## Preparing for <br> Supercomputing's Sixth Wave

Jeffrey S. Vetter



## Presented to

ACM Symposium on High-Performance Parallel and Distributed Computing

2 June 2016
Kyoto


## 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




## End of Moore's Law Really ??

## Contemporary devices are approaching fundamental limits



Figure $1 \mid$ As a metal oxide-semiconductor field effect transistor (MOSFET) shrinks, the gate dielectric (yellow) thickness approaches several atoms ( 0.5 nm at the $22-\mathrm{nm}$ 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
I.L. Markov, "Limits on fundamental limits to computation," Nature, 512(7513):147-54, 2014, doi:10.1038/nature13570.

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



## 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 LOGIN TO RATE
10/3/2012 03:00

## f Like $4 \quad y$ Tweet in Share $\mathbf{G + 1} 0$

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 $450-\mathrm{mm}$ manufacturing will achieve the goal of 3 r percent cost reduction. But that equates to only three or four years of ir reasing equipment costs, and consequently, delays the inevitable, G Aner said. It is also possible that new technologies will emerge that will s'Jw the rate of cost increases, according to the firm.

According to Gartner, the costs of manufacturing $\varsigma$ quipment needed for leading edge semiconductor manufacturing are increasir $g$ at a rate between 7 percent and 10 percent per year, depending on the bp sic process.
By 2020, current cost trends will lead tc 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 budger 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. 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)

| $\begin{aligned} & \hline 2013 \\ & \text { Rank } \\ & \hline \end{aligned}$ | $\begin{aligned} & 2012 \\ & \text { Rank } \\ & \hline \end{aligned}$ | Company | Foundry Type | Location | $\begin{gathered} 2011 \text { Sales } \\ \text { (SM) } \\ \hline \end{gathered}$ | $\begin{gathered} 2012 \text { Sales } \\ \text { (SM) } \end{gathered}$ | $\begin{array}{\|c\|} \hline 2012 / 2011 \\ \text { Change (\%) } \\ \hline \end{array}$ | $\begin{array}{\|c\|} \hline 2013 \text { Sales } \\ \text { (SM) } \end{array}$ | $\begin{gathered} \hline 2013 / 2012 \\ \text { Change }(\%) \\ \hline \end{gathered}$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 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\% |
| - | - | Top 13 Total | - | - | 28,464 | 33,951 | 19\% | 38,920 | 15\% |
| - | - | Top 13 Share | - | - | 89\% | 90\% |  | 91\% |  |
| - | - | Other Foundry | - | - | 3,446 | 3,669 | 6\% | 3,920 | 7\% |
| - | - | Total Foundry | - | - | 31,910 | 37,620 | 18\% | 42,840 | 14\% |

Source: IC Insights, compary reports Wuhan Xinxin (now XMC) for 2012 or 2013. .
OAK RIDGE
*-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
NO RATINGS
LOGIN TO RATE

## 12/2/2015 10:13 AM EST

- comments post a comment
$f$ Like $10 \quad$ Tweet in share 81 G+1 1
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
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 and acquisition-happy industry's DNA from the get go. But, even in and acquisionesure.
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
thase imosisiswer
Avago Agrees to Buy Broadcom for $\$ 37$ Billion
(a) Reuters

ву Mchase J.de a merceb and chad bray mar 28,2015


SANDISK COMPLETES ACQUISITION OF FUSION

## JUL 23, 2014

ACQUIIITION TO BOOST SANDISK'S ENTERPRISE GROWTH


Western Digital Now A Storage Powerhouse With the Fusion-io team will accelerate our efforts to enable the fl the industry:

## SanDisk Acquisition

the industry:


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


## 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


Sources: Intel; press reports; The Economist 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


类 NAK RItional Laboratory National Laboratory

## 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 demonstrations 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.





## 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



## Sixth Wave 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


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


g.00GLe buILT ITS very Own CIIIPS TO POWER ITS AI BOTS


回 0006 LE
godgle 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

## \#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



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



## Improved Stacking, Vias, Communication Techniques

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

$\square$ Wafer thinning has been stuck at $\sim 40 \mu$ 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.
$\square$ 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 $4 \mu \mathrm{~m}$ using $300-\mathrm{mm}$ Wafer proven by $40-\mathrm{nm}$ 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)

Communication is via magnetic field

$\mu_{\mathrm{si}}=\mu_{\mathrm{siO}_{2}}=1$
Can easily induce a 200 mV signal in receiver coil.

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

## \#4: Software and applications will struggle to survive

| System attributes | NERS <br> C <br> Now | OLCF <br> Now | ALCF Now | NERSC <br> Upgrade | OLCF Upgrade | ALCF Upgrades |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Planned Installation | Edison | TITAN | MIRA | $\begin{aligned} & \text { Cori } \\ & 2016 \end{aligned}$ | $\begin{aligned} & \text { Summit } \\ & \text { 2017-2018 } \end{aligned}$ | $\begin{aligned} & \text { Theta } \\ & 2016 \end{aligned}$ | $\begin{gathered} \text { Aurora } \\ \text { 2018-2019 } \end{gathered}$ |
| 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 | $>1.74 \mathrm{~PB}$ <br> DDR4 + <br> HBM + 2.8 <br> PB <br> persistent memory | >480 TB DDR4 <br> + High <br> Bandwidth <br> Memory (HBM) | > 7 PB High Bandwidth OnPackage 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 <br> Opter <br> on <br> Nvidia <br> Kepler | 64-bit PowerP C 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) | $\begin{aligned} & 5,600 \\ & \text { nodes } \end{aligned}$ | $\begin{gathered} 18,68 \\ 8 \\ \text { nodes } \end{gathered}$ | 49,152 | 9,300 nodes 1,900 nodes in data partition | $\begin{gathered} \sim 3,500 \\ \text { nodes } \end{gathered}$ | >2,500 nodes | >50,000 nodes |
| System Interconnect | Aries | Gemin i | $\begin{gathered} \text { 5D } \\ \text { Torus } \end{gathered}$ | Aries | Dual Rail EDR-IB | Aries | $2^{\text {nd }}$ Generation Intel Omni-Path Architecture |
| File System | $\begin{aligned} & \text { 7.6 PB } \\ & 168 \\ & \text { GB/s, } \\ & \text { Lustre }^{\circledR} \end{aligned}$ | $\begin{gathered} 32 \mathrm{~PB} \\ 1 \\ \mathrm{~TB} / \mathrm{s}, \\ \text { Lustre } \end{gathered}$ | $\begin{gathered} 26 \mathrm{~PB} \\ 300 \\ \mathrm{~GB} / \mathrm{s} \\ \mathrm{GPFS}^{\mathrm{TM}} \end{gathered}$ | $\begin{gathered} 28 \mathrm{~PB} \\ 744 \mathrm{~GB} / \mathrm{s} \\ \text { Lustre }{ }^{\circledR} \end{gathered}$ | $\begin{gathered} 120 \mathrm{~PB} \\ 1 \mathrm{~TB} / \mathrm{s} \\ \mathrm{GPFS}^{\mathrm{TM}} \end{gathered}$ | 10PB, 210 GB/s Lustre initial | $\begin{aligned} & 150 \mathrm{~PB} \\ & 1 \mathrm{~TB} / \mathrm{s} \\ & \text { Lustre }^{\circledR} \end{aligned}$ |

- 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!



## 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



## \#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


Fig. 4. (a) A typical ITIR structre of RRAM with Hfo (b) HR TEA mage of the TiN/Ti/Hfo,/Tin stacked layer; the Hforckess of the $\mathrm{HfO}_{2}$ is 20 nm .
H.S.P. Wong H Y Lee S. Yuet al "Metal-oxide RRAM" H.S.P. Wong, H.Y. Lee, S. Yu et al., Metal-oxide
Proceedings of the IEEE, $100(6)$ :1951-70, 2012.
https://www.micron.com/~/media/track-2-2images/content-images/content image hmc.jpg?la=en


## Emerging Memory Technologies : Nonvolatile Memory

|  | SRAM | DRAM | eDRAM | 2D <br> NAND <br> Flash | 3D <br> NAND <br> Flash | PCRAM | STTRAM | 2D <br> ReRAM | 3D <br> ReRAM |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Data Retention | N | N | N | Y | Y | Y | Y | Y | Y |
| Cell Size $\left(\mathrm{F}^{2}\right)$ | $50-200$ | $4-6$ | $19-26$ | $2-5$ | $<1$ | $4-10$ | $8-40$ | 4 | $<1$ |
| Minimum F demonstrated <br> $(\mathrm{nm})$ | 14 | 25 | 22 | 16 | 64 | 20 | 28 | 27 | 24 |
| Read Time (ns) | $<1$ | 30 | 5 | $10^{4}$ | $10^{4}$ | $10-50$ | $3-10$ | $10-50$ | $10-50$ |
| Write Time (ns) | $<1$ | 50 | 5 | $10^{5}$ | $10^{5}$ | $100-300$ | $3-10$ | $10-50$ | $10-50$ |
| Number of Rewrites | $10^{16}$ | $10^{16}$ | $10^{16}$ | $10^{4}-10^{5}$ | $10^{4}-10^{5}$ | $10^{8}-10^{10}$ | $10^{15}$ | $10^{8}-10^{12}$ | $10^{8}-10^{12}$ |
| 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 VNAND Flash
Kevin Gibb, Product Line Manager, Techlnsights
/6/2016 04:40 PM ED
Q comments post a comment
f Like 16 Tweet in share


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

Samsung had announced its 256 Gb 3-bit multi-level cell


[^0]
## designlines WIRELESS \& NETWORKING

## Slideshow

Facebook Likes Intel's 3D XPoint Google joins open hardware effort Rick Merritt
$3 / 101201607 \cdot 56$
7 3/101/2016 07:56 AM ES
7 comments


SAN JOSE, Calif.-Facebook said it hopes to use Intel's emerc SD XPoint memories in its data centers. Meanwhile Google join its archrival's open hardware efforts to drive standards ranging
high-power compute racks to giant form factors for disk drives. The two moves were likely the highest impact announcements
the annual event of the Facebook-led Open Compute Proiect (OCP) here Among other news, Intel showed a new 16 -core (OCP) here. Among other news, Intel showed a new 16 -core $x$
SoC with dual 10 G Ethernet controllers and a prototype chip merging Xeon with an Arria FPGA in a single package.
would be used ir on the market in Figure 1 . volatile memories $\mathrm{b} \epsilon$
e endorsement," saii Colatie memories
cendorsement," sa
cher Insight64 endorsement,

May 18,2016
IBM Puts 3D XPoint on Notice with 3 Bits/Cell PCM Breakthrough Tiffany Trader

scientists have broken new ground in the development of a pla change memory technology (PCM) that puts a target on competin er cell in a 64 k-cell array that had been pre-cycled 1 millilion times exposed to temperatures up to $75 \circ \mathrm{C}$. A paper describing the adva was pre
Paris.

Phase-change memory is an up-and-coming non-volatile memory technology - a storage-class memory that bridges the divide betw
expensive performant, volatile memory (namely DRAM), and slower persistent storage (flash or hard disk drive According to $1 B M$, having the ability to eriably fit 3 bits per cell is what will make this technology price-compettivy llash. or keeping costs under control
Using a combination of electrical sensing tec chiques and signal processing technologies, the researchers have
shown for the first ime the the viabity o Triple shown for the first time the the viability of Triple-Level-Cell ( (TLC) storage in phase-change memory cells. The endurance cycling with two innovative enabing technologies
(a) an advanced, nonresistance cell-state metric that exhibits robustress to dritt and PCM noise, and (b) an ad level-detection and modulation-coding framework that enables further resilience to dritit, noise and temperature variation effects.

## success, the ecmands riding the tide of big data, phase change memory has a lot to recommend it but to be a

 Iockeeping costs undes must work, say the authors, and being able to store multiple bits per memory cell is es researchers addressed challenges related to multi-bib PCM including drift, variability, temperature sensitivity andOniginal URL: http//mww theregister co.uk2013/11/01/hp memistor 2018/
HP 100TB Memristor drives by 2018 - if you're lucky, admits tech titan Universal memory slow in coming By Chris Mellor
restedine 1st November 201302.28 GMT
Blocks and Files HP has wamed $E$ I Reg not to get its hopes up too high after the tech titan's CTO
Mattin Fink suggested StoreSen arays could be packed with 10008 Memistor dives come 2018 .
In five years, accorring to Fink, DRAM and NAND scaling will hita wall, limiting the maximum capacity

The HP answe to this scaling wall is Memistor, it flavour of resistive RAM technology that is supposed
to have DRAM-Ike speed and better-1han-NAND storage density Fink claimed at an HP
Discover event in Las Vegas that Memistor deices will be ready by the time flash NAND hits its sinitit five years. He

-28,2015@2:46 PM 7,391 vews
Intel And Micron Jointly Announce Game-Changing 3D XPoint Memory Technology
$\qquad$

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

## Caches

- 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 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 ( $\sim 100$ Kb per core) to store synapses and neuron parameters
- Can be scaled with inter-chip communication interface
- 70 mW 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



## 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, packetswitched 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

First cxperimental dcmonstration of
a quantum algorithm a quantum algoriilhm Jonathan A. Jones, Michele Mosca A working 2-qubil NMR quantum compucr
used to solve Deulschs problem.(1988)


## 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/MC.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| Depa <br> Emai | Cognitive Computing Programming Paradigm: A Corelet Language for Composing Networks of Neurosynaptic Cores |  |
| :---: | :---: | :---: |
|  | Arnon Amir, Pallab Datta, William P. R Steve K. Esser, Alexander Andreopoul Rodrigo Alvarez-Icaza, Emmett McQuinn, Be IBM Research - Alma | Andrew S. Cassidy, Jeffrey A. Kusnitz, , Theodore M. Wong, Myron Flickner, Shaw, Norm Pass, and Dharmendra S. Modha n, San Jose, CA 95120 |
| surpr <br> langu <br> techn <br> quan | 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 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 $[\overline{8}]$ and Communications of the $A C M$ [ $]$ ]. <br> 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 cog nitive systems that is effective for the programmer and ef |

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

[^1]
## D-wave has operational Adiabatic Quantum Computer for solving optimization applications

## Adiabatic Quantum Annealing

Problem: find the ground state of

$$
H_{\text {Ising }}=\sum_{j} h_{j} \sigma_{j}^{z}+\sum_{(i, j) \in E} J_{i j} \sigma_{i}^{z} \sigma_{j}^{z}
$$

Shown by Barahona (1982) to be NP-hard in 2D, $J_{i j}= \pm, h_{j} \neq 0$.
Use adiabatic interpolation from transverse field (Farti etal. 2000)

$$
\begin{array}{ll}
H(t)=A(t) \sum_{j} \sigma_{j}^{x}+B(t) H_{\text {Ising }} \\
t \in\left[0, t_{f}\right] & \text { Program }\left\{h_{i}\right\},\left\{J_{i j}\right\}
\end{array}
$$



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


School of Enginecring University of Southern California
B. Lucas, ISI, 2015

48 We are evaluating these platforms now: http://www.csm.ornl.gov/workshops/ascrqcs2015/index.html

## Summary

卷OAK RIDGE
National Laboratory

## Disruption in Computing Stack (== research opportunities)

| Layer | Switch, 3D | NVM | Approximate | Neuro | Quantum |
| :--- | ---: | ---: | ---: | ---: | ---: | ---: |
| Application | 1 | 1 | 2 | 2 | 3 |
| Algorithm | 1 | 1 | 2 | 3 | 3 |
| Language | 1 | 2 | 2 | 3 | 3 |
| API | 1 | 2 | 2 | 3 | 3 |
| Arch | 1 | 2 | 2 | 3 | 3 |
| ISA | 1 | 2 | 2 | 3 | 3 |
| Microarch | 2 | 3 | 2 | 3 | 3 |
| FU | 2 | 3 | 2 | 3 | 3 |
| Logic | 3 | 3 | 2 | 3 | 3 |
| Device | 3 | 3 | 2 | 3 | 3 |

## 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
6. Interconnect and memory bandwidth and capacity will need to improve
7. Our community should aggressively pursue disruptive technologies
8. Some technologies could disrupt entire stack
9. Tremendous challenges in deploying these technologies with existing software
10. Many opportunities to provide new software frameworks for fundamental computer science problems: resource management, mapping, programming models, portability, etc.
11. Start talking to your colleagues in physics, chemistry, electrical engineering, etc
12. If applications suffer, so will we!

## 2016 Post-Moores Era Supercomputing Workshop @ SC16

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



## 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/computerscience/
- 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


[^0]:    htpo//www.eetasia.com/STATC/ARTICLE IMAGES/2001212EEOL 2012DEC28 STOR MFG NT 01 ip

[^1]:    Ali JavadiAbhari*, Shruti Patil ${ }^{*}$, Daniel Kudrow ${ }^{\dagger}$, Jeff Heckey ${ }^{\dagger}$, Alexey Lvov ${ }^{* *}$ Frederic T. Chong ${ }^{\dagger}$, Margaret Martonosi* ${ }^{*}$

