OAK RIDGE National Laboratory

The Coming Age of Extreme Heterogeneity

Jeffrey S. Vetter

With many contributions from FTG Group and Colleagues

ATPESC 2019
Chicago
30 Jul 2019

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

U.S. DEPARTMENT OF ENERGY

Future Technologies Group ornl

http://ft.ornl.gov vetter@computer.org

OAK RIDGE National Laboratory

Alternate title

Does ‘A New Golden Age for Computer Architecture’ equal ‘Dark Ages for Software, Algorithms, and Applications?’

Jeffrey S. Vetter

With many, many contributions from workshop participants, FTG Group, and Colleagues

ATPESC 2019
Chicago
30 Jul 2019

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

U.S. DEPARTMENT OF ENERGY

Time for a short poll…

Q: Think back 10 years. How many of you would have predicted that many of our top HPC systems would be GPU-based architectures?

a) Yes

b) No

c) Waffle ☺

Q: Think forward 10 years. How many of you predict that most of our top HPC systems will have the following architectural features?

a) X86 multicore CPU

b) GPU

c) FPGA/Reconfigurable processor

d) Neuromorphic processor

e) Deep learning processor

f) Quantum processor

g) RISC-V processor

h) Some new unknown processor

i) All/some of the above in one SoC

Q: Now imagine you are building a new application with ~3M LOC and 20 team members over the next 10 years. What on-node programming model/system do you use?

a) C, C++, Fortran

b) C++ templates, policies, etc (e.g., AMP, Kokkos, RAJA,)

c) CUDA, cu***, HIP

d) OpenCL, SYCL

e) OpenMP or OpenACC

f) R, Python, Matlab, etc

g) A Domain Specific Language (e.g., Claw, PySL)

h) A Domain Specific Framework (e.g., PetSc)

i) Some new unknown programming approach

j) All/some of the above

Motivating Trends

Contemporary devices are approaching fundamental limits

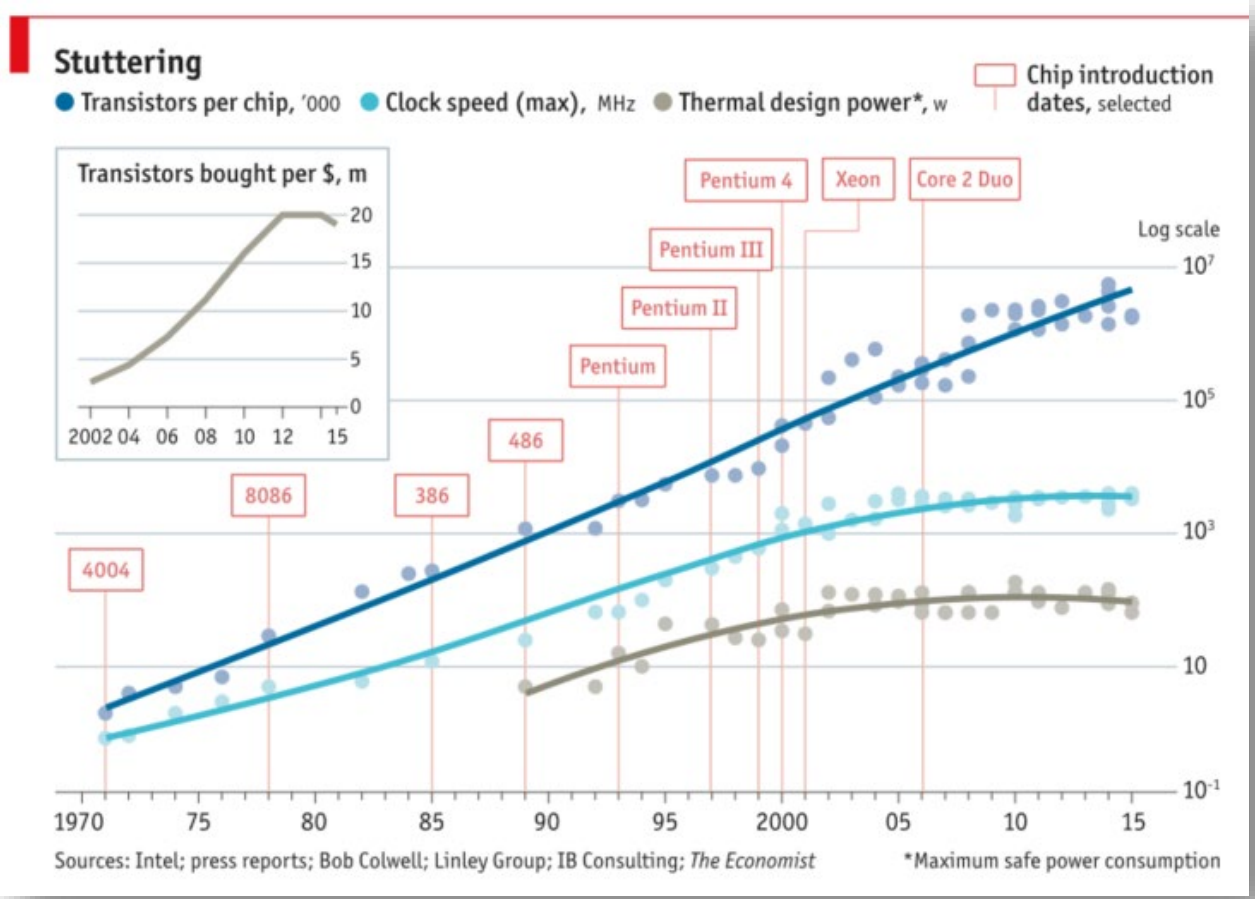

Economist, Mar 2016

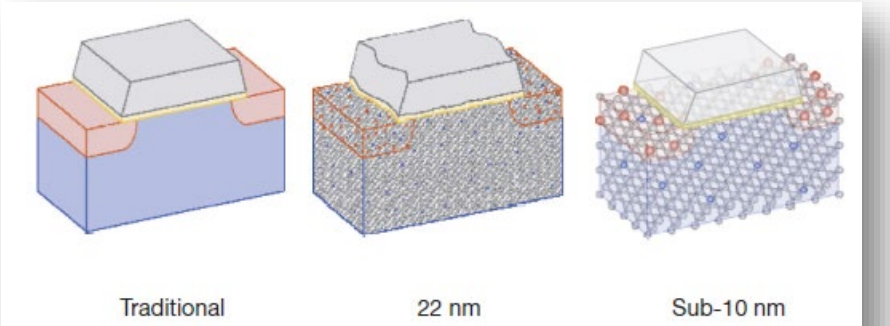

Figure 1 | As a metal oxide–semiconductor field effect transistor (MOSFET) shrinks, the gate dielectric (yellow) thickness approaches several atoms (0.5 nm at the 22-nm technology node). Atomic spacing limits the

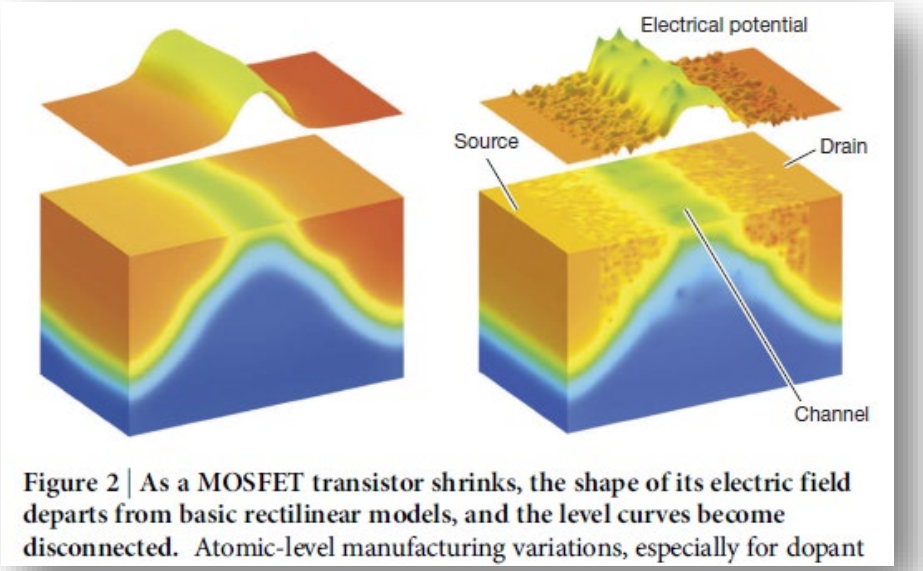

Figure 2 | As a MOSFET transistor shrinks, the shape of its electric field departs from basic rectilinear models, and the level curves become disconnected. Atomic-level manufacturing variations, especially for dopant

Dennard scaling has already ended. Dennard observed that voltage and current should be proportional to the linear dimensions of a transistor: 2x transistor count implies 40% faster and 50% more efficient.

R.H. Dennard, F.H. Gaensslen, V.L. Rideout, E. Bassous, and A.R. LeBlanc, "Design of ion-implanted MOSFET's with very small physical dimensions," *IEEE Journal of Solid-State Circuits, 9(5):256-68, 1974,*

I.L. Markov, "Limits on fundamental limits to computation," *Nature, 512(7513):147-54, 2014, doi:10.1038/nature13570.*

designlines AUTOMOTIVE

News & Analysis

Foundries' Sales Show Hard Times Continuing

Peter Clarke

5/23/2016 09:33 PM EDT

2 comments

Like 6 Tweet Share 43

SEMICONDUCTOR ENGINEERING

Home > Manufacturing, Design & Test > Uncertainty Grows For 5nm, 3nm

MANUFACTURING, DESIGN & TEST

Uncertainty Grows For 5nm, 3nm

797 74

Nanosheets and nanowire FETs under development, but costs are skyrocketing. New packaging options could

eetasia.com

GlobalFoundries Forfeit 7nm Manufacturing - EE Times Asia

Samsung to Invest $115 Billion in Foundry & Chip Businesses by 2030

Intel's 10nm Is Broken, Delayed Until 2019

37 COMMENTS

by Paul Alcorn April 26, 2018 at 6:30 PM

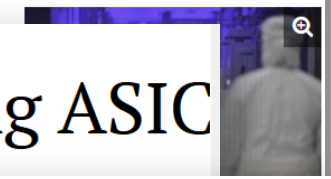

DESIGNLINES | WIRELESS AND NETWORKING DESIGNLINE

GlobalFoundries Selling ASIC Business to Marvell

By Dylan McGrath, 05.20.19 1

Share Post

Another Step Toward the End of Moore's Law

Samsung and TSMC move to 5-nanometer manufacturing

Number of Foundries with a Cutting Edge Logic Fab

SilTerra										
X-FAB										
Dongbu HiTek										
ADI	ADI									
Atmel	Atmel									
Rohm	Rohm									
Sanyo	Sanyo									
Mitsubishi	Mitsubishi									
ON	ON									
Hitachi	Hitachi									
Cypress	Cypress	Cypress								
Sony	Sony	Sony								
Infineon	Infineon	Infineon								
Sharp	Sharp	Sharp								
Freescale	Freescale	Freescale								
Renesas (NEC)	Renesas	Renesas	Renesas	Renesas						
SMIC	SMIC	SMIC	SMIC	SMIC						
Toshiba	Toshiba	Toshiba	Toshiba	Toshiba						
Fujitsu	Fujitsu	Fujitsu	Fujitsu	Fujitsu						
TI	TI	TI	TI	TI						
Panasonic	Panasonic	Panasonic	Panasonic	Panasonic	Panasonic					
STMicroelectronics	STM	STM	STM	STM	STM					
UMC	UMC	UMC	UMC	UMC	UMC					
IBM	IBM	IBM	IBM	IBM	IBM	IBM				
AMD	AMD	AMD	GlobalFoundries	GF	GF	GF	GF			
Samsung	Samsung	Samsung	Samsung	Samsung	Samsung	Samsung	Samsung	Samsung	Samsung	
TSMC	TSMC	TSMC	TSMC	TSMC	TSMC	TSMC	TSMC	TSMC	TSMC	
Intel	Intel	Intel	Intel	Intel	Intel	Intel	Intel	Intel	Intel	Future
180 nm	130 nm	90 nm	65 nm	45 nm/40 nm	32 nm/28 nm	22 nm/20 nm	16 nm/14 nm	10 nm	7 nm	5 nm

Business climate reflects this uncertainty, cost, complexity, consolidation

NVIDIA Buys Mellanox To Bring HPC Scaling To Data Centers

Kevin Krewell Contributor
Tirias Research Contributor Group
Enterprise & Cloud

The 2019 semiconductor merger and acquisition season has officially been kicked
technol
offer aft
and Xili
said in a
on the c
center p
perform

nytimes.com

Hewlett Packard Enterprise to Acquire Supercomputer Pioneer Cray

5-6 minutes

Technology|Hewlett Packard Enterprise to Acquire Supercomputer Pioneer Cray

Hewlett Packard Enterprise will pay about $1.4 billion to acquire Cray, which has designed some of the most powerful computer systems in use.CreditPaco Freire/SOPA Images, via LightRocket and Getty Images

Intel to acquire Altera for $54 a share

Monday, 1 Jun 2015 |
REUTE

Broadcom acquires Brocade in $5.9 billion deal

Posted 1 hour ago by Ron Miller (@ron_miller)

Avago Agrees to Buy Broadcom for $37 Billion

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

Tech giant ARM Holdings sold to Japanese firm for £24bn

Britain's large
deal includin

SoftBank to sell 25% of Arm to Saudi-backed fund

Son puts stake worth $8bn in UK's largest tech company into $100bn Vision Fund

Britain's bigge
ARM Holding
purchased only six months ago, placing 25 per cent of Britain's largest technology company into a new, Saudi-backed $100bn investment fund.

EXCLUSIVE

Amazon Is Becoming an AI Chip Maker, Speeding Alexa Responses

By Aaron Tilley Feb. 12, 2018 7:00 AM PST · Comments by Yonatan Raz-Fridman and Mohammad Musa

Subscribe now

Amazon.com is developing a chip designed for artificial intelligence to work on the Echo and other hardware powered by Amazon's Alexa virtual assistant, says a person familiar with Amazon's plans. The chip should allow Alexa-powered devices to respond more quickly to commands, by allowing more data processing to be handled on the device than in the cloud.

Amazon's Echo devices, powered by Alexa. Photo by Bloomberg.

The effort makes Amazon the latest major tech company, after Google and Apple, to design its own AI chips, in hopes of differentiating their products from those of rivals. That strategy has major ramifications for chip companies like Intel and Nvidia, which are now competing with companies that previously

SANDISK COMPLETES ACQUISITION OF FUSION IO

Toshiba to sell 'minority stake' in chip business to Western Digital

In April-June 2016, Toshib

The Dollar Business B

Japanese conglomerat
performing flash mem
funds, according to a s

Earlier, in a statement
memory business, ever

The laptops-to-enginee
surprised its investors
acquired in 2015, which

Western Digital Now A Storage Powerhouse With SanDisk Acquisition

DESIGNLINES | MEMORY DESIGNLINE

Q1 Chip Sales Drop Among Largest on Record

By Dylan McGrath, 05.01.19 0

Share Post

SAN FRANCISCO — Global chip sales sank by 15.5% sequentially in the first quart among the largest quarter-to-quarter declines for the industry in the last 35 years

Chip sales totaled $96.8 billion in the first quarter, down from $114.7 billion last y according to the World Semiconductor Trade Statistics (WSTS) organization, whic sales data from chipmaker member companies. On a year-over-year basis, first-qu

Sixth Wave of Computing

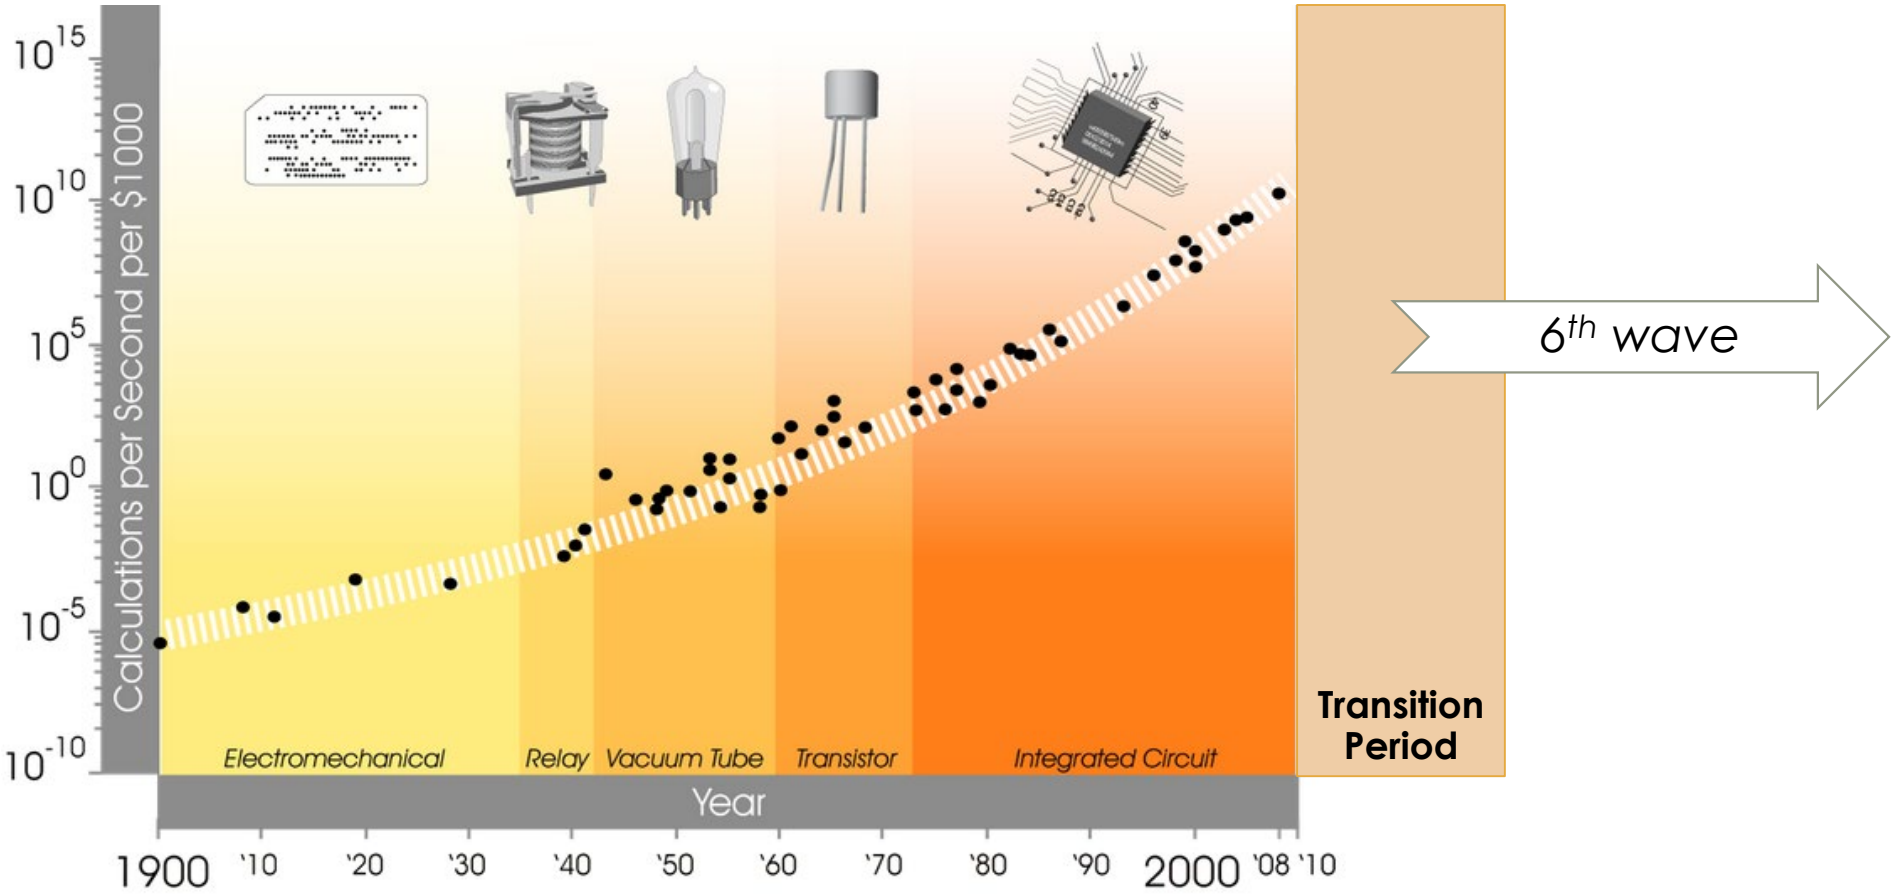

http://www.kurzweilai.net/exponential-growth-of-computing

Q: when was the field effect transistor patented?

Lilienfeld patents field effect transistor, October 8, 1926

Jessica MacNeil -October 08, 2018

6 Comments

On this day in tech history, JE Lilienfeld filed a patent for a three-electrode structure using copper-sulfide semiconductor material, known today as a field-effect transistor.

Lilienfeld's patent for a "**method and apparatus for controlling electric currents**" was granted on January 28, 1930.

According to the patent, his invention was for controlling the flow of electric current between two terminals of an electrically conducting solid by establishing a third potential between the terminals, particularly for the amplification of oscillating currents like those in radio communication.

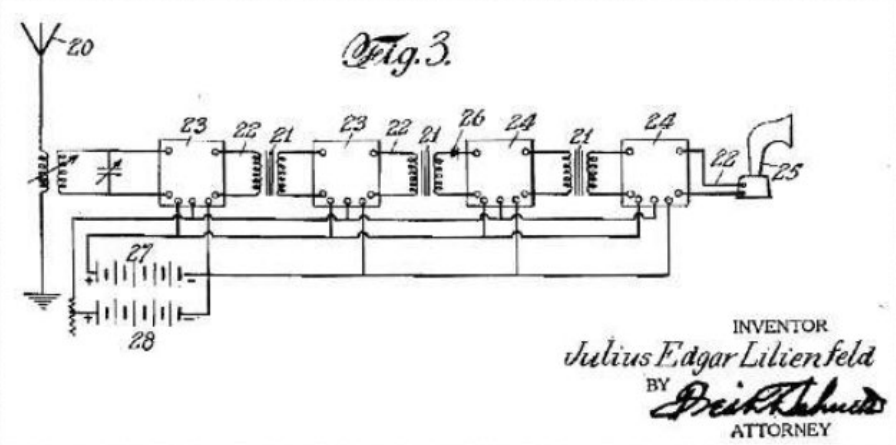

https://www.edn.com/electronics-blogs/edn-moments/4422371/Lilienfeld-patents-field-effect-transistor--October-8--1926

Predictions for Transition Period

Optimize Software and Expose New Hierarchical Parallelism

- Redesign software to boost performance on upcoming architectures
- Exploit new levels of parallelism and efficient data movement

Architectural Specialization and Integration

- Use CMOS more effectively for specific workloads
- Integrate components to boost performance and eliminate inefficiencies
- Workload specific memory+storage system design

Emerging Technologies

- Investigate new computational paradigms
  - Quantum
  - Neuromorphic
  - Advanced Digital
  - Emerging Memory Devices

Predictions for Transition Period

Optimize Software and Expose New Hierarchical Parallelism

- Redesign software to boost performance on upcoming architectures
- Exploit new levels of parallelism and efficient data movement

Architectural Specialization and Integration

- Use CMOS more effectively for specific workloads
- Integrate components to boost performance and eliminate inefficiencies
- Workload specific memory+storage system design

Emerging Technologies

- Investigate new computational paradigms
  - Quantum
  - Neuromorphic
  - Advanced Digital
  - Emerging Memory Devices

Predictions for Transition Period

Optimize Software and Expose New Hierarchical Parallelism

- Redesign software to boost performance on upcoming architectures
- Exploit new levels of parallelism and efficient data movement

Architectural Specialization and Integration

- Use CMOS more effectively for specific workloads
- Integrate components to boost performance and eliminate inefficiencies
- Workload specific memory+storage system design

Emerging Technologies

- Investigate new computational paradigms
  - Quantum
  - Neuromorphic
  - Advanced Digital
  - Emerging Memory Devices

Quantum computing: Qubit design and fabrication have made recent progress but still face challenges

Science 354, 1091 (2016) – 2 December

A bit of the action

In the race to build a quantum computer, companies are pursuing many types of quantum bits, or qubits, each with its own strengths and weaknesses.

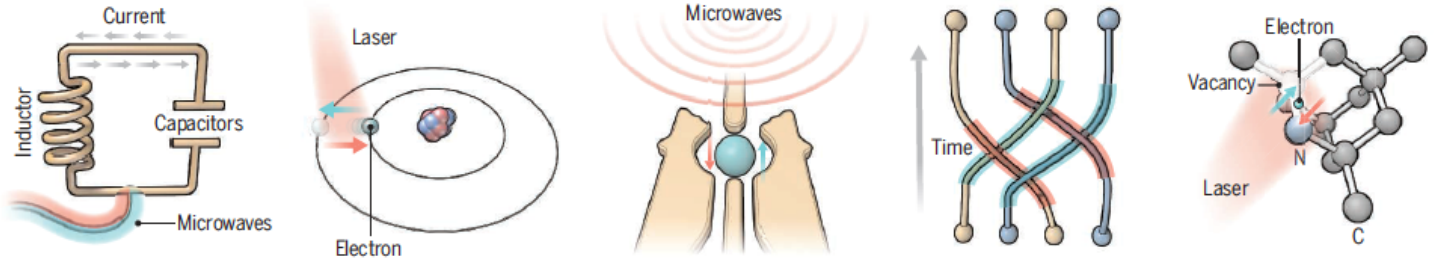

	Superconducting loops	Trapped ions	Silicon quantum dots	Topological qubits	Diamond vacancies
	A resistance-free current oscillates back and forth around a circuit loop. An injected microwave signal excites the current into super-position states.	Electrically charged atoms, or ions, have quantum energies that depend on the location of electrons. Tuned lasers cool and trap the ions, and put them in superposition states.	These "artificial atoms" are made by adding an electron to a small piece of pure silicon. Microwaves control the electron's quantum state.	Quasiparticles can be seen in the behavior of electrons channeled through semi-conductor structures. Their braided paths can encode quantum information.	A nitrogen atom and a vacancy add an electron to a diamond lattice. Its quantum spin state, along with those of nearby carbon nuclei, can be controlled with light.
Longevity (seconds)	0.00005	>1000	0.03	N/A	10
Logic success rate	99.4%	99.9%	~99%	N/A	99.2%
Number entangled	9	14	2	N/A	6
Company support	Google, IBM, Quantum Circuits	ionQ	Intel	Microsoft, Bell Labs	Quantum Diamond Technologies
⊕ **Pros**	Fast working. Build on existing semiconductor industry.	Very stable. Highest achieved gate fidelities.	Stable. Build on existing semiconductor industry.	Greatly reduce errors.	Can operate at room temperature.
⊖ **Cons**	Collapse easily and must be kept cold.	Slow operation. Many lasers are needed.	Only a few entangled. Must be kept cold.	Existence not yet confirmed.	Difficult to entangle.

Note: Longevity is the record coherence time for a single qubit superposition state, logic success rate is the highest reported gate fidelity for logic operations on two qubits, and number entangled is the maximum number of qubits entangled and capable of performing two-qubit operations.

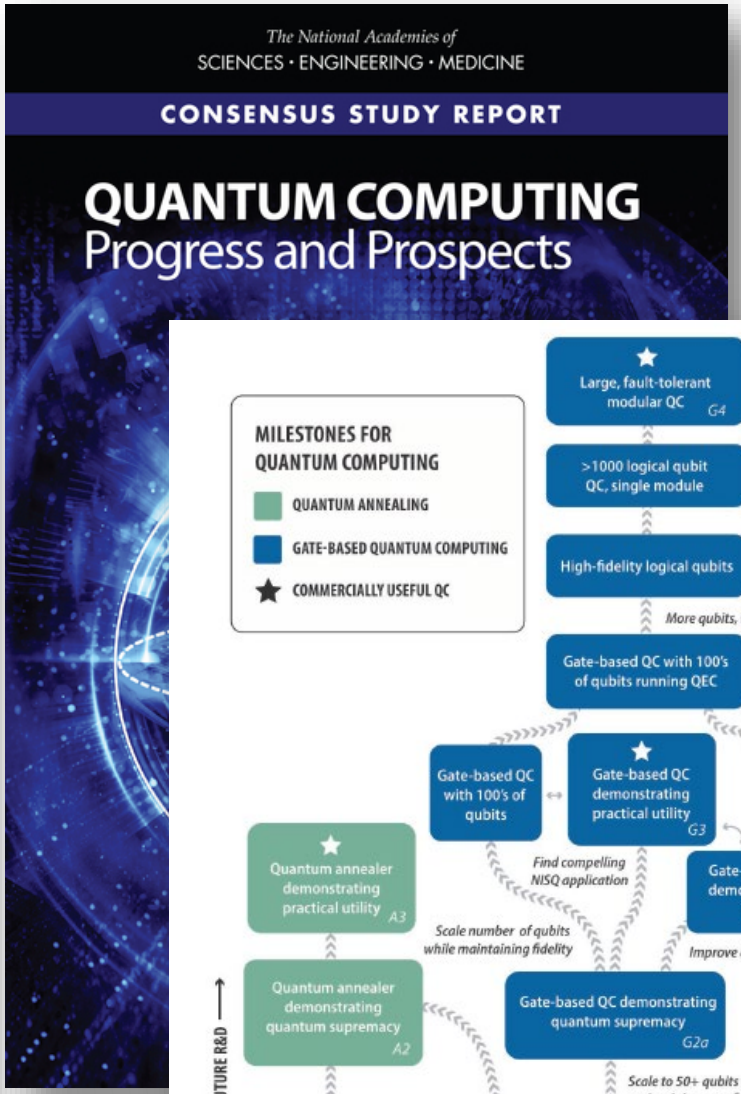

FIGURE 7.4 An illustration of potential milestones of progress in quantum computing. The arrangement of milestones corresponds to the order in which the committee thinks they are likely to be achieved; however, it is possible that some will not be achieved, or that they will not be achieved in the order indicated.

http://nap.edu/25196

Neuromorphic (Brain Inspired) Computing

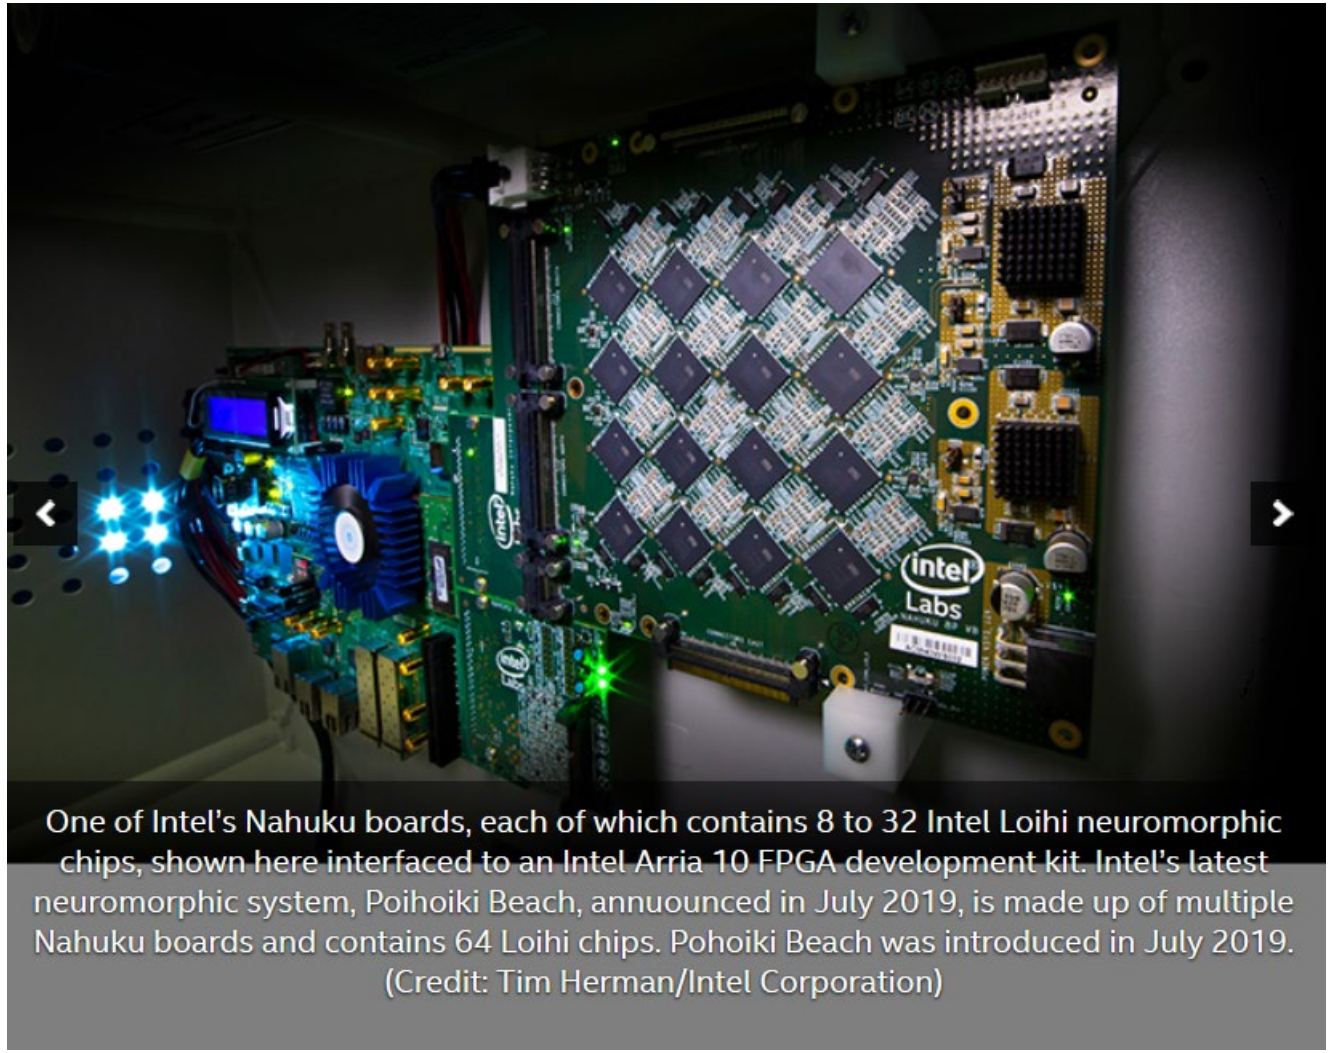

One of Intel's Nahuku boards, each of which contains 8 to 32 Intel Loihi neuromorphic chips, shown here interfaced to an Intel Arria 10 FPGA development kit. Intel's latest neuromorphic system, Poihoiki Beach, annuounced in July 2019, is made up of multiple Nahuku boards and contains 64 Loihi chips. Pohoiki Beach was introduced in July 2019. (Credit: Tim Herman/Intel Corporation)

https://newsroom.intel.com/news/intels-pohoiki-beach-64-chip-neuromorphic-system-delivers-breakthrough-results-research-tests/

- SpiNNaker
- IBM True North
- BrainScaleS
- DANNA
- Others…

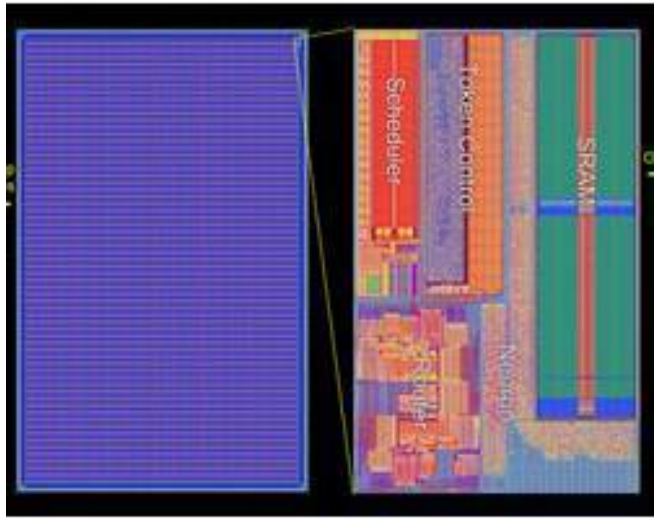

https://m-cacm.acm.org/news/201072-the-future-of-microchips

OAK RIDGE National Laboratory

Emerging Memory Devices

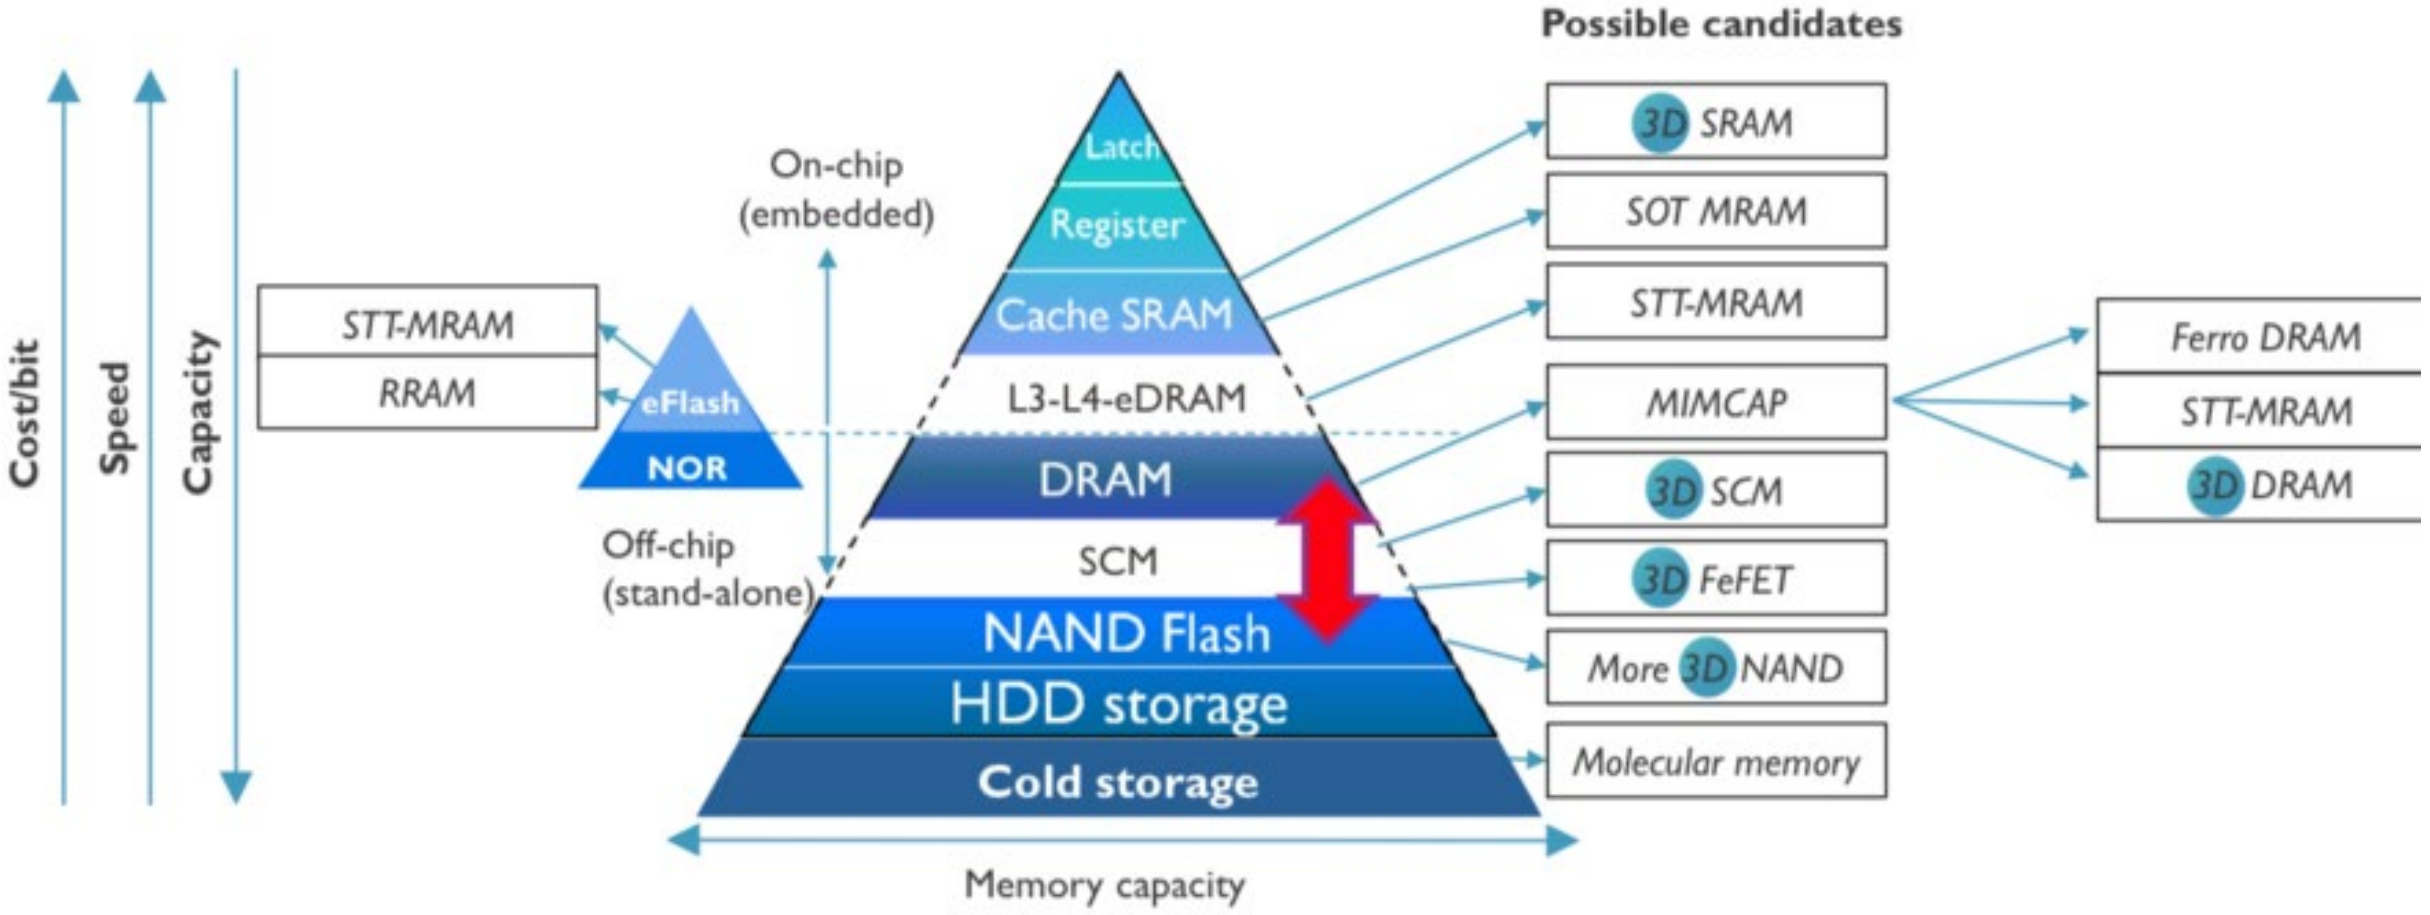

Image Source: IMEC

New devices: Carbon Nanotube Transistors and Circuits

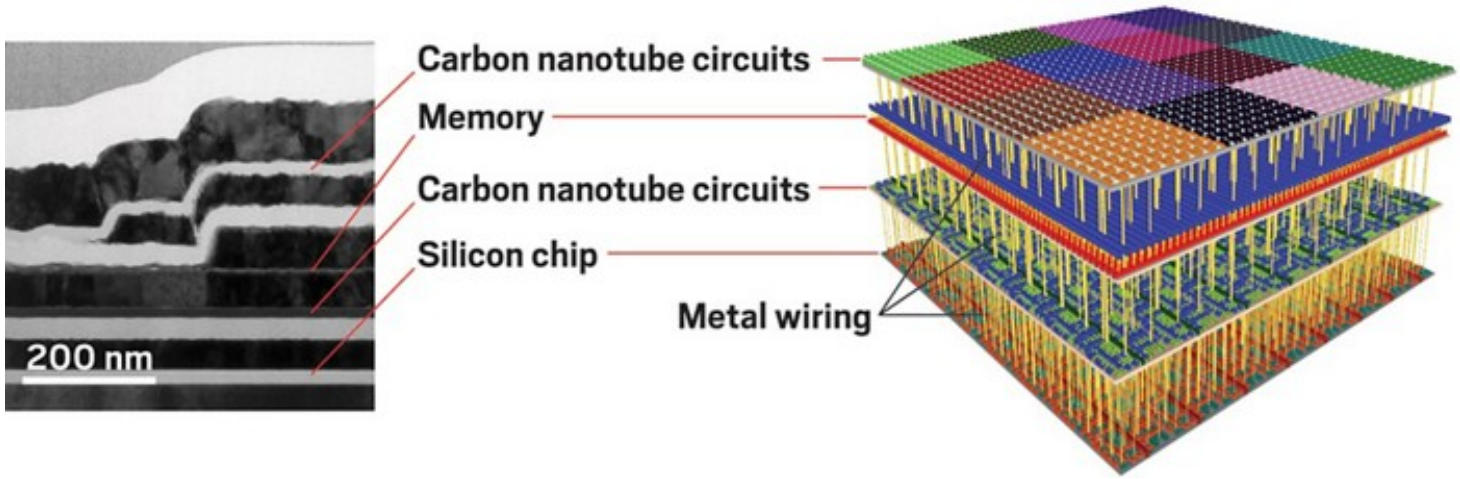

https://cen.acs.org/materials/electronic-materials/Carbon-nanotube-computers-face-makebreak/97/i8

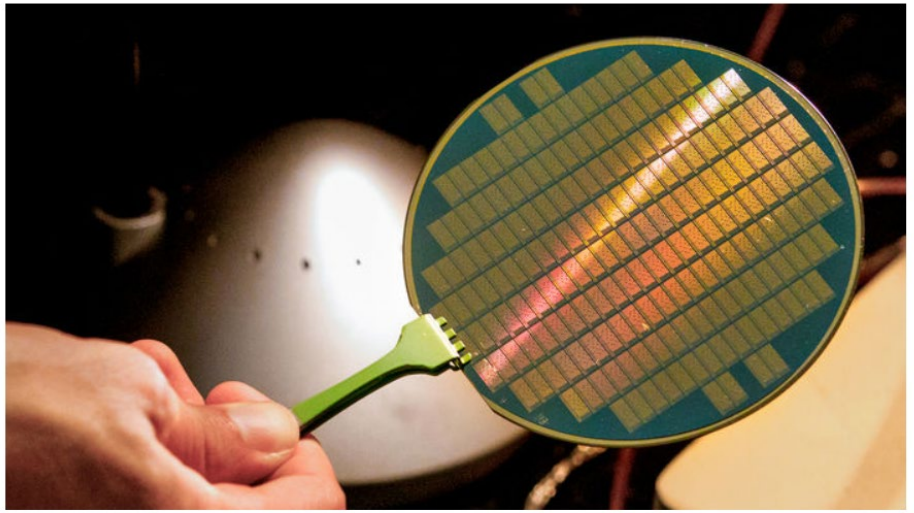

A wafer contains hundreds of tiny computer chips made from carbon nanotubes, which switch faster and more efficiently than transistors made from silicon. STANFORD ENGINEERING

Beyond silicon: $1.5 billion U.S. program aims to spur new types of computer chips

By Robert F. Service | Jul. 24, 2018 , 8:30 AM

Predictions for Transition Period

Optimize Software and Expose New Hierarchical Parallelism

- Redesign software to boost performance on upcoming architectures
- Exploit new levels of parallelism and efficient data movement

Architectural Specialization and Integration

- Use CMOS more effectively for specific workloads
- Integrate components to boost performance and eliminate inefficiencies
- Workload specific memory+storage system design

Emerging Technologies

- Investigate new computational paradigms
  - Quantum
  - Neuromorphic
  - Advanced Digital
  - Emerging Memory Devices

Pace of Architectural Specialization is Quickening

- Industry, lacking Moore's Law, will need to continue to differentiate products (to stay in business)
  - Use the same transistors differently to enhance performance
- Architectural design will become extremely important, critical
  - Dark Silicon
  - Address new parameters for benefits/curse of Moore's Law
- 50+ new companies focusing on hardware for Machine Learning

http://www.theinquirer.net/inquirer/news/2477796/intels-nervana-ai-platform-takes-aim-at-nvidias-gpu-techology

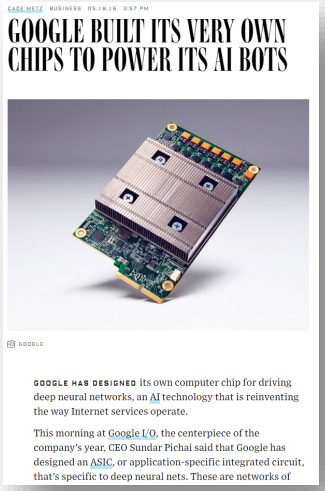

http://www.wired.com/2016/05/google-tpu-custom-chips/

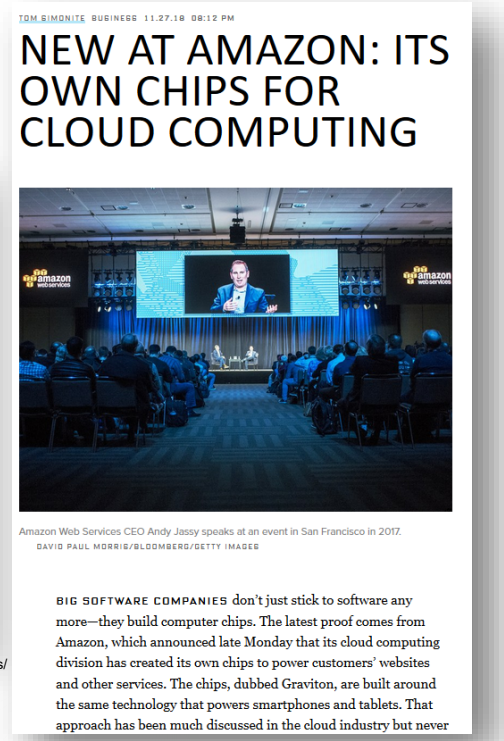

D.E. Shaw, M.M. Deneroff, R.O. Dror *et al., "Anton, a special-purpose machine for molecular dynamics simulation," Communications of the ACM, 51(7):91-7, 2008.*

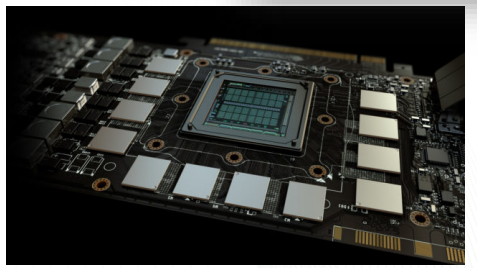

https://fossbytes.com/nvidia-volta-gddr6-2018/

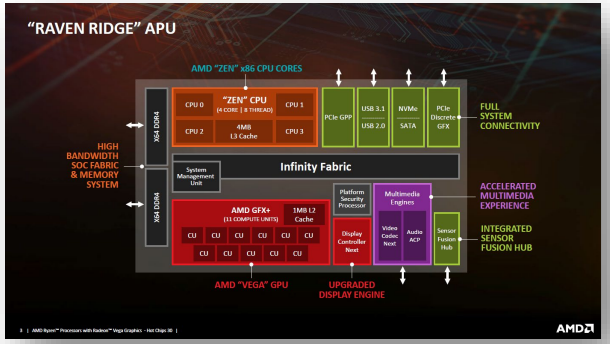

HotChips 2018

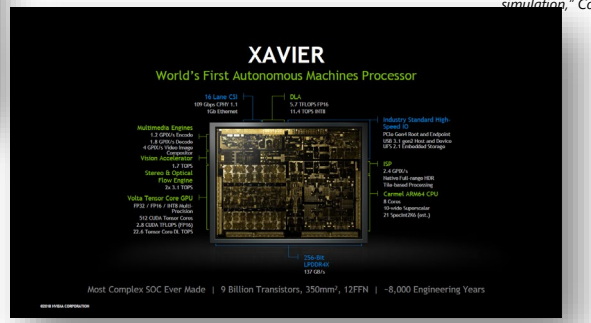

HotChips 2018

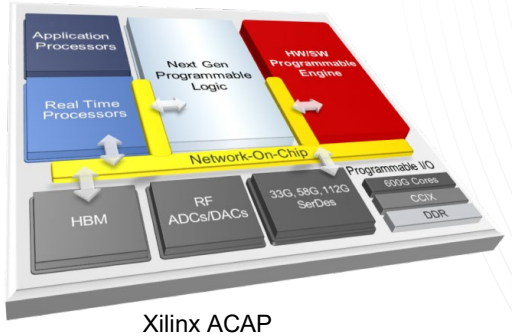

Xilinx ACAP

https://www.thebroadcastbridge.com/content/entry/1094/altera-announces-arria-10-devices-ddr4-pcie-gen3-capability

Analysis of Apple A-* SoCs

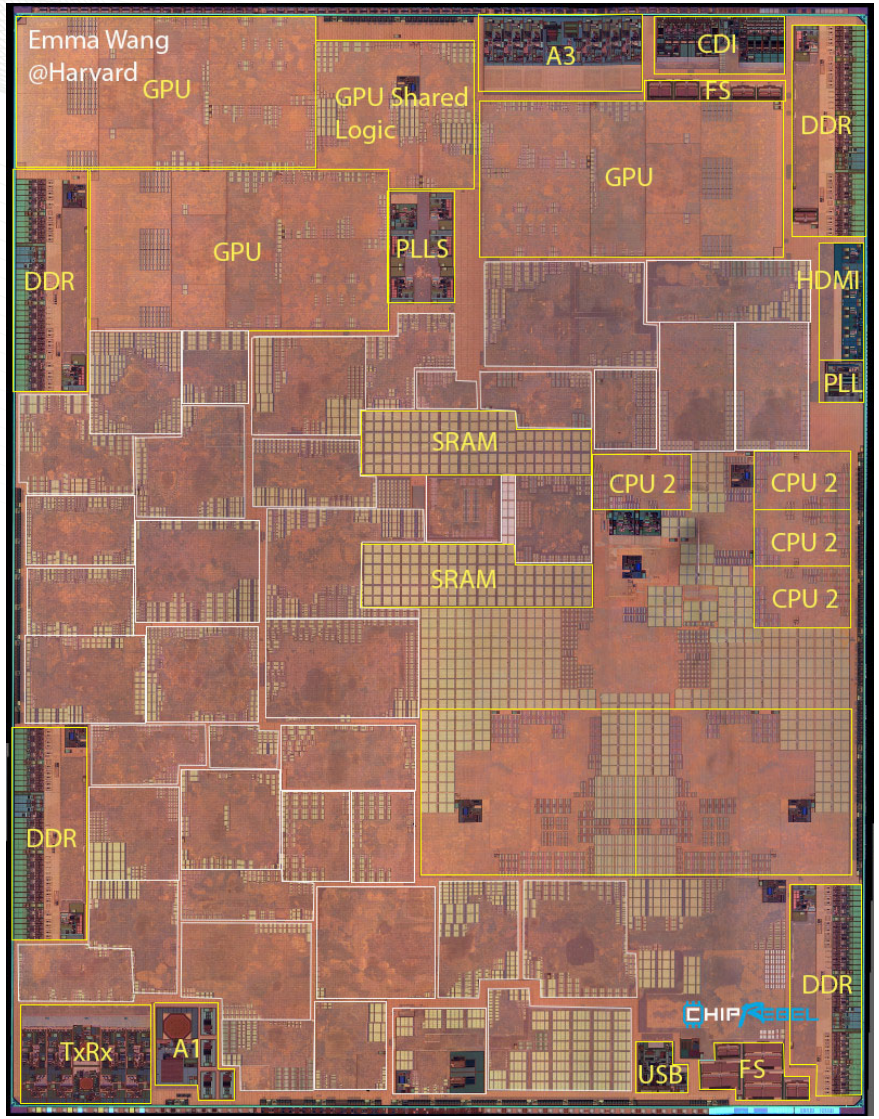

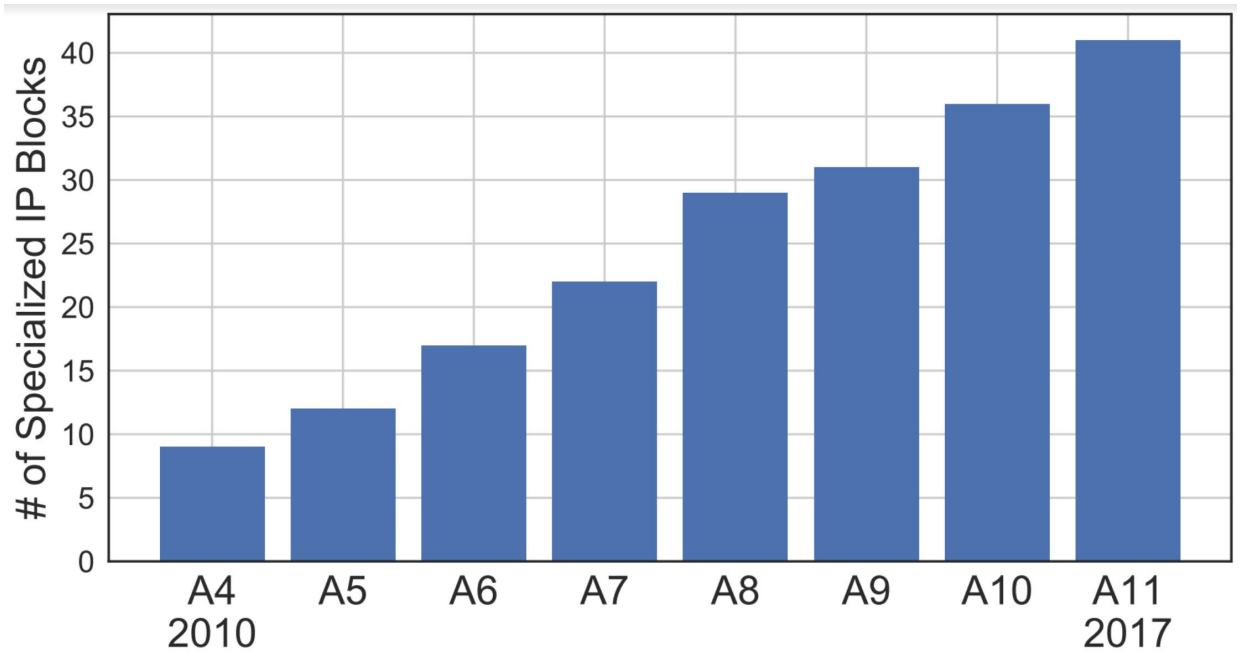

http://vlsiarch.eecs.harvard.edu/accelerators/die-photo-analysis

Growing Open Source Hardware Movement Enables Rapid Chip Design

RISC-V Ecosystem

Software

Open-source software:
Gcc, binutils, glibc, Linux, BSD, LLVM, QEMU, FreeRTOS, ZephyrOS, LiteOS, SylixOS, …

Commercial software:
Lauterbach, Segger, Micrium, ExpressLogic, …

RISC-V Foundation

ISA specification | Golden Model | Compliance

Hardware

Open-source cores:
Rocket, BOOM, RI5CY, Ariane, PicoRV32, Piccolo, SCR1, Hummingbird, …

Commercial core providers:
Andes, Bluespec, Cloudbear, Codasip, Cortus, C-Sky, Nuclei, SiFive, Syntacore, …

Inhouse cores:
Nvidia, +others

7

DARPA ERI Programs Aiming for Agile (and Frequent) Chip Creation

IDEA/POSH End State – A Universal Hardware Compiler

```
$ git clone https://github.com/darpa/idea
$ git clone https://github.com/darpa/posh
$ cd posh
$ make soc42
```

Distribution Statement "A" (Approved for Public Release, Distribution Unlimited) 23

A. Olofsson, 2018

Transition Period will be Disruptive – Opportunities and Pitfalls Abound

- New devices and architectures may not be hidden in traditional levels of abstraction

- Examples
  - A new type of CNT transistor may be completely hidden from higher levels
  - A new paradigm like quantum may require new architectures, programming models, and algorithmic approaches

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

Adapted from IEEE Rebooting Computing Chart

Department of Energy (DOE) Roadmap to Exascale Systems

An impressive, productive lineup of *accelerated node* systems supporting DOE's mission

Pre-Exascale Systems [Aggregate Linpack (Rmax) = 323 PF!]

First U.S. Exascale Systems

2012	2016	2018	2020	2021-2023

ORNL
Cray/AMD/NVIDIA

ORNL
IBM/NVIDIA

Heterogeneous Cores

ORNL
AMD/Cray

ANL
IBM BG/Q

ANL
Cray/Intel KNL

Deep Memory incl NVM

ANL
Intel/Cray

LBNL
Cray/Intel Xeon/KNL

LBNL
Cray/AMD/NVIDIA

Plateauing I/O Performance

LLNL
IBM BG/Q

LANL/SNL
Cray/Intel Xeon/KNL

LLNL
IBM/NVIDIA

LANL/SNL
TBD

LLNL
TBD

OAK RIDGE National Laboratory

Jan 2018

Frontier | 7 May 2019

U.S. Department of Energy and Cray to Deliver Record-Setting Frontier Supercomputer at ORNL

Exascale system expected to be world's most powerful computer for science and innovation

Topic: Supercomputing

May 7, 2019

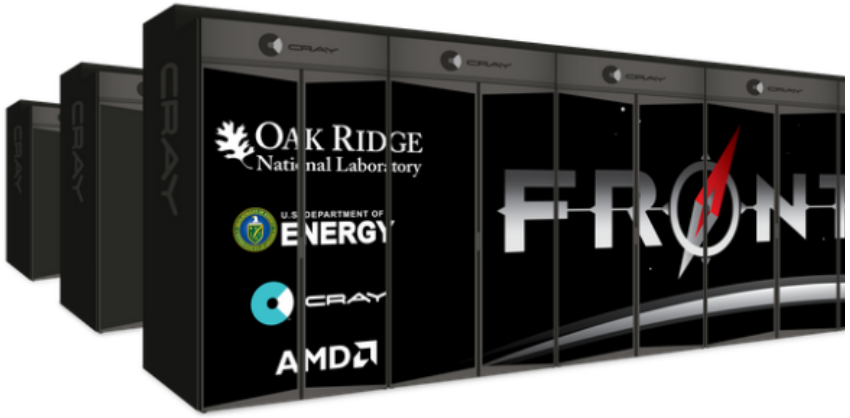

OAK RIDGE, Tenn., May 7, 2019—The U.S. Department of Energy today anr contract with Cray Inc. to build the Frontier supercomputer at Oak Ridge Nati which is anticipated to debut in 2021 as the world's most powerful computer v performance of greater than 1.5 exaflops.

Scheduled for delivery in 2021, Frontier will accelerate innovation in science and technology and maintain U.S. leadership in high-performance computing and artificial intelligence. The total contract award is valued at more than $600 million for the system and technology development. The system will be based on Cray's new Shasta architecture and Slingshot interconnect and will feature high-performance AMD EPYC CPU and AMD Radeon Instinct GPU technology.

Peak Performance	>1.5 EF
Footprint	> 100 cabinets
Node	1 HPC and AI Optimized AMD EPYC CPU 4 Purpose Built AMD Radeon Instinct GPU
CPU-GPU Interconnect	AMD Infinity Fabric Coherent memory across the node
System Interconnect	Multiple Slingshot NICs providing 100 GB/s network bandwidth Slingshot dragonfly network which provides adaptive routing, congestion management and quality of service.
Storage	2-4x performance and capacity of Summit's I/O subsystem. Frontier will have near node storage like Summit.

OAK RIDGE National Laboratory

ASCR Extreme Heterogeneity Workshop

January 23-25, 2018 Virtual Meeting

https://orau.gov/exheterogeneity2018/
https://doi.org/10.2172/1473756

- **Goal: Identify Priority Research Directions for Computer Science needed to make future supercomputers usable, useful and secure for science applications in the 2025-2040 timeframe**
  - Note that quantum computing was defined as out of scope by ASCR.
- Primary focus on the software stack and programming models/environments/tools
- 150+ participants: DOE labs, academia, and industry
- White papers solicited (106 received!) to contribute to the FSD, identify potential participants, and help refine the agenda
- First ASCR workshop to use Basic Research Needs format (BES inspired)
  - Summit, Summit report, Factual Status Document, whitepapers, BRN/PRD result
- Organizing Committee
  - Jeffrey Vetter (ORNL), Lead Organizer and Program Committee Chair
  - Ron Brightwell (Sandia-NM), Pat McCormick (LANL), Rob Ross (ANL), John Shalf (LBNL)
  - Lucy Nowell, ASCR Program Manager
- Program Committee Members
  - Katy Antypas (LBNL, NERSC), David Donofrio (LBNL), Maya Gokhale (LLNL), Travis Humble (ORNL), Catherine Schuman (ORNL), Brian Van Essen (LLNL), Shinjae Yoo (BNL)

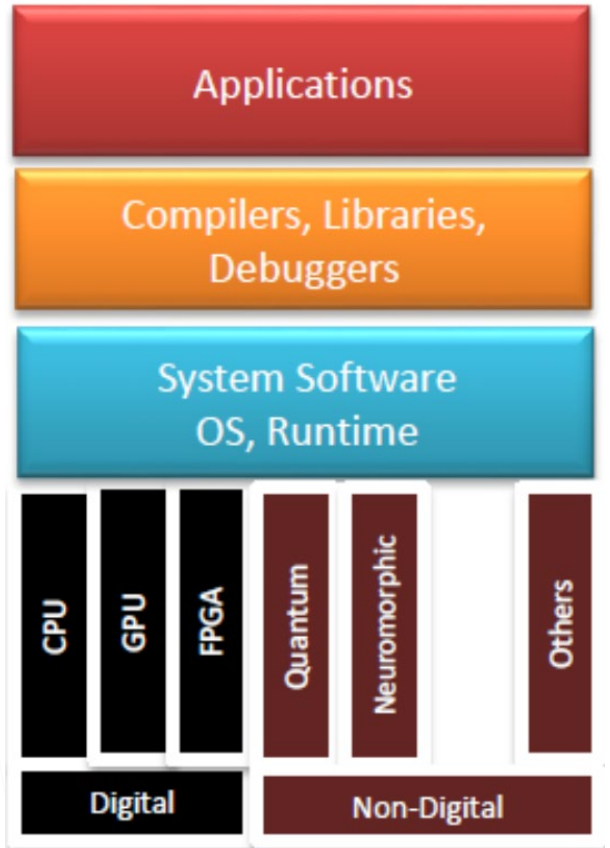

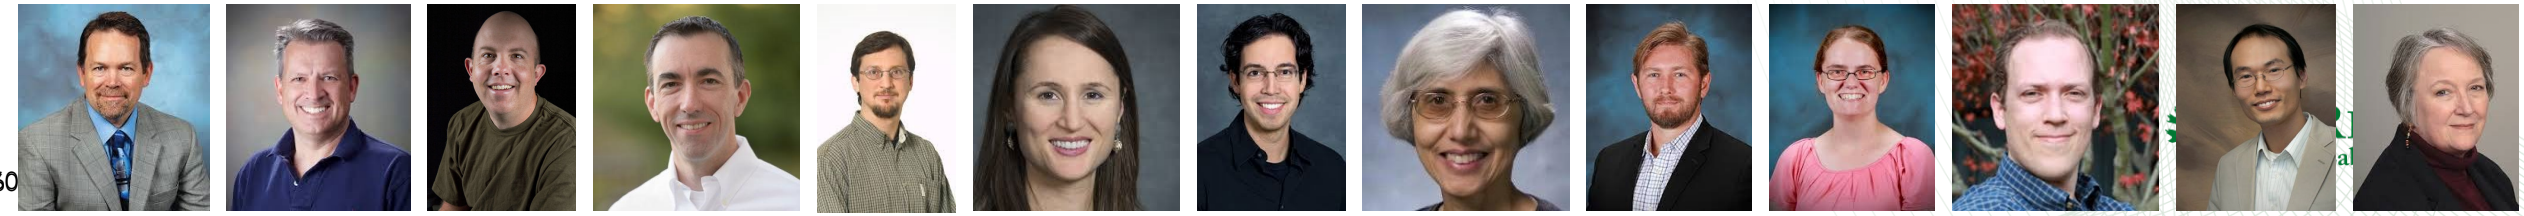

EH Priority Research Directions (PRDs)

Maintaining and improving programmer productivity	• Flexible, expressive, programming models and languages • Intelligent, domain-aware compilers and tools • Composition of disparate software components
Managing resources intelligently	• Automated methods using introspection and machine learning • Optimize for performance, energy efficiency, and availability
Modeling & predicting performance	• Evaluate impact of potential system designs and application mappings • Model-automated optimization of applications
Enabling reproducible science despite non-determinism & asynchrony	• Methods for validation on non-deterministic architectures • Detection and mitigation of pervasive faults and errors
Facilitating Data Management, Analytics, and Workflows	• Mapping of science workflows to heterogeneous hardware and software services • Adapting workflows and services to meet facility-level objectives through learning approaches

Future Technologies Group (FTG)

Jeffrey S. Vetter, Group Leader

The Future Technologies Group performs research in core technologies for emerging generations of high-end computing architectures, including prototype computer architectures and experimental software systems. We investigate these technologies with the goal of improving the performance, energy efficiency, reliability, and productivity of these architectures for our sponsors and applications teams. See http://ft.ornl.gov.

https://www.thebroadcastbridge.com/content/entry/1094/altera-announces-arria-10-2666mbps-ddr4-memory-fpg

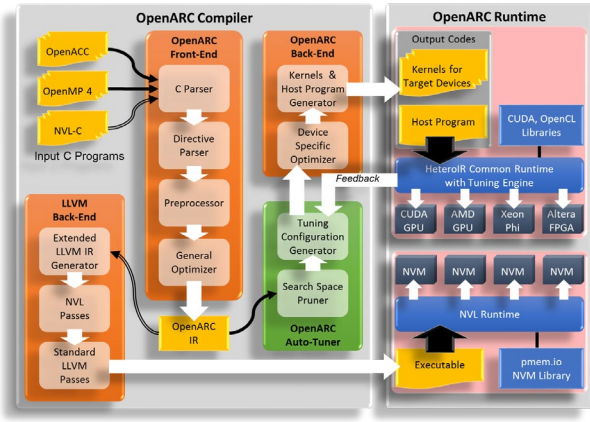

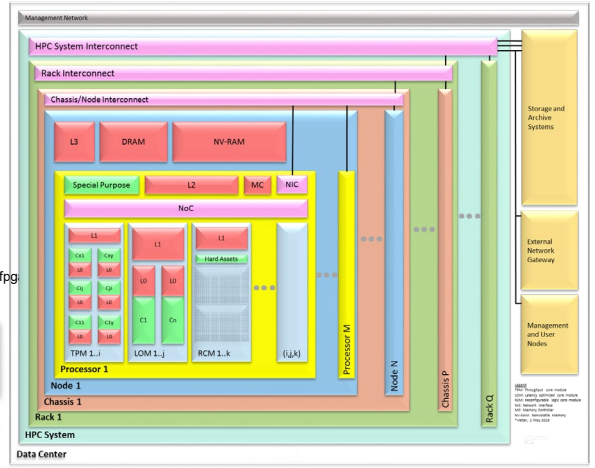

Key Technical Areas

- Heterogeneous architectures
- Deep memory hierarchies including non-volatile memory
- Performance measurement, analysis, simulation, and modeling of emerging architectures.
- Programming systems to address emerging architectures
- Beyond Moore's Computing

Software Artifacts

- Scalable Heterogeneous Computing Benchmarks (SHOC)
- mpiP
- DESTINY
- Aspen
- OpenARC
- Papyrus
- NVL-C
- Oxbow
- LLVM Clacc and Parallel IR
- DRAGON
- RISC-V Extensions

Sponsors

- DOE ASCR, BER
- DOE Exascale Computing Project
- DOE SciDAC
- DARPA
- ORNL LDRD
- National Science Foundation
- Department of Defense
- NIH

Impact

- Publications in SC, ICS, HPDC, TPDS, DATE, PLDI, IPDPS, Trans VLSI, etc.
- Two Gordon Bell awards
- NSF Keeneland
- DOE Titan
- IEEE TCHPC Early Career
- IEEE Fellows
- ~100 interns
- ~130 FTG seminars

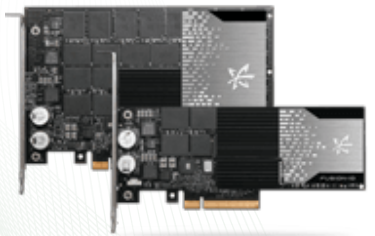

FusionIO

Progression of Experimental Computing Technologies

TRL 1-3 Basic Concepts

- *Examples: carbon-nanotube computing, memristor-based neuromorphic computing, chip-level silicon photonics, universal quantum computing*

TRL 4-6 Emerging

- *Examples: FPGAs in HPC, TrueNorth, SpiNNaker, D-Wave, Emu, many SoC-based systems, TPU, Gen-Z NoCs, near-memory computing*

TRL 7-9 Operational

- *Examples: Titan, Cori, Mira, Summit, BlueWaters, Keeneland, Stampede, Tsubame2.5*

Evaluate, Select, and Improve Emerging Computing Technologies

Production

Experimental Prototype

Limited Access Testbed

"Bench" System

CS & Math Research

Programming	Assembly language, or less	Few, if any, development tools	Language support and compilers.
OS-R	Manual	Specialized programming environments and OSs	Commodity OS & runtime systems
Scale	Small collections of devices	Single to hundreds of engineered processing elements	>10,000 processing elements
Performance	Analytical projections based on device empirical evaluation.	Analytical projections or simulation based on component or pilot system empirical evaluation.	Empirical evaluation of prototype and final systems.
Apps	Small encoded kernels	Architecture-aware algorithms; Mini-apps; Small applications	Numerical libraries; Full scale applications
Example	GPUs invented in 1999	OpenGL in 2001; CUDA in 2007; OpenCL in 2008; OpenACC in 2010; DP in 2010; ECC in 2012	GPUs are a fully supported compute technology in the HPC ecosystem.

ORNL ExCL Model

https://excl.ornl.gov

- Provide low-level access to emerging computer architectures to encourage experimentation and prototyping of new hardware and software solutions.
- Not just testbeds, but staff and software environments to support this mode of operation.

ExCL Common Infrastructure

Project and User management
- Accounts
- Projects and Proposals
- Help

Community
- Workshops
- Online discussions forums and issues
- Consolidated
- News

Shared Login and Gateway Nodes
- Gateway nodes
- Data transfer nodes
- Consistent and secure access to private network compartments

Authentication and Authorization
- Secure operations
- Partition access to specific compartments
- System and account lifecycles
- Experience with management of export controlled and proprietary systems

Shared Filesystems and Databases
- Secure access to filesystems across pillars

Monitoring and control systems
- Manage access to shared resources
- Manage privileged access levels
- Lights out operation

Source Code and Data sets
- Source Code repos
- Performance databases for applications and architectures

Web
- Educational and reference materials
- Outreach
- Both Open and Controlled access

ExCL Technology Pillars

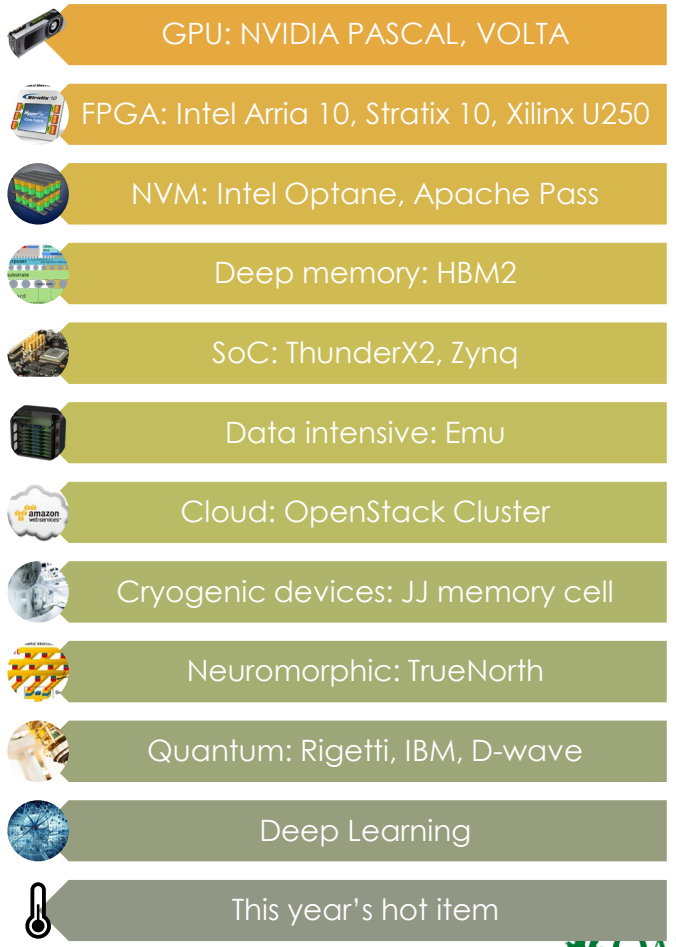

Per pillar expert collaboration

Pillars refreshed annually

OAK RIDGE National Laboratory

Learn more about ExCL or Apply for Access

https://excl.ornl.gov

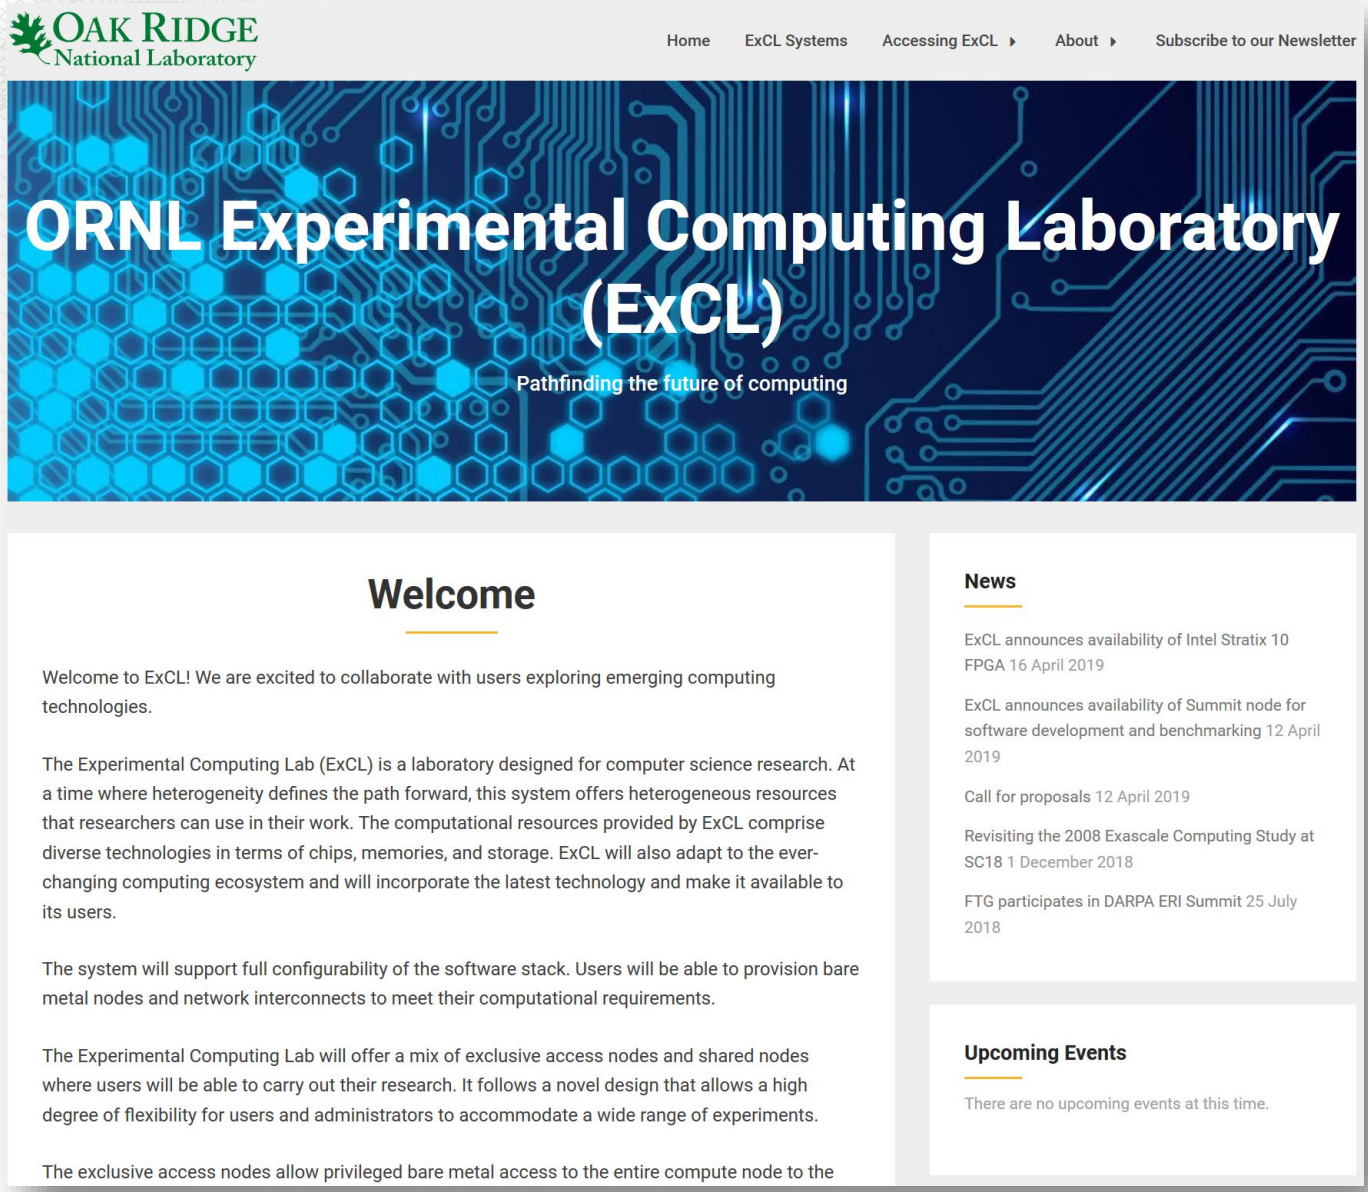

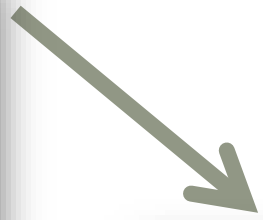

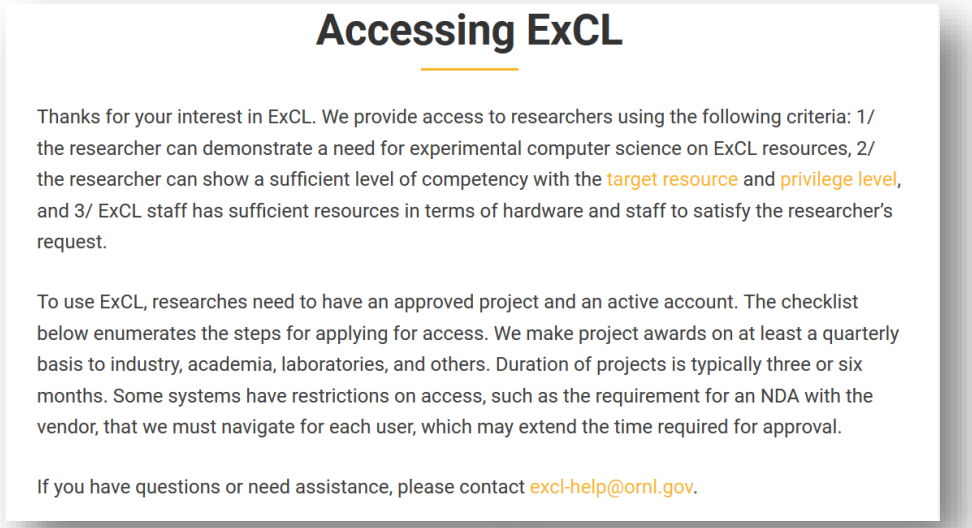

https://excl.ornl.gov/accessing-excl/

91

Take Away Messages

1. Moore's Law as we know it is definitely ending for either economic or technical reasons in the next by 2025
2. CMOS continues indefinitely – incredible technology!
    1. Specialization – use the same transistors differently
    2. Architecting effective solutions will be critical for industry, HPC
3. 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
4. Our community must aggressively explore emerging technologies now!
    1. Some technologies will disrupt entire stack
5. Tremendous opportunities and challenges in designing and deploying these new technologies with massive existing software
    1. Many opportunities to provide new software frameworks for fundamental computer science problems: resource management, mapping, programming models, portability, algorithms, etc.
6. Start exploring these new technologies!
7. Talk to your colleagues in physics, chemistry, electrical engineering, math, etc

Recap

- Recent trends in computing paint an ambiguous future (for HPC and broader community)
  - Contemporary systems provide evidence that power constraints are driving architectures to change rapidly (e.g., Dennard, Moore)
  - Multiple architectural dimensions are being (dramatically) redesigned: Processors, node design, memory systems, I/O
- Major transition point for computing
  - New devices
  - New architectures
  - New programming systems
- Complexity and uncertainty are ubiquitous
- Programming systems must provide performance portability (in addition to functional portability)!!
- In near term, rate of change will accelerate and grow more diverse

- Visit us
  - We host interns and other visitors year round
    - Faculty, grad, undergrad, high school, industry
- Jobs in FTG
  - Postdoctoral Research Associate in Computer Science
  - Software Engineer
  - Computer Scientist
  - Visit https://jobs.ornl.gov
- Contact me vetter@ornl.gov

Acknowledgements

- Contributors and Sponsors
  - Future Technologies Group: http://ft.ornl.gov
  - US Department of Energy Office of Science
    - Exascale Computing Project
    - DOE Vancouver Project: https://ft.ornl.gov/trac/vancouver
    - DOE Blackcomb Project: https://ft.ornl.gov/trac/blackcomb
    - SciDAC RAPIDS Project
  - US DARPA

Bonus Material

Memory Systems Started Diversifying Several Years Ago

- Architectures
  - HMC, HBM/2/3, LPDDR4, GDDR5X, etc
  - 2.5D, 3D Stacking
- Configurations
  - Unified memory
  - Scratchpads
  - Write through, write back, etc
  - Consistency and coherence protocols
  - Virtual v. Physical, paging strategies
- New devices
  - ReRAM, PCRAM, STT-MRAM, 3D-Xpoint
- Integrating compute and memory
  - PIM, CIM, In-mem

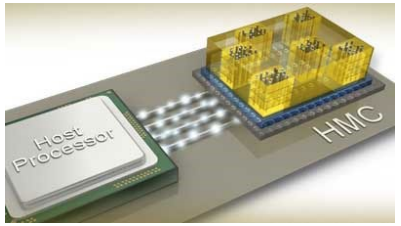

https://www.micron.com/~/media/track-2-images/content-image

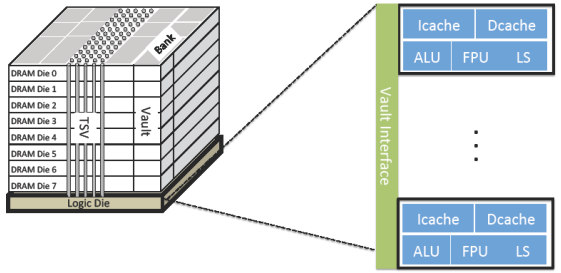

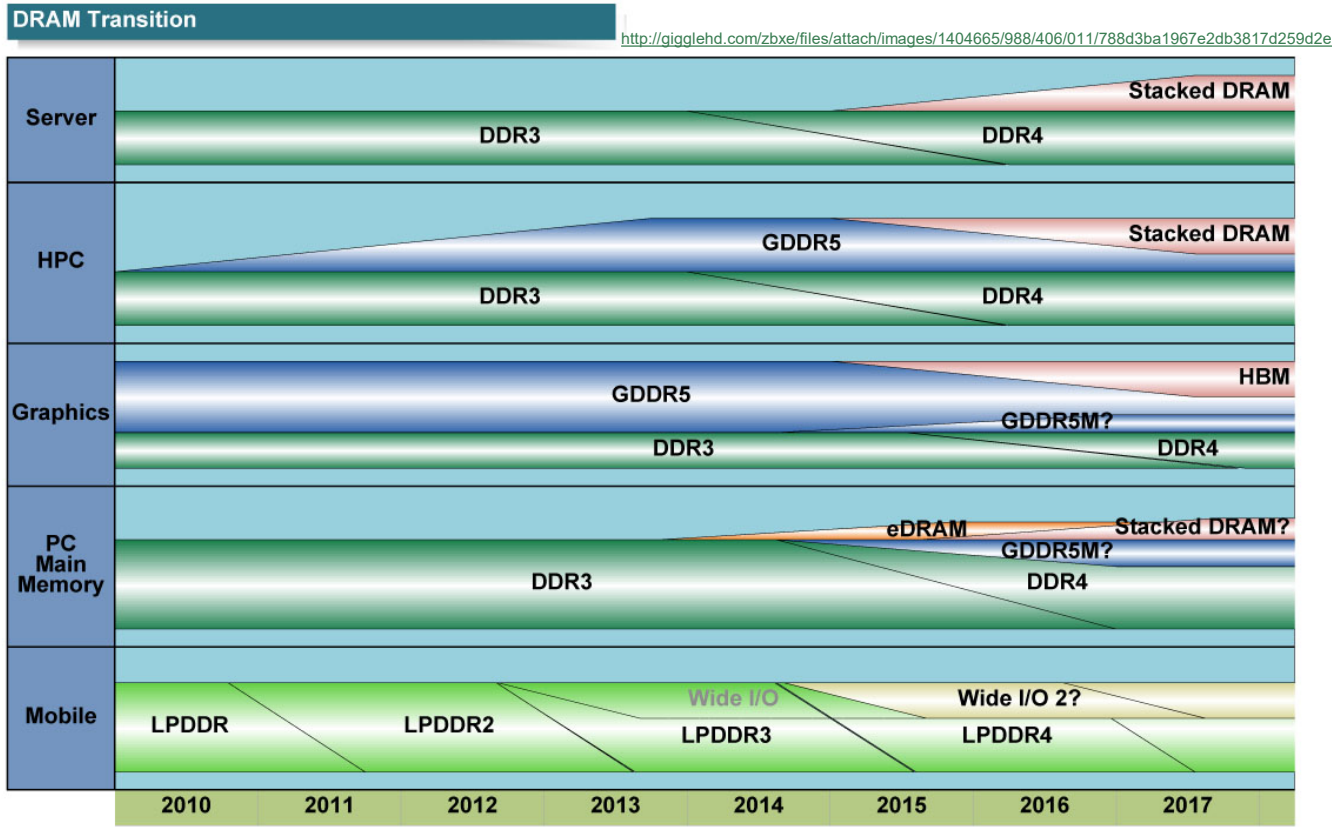

http://gigglehd.com/zbxe/files/attach/images/1404665/988/406/011/788d3ba1967e2db3817d259d2e

	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^2)	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	10^3	10^5	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	10^8-10^9	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									

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

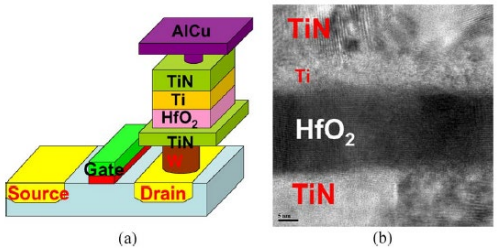

Fig. 4. *(a) A typical 1T1R structure of RRAM with HfO_x; (b) HR-TEM image of the TiN/Ti/HfO_x/TiN stacked layer; the thickness of the HfO_2 is 20 nm.*

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

NVRAM Technology Continues to Improve – Driven by Broad Market Forces

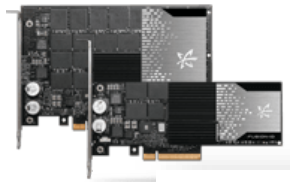

designlines MEMORY

Blog

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

Kevin Gibb, Product Line Mar
TechInsights
4/6/2016 04:40 PM EDT
9 comments post a comment

Like 16 Tweet Shar

The highly anticipated Samsu memory is out in the market, first look.

Samsung had announced its 25 K9AFGY8S0M 3D V-NAND as would be used in a variety of so on the market in early 2016. True to their word, we managed to find them in their 2 TB capacity, mSATA, T3 port Figure 1.

designlines WIRELESS & NETWOR

Slideshow

Facebook Likes Intel's 3D XPoint

Google joins open hardware effort

Rick Merritt

NO RATINGS
LOGIN TO RATE

Facebook's support for the 3D XPoint non-volatile me co-developed by Intel and Micron is "a huge endorse

May 18, 2016

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

Tiffany Trader

IBM scientists have broken new ground in the c change memory technology (PCM) that puts a XPoint technology from Intel and Micron. IBM s

Original URL: 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 NANI of the technologies: process shrinks will come to off a cliff as a side effect of reducing the size of e

The HP answer to this scaling wall is Memristor, to have DRAM-like speed and better-than-NAND in Las Vegas that Memristor devices will be ready also showed off a Memristor wafer, adding that it

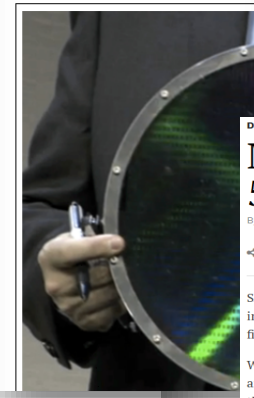

designlines MEMORY

News & Analysis

Samsung Debuts 3D XPoint Killer

3D NAND variant stakes out high-end SSDs

Rick Merritt
8/11/2016 00:01 AM EDT
5 comments

NO RATINGS
1 saves
LOGIN TO RATE

Like 56 Tweet Share 212 G+1 4

SANTA CLARA, Calif. – Samsung lobbed a new variant of its 3D

tom's HARDWARE PRODUCT REVIEWS NEWS DEALS FORUM

Samsung's 10-Year Plan Starts With 128TB QLC SSD, 960 Successor

22 COMMENTS

by Chris Ramseyer August 8, 2017 at 12:30 PM

designlines MEMORY

News & Analysis

3D NAND Flash at 2 Cents per GB

BeSang wants to lower barrier to 3D NAND flash

R. Colin Johnson
7/18/2016 07:10 PM EDT
14 comments

LOGIN TO RATE

chip technology
ms to have since
cture for NAND
mentation of its
door to
s offering to
ice that reduces
ut 2¢ per

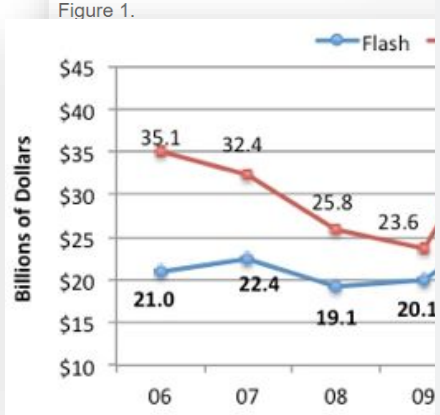

http://www.eetasia.com/STATIC/ARTICLE_IMAGES/201212

ars TECHNICA BIZ & IT TECH SCIENCE POLICY CARS GAMING & CULTURE

THE REVOLUTION IS HERE —

Intel at last announces Optane memory: DDR4 that never forgets

New memory offers huge capacities and persistence, but fits in a DDR4 slot.

PETER BRIGHT - 5/30/2018, 8:45 PM

SanDisk 1TB
Extreme UHS-I ...
$449.99
B&H Photo

46 PM 7,391 VIEWS

nd Micron Jointly Announce
-Changing 3D XPoint Memory
ology

DESIGNLINES | MEMORY DESIGNLINE

Memory Forecast to Account fo 53% of Semiconductor Capex

By Dylan McGrath, 08.29.18

Share Post Share on Facebook Share on Twitter G+ in

SAN FRANCISCO — Capital spending for memory chips is expected to account for 53% industry capex of $102 billion this year, nearly twice the percentage that memory accou five years ago, according to market research firm IC Insights.

With all NAND flash vendors ramping up 3D NAND capacity, NAND-related capital expe are forecast to total more than $31 billion, 31% of the semiconductor industry total, acc the latest edition of IC Insights' McClean Report. The total for NAND capex would repre increase of 13% over 2017, when NAND flash capex grew by 91%.

Meanwhile, the report forecasts that capital spending for DRAM and SRAM will increase any other industry segment, growing 41% in 2018 after an 82% increase last year. DRAM capex is expected to total $22.9 billion, 22% of the industry-wide total, according to the

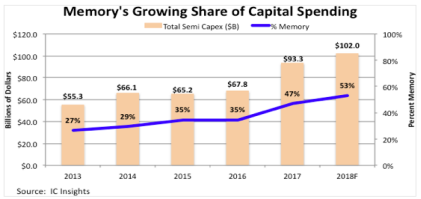

The forecasted total of $102 billion for the overall semiconductor industry — including upgrades to existing wafer fab lines and brand new manufacturing facilities — would ma