

# Preparing for Supercomputing's Sixth Wave

Jeffrey S. Vetter



Presented to  
**ACM Symposium on High-Performance Parallel  
and Distributed Computing**


2 June 2016  
Kyoto



<http://ft.ornl.gov> [vetter@computer.org](mailto:vetter@computer.org)

ORNL is managed by UT-Battelle  
for the US Department of Energy

 **OAK RIDGE**  
National Laboratory




# Overview

- Evidence demonstrates that we are nearing the end of Moore's Law
- Put differently, we are ending the fifth wave of computing, and entering the sixth wave
- Predictions for next decade
- Possible sixth wave technologies
- Implications, warnings, and opportunities



<http://spectrum.ieee.org/geek-life/history/moores-law-milestones>



<http://spectrum.ieee.org/img/NEW31965AprIntel300-1430139827044.jpg>

# End of Moore's Law Really ??

# Contemporary devices are approaching fundamental limits

## Stuttering





Figure 1 | As a metal oxide–semiconductor field effect transistor (MOSFET) shrinks, the gate dielectric (yellow) thickness approaches several atoms (0.5 nm at the 22-nm technology node). Atomic spacing limits the


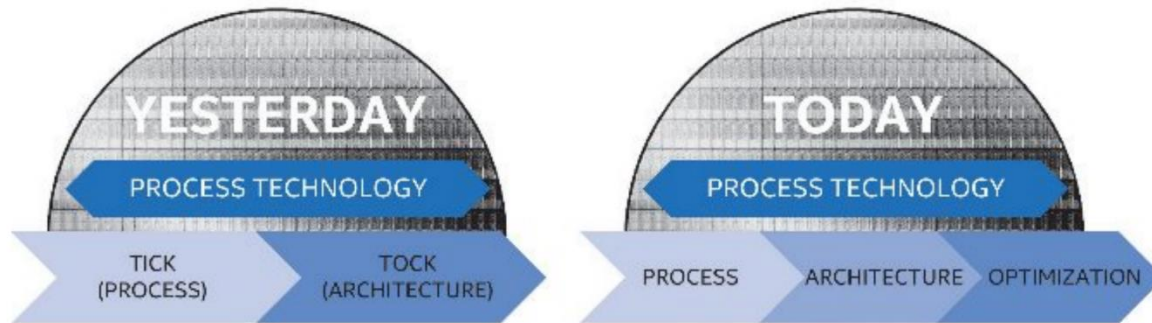
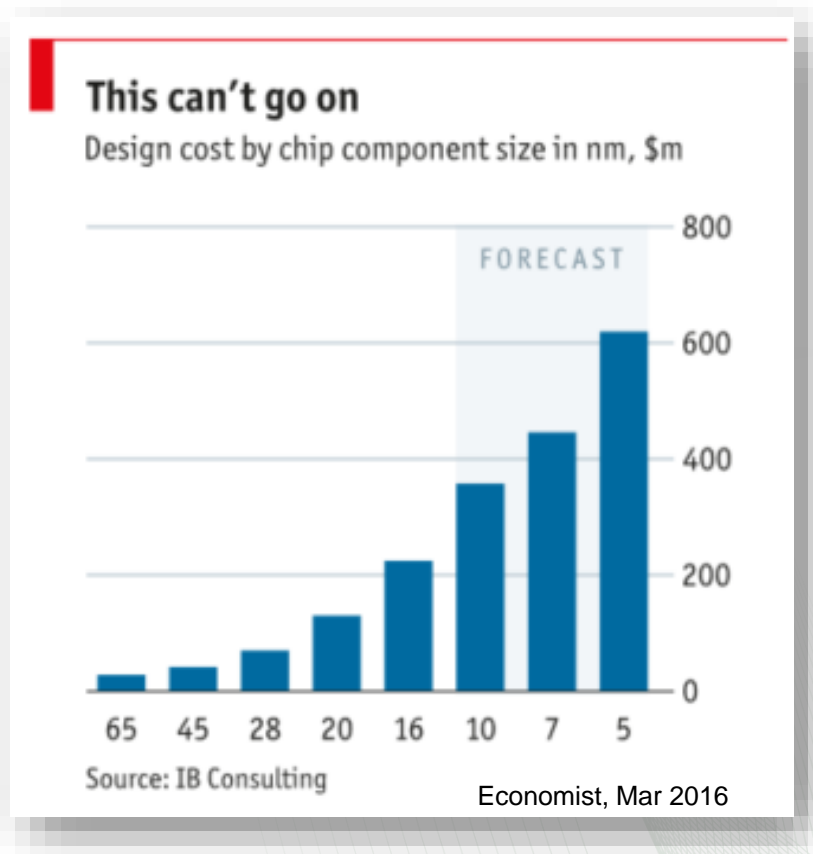


Figure 2 | As a MOSFET transistor shrinks, the shape of its electric field departs from basic rectilinear models, and the level curves become disconnected. Atomic-level manufacturing variations, especially for dopant

# Semiconductors are taking longer and cost more to design and produce



<http://www.anandtech.com/show/10183/intels-tick-tock-seemingly-dead-becomes-process-architecture-optimization>



# Semiconductor business is highly process-oriented, optimized, and growing extremely capital intensive

designlines INDUSTRIAL CONTROL

## News & Analysis

### Semi industry fab costs limit industry growth

Nicolas Mokhoff

10/3/2012 03:00 PM EDT

Post a comment

NO RATINGS  
LOGIN TO RATE

Like 4 Tweet Share G+ 0

MANHASSET, N.Y. -- The fundamental economics of the semiconductor industry may start changing sooner rather than later, according to market research firm Gartner Inc.

The costs of staying at the leading edge in semiconductor manufacturing are rising. Semiconductor manufacturers need to plan on equipment costs increasing at about 15 percent for each new node, according to Gartner (Stamford, Conn.).

It's possible that 450-mm manufacturing will achieve the goal of 30 percent cost reduction. But that equates to only three or four years of increasing equipment costs, and consequently, delays the inevitable, Gartner said. It is also possible that new technologies will emerge that will slow the rate of cost increases, according to the firm.

According to Gartner, the costs of manufacturing equipment needed for leading-edge semiconductor manufacturing are increasing at a rate between 7 percent and 10 percent per year, depending on the basic process.

By 2020, current cost trends will lead to an average cost of between \$15 billion and \$20 billion for a leading-edge fab, according to the report. By 2016, the minimum capital expenditure budget needed to justify the building of a new fab will range from \$8 billion to \$10 billion for logic, \$3.5 billion to \$4.5 billion for DRAM and \$6 billion to \$7 billion for NAND flash, according to the report.

The Gartner report predicts that at current spending rates, only eight companies could afford to build fabs in the next few years.

By 2020, current cost trends will lead to an average cost of between \$15 billion and \$20 billion for a leading-edge fab, according to the report. By 2016, the minimum capital expenditure budget needed to justify the building of a new fab will range from \$8 billion to \$10 billion for logic, \$3.5 billion to \$4.5 billion for DRAM and \$6 billion to \$7 billion for NAND flash, according to the report.

Context: Intel Reports Full-Year Revenue of \$55.4 Billion, Net Income of \$11.4 Billion (Intel SEC Filing for FY2015)

#### Major 2013 IC Foundries (Pure-Play and IDM)

2013 Rank	2012 Rank	Company	Foundry Type	Location	2011 Sales (\$M)	2012 Sales (\$M)	2012/2011 Change (%)	2013 Sales (\$M)	2013/2012 Change (%)
1	1	TSMC	Pure-Play	Taiwan	14,299	16,951	19%	19,850	17%
2	2	GlobalFoundries	Pure-Play	U.S.	3,195	4,013	26%	4,261	6%
3	3	UMC	Pure-Play	Taiwan	3,760	3,730	-1%	3,959	6%
4	4	Samsung	IDM	South Korea	2,192	3,439	57%	3,950	15%
5	5	SMIC*	Pure-Play	China	1,320	1,542	17%	1,973	28%
6	8	Powerchip**	Pure-Play	Taiwan	374	625	67%	1,175	88%
7	9	Vanguard	Pure-Play	Taiwan	520	582	12%	713	23%
8	6	Huahong Grace***	Pure-Play	China	619	677	9%	710	5%
9	10	Dongbu	Pure-Play	South Korea	500	540	8%	570	6%
10	7	TowerJazz	Pure-Play	Israel	611	639	5%	509	-20%
11	11	IBM	IDM	U.S.	420	432	3%	485	12%
12	12	MagnaChip	IDM	South Korea	350	400	14%	411	3%
13	13	WIN	Pure-Play	Taiwan	304	381	25%	354	-7%
—	—	<b>Top 13 Total</b>	—	—	<b>28,464</b>	<b>33,951</b>	<b>19%</b>	<b>38,920</b>	<b>15%</b>
—	—	<b>Top 13 Share</b>	—	—	<b>89%</b>	<b>90%</b>	—	<b>91%</b>	—
—	—	<b>Other Foundry</b>	—	—	<b>3,446</b>	<b>3,669</b>	<b>6%</b>	<b>3,920</b>	<b>7%</b>
—	—	<b>Total Foundry</b>	—	—	<b>31,910</b>	<b>37,620</b>	<b>18%</b>	<b>42,840</b>	<b>14%</b>

\*Does not include Wuhan Xinxin (now XMC) for 2012 or 2013.

Source: IC Insights, company reports

\*\*Powerchip transitioned from an IDM foundry to a pure-play foundry in 2013.

\*\*\*Hua Hong NEC and Grace merged in 2012 (excludes Shanghai Huali).

# Business climate reflects this uncertainty, cost, complexity: consolidation

designlines WIRELESS & NETWORKING

Blog

## IC Merger Mania Hits Fever Pitch

Dylan McGrath, Contributing Editor

12/2/2015 10:13 AM EST

1 comments post a comment

Like 10 Tweet Share 81 G+ 1

NO RATINGS  
LOGIN TO RATE

**With the announcement that Microsemi outbid Skyworks for PMC-Sierra, the total value of semiconductor industry acquisitions announced in 2015 eclipsed the \$100 billion mark.**

The wave of consolidation that has run rampant in the semiconductor industry for more than a year has reached new heights as buyers scramble to land a shrinking number of potential acquisition targets.

Case in point: Microsemi Corp.'s agreement last week to acquire mixed-signal storage IC house PMC-Sierra Inc. for \$2.5 billion in cash and stock. The purchase price is a full 25% higher than the \$2 billion that Skyworks Solutions Inc. agreed to pay for PMC less than two months earlier. It is also 77% higher than the closing price of PMC's stock on Sept. 30, a few days before the Skyworks deal was announced.

Mergers and acquisitions have been part of the semiconductor industry's DNA from the get go. But, even in an acquisition-happy industry, 2015 has been an extraordinary year by any measure. The total value of semiconductor industry acquisitions announced this year now stands at a little over \$102 billion, more than eight times the average annual value of acquisitions by semiconductor companies over the past five years, according to market research firm IC Insights Inc.

Faced with an industry that is growing far too slowly to satisfy Wall Street, chip companies have no choice but to dig deep to fund acquisitions as a way to grow revenue and scale.

"It all comes back to the pressure these companies are facing in trying to grow their sales faster than the [semiconductor] market is growing," said Rob Lineback, a senior market research analyst at IC Insights.

## Intel to acquire Altera for \$54 a share

Monday, 1 Jun 2015 | 8:33 AM ET

REUTERS

## SANDISK COMPLETES ACQUISITION OF FUSION IO

JUL 23, 2014

### ACQUISITION TO BOOST SANDISK'S ENTERPRISE GROWTH

MILPITAS, Calif., July 23, 2014 - SanDisk Corporation (NASDAQ: SNDK), a global leader in flash storage solutions, today announced it has completed the previously announced acquisition of Fusion-io, a leading provider of high-performance hardware and software solutions that enhance application performance.

"I am delighted to welcome the employees of Fusion-io to SanDisk. The Fusion-io team will accelerate our efforts to enable the flash storage market to reach its full potential. Together we will offer the best solutions to our customers and drive growth in the industry."

## Avago Agrees to Buy Broadcom for \$37 Billion

By MICHAEL J. de la MERCED and CHAD BRAY MAY 28, 2015



## Western Digital Now A Storage Powerhouse With SanDisk Acquisition



(SanDisk)

BRIAN DEAGON | 3/16/2016

f t in e+ Rss

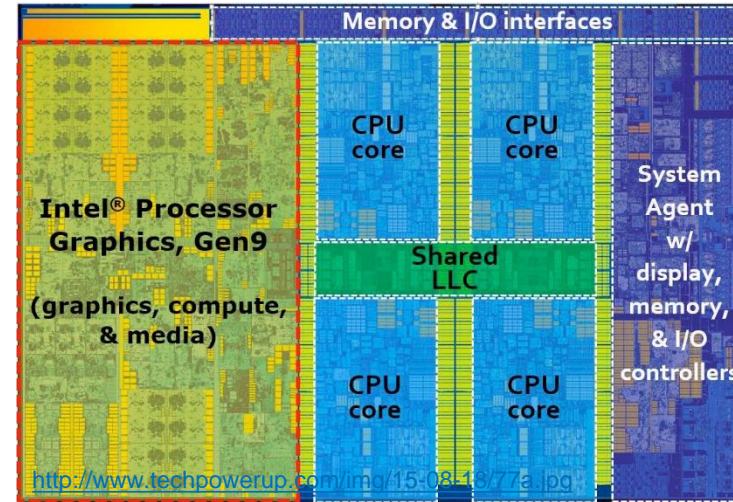
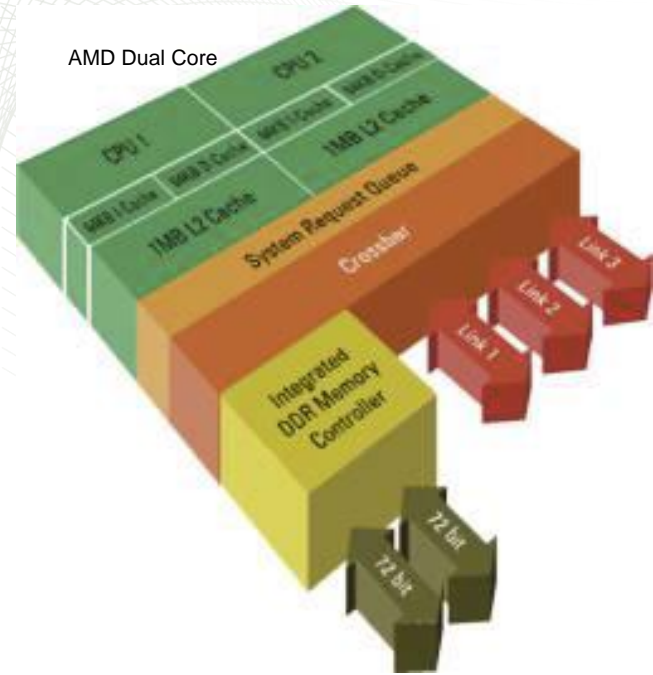
**W**estern Digital (WDC) shareholders approved a \$16 billion deal to acquire SanDisk (SNDK), creating a formidable competitor in both disk drives and flash-chip storage.

"We like the merger," Monika Garg, an analyst at Pacific Crest Securities, wrote in a research note. "The deal makes strategic sense for both companies. For Western Digital, it removes an overhang from the flash-chip impact on hard disk drives. Western Digital would then own the best flash assets and supply for its solid-state drive business."

OAK RIDGE  
National Laboratory

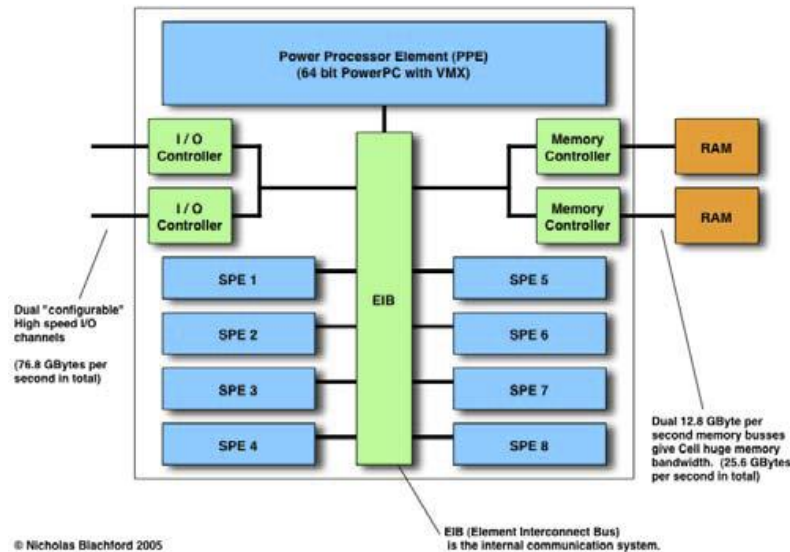
# Past decade has seen an increase in processor/node architectural complexity in order to provide more performance per watt

AMD Dual Core




Intel Skylake

Cell Processor Architecture




NVIDIA Pascal





# End of Moore's Law ??


- Device level physics will prevent much smaller level design of current transistor technologies
- Business trends indicate asymptotic limits of both manufacturing capability and economics
- Architectural complexity is growing in unbounded (and often unfortunate) directions
- Fortunately, our HPC community has been driving the additional dimension of parallelism for several decades



Put differently, we are ending  
the fifth wave of computing  
and entering the sixth wave



# Five Waves of Computing




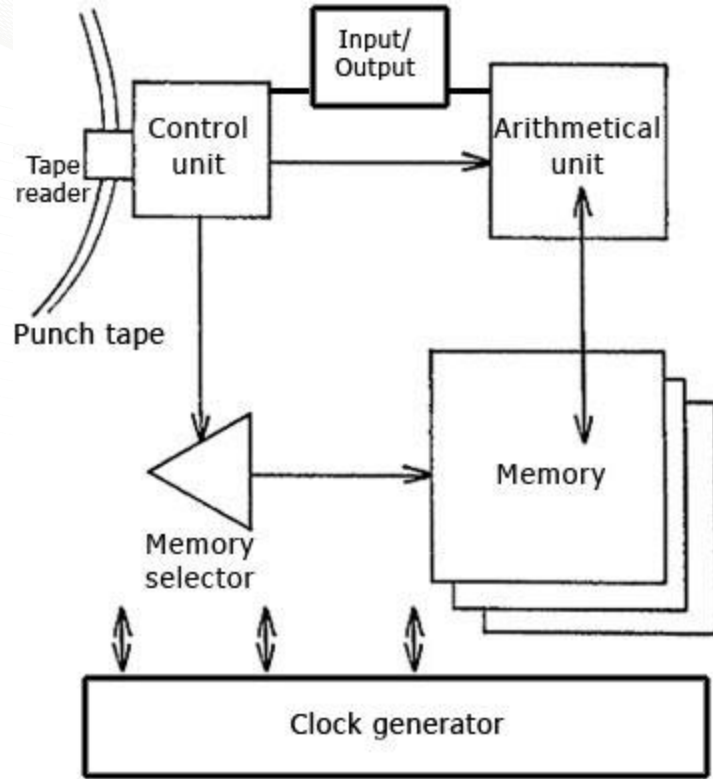
<http://www.kurzweilai.net/exponential-growth-of-computing>

# Mechanical



- Why replaced?
  - Performance
  - Fixed program


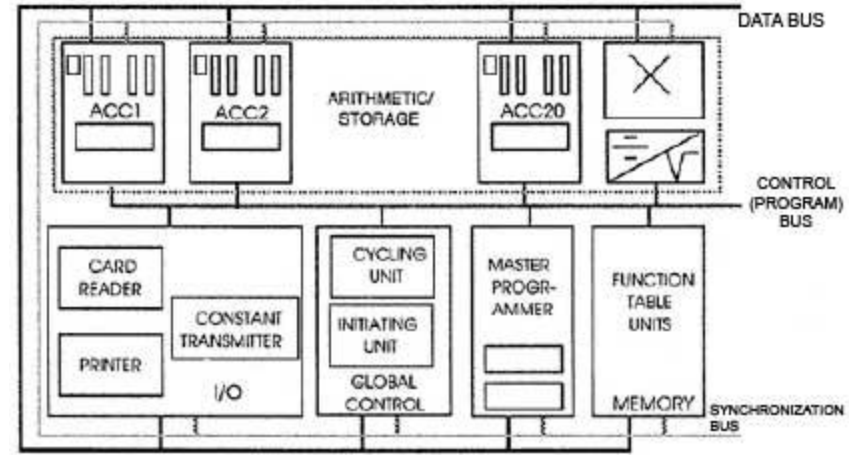
# Electromechanical Relay – Z1 by Konrad Zuse – circa 1936



# Electronic Computer : ENIAC

It could perform 5000 addition cycles a second and do the work of 50000 people working by hand. In thirty seconds, ENIAC could calculate a single trajectory, something that would take twenty hours with a desk calculator or fifteen minutes on the Differential Analyzer.

ENIAC required 174 kilowatts of power to run. It contained 17468 vacuum tubes, 1500 relays, 500000 soldered joints, 70000 resistors and 10000 capacitors-circuitry. The clock rate was 100 kHz. Input and output via an IBM card reader and card punch and tabulator.



# Transistor



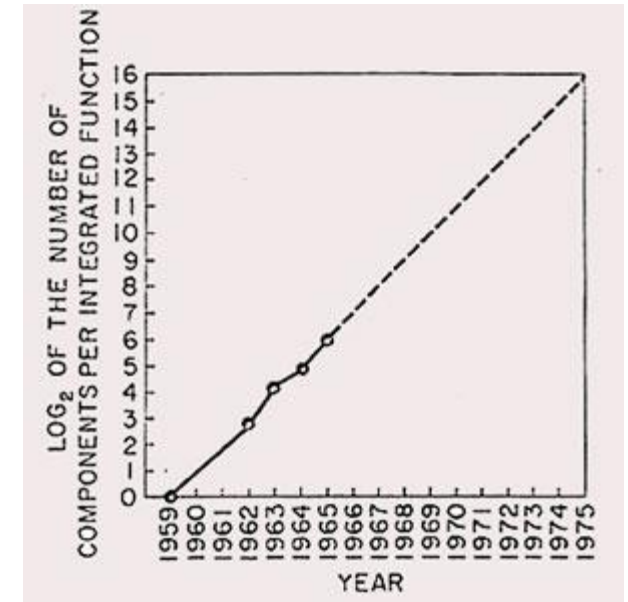
- Bell Labs demonstrates first transistor in 1947.
- University of Manchester demonstrates first operational Transistor Computer in 1953
  - 92 point-contact transistors and 550 diodes
- IBM, Philco, Olivetti commercialize transistorized large-scale computers
- Non IC Architectures popular through 1968

# Integrated Circuits

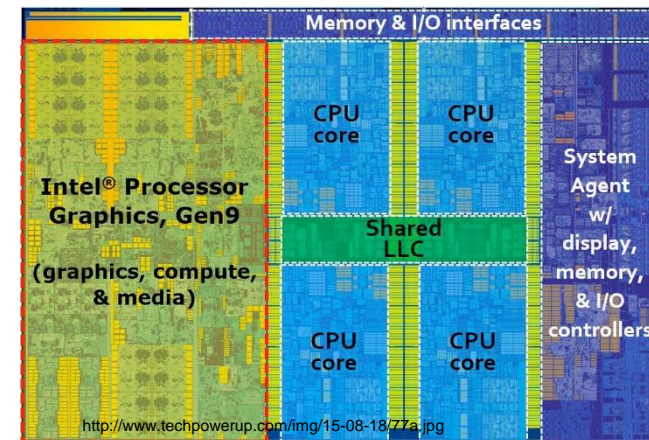
- Demonstrated by Kilby in 1958
  - Kilby won the 2000 Nobel Prize in Physics
- Fairchild Semiconductor
  - developed own idea of an integrated circuit that solved many practical problems Kilby's had not (silicon v. germanium)
  - was also home of the first silicon-gate IC technology with self-aligned gates, the basis of all modern CMOS computer chips.



<http://spectrum.ieee.org/geek-life/history/moores-law-milestones>



<http://spectrum.ieee.org/img/NEW31965AprIntel300-1430139827044.jpg>



<http://www.techpowerup.com/img/15-08-18/77a.jpg>




# Remember that our community (Supercomputing) has added the dimension of scalable parallelism beyond Moore's Law

– several times over


- Supercomputers have pushed the boundaries using parallelism
  - CDC
  - Cray Vector
  - MPP
  - Clusters (killer micro)
  - Shared memory multicore
  - Heterogeneous
- In fact, parallelism has provided most of our recent progress

## From Giga to Exa, via Tera & Peta



Shekhar Borkar, Intel

# Sixth Wave of Computing




<http://www.kurzweilai.net/exponential-growth-of-computing>

What is in store for the Transition  
Period (Next Decade) ??

# #1: Architectural specialization will continue and accelerate

- Accelerators and SoCs already dominate multiple markets
- Vendors, lacking Moore's Law, will need to continue to offer new products (to stay in business)
  - Grant that advantage of better CMOS process stalls
  - Use the same transistors differently to enhance performance
- Architectural design will become extremely important, critical
  - Address new parameters for benefits/curse of Moore's Law

## No single architecture solves all power problems



- Industry has debated merits of each architecture for decades...
- Combination of all approaches optimizes power and performance


©2010 IEEE International Solid-State Circuits Conference

©2010 IEEE

Source: Delagi,  
ISSCC 2010

# Co-designing architectures for very specific applications can produce profound performance improvements: Anton can offer 100-1000x

NVIDIA Pascal



D.E. Shaw, M.M. Deneroff, R.O. Dror et al., "Anton, a special-purpose machine for molecular dynamics simulation," *Communications of the ACM*, 51(7):91-7, 2008.

CADE.METZ BUSINESS 05.18.16 3:57 PM

## GOOGLE BUILT ITS VERY OWN CHIPS TO POWER ITS AI BOTS



GOOGLE

**GOOGLE HAS DESIGNED** its own computer chip for driving deep neural networks, an AI technology that is reinventing the way Internet services operate.

This morning at Google I/O, the centerpiece of the company's year, CEO Sundar Pichai said that Google has designed an ASIC, or application-specific integrated circuit, that's specific to deep neural nets. These are networks of

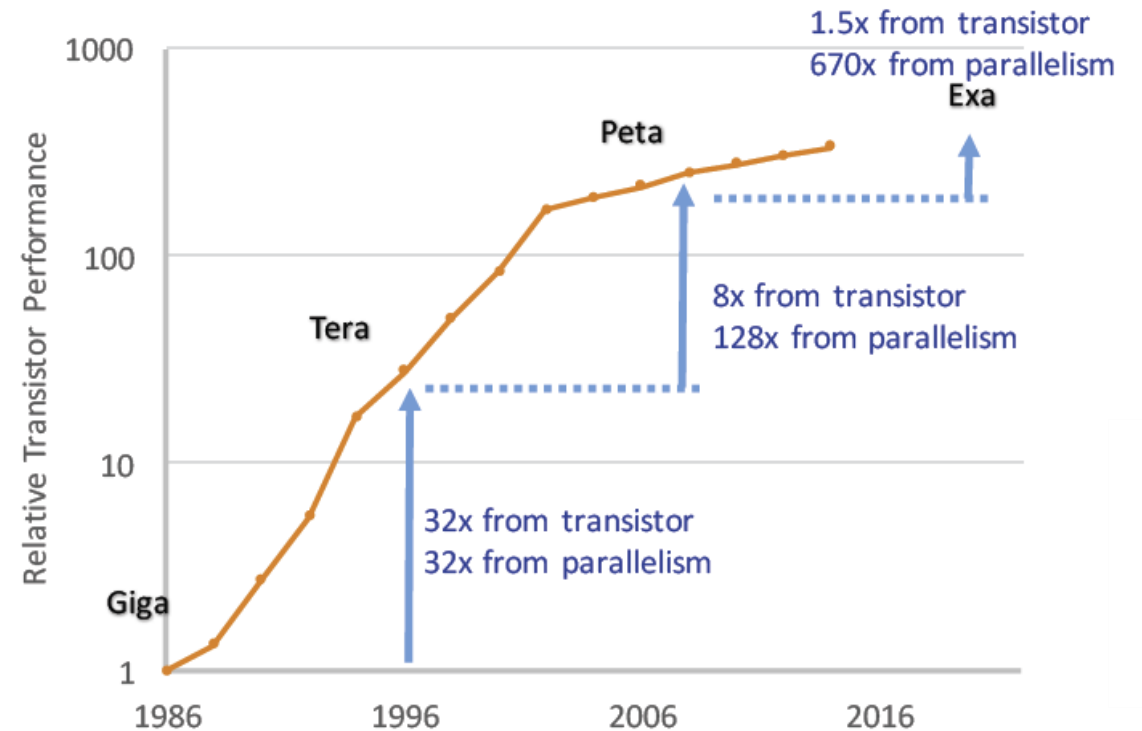
OAK RIDGE  
National Laboratory

<http://www.wired.com/2016/05/google-tpu-custom-chips/>

## #2: Continue finding new opportunities for hierarchical parallelism

- Expect no gain from transistors
- Specialization and commodity systems will need to use parallelism at all levels effectively
  - Continuing this trend!
- However, interconnection networks and memory systems must increase capacity and bandwidth (with no real improvements in latency)
  - Optical networks
  - Silicon photonics
  - Non-volatile memory

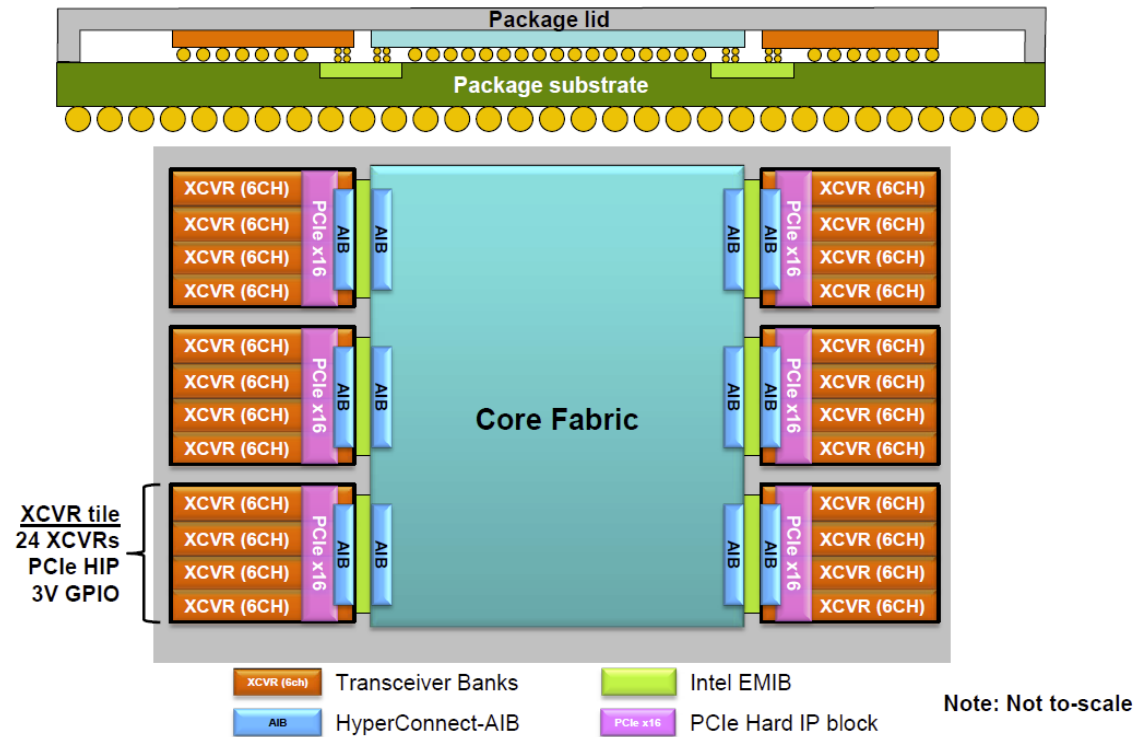
### From Giga to Exa, via Tera & Peta



Shekhar Borkar, Intel

# #3: Tighter integration and manufacturing of components will provide some benefits: components with different processes, functionality; local bandwidth



## System-in-Package Construction



# Improved Stacking, Vias, Communication Techniques

## Ultra-Thin 4 $\mu$ wafer breakthrough

- ❑ Wafer thinning has been stuck at ~40 $\mu$  due to “Gettering problem”
  - ❑ Barrier was due in part to loss of the “gettering effect” at smaller dimensions when performing back grinding, causing impurities affecting device performance (particularly leakage) and yield.
- ❑ DISCO Corporation solution can now thin to a few microns
  - ❑ DISCO introduced a “Gettering Dry Polish” wheel which forms gettering sites while grinding, allowing thinning of wafer silicon to a few microns without device damage. [35]
- ❑ Example: DRAM silicon thinned to 4 microns
  - ❑ See “Ultra Thinning down to 4 $\mu$ m using 300-mm Wafer proven by 40-nm Node 2 Gb DRAM for 3D Multi-stack WOW Applications.”[36] They concluded “No degradation in terms of retention characteristics and distribution employing 2 Gb DRAM wafer was found after ultra-thinning.”





Ultra-thin wafers can be handled (from DISCO website)

August 11, 2014

Hot Chips 26 – ThruChip Wireless Connections

10

## Communication is via magnetic field



$$\mu_{Si} = \mu_{SiO_2} = 1$$

Can easily induce a 200 mV signal in receiver coil.

Magnetic field can pass through silicon, including over active circuitry.

August 11, 2014

Hot Chips 26 – ThruChip Wireless Connections


13



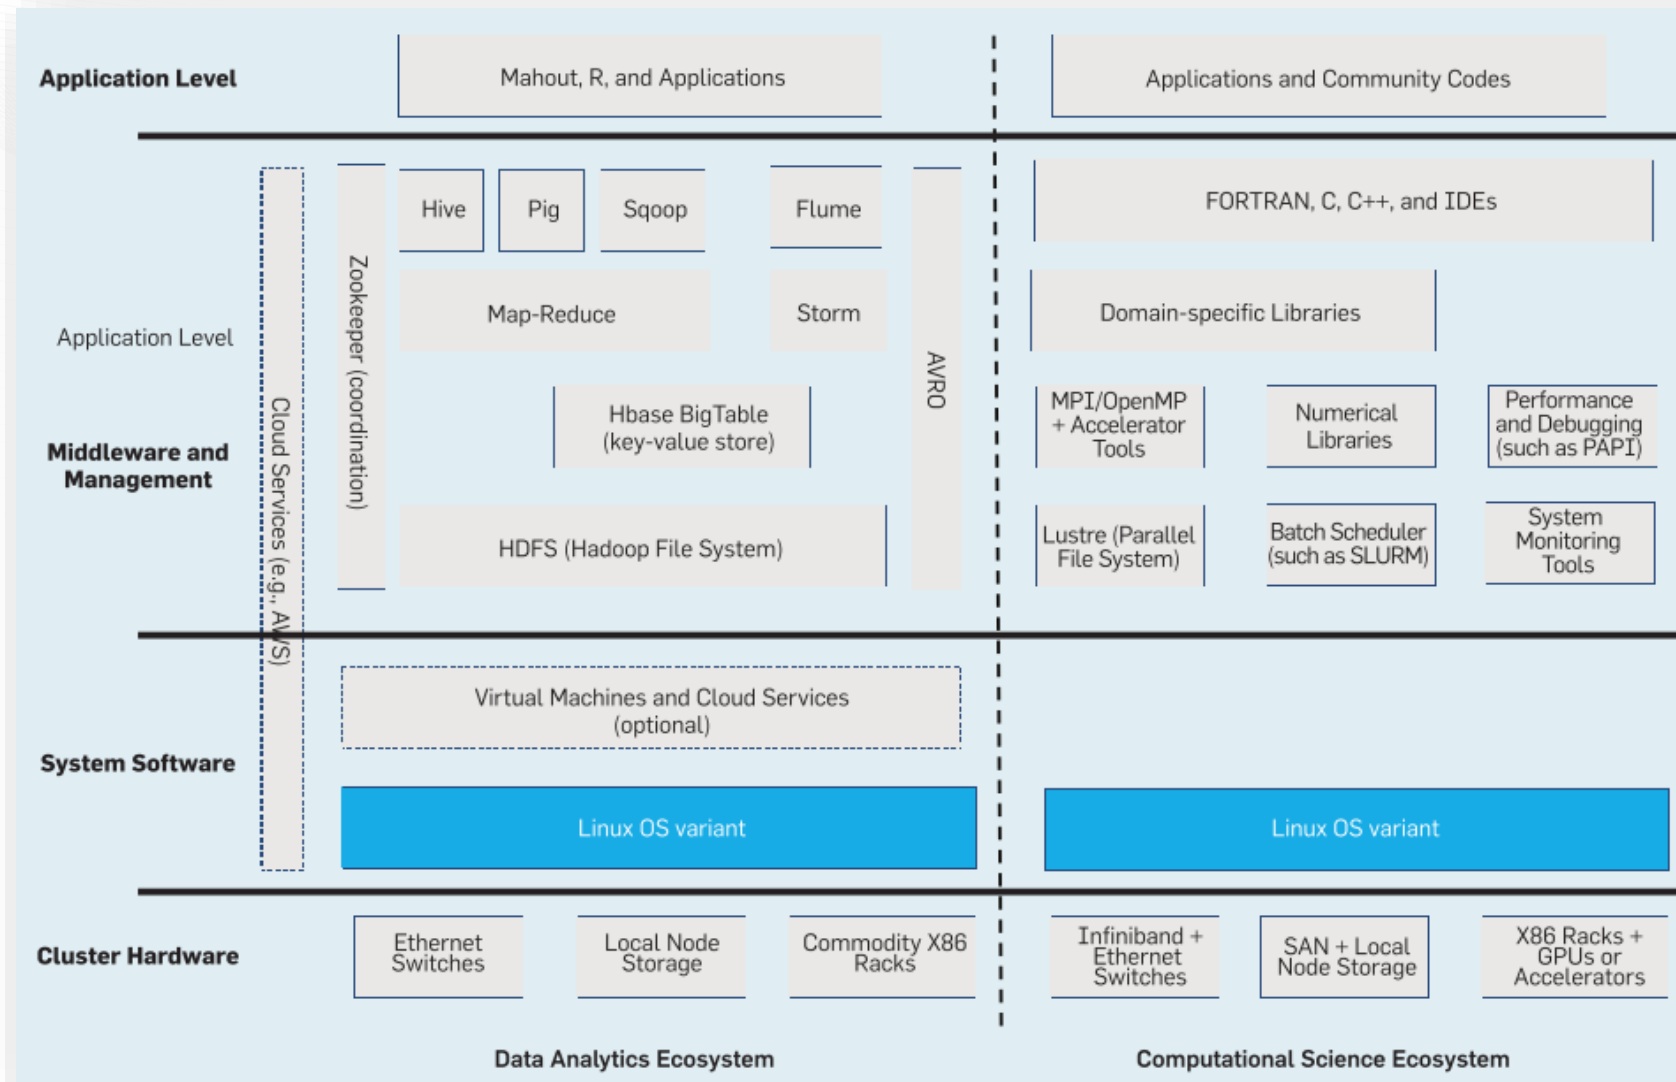
# #4: Software and applications will struggle to survive

System attributes	NERS C Now	OLCF Now	ALCF Now	NERSC Upgrade	OLCF Upgrade	ALCF Upgrades	
Planned Installation	<b>Edison</b>	<b>TITAN</b>	<b>MIRA</b>	<b>Cori 2016</b>	<b>Summit 2017-2018</b>	<b>Theta 2016</b>	<b>Aurora 2018-2019</b>
System peak (PF)	2.6	27	10	> 30	150	>8.5	180
Peak Power (MW)	2	9	4.8	< 3.7	10	1.7	13
Total system memory	357 TB	710TB	768TB	~1 PB DDR4 + High Bandwidth Memory (HBM)+1.5PB persistent memory	> 1.74 PB DDR4 + HBM + 2.8 PB persistent memory	>480 TB DDR4 + High Bandwidth Memory (HBM)	> 7 PB High Bandwidth On-Package Memory Local Memory and Persistent Memory
Node performance (TF)	0.460	1.452	0.204	> 3	> 40	> 3	> 17 times Mira
Node processors	Intel Ivy Bridge	AMD Opteron on Nvidia Kepler	64-bit PowerPC A2	Intel Knights Landing many core CPUs Intel Haswell CPU in data partition	Multiple IBM Power9 CPUs & multiple Nvidia Voltas GPUS	Intel Knights Landing Xeon Phi many core CPUs	Knights Hill Xeon Phi many core CPUs
System size (nodes)	5,600 nodes	18,688 nodes	49,152	9,300 nodes 1,900 nodes in data partition	~3,500 nodes	>2,500 nodes	>50,000 nodes
System Interconnect	Aries	Gemin i	5D Torus	Aries	Dual Rail EDR-IB	Aries	2 <sup>nd</sup> Generation Intel Omni-Path Architecture
File System	7.6 PB 168 GB/s, Lustre®	32 PB 1 TB/s, Lustre®	26 PB 300 GB/s GPFS™	28 PB 744 GB/s Lustre®	120 PB 1 TB/s GPFS™	10PB, 210 GB/s Lustre initial	150 PB 1 TB/s Lustre®

- Even today, we do not have a portable solution for applications scientists to prepare for systems arriving soon
- No solutions for portable use of NVM, threading, HBM, ...
- Scientists may have decades of investment in existing software
  - DOE Climate modeling application is nearly 3M lines of code!
  - Must run across available architectures

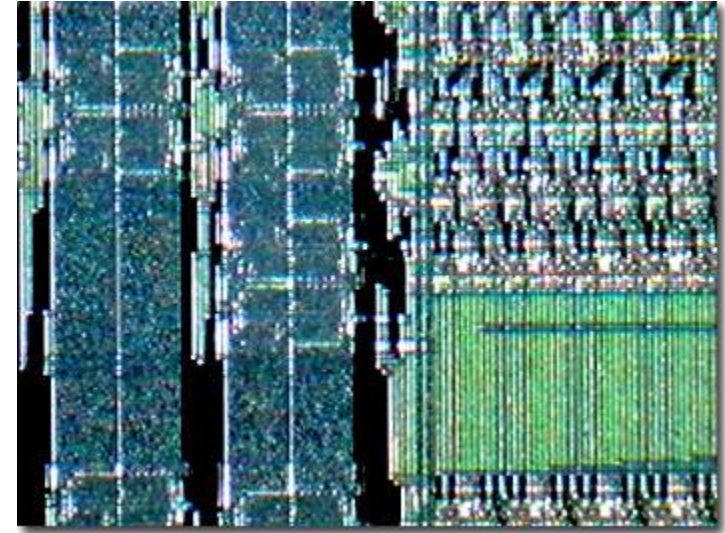


# This challenge will impact all areas of computing: HPC, Cloud, Laptop,...



# #5: Exploration of alternative, potentially disruptive technologies will thrive

- Three decades of alternative technologies have fallen victim to 'curse of Moore's law': general CPU performance improvements without any software changes
  - Weitek Floating Point accelerator (circa 1988)
  - Piles of other types of processors: clearspeed,
  - FPGAs
- Some of these technologies found a specific market to serve
  - But most failed
- Now, the parameters have changed!



<https://micro.magnet.fsu.edu/optics/olympusmicd/galleries/chips/weitekmathmedium.html>



<http://www.clearspeed.com>

# Sixth Wave of Supercomputing: Possible Technology Pathways

# Candidates are flourishing

- **New digital electronics**
  - CNT, memristors, etc
  - Mass customization
  - Reconfigurable computing
  - Millivolt Switches
  - Superconducting electronics (Cryoelectronics)
- **Alternative memory systems including non-volatile memory**
- **Spintronics**
- **Silicon photonics and optical networks**
- **Neuromorphic and brain-inspired computing**
- **Quantum computing**
- **Probabilistic and stochastic computing**
- **Approximate computing**
- **etc**


- **Focus on DOE interests and investments**
  - Non-volatile memory
  - Neuromorphic computing
  - Quantum computing




<http://blogs-images.forbes.com/erikkain/files/2012/03/avatar.jpg>

# #1: Memory Systems


- HMC, HBM/2/3, LPDDR4, GDDR5X, WIDEIO2, etc
- 2.5D, 3D Stacking
- New devices (ReRAM, PCRAM, STT-MRAM, Xpoint)
- Configuration diversity
  - Fused, shared memory
  - Scratchpads
  - Write through, write back, etc
  - Consistency and coherence protocols
  - Virtual v. Physical, paging strategies



[http://gigglehd.com/zbxe/files/attach/images/1404665/988/406/011/788d3ba1967e2db3817d259d2e83c88e\\_1.jpg](http://gigglehd.com/zbxe/files/attach/images/1404665/988/406/011/788d3ba1967e2db3817d259d2e83c88e_1.jpg)  
Copyright (c) 2014 Hiroshige Goto All rights reserved.



[https://www.micron.com/~media/track-2-images/content-images/content\\_image\\_hmc.jpg?la=en](https://www.micron.com/~media/track-2-images/content-images/content_image_hmc.jpg?la=en)



**Fig. 4.** (a) A typical 1T1R structure of RRAM with HfO<sub>2</sub>; (b) HR-TEM image of the TiN/Ti/HfO<sub>2</sub>/TiN stacked layer; the thickness of the HfO<sub>2</sub> is 20 nm.

H.S.P. Wong, H.Y. Lee, S. Yu et al., "Metal-oxide RRAM," *Proceedings of the IEEE*, 100(6):1951-70, 2012.

	SRAM	DRAM	eDRAM	2D NAND Flash	3D NAND Flash	PCRAM	STTRAM	2D ReRAM	3D ReRAM
Data Retention	N	N	N	Y	Y	Y	Y	Y	Y
Cell Size (F <sup>2</sup> )	50-200	4-6	19-26	<1	<1	4-10	8-40	4	<1
Minimum F demonstrated (nm)	14	25	22	16	64	20	28	27	24
Read Time (ns)	<1	30	5	10 <sup>2</sup>	10 <sup>2</sup>	10-50	3-10	10-50	10-50
Write Time (ns)	<1	50	5	10 <sup>2</sup>	10 <sup>2</sup>	100-500	3-10	10-50	10-50
Number of Rewrites	10 <sup>10</sup>	10 <sup>10</sup>	10 <sup>10</sup>	10 <sup>5-10</sup>	10 <sup>5-10</sup>	10 <sup>5-10</sup>	10 <sup>7</sup>	10 <sup>5-10</sup>	10 <sup>5-10</sup>
Read Power	Low	Low	Low	High	High	Low	Medium	Medium	Medium
Write Power	Low	Low	Low	High	High	High	Medium	Medium	Medium
Power (other than R/W)	Leakage	Refresh	Refresh	None	None	None	None	Sneak	Sneak
Maturity									

J.S. Vetter and S. Mittal, "Opportunities for Nonvolatile Memory Systems in Extreme-Scale High Performance Computing," *CISE*, 17(2):73-82, 2015.

# Emerging Memory Technologies : Nonvolatile Memory

	SRAM	DRAM	eDRAM	2D NAND Flash	3D NAND Flash	PCRAM	STTRAM	2D ReRAM	3D ReRAM
Data Retention	N	N	N	Y	Y	Y	Y	Y	Y
Cell Size (F <sup>2</sup> )	50-200	4-6	19-26	2-5	<1	4-10	8-40	4	<1
Minimum F demonstrated (nm)	14	25	22	16	64	20	28	27	24
Read Time (ns)	< 1	30	5	10 <sup>4</sup>	10 <sup>4</sup>	10-50	3-10	10-50	10-50
Write Time (ns)	< 1	50	5	10 <sup>5</sup>	10 <sup>5</sup>	100-300	3-10	10-50	10-50
Number of Rewrites	10 <sup>16</sup>	10 <sup>16</sup>	10 <sup>16</sup>	10 <sup>4</sup> -10 <sup>5</sup>	10 <sup>4</sup> -10 <sup>5</sup>	10 <sup>8</sup> -10 <sup>10</sup>	10 <sup>15</sup>	10 <sup>8</sup> -10 <sup>12</sup>	10 <sup>8</sup> -10 <sup>12</sup>
Read Power	Low	Low	Low	High	High	Low	Medium	Medium	Medium
Write Power	Low	Low	Low	High	High	High	Medium	Medium	Medium
Power (other than R/W)	Leakage	Refresh	Refresh	None	None	None	None	Sneak	Sneak
Maturity									

IBM TLC PC  
Intel/Micron Xpoint?

# NVRAM Technology Continues to Improve – Driven by Market Forces



## designlines MEMORY

### Blog

## First Look at Samsung's 48L 3D V-NAND Flash

Kevin Gibb, Product Line Manager, TechInsights

4/6/2016 04:40 PM EDT

9 comments post a comment

Like 16 Tweet in Share 61 G+ 2

NO RATINGS  
LOGIN TO RATE

The highly anticipated Samsung's 48 layer V-NAND 3D flash memory is out in the market, and we at TechInsights have the first look.

Samsung had announced its 256 Gb 3-bit multi-level cell K9AFGY8S0M 3D V-NAND as early as August 2015, stating that it would be used in their 2 T Figure 1.



[http://www.eetasia.com/STATIC/ARTICLE\\_IMAGES/2012/EEOL\\_2012DEC28\\_STOR\\_MFG\\_NT\\_01.jpg](http://www.eetasia.com/STATIC/ARTICLE_IMAGES/2012/EEOL_2012DEC28_STOR_MFG_NT_01.jpg)

## designlines WIRELESS & NETWORKING

### Slideshow

## Facebook Likes Intel's 3D XPoint

Google joins open hardware effort

Rick Merritt

3/10/2016 07:56 AM EST

7 comments

NO RATINGS  
LOGIN TO RATE

Like 115 Tweet in Share 46 G+ 3

SAN JOSE, Calif.—Facebook said it hopes to use Intel's emerging 3D XPoint memories in its data centers. Meanwhile Google joins its archival's open hardware efforts to drive standards ranging from high-power compute racks to giant form factors for disk drives.


The two moves were likely the highest impact announcements at the annual event of the Facebook-led Open Compute Project (OCP) here. Among other news, Intel showed a new 16-core Xeon SoC with dual 10G Ethernet controllers and a prototype chip merging Xeon with an Arria FPGA in a single package.

volatile memories be endorsement," said TechInsights

May 18, 2016

## IBM Puts 3D XPoint on Notice with 3 Bits/Cell PCM Breakthrough

Tiffany Trader



IBM scientists have broken new ground in the development of a phase change memory technology (PCM) that puts a target on competing XPoint technology from Intel and Micron. IBM successfully stored 3 per cell in a 64k-cell array that had been pre-cycled 1 million times: exposed to temperatures up to 75°C. A paper describing the advance was presented this week at the IEEE International Memory Workshop in Paris.

Phase-change memory is an up-and-coming non-volatile memory technology—a storage-class memory that bridges the divide between expensive, volatile memory (namely DRAM), and slower persistent storage (flash or hard disk drives). According to IBM, having the ability to reliably fit 3 bits per cell is what will make this technology price-competitive flash.

With memory demands riding the tide of big data, phase change memory has a lot to recommend it but to be a success, the economics must work, say the authors, and being able to store multiple bits per memory cell is essential for keeping costs under control.

Using a combination of electrical sensing techniques and signal processing technologies, the researchers have shown for the first time the viability of Triple-Level-Cell (TLC) storage in phase-change memory cells. The researchers addressed challenges related to multi-bit PCM including drift, variability, temperature sensitivity and endurance cycling with two innovative enabling technologies:

(a) an advanced, nonresistance cell-state metric that exhibits robustness to drift and PCM noise, and (b) an advanced level-detection and modulation-coding framework that enables further resilience to drift, noise and temperature variation effects.

Original URL: [http://www.theregister.co.uk/2013/11/01/hp\\_memristor\\_2018/](http://www.theregister.co.uk/2013/11/01/hp_memristor_2018/)

HP 100TB Memristor drives by 2018 – if you're lucky, admits tech titan  
Universal memory slow in coming


By Chris Mellor

Posted in Storage, 1st November 2013 02:28 GMT

**Blocks and Files** HP has warned *El Reg* not to get its hopes up too high after the tech titan's CTO Martin Fink suggested StoreServ arrays could be packed with 100TB Memristor drives come 2018.

In five years, according to Fink, DRAM and NAND scaling will hit a wall, limiting the maximum capacity of the technologies: process shrinks will come to a shuddering halt when the memories' reliability drops off a cliff as a side effect of reducing the size of electronics on the silicon dies.

The HP answer to this scaling wall is Memristor, its flavour of resistive RAM technology that is supposed to have DRAM-like speed and better-than-NAND storage density. Fink claimed at an HP Discover event in Las Vegas that Memristor devices will be ready by the time flash NAND hits its limit in five years. He also showed off a Memristor wafer, adding that it could have a 1.5PB capacity by the end of the decade.



## Forbes / Tech

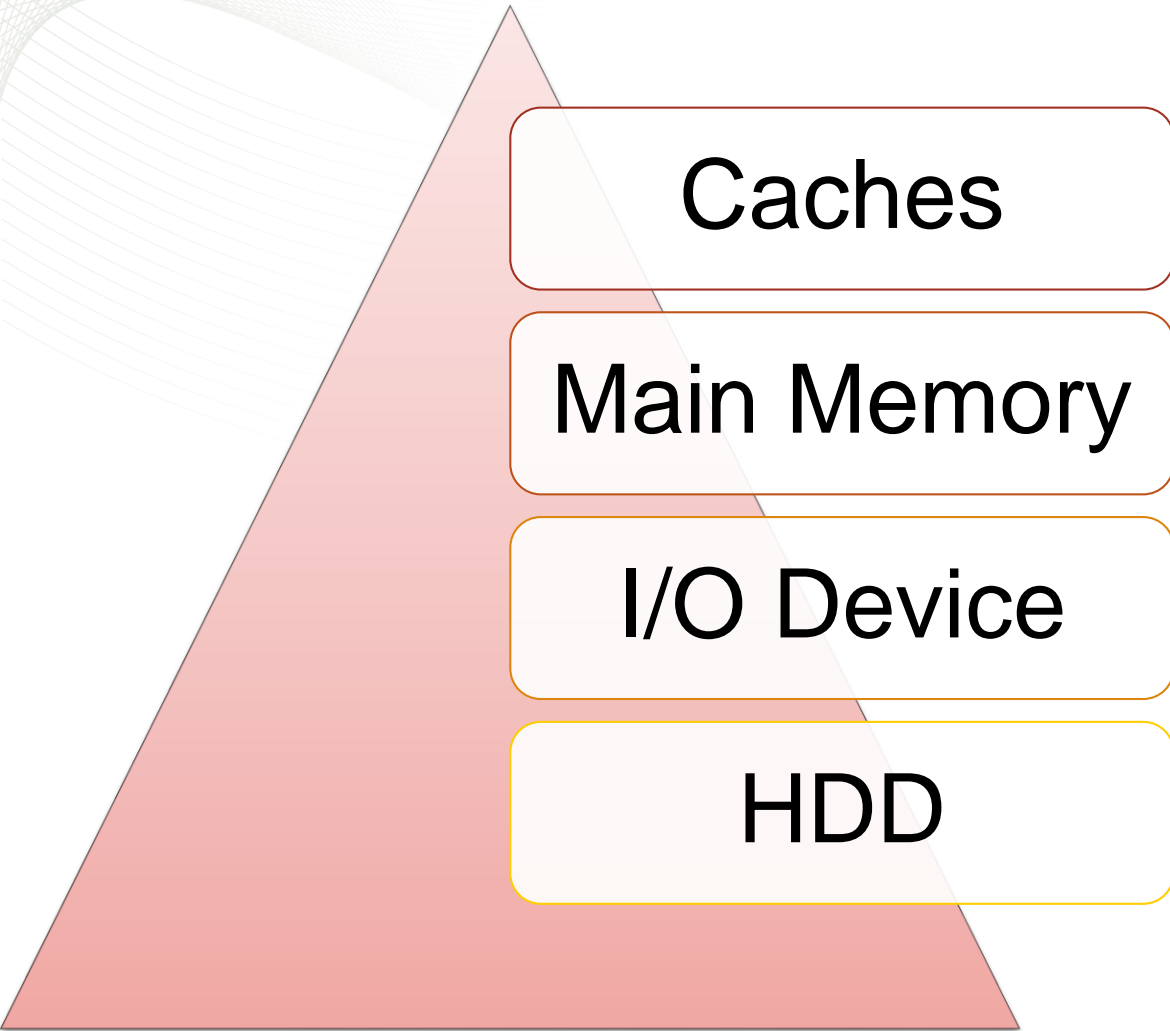
JUL 28, 2015 @ 2:46 PM 7,391 VIEWS

## Intel And Micron Jointly Announce Game-Changing 3D XPoint Memory Technology






# As NVM improves, it is working its way toward the processor core



- Newer technologies improve
  - density,
  - power usage,
  - durability
  - r/w performance
- In scalable systems, a variety of architectures exist
  - NVM in the SAN
  - NVM nodes in system
  - NVM in each node
- Expect energy efficient, cheap, vast capacity

## #2: Brain-inspired or Neuromorphic Computing

- Concept and term developed by Carver Mead in late 1980s describing use of electronic (analog) circuits to mimic neurobiological architectures
  - Neurons and synapses
- Why?
  - Energy, space efficiency
  - Plasticity, Flexible of dynamic learning
- Examples of recent work
  - IBM True North chip
  - Human Brain Project: SpiNNaker chip





P.A. Merolla, J.V. Arthur, R. Alvarez-Icaza et al., "A million spiking-neuron integrated circuit with a scalable communication network and interface," *Science*, 345(6197):668-73, 2014, doi:10.1126/science.1254642.

# IBM True North

- Simulates complex neural networks
  - 5.4B transistors / CMOS
  - One million individually programmable neurons-- sixteen times more than the current largest neuromorphic chip
  - 256 million individually programmable synapses on chip which is a new paradigm
  - 5.4B transistors. By device count, largest IBM chip ever fabricated, second largest (CMOS) chip in the world
  - 4,096 parallel and distributed cores, interconnected in an on-chip mesh network
  - Over 400 million bits of local on-chip memory (~100 Kb per core) to store synapses and neuron parameters
- Can be scaled with inter-chip communication interface
- 70mW total power while running a typical recurrent network at biological real-time, four orders of magnitude lower than a conventional computer running the same network
- NN trained offline

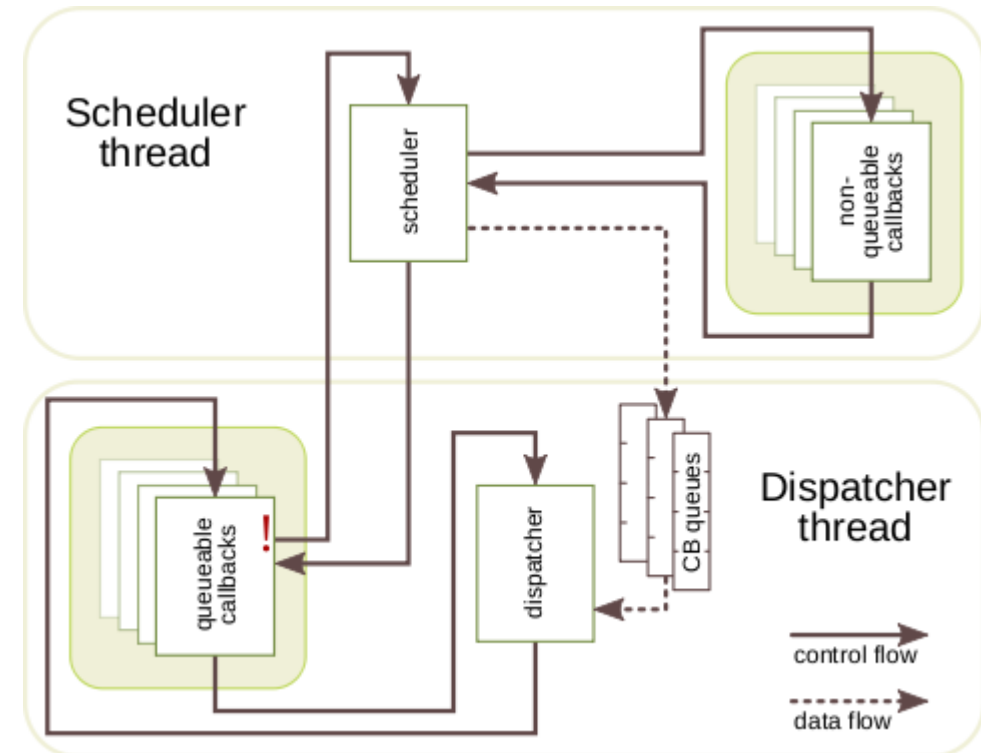

<http://www.research.ibm.com/articles/brain->



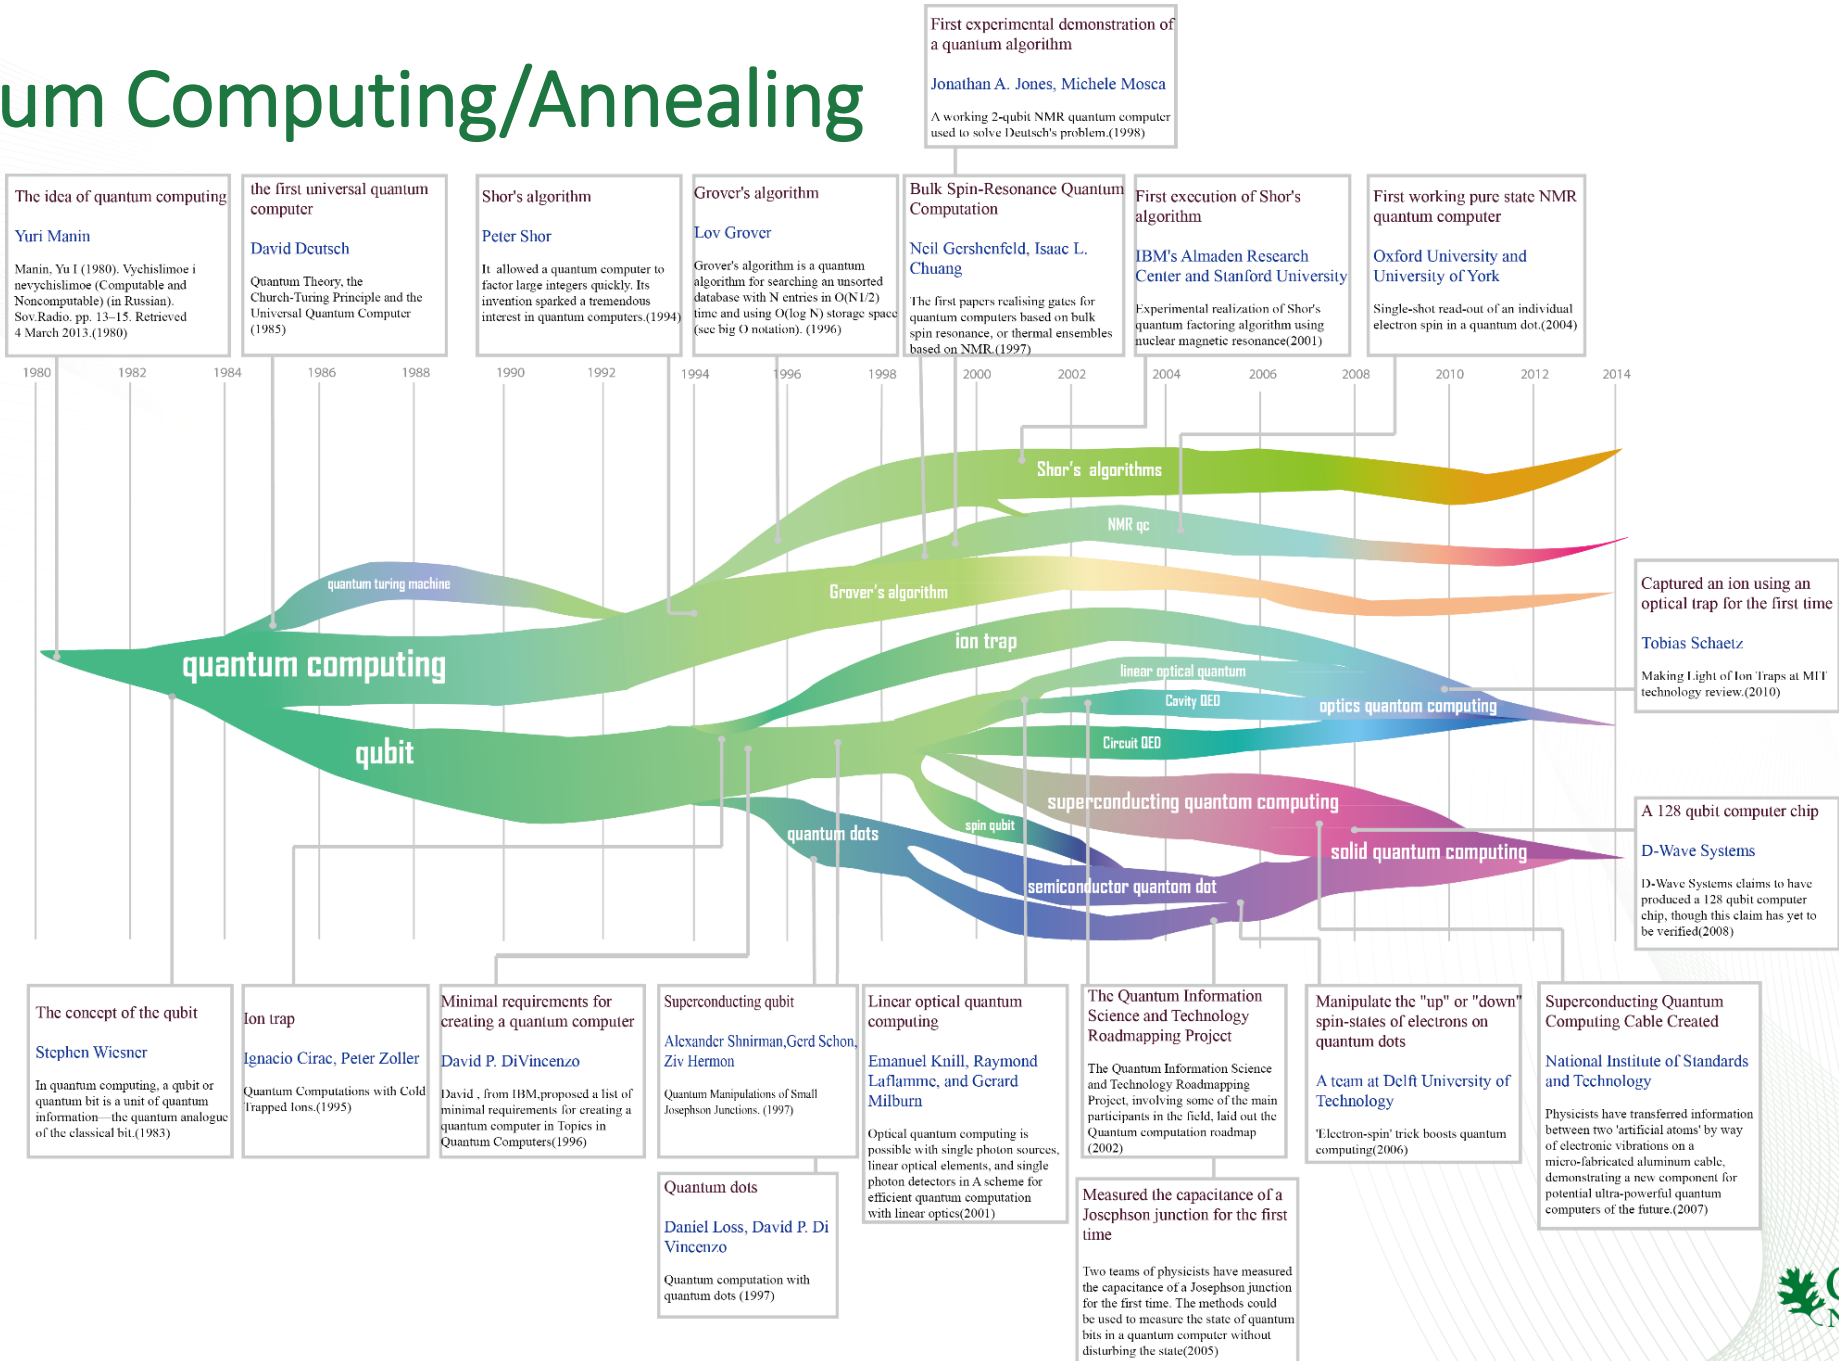
By DARPA SyNAPSE - <http://www.darpa.mil/NewsEvents/Releases/2014/08/07.aspx>, Public Domain  
<https://commons.wikimedia.org/w/index.php?curid=34614979>

# SpiNNaker of Human Brain Project

- Modelling spiking neural networks
- Excellent energy efficiency
- Globally Asynchronous Locally Synchronous (GALS) system with 18 ARM968 processor nodes residing in synchronous islands, surrounded by a light-weight, packet-switched asynchronous communications infrastructure.
- Eventual goal is to be able to simulate a single network consisting of one billion simple neurons, requiring a machine with over 50,000 chips.
- Programmed with Neural Engineering Framework
- Demonstrated on vision, robotic tasks



# #3: Quantum Computing/Annealing



# Designs strategies abound

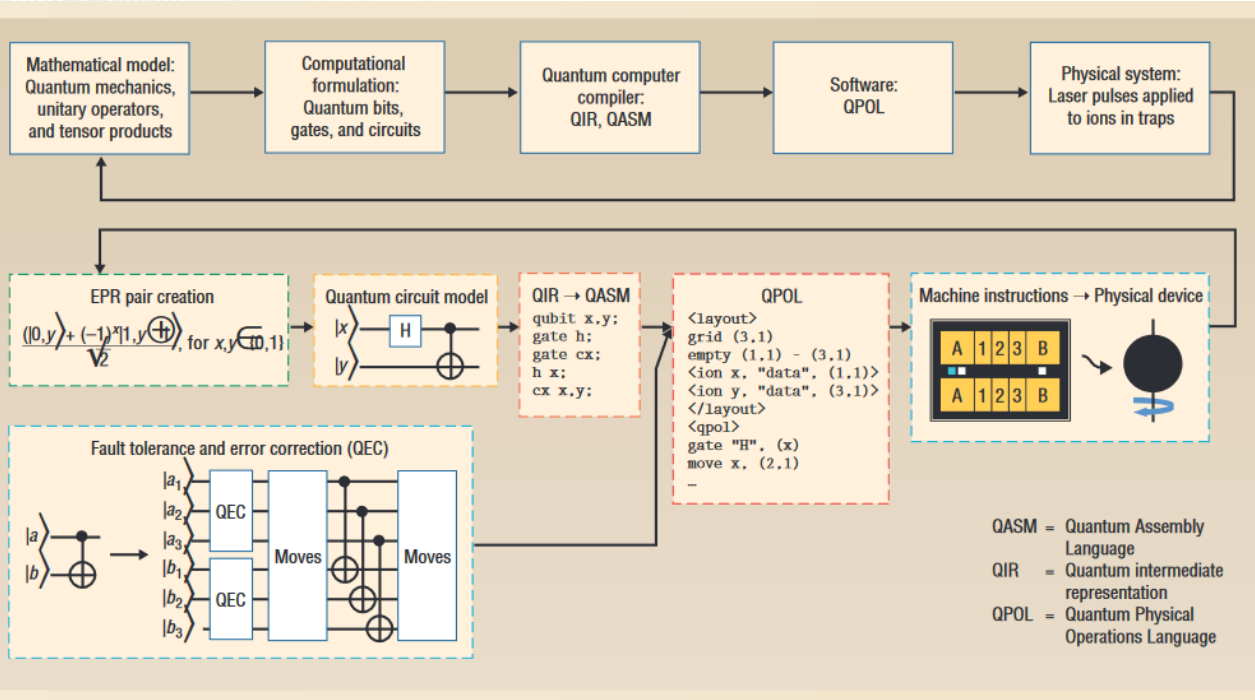


Figure A. Using quantum information processing to control live physical systems. Proposed four-phase design flow, detailed for EPR pair creation on a trapped-ion computer with machine instructions translated into a sequence of laser pulses that perform a CNOT gate. A feedback loop allows for repetition of earlier phases.

K.M. Svore, A.V. Aho, A.W. Cross, I. Chuang, and I.L. Markov, "A layered software architecture for quantum computing design tools," *IEEE Computer*, 39(1):74-83, 2006, doi:10.1109/IC.2006.4.

*Math. Struct. in Comp. Science* (2006), vol. 16, pp. 581–600. © 2006 Cambridge University Press  
doi:10.1017/S0960129506005378 Printed in the United Kingdom

## Quantum programming languages: survey and bibliography

SIMON J. GAY

### Depa: Cognitive Computing Programming Paradigm: A Corelet Language for Composing Networks of Neurosynaptic Cores

Recei

Arnon Amir, Pallab Datta, William P. Risk, Andrew S. Cassidy, Jeffrey A. Kusnitz, Steve K. Esser, Alexander Andreopoulos, Theodore M. Wong, Myron Flickner, Rodrigo Alvarez-Icaza, Emmett McQuinn, Ben Shaw, Norm Pass, and Dharmendra S. Modha  
IBM Research - Almaden, San Jose, CA 95120

The f  
surpr  
langu  
techn  
quant

**Abstract**—Marching along the DARPA Synapse roadmap, IBM unveils a trilogy of innovations towards the TrueNorth cognitive computing system inspired by the brain's function and efficiency. The sequential programming paradigm of the von Neumann architecture is wholly unsuited for TrueNorth. Therefore, as our main contribution, we develop a new programming paradigm that permits construction of complex cognitive

TrueNorth architecture—that was featured on the covers of *Science* [8] and *Communications of the ACM* [1].

We unveil a series of interlocking innovations in a set of three papers. In this paper, we present a programming paradigm for hierarchically composing and configuring cognitive systems that is effective for the programmer and ef-

## ScaffCC: A Framework for Compilation and Analysis of Quantum Computing Programs

Ali JavadiAbhari\*, Shruti Patil\*, Daniel Kudrow†, Jeff Heckey†, Alexey Lvov\*\*, Frederic T. Chong†, Margaret Martonosi\*

# D-wave has operational Adiabatic Quantum Computer for solving optimization applications

## Adiabatic Quantum Annealing



**Problem:** find the ground state of

$$H_{\text{Ising}} = \sum_j h_j \sigma_j^z + \sum_{(i,j) \in E} J_{ij} \sigma_i^z \sigma_j^z$$


Shown by Barahona (1982) to be NP-hard in 2D,  $J_{ij} = \pm 1$ ,  $h_j \neq 0$ .

**Use adiabatic interpolation from transverse field** (Farhi et al., 2000)

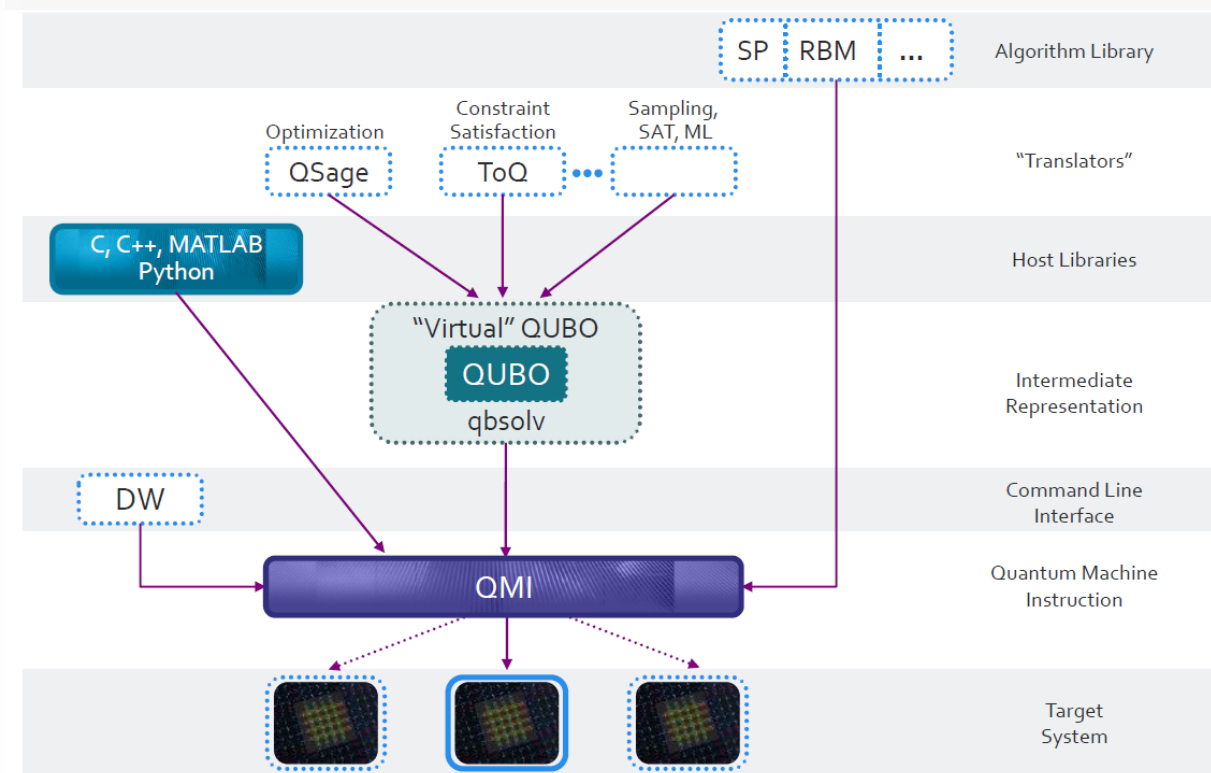
$$H(t) = A(t) \sum_j \sigma_j^x + B(t) H_{\text{Ising}}$$

$t \in [0, t_f]$

Program  $\{h_i, \{J_{ij}\}$



Graph Embedding implemented on DW-1 via Chimera graph retains NP-hardness (V. Choi, 2010)



USC

School of Engineering

Introduction | 3

University of Southern California

B. Lucas, ISI, 2015

D:wave  
The Quantum Computing Company™

© 2016 D-Wave Systems Inc. All Rights Reserved | 2

# Summary



# Disruption in Computing Stack (== research opportunities)

Layer	Switch, 3D	NVM	Approximate	Neuro	Quantum
<i>Application</i>	1	1	2	2	3
<i>Algorithm</i>	1	1	2	3	3
<i>Language</i>	1	2	2	3	3
<i>API</i>	1	2	2	3	3
<i>Arch</i>	1	2	2	3	3
<i>ISA</i>	1	2	2	3	3
<i>Microarch</i>	2	3	2	3	3
<i>FU</i>	2	3	2	3	3
<i>Logic</i>	3	3	2	3	3
<i>Device</i>	3	3	2	3	3

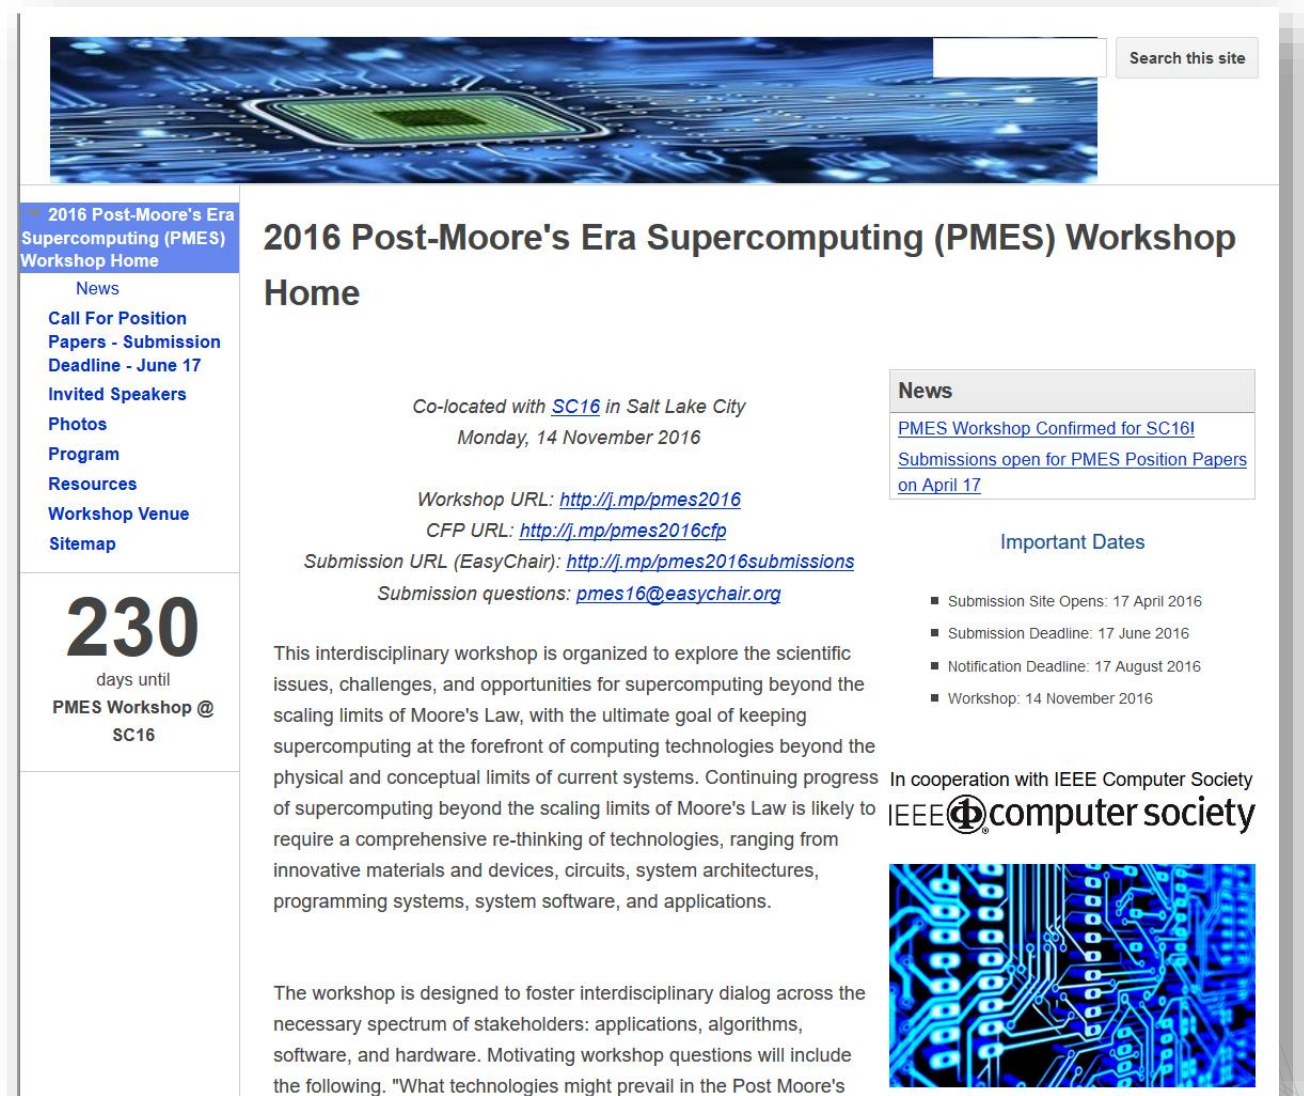
Adapted from IEEE Rebooting Computing Chart

# Take Away Messages

1. Moore's Law is definitely ending for either economic or technical reasons
2. Specialization – use the same transistors differently
3. Architecting effective solutions will be critical
4. CMOS continues indefinitely
5. Parallelism – our area of expertise – will continue to be the major contributor to performance improvements in HPC, enterprise for moving forward for the next decade
  1. Interconnect and memory bandwidth and capacity will need to improve
6. Our community should aggressively pursue disruptive technologies
  1. Some technologies could disrupt entire stack
7. Tremendous challenges in deploying these technologies with existing software
  1. Many opportunities to provide new software frameworks for fundamental computer science problems: resource management, mapping, programming models, portability, etc.
8. Start talking to your colleagues in physics, chemistry, electrical engineering, etc
9. If applications suffer, so will we!

# 2016 Post-Moores Era Supercomputing Workshop @ SC16

- <https://j.mp/pmes2016>
- @SC16
- Position papers due June 17



The screenshot shows the workshop's home page. At the top is a banner image of a green microchip on a blue circuit board. Below the banner is a search bar. The main heading is "2016 Post-Moore's Era Supercomputing (PMES) Workshop Home". To the left is a navigation menu with links for News, Call For Position Papers - Submission Deadline - June 17, Invited Speakers, Photos, Program, Resources, Workshop Venue, and Sitemap. A large "230" indicates the number of days until the workshop. The main content area includes the location (Salt Lake City, SC16), dates (November 14, 2016), and various URLs for the workshop, CFP, and submissions. A "News" section contains two items: "PMES Workshop Confirmed for SC16!" and "Submissions open for PMES Position Papers on April 17". An "Important Dates" section lists submission site opening, submission deadline, notification deadline, and the workshop date. The page also features the IEEE Computer Society logo and a description of the workshop's interdisciplinary focus.

2016 Post-Moore's Era Supercomputing (PMES) Workshop Home

News

Call For Position Papers - Submission Deadline - June 17

Invited Speakers

Photos

Program

Resources

Workshop Venue

Sitemap

230 days until PMES Workshop @ SC16

## 2016 Post-Moore's Era Supercomputing (PMES) Workshop Home

Co-located with [SC16](#) in Salt Lake City  
Monday, 14 November 2016

Workshop URL: <http://j.mp/pmes2016>  
CFP URL: <http://j.mp/pmes2016cfp>  
Submission URL (EasyChair): <http://j.mp/pmes2016submissions>  
Submission questions: [pmes16@easychair.org](mailto:pmes16@easychair.org)

This interdisciplinary workshop is organized to explore the scientific issues, challenges, and opportunities for supercomputing beyond the scaling limits of Moore's Law, with the ultimate goal of keeping supercomputing at the forefront of computing technologies beyond the physical and conceptual limits of current systems. Continuing progress of supercomputing beyond the scaling limits of Moore's Law is likely to require a comprehensive re-thinking of technologies, ranging from innovative materials and devices, circuits, system architectures, programming systems, system software, and applications.

The workshop is designed to foster interdisciplinary dialog across the necessary spectrum of stakeholders: applications, algorithms, software, and hardware. Motivating workshop questions will include the following. "What technologies might prevail in the Post Moore's


News

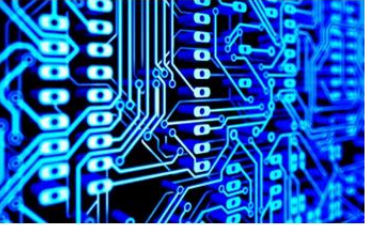
[PMES Workshop Confirmed for SC16!](#)  
[Submissions open for PMES Position Papers on April 17](#)

Important Dates

- Submission Site Opens: 17 April 2016
- Submission Deadline: 17 June 2016
- Notification Deadline: 17 August 2016
- Workshop: 14 November 2016

In cooperation with IEEE Computer Society

IEEE  computer society



# Acknowledgements



- Contributors and Sponsors

- Future Technologies Group: <http://ft.ornl.gov>
- US Department of Energy Office of Science
  - DOE Vancouver Project: <https://ft.ornl.gov/trac/vancouver>
  - DOE Blackcomb Project: <https://ft.ornl.gov/trac/blackcomb>
  - DOE ExMatEx Codesign Center: <http://codesign.lanl.gov>
  - DOE Cesar Codesign Center: <http://cesar.mcs.anl.gov/>
  - DOE Exascale Efforts:  
<http://science.energy.gov/ascr/research/computer-science/>
- Scalable Heterogeneous Computing Benchmark team:  
<http://bit.ly/shocmarx>
- US National Science Foundation Keeneland Project:  
<http://keeneland.gatech.edu>
- US DARPA
- NVIDIA CUDA Center of Excellence



# Bonus Material