# Future Processor Architectures

Chris Green

DAQ R&D Workshop

Wednesday, October 11, 2017

# Goal in terms of workshop

- Session: Promising Technologies and Techniques
- From charge item (4):
  - "To identify future trends in electronics, computing, and other data acquisition technologies that may be promising areas of research for use in future experiments and ongoing detector research."
- This talk addresses item (4) in the charge
  - "ideally present the breadth of possibilities within the specified topic area."
- Covering roadmap topic areas

# Outline

- Focus and motivation
- Current technology drivers and directions
- Survey of current and upcoming technologies
- Conclusion

# Focus and motivation

- Triggering / filtering within time constraints will continue to be a demanding challenge, but much HPC-caliber hardware is already commodity.
- Event-building on commodity hardware is already happening: high-speed, low latency local networking is already here, affordable and only getting better.
- Integrated, heterogeneous systems improving rapidly – PFLOP racks.
- Moore's Law still (mostly) good for a few years (5nm limit in ~2020) but Dennard scaling has been dead for >10y. Many novel uses for all those transistors which will require planning and ingenuity to take advantage of: need to start now!
- What will be available in 10+ years and how to plan to use it?

# What is driving changes in computing architecture?

- DOE's ASCR program looks to be driving a good amount of future technology
- This recent slide provides a good summary of computing drivers

## What does the Future Hold: Strategic Vision for ASCR's Research Program

Emerging trends are pointing to a future that is increasingly

1. **Instrumented**: Sensors, satellites, drones, offline repositories
2. **Interconnected**: Internet of Things, composable infrastructure, heterogeneous resources
3. **Automated**: Complexity, real-time, machine learning
4. **Accelerated**: Faster &flexible research pathways for science & research insights

*What is the role of ASCR's Research Program in transforming the way we carry out energy & science research?*

1. **Post-Moore technologies**: Need basic research in new algorithms, software stacks, and programming tools for quantum and neuromorphic systems
2. **Extreme Heterogeneity**: Need new software stacks, programming models to support the heterogeneous systems of the future
3. **Adaptive Machine Learning, Modeling, & Simulation for Complex Systems**: Need algorithms and tools that support automated decision making from intelligent operating systems, in situ workflow management, improved resilience and better computational models.
4. **Uncertainty Quantification:** Need basic research in uncertainty quantification and artificial intelligence to enable statistically and mathematically rigorous foundations for advances in science domain-specific areas.
5. **Data Tsunami:** Need to develop the software and coordinated infrastructure to accelerate scientific discovery by addressing challenges and opportunities associated with research data management, analysis, and reuse.

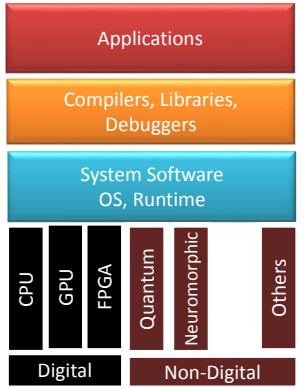


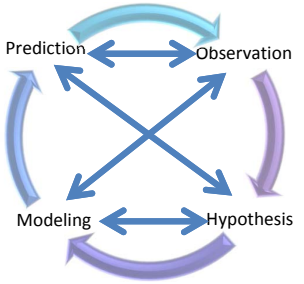

# Post-Moore directions

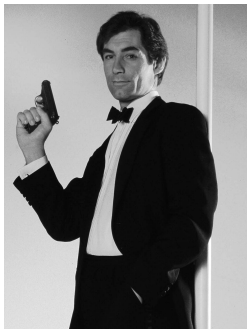

Not that post-Moore …

**Numerous Opportunities to Continue Moore's Law Technology!**

*(but winning solution is unclear)*

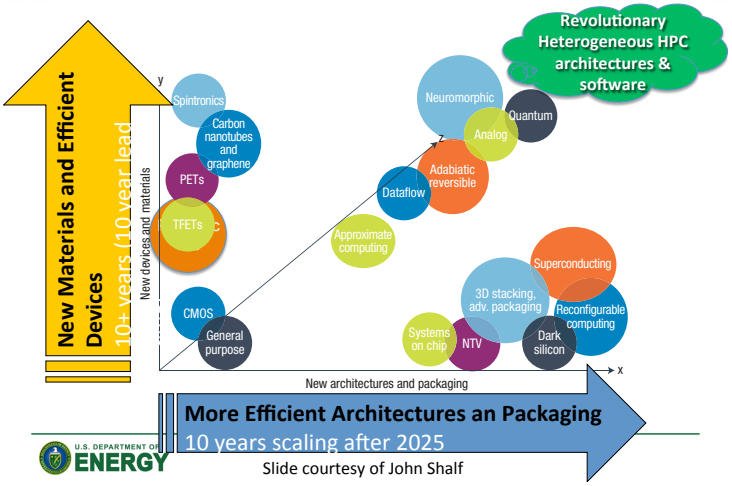


Slide courtesy of John Shalf

## Computing Beyond Moore's Law

**TABLE 1.** Summary of techology options for extending digital electronics.

| Improvement Class | Technology | Timescale | Complexity | Risk | Opportunity |
|---|---|---|---|---|---|
| Architecture and software advances | Advanced energy management | Near-Term | Medium | Low | Low |
| | Advanced circuit design | Near-Term | High | Low | Medium |
| | System-on-chip specialization | Near-Term | Low | Low | Medium |
| | Logic specialization/dark silicon | Mid-Term | High | High | High |
| | Near threshold voltage (NTV) operation | Near-Term | Medium | High | High |
| 3D integration and packaging | Chip stacking in 3D using thru-silicon vias (TSVs) | Near-Term | Medium | Low | Medium |
| | Metal layers | Mid-Term | Medium | Medium | Medium |
| | Active layers (epitaxial or other) | Mid-Term | High | Medium | High |
| Resistance reduction | Superconductors | Far-Term | High | Medium | High |
| | Crystaline metals | Far-Term | Unknown | Low | Medium |
| Millivolt switches (a better transistor) | Tunnel field-effect transistors (TFETs) | Mid-Term | Medium | Medium | High |
| | Heterogeneous semiconductors/strained silicon | Mid-Term | Medium | Medium | Medium |
| | Carbon nanotubes and graphene | Far-Term | High | High | High |
| | Piezo-electric transistors (PFETs) | Far-Term | High | High | High |
| Beyond transistors (new logic paradigms) | Spintronics | Far-Term | Medium | High | High |
| | Topological insulators | Far-Term | Medium | High | High |
| | Nanophotonics | Near/Far-Term | Medium | Medium | High |
| | Biological and chemical computing | Far-Term | High | High | High |

U.S. DEPARTMENT OF ENERGY

Slide courtesy of John Shalf

Nowell – SSDBM, June 29, 2017

Fermilab

# What is our trajectory? Here is the story.

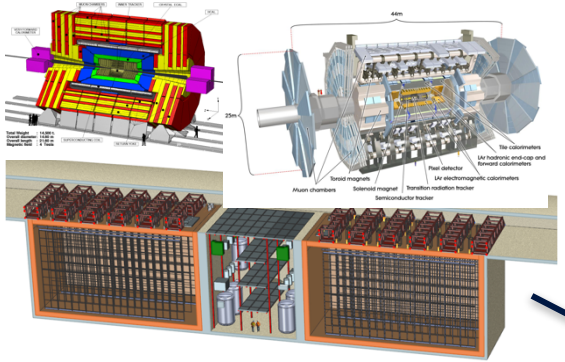

- Seems that given the deluge of data and the end of Moore's Law, that a major computing problem is imminent.
- Thanks in part to AI and post-Moore technological advances, there is hope.

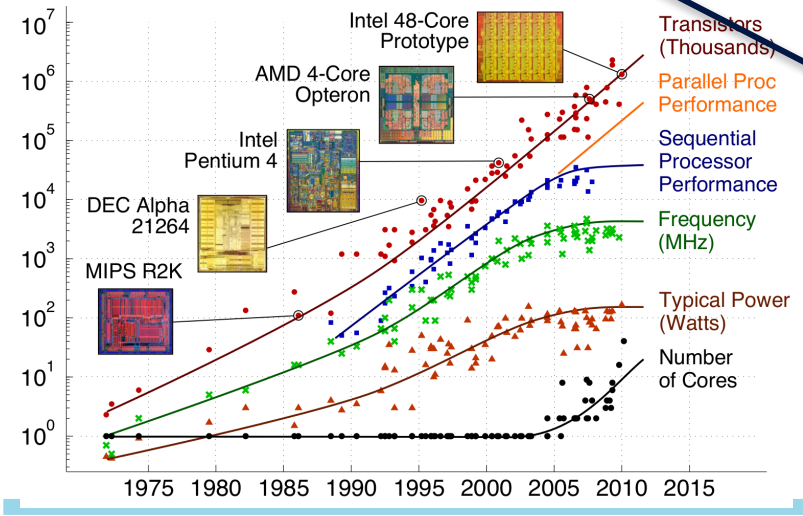


Exascale

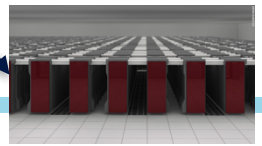

AI

Saved

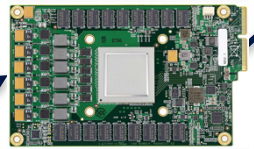

Post-Moore

Machine takeover

Likely final state

The Singularity

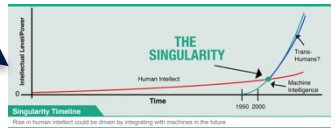




Fermilab

# Xeon Scalable Processor (SP) platform

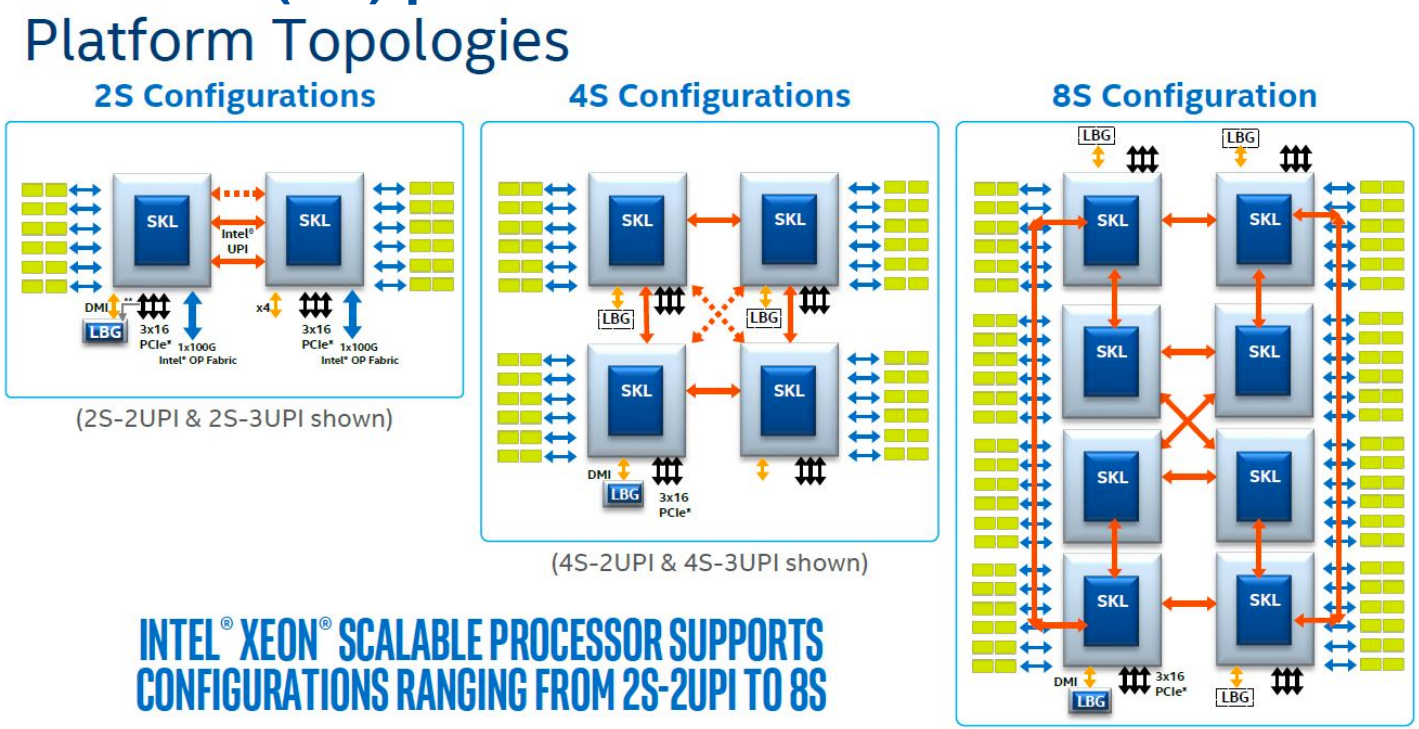


- Looks like Intel has repackaged their standard processor line

# Xeon SP roadmap

- Moving slowly ahead
- Not many surprises here: architecture changes unlikely to be earth-shattering.

**The Processes And Time Between PCs And Servers At Intel**

| Chip Family | Process | PC Chip | Launch Date | Server Chip | Launch Date | Months |
|---|---|---|---|---|---|---|
| Nehalem | 45 nm | Core i7 | Nov-08 | Xeon 5500 | Mar-09 | 4 |
| Westmere | 32 nm | Core i7 | Jul-10 | Xeon 5600 | Feb-11 | 7 |
| Sandy Bridge | 32 nm | Core i7 | Jan-11 | Xeon E5 | Mar-12 | 14 |
| Ivy Bridge | 22 nm | Core i7 | Apr-12 | Xeon E5 v2 | Sep-13 | 17 |
| Haswell | 22 nm | Core i7 | Jun-13 | Xeon E5 v3 | Sep-14 | 15 |
| Broadwell | 14 nm | Core i7 | Jan-15 | Xeon E5 v4 | Mar-16 | 14 |
| Skylake | 14+ nm | Core i7 | Sep-15 | Xeon E5 v5 | *Jul-17* | *22* |
| Kaby Lake | 14++ nm | Core i7 | Jan-17 | Xeon E5 v6 | *Mar-18* | *14* |
| Cannonlake | 10 nm | Core i7 | *Nov-17* | Xeon E5 v7 | *Nov-18* | *12* |
| Icelake | 10+ nm | Core i7 | *Jun-18* | Xeon E5 v8 | *Jan-19* | *7* |
| Tigerlake | 10++ nm | Core i7 | *Jun-19* | Xeon E5 v9 | *Sep-19* | *3* |
| ??? | 7 nm | Core i7 | *Jun-20* | Xeon E5 v10 | *Jan-20* | *-5* |
| ??? | 7+ nm | Core i7 | *Jun-21* | Xeon E5 v11 | *Jan-21* | *-5* |
| ??? | 7++ nm | Core i7 | *Jun-22* | Xeon E5 v12 | *Jan-22* | *-5* |

# AMD Zen

- “The Epyc 7000s are system-on-chip designs, which means they do not require an external chipset to operate in single-socket mode or two expand out to two-sockets with NUMA interconnects. All the necessary I/O for linking the sockets is on the chips, as are all of the controllers to link out to memory and peripherals.”
  - https://www.nextplatform.com/2017/06/20/competition-returns-x86-servers-epyc-fashion/

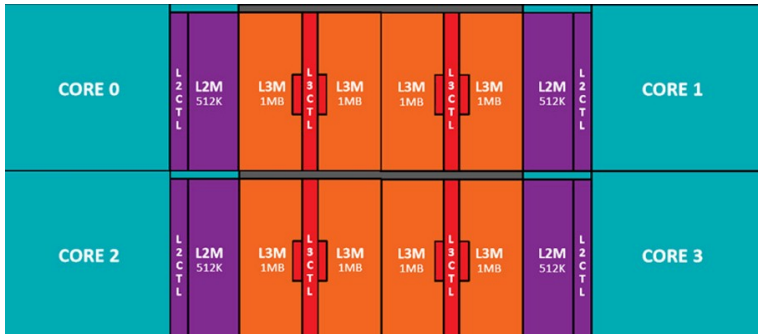


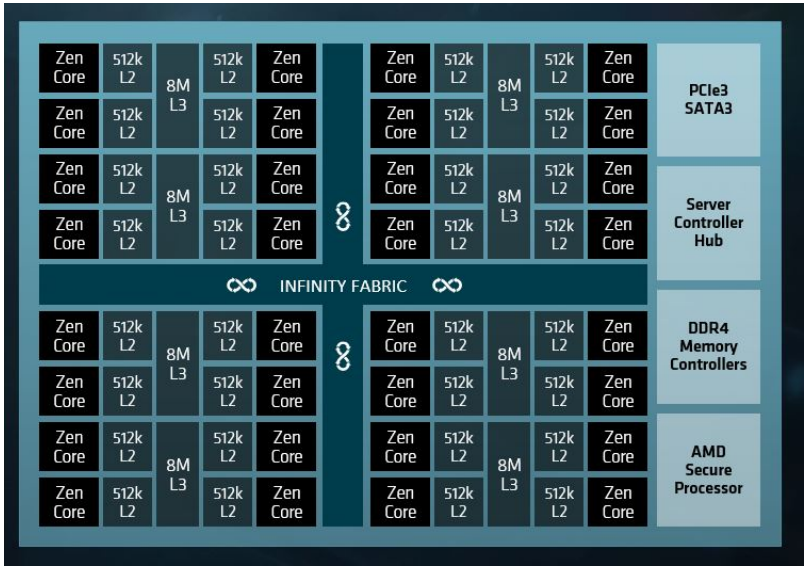

# AMD Zen (2)

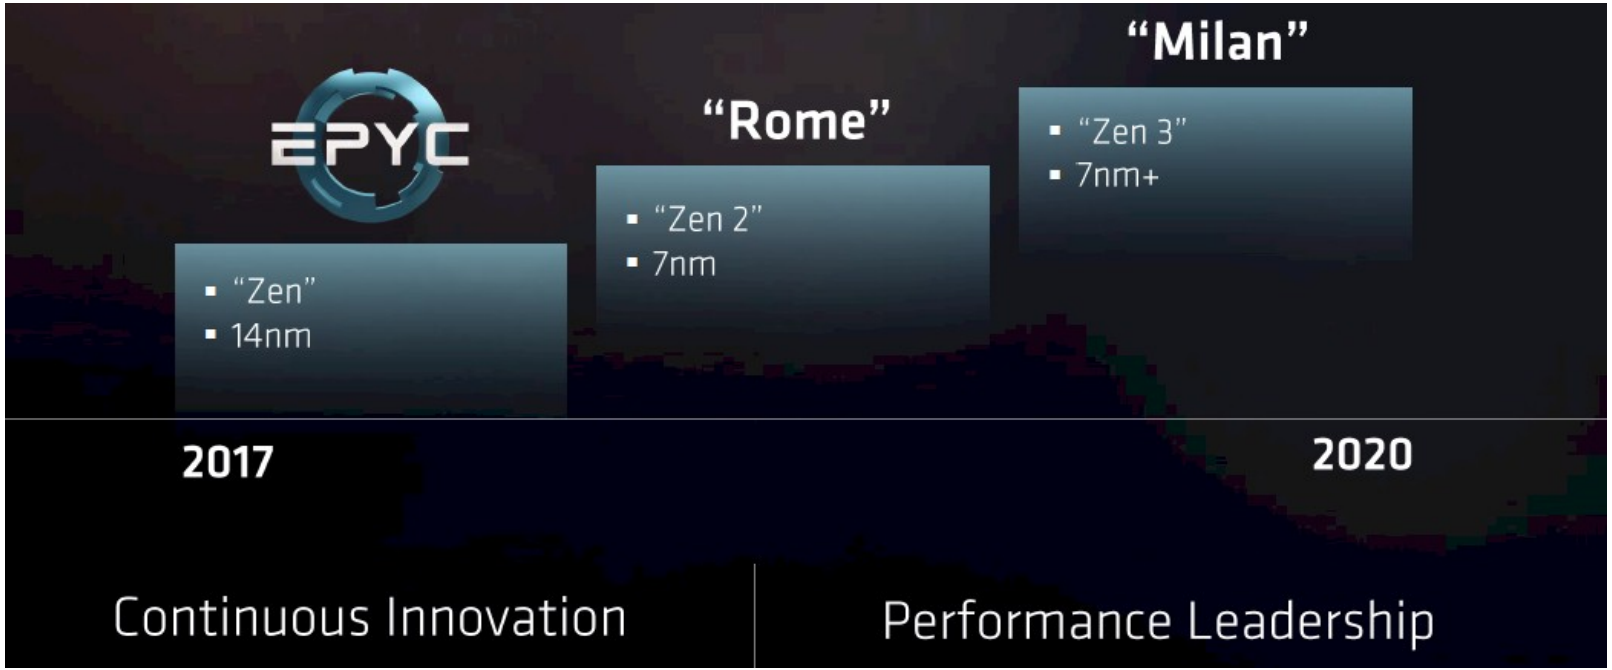
EPYC
"Zen"
14nm
"Rome"
"Zen 2"
7nm
"Milan"
"Zen 3"
7nm+
2017
2020
Continuous Innovation
Performance Leadership

# IBM Power

Power Systems

## POWER Processor Roadmap

IBM

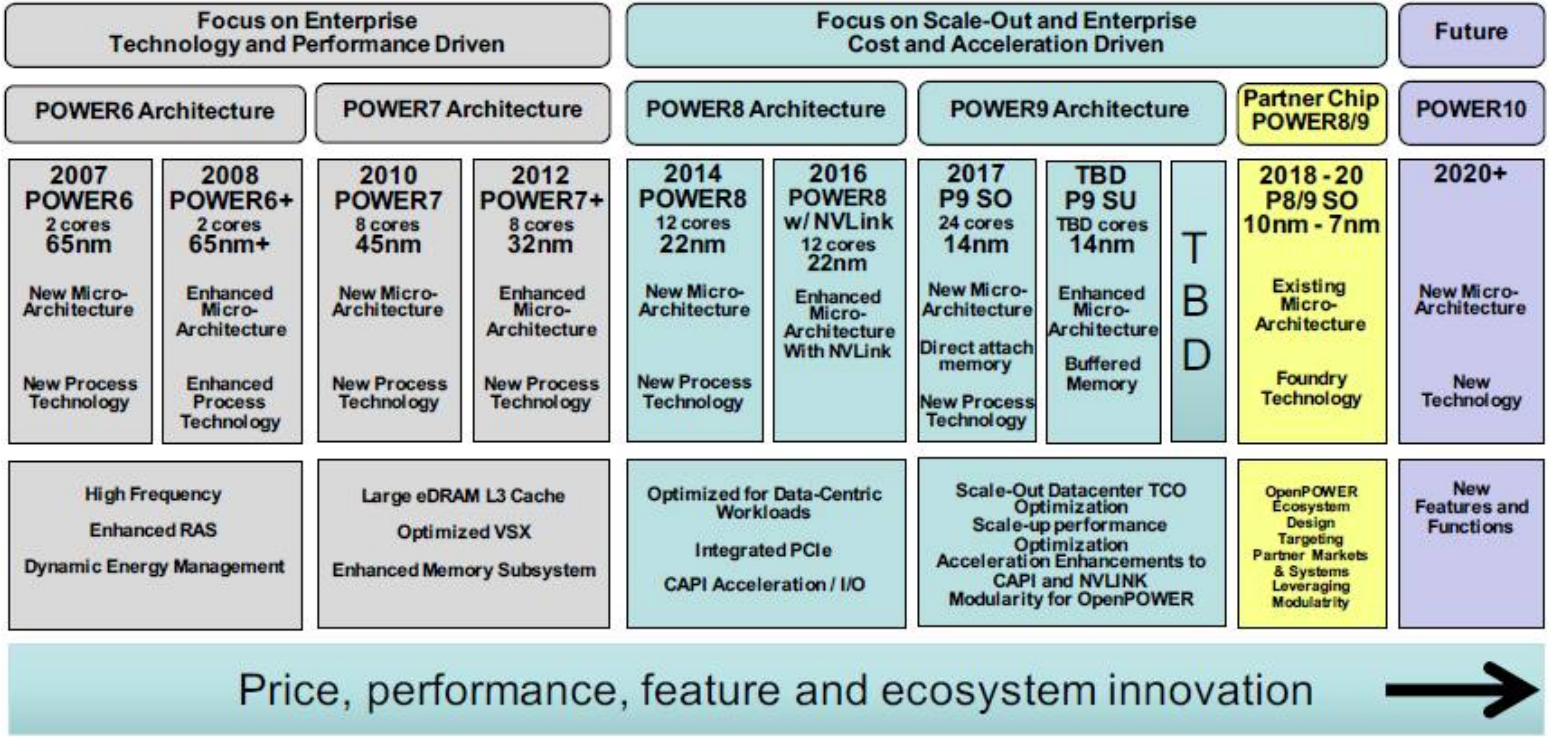


- "Power9 boasts speeds and technologies unseen in any chip before."
  - https://baseline-data.com/blog/disaster-recovery-news/ibm-power9-chip-processor/

Fermilab

# IBM Power (2)

**Four targeted implementations**

| SMP scalability / Memory subsystem | Core Count / Size: SMT4 Core | SMT8 Core |
|---|---|---|
| | 24 SMT4 Cores / Chip<br>Linux Ecosystem Optimized | 12 SMT8 Cores / Chip<br>PowerVM Ecosystem Continuity |
| **Scale-Out – 2 Socket Optimized**<br>Robust 2 socket SMP system<br>Direct Memory Attach<br>• Up to 8 DDR4 ports<br>• Commodity packaging form factor | 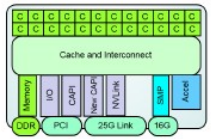 | 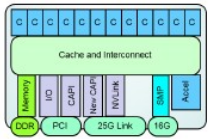 |
| **Scale-Up – Multi-Socket Optimized**<br>Scalable System Topology / Capacity<br>• Large multi-socket<br>Buffered Memory Attach<br>• 8 Buffered channels | 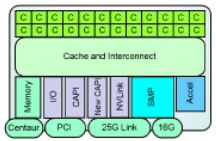 | 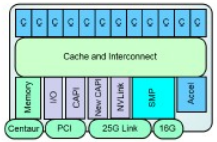 |

**Modular Execution Slices**

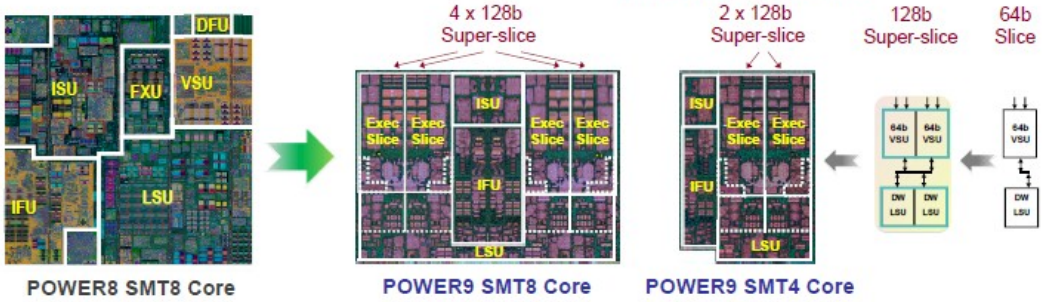


**Re-factored Core Provides Improved Efficiency & Workload Alignment**

- Enhanced pipeline efficiency with modular execution and intelligent pipeline control
- Increased pipeline utilization with symmetric data-type engines: Fixed, Float, 128b, SIMD
- Shared compute resource optimizes data-type interchange



Fermilab

# Intel Xeon Phi - KNL

## Knights Landing Overview

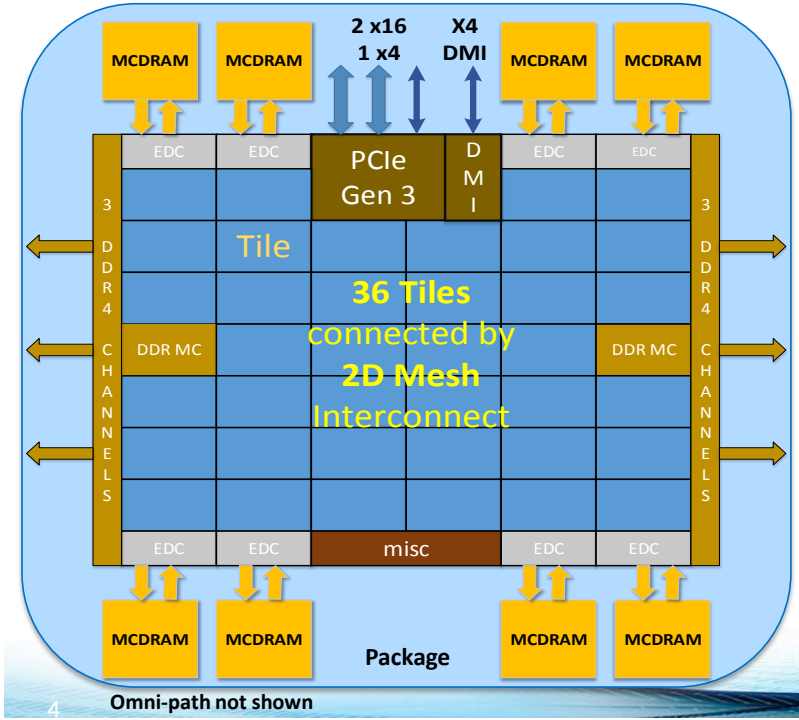


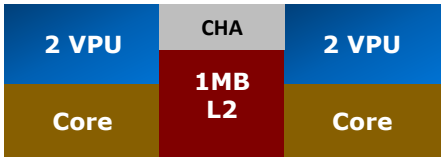


**Chip: 36 Tiles** interconnected by **2D Mesh**
**Tile**: 2 Cores + 2 VPU/core + 1 MB L2

**Memory: MCDRAM:** 16 GB on-package; High BW
**DDR4**: 6 channels @ 2400 up to 384GB
**IO:** 36 lanes PCIe Gen3. 4 lanes of DMI for chipset
**Node**: 1-Socket only
**Fabric:** Omni-Path on-package (not shown)

**Vector Peak Perf**: 3+TF DP and 6+TF SP Flops
**Scalar Perf**: ~3x over Knights Corner
**Streams Triad (GB/s)**: MCDRAM : 400+; DDR: 90+

Source Intel: All products, computer systems, dates and figures specified are preliminary based on current expectations, and are subject to change without notice. KNL data are preliminary based on current expectations and are subject to change without notice. 1Binary Compatible with Intel Xeon processors using Haswell Instruction Set (except TSX). 2Bandwidth numbers are based on STREAM-like memory access pattern when MCDRAM used as flat memory. Results have been estimated based on internal Intel analysis and are provided for informational purposes only. Any difference in system hardware or software design or configuration may affect actual performance.



Fermilab

# Intel Xeon Phi - KNL (2)

- KNL has interesting memory and vector processing features for high performance applications

## KNL ISA

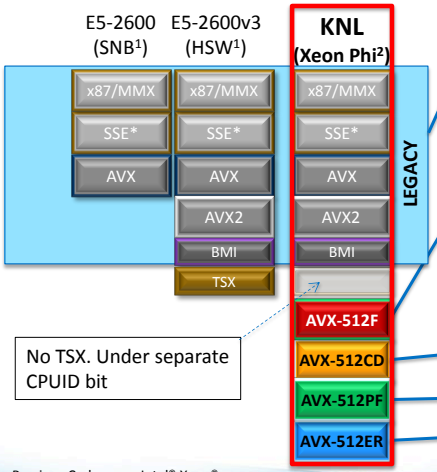


**KNL implements all legacy instructions**
- Legacy binary runs w/o recompilation
- KNC binary requires recompilation

**KNL introduces AVX-512 Extensions**
- 512-bit FP/Integer Vectors
- 32 registers, & 8 mask registers
- Gather/Scatter

**Conflict Detection**: Improves Vectorization

**Prefetch**: Gather and Scatter Prefetch

**Exponential and Reciprocal** Instructions

1. Previous Code name Intel® Xeon® processors
2. Xeon Phi = Intel® Xeon Phi™ processor

## Memory Modes

**Three** Modes. Selected at boot

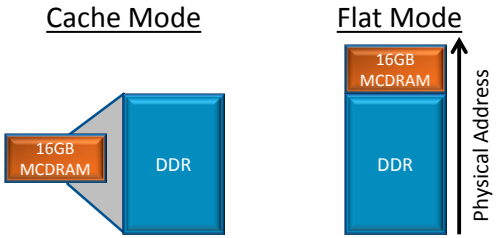


- SW-Transparent, Mem-side cache
- Direct mapped. 64B lines.
- Tags part of line
- Covers whole DDR range

- MCDRAM as regular memory
- SW-Managed
- Same address space

- Part cache, Part memory
- 25% or 50% cache
- Benefits of both

# Xeon Phi – deployments

- Cori-2 is already KNL
- Recently announced Aurora will feature Knights Hill

## ALCF 2021 EXASCALE SUPERCOMPUTER – A21

Intel/Cray Aurora supercomputer planned for 2018 shifted to 2021

Scaled up from **180 PF to over 1000 PF**


Support for three “pillars”

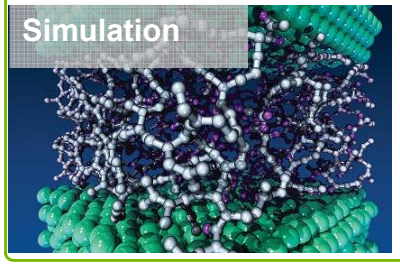


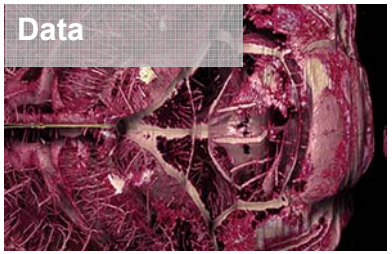


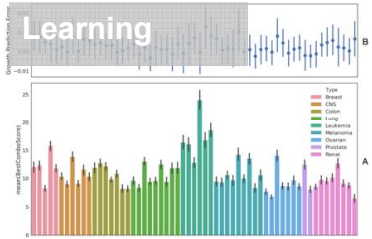


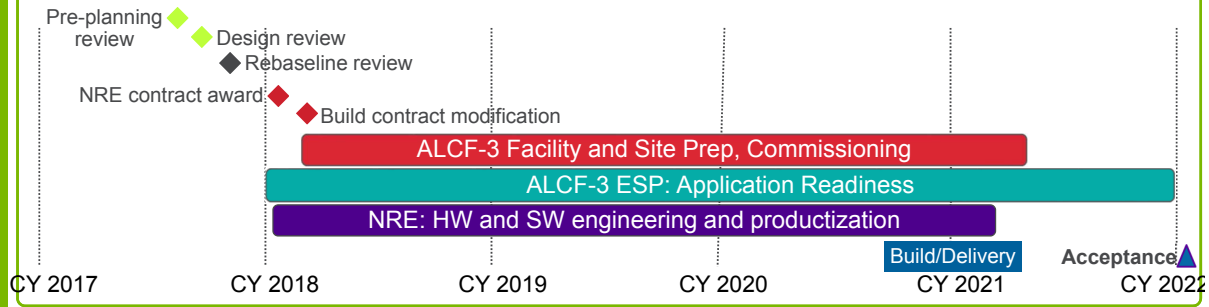

# NVIDIA Telsa

| Tesla Product | Tesla K40 | Tesla M40 | Tesla P100 | Tesla V100 |
|---|---|---|---|---|
| GPU | GK180 (Kepler) | GM200 (Maxwell) | GP100 (Pascal) | GV100 (Volta) |
| SMs | 15 | 24 | 56 | 80 |
| TPCs | 15 | 24 | 28 | 40 |
| FP32 Cores / SM | 192 | 128 | 64 | 64 |
| FP32 Cores / GPU | 2880 | 3072 | 3584 | 5120 |
| FP64 Cores / SM | 64 | 4 | 32 | 32 |
| FP64 Cores / GPU | 960 | 96 | 1792 | 2560 |
| Tensor Cores / SM | NA | NA | NA | 8 |
| Tensor Cores / GPU | NA | NA | NA | 640 |
| GPU Boost Clock | 810/875 MHz | 1114 MHz | 1480 MHz | 1530 MHz |
| Peak FP32 TFLOPS[1] | 5 | 6.8 | 10.6 | 15.7 |
| Peak FP64 TFLOPS[1] | 1.7 | .21 | 5.3 | 7.8 |
| Peak Tensor TFLOPS[1] | NA | NA | NA | 125 |
| Texture Units | 240 | 192 | 224 | 320 |
| Memory Interface | 384-bit GDDR5 | 384-bit GDDR5 | 4096-bit HBM2 | 4096-bit HBM2 |
| Memory Size | Up to 12 GB | Up to 24 GB | 16 GB | 16 GB |
| L2 Cache Size | 1536 KB | 3072 KB | 4096 KB | 6144 KB |
| Shared Memory Size / SM | 16 KB/32 KB/48 KB | 96 KB | 64 KB | Configurable up to 96 KB |
| Register File Size / SM | 256 KB | 256 KB | 256 KB | 256KB |
| Register File Size / GPU | 3840 KB | 6144 KB | 14336 KB | 20480 KB |
| TDP | 235 Watts | 250 Watts | 300 Watts | 300 Watts |
| Transistors | 7.1 billion | 8 billion | 15.3 billion | 21.1 billion |
| GPU Die Size | 551 mm² | 601 mm² | 610 mm² | 815 mm² |
| Manufacturing Process | 28 nm | 28 nm | 16 nm FinFET+ | 12 nm FFN |

[1] Peak TFLOPS rates are based on GPU Boost Clock

## NVLink high speed interconnects

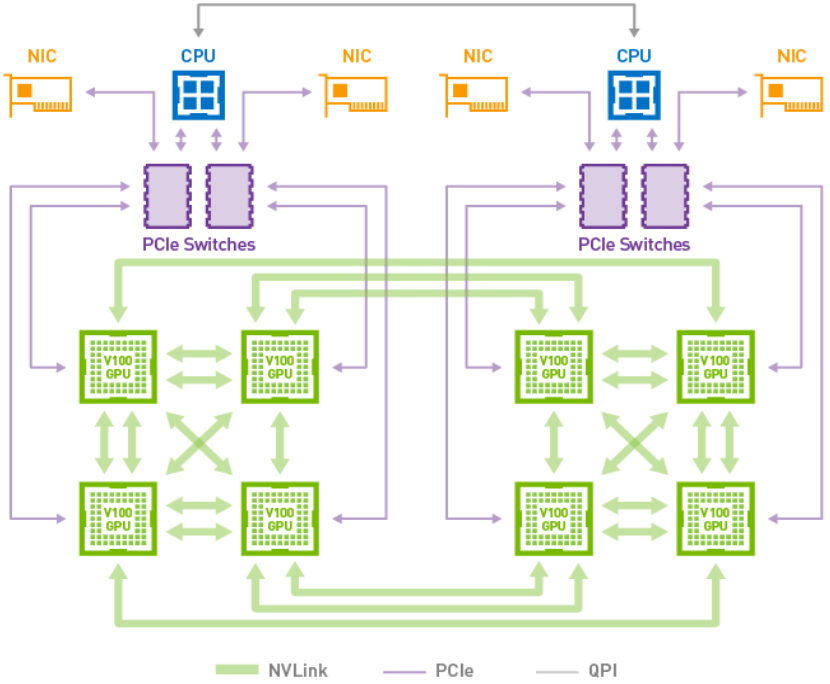


Figure 12. Hybrid Cube Mesh NVLink Topology as used in DGX-1 with V100



Fermilab

# NVIDIA Volta architecture

- Improved SIMT model
- Multiprocessing capabilities
- Tensor cores



Figure 4. Volta GV100 Full GPU with 84 SM Units



Fermilab

# NVIDIA - deployments Pascal to Volta

summit Begins at OLCF

OAK RIDGE National Laboratory | OAK RIDGE LEADERSHIP COMPUTING FACI

Start of Installation

Installation Complete

| Feature | Titan | Summit |
|---|---|---|
| Application Performance | Baseline | 5-10x Titan |
| Nodes | 18,688 | ~4,600 |
| Node performance | 1.4 TF | > 40 TF |
| Memory per Node | 32 GB DDR3 + 6 GB GDDR5 | 512 GB DDR4 + 96 GB HBM |
| NV memory per Node | 0 | 1600 GB |
| System Interconnect | Gemini (6.4 GB/s) | Dual Rail EDR-IB (23 GB/s) |
| Interconnect Topology | 3D Torus | Non-blocking Fat Tree |
| Processors | 1 AMD Opteron™ 1 NVIDIA Kepler™ | 2 IBM POWER9™ 6 NVIDIA Volta™ |
| File System | 32 PB, 1 TB/s, Lustre® | 250 PB, 2.5 TB/s, GPFS™ |

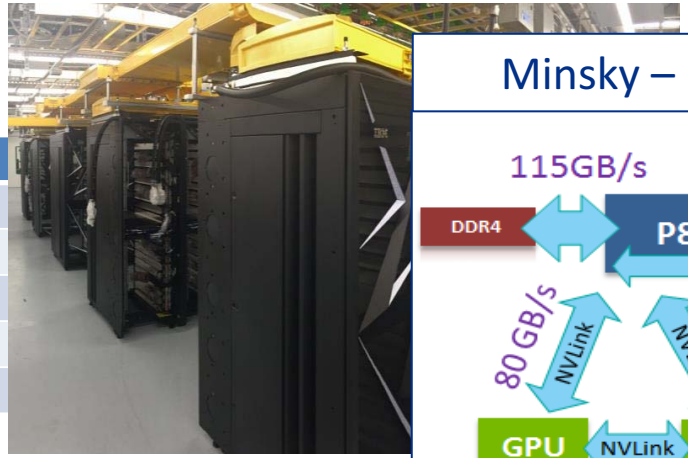

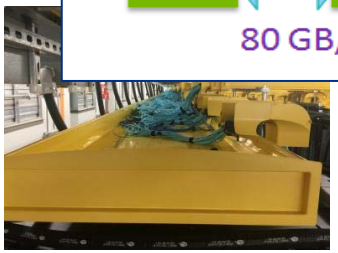

## Minsky – precursor to Summit

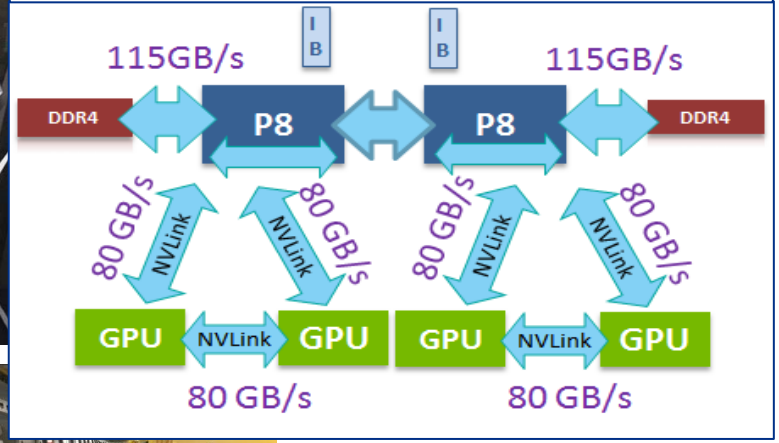




Fermilab

# Intel Knights Mill (KNM)

- Expected release in Q4 2017
- Based on KNLs, but optimized for deep learning optimizations.

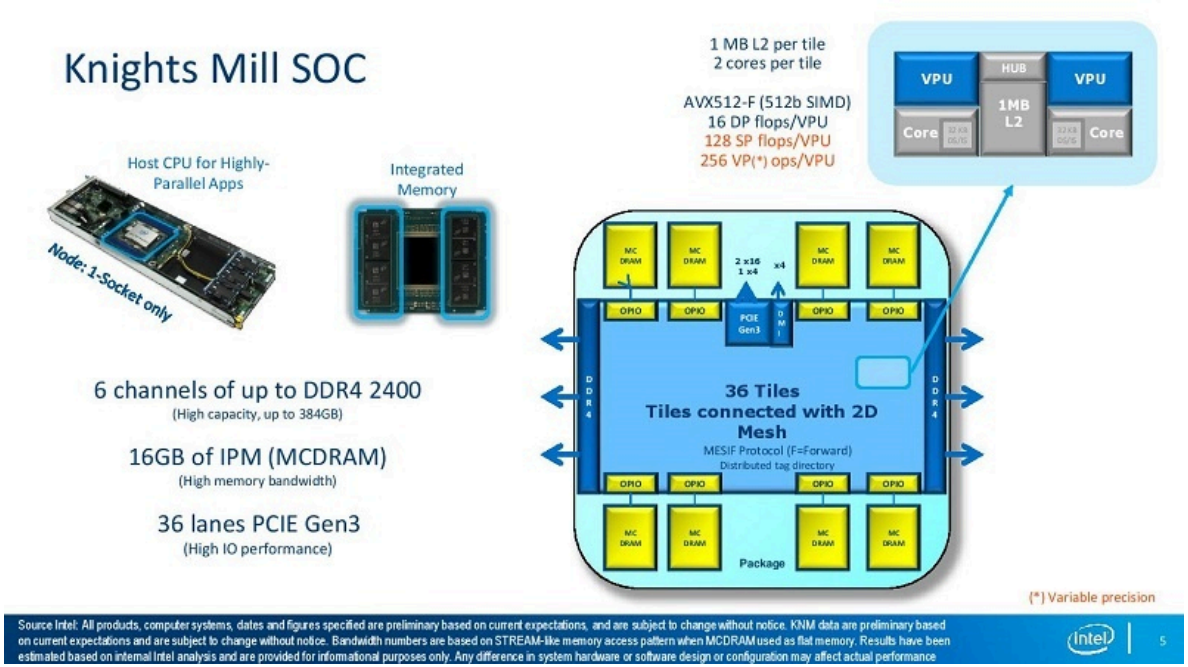

# NVIDIA Tesla P40 Inferencing Accelerator

- Based on Pascal

## FEATURES

| |
|---|
| The world's fastest processor for inference workloads |
| 47 TOPS of INT8 for maximum inference throughput and responsiveness |
| Hardware-decode engine capable of transcoding and inferencing 35 HD video streams in real time |

## SPECIFICATIONS

| | |
|---|---|
| GPU Architecture | **NVIDIA Pascal™** |
| Single-Precision Performance | **12 TeraFLOPS*** |
| Integer Operations (INT8) | **47 TOPS* (Tera-Operations per Second)** |
| GPU Memory | **24 GB** |
| Memory Bandwidth | **346 GB/s** |
| System Interface | **PCI Express 3.0 x16** |
| Form Factor | **4.4" H x 10.5" L, Dual Slot, Full Height** |
| Max Power | **250 W** |
| Enhanced Programmability with Page Migration Engine | **Yes** |
| ECC Protection | **Yes** |
| Server-Optimized for Data Center Deployment | **Yes** |
| Hardware-Accelerated Video Engine | **1x Decode Engine, 2x Encode Engine** |

* With Boost Clock Enabled

# AMD Instinct

- https://instinct.radeon.com/en/
- AMD's ML-focussed GPGPU: 16, 32-bit FP.
- Integrated into Project 47, "PFLOP in a rack" for Q4 2017:
  - 20 * (Epyc 7601 + 4 * Instinct MI25).
  - 30 GFLOP/W.

# Catapult: datacenter-level FPGA integration

- Microsoft has FPGA available in some of their large-scale cloud resources
- Alternative high-speed network presented as TCP/IP to OS
- Added portion of distributed search engine filtering into user space on FPGA

Catapult v2 Mezzanine card

WCS Gen4.1 Blade with NIC and Catapult FPGA

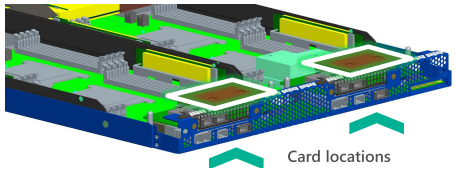


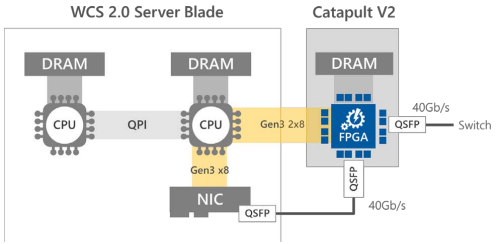


bing

- 2x Faster at 2x higher load
- Much lower variance

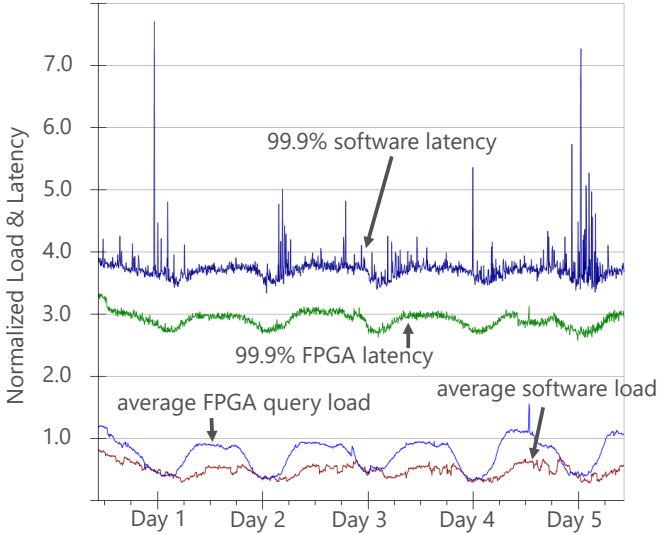


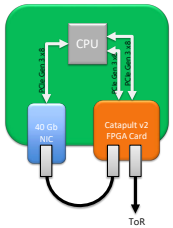


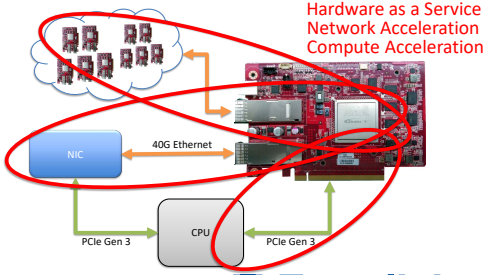




Fermilab

# Google Tensor Processing Units (TPU)

- Custom ASIC for Google's TensorFlow framework
- First generation: 8-bit matrix multiply engine, clock speed 700MHz, 8 GiB of dual-channel 2133 MHz DDR3 SDRAM, 34GB/s of bandwidth
  - Up to 128K operations per cycle
  - **83X** better performance per watt ratio compared with contemporary CPUs and a **29X** better ratio than contemporary GPUs.
- Second generation: both training and inference in a 180 TFLOP system board, complete with custom network to lash several together into "TPU pods" that can deliver Top 500-class supercomputing at up to 11.5 petaflops of peak performance.

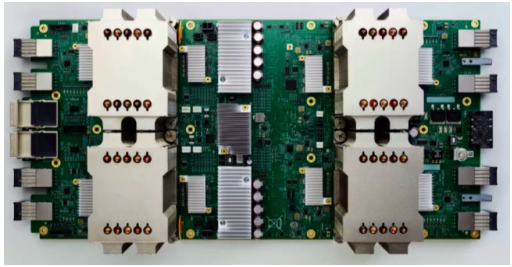

# Neuromorphic chips

- IBM's TrueNorth:
  - 1 million neurons and 256 million synapses
  - 5.4 billion transistors, and an on-chip network of 4,096 neurosynaptic cores
  - 70mW during real-time operation
  - The Artificial Brain built by IBM now has 64 Million Neurons (using 64 of the chips), using 10 watts to power all 64
  - 10 Billion neurons by 2020
- Intel's Loihi
  - Test chip, consists of 128 computing cores
  - Each core has 1,024 artificial neurons
    - a total of more than 130,000 neurons on chip, and
    - 130 million synaptic connections on chip
  - currently limited tests using FPGA for application like path planning
  - First test chip will be available in Nov 2017
  - Will be available to research groups in 2018
- Qualcomm's Zeroth
  - Neural Processing Unit
  - Making deep learning available to mobile devices
    - 10nm Snapdragon 835 processor

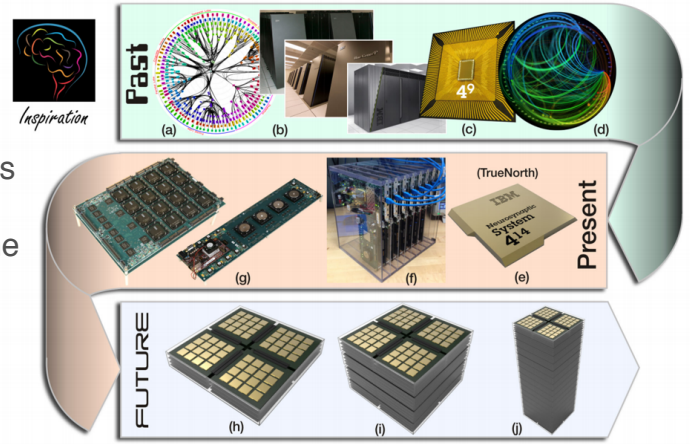


# Quantum computers

- D-Wave
  - Quantum annealing device
  - 2000 Q system in Jan 2017, current buyers: Temporal Defense Systems Inc, Google/NASA/USRA, Volkswagen Group, Virginia Tech
  - The world's first commercial quantum-computer processor is smaller than a wristwatch and can evaluate more possibilities simultaneously than there are atoms in the visible universe.
- Google http://www.quantumplayground.net/#/home
  - 22 qubit quantum processing unit (QPU) in its wiring mount, currently in test.
  - 50-qbit next target for early 2018
- IBM https://quantumexperience.ng.bluemix.net/qx/community
  - QX is a five qubit quantum processor and matching simulator
  - A 16 qubit processor was made available in May 2017

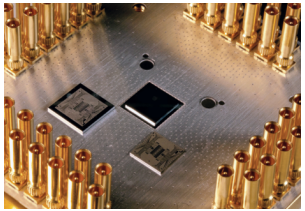

**Qubit**: a unit of quantum information, a two-state quantum-mechanical system, which allows the qubit to be in a superposition of both states.

# Automata Processing

- Reconfigurable processing architecture that enables programmers to easily exploit massive parallelism.
- FNAL LDRD to investigate pattern recognition / waveform recognition (M. Wang, et al.)

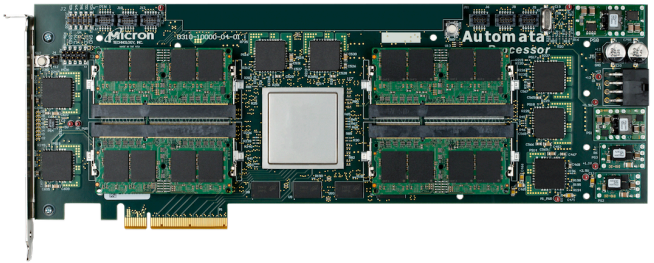


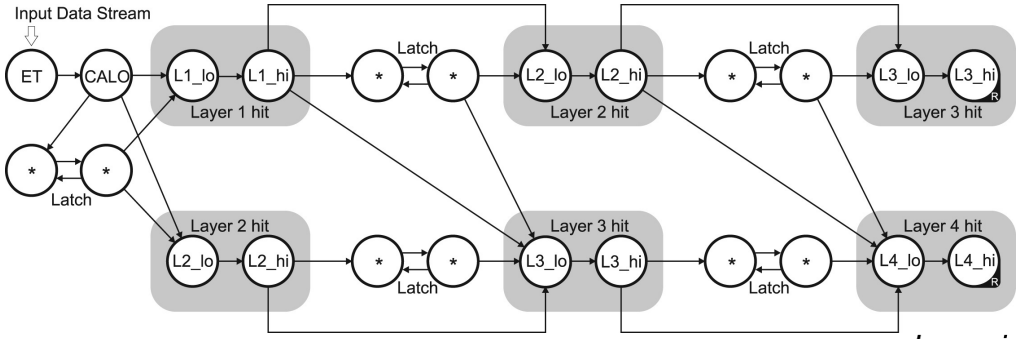


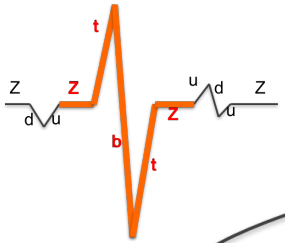


[^X]: anything but a "X"

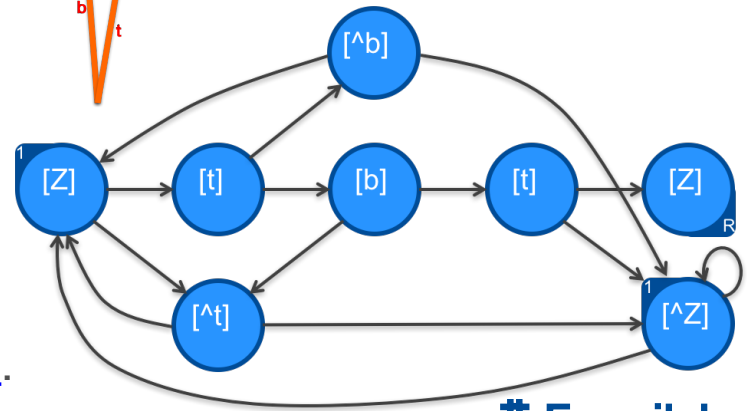


*Incoming data*

- Currently only available via CAP.
- Previously Micron, spun off to Natural Intelligence.



Fermilab

# Probabilistic and Approximate computing

- **Probabilistic** computing: less *accurate*.
- **Approximate** computing: less *precise*.
- Implications for fault tolerance / power consumption.
- Examples:
  - Intel PCMOS (went quiet).
  - Lyric Computing (went quiet).
  - Rice U / I-Slate (went quiet).
- Might see some low significance bit decay or truncation at lower power modes in future processors.

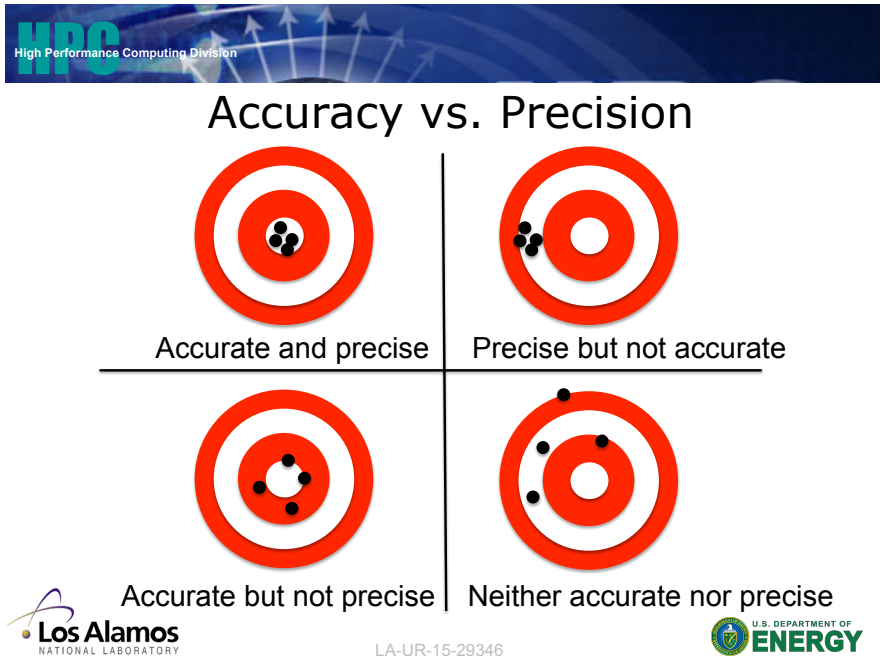


Slide from L. Monroe talk, 2015-12-10

# Memory and caching storage

- "Holy Grail:" non-volatile (NV), DRAM speeds & latencies, HDD+ capacities.
- DDR4-2666: V, 13.5ns latency, 60GiB/s, $5-$9/GiB. DDR5 coming soon. STT MRAM, "tweaks?"
- Intel 3D-Xpoint (Optane): NV, 9us latency, ~2GiB/s, $4/GiB.
- NRAM (carbon nanotubes, Nantero): NV, <50ns(?), ?? GiB/s, $??/GiB. Fujitsu to mass-produce in 2018, initially aimed at flash-style packages, capacity, density to increase rapidly. RAM to come.
- ReRAM (e.g. Crossbar): NV, 1T1R (low latency) vs 1TnR (higher density, latency).
  - SMIC presentation Thursday in Beijing!

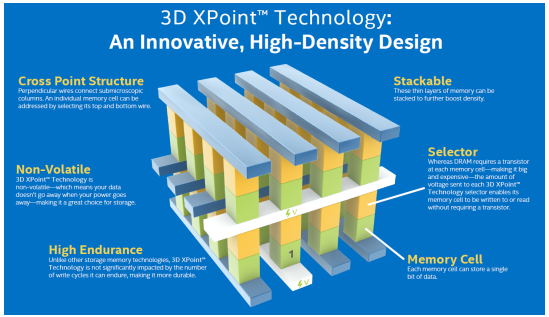


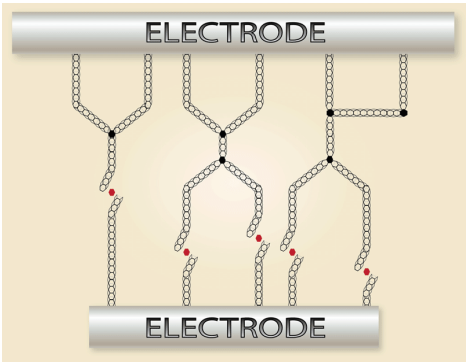


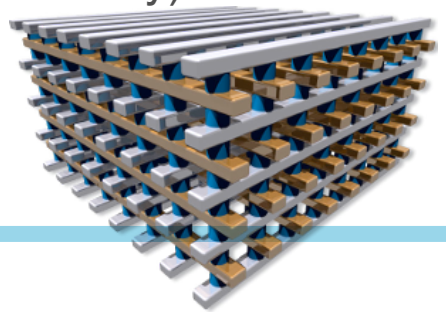



Fermilab

# ”File” storage

- SSD (3D-NAND): NV, 20us latency, ~3.5GiB/s, $0.90/GiB. Likely to be eclipsed within 10y by NRAM / ReRAM.
- HDD: NV, 4ms latency, ~2GiB/s, $0.06/GiB for 12TB, HAMR -> capacity, throughput increases by factor 10-40 (?) over the next ten years. Possibly eclipsed by NRAM or ReRAM within 10y?
- Tape storage: NV, loooonnnnggg latency, 350MiB/s, $0.006/GiB for 15TB. Just announced 330TB, capacities, throughput to double every 2 years for the next 10.

# Networking: current

- Ethernet is the reliable standby
  - Commodity high-throughput networking: $10K for 4 x 40 Gb/s ports, or 16 x 10 Gb/s ports
  - RoCEv2 latency as low as 1.3 microseconds
- InfiniBand is commodity
  - $13k for 36-port EDR 100 Gb/s IB, 7 Tb/s total throughput, 90 ns latency
- OmniPath was new in 2015
  - $8k for 24-port 100 Gb/s, 110ns latency
- Wireless 5GHz (IEEE 802.11ac)
  - up to 1.3 Gb/s on 5 GHz band; 1 ms latency
  - My Mac laptop has one of these!

# Networking: near future

- Ethernet: 200 Gb/s planned for 2018/19.
  - 50 Gb/s serial links x 4
  - reduced power consumption w/r/t 100 Gb/s, using 10 lanes @10 GB/s
- InfiniBand: 50 Gb/s lanes in 2017
- 200 Gb/sec Omni-Path 2 *was* slated for ANL Aurora in 2018
  - this machine is being both re-imagined and delayed
- Intel: Moving the network (OmniPath) into the CPU
- Mellanox: Moving computing (CPUs) into the network

# Networking: 10+ years out

- Ethernet

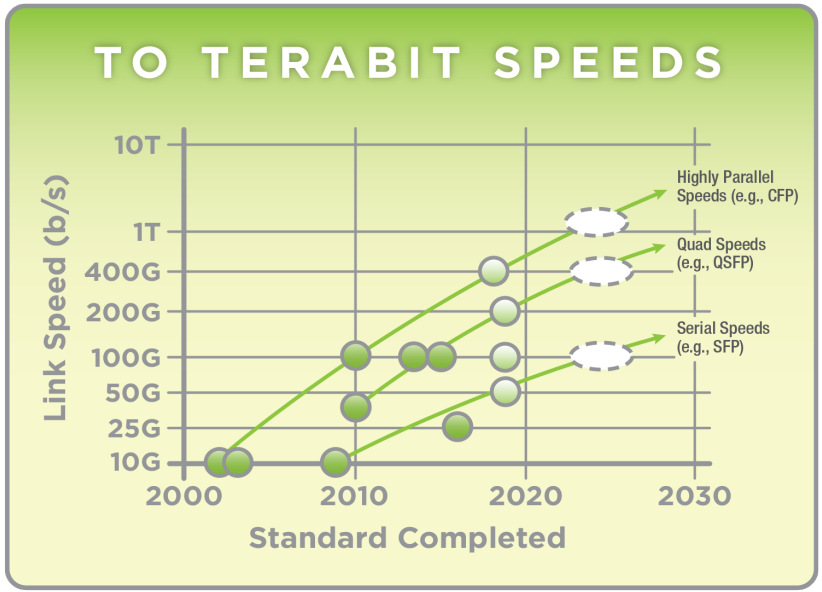


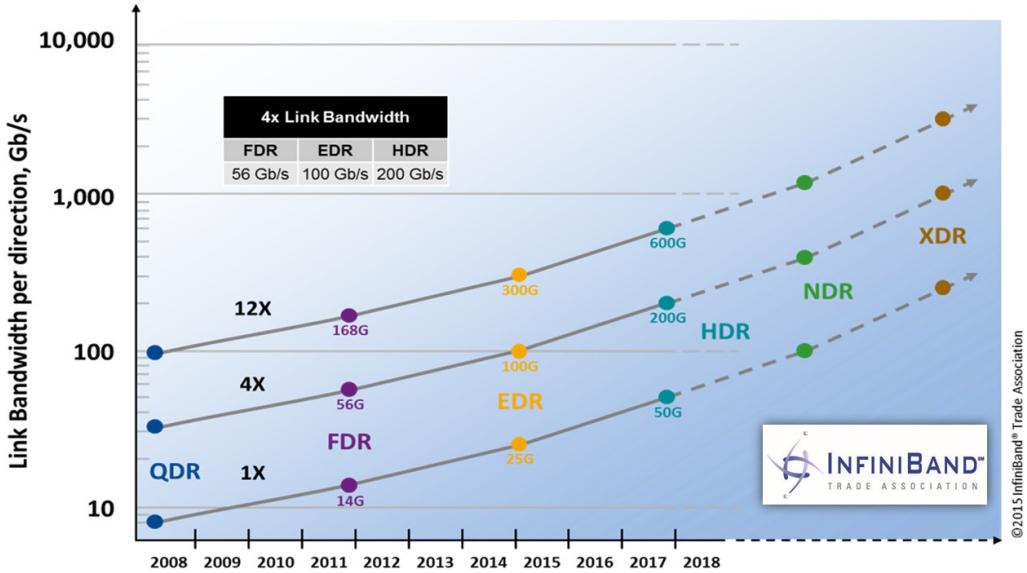


- InifiniBand
- Omnipath?



Fermilab

# Bus technologies

- https://www.nextplatform.com/2017/07/14/system-bottleneck-shifts-pci-express/
  - PCIe4.0 in Power9, not Kaby Lake Xeon or Epyc.
  - PCIe5.0 spec not complete until 2019, at least a couple of years beyond that.
  - Server systems already hitting bus bottlenecks in some configurations – hence NVLink, etc.
- Anything else is blue-sky: photonics, surface plasmons?
- Moving toward system-on-a-chip (SOC) – on-die chipset, networking, main memory … ?

| | RAW BIT RATE | LINK BW | BW/ LANE/WAY | TOTAL BW X16 |
|---|---|---|---|---|
| PCIe 1.x | 2.5GT/s | 2Gb/s | 250MB/s | 8GB/s |
| PCIe 2.x | 5.0GT/s | 4Gb/s | 500MB/s | 16GB/s |
| PCIe 3.x | 8.0GT/s | 8Gb/s | ~1GB/s | ~32GB/s |
| PCIe 4.0 | 16GT/s | 16Gb/s | ~2GB/s | ~64GB/s |
| PCIe 5.0 | 32GT/s | 32Gb/s | ~4GB/s | ~128GB/s |

# Conclusion

- Commodity now is different from commodity 10 or even 5 years ago.
- Moore's Law demise is being addressed with innovative new uses of transistors (SIMD, SIMT, TPU, AP, neuromorphic, etc.) and entirely new ideas: QuXXX.
- Evolutionary advances in file storage, networking, bus.
- Possibility (likelihood?) of revolution in fast non-volatile memory -> new board architectures, SOC, etc.
- Integrated systems showing much promise for supercomputers: Aurora, Summit, AMD Project 47: "Commodity" supercomputing on the way?
- Will need to think very carefully about programming / algorithms to take advantage of all this innovation: Intel TBB, C++ standard, HPX, MPI, OpenACC, OpenMP, Numpy, Julia.