:85–117.
- Müller B, Reinhardt J, Strickland MT. *Neural networks: an introduction*. Berlin: Springer; 2012. p. 14–5.
- Yann L, Yoshua B, Geoffrey H. Deep learning. *Nature.* 2015;521:436–44.
- Li D, Dong Y. Deep learning: methods and applications. *Found Trends Signal Process.* 2014;7:3–4.
- Jianqing F, Cong M, Yiqiao Z. A selective overview of deep learning. *arXiv:1904.05526[stat.ML]*. 2019.
- Wenyan L, et al. FlexFlow: a flexible dataflow accelerator architecture for convolutional neural networks. In: *IEEE international symposium on high performance computer architecture*. 2017.
- Katrik H, et al. Morph: flexible acceleration for 3D CNN-based Video Understanding. In *51st annual IEEE/ACM international symposium on microarchitecture (MICRO)*. 2018.
- Tianshi R, et al. DianNao: a small-footprint high-throughput accelerator for ubiquitous machine-learning. In: *ACM SIGARCH Computer Architecture News*; 2014.
- Zidong D, et al. ShiDianNao: Shifting Vision Processing Closer to the Sensor. In: *ACM/IEEE 42nd annual international symposium on computer architecture (ISCA)*. 2015.
- Chen Y-H, Emer J, Sze V. Eyeriss: A spatial architecture for energy-efficient dataflow for convolutional neural networks. *ACM SIGARCH Computer Architect News.* 2016;44:367–79.
- Michael P, et al. Buffets: an efficient and composable storage idiom for explicit decoupled data orchestration. In: *ASPLOS '19: proceedings of the twenty-fourth international conference on architectural support for programming languages and operating systems*. 2019, P 137.
- Xia X, Marcin W, Fan X, Damasevicius R, Li Y. Multi-sink distributed power control algorithm for Cyber-physical-systems in coal mine tunnels. *Comput Netw.* 2019;161:210–9.
- Song H, Li W, Shen P, Vasilakos A. Gradient-driven parking navigation using a continuous information potential field based on wireless sensor network. *Inf Sci.* 2017;408(2):100–14.
- Bin Z, Dawid P, Marcin W. A regional adaptive variational PDE model for computed tomography image reconstruction. *Pattern Recogn.* 2019;92:64–81. <https://doi.org/10.1016/j.patcog.2019.03.009>.
- Hyoukjun K, Ananda S, Tushar K. MAERI: enabling flexible dataflow mapping over DNN accelerators via reconfigurable interconnects. In: *ASPLOS '18, Proceedings of the twenty-third international conference on architectural support for programming languages and operating systems*.
- Uktu A, Shane O, Davor C, Andrew C. L, and Gordon R. C. An OpenCL deep learning accelerator on Arria 10. In: *FPGA '17, Proceedings of the 2017 ACM/SIGDA international symposium on field programmable gate arrays*, pp. 55–64.
- Hyoukjun K, Michael P, Tushar K. MAESTRO: an open-source infrastructure for modeling dataflows within deep learning accelerators. *arXiv:1805.02566*. 2018.
- Karen S, Andrew Z. Very deep convolutional network for large-scale image recognition. *arXiv:1409.1556*. 2015
- NVDLA Deep Learning Accelerator, <https://nvdla.org>. 2017.
- George SE. *The anatomy and physiology of the human stress response. A clinical guide to the treatment of the human stress responses*. Berlin: Springer, pp 19–56.

21. Christian S, Alexander T, Dumitru E. Deep neural networks for object detection. *Advances in neural information processing systems* 26, NIPS 2013.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Submit your manuscript to a SpringerOpen[®] journal and benefit from:

- ▶ Convenient online submission
- ▶ Rigorous peer review
- ▶ Open access: articles freely available online
- ▶ High visibility within the field
- ▶ Retaining the copyright to your article

Submit your next manuscript at ▶ [springeropen.com](https://www.springeropen.com)
