

NTX: An Energy-efficient Streaming Accelerator for Floating-point Generalized Reduction Workloads in 22 nm FD-SOI

Fabian Schuiki
IIS, ETH Zürich
 Zürich, Switzerland
 fschuiki@iis.ee.ethz.ch

Michael Schaffner
IIS, ETH Zürich
 Zürich, Switzerland
 schaffner@iis.ee.ethz.ch

Luca Benini
IIS, ETH Zürich
 Zürich, Switzerland
DEI, University of Bologna
 Bologna, Italy
 lbenini@iis.ee.ethz.ch

Abstract—Specialized coprocessors for Multiply-Accumulate (MAC) intensive workloads such as Deep Learning are becoming widespread in SoC platforms, from GPUs to mobile SoCs. In this paper we revisit NTX (an efficient accelerator developed for training Deep Neural Networks at scale) as a generalized MAC and reduction streaming engine. The architecture consists of a set of 32 bit floating-point streaming co-processors that are loosely coupled to a RISC-V core in charge of orchestrating data movement and computation. Post-layout results of a recent silicon implementation in 22 nm FD-SOI technology show the accelerator’s capability to deliver up to 20 Gflop/s at 1.25 GHz and 168 mW. Based on these results we show that a version of NTX scaled down to 14 nm can achieve a 3× energy efficiency improvement over contemporary GPUs at 10.4× less silicon area, and a compute performance of 1.4 Tflop/s for training large state-of-the-art networks with full floating-point precision. An extended evaluation of MAC-intensive kernels shows that NTX can consistently achieve up to 87% of its peak performance across general reduction workloads beyond machine learning. Its modular architecture enables deployment at different scales ranging from high-performance GPU-class to low-power embedded scenarios.

Index Terms—Processor Architecture, Accelerator, Deep Learning, VLSI, Linear Algebra

I. INTRODUCTION

Specialized accelerators for parallel MAC intensive workloads are becoming essential platforms ranging from mobile SoCs to high-performance GPUs, due to the widespread diffusion of Deep Neural Networks (DNNs) into various classification and recognition tasks [1], [2]. Yet most such accelerators are narrowly specialized for inference only [2]–[5]. Acceleration of training and more general MAC-intensive workloads has only received moderate attention thus far [6]–[8] and is still mainly carried out using GPUs [9].

This work has been supported by Microsoft Research under the project “Enabling Practical, Efficient and Large-Scale Computation Near Data to Improve the Performance and Efficiency of Data Center and Consumer Systems” with MRL contract number 2017-044.

This work has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 732631, project “OPRECOMP”.

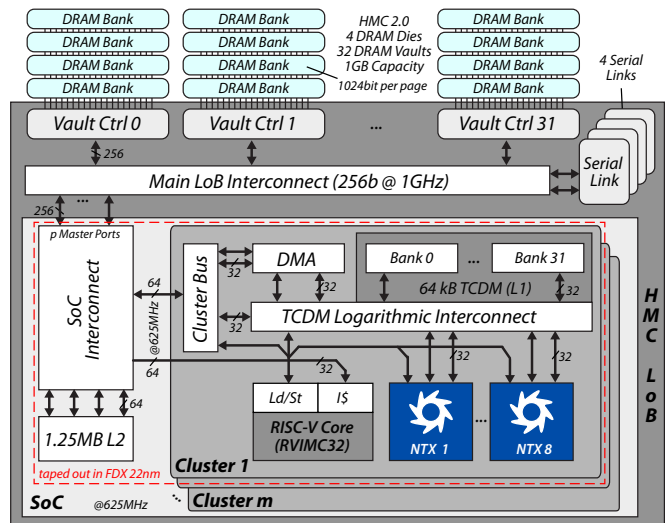


Figure 1. Top-level block diagram of one HMC enhanced with m processing clusters. The LoB contains the vault controllers, main interconnect, and the four serial links that lead off-cube. The proposed processing clusters attach directly to the main interconnect and gain full access to the HMC’s memory space and the serial links. Each cluster consists of a DMA unit, a TCDM, and one or more RISC-V processor cores augmented with NTX streaming co-processors.

Training of DNNs and general MAC-intensive workloads incur additional complexity due to additional data dependencies and higher accuracy requirements when compared to inference. At the same time the parameter memory footprint of typical DNNs has grown rapidly from a few MB to several tens or hundreds of MB over the last years [10], [11]. The corresponding training data sets are bound to grow as well, since larger networks require more training data in order to reach good generalization performance. These observations suggest that Processor-in-Memory (PiM) architectures that leverage lower access latencies and efficient data movement mechanisms should be closer investigated in the context of DNN training.

In this work we revisit NTX, a recently proposed PiM

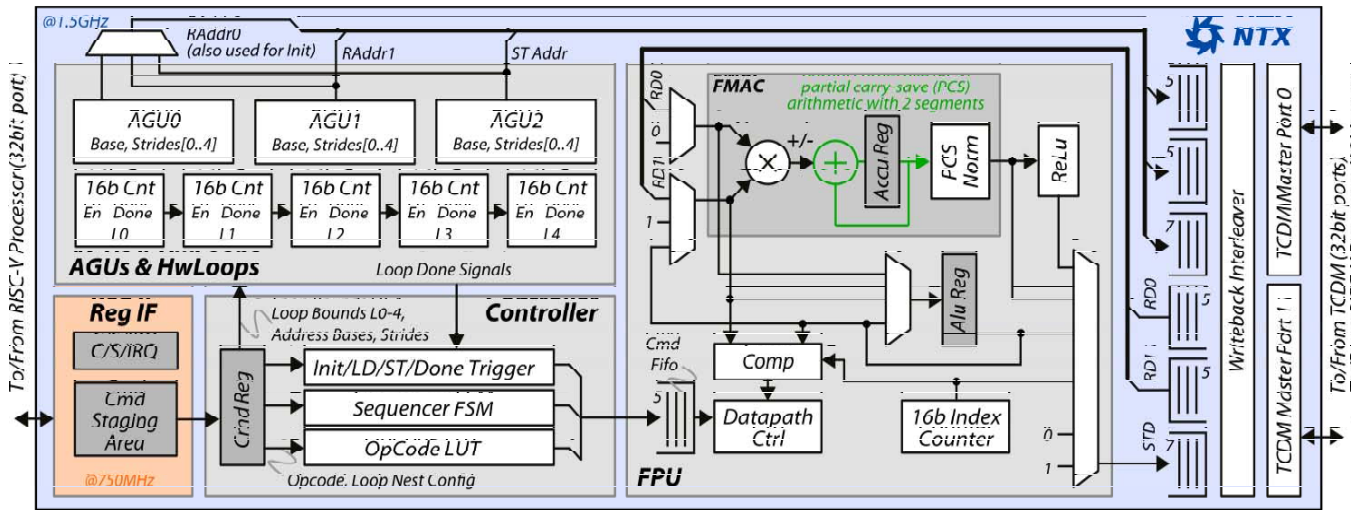


Figure 2. Block diagram of the NTX accelerator. It contains 5 hardware loops; 3 address generator units; a double-buffered command staging area in the register interface; a main controller; and a FPU with a comparator, index counter (for argmax calculations) and a fast FMAC unit. The employed depths for all FIFOs are indicated and have been determined in simulations for a TCDM read-latency of 1 cycle.

accelerator for deep learning applications [12]¹. In particular we provide results of a recently taped out variant that consists of a single processing cluster with 1 RISC-V core and 8 NTX accelerators, capable of delivering 20 Gflop/s at 1.25 GHz in 22 nm FD-SOI technology. Based on post-layout simulation results, we show that a HMC extended with NTX can achieve a compute performance of 1.4 Tflop/s at an efficiency of 63.5 Gflop/s W while training large state-of-the-art networks at full 32 bit floating-point precision. In addition, we investigate NTX’s applicability to and efficiency operating on more general kernels beyond machine learning, such as basic linear algebra subprograms and stencil codes. These are important methods in many domains, e.g. for solving least squares or finite difference problems in image and signal processing applications such as simultaneous localization and [13], optical flow estimation and inpainting [14], [15], and weather and seismic modeling [16], [17]. The contributions of this paper are:

- post-layout timing and power results based on a design recently taped out in 22 nm technology; and
- an extended performance analysis of general reduction kernels beyond deep learning.

We find that NTX is a competitive solution for general reduction applications. Its modular architecture enables deployment at scales other than training in a data center, for example as an accelerator for low power data analytics on edge devices.

II. ARCHITECTURE

In the following we provide a short overview of the NTX architecture outlined in more details in [12]. The Logic Base

¹This invited paper summarizes the *IEEE Transactions on Computers* article [12] (in press), and extends it with results from a recent tape out in 22 nm and additional kernel benchmarks beyond the realm of deep learning.

(LoB) of a HMC offers a unique opportunity to introduce a PiM as depicted in Figure 1. The memory dies are subdivided into vertically connected vaults, with corresponding memory controllers on the LoB connected to the serial links via the main interconnect. Our architecture consists of multiple processing clusters attached to this interconnect which thus gain full access to the entire HMC memory space, and sibling HMCs attached via the serial links.

A. Processing Cluster

We combine a small 32 bit RISC-V processor core (RV32IMC) [18] with multiple NTX co-processors. Both operate on shared 64 kB TCDM (reduced from 128 kB in [12]). The memory is divided into 32 banks that are connected to the processors via an interconnect offering single-cycle access latency. An additional DMA engine allows the transfer of two-dimensional data planes between the TCDM and the HMC’s memory space. The RISC-V processors perform address calculation and control data movement via the DMA. Actual computation is performed on the data in the TCDM by the NTX co-processors which we describe in the next section. An additional explicitly managed 1.25 MB of memory outside the clusters holds the RISC-V binary executed by the processors and may be used by the program to cache frequently used data and shared variables.

B. Network Training Accelerator (NTX)

The computations involved in DNNs training and many stencil codes are highly regular and can be broken down into a collection of reduction operations. The NTX co-processor is capable of performing thousands of FMAC cycles directly on the TCDM without any RISC-V core intervention or explicit load or store instructions. The architecture of NTX is depicted in Figure 2. It consists of four main blocks: the FPU containing

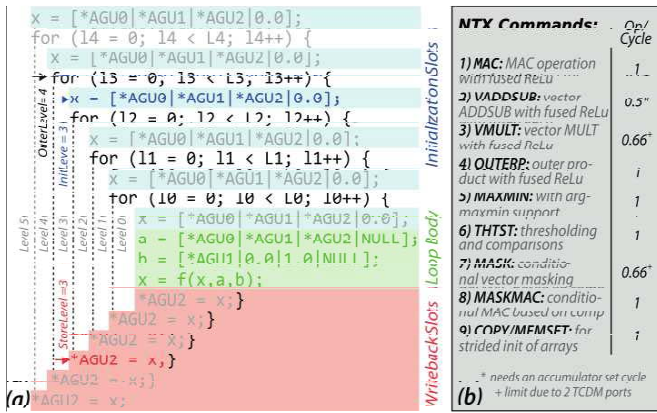


Figure 3. The structure of nested loops that can be directly offloaded to NTX (a), and an overview of the supported commands and their throughput (b). [...] indicates a choice of options; *AGU0 indicates memory access at address AGU0.

the main data path, the register interface for command offloading, the controller that decodes the commands and issues micro-instructions to the FPU, and the address generators and hardware loops.

C. FMAC and FPU

The floating-point unit (FPU) in NTX can perform fast FMAC operations with single-cycle throughput. It is based on a Partial Carry-Save (PCS) accumulator which aggregates the 48 bit multiplication result at full fixed-point precision (≈ 300 bit). After accumulation the partial sums are reduced in multiple pipelined segments. The employed format has been aligned with IEEE 754 32 bit floats. The wide accumulator and deferred rounding allows NTX to achieve higher precision than conventional FPUs. Analysis on a DNN convolution layer has shown NTX's Root Mean Squared Error (RMSE) to be $1.7\times$ lower than that of a 32 bit FPU.

The FMAC unit allows NTX to handle common matrix operations such as inner/outer products and vector additions/multiplications. An additional comparator, index counter, and ALU register enable various additional commands such as finding minima/maxima, ReLU, thresholding and masking, and memcopy/memset [12].

D. Hardware Loops and Address Generation

At the core of address generation in NTX are the five Hardware Loops (HWLs). Each loop maintains a 16 bit counter that has a programmable maximum count and can be enabled or disabled. The counters form a cascade to implement nested loops such that a loop wrapping from its maximum count to zero will increment the next higher loop. Three Address Generation Units (AGUs) allow NTX to keep track of three pointers into memory. Each unit consists of a 32 bit register holding the address and an adder. The address is incremented every cycle by one of five programmable step sizes chosen based on the outermost loop enabled in that cycle.

Figure 3a shows the pseudo code structure of nested loops that NTX can natively perform. The number of loops (*outer*

level), position of the accumulator initialization (*init level*), and position of the accumulator write back (*store level*) are fully programmable. The AGUs provide addresses for the memory reads and writes depicted, thus removing the need for the majority of explicit load/store instructions. The operation performed by the FPU always occurs in the innermost loop and can be set to one of the commands listed in Figure 3b.

E. Offloading

Each NTX has a set of configuration registers that are mapped into the memory space of the associated RISC-V core. As such the program can directly access and modify these registers, specifying the base address, strides, loop iterations, and command to be executed. Writing to the command register causes the current configuration to be copied into an internal buffer and executed, allowing the CPU to prepare the next command in parallel. All NTX attached to a core are aliased to a broadcast address, allowing efficient setting of common configuration values. This offloading scheme has proven to be very lean and efficient [12], allowing each NTX to run independently for thousands of cycles during which the RISC-V core can perform other tasks such as data movement.

We subdivide kernels to be executed into tiles. The DMA engine is used to copy input data into and results out of the TCDM in a double buffering scheme, allowing the NTX coprocessors to operate on one buffer while the DMA operates on another. This allows us to decouple and overlap computation and data movement, thus hiding memory latency and fully utilizing the available memory bandwidth.

III. EVALUATION AND RESULTS

A. Silicon Results

We have implemented and taped out an NTX cluster in GLOBALFOUNDRIES' 22FDX, a 22 nm FD-SOI technology. Table I summarizes the figures of merit of our implementation. Post-layout timings were extracted from Cadence Innovus and used in a back-annotated gate-level simulation to obtain a trace of the cluster performing computation and DMA operation. This trace was then used in Innovus alongside the design to estimate the power consumption. The cluster consists of one RISCY [18] processor core and eight NTX coprocessors which operate on 64 kB of TCDM. A 2 kB instruction cache with linear prefetching is located between the processor and the memory interface. The TCDM and NTX operate at 1.25 GHz in the worst case (0.72 V, 125 °C/−40 °C, SSG), while the RISC-V processor and remaining cluster runs at half the speed, 625 GHz. In this corner the cluster occupies 0.51 mm² at 59% placement density while achieving a compute performance of 20 Gflop/s and a memory bandwidth of 5 GB/s. Assuming typical silicon (0.8 V, 25 °C, TT) the cluster consumes 186 mW of power while performing a 3×3 convolution, which yields an energy efficiency of 108 Gflop/s W or conversely 9.3 pJ/flop.

Table I
FIGURES OF MERIT OF ONE NTX CLUSTER IMPLEMENTED IN 22FDX.

Processors:	1 RISC-V 8 NTX	Memory:	64 kB TCDM 2 kB ICACHE
Frequency:	1.25 GHz NTX 625 MHz Cluster	Area:	0.51 mm ² 59% density
Peak Perf.:	20 Gflop/s 5 GB/s	Power:	186 mW
		Efficiency:	108 Gflop/s W 9.3 pJ/flop




Figure 4. Floorplan of the NTX cluster as implemented in 22FDX. Highlighted are the main area contributors: TCDM, logarithmic interconnect, RISC-V processor and peripherals, NTX coprocessors, and the instruction cache.

B. Evaluated Kernels

We estimate the execution time of a kernel based on the model presented in [12]. The data is assumed to initially reside outside the cluster, e.g. in a DRAM attached to the AXI port.

1) *Basic Linear Algebra Subprograms*: The AXPY ($y = a \cdot x + y$), matrix-vector product GEMV, and matrix-matrix product GEMM are taken from the BLAS 1, 2, and 3 set of kernels, respectively. For AXPY and GEMV the input data is split into tiles that fit into the cluster’s TCDM memory, which are then processed tile-by-tile. Data reuse, which manifests itself as increased operational intensity, is limited by the kernel itself as well as the size of the largest tile that fits into the TCDM. For GEMM we use a block matrix multiplication to subdivide the input matrices.

2) *Convolutions*: We evaluate 3×3 , 5×5 , and 7×7 convolutions as they commonly appear in DNNs [10]. Reuse factors per image pixel are 9, 25, and 49, respectively. Larger convolution kernels exhibit higher operational intensity since input image pixels are reused for more operations, thus allowing NTX to operate even further in the compute-bound regime.

3) *Stencils*: Stencil codes are common in High Performance Computing (HPC). We evaluate the Discrete Laplace Operator [19] in one, two, and three dimensions with three, five, and seven coefficients, respectively. Its star shaped access pattern allows it to be computed efficiently on NTX by decomposing the kernel into its separate dimensions. Furthermore, we also consider the diffusion kernel presented as an example in [16] which has 13 coefficients and can be decomposed into three separate NTX instructions with nine, two, and

two coefficients each. Together with convolutions these are a representative sample of the common five- and nine-point stencil shapes and beyond.

C. Roofline Model

The roofline model of one NTX cluster is depicted in Figure 5. The eight NTX co-processors at 1.25 GHz achieve 20 Gflop/s of peak performance, while the AXI memory port at 625 MHz can carry 5 GB/s of peak traffic. We estimate the performance of different kernels by extrapolation of a gate-level simulation of the 3×3 convolution. For the three BLAS kernels AXPY, GEMV, and GEMM the NTX cluster achieves close to maximum performance with a sufficiently large problem size. AXPY and GEMV are memory bound in all configurations, while GEMM quickly becomes compute bound as operational intensity increases due to better amortization of constant setup and write back overheads. The investigated convolution kernels are all compute bound and achieve close to maximum performance. The three Discrete Laplace Operator [19] and the diffusion stencil [16] are all memory bound, yet achieve close to maximum bandwidth utilization since their regular structure is highly amenable to execution on NTX.

We observe in simulations that the practically achievable compute performance is limited by the probability of a banking conflict in the TCDM interconnect, which causes an NTX stall. This probability was measured to be around 13%, which puts the maximum performance achievable in practice at around 17.4 Gflop/s. This also limits the expected maximum memory bandwidth of the system for memory-bound kernels to around 4.35 GB/s.

The memory bound of the roofline plot is dictated by the width of the AXI port of the cluster, which is a design parameter. It was set to 64 bit to accommodate the bandwidth requirements of DNN training and to facilitate system integration also with lower-end devices. This parameter could be increased to 128 or 256 bit, raising the bandwidth limit to 10 GB/s and 20 GB/s, respectively. This would allow the cluster to sustain very high utilization also for operational intensities down to 2 flop/B and 1 flop/B, respectively.

D. Neural Network Training Efficiency

To compare DNN training performance against other accelerators we reproduce Table II and Figure 6 from [12] with updated numbers based on our implementation in 22FDX, which provides a more accurate estimate of the performance achievable with the resulting hardware. Among the custom accelerators, DaDianNao has an efficiency of 65.8 Gop/s W with fixed-point arithmetic, which is similar to the computationally equivalent NTX 128. ScaleDeep has an efficiency of 100.8 Gflop/s W which is $1.3 \times$ higher than NTX 512, the largest configuration considered by us. Furthermore our architecture can achieve significantly higher energy efficiency than a GPU at a comparable technology node (see Figure 6). Considering the largest NTX configurations that do not require additional LiMs, we achieve an efficiency increase of $2.5 \times$

Table II

COMPARISON BETWEEN DIFFERENT CONFIGURATIONS OF THE ARCHITECTURE PROPOSED IN THIS WORK, RELATED CUSTOM ACCELERATORS, AND GPUS. THE ENERGY EFFICIENCIES REPORTED ARE WITH RESPECT TO TRAINING DIFFERENT DNNs. REFER TO [12] FOR A MORE DETAILED TREATMENT.

Platform	Characteristics							Energy Efficiency [Gop/s W]						
	Logic [nm]	DRAM [nm]	Area [mm ²]	LiM	Freq. [GHz]	Peak Top/s	Arithmetic	AlexNet [20]	GoogLeNet [10]	Incep. v3 [21]	ResNet 34 [11]	ResNet 50 [11]	ResNet 152 [11]	Geom. Mean
This Work														
NTX (16×)	22	50	4.8	0	2.50	0.640	(a)	19.8	23.7	24.3	21.7	21.4	23.6	22.5
NTX (32×)	22	50	9.6	0	1.90	0.973	(a)	25.8	30.9	31.6	28.2	27.9	30.8	29.3
NTX (64×)	22	50	19.3	1	1.43	1.466	(a)	32.3	38.8	39.7	35.4	35.0	38.6	36.7
NTX (16×)	14	30	1.9	0	3.50	0.896	(a)	31.6	37.9	38.8	34.6	34.2	37.7	35.9
NTX (32×)	14	30	3.9	0	2.66	1.362	(a)	41.8	50.1	51.3	45.8	45.2	49.9	47.5
NTX (64×)	14	30	7.7	0	1.88	1.920	(a)	53.2	63.8	65.3	58.3	57.6	63.5	60.4
NTX (128×)	14	30	15.4	1	0.94	1.920	(a)	62.1	74.6	76.2	68.1	67.2	74.2	70.6
NTX (256×)	14	30	30.8	2	0.47	1.920	(a)	66.9	80.3	82.1	73.3	72.4	79.8	76.0
NTX (512×)	14	30	61.6	3	0.23	1.920	(a)	69.3	83.2	85.0	75.9	75.0	82.7	78.7
Custom Accelerators														
NS (16×) [4]	28	50	9.3	—	1.0	0.256	(a)	10.2	15.1	14.6	13.1	12.9	14.2	13.0
DaDianNao [7]	28	28	67.7	—	0.6	2.09	(b)	—	—	—	—	—	—	65.8
ScaleDeep [6]	14	—	—	—	0.6	680	(c)	87.7	83.0	—	139.2	—	—	100.8
GPUs														
Tesla K80	28	40	561	—	0.59	8.74	(a)	—	4.5	3.5	—	3.7	8.8	4.7
Tesla M40	28	30	601	—	1.11	7.00	(a)	—	11.3	—	—	—	—	11.3
Titan X	28	30	601	—	1.08	7.00	(a)	12.8	9.9	—	17.6	8.5	12.2	11.8
Tesla P100	16	21	610	—	1.3	10.6	(a)	—	19.8	19.5	—	18.6	24.18	20.4
GTX 1080 Ti	16	20	471	—	1.58	11.3	(a)	20.1	16.6	—	27.6	13.4	19.56	18.9

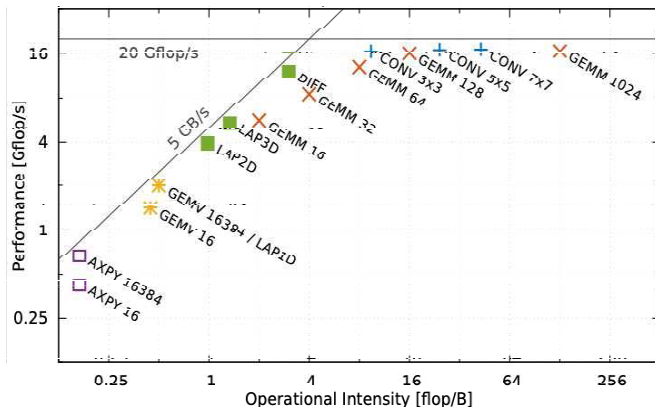


Figure 5. Roofline model of NTX for different kernels. For AXPY, GEMV, and GEMM the vector length and matrix side length are indicated. For CONV the size of the convolution kernel is indicated. LAP are discrete Laplace operators in 1D, 2D, and 3D. DIFF is the diffusion stencil used as an example in [16]. Note that LAP1D coincides with GEMV 16384.

from 11.8 Gop/s W to 29.5 Gop/s W in 22 nm, and an increase of 3× from 20.4 Gop/s W to 63.5 Gop/s W in 14 nm.

We also reproduce Figure 7 from [12] with updated numbers for 22FDX to compare the Gop/s of compute performance per deployed amount of silicon for NTX and GPUs. Our solution requires 10.4× less area to achieve the same compute performance as a GPU.

IV. RELATED WORK

Acceleration of DNNs is a well researched field with a rich literature. Goodfellow, et al. [1] provide a good coverage of




Figure 6. Comparison of energy efficiency when training the networks listed in Table II (geometric mean), with GPUs, NS [4], and the largest NTX configurations that do not require additional LiMs. NTX 32 in 22 nm achieves a 2.5× increase, and NTX 64 in 14 nm a 3× increase in efficiency over GPUs in similar technology nodes.




Figure 7. Comparison of the Gop/s of compute performance per deployed area of silicon, for GPUs, NS [4], and the largest NTX configurations that do not require additional LiMs. NTX 32 in 22 nm achieves a 6.5× increase, and NTX 64 in 14 nm a 10.4× increase in area efficiency over GPUs in similar technology nodes.

the mathematical background of Deep Learning, and [2] offer an overview of techniques for efficient DNN inference and the challenges involved. Dedicated DNN accelerators have mainly focused on inference [3]–[5]. The increasing size of parameter and training data of state-of-the-art networks [10], [11] provide a compelling reason for PiM solutions. We observe that fewer architectures that support both inference and training have been proposed so far [6]–[8]. DaDianNao [7] is a multi-node system achieving an energy efficiency of around 350 Gop/s W for 16 bit fixed point arithmetic. ScaleDeep [6] is a multi-node system with heterogeneous chips assembled from memory-heavy and compute-heavy tiles that distribute the DNN state across several chips and nodes, achieving a very high energy efficiency around 332 Gflop/s W in a 14 nm technology. A more detailed treatment of related accelerators can be found in [12].

GPUs are commonly used for both inference and training, where recent implementations on the GTX 780 and GTX Titan reach 1650 Gflop/s at 250 W and 999 Gflop/s at 240 W, which corresponds to 6.6 and 4.2 Gflop/s W, respectively [4], [22]. Embedded GPUs like the Tegra K1 have lower absolute throughput, but reach a energy efficiencies of around 7 Gflop/s W [22]. Newer GPU generations such as Pascal offer High Bandwidth Memory (HBM) and 16 bit floating-point (FP) support allowing for higher peak throughput and efficiency of up to 10.6 Tflop/s and 20 Gflop/s W, respectively [23], [24].

In the HPC domain stencil codes and general linear algebra are crucial building blocks of many applications. The increasing complexity and data volume of state-of-the-art problems requires dedicated acceleration engines to keep power consumption manageable. Green Wave [17] for example focuses on solving 8th order Laplacian stencils for seismic modeling applications using a large array of dedicated streaming nodes, reaching 82.5 Gflop/s at 1.25 Gflop/s W. A GPU executing the same stencil reaches 145 Gflop/s at 0.33 Gflop/s W, a $1.7\times$ increased performance but at the cost of $3.5\times$ less energy efficiency [17]. We estimate NTX 16 to achieve 130 Gflop/s at 11 Gflop/s W on the same stencil code. This suggests that dedicated streaming-based accelerators for stencil codes and linear algebra are an attractive proposition to reduce the energy footprint of HPC applications.

V. CONCLUSION

We have presented an evaluation of the NTX floating point co-processor [12] based on a concrete implementation taped out in GLOBALFOUNDRIES' 22FDX technology. Power analysis based on post-layout simulation confirms the estimates in previous work. The hardware loop and FMAC capabilities of NTX apply well to kernels beyond DNNs such as stencil codes prevalent in HPC, while allowing NTX to achieve very high utilization of its peak performance, and suggests that the co-processor is capable of handling other well-structured problems to be investigated further.

REFERENCES

- [1] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016, <http://www.deeplearningbook.org>.
- [2] V. Sze, Y.-H. Chen, T.-J. Yang, and J. Emer, "Efficient processing of deep neural networks: A tutorial and survey," *arXiv:1703.09039*, 2017.
- [3] M. Gao, J. Pu, X. Yang, M. Horowitz, and C. Kozyrakis, "TETRIS: Scalable and Efficient Neural Network Acceleration with 3D Memory," in *ASPLOS*, 2017.
- [4] E. Azarkhish, D. Rossi, I. Loi, and L. Benini, "Neurostream: Scalable and Energy Efficient Deep Learning with Smart Memory Cubes," *IEEE TPDS*, vol. PP, no. 99, 2017.
- [5] Z. Du, R. Fasthuber, T. Chen, P. Ienne, L. Li, T. Luo, X. Feng, Y. Chen, and O. Temam, "ShiDianNao: Shifting vision processing closer to the sensor," in *ACM SIGARCH CAN*, vol. 43, no. 3, 2015, pp. 92–104.
- [6] S. Venkataramani, A. Ranjan, S. Banerjee, D. Das, S. Avancha, A. Jagannathan, A. Durg, D. Nagaraj, B. Kaul, P. Dubey *et al.*, "ScaleDeep: A Scalable Compute Architecture for Learning and Evaluating Deep Networks," in *ISCA*, 2017, pp. 13–26.
- [7] T. Luo, S. Liu, L. Li, Y. Wang, S. Zhang, T. Chen, Z. Xu, O. Temam, and Y. Chen, "DaDianNao: a neural network supercomputer," *IEEE TC*, vol. 66, no. 1, pp. 73–88, 2017.
- [8] D. Kim, J. Kung, S. Chai, S. Yalamanchili, and S. Mukhopadhyay, "Neurocube: A programmable digital neuromorphic architecture with high-density 3D memory," in *ISCA*, 2016, pp. 380–392.
- [9] NVidia, "Artificial Intelligence Architecture | NVIDIA Volta," <https://www.nvidia.com/en-us/data-center/volta-gpu-architecture/>, 2017, acc.: July 2017.
- [10] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," in *CVPR*, 2015.
- [11] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *CVPR*, 2016.
- [12] F. Schuike, M. Schaffner, F. K. Gürkaynak, and L. Benini, "A scalable near-memory architecture for training deep neural networks on large in-memory datasets," *IEEE Transactions on Computers*, 2018, in press.
- [13] C. Kerl, J. Sturm, and D. Cremers, "Dense Visual SLAM for RGB-D Cameras," in *IEEE IROS*, Nov 2013.
- [14] D. Krishnan and R. Szeliski, "Multigrid and Multilevel Preconditioners for Computational Photography," in *ACM TOG*, vol. 30, no. 6, 2011.
- [15] I. Koutis, G. L. Miller, and D. Tolliver, "Combinatorial Preconditioners and Multilevel Solvers for Problems in Computer Vision and Image Processing," *JCVIU*, vol. 115, no. 12, pp. 1638–1646, 2011.
- [16] T. Gysi, T. Grosser, and T. Hoefler, "Modesto: Data-centric analytic optimization of complex stencil programs on heterogeneous architectures," in *Proceedings of the 29th ACM on International Conference on Supercomputing*. ACM, 2015, pp. 177–186.
- [17] J. Krueger, D. Donofrio, J. Shalf, M. Mohiyuddin, S. Williams, L. Oliker, and F.-J. Pfreund, "Hardware/software co-design for energy-efficient seismic modeling," in *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis*. ACM, 2011, p. 73.
- [18] M. Gautschi, P. D. Schiavone, A. Traber, I. Loi, A. Pullini, D. Rossi, E. Flamand, F. K. Gürkaynak, and L. Benini, "Near-Threshold RISC-V Core With DSP Extensions for Scalable IoT Endpoint Devices," *TVLSI*, 2017.
- [19] M. Wardetzky, S. Mathur, F. Kälberer, and E. Grinspun, "Discrete laplace operators: no free lunch," in *Symposium on Geometry processing*. Aire-la-Ville, Switzerland, 2007, pp. 33–37.
- [20] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *NIPS*, 2012, pp. 1097–1105.
- [21] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the inception architecture for computer vision," in *CVPR*, 2016, pp. 2818–2826.
- [22] L. Cavigelli, M. Magno, and L. Benini, "Accelerating real-time embedded scene labeling with convolutional networks," in *DAC*, 2015, pp. 1–6.
- [23] J. Johnson, "cnn-benchmarks," <https://github.com/jcjohnson/cnn-benchmarks>, acc.: September 2017, Commit hash 83d441f.
- [24] "Tensorflow Benchmarks," <https://www.tensorflow.org/performance/benchmarks>, August 2017, acc.: September 2017.