HETEROGENEOUS SUPERCOMPUTING AND THE POWER9 PROCESSOR

H. Peter Hofstee, Ph.D.
IBM & TU Delft

March 28, 2018

SCAsia 2018 Supercomputing

Agenda

- Motivation
- POWER9 - Made for acceleration/cooperation
- Acceleration, Network & Storage
- IBM AC922 & Rackspace/Google Zaius OCP
- HPC - Coral system & Posits
- Big Data - GPU-based sort & Arrow/Fletcher
- AI/Cognitive - Large model support
- Conclusions

Fundamental forces are accelerating change in our industry

IBM

IT innovation can no longer come from just the processor

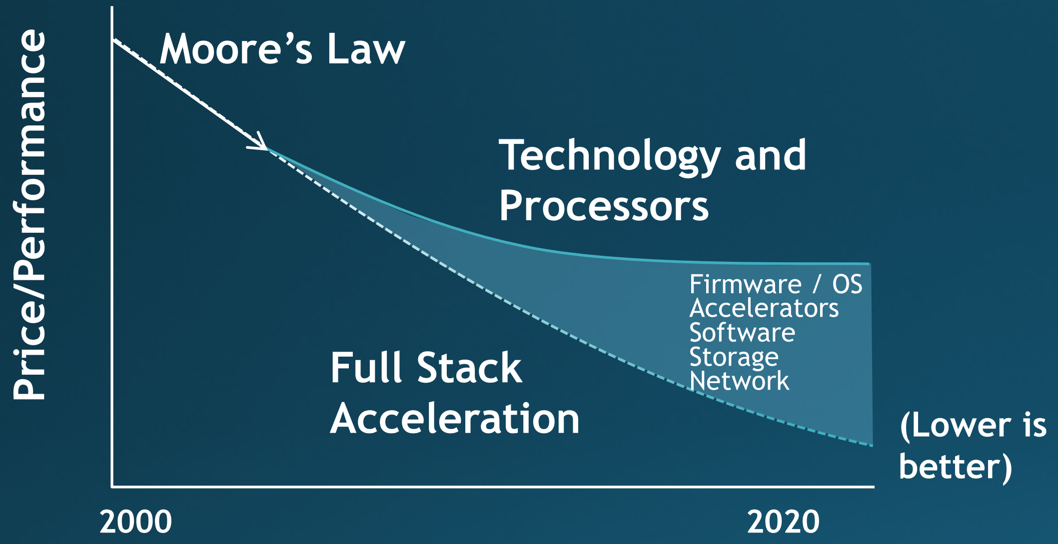

Full system stack innovation required

IT consumption models are expanding

- Cognitive
- Custom Hyperscale Data Centers
- Hybrid Cloud
- Open Solutions

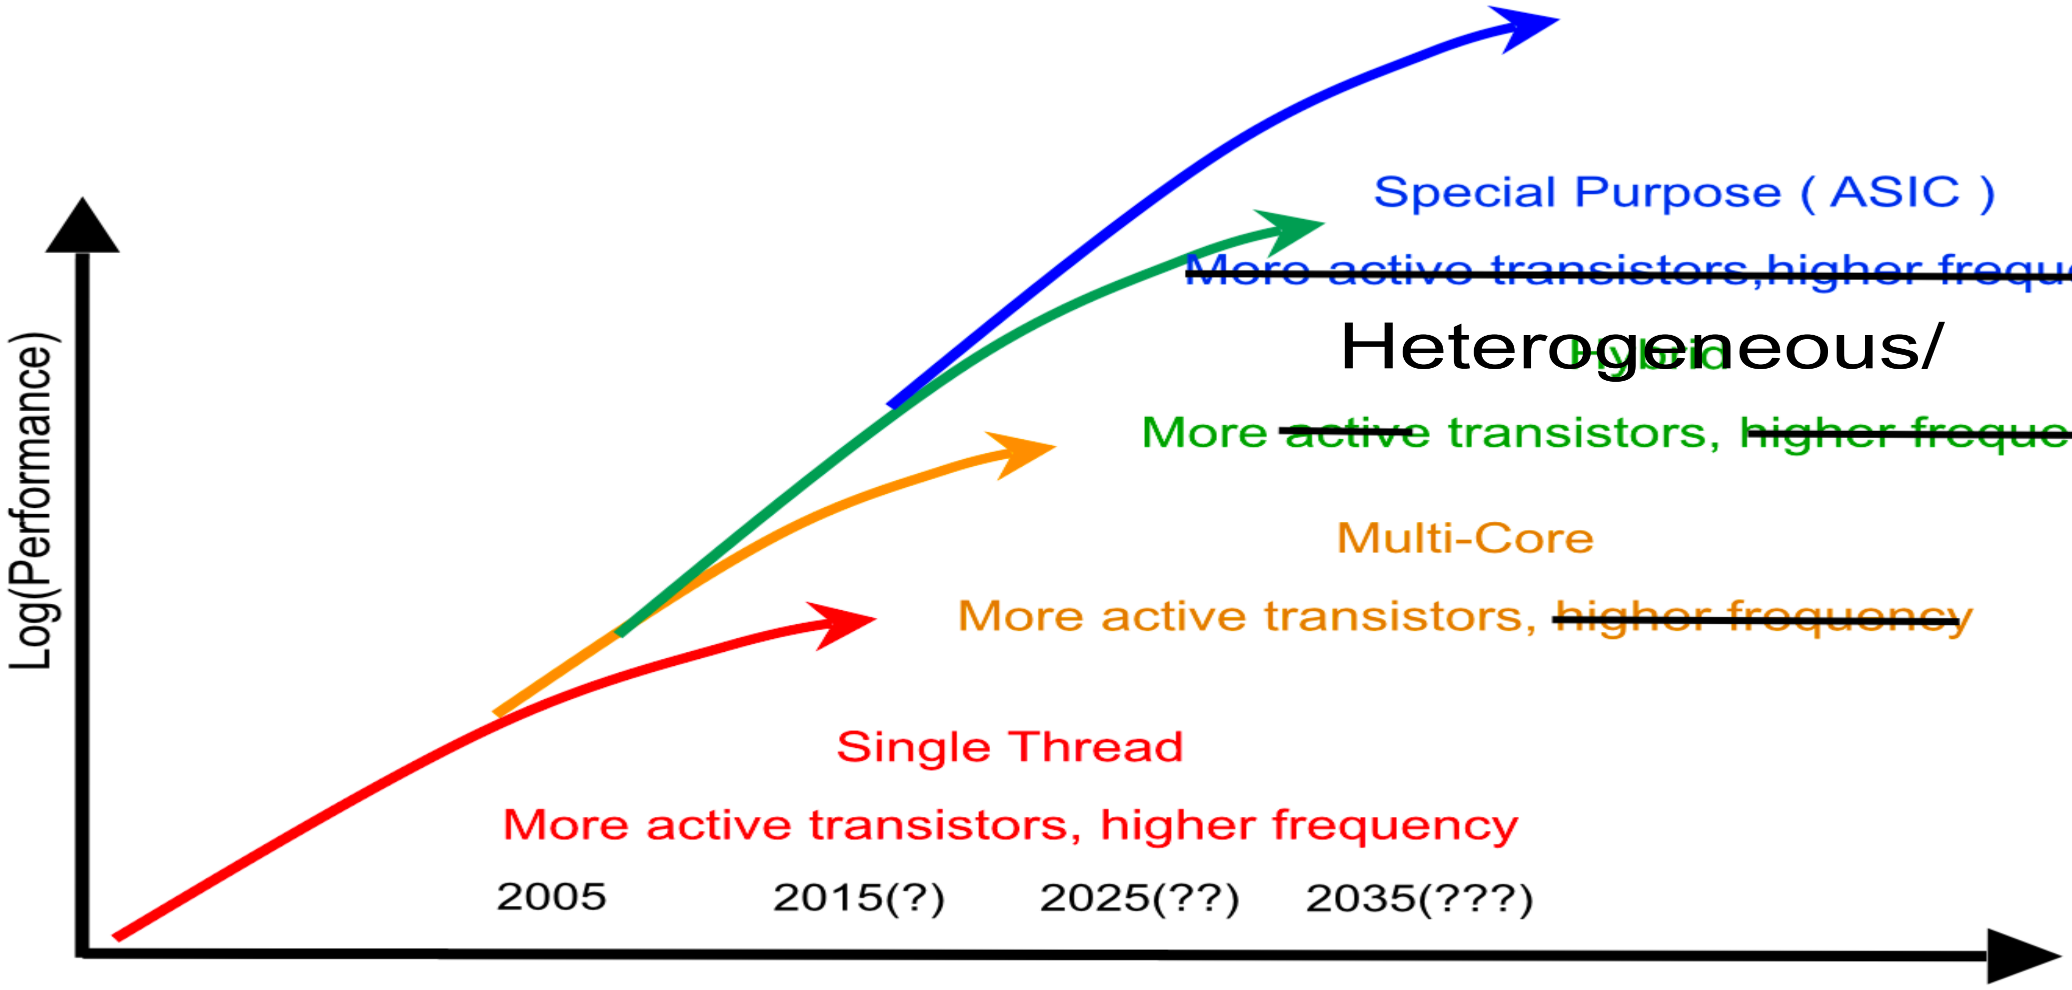
Special Purpose (ASIC)
More active transistors,higher frequency
Heterogeneous/
Hybrid
More active transistors, higher frequency
Multi-Core
More active transistors, higher frequency
Single Thread
More active transistors, higher frequency
2005
2015(?)
2025(??)
2035(???)
Log(Performance)

x

Network, Storage, & DRAM trends

Network
Storage
Grow Fast

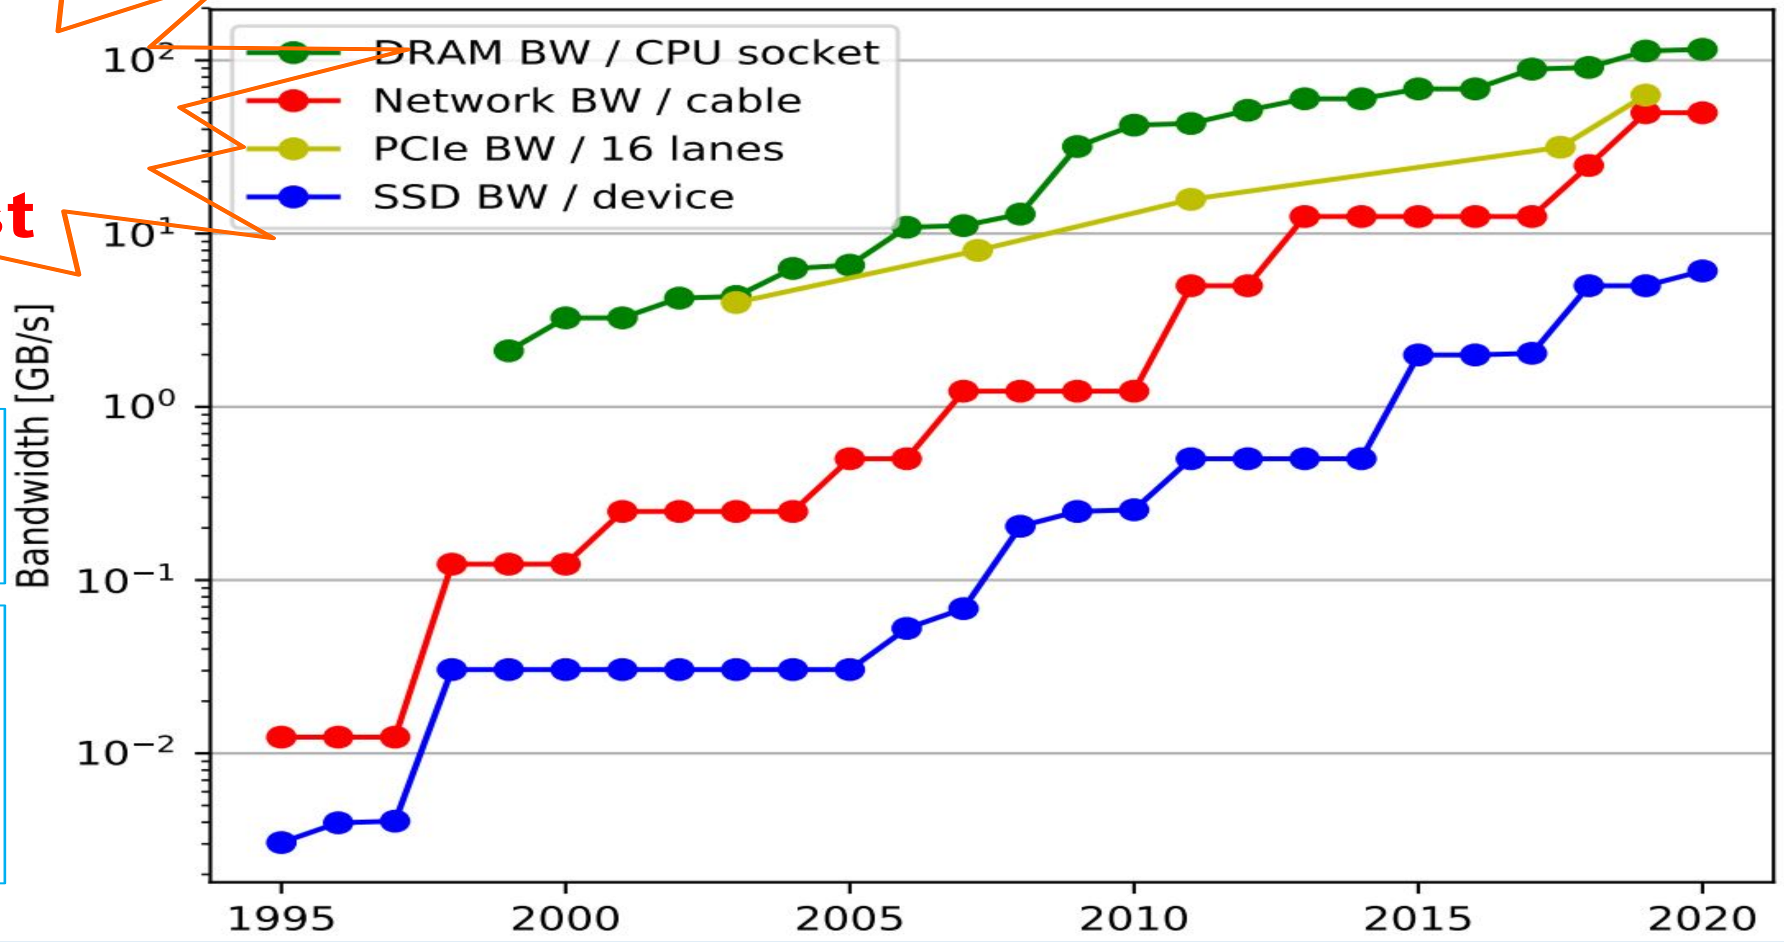

Use DRAM Bandwidth as a proxy for CPU throughput

Reasonable approximation for DMA and poor cache performance workloads (e.g. Storage)

Source of data: Sandisk IT blog

Proposed POWER Processor Technology and I/O Roadmap

	POWER7 Architecture		POWER8 Architecture		POWER9 Architecture			POWER10
	2010 **POWER7** 8 cores **45nm** New Micro-Architecture New Process Technology	**2012** **POWER7+** 8 cores **32nm** Enhanced Micro-Architecture New Process Technology	**2014** **POWER8** 12 cores **22nm** New Micro-Architecture New Process Technology	**2016** **POWER8 w/ NVLink** 12 cores **22nm** Enhanced Micro-Architecture With NVLink	**2017** **P9 SO** 24 cores **14nm** New Micro-Architecture Direct attach memory New Process Technology	**2018** **P9 SU** 24 cores **14nm** Enhanced Micro-Architecture Buffered Memory	**2019** **P9 w/ Adv. I/O** 24 cores **14nm** Enhanced Micro-Architecture New Memory Subsystem	**2020+** **P10** TBD cores New Micro-Architecture New Technology
Sustained Memory Bandwidth	Up To 65 GB/s	Up To 65 GB/s	Up To 210 GB/s	Up To 210 GB/s	Up To 150 GB/s	Up To 210 GB/s	Up To 350 GB/s	Up To 435 GB/s
Standard I/O Interconnect	PCIe Gen2	PCIe Gen2	PCIe Gen3	PCIe Gen3	PCIe Gen4 x48	PCIe Gen4 x48	PCIe Gen4 x48	PCIe Gen5
Advanced I/O Signaling	N/A	N/A	N/A	20 GT/s 160GB/s	25 GT/s 300GB/s	25 GT/s 300GB/s	25 GT/s 300GB/s	32 & 50 GT/s
Advanced I/O Architecture	N/A	N/A	CAPI 1.0	CAPI 1.0 , NVLink 1.0	CAPI 2.0, OpenCAPI3.0, NVLink2.0	CAPI 2.0, OpenCAPI3.0, NVLink2.0	CAPI 2.0, OpenCAPI4.0, NVLink3.0	TBD

Statement of Direction, Subject to Change

POWER9 – Premier Acceleration Platform

- **Extreme Processor / Accelerator Bandwidth and Reduced Latency**
- **Coherent Memory and Virtual Addressing Capability for all Accelerators**
- **OpenPOWER Community Enablement – Robust Accelerated Compute Options**

- **State of the Art I/O and Acceleration Attachment Signaling**
  - **PCIe Gen 4** x 48 lanes – 192 GB/s duplex bandwidth
  - **25G Link** x 48 lanes – 300 GB/s duplex bandwidth
- **Robust Accelerated Compute Options with OPEN standards**
  - **On-Chip Acceleration** – Gzip x1, 842 Compression x2, AES/SHA x2
  - **CAPI 2.0** – 4x bandwidth of POWER8 using *PCIe Gen 4*
  - **NVLink 2.0** – Next generation of GPU/CPU bandwidth and integration
  - **OpenCAPI** – High bandwidth, low latency and open interface using *25G Link*

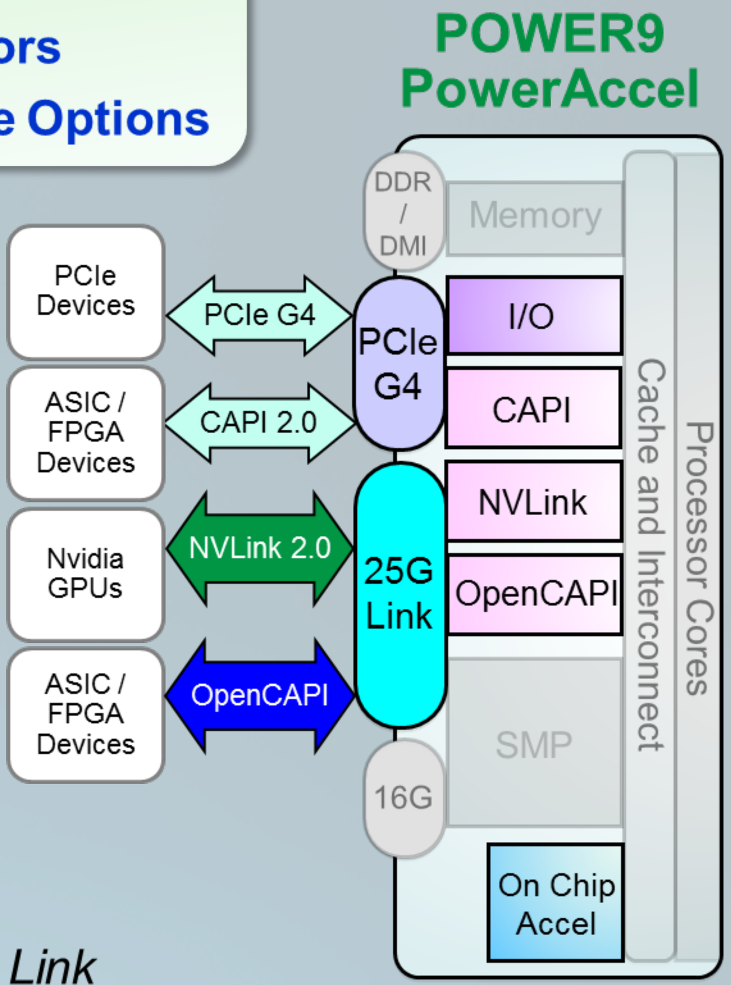

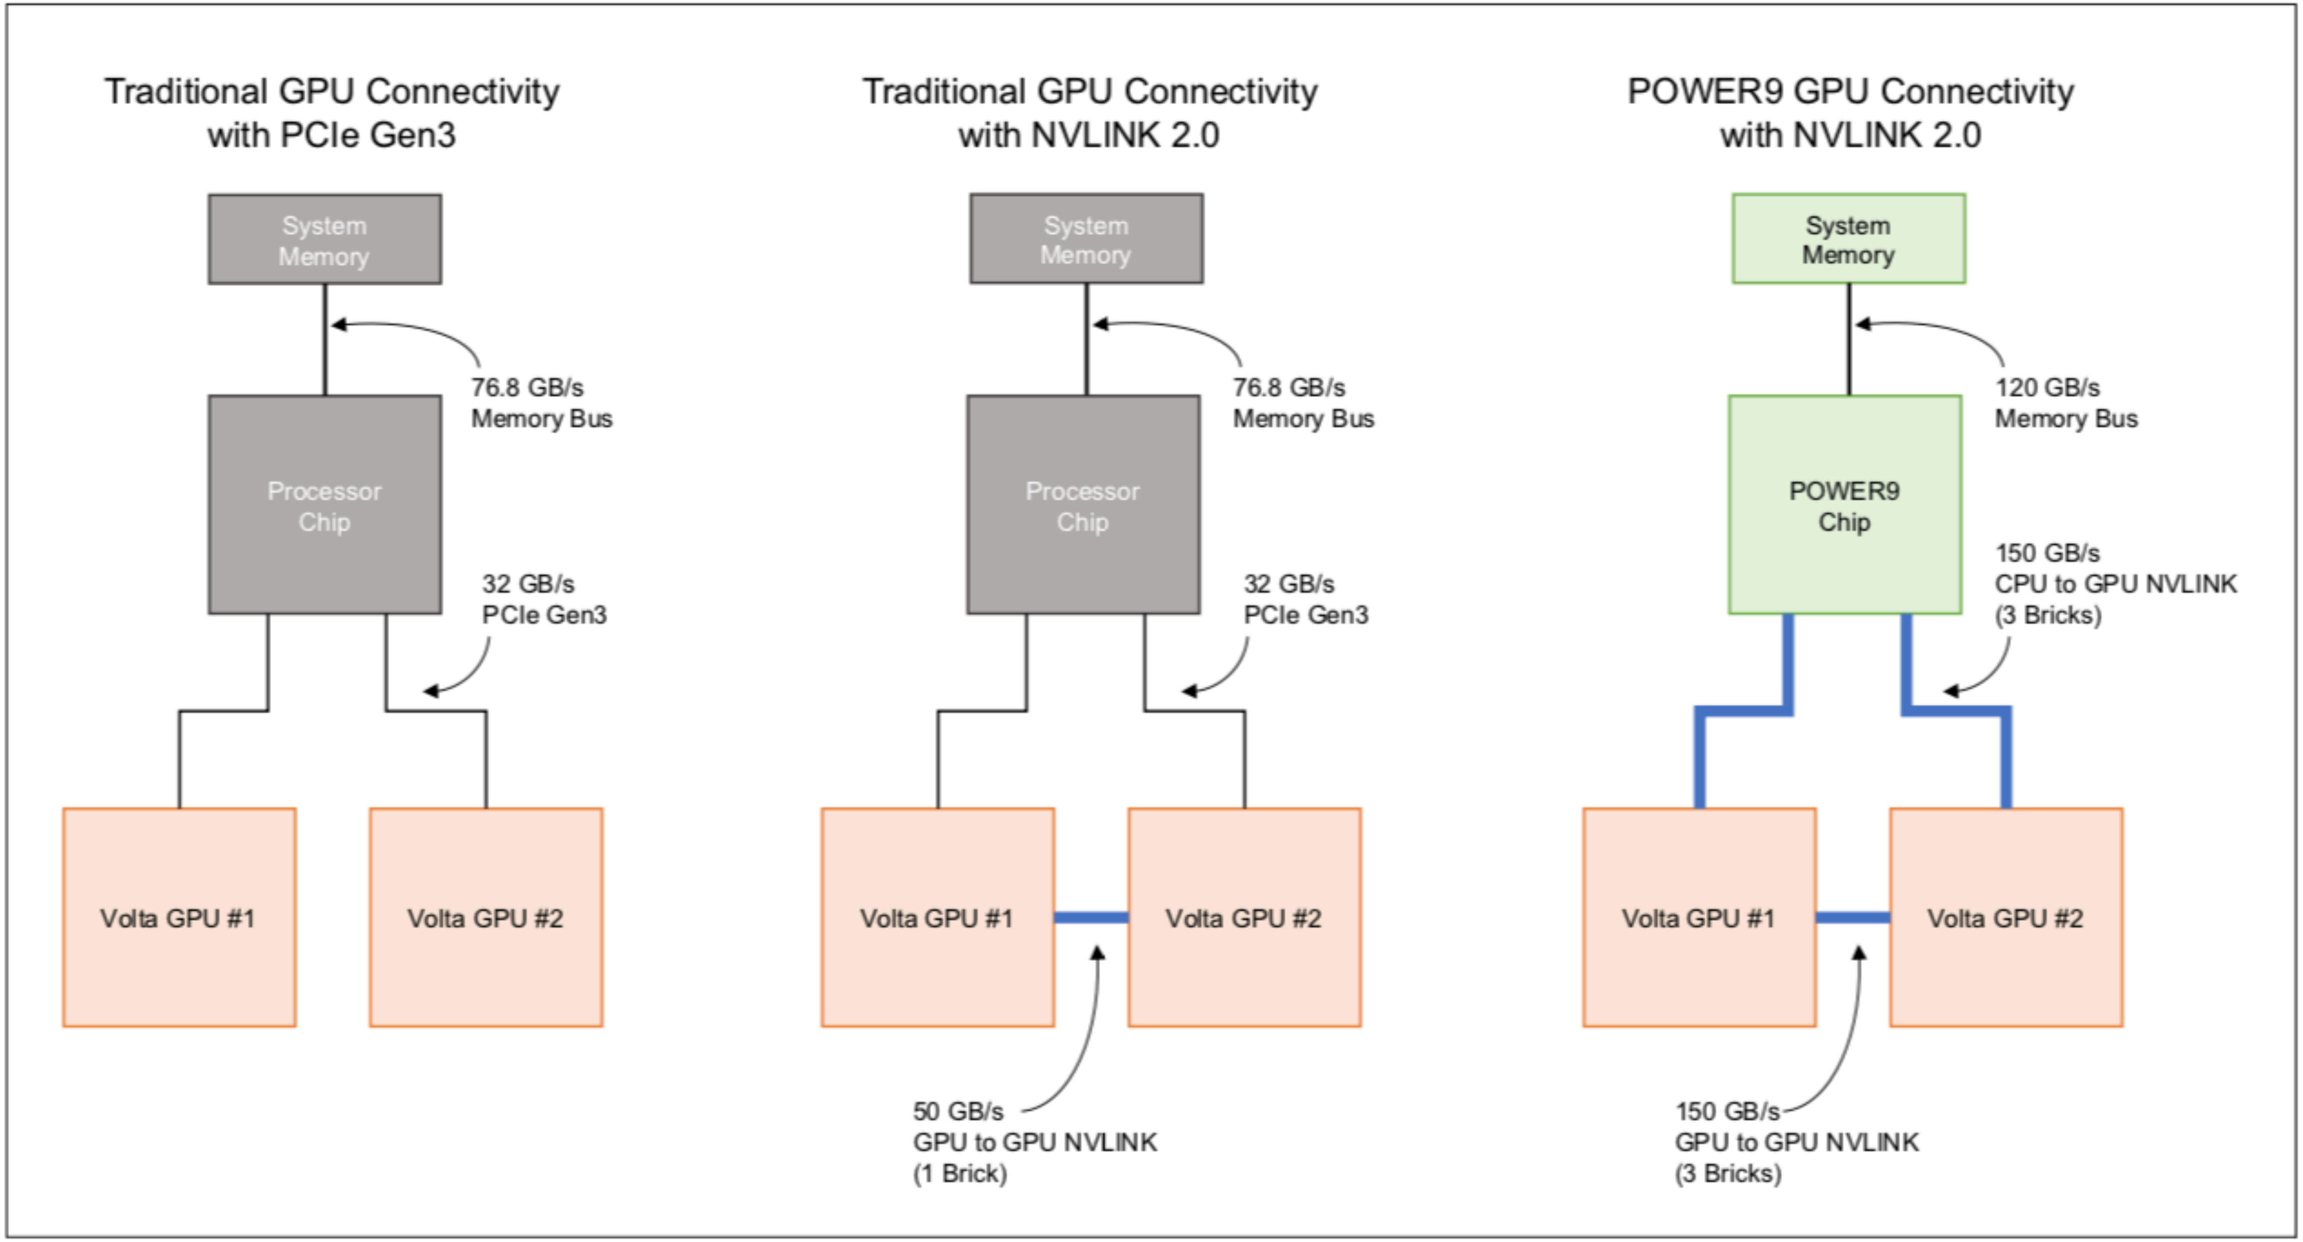
Traditional GPU Connectivity with PCIe Gen3
System Memory
76.8 GB/s Memory Bus
Processor Chip
32 GB/s PCIe Gen3
Volta GPU #1
Volta GPU #2
Traditional GPU Connectivity with NVLINK 2.0
System Memory
76.8 GB/s Memory Bus
Processor Chip
32 GB/s PCIe Gen3
Volta GPU #1
Volta GPU #2
50 GB/s GPU to GPU NVLINK (1 Brick)
POWER9 GPU Connectivity with NVLINK 2.0
System Memory
120 GB/s Memory Bus
POWER9 Chip
150 GB/s CPU to GPU NVLINK (3 Bricks)
Volta GPU #1
Volta GPU #2
150 GB/s GPU to GPU NVLINK (3 Bricks)


IBM


Coherence across CPU, GPU & Memory
Up to 2TB DDR4 industry standard memory with 340GB/s bandwidth
POWER9 Processor with 2nd Gen NVLink, up to 44 Cores (176 Threads)

CAPI Overview

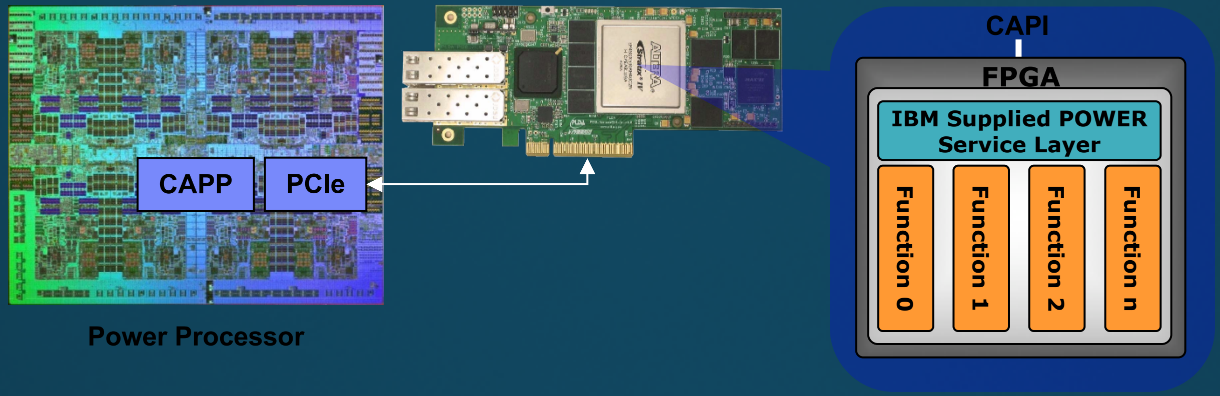

Typical I/O Model Flow

DD Call → Copy or Pin Source Data → MMIO Notify Accelerator → Acceleration → Poll / Int Completion → Copy or Unpin Result Data → Ret. From DD Completion

Flow with a Coherent Model

Shared Mem. Notify Accelerator → Acceleration → Shared Memory Completion

Advantages of Coherent Attachment Over I/O Attachment

- **Virtual Addressing & Data Caching**
  - Shared Memory
  - Lower latency for highly referenced data
- **Easier, More Natural Programming Model**
  - Traditional thread level programming
  - Long latency of I/O typically requires restructuring of application
- **Enables Applications Not Possible on I/O**
  - Pointer chasing, etc…

CAPI/OpenCAPI FPGA ACCELERATION

(example later)

	CAPI	Xilinx® FPGA	Memory	I/O Interfaces
ADM-PCIE-9V3 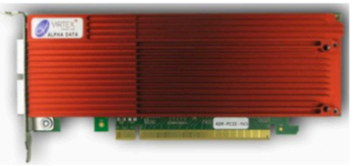	OpenCAPI 3.0 (PCIe G4) 2.0	UltraScale+™ VU3P 862K Logic Cells 2,280 DSP PCIe G3x16 / G4x8	16GB ECC (32GB option) DDR4-2400	Dual QSFP28 SlimSAS (25G x8) USB Board Management (JTAG built in) Customizable front GPIO
ADM-PCIE-8K5 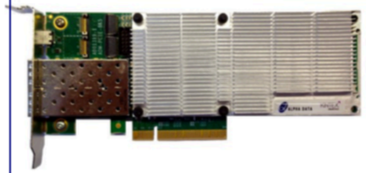	(PCIe G3) 1.1	UltraScale™ KU115 1,161K Logic Cells 5,520 DSP PCIe G3x8	16GB ECC (32GB option) DDR4-2400	Dual SFP+ Dual Firefly (16 x 16Gbps) USB Board Management (JTAG built in) Customizable front GPIO
ADM-PCIE-KU3 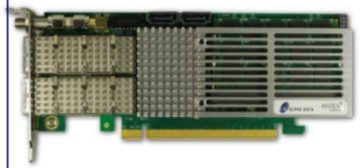	(PCIe G3) 1.1	UltraScale™ KU060 580K Logic Cells 2,760 DSP PCIe G3x8 (dual)	16GB ECC (32GB option) DDR3-1600	Dual QSFP+ Dual SATA GPIO / Timing
ADM-PCIE-7V3 	(PCIe G3) 1.1	Virtex-7 VX690T 693K Logic Cells 3,600 DSP PCIe G3x8	16GB ECC DDR3-1333	Dual SFP+ Dual SATA

OpenPOWER READY

OpenPOWER™

ALPHA DATA
e: sales@alpha-data.com
w: www.alpha-data.com

XILINX®

CAPI 3.0 in AC922 (prototype)

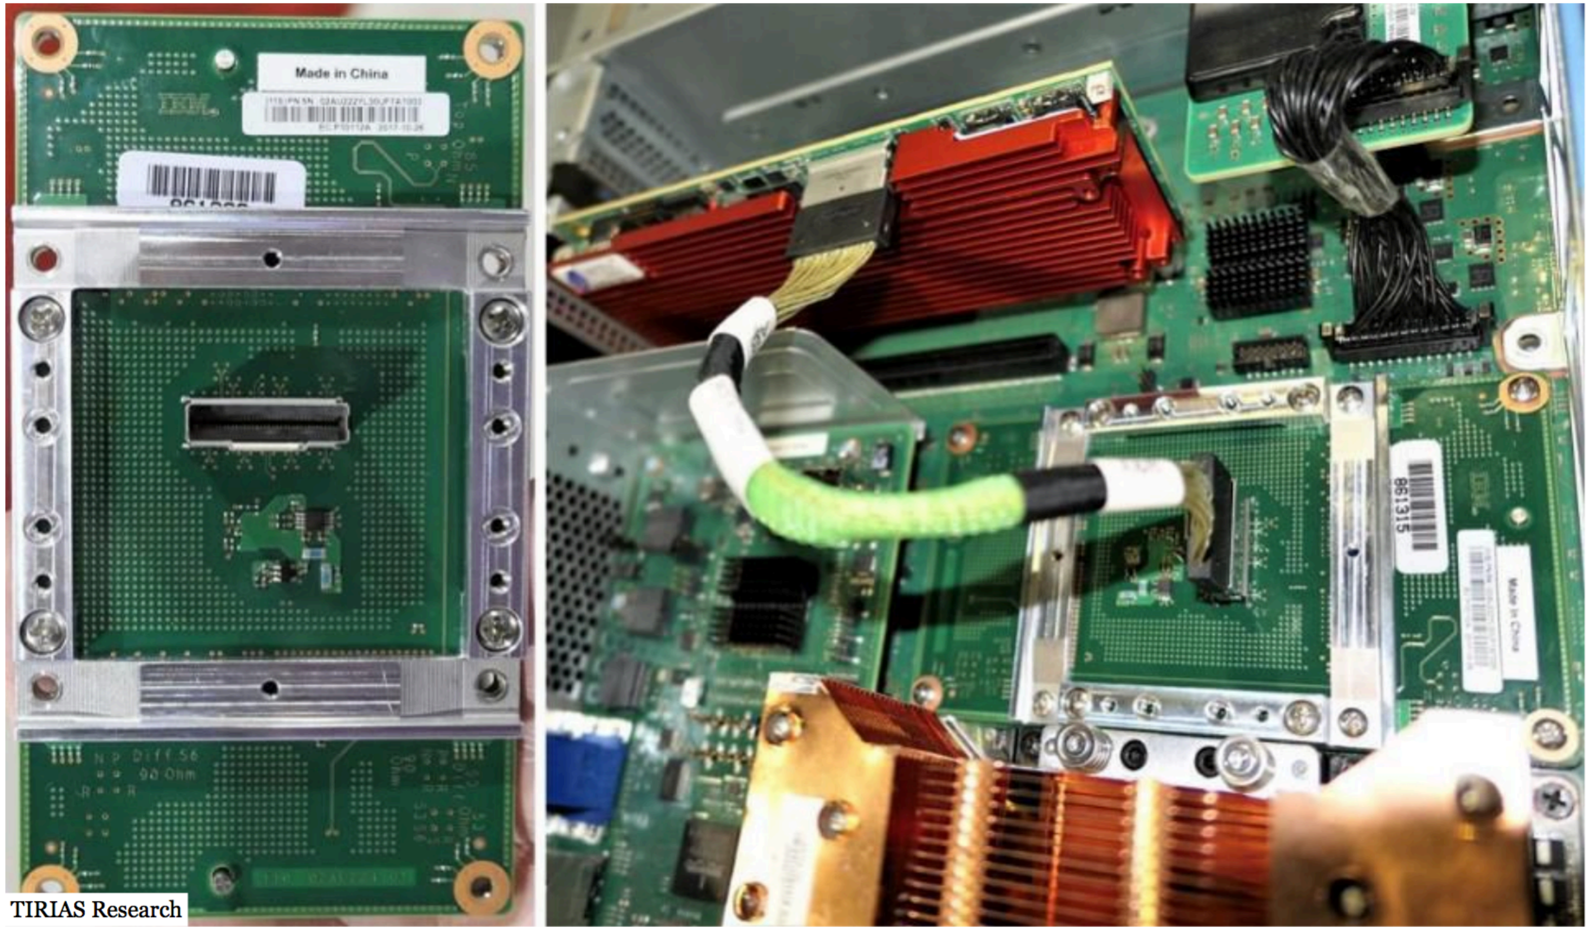

NVLink to OpenCAPI converter (left) and OpenCAPI attached Alpha Data ADM-PCIE-9V3 FPGA card (right) [+]

Source: Forbes

Networking w. OpenCAPI

Networking: Rackspace/Google Zaius OCP (1Tb/s)

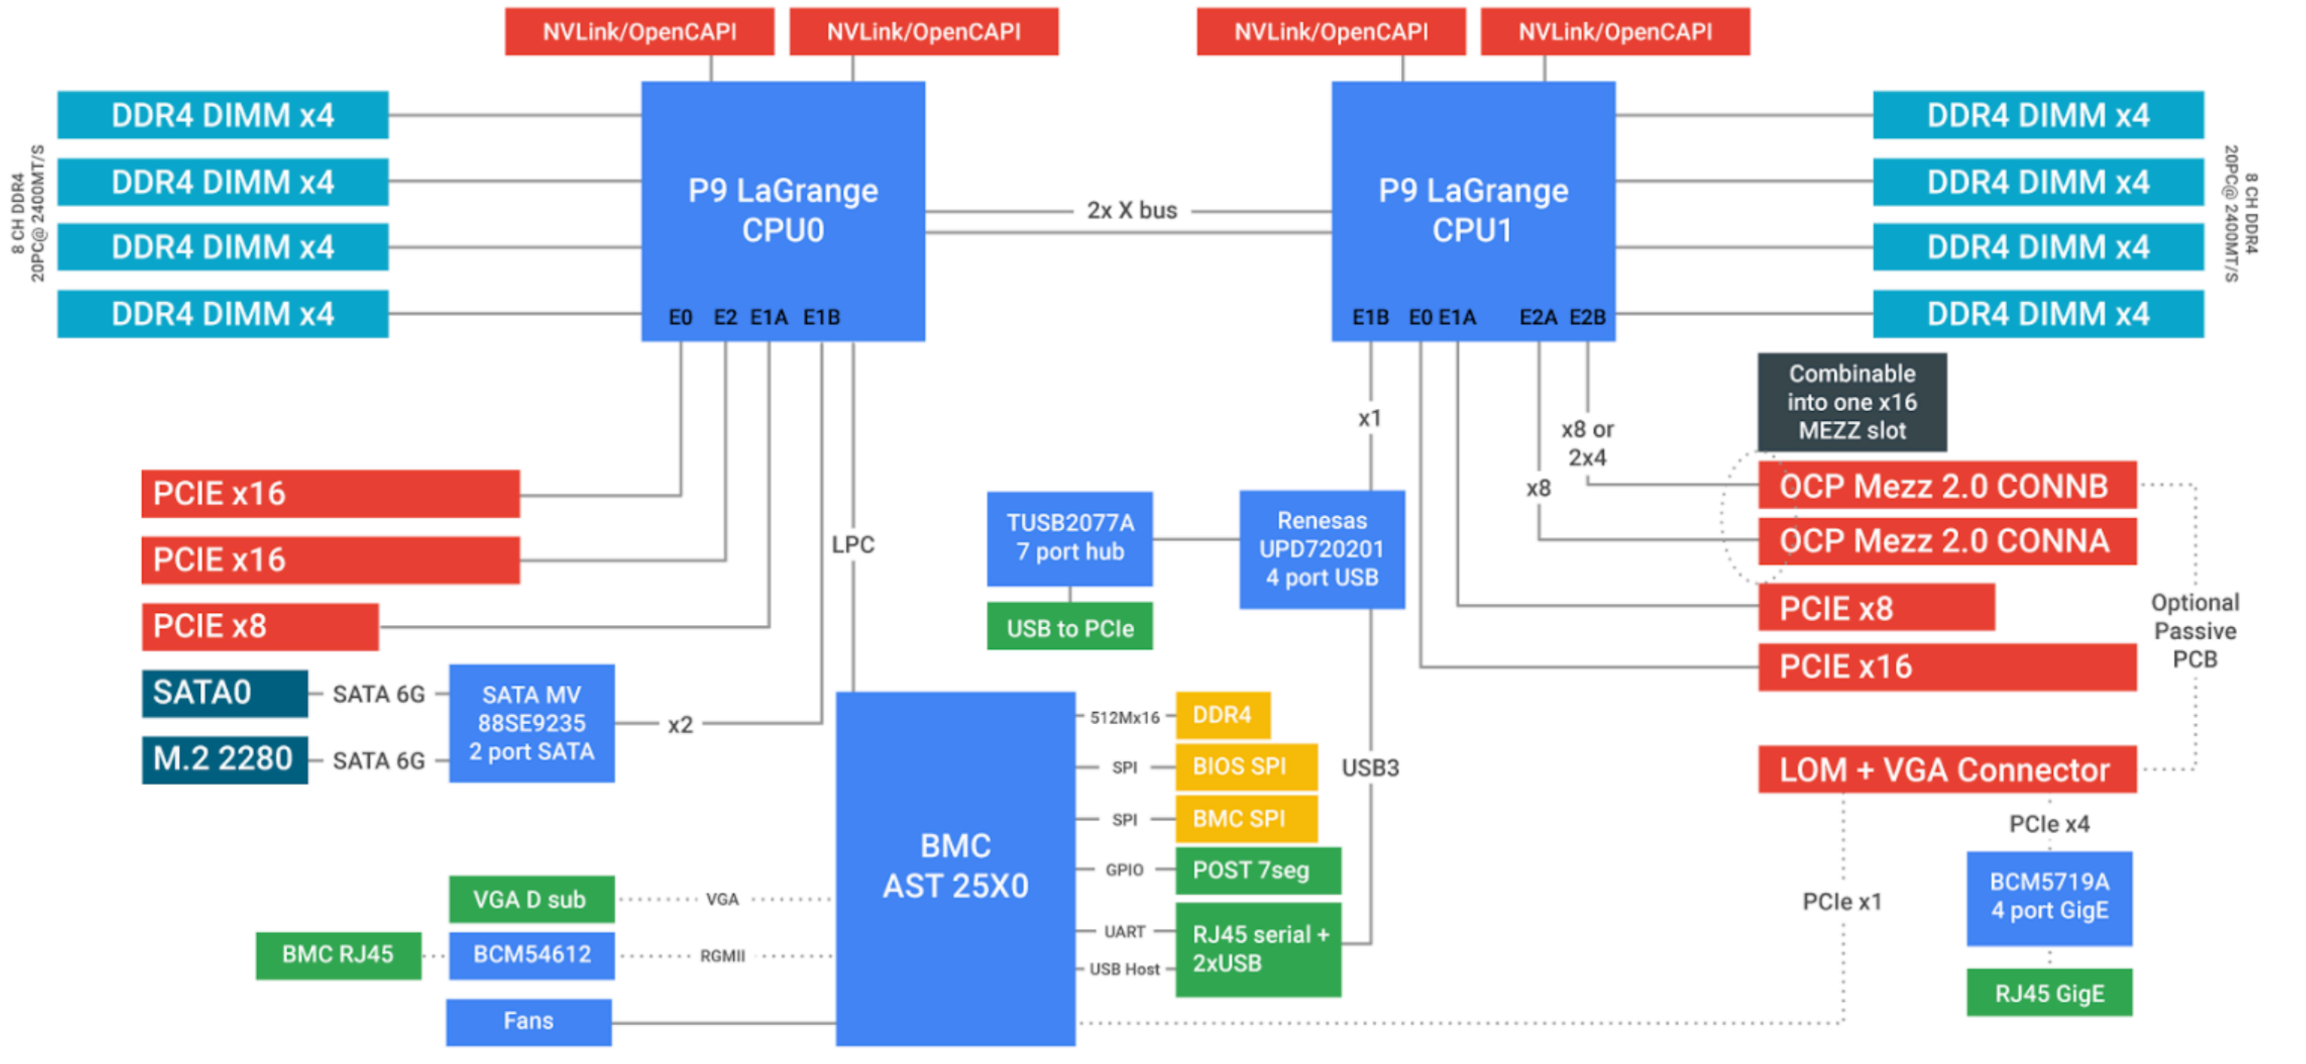

Zaius Block Diagram

Nallatech/IBM CAPI NVMe Flash Accelerator (FlashGT)

IBM

2016 Flash Adapter (CAPI 1.0)

- FPGA Controller
- 2x 960GB M.2 SSDs
- Supports Surelock KVS and Block APIs + Linux CAPI filesystem
- ~4x reduction in CPU overhead compared to NVME

Opportunities for Further Innovation:

- CAPI 2.0: Coherent Flash at PCIe Gen4 speed
- OpenCAPI: Extreme bandwidth & scaling of coherent flash
- Additional software exploitation

Available Now!

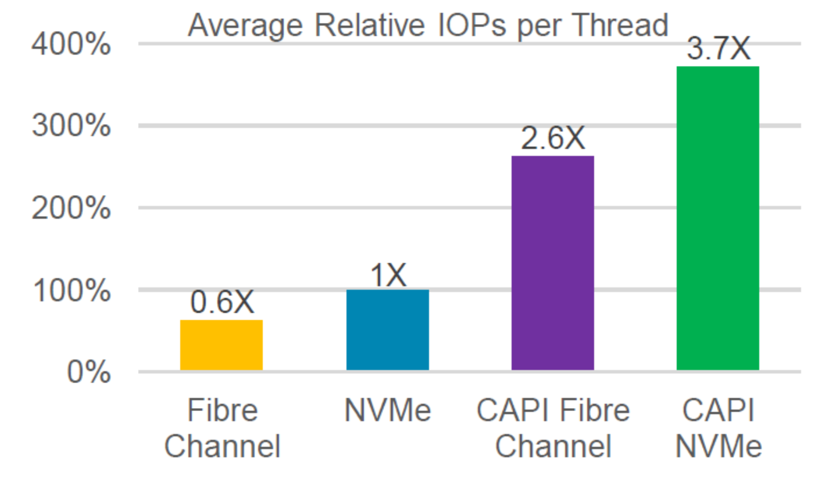

PCIe Gen4 / CAPI 2.0 to NVMe (2.9M IOPs)

IBM

Near Storage Accelerator

Shipping Now!

XILINX
ALL PROGRAMMABLE.

250S+

High-performance PCIe- based Flash SSD with localized FPGA acceleration capability

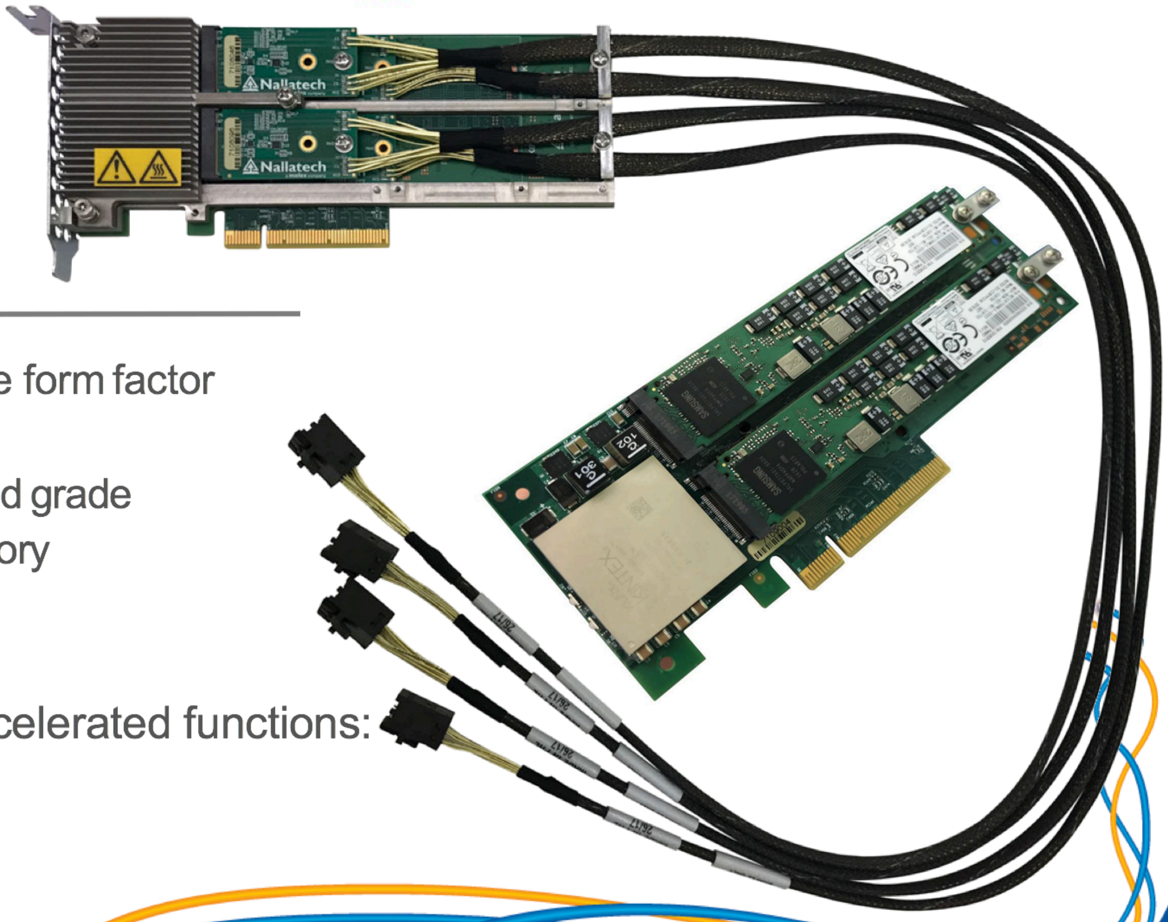

- **NIC** Form Factor - Low Profile, Half Length PCIe form factor
- PCIe Gen 3 or Gen 4 8-lane
- (1) **Xilinx** XCKU15P-2FFVA1156I FPGA-2 speed grade
- (1) bank of 4GByte 2400MTPS x80 DDR4 memory
- (4) M.2 connectors
- (4) M.2 to OCULink678mm cables
- Available pre-configured with NVMe-oF accelerated functions:
  - PCIe Gen 4 Host Bus Adaptor (HBA)

Nallatech
a molex company

FPGA Accelerated Computing

Molex "Sawmill"

CAPI/OpenCAPI STORAGE ACCELERATION

4x 20+GB/s !

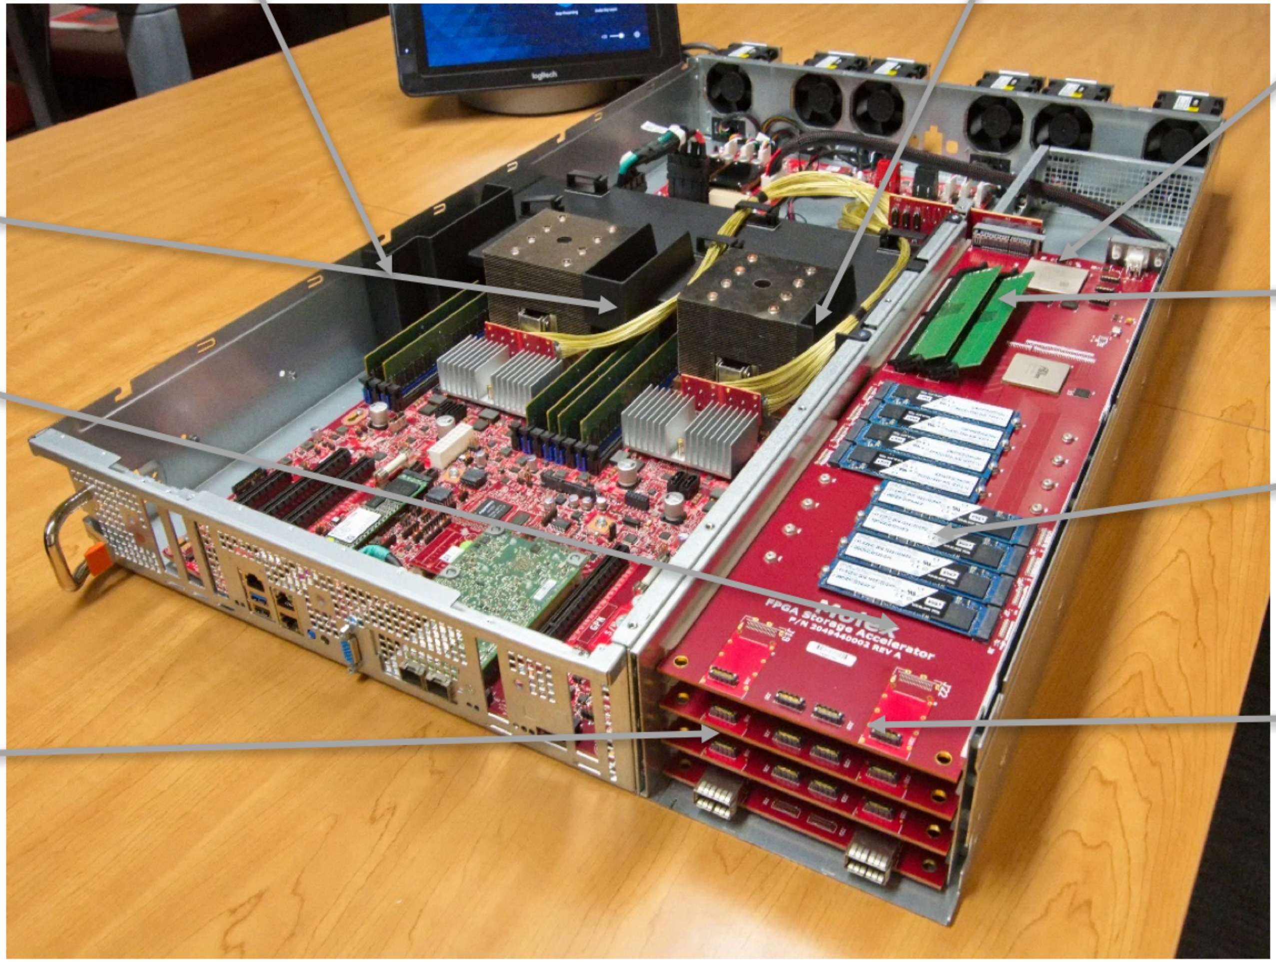

molex

5x Faster Data Communication with Unique CPU-GPU NVLink High-Speed Connection

Store Large Models in System Memory

Fast Transfer via NVLink

Operate on One Layer at a Time

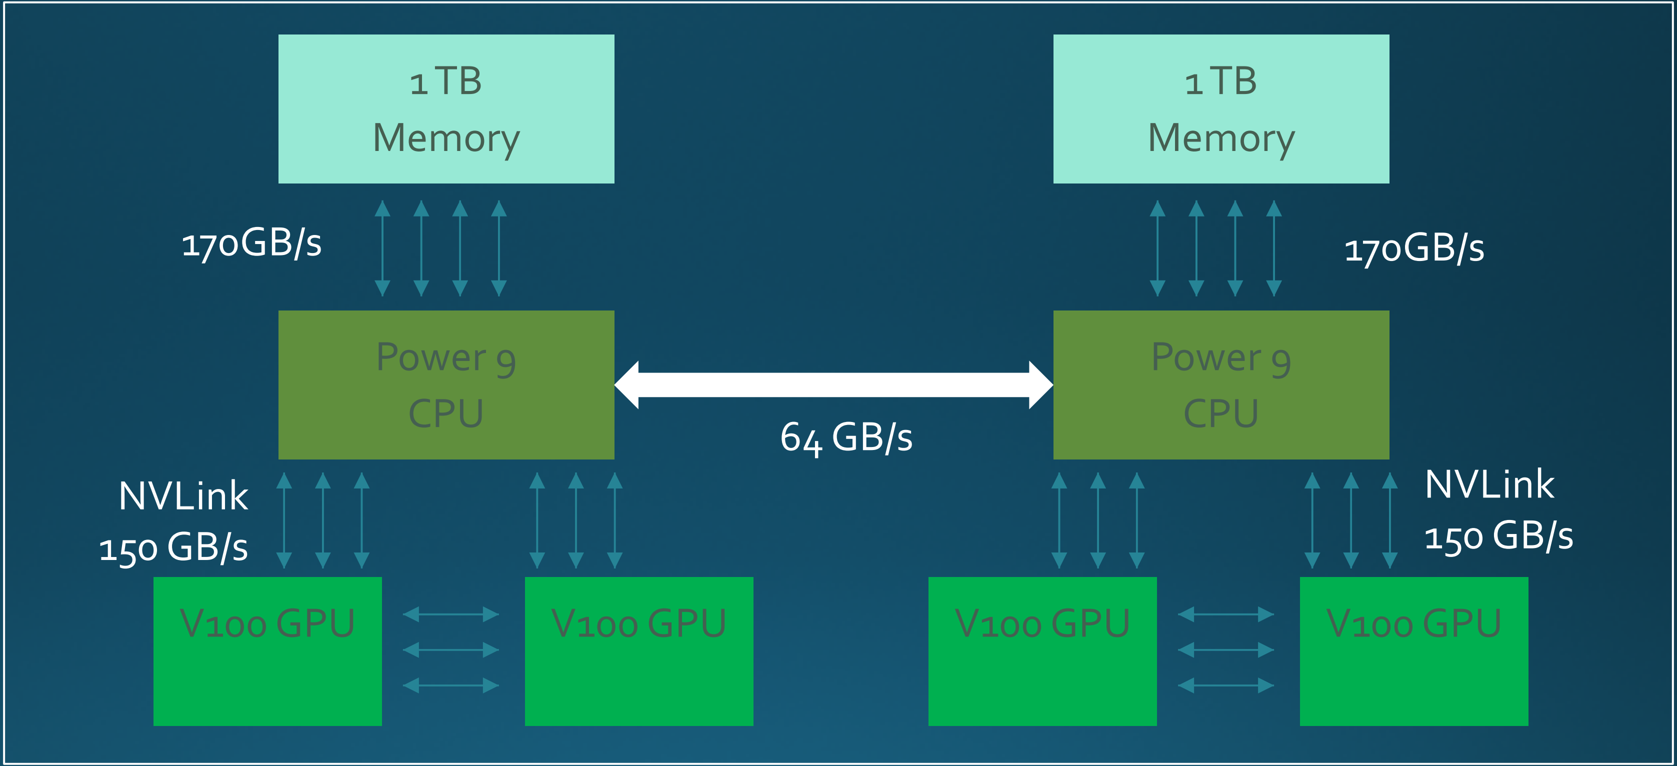

IBM AC922 Power System

Deep Learning Server (4-GPU Config)

AI at Unrivaled Scale:

Trusted as the building block for CORAL

Born of collaboration

The P9 architecture was developed by IBM, in collaboration with members of the OpenPOWER Foundation.

An AI Pioneer

CORAL in aggregate is likely to become the most powerful supercomputer in the world when completed. It’s on track to deliver 300+ PetaFlops of HPC and 3 ExaFlops of AI as a service performance.

Deploy your own Mini CORAL

An advantage of this collaborative approach is the repeatable building block which organizations will be able to leverage for the raw HPC horsepower and cutting-edge AI performance, in their own organizations.

AC922 HPC example:

2.9X faster running CPMD compared to tested x86 systems

- IBM Power System AC922 delivers **2.9X reduction in execution time** of tested x86 systems
- POWER9 with NVLink 2.0 unlocks the performance of GPU-accelerated version of CPMD by enabling lightning fast CPU-GPU data transfers
- **3.3TB of data movement required between CPU and GPU**
  - **70 seconds for NVLink 2.0 transfer time vs 300+ seconds for traditional PCIe bus transfer time**

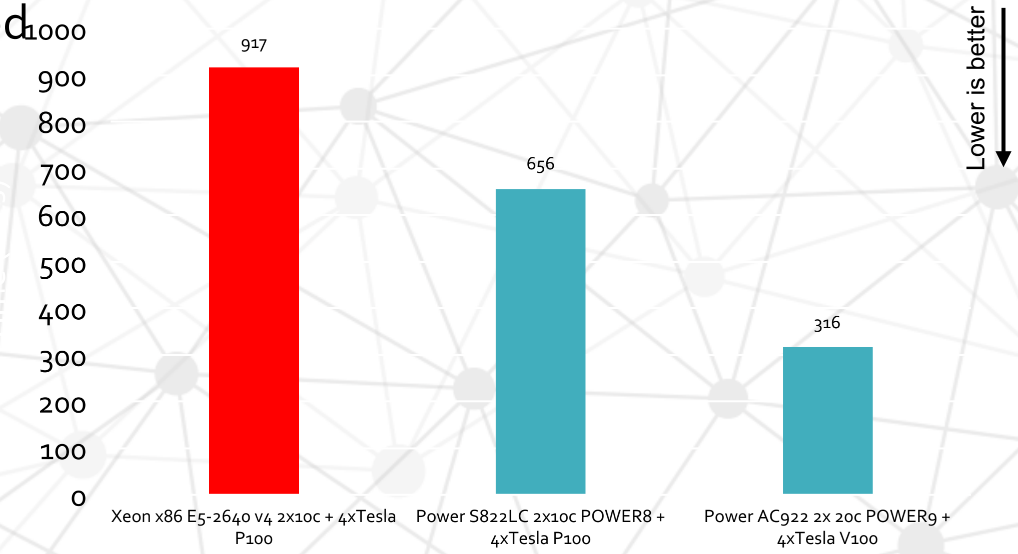

- All results are based on running CPMD, a parallelized plane wave / pseudopotential implementation of Density Functional Theory Application. A Hybrid version of CPMD (e.g. MPI + OPENMP + GPU + streams) was implemented with runs are made for 256-Water Box, RANDOM initialization. Results are reported in Execution Time (seconds).. Effective measured data rate on PCIe bus of 10 GB/s and on Nvlink 2.0 of 50GB/s.
- IBM Power AC922; 40 cores (2 x 20c chips), POWER9 with NVLink 2.0; 2.25 GHz, 1024 GB memory, 4xTesla V100 GPU; ; Red Hat Enterprise Linux 7.4 for Power Little Endian (POWER9) with ESSL PRPQ; Spectrum MPI: PRPQ release, XLF: 15.16, CUDA 9.1
- IBM Power System S822LC for HPC; 20 cores (2 x 10c chips) / 160 threads, POWER8 with NVLink; 2.86 GHz, 256 GB memory, 2 x 1TB SATA 7.2K rpm HDD, 2-port 10 GbEth, 4xTesla P100 GPU; RHEL 7.4.with ESSL 5.3.2.0; PE2.2; XLF: 15.1, CUDA 8.0
- 2x Xeon E5-2640 v4; 20 cores (2 x 10c chips) / 40 threads; Intel Xeon E5-2640 v4; 2.4 GHz; 256 GB memory, 1 x 2TB SATA 7.2K rpm HDD, 2-port 10 GbEth; , 4xTesla P100 GPU; Ubuntu 16.04 with OPENBLAS 0.2.18, OpenMPI: 1.10.2, GNU-5.4.0, CUDA-8.0

Big Data: Sorting Large Datasets: sortbenchmark.org

Top Results

	Daytona	Indy
Gray	2016, 44.8 TB/min **Tencent Sort** 100 TB in 134 Seconds 512 nodes x (2 OpenPOWER 10-core POWER8 2.926 GHz, 512 GB memory, 4x Huawei ES3600P V3 1.2TB NVMe SSD, 100Gb Mellanox ConnectX4-EN) Jie Jiang, Lixiong Zheng, Junfeng Pu, Xiong Cheng, Chongqing Zhao Tencent Corporation Mark R. Nutter, Jeremy D. Schaub	2016, 60.7 TB/min **Tencent Sort** 100 TB in 98.8 Seconds 512 nodes x (2 OpenPOWER 10-core POWER8 2.926 GHz, 512 GB memory, 4x Huawei ES3600P V3 1.2TB NVMe SSD, 100Gb Mellanox ConnectX4-EN) Jie Jiang, Lixiong Zheng, Junfeng Pu, Xiong Cheng, Chongqing Zhao Tencent Corporation Mark R. Nutter, Jeremy D. Schaub
Cloud	2016, $1.44 / TB **NADSort** 100 TB for $144 394 Alibaba Cloud ECS ecs.n1.large nodes x (Haswell E5-2680 v3, 8 GB memory, 40GB Ultra Cloud Disk, 4x 135GB SSD Cloud Disk) Qian Wang, Rong Gu, Yihua Huang Nanjing University Reynold Xin Databricks Inc. Wei Wu, Jun Song, Junluan Xia Alibaba Group Inc.	2016, $1.44 / TB **NADSort** 100 TB for $144 394 Alibaba Cloud ECS ecs.n1.large nodes x (Haswell E5-2680 v3, 8 GB memory, 40GB Ultra Cloud Disk, 4x 135GB SSD Cloud Disk) Qian Wang, Rong Gu, Yihua Huang Nanjing University Reynold Xin Databricks Inc. Wei Wu, Jun Song, Junluan Xia Alibaba Group Inc.
Minute	2016, 37 TB **Tencent Sort** 512 nodes x (2 OpenPOWER 10-core POWER8 2.926 GHz, 512 GB memory, 4x Huawei ES3600P V3 1.2TB NVMe SSD, 100Gb Mellanox ConnectX4-EN) Jie Jiang, Lixiong Zheng, Junfeng Pu, Xiong Cheng, Chongqing Zhao Tencent Corporation Mark R. Nutter, Jeremy D. Schaub	2016, 55 TB **Tencent Sort** 512 nodes x (2 OpenPOWER 10-core POWER8 2.926 GHz, 512 GB memory, 4x Huawei ES3600P V3 1.2TB NVMe SSD, 100Gb Mellanox ConnectX4-EN) Jie Jiang, Lixiong Zheng, Junfeng Pu, Xiong Cheng, Chongqing Zhao Tencent Corporation Mark R. Nutter, Jeremy D. Schaub
Joule 10^{10} recs	2013, 168,242 Joules **NTOSort** 59,444 records sorted / joule Intel i7-3770K, 16GB RAM, Nsort, Windows 8, 16 Samsung 840 Pro 256GB SSDs, 1 Samsung 840 Pro 128GB SSD Andreas Ebert Microsoft	2013, 168,242 Joules **NTOSort** 59,444 records sorted / joule Intel i7-3770K, 16GB RAM, Nsort, Windows 8, 16 Samsung 840 Pro 256GB SSDs, 1 Samsung 840 Pro 128GB SSD Andreas Ebert Microsoft

Dual Socket
POWER8
100 TB
~100 sec
~500 systems
100Gb/s network
4x NVMe

~2 TB/s/node

Recent work: w. G. Fossum & T. Wang, IBM
IBM AC922

- One node, 4GPU, from memory
  - partitioner > 40GB/s
  - partition & sort > 20GB/s
- To achieve same on a cluster
  - Network must match 2nd (or 1st) phase throughput
  - i.e. 50GB/s … 400Gb/s
  - Should be no problem
    - 2x 200Gb adapter (x16 PCIe Gen4 or CAPI 2.0)
- Sortbenchmark.org rules require read/write to persistent store
  - Even that should be doable
  - 32x NVMe solution from Nallatech/Molex
- Net: 10x per node is in the cards …
  - Beat current 512 systems with 64 systems

Large AI Models Train ~4 Times Faster

POWER9 Servers with NVLink to GPUs
vs
x86 Servers with PCIe to GPUs

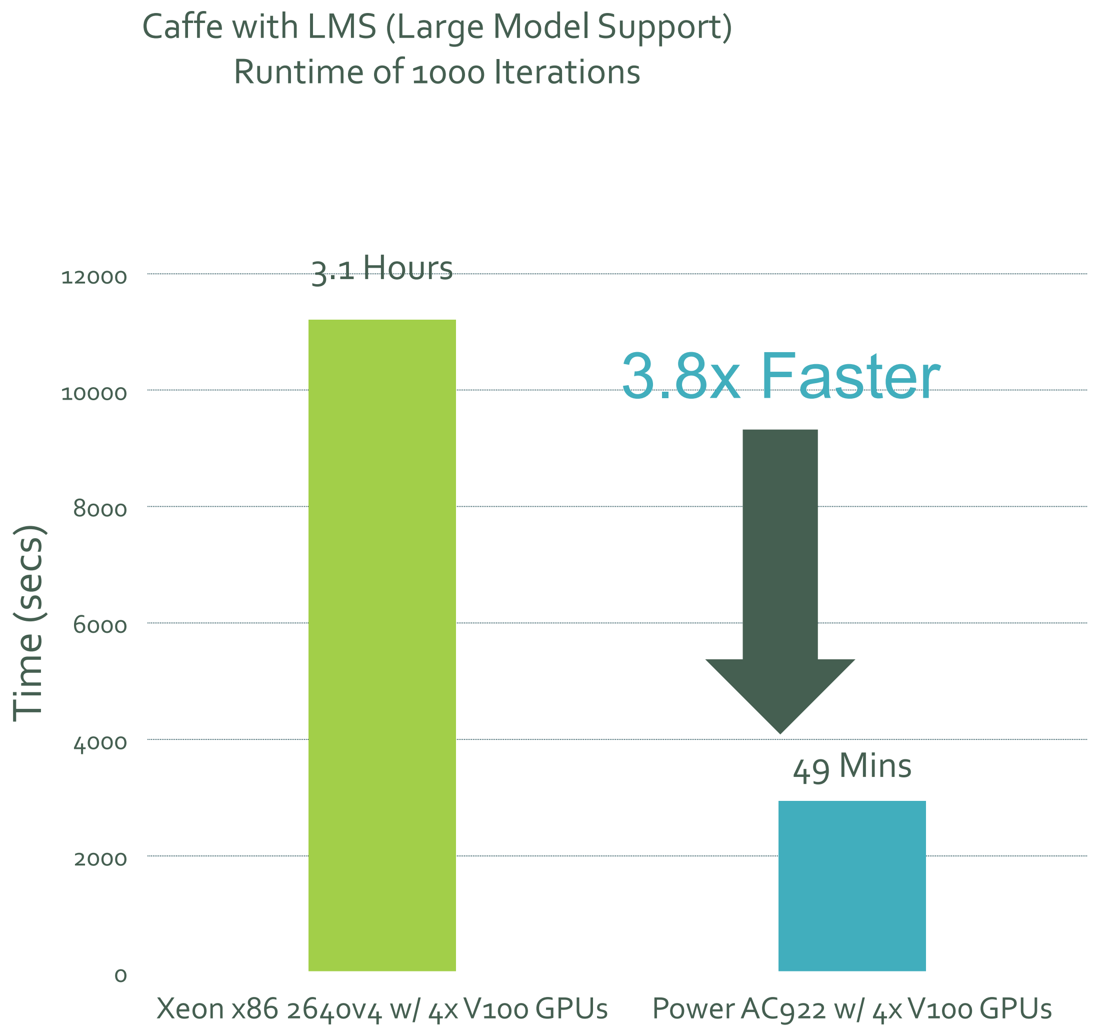

GoogleNet model on Enlarged ImageNet Dataset (2240x2240)

Distributed Deep Learning (DDL)

- Deep learning training takes days to weeks
- Limited scaling to
- multiple x86 servers
- PowerAI with DDL enables scaling to 100s of servers

16 Days Down to 7 Hours

58x Faster

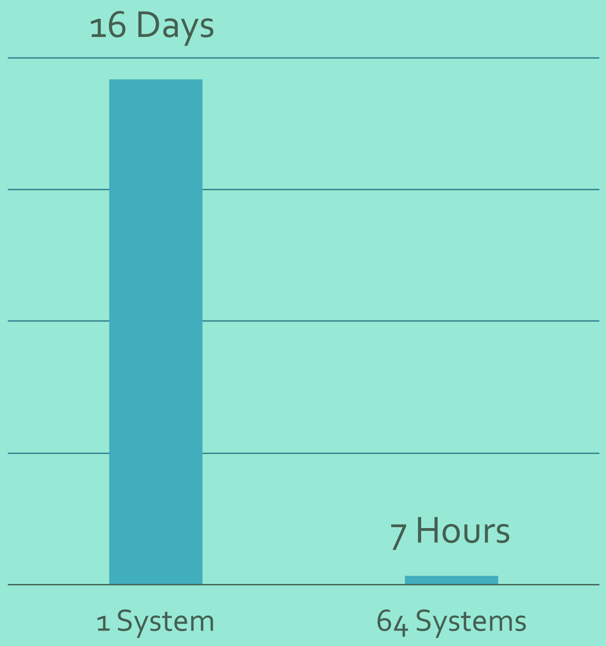

ResNet-101, ImageNet-22K

Near Ideal Scaling to 256 GPUs

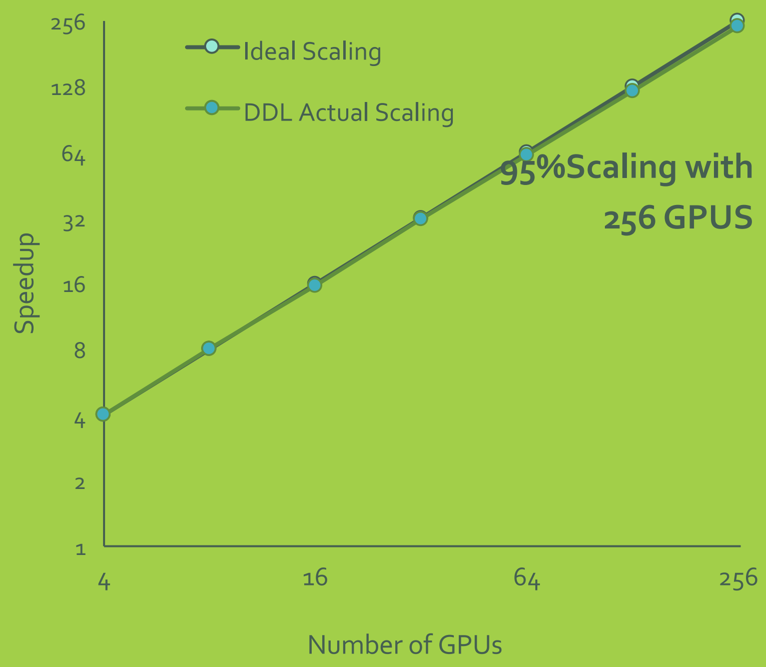

ResNet-50, ImageNet-1K

Caffe with PowerAI DDL, Running on Minsky (S822Lc) Power System

Tera-scale Computational Advertising Application

Criteo Releases Industry's Largest-Ever Dataset for Machine Learning to Academic Community

New York – June 18, 2015 – Criteo (NASDAQ: CRTO), the performance marketing technology company, today announced the release of the largest public machine learning dataset ever issued to the open source community, with the goal of supporting academic research and innovation in distributed machine learning algorithms.

* Criteo Labs. 2015. Criteo Releases Industry□s Largest-Ever Dataset for Machine Learning to Academic Community. h□ps://www.criteo.com/news/press-releases/2015/07/criteo-releases-industrys-largest-ever-dataset/

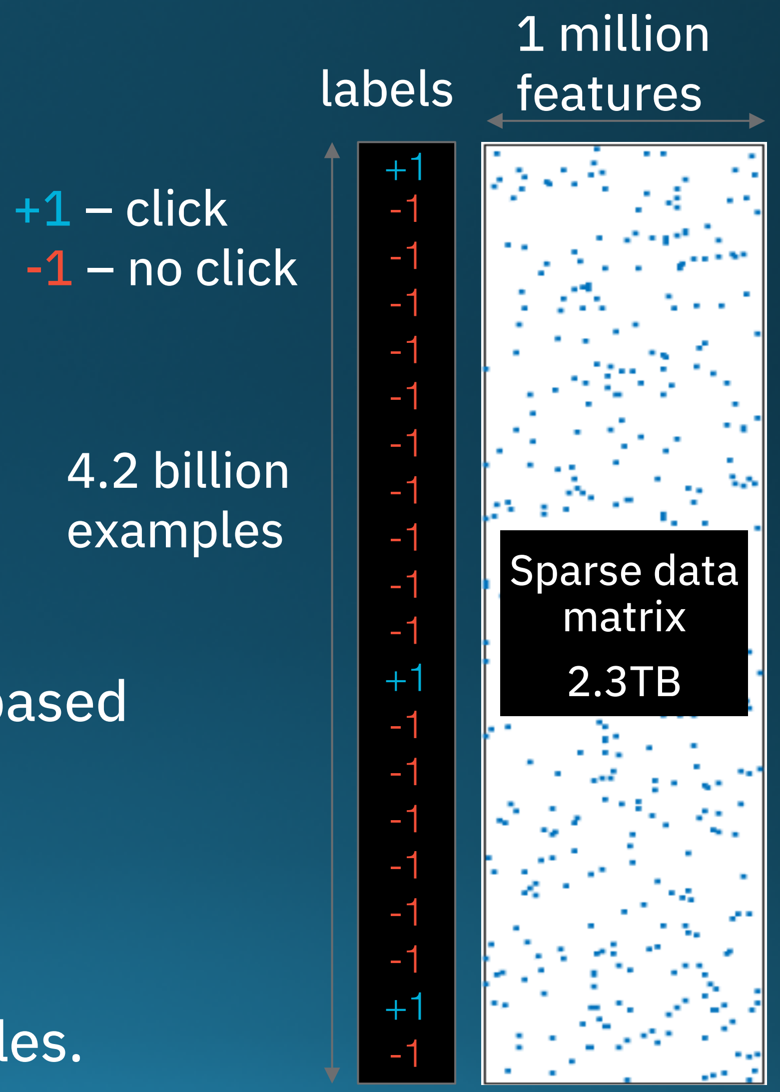

Goal: Predict whether a user will click on a given advert based on an anonymized set of features.

Train: Fit model parameters using **4.2 billion** examples.

Inference: Evaluate model on **180 million** unseen examples.

SNAP ML: 3 Key Breakthroughs

GPU Acceleration

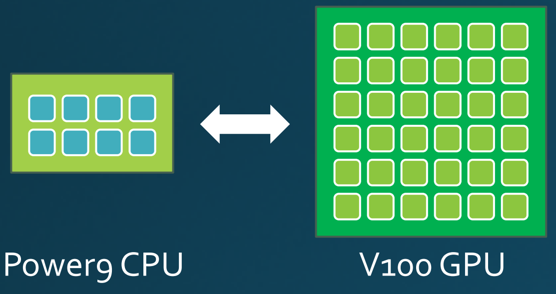

Dynamic Optimized Memory Management

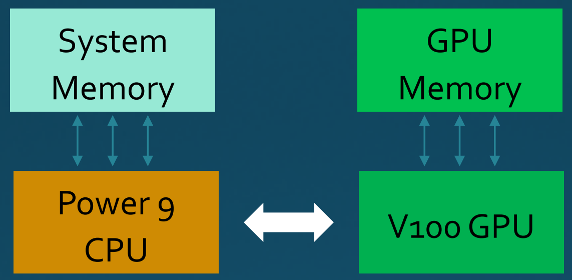

Efficient Cluster Scaling

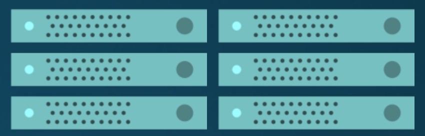

- C. Duenner, S. Forte, M. Takac, and M. Jaggi. "Primal-Dual Rates and Certificates." In *International Conference on Machine Learning (ICML 2016)*, pp. 783-792. 2016.
- T. Parnell, C. Duenner, K. Atasu, M. Sifalakis and H. Pozidis, "Large-scale stochastic learning using GPUs," *2017 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, Lake Buena Vista, FL, 2017, pp. 419-428.
- C. Duenner, T. Parnell, K. Atasu, M. Sifalakis and H. Pozidis, “Understanding and Optimizing the Performance of Distributed Machine Learning Applications on Apache Spark”, *poster presentation at NIPS 2016 ML Systems workshop, IEEE Big Data 2017*
- C. Duenner, T. Parnell, M. Jaggi, “Efficient Use of Limited-Memory Resources to Accelerate Linear Learning”, proceedings of 2017 *Neural Information Processing Systems (NIPS 2017)*

snap ML: Tera-scale ML benchmark

Criteo Terabyte Click Logs Benchmark

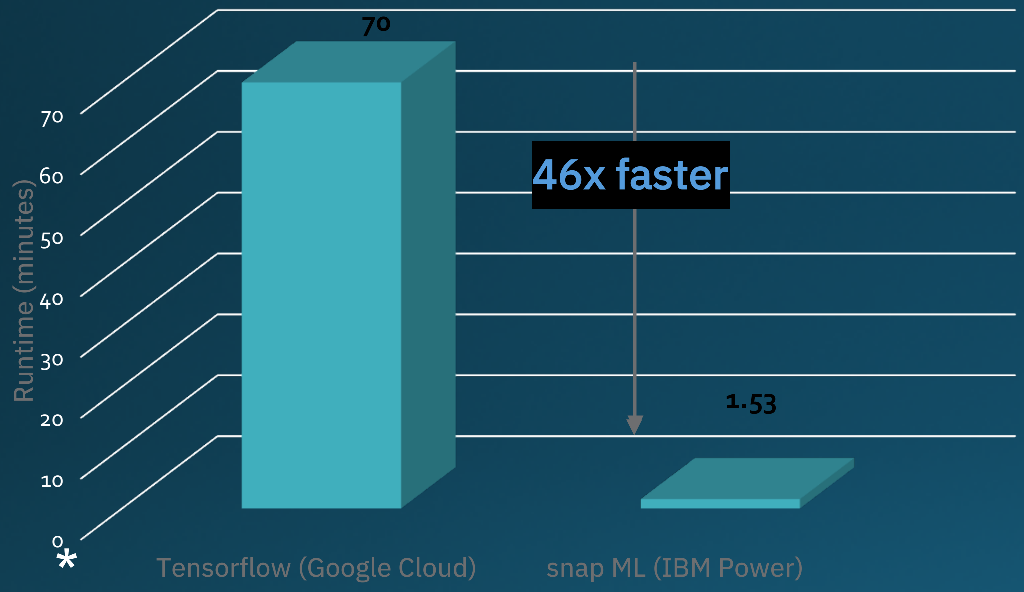

* https://cloud.google.com/blog/big-data/2017/02/using-google-cloud-machine-learning-to-predict-clicks-at-scale

Comparison of Tensorflow** on Google Cloud with SNAP ML on POWER9* (AC922) cluster

Workload: Click-through-rate prediction for computational advertising, using Logistic Regression

Dataset: Criteo Terabyte Click Logs (http://labs.criteo.com/2013/12/download-terabyte-click-logs/)

Dataset: 4.2 billion training examples, 1 million features
Model: Logistic Regression
Test LogLoss: 0.1293 (Tensorflow), 0.1292 (snap ML)
Platform: 89 machines (Tensorflow),
8 Power9 CPUs+16 NVIDIA® Tesla™ V100 GPUs (snap ML)

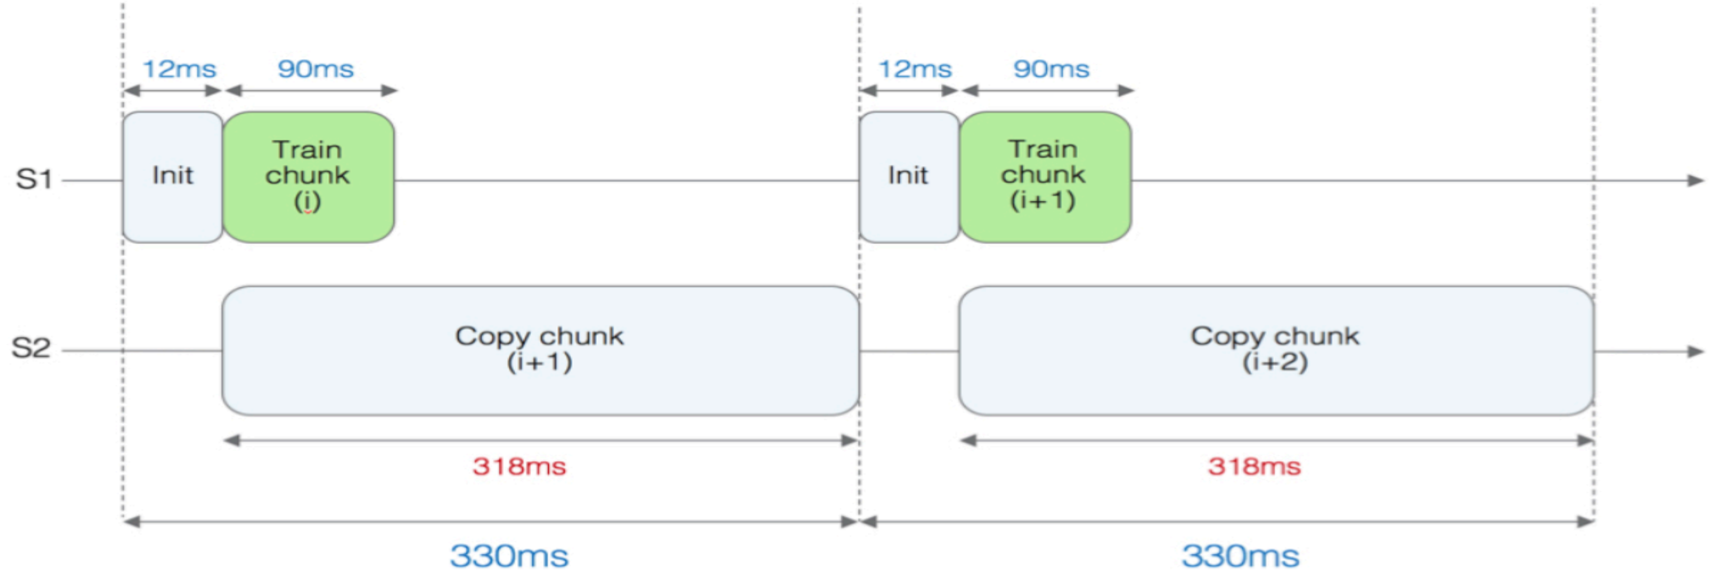

(a) Runtime profile on Intel x86 + PCIe Gen 3.0

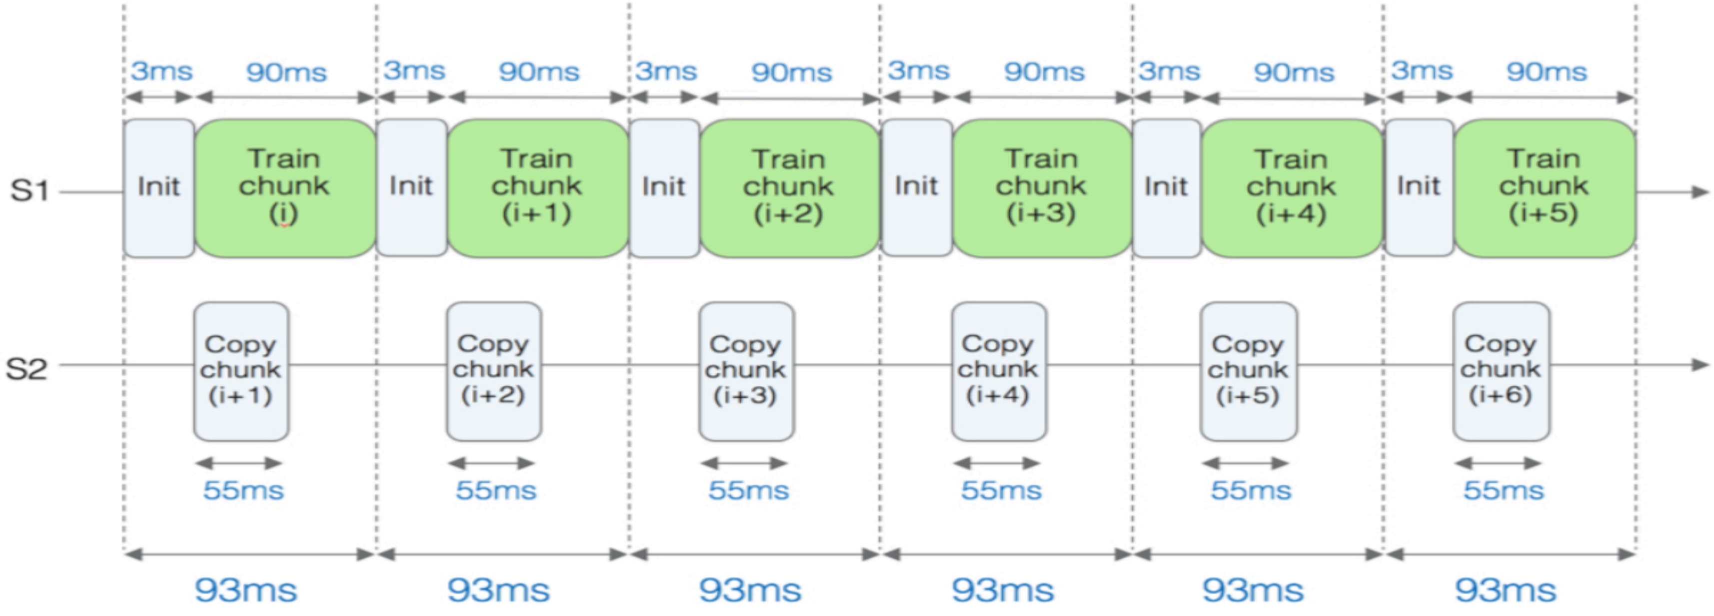

(b) Runtime profile on POWER9 + NVLINK 2.0

FPGA Acceleration & Architectural Exploration

- “Completed” example
  - Gzip – prototyped (FPGA) on P7, productized on P8 (FPGA), integrated in P9

- Some current examples
  - “Fletcher” an open source frameworks for processing Arrow files with FPGAs
    J. Peltenburg e.a., TU Delft, Netherlands
  - 16 Gpop, 128x128 “32b posit” matrix multiply,
    J. Chen e.a., TU Delft, Netherlands

Old Way   Apache Arrow & Fletcher

IBM

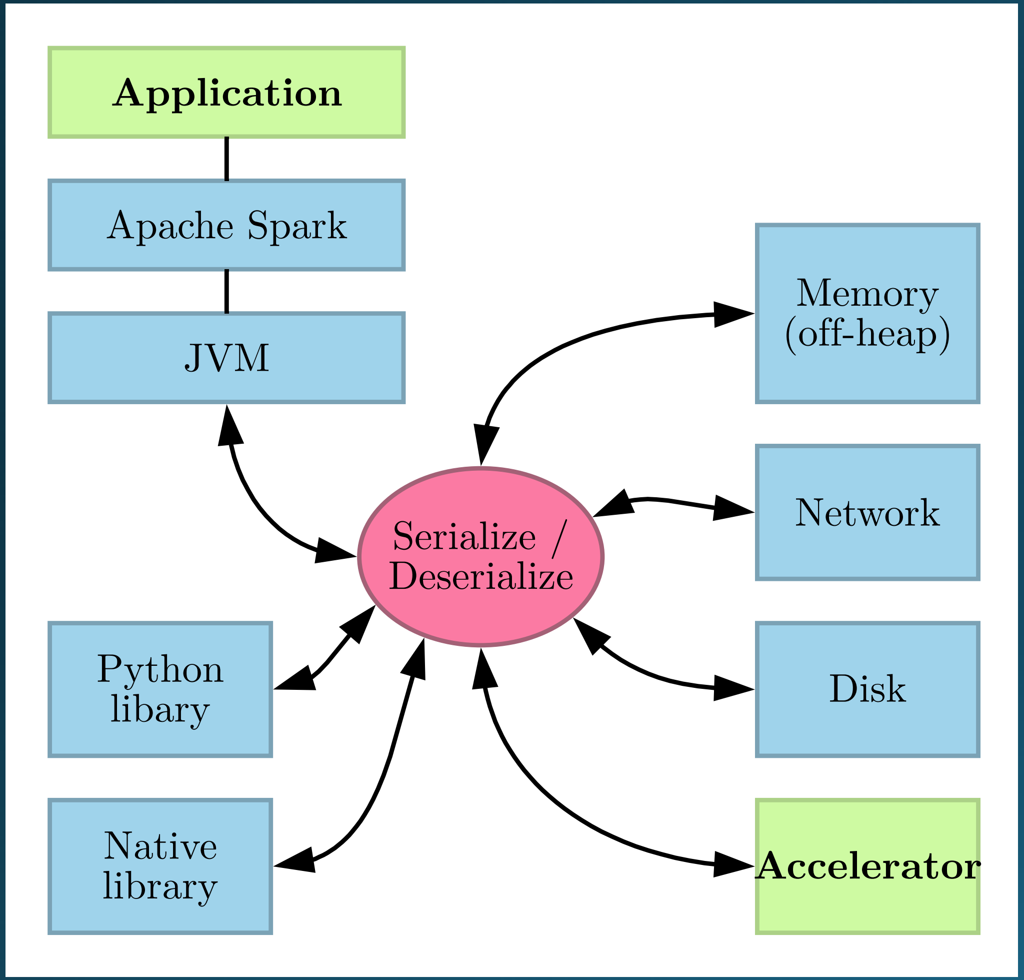

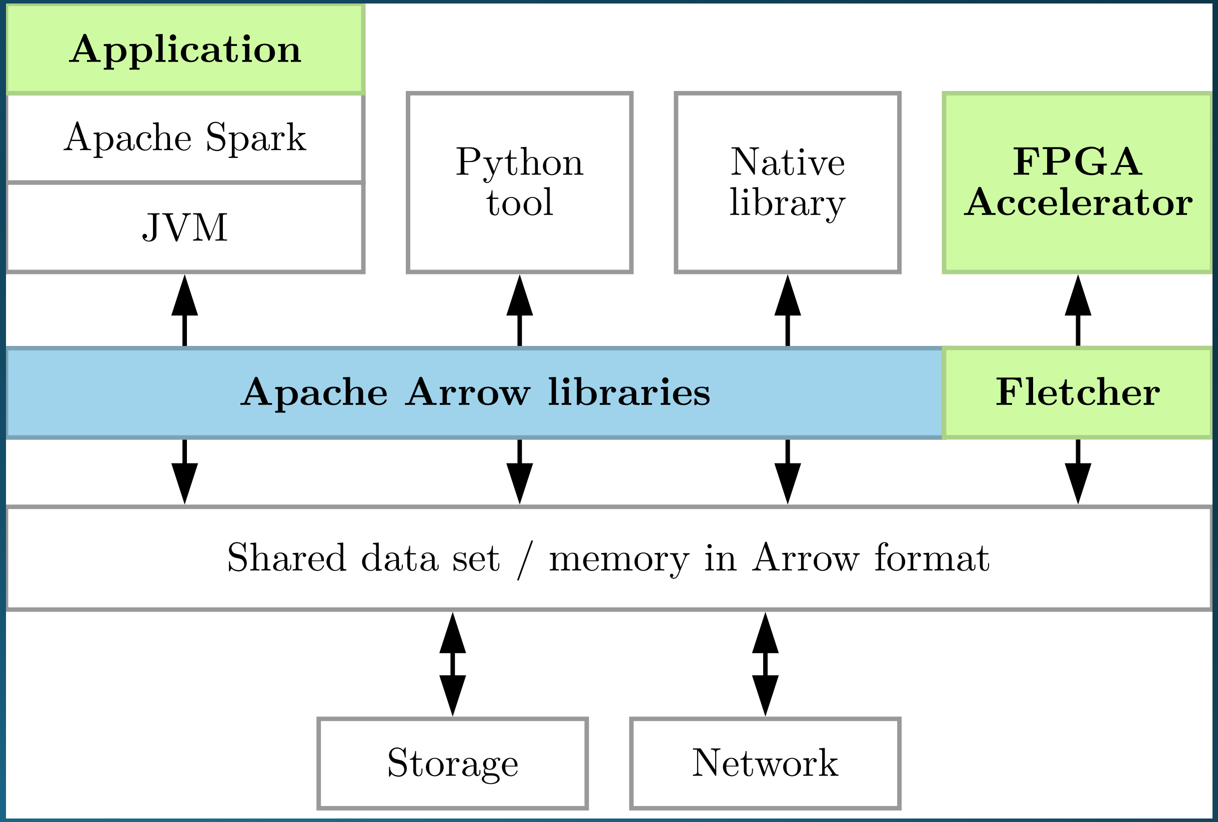

J. Peltenburg, e.a., TU Delft (OpenPOWER Summit USA 2018)

Regular expression matching

R=16 different regular expressions per unit

AWS EC2 F1:
- Virtex Ultrascale+
- N=16 regex units
- 256 regexes being matched in parallel

POWER8 CAPI (Supervessel, & soon at Nimbix):
- AlphaData KU3 (Kintex Ultrascale)
- N=8 regex units
- 128 regex being matched in parallel

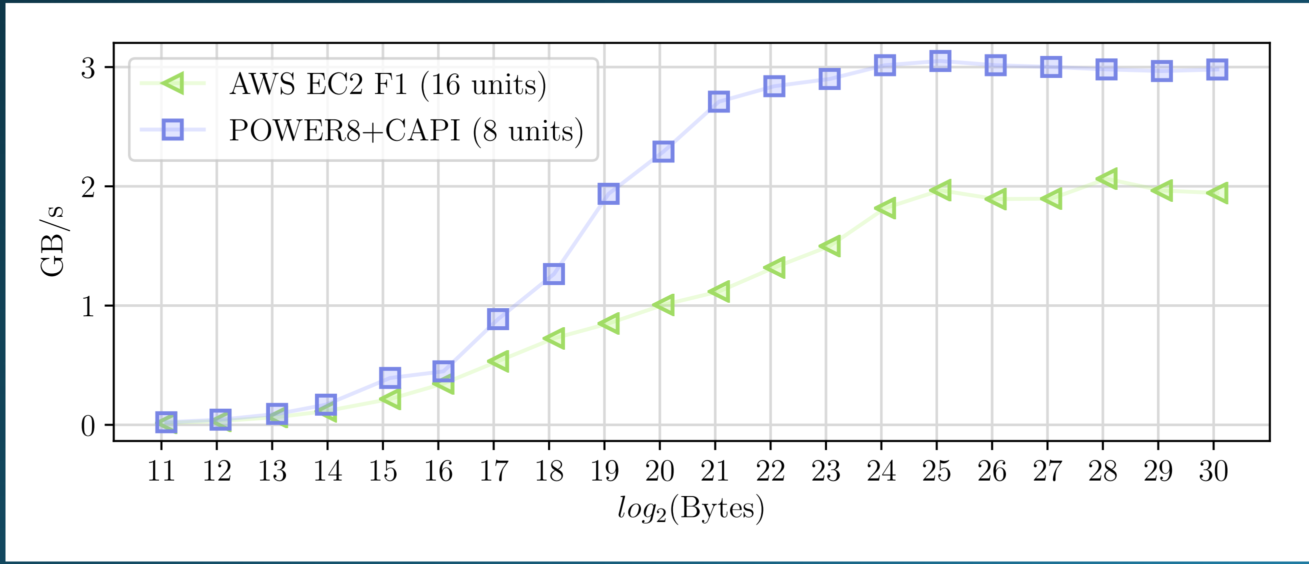

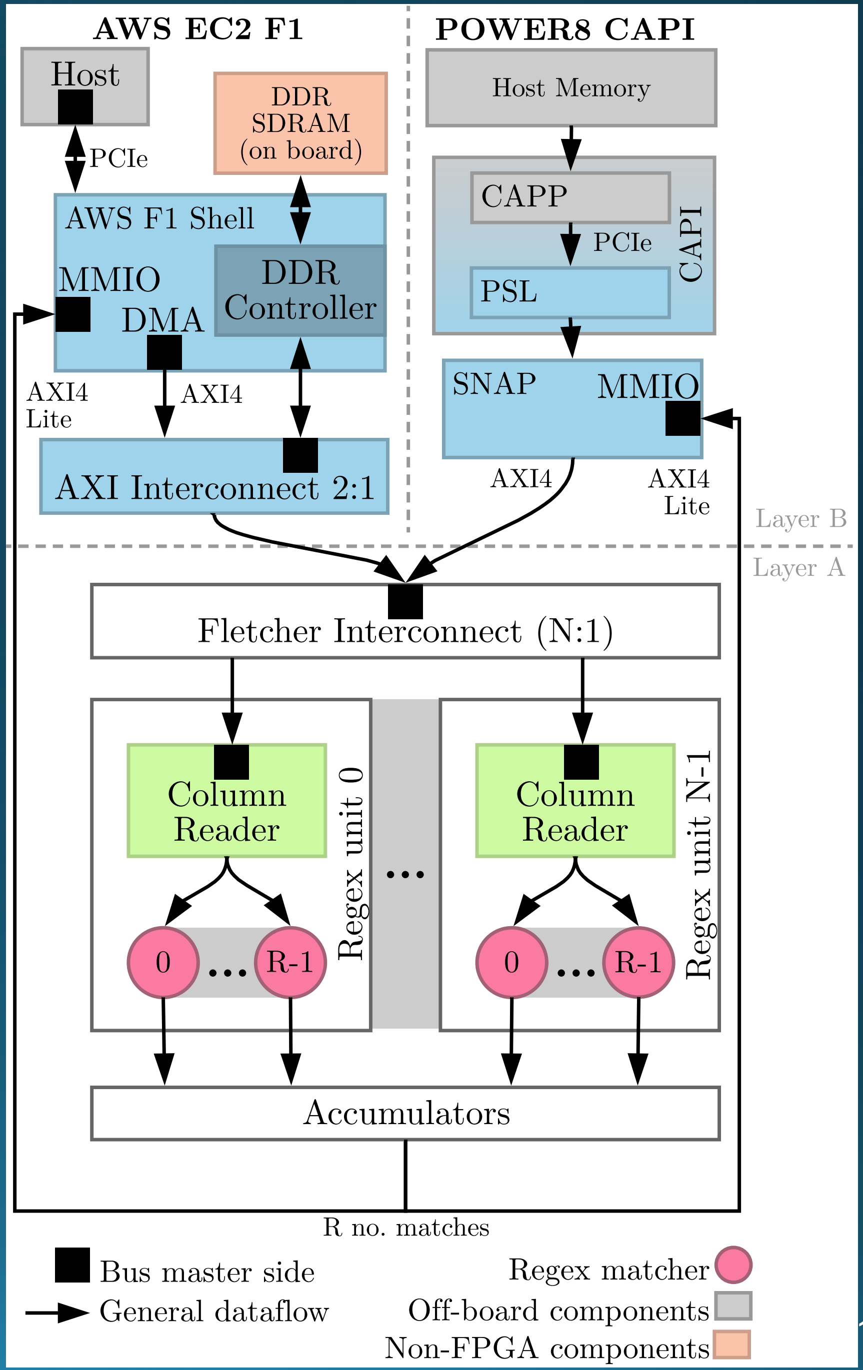

Posit Matrix-Multiply

- New proposed format for floating-point by Dr. J. Gustafson
  - Fixed length representation, but variable length mantissa and exponent

- Nicer properties than conventional floating-point
  - Symmetry, overflow, …
  - Often more accuracy with fewer bits

- Built a CAPI 1.0 matrix-multiply unit (Jianyu Chen e.a., TU Delft)
  - Uses wide "quire" register (accumulator) for dot products
  - Just pass pointer to matrix A, B, C and array dimensions
    - CAPI accelerator has full access to (effective/virtually addressed) host memory

- 16 Gpops (streaming 128x128 MMuls, CAPI 1.0 AlphaData '7V3)
  - Accessible free to academics at TACC (USA), working to get next one in Singapore
  - Should scale to 32 & 64 Gpops CAPI 2.0 & OpenCAPI

- Next step is to use for application studies
  - Let me (or Dr. Gustafson) know if you're interested!

Conclusions

- It's about more than the CPU cores
  - Even though POWER9 cores are very good too!
- Investment in IO & OpenPOWER collaborations pays off
  - Better acceleration – better BW, latency, CPU utilization with GPU & FPGA
  - Better networking – better BW (1Tb/s demo), lower latency, lower CPU
  - Better storage – better BW, lower latency, lower CPU
- Use examples:
  - HPC – Coral system
  - Big Data – sort (10x per node of current sortbenchmark.org leader)
  - AI – large models (3.5-4x faster on large models)
- Architectural exploration:
  - Posits
  - Arrow/Fletcher

Legal notices

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 1 0504- 785
U.S.A.

Information and trademarks

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

This presentation and the claims outlined in it were reviewed for compliance with US law. Adaptations of these claims for use in other geographies must be reviewed by the local country counsel for compliance with local laws.

Special notices

This document was developed for IBM offerings in the United States as of the date of publication. IBM may not make these offerings available in other countries, and the information is subject to change without notice. Consult your local IBM business contact for information on the IBM offerings available in your area.

Information in this document concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. Send license inquires, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

The information contained in this document has not been submitted to any formal IBM test and is provided "AS IS" with no warranties or guarantees either expressed or implied.

All examples cited or described in this document are presented as illustrations of the manner in which some IBM products can be used and the results that may be achieved. Actual environmental costs and performance characteristics will vary depending on individual client configurations and conditions.

IBM Global Financing offerings are provided through IBM Credit Corporation in the United States and other IBM subsidiaries and divisions worldwide to qualified commercial and government clients. Rates are based on a client's credit rating, financing terms, offering type, equipment type and options, and may vary by country. Other restrictions may apply. Rates and offerings are subject to change, extension or withdrawal without notice.

IBM is not responsible for printing errors in this document that result in pricing or information inaccuracies.

All prices shown are IBM's United States suggested list prices and are subject to change without notice; reseller prices may vary.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Some measurements quoted in this document may have been estimated through extrapolation. Users of this document should verify the applicable data for their specific environment.