

In-Data Center Performance Analysis of a Tensor Processing UnitTM *

David Patterson and the Google TPU Team

davidpatterson@google.com

April 5, 2017

*4/5/17 Google published a blog on the TPU. A 17-page technical paper with same title will be on arXiv.org. (Paper will also appear at the *International Symposium on Computer Architecture* on June 26, 2017.)

A Golden Age in Microprocessor Design


- Stunning progress in microprocessor design 40 years $\approx 10^6$ x faster!
- Three architectural innovations (~ 1000 x)
 - Width: 8- \rightarrow 16- \rightarrow 32 - \rightarrow 64 bit (~ 8 x)
 - Instruction level parallelism:
 - 4-10 *clock cycles per instruction* to 4+ *instructions per clock cycle* (~ 10 -20x)
 - Multicore: 1 processor to 16 cores (~ 16 x)
- Clock rate: 3 to 4000 MHz (~ 1000 x thru technology & architecture)
- Made possible by IC technology:
 - **Moore's Law:** growth in transistor count (2X every 1.5 years)
 - **Dennard Scaling:** power/transistor shrinks at same rate as transistors are added (constant per mm^2 of silicon)

Changes Converge

- Technology
 - End of Dennard scaling: power becomes the key constraint
 - Slowdown (retirement) of Moore's Law: transistors cost
- Architectural
 - Limitation and inefficiencies in exploiting instruction level parallelism end the uniprocessor era in 2004
 - Amdahl's Law and its implications end "easy" multicore era
- Products
 - PC/Server \Rightarrow Client/Cloud

End of Growth of Performance?

40 years of Processor Performance



Based on SPECintCPU. Source: John Hennessy and David Patterson, Computer Architecture: A Quantitative Approach, 6/e. 2018

What's Left?

Since


- Transistors not getting much better
- Power budget not getting much higher
- Already switched from 1 inefficient processor/chip to N efficient processors/chip

Only path left is *Domain Specific Architectures*

- Just do a few tasks, but extremely well



What is Deep Learning?

- Loosely based on (what little) we know about the brain



The Artificial Neuron

$$y = F \left(\sum_i w_i x_i \right)$$

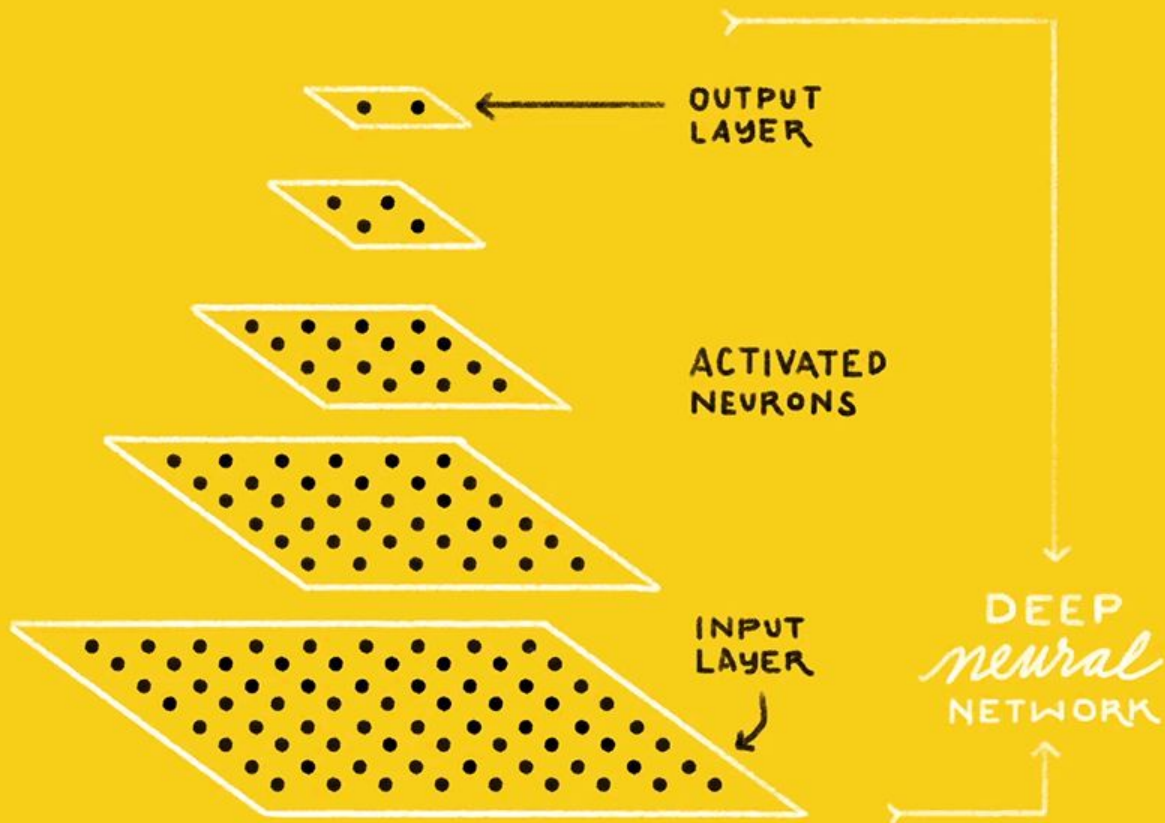


F : a nonlinear
differentiable
function

IS THIS A
CAT or DOG?



CAT DOG




Key NN Concepts for Architects

- *Training* or learning (development)
vs. *Inference* or prediction (production)
- *Batch size*
 - Problem: DNNs have millions of weights that take a long time to load from memory (DRAM)
 - Solution: Large batch \Rightarrow Amortize weight-fetch time by inferring (or training) many input examples at a time
- Floating-Point vs. Integer ("*Quantization*")
 - Training in Floating Point on GPUs popularized DNNs
 - Inferring in Integers faster, lower energy, smaller

- 2013: Prepare for success-disaster of new DNN apps
 - Scenario with users speaking to phones 3 minutes per day:
If only CPUs, need 2X-3X times whole fleet
 - Unlike some hardware targets, DNNs applicable to a wide range of problems, so can reuse for solutions in speech, vision, language, translation, search ranking, ...
- Custom hardware to reduce the TCO of DNN inference phase by 10X vs. GPUs
 - Must run existing apps developed for CPUs and GPUs
- A very short development cycle
 - Started project 2014, running in datacenter 15 months later:
Architecture invention, compiler invention, hardware design, build, test, deploy
- Google CEO Sundar Pichai reveals Tensor Processing Unit at Google I/O on May 18, 2016 as “10X performance/Watt”
cloudplatform.googleblog.com/2016/05/Google-supercharges-machine-learning-tasks-with-custom-chip.html

- TPU Card to replace a disk
- Up to 4 cards / server

TPU Card & Package



3 Types of NNs

1. Multilayer Perceptrons

- Each new layer applies nonlinear function F to weighted sum of all outputs from prior layer (“fully connected”) $x_n = F(Wx_{n-1})$

2. Convolutional Neural Network

- Like MLPs, but same weights used on nearby subsets of outputs from prior layer

3. Recurrent NN/“Long Short-Term Memory”

- Each new layer a NL function of weighted sums of past *state* and prior outputs; same weights used across time steps

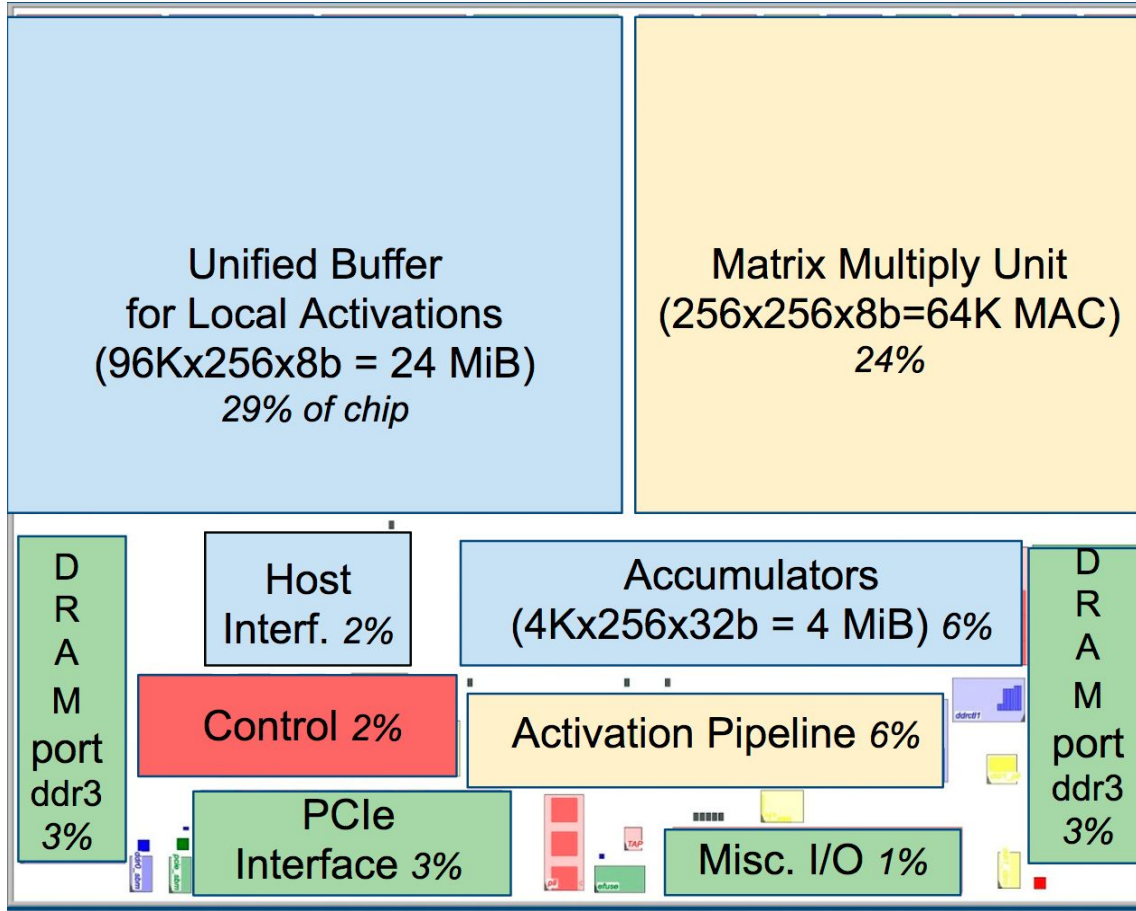
Inference Datacenter Workload (95%)

<i>Name</i>	<i>LOC</i>	<i>Layers</i>					<i>Nonlinear function</i>	<i>Weights</i>	<i>TPU Ops / Weight Byte</i>	<i>TPU Batch Size</i>	<i>% Deployed</i>
		<i>FC</i>	<i>Conv</i>	<i>Vector</i>	<i>Pool</i>	<i>Total</i>					
MLP0	0.1k	5				5	ReLU	20M	200	200	61%
MLP1	1k	4				4	ReLU	5M	168	168	
LSTM0	1k	24		34		58	sigmoid, tanh	52M	64	64	29%
LSTM1	1.5k	37		19		56	sigmoid, tanh	34M	96	96	
CNN0	1k		16			16	ReLU	8M	2888	8	5%
CNN1	1k	4	72		13	89	ReLU	100M	1750	32	

TPU Architecture and Implementation

- Add as accelerators to existing servers
 - So connect over I/O bus (“PCIe”)
 - TPU \approx matrix accelerator on I/O bus
- Host server sends it instructions like a Floating Point Unit
 - Unlike GPU that fetches and executes own instructions

TPU: a Neural Network Accelerator Chip




TPU Architecture, programmer's view

- 5 main (CISC) instructions
 - Read_Host_Memory
 - Write_Host_Memory
 - Read_Weights
 - MatrixMultiply/Convolve
 - Activate (ReLU, Sigmoid, Maxpool, LRN, ...)
- Average Clock cycles per instruction: >10
- 4-stage overlapped execution, 1 instruction type / stage
 - Execute other instructions while matrix multiplier busy
- Complexity in SW: No branches, in-order issue, SW controlled buffers, SW controlled pipeline synchronization

Systolic Execution in Matrix Array

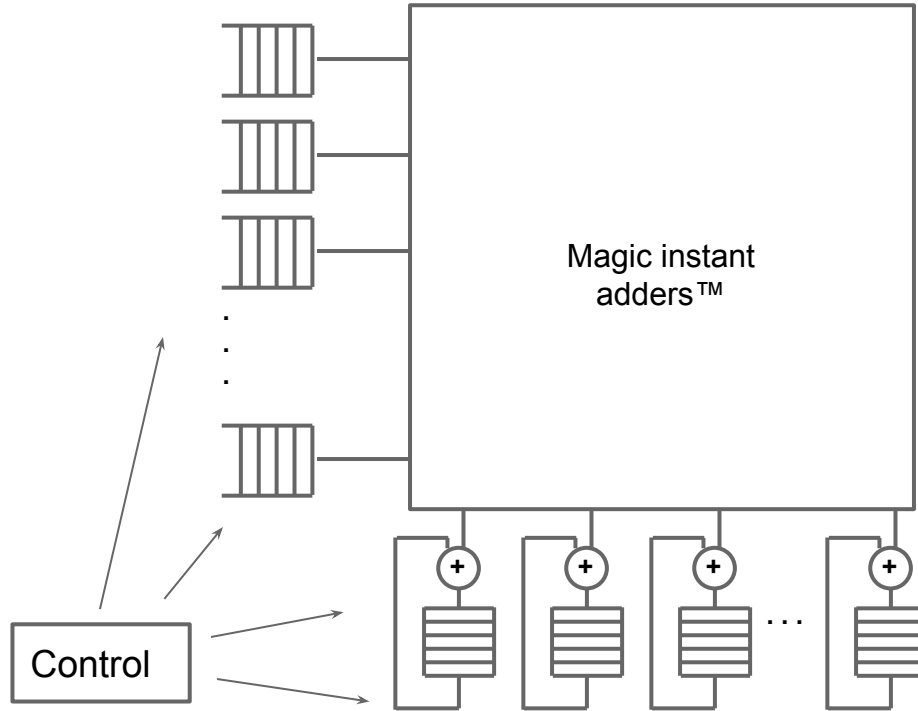
- Problem: energy/ time for repeated SRAM accesses of matrix multiply
- Solution: “Systolic Execution” to compute data on the fly in buffers by pipelining control and data
 - Relies on data from different directions arriving at cells in an array at regular intervals and being combined

Systolic Execution: Control and Data are pipelined



Can now ignore pipelining in matrix

Pretend each 256B input read at once, & they instantly update 1 location of each of 256 accumulator RAMs.



Relative Performance: 3 Contemporary Chips

<i>Processor</i>	<i>mm²</i>	<i>Clock MHz</i>	<i>TDP Watts</i>	<i>Idle Watts</i>	<i>Memory GB/sec</i>	<i>Peak TOPS/chip</i>	
						<i>8b int.</i>	<i>32b FP</i>
CPU: Haswell (18 core)	662	2300	145	41	51	2.6	1.3
GPU: Nvidia K80 (2 / card)	561	560	150	25	160	--	2.8
TPU	<331*	700	75	28	34	91.8	--

*TPU is less than half die size of the Intel Haswell processor

K80 and TPU in 28 nm process; Haswell fabbed in Intel 22 nm process

These chips and platforms chosen for comparison because widely deployed in Google data centers

GPUs and TPUs added to CPU server

Relative Performance: 3 Platforms

<i>Processor</i>	<i>Chips/ Server</i>	<i>DRAM</i>	<i>TDP Watts</i>	<i>Idle Watts</i>	<i>Observed Busy Watts in datacenter</i>
CPU: Haswell (18 cores)	2	256 GB	504	159	455
NVIDIA K80 (13 cores) (2 die per card; 4 cards per server)	8	256 GB (host) + 12GB x 8	1838	357	991
TPU (1 core) (1 die per card; 4 cards per server)	4	256GB (host) + 8GB x 4	861	290	384

These chips and platforms chosen for comparison because widely deployed in Google datacenters

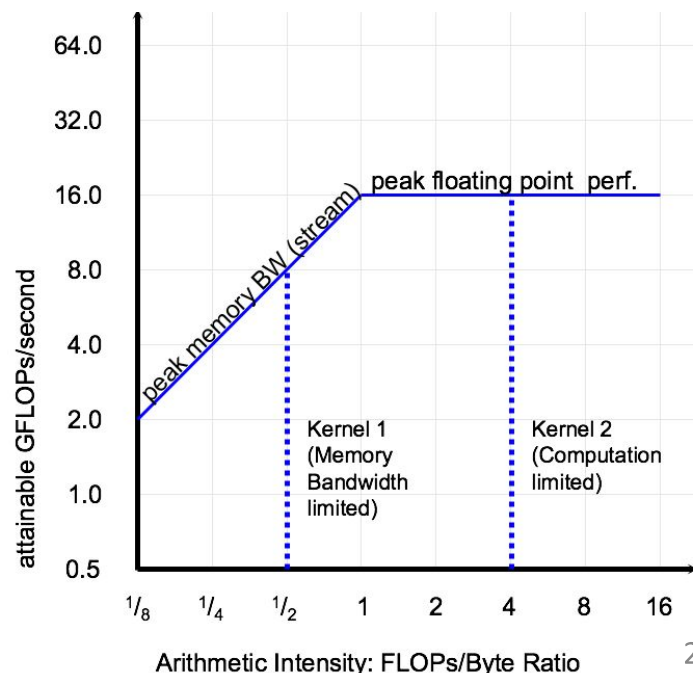
2 Limits to performance:

1. Peak Computation
2. Peak Memory Bandwidth
(For apps with large data that don't fit in cache)


Arithmetic Intensity (FLOP/byte or reuse) determines which limit
Weight-reuse = Arithmetic Intensity for DNN roofline

Roofline Visual Performance Model

$$\text{GFLOP/s} = \text{Min}(\text{Peak GFLOP/s}, \text{Peak GB/s} \times \text{AI})$$




TPU Die Roofline



Haswell (CPU) Die Roofline

Haswell Log-Log



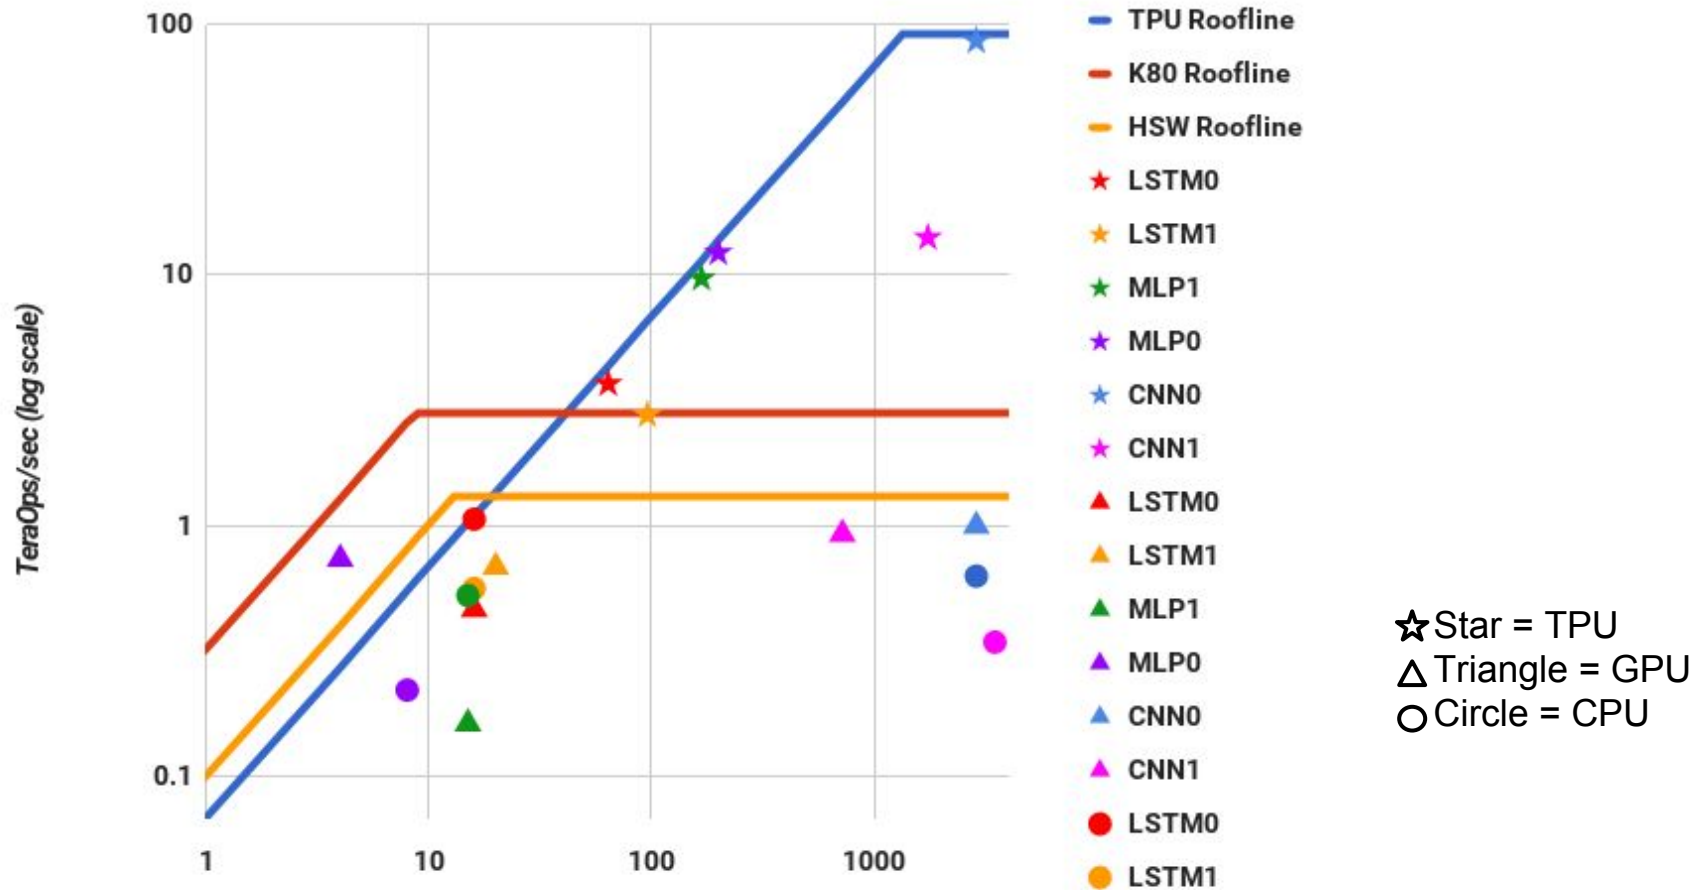
K80 (GPU) Die Roofline



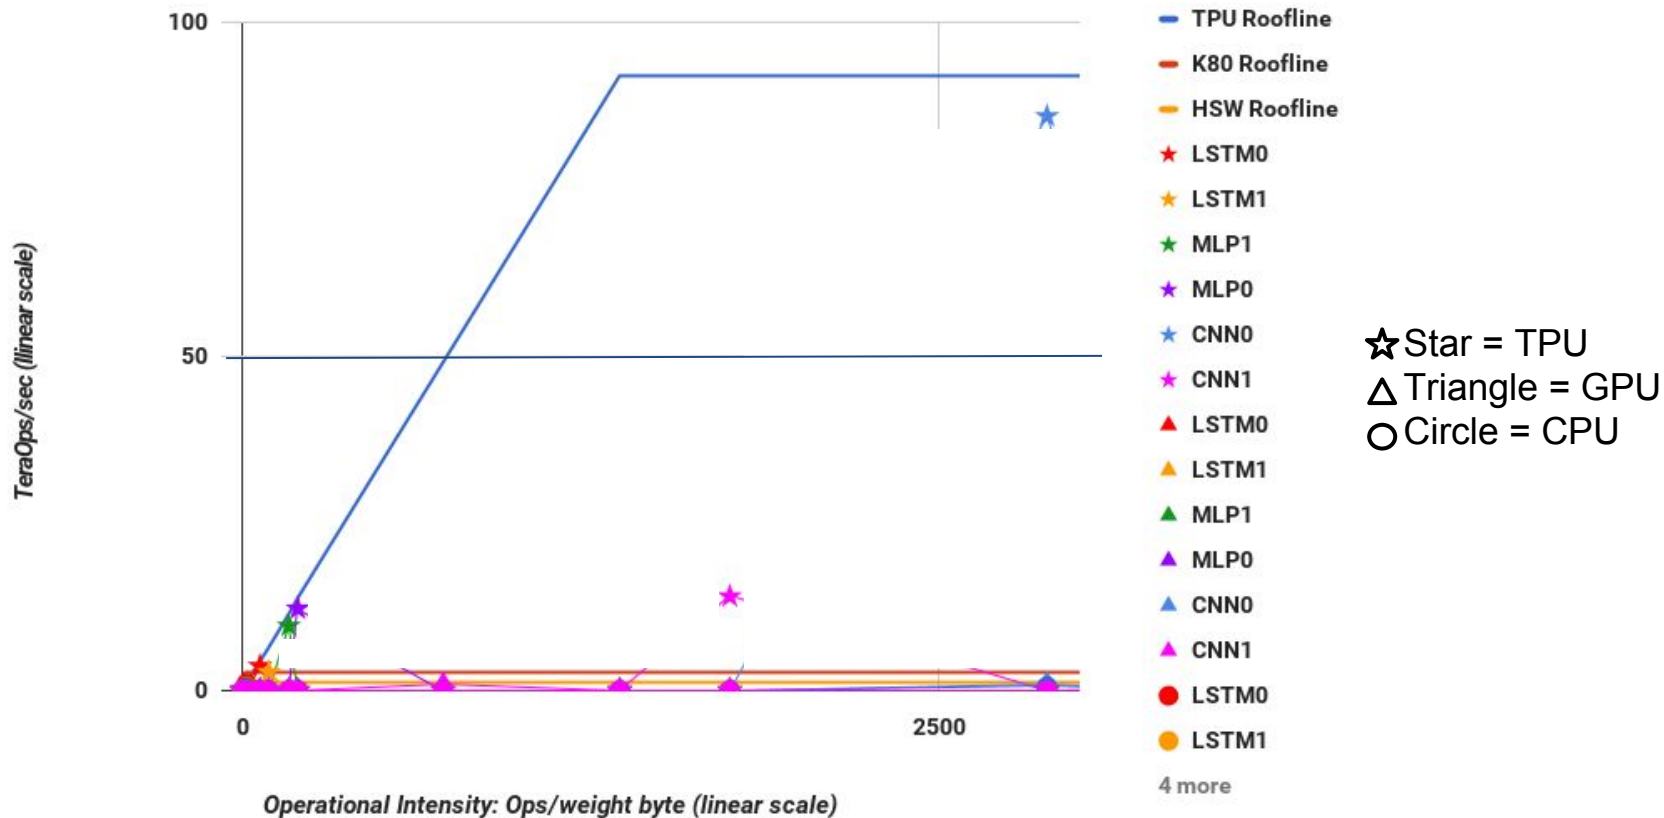
Why so far below Rooflines? (MLP0)

<i>Type</i>	<i>Batch</i>	<u><i>99th% Response</i></u>	<i>Inf/s (IPS)</i>	<i>% Max IPS</i>
CPU	16	7.2 ms	5,482	42%
CPU	64	21.3 ms	13,194	100%
GPU	16	6.7 ms	13,461	37%
GPU	64	8.3 ms	36,465	100%
TPU	200	7.0 ms	225,000	80%
TPU	250	10.0 ms	280,000	100%

Log Rooflines for CPU, GPU, TPU



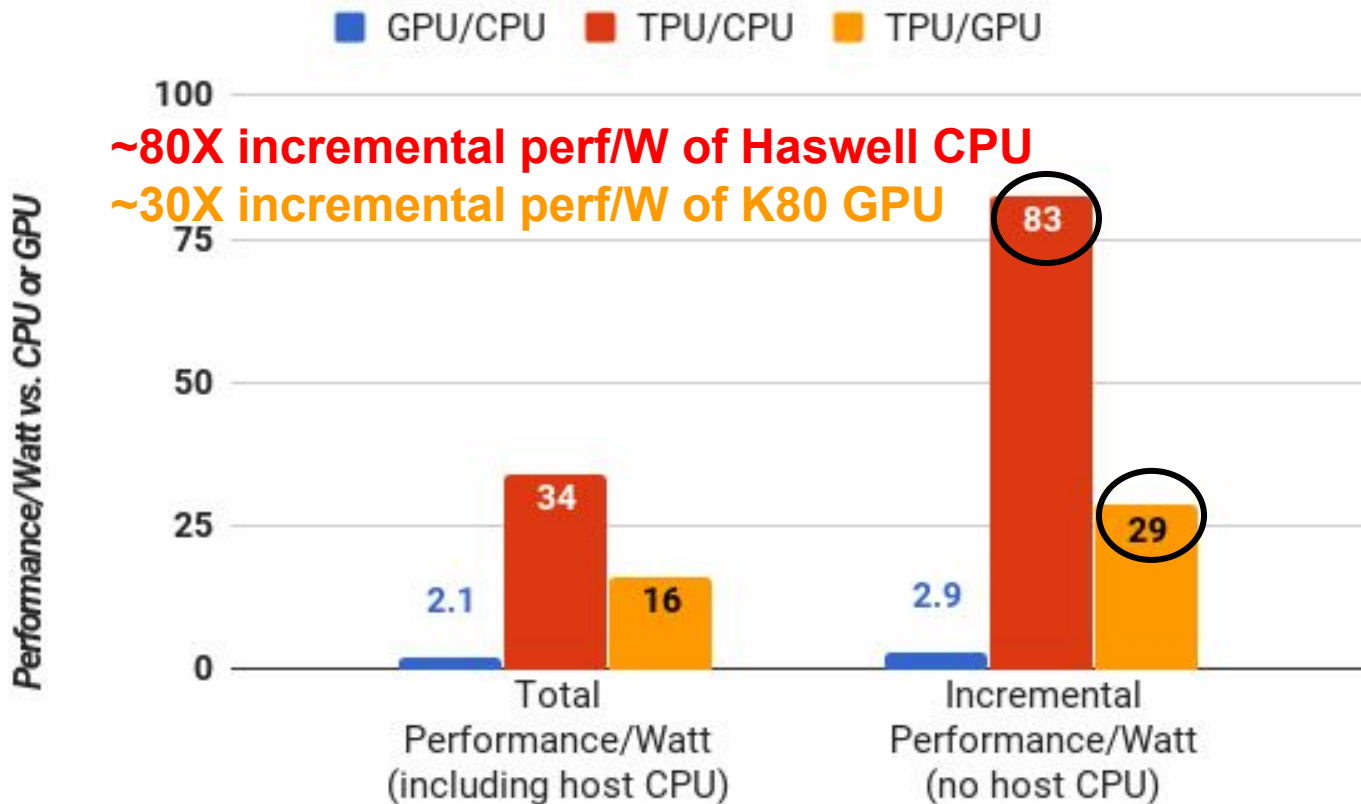
Linear Rooflines for CPU, GPU, TPU



TPU & GPU Relative Performance to CPU

<i>Type</i>	<i>MLP</i>		<i>LSTM</i>		<i>CNN</i>		<i>Weighted Mean</i>
	<i>0</i>	<i>1</i>	<i>0</i>	<i>1</i>	<i>0</i>	<i>1</i>	
GPU	2.5	0.3	0.4	1.2	1.6	2.7	1.9
TPU	41.0	18.5	3.5	1.2	40.3	71.0	29.2
Ratio	16.7	60.0	8.0	1.0	25.4	26.3	15.3

Perf/Watt TPU vs CPU & GPU

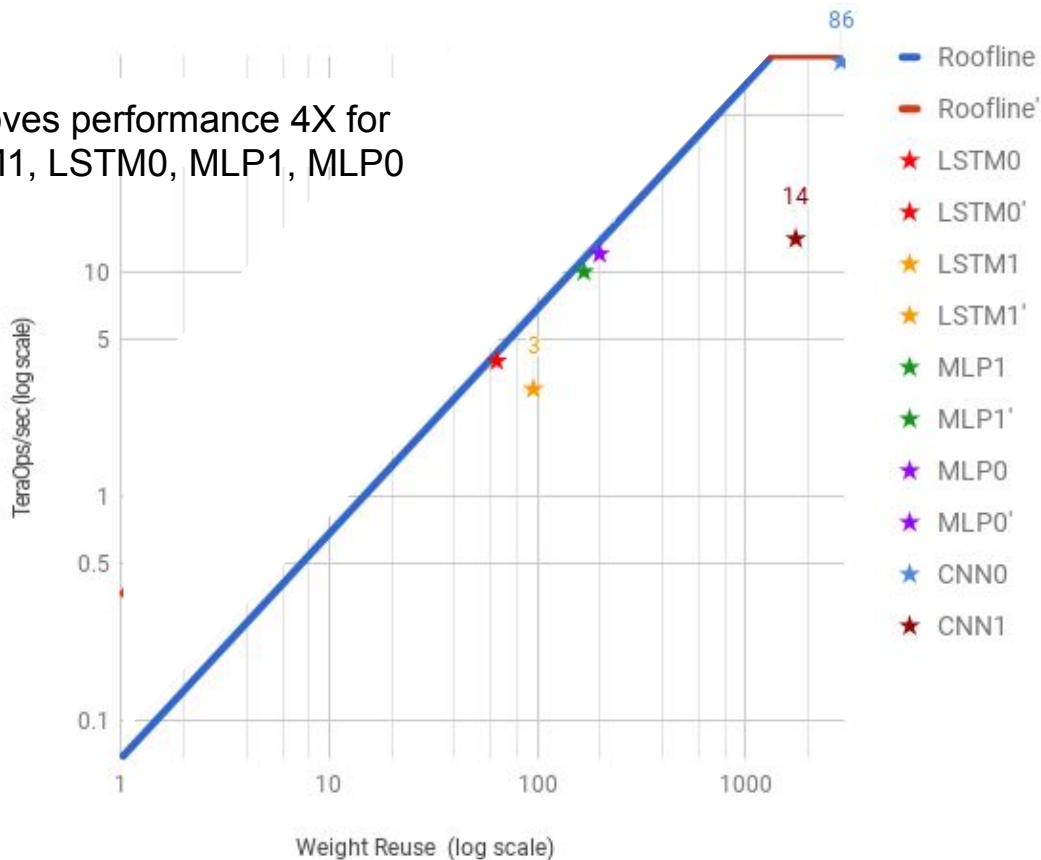


Improving TPU: Move “Ridge Point” to the left


- Current DRAM
 - 2 DDR3 2133 \Rightarrow 34 GB/s
- Replace with GDDR5 like in K80 \Rightarrow 180 GB/s
 - Move Ridge Point from 1400 to 256

Revised TPU Raises Roofline

Improves performance 4X for LSTM1, LSTM0, MLP1, MLP0



Perf/Watt Original & Revised TPU



Related Work

Two survey articles document that custom NN ASICs go back at least 25 years [Irn96][Asa02]. For example, CNAPS chips contained a 64 SIMD array of 16-bit by 8-bit multipliers, and several CNAPS chips could be connected together to a sequencer [Ham90]. The Synapse-1 system was based on a custom systolic multiply-accumulate chip called the MA-16, which performed sixteen 16-bit multiplies at a time [Ram91]. The system concatenated several MA-16 chips together and had custom hardware to do activation functions.

Twenty-five SPERT-II workstations, accelerated by the T0 custom ASIC, were deployed starting in 1995 to do both NN training and inference for speech recognition [Asa98]. The 40-MHz T0 added vector instructions to the MIPS instruction set architecture. The eight-lane vector unit could produce up to sixteen 32-bit arithmetic results per clock cycle based on 8-bit and 16-bit inputs, making it 25 times faster at inference and 20 times faster at training than a SPARC-20 workstation. They found that 16 bits were insufficient for training, so they used two 16-bit words instead, which doubled training time. To overcome that drawback, they introduced “bunches” (batches) of 32 to 1000 data sets to reduce time spent updating weights, which makes it faster than training with one word but no batches.

To do more recent DianNao family of NN architectures minimizes memory accesses both on the chip and to external DRAM by having efficient architectural support for the memory access patterns that appear in NN applications [Keu16]. [Che16a]. All use 16-bit integer operations and all designs do down to layout, but no chips were fabricated. The original DianNao uses an array of 16-bit integer multiply-accumulate units with 44 kB of on-chip memory and is estimated to be 3 mm² (65 nm), run at 1 GHz, and to consume 0.5W [Che14a]. Most of this energy went to DRAM accesses for weights, so one successor DianNao (“big computer”) includes eDRAM of 365 MiB of weights on chip [Che14b]. The goal was to have enough memory in a multiplex system to avoid external DRAM accesses. The follow-on PuDianNao (“general computer”) is aimed at more traditional machine learning algorithms beyond DNNs, such as support vector machines [Liu15]. Another offshoot is ShiDianNao (“vision computer”) aimed at CNNs, which avoids DRAM accesses by connecting the accelerator directly to the sensor [Dai15].

The Convolution Engine is also focused on CNNs for image processing [Qad13]. This design deploys 64 16-bit multiply-accumulator units and customizes a Tensorflow processor estimated to run at 800 MHz in 45 nm. It is projected to be 8X to 15X more energy-efficient than an SIMD processor, and within 2X to 3X of custom hardware designed just for a specific kernel.

The Fatfish benchmark paper recently reports results contradictory to ours, with the GPU training inference much faster than the CPU [Ada16]. However, their CPU and GPU are not server-class, the CPU has only four cores, the applications do not use the CPU’s AVX instructions, and there is no response-time cutoff (see Table 4) [Bro16].

Catapult is the most widely deployed example of using reconfigurability to support DNNs, which many have proposed [Far09][Choi10][F11][Pec13][Cav15][Zha15]. These chips FPGAs over GPUs to reduce power as well as the risk that latency-sensitive applications wouldn’t map well to FPGAs. FPGAs can also be reprogrammed, such as for search, compression, and network interface cards [Put15]. The TPU project actually began with FPGAs, but we abandoned them when we saw that the FPGAs of that time were not competitive in performance compared to the GPUs of that time, and the TPU could be much lower power than GPUs while being as fast or faster, giving it potentially significant benefits over both of FPGAs and GPUs.

Although first published in 2014 [Put14], Catapult is a TPU contemporary since it deployed 28-nm Stratix V FPGAs into datacenters concurrently with the TPU in 2015. Catapult has a 200 MHz clock, 3,926 18-bit MACs, 5 MiB of on-chip memory, 11 GB/s memory bandwidth, and uses 25 Watts. The TPU has a 700 MHz clock, 65,536 8-bit MACs, 28 MB, 34 GB/s, and typically uses 40 Watts. A revised version of Catapult uses newer FPGAs and was called out at large scale in [Ada16].

Catapult V1 runs CNNs—and a systolic matrix multiplier—2.3X as fast as a T16, 16-core, 164-dual-socket server [Ovi15a]. Using the next generation of FPGAs (14-nm Artix 10) of Catapult V2, performance might go up to 7X, and perhaps even 17X with more careful floorplanning [Ovi15b]. Although it’s apples versus oranges, a current TPU die runs its CNNs 40X to 70X versus a somewhat faster server (Tables 2 and 6). Perhaps the biggest difference is that to get the best performance the user must write long programs in the low-level hardware-design-language Verilog [Mel16][Put16] versus writing short programs using the high-level Tensorflow framework. That is, reprogrammability comes from software for the TPU rather than from firmware for the FPGA.

Recent research, which appeared after the TPU was deployed, accelerates DNNs by optimizing the cases when weights and data are very small or zero. Our tight schedule precluded such optimizations in the TPU, but we saw the same opportunity in our studies. The Efficient Inference Engine is based on a first pass that reduces the number of weights by about a factor of 10 [Han15] as a separate step by filtering out very small values and then uses Huffman encoding to shrink the data even further to improve inference performance [Han16]. Cnvlutin [Aib16] avoids multiplies when an activation input is zero—which it is 44% of the time, presumably in part due to ReLU nonlinear function that transforms negative values to zero—to improve performance by an average 1.4 times.

Eyeris is a novel, low-power dataflow architecture that takes advantage of zeros by run-length encoding data to reduce the memory footprint and saves power by avoiding computations when an input is zero [Che16a]. Using Eyeris terminology, a TPU convolutional layer maps C and M to the rows and columns of the matrix unit, taking HWN cycles to perform one pass. With high C/M, it takes RS passes to process the layer; for low C/M, a number of techniques reduce passes and improve utilization. (More can be found in the online references [Ros15a][Ros15b][Ros15c][Ros15d][Tho15].)

Minerva is a cross-layer system that crosses algorithm, architecture, and circuit disciplines to reduce power by 8X in a pass by pruning activation data with small values and in part by quantizing the data [Rea16]. [Cup15] looks at 16-bit fixed-point arithmetic for training instead of for inference. Others leverage the lower precision of DNN calculations by utilizing analog circuits during the computation to improve energy and performance [LJK16] [Sha16]. By tailoring an instruction set to DNNs, Cambridge reduces code size [Liu16]. Recent work looked at processor-in-memory architectures for NNs [Chil16][Kim16].

Related Work

Comparing the TPU to some of these architectures:

- [Che14a] DMAAs data from DRAM to input and weight buffers. They are read by the 3-stage pipelined NFU that performs multiplies, adds, and non-linear-functions; the results go to the output buffer, and then to DRAM. The NFU has no storage and isn’t systolic.
- [Gup15] appears to stream both matrix inputs while storing partial sums in the systolic array; the TPU stores the weights in a stochastic round robin while streaming the other input and the pre-activation partial sums. The TPU doesn’t support stochastic rounding.
- [Zha15] is built out of computation units equivalent to a 4x2 version of the TPU matrix unit. In an ASIC, the wiring cost of the crossbars that connect input and output buffers to these compute engines would be significant. We are surprised that we didn’t see architectural support for additional reductions to combine results from compute engines in [Zha15].

All three of [Gup15][Che14a][Zha15] store activations in DRAM during computation; the TPU’s Unified Buffer is sized so that no DRAM spilling or reloading happens in normal operation.

References

[Aba16] Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G.S., Davis, A., Dean, J., Devin, M., Ghemawat, S., 2016. TensorFlow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*.

[Aib16] Alberico, J., Judd, P., Hetherington, T., Aamodt, T., Jeger, N.E., and Moshovos, A. 2016. Convolution: Inefficient-Nexus-Free Deep Neural Network Computing. *Proceedings of the 43rd International Symposium on Computer Architecture*.

[Ada16] Adadi, R., Banaś, S., Reagen, B., Wu, C.Y., and Brooks, D., 2016, September. Fatfish: reference workloads for modern deep learning methods. *IEEE International Symposium on Workload Characterization (ISWC)*.

[Ara02] Awasaki, K. 2002. Programmable Neurocomputing. In *The Handbook of Brain Theory and Neural Networks: Second Edition*, M. A. Arbib (Ed.), MIT Press, ISBN 0-262-01197-2, November 2002. <http://people.cs.berkeley.edu/~arbib/handbook-neurocomp.pdf>

[Asa98] Awasaki, K., 1998. Awasaki, K., Beck, Johnson, J., Wazarynek, J., Kingstbury, B. and Morgan, N., November 1998. Training Neural Networks with Spars-B. Chapter 11 in *Parallel Architectures for Artificial Neural Networks: Paradigms and Implementations*, N. Sandhu and P. Sarachandran (Eds.), IEEE Press, ISBN 0-7621-0197-2, October 1998. <http://people.cs.berkeley.edu/~arbib/handbook-neurocomp.pdf>

[Bar16] Barua, L.A. and Hilde, U., 2007. The case for energy-proportional computing. *IEEE Computer*, vol. 40.

[Bar16] Bar, J. September 26, 2016. New IP2: Intense Type for Amazon EC2 up-to 16 GPUs. <http://aws.amazon.com/gpu/accelerated-ip2-optimized-type-for-amazon-ec2-up-to-16-gpu/>

[Bro16] Brooks, D. November 4, 2016. Private communication.

[Cav16] Caulfield, A.M., Chang, E.S., Patnam, A., Haselman, H.A.J.F.M., Humphrey, S.H.M., Danzig, P.K.J.V.K., Oviath, L.T.M.K., Lanka, M.P. S.S. and Biagio, D.C.D. 2016. A 16k-Scale Acceleration Architecture. *MF2016 conference*.

[Caw15] Casagelli, L., Gheung, D., Mayer, C., Willis, S., Mariani, B. and Benini, L. 2015, May. On-chip convolutional neural networks. *Proceedings of the 25th edition of Great Lakes Symposium on VLSI*.

[Chai16] Chai, Y., 2016. *Chai: A Dynamically Configurable Coprocessor for Convolutional Neural Networks*. *Proceedings of the 37th International Symposium on Computer Architecture*.

[Che14a] Chen, T., Dao, Z., Sun, N., Zhang, J., Wu, C., Wang, Y., and Temam, O., 2014. DianNao: A small-footprint high-throughput accelerator for ubiquitous machine-learning. *Proceedings of the 77th Annual International Symposium on Microarchitecture*.

[Che14b] Chen, Y.H., Chen, T., Sun, N., Wang, J., Wu, C., Wang, Y., and Temam, O., 2014. Efficient Architecture for Energy-Efficient Dataflow for Convolutional Neural Networks. *Proceedings of the 43rd International Symposium on Computer Architecture*.

[Che14c] Chen, T., Chen, T., Xu, S., Sun, N., Zhang, S., and Temam, O., 2016. DianNao Family: Energy-Efficient Hardware Accelerators for Machine Learning. *Proceedings of the Highlights, Communications of the ACM*, 9(11), 1-10.

[Chi16] Chi, P., Li, S., Gu, P., Xu, C., Zhang, T., Zhao, J., Liu, Y., Wang, Y., and Xie, Y., 2016. PRIME: A Novel Processing-in-Memory Architecture for Neural Network Computation in ReRAM-Based Main Memory. *Proceedings of the 43rd International Symposium on Computer Architecture*.

[Choi15] Clark, J. October 26, 2015. Google Training Its Lustrative Web Search Over AI at Machines. *Blockchain Technology*, www.blocknet.com.

[Dal16] Dalby, W. February 9, 2016. High Performance Hardware for Machine Learning. *Caltech ENN Summit*.

[Dea13] Dea, J. and Hwang, L.-A., 2013. The fall of scale. *Communications of the ACM*, 56(2).

[Dean14] Dean, J., July 7, 2014. Large Scale Deep Learning with TensorFlow. *Systems, ACM Webinar*.

[Dai15] Dai, Z., Fadhler, R., Chen, T., Imani, P., Li, L., Luo, T., Feng, X., Chen, Y., and Temam, O., 2015. The ShiDianNao: Shifting system complexity to the user. *Proceedings of the 42nd International Symposium on Computer Architecture*.

[Far11] Farabet, C., Puaud, C., Han, J.Y. and Lecun, Y., 2009. August. Giga-FA: An FPGA-based accelerator for convolutional networks. 2009 International Conference on Field Programmable Logic and Applications.

[Far11] Farabet, C., Marin, B., Corda, B., Akelof, P., Culevischi, F., and Lecun, Y., 2011. June. *NeuFlow: A runtime reconfigurable dataflow processor for vision*. In *CVPR 2011 Workshop*.

[Gup15] Gupta, S., Agrawal, A., Goldkornik, K. and Narayanan, P., 2015, July. Deep Learning with Limited Numerical Precision. *ICML*.

[Ham90] Hammar, D., 1990. June. A VLSI architecture for high-performance, low-cost, on-chip learning. 1990 *ICNN International Neural Network Conference*.

[Han15] Han, S., Pool, J., Tran, J., and Dally, W. 2015. Learning both weights and connections for efficient neural networks. In *Advances in Neural Information Processing Systems*.

[Han16] Han, S., Liu, X., Liu, H., Pu, J., Pedram, A., Hosoi, M.A. and Dally, W., 2016. EE: efficient inference engine on compressed deep neural network. *Proceedings of the 43rd International Symposium on Computer Architecture*.

[He16] He, K., Zhang, X., Ren, S. and Sun, J., 2016. Identity mappings in deep residual networks. Also in *arXiv preprint arXiv:1603.06027*.

[Hetherington, J.I. and Patterson, D.A., 2014. *Computer architecture: a quantitative approach*, 4th edition, Elsevier.

[Hilde16] Hilde, U. and Barua, L., 2009. *The datacenter as a computer*. Morgan and Claypool.

[Irn96] Imani, P., Coen, T. and Kuhn, G., 1996. Special-purpose digital hardware for neural networks: An architectural survey. *Journal of VLSI signal processing systems for robotics, I*, 1-11.

[Intel16] Intel, 2016. Intel® Xeon® Processor E3-4609 v5. http://ark.intel.com/products/67560/intel_xeon_processor_e3-4609_v5-4M-Cache-2_16-GHz

[Joep16] Joep, N. May 18, 2016. Google supercharges machine learning tasks with TPU custom chip. <http://cloudplatform.googleblog.com/2016/05/Google-TPU.html>

[Katz16] Katz, K., 2011. On-chip programmable logic. *Communications of the ACM*, 54(7), 981-11.

[Kim16] Kim, D., Kang, H.H., Cha, S., Vatanparast, S. and Makhadmeh, S., 2016. NeuroArch: A Programmable Digital Neuromorphic Architecture with High-Density 3D Memory. *Proceedings of the 43rd International Symposium on Computer Architecture*.

[Kishino12] Kishino, A., Saitoh, E. and Hirota, G., 2012. Efficient classification with deep convolutional neural networks. *Advances in neural information processing systems*.

[Kuro10] Kung, H.T. and Leiserson, C.E., 1980. Algorithms for VLSI processor arrays. *Introduction to VLSI systems*.

[Lang16] Lang, J.D., 2009. Identifying shades of gray. *The SPECpower benchmark*. <http://www.spec.com/powerbench/>

[Lar16] Landau, M. March 10, 2016. Google Looks To Open Up StreamCloud To Make GPU Programming Easier, Phoronix, http://www.phoronix.com/scan.php?page=news_item&catid=google-StreamCloud-Easier-Parallel.

[LJK16] Li, Kan-Wei, R. Hou, Y., Guo, J., Polansky, M. and Zhong, L., 2016. RedEye: Analog ConvNet Image Sensor Architecture for Continuous Mobile Vision. *Proceedings of the 43rd International Symposium on Computer Architecture*.

[Liu15] Liu, Q., Chen, T., Liu, S., Zhao, Z., Temam, O., Feng, X., Zhou, S., and Chen, Y., 2015. Masch. Padmanava: A polyvalent machine learning accelerator. *Proceedings of the 42nd International Symposium on Computer Architecture*.

[Liu16] Liu, S., Dao, Z.D., Tao, J.H., Han, D., Luo, T., Xie, Y., Chen, Y. and Chen, T., 2016. Cambridge: An instruction set architecture for neural networks. *Proceedings of the 43rd International Symposium on Computer Architecture*.

[Met16] Metz, C. September 26, 2016. Microsoft Beta Its Future On A Reconfigurable Computer Chip. *Wired Magazine*, <http://www.wired.com/2016/09/microsoft-beta-future-chip-reprogramm-0/>

[Nvi15] Nvidia, January 2015. Tesla K80 GPU Accelerator. *Board Specification* <http://nvidia.com/content/pdf/tesla-k80-board-spec-0711-001-005.pdf>

[Nvidia] Nvidia, 2016. Tesla GPU Accelerator Servers. <http://www.nvidia.com/object/tesla-servers.html>.

[Ovi15a] Oviath, K., Rawase, O., Kim, J.Y., Fowers, J., Straus, K. and Chang, E.S., February 2, 2015. Accelerating deep convolutional neural networks using specialized hardware. Microsoft Research Whitepaper, <http://www.microsoft.com/en-us/research/publication/accelerating-deep-convolutional-neural-networks-using-specialized-hardware/>

[Ovi15b] Oviath, K., Rawase, O., Kim, J.Y., Fowers, J., Straus, K. and Chang, E.S., 2015, August. Toward accelerating deep learning at scale using specialized hardware in the datacenter. 2015 *IEEE Hot Chips 27 Symposium*.

[Pata04] Patterson, D.A., 2004. Latency-lags hardware. *Communications of the ACM*, 47(10).

[Pec13] Peenen, M., Setis, A.A., Mesnam, B. and Corporal, H., 2013, October. Memory-centric design for convolutional neural networks. In 2013 *IEEE/ACM International Conference on Computer Design (ICCD)*.

[Put14] Patnam, A., Caulfield, A.M., Chang, E.S., Chio, D., Costantinescu, K., Demme, J., Farnazizadeh, H., Fowers, J., Gopal, G.P., Gray, J., Haselman, M., Hauck, S., Heil, S., Horns, A., Kim, J.Y., Lanka, S., Laria, J., Peterson, E., Pope, S., Smith, A., Thong, J., Xiao, P.Y., Burger, D., 2014, June. A reconfigurable fabric for accelerating large-scale datacenter services. 41st *International Symposium on Computer Architecture*.

[Put16] Patnam, A., Caulfield, A.M., Chang, E.S., Chio, D., Costantinescu, K., Demme, J., Farnazizadeh, H., Fowers, J., Gopal, G.P., Gray, J., Haselman, M., Hauck, S., Heil, S., Horns, A., Kim, J.Y., Lanka, S., Laria, J., Peterson, E., Pope, S., Smith, A., Thong, J., Xiao, P.Y., Burger, D., 2016. A Reconfigurable Fabric for Accelerating Large-Scale Datacenter Services. *Communications of the ACM*, 9(11), 1-10.

[Qad13] Qader, W., Hamed, R., Shachan, O., Venkatesan, P., Kozuyk, C. and Horowitz, M.A., 2013, June. Convolution engine: balancing efficiency and flexibility in specialized computing. *Proceedings of the 40th International Symposium on Computer Architecture*.

[Rama11] Ramesher, U., Reichter, J., Rahn, W., Andujar, J., Bruns, N., Hachmann, U. and Westeling, M., 1991. Design of a 1st Generation Neurocomputer. In *VLSI Design of Neural Networks*. Springer.

[Reag16] Reagen, B., Whittington, P., Adluri, B., Rama, S., Lee, H., Lee, S.K., Hernandez-Lobato, J.M., Wei, G.Y. and Brooks, D., 2016. Minerva: Enabling low-power, highly-accurate deep neural network acceleration. *Proceedings of the 43rd International Symposium on Computer Architecture*.

[Ros15a] Ross, J., Joop, N., Phelps, A., Young, S., Norman, A., Thorogood, G., Liu, D., 2015. Neural Network Processor. Patent Application No. 62/164,931.

[Ros15] Ross, J., Phelps, A., 2015. Computing Convolutions Using a Neural Network Processor. Patent Application No. 62/164,902.

[Ros15] Ross, J., 2015. Prefetching Weights for a Neural Network Processor. Patent Application No. 62/164,981.

[Ros16] Ross, J., Thomson, G., 2015. Rotating Data for Neural Network Computations. Patent Application No. 62/164,908.

[Rou15] Rousavsky, O., Deng, J., Su, H., Krnizek, S., Safirsch, S., Ma, S., Huang, Z., Kapurthy, A., Khosla, A., Benetman, M. and Berg, A.C., 2015. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(1).

[Schu16] Schuman, E. and Brubaker, J., 2009, June. The user and business impact of server delays, additional bytes, and HTTP chunking in web search. In *Web Performance and Optimization Conference*.

[Sha16] Shafiq, A., Nig, A., Muralimohan, N., Balasubramanian, R., Strachan, J.P., Hu, M., Williams, R.S. and Srikanan, V., 2016. ISAAC: A Convolutional Neural Network Accelerator with In-Situ Analog Arithmetic in Crossbars. *Proceedings of the 43rd International Symposium on Computer Architecture*.

[Sil16] Silver, D., Huang, A., Maddison, C.J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneerselvam, V., Lanctot, M. and Dieleman, S., 2016. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587).

[Stein12] Steinberg, J.E., 1982. April. Decoupled accelerators for convolutional networks. *Proceedings of the 11th International Symposium on Computer Architecture*.

[Stein15] Steinberg, D., 2015. Full-Chip Simulations, Keys to Success. *Proceedings of the Synopsys User Group (SUGUG) Silicon Valley 2015*.

[Tho15] Theogly, C., Liu, X., Wang, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V. and Rabinovich, A., 2015. Going deeper with convolutions. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.

[Tho15] Thonson, G., Clark, C., Liu, D., 2015. Vector Computation Unit in a Neural Network Processor. Patent Application No. 62/165,022.

[Wom09] Williams, S., Waterman, A. and Patterson, D., 2009. RedEye: an insightful visual processor model for multicore architectures. *Communications of the ACM*.

TPU succeeded because of

Conclusions (1/2)

- Large matrix multiply unit
- Substantial software-controlled on-chip memory
- Run whole inference models to reduce host CPU
- Single-threaded, deterministic execution model
good match to 99th-percentile response time
- Enough flexibility to match NNs of 2017 vs. 2013
- Omission of GP features \Rightarrow small, low power die
- Use of 8-bit integers in the quantized apps
- Apps in TensorFlow, so easy to port at speed

Conclusions (2/2)

- Inference prefers latency over throughput
- K80 GPU relatively poor at inference (vs. training)
- Small redesign improves TPU at low cost
- 15-month design & live on I/O bus yet TPU
15X-30X faster Haswell CPU, K80 GPU (inference),
<1/2 die size, 1/2 Watts
 - 65,536 (8-bit) TPU MACs cheaper, lower energy, &
faster 576 (32-bit) CPU MACs, 2496 GPU (32-bit) MACs
- 10X difference in computer products are rare

Questions?

*4/5/17 Google published a blog on the TPU. A 17-page technical paper with same title will be on arXiv.org. (Paper will also appear at the *International Symposium on Computer Architecture* on June 26, 2017.)

<https://cloudplatform.googleblog.com/2017/04/quantifying-the-performance-of-the-TPU-our-first-machine-learning-chip.html>