ASCR Overview and Perspective on Semiconductor Technologies

Bill Harrod
DOE / ASCR
March 24, 2016

U.S. DEPARTMENT OF ENERGY | Office of Science

Advanced Scientific Computing Research (ASCR) at a Glance

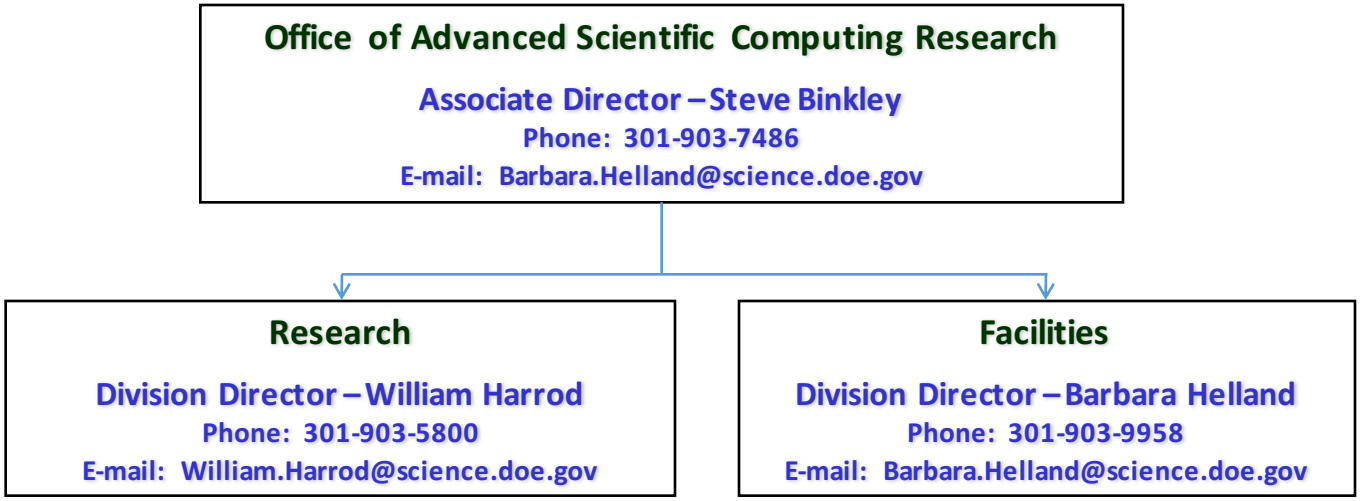

Relevant Websites

ASCR: science.energy.gov/ascr/

ASCR Workshops and Conferences: science.energy.gov/ascr/news-and-resources/workshops-and-conferences/

SciDAC: www.scidac.gov

INCITE: science.energy.gov/ascr/facilities/incite/

U.S. DEPARTMENT OF ENERGY | Office of Science

ASCR Research Division

- Applied Mathematics
  - Emphasizes scalable numerical methods for complex systems, uncertainty quantification, large-scale data analysis and exascale algorithms;
- Computer Science
  - Exascale computing (architecture, parallelism, power aware, fault tolerance), operating systems, compilers, performance tools, productivity, scientific data management, analysis and visualization for petabyte to exabyte data sets;
- Partnerships
  - Co-Design and partnerships to pioneer the future of scientific applications;
- Next Generation Networks for Science
  - Tools for the future of distributed science
- Research and Evaluation Prototypes
  - Fast Forward and Design Forward partnerships with Industry and Non-Recurring Engineering for the planned facility upgrades

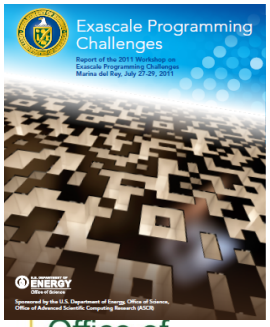

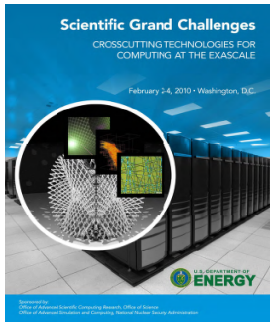

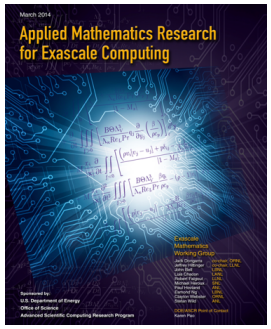

U.S. DEPARTMENT OF ENERGY | Office of Science

NATIONAL STRATEGIC COMPUTING INITIATIVE July 29, 2015

EXECUTIVE ORDER

- - - - - - -

CREATING A NATIONAL STRATEGIC COMPUTING INITIATIVE

By the authority vested in me as President by the Constitution and the laws of the United States of America, and to maximize benefits of high-performance computing (HPC) research, development, and deployment, it is hereby ordered as follows:

The NSCI is a whole-of-government effort designed to create a cohesive, multi-agency strategic vision and Federal investment strategy, executed in collaboration with industry and academia, to maximize the benefits of HPC for the United States.

https://www.whitehouse.gov/the-press-office/2015/07/29/executive-order-creating-national-strategic-computing-initiative
https://www.whitehouse.gov/sites/default/files/microsites/ostp/nsci_fact_sheet.pdf

NSCI Intent

- **National**
  - "Whole-of-government" and "whole-of-Nation" approach
  - Public/private partnership with industry and academia
- **Strategic**
  - Leverage beyond individual programs (a key "platform" technology)
  - Long time horizon (decade or more)
- **Computing**
  - HPC = most advanced, capable computing technology available in a given era
  - Multiple styles of computing and all necessary infrastructure
  - Scope includes everything necessary for a fully integrated capability
    - Theory and practice, software and hardware
- **Initiative**
  - Above baseline effort
  - Link and lift efforts

Enhance U.S. strategic advantage in HPC for economic competitiveness and scientific discovery

Key Themes

- Strive for convergence of numerically intensive and data-intensive computing
- Keep the U.S. at the forefront of HPC capabilities
- Streamline HPC application development
- Make HPC readily usable and accessible
- Establish hardware technology for future HPC systems

Convergence of Compute and Data-intensive Science
Critical to 21st Century Science

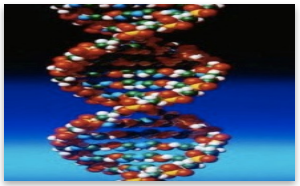

Genomics
- Sequencer data volume increasing 12x,next 3 years
- Sequencer cost deceasing by 10 over same period

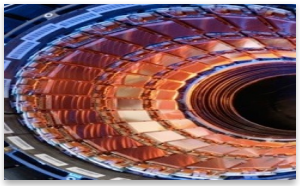

High Energy Physics
- LHC Experiments produce petabytes of data/year
- Peak data rates increase 3-5x over 5 years

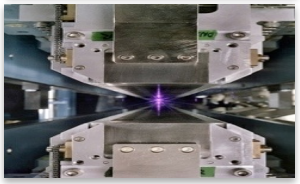

Light Sources
- Many detectors on Moore's Law curve
- Data volumes rendering previous models obsolete

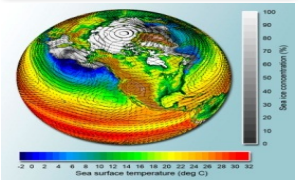

Climate
- By 2020, climate data expected to be exabytes
- Significant challenges in data management & analysis

- **DOE missions require computational environments that address both compute and data-intensive simultaneously**
- **Data-intensive science faces many of the same technology challenges of extreme-computing**
  - Some are even worse for "big-data"
  - Energy use is the grand challenge (e.g. the square kilometer array estimates 100MW needed for computing)

"Very few large scale applications of practical importance are NOT data intensive."
Alok Choudhary, IESP, Kobe, Japan, April 2012

System attributes	NERSC Now	OLCF Now	ALCF Now	NERSC Upgrade	OLCF CORAL Upgrade	ALCF CORAL Upgrades	
Name Installation	Edison	TITAN	MIRA	Cori 2016	Summit 2017-2018	Theta 2016	Aurora 2018-2019
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	710 TB	768 TB	~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 Perf. (TF)	0.460	1.452	0.204	>3	>40	>3	>17 times Mira
Node processors	Intel Ivy Bridge	AMD Opteron Nvidia Kepler	64-bit PowerPC A2	Intel Knights Landing Xeon Phi Intel Haswell CPU in data partition	IBM Power9 CPU Nvidia Voltas GPUS	Intel Knights Landing Xeon Phi	Intel Knights Hill Xeon Phi
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	Gemini	5D Torus	Aries	Dual Rail EDR-IB	Aries	2nd 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™	10 PB 210 GB/s Lustre®	150 PB 1 TB/s Lustre®

High Performance Computing (HPC)

US Federal Government Investments

- US federal High Performance Computing (HPC) investments
  - Made pivotal investments in the computer industry at critical times
  - During stable times, no investment is required or requested
  - Today, we have reached a critical period: *confluence of digitalization of our economy and society and the end of Dennard Scaling*
- Previous US Federal HPC investments

  Fueled major HPC advances
  - 1946 ENIAC: start of electronic digital computing
  - 1951 ERA-1101: technical computing
  - 1972 ILLIAC IV: parallel computing
  - 1993 Cray T3D: massively parallel computing
  - 2004 IBM BG: low power computing
  - 2011 Cray XC30 & IBM POWER7: productivity computing
  - 2023 Exascale: energy efficiency computing

Uncertainty Threatens US Economic Growth

- The world has changed – technology is changing at a dramatic rate
  - Dennard Scaling has ended
  - End of Moore's Law looming
- The IT marketplace is also changing dramatically
  - PC sales have flattened
  - Handhelds dominate growth, H/W and S/W
  - HPC vendor uncertainty
- Need to drive innovations at all levels of technology
  - Node and system designs
  - System and development software
  - Workflows
  - Algorithms

IDC Worldwide IT Data

IT Spending ($m)					
Row Labels	**2015**	**2016**	**2017**	**2018**	**2019**
Devices	**795,402**	**807,257**	**810,643**	**814,677**	**809,947**
Enterprise Hardware	**249,493**	**256,669**	**264,770**	**272,898**	**280,366**
Software	**434,727**	**464,242**	**496,264**	**530,551**	**568,049**
Services	**668,305**	**690,734**	**713,829**	**738,019**	**762,725**
Total IT	**2,147,927**	**2,218,902**	**2,285,507**	**2,356,145**	**2,421,087**

HPC Only:	2015	2016	2017	2018	2019
Revenue $M	**11,434**	**12,327**	**13,286**	**14,160**	**15,262**

Source: IDC 2016

Challenges – across the IT marketspace

- **Increasing Performance / Value**
- **Efficiency**
  - Energy efficiency: reduce energy per operation (pJ/op)
  - Hardware efficiency: massively parallel architectures, down to the processor level
  - Software efficiency: effectively exploit H/W parallelism
    - Current efficiencies on conventional machines are < 10% for many real-world applications
- **Memory / Storage**
  - Make effective use of data movement (this is the dominant energy cost)
- **Reliability**
  - Successfully complete execution through system failures
- **Productivity**
  - Programming environment that makes HPC machines accessible to everyone
  - Reduce time to solution
- **Cost / Affordability**

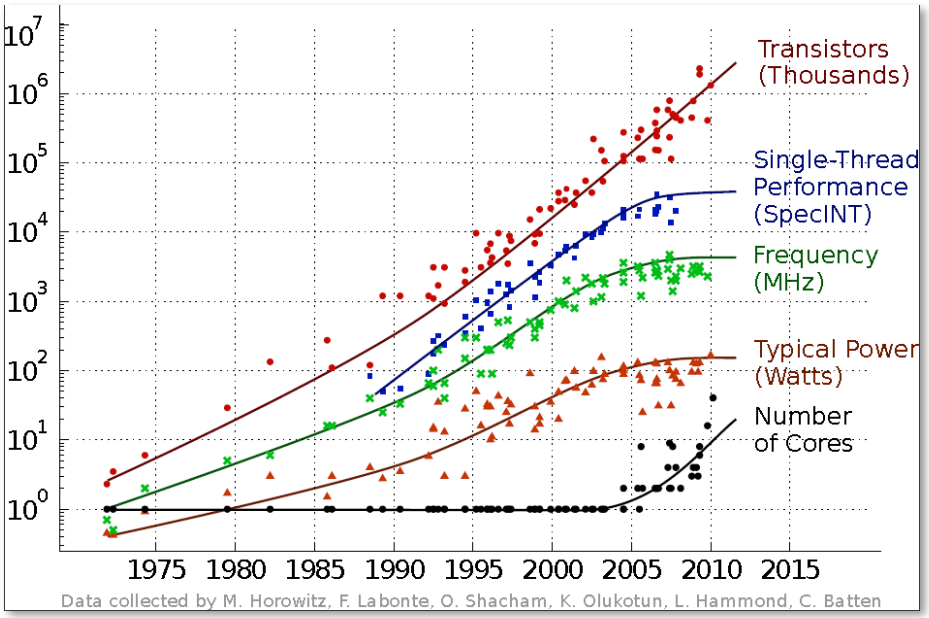

From Giga to Exa, via Tera & Peta

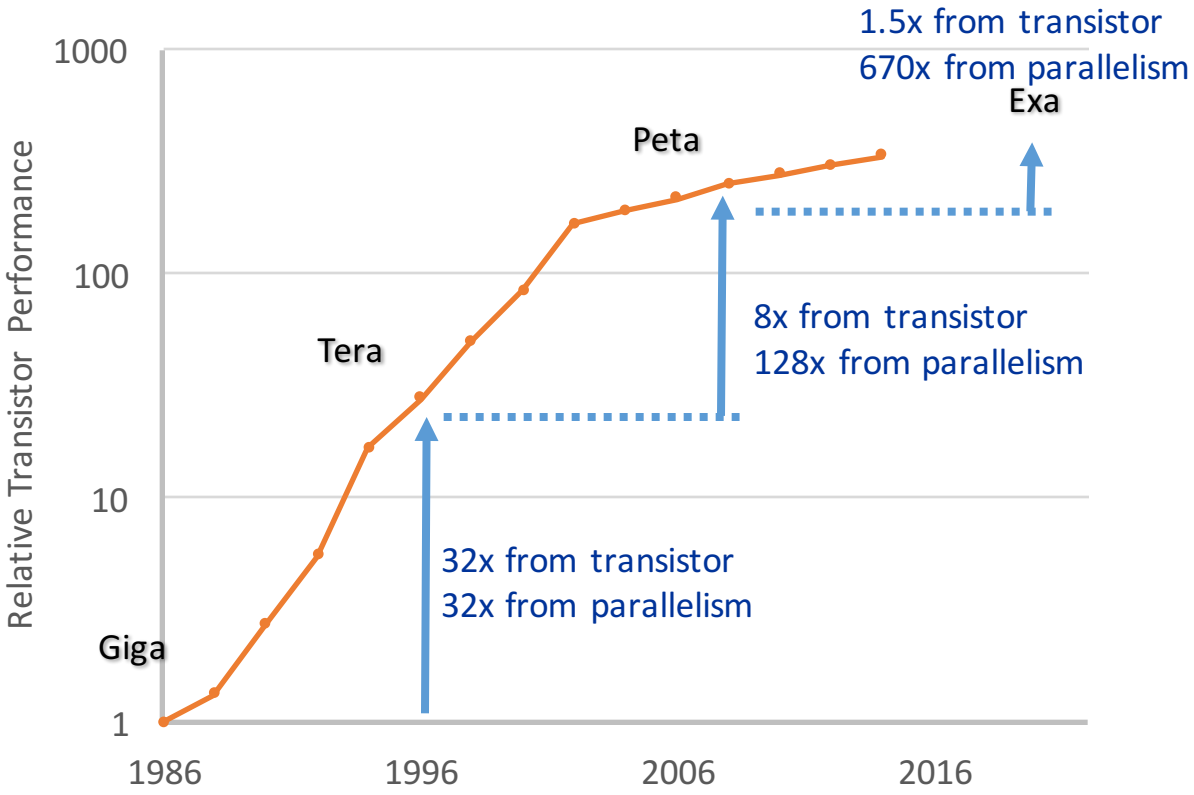

Shekhar Borkar, Intel

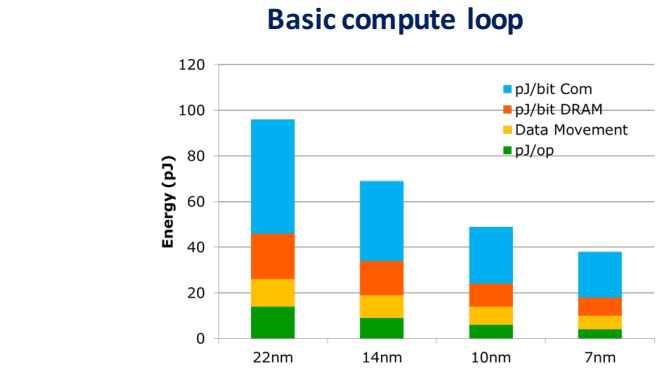

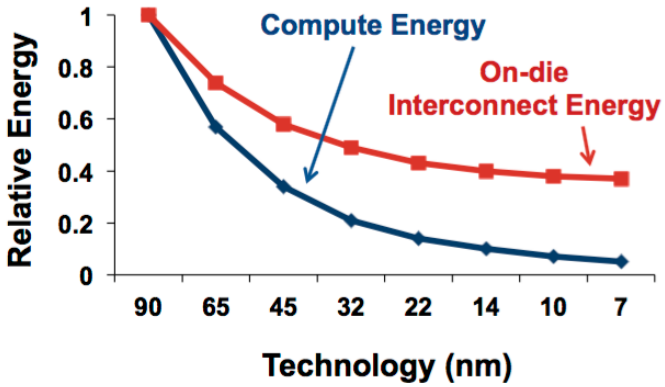

Shekhar Borkar, Journal of Lightwave Technology, 2013

Performance Factors - SLOWER

Thomas Sterling, IU

$$P = s(S) \times e(L, O, W) \times U(E) \times a(R)$$

P – average performance (ops)
e – efficiency (0 < e < 1)
s – application's average parallelism,
a – availability (0 < a < 1)
U – normalization factor/compute unit
E – watts per average compute unit
R – reliability (0 < R < 1)

Definition of SLOWER terms

- **S**tarvation
  - Insufficiency of concurrency of work
  - Impacts scalability and latency hiding
  - Effects programmability
- **L**atency
  - Time measured distance for remote access and services
  - Impacts efficiency
- **O**verhead
  - Critical time additional work to manage tasks & resources
  - Impacts efficiency and granularity for scalability
- **W**aiting for contention resolution
  - Delays due to simultaneous access requests to shared physical or logical resources

Hardware Architecture Impact on SLOWER metrics

- **Starvation**
  - Support for fine grain parallelism and light weight messaging, eliminate global barriers
- **Latency**
  - Put memory and computational elements in close proximity, support message driven computation
- **Overhead**
  - Reduce times for thread creation and context switching, support for GAS
- **Waiting for contention resolution**
  - Increase bandwidths for memory, networks, and ALUs with adaptive scheduling, routing, and resource allocation

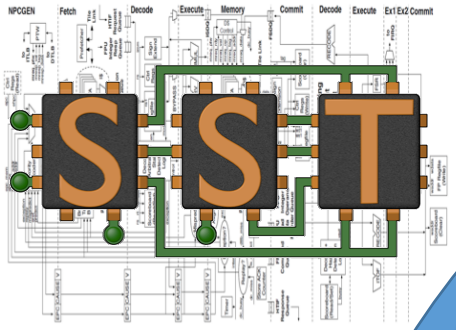

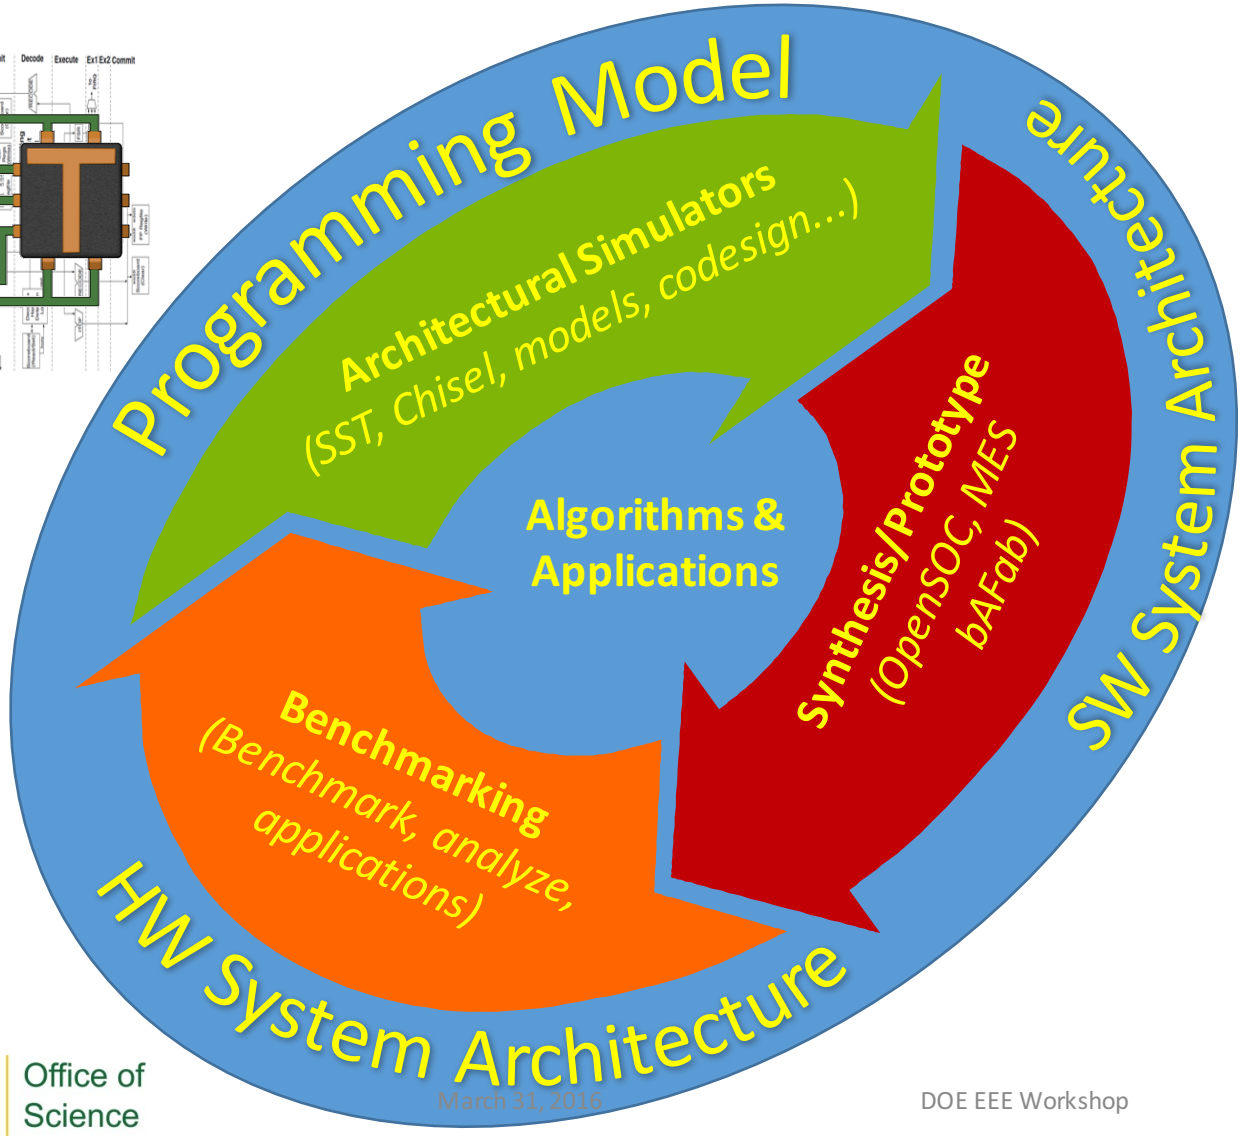

I'm supposed to be a scientific person but I use intuition more than logic in making basic decisions.
Seymour Cray

Read more at:
http://www.azquotes.com/quote/729170

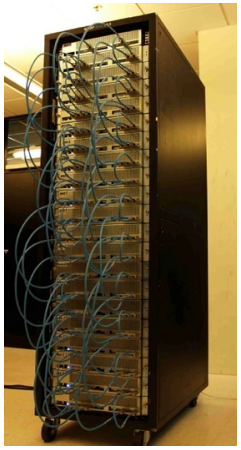

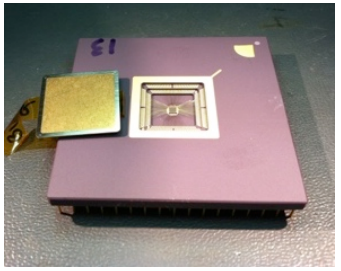

The post-2025 applications, software and hardware efforts desperately require the utilization of an application-driven co-design process

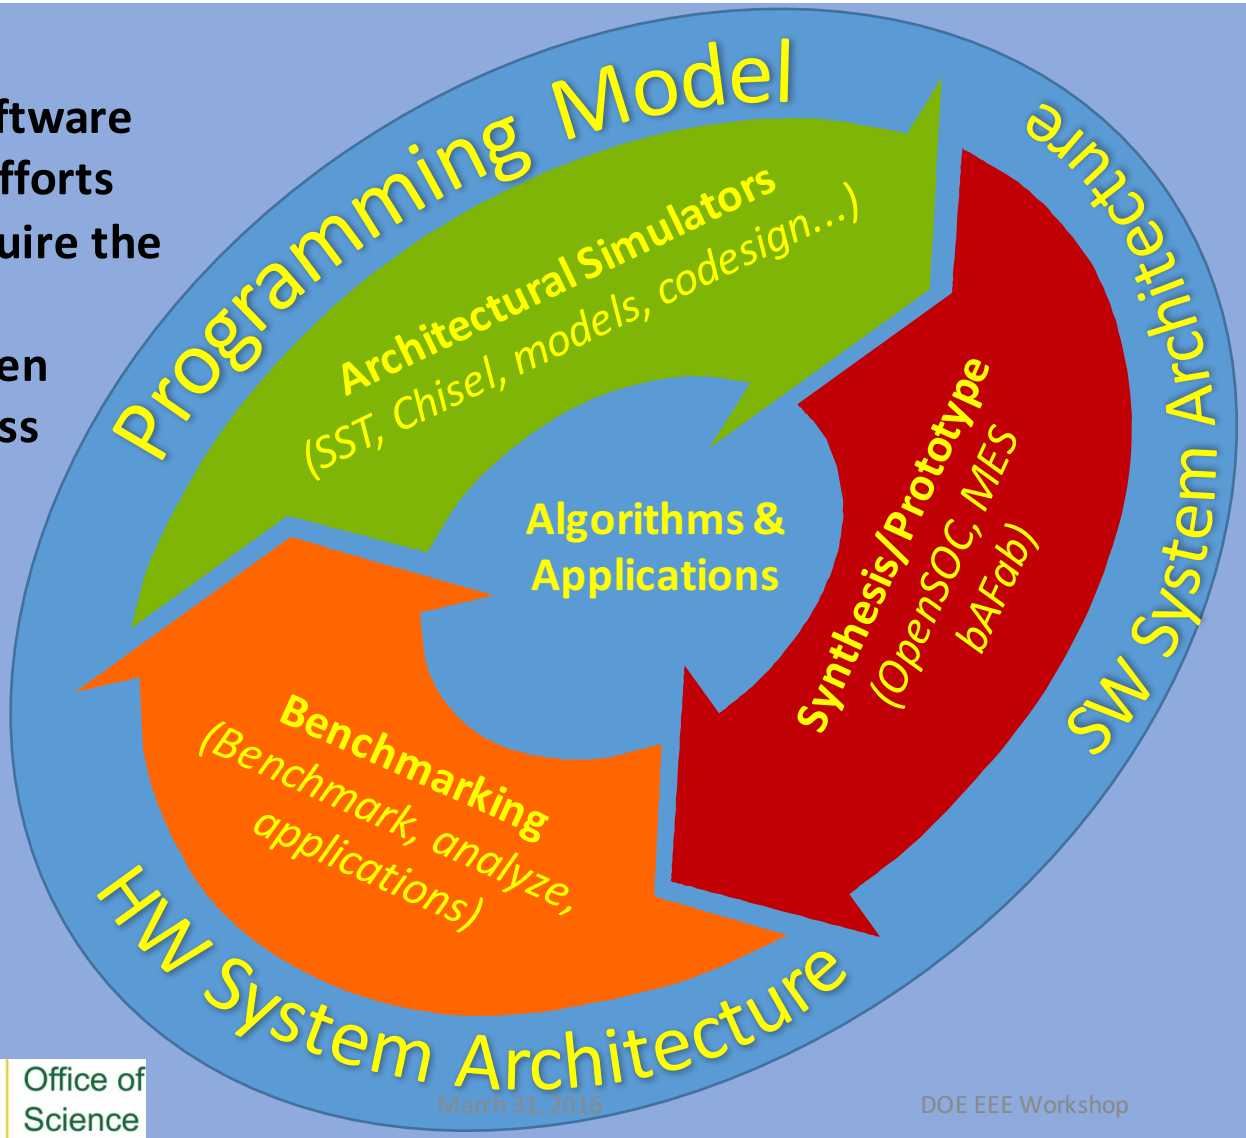

Application-driven co-design is the process by which:

- Scientific problems requirements guide the computer architecture and system software design
- Technology capabilities and constraints inform formulation and design of algorithms, applications and software

U.S. DEPARTMENT OF ENERGY | Office of Science

Hardware Architecture Research Areas

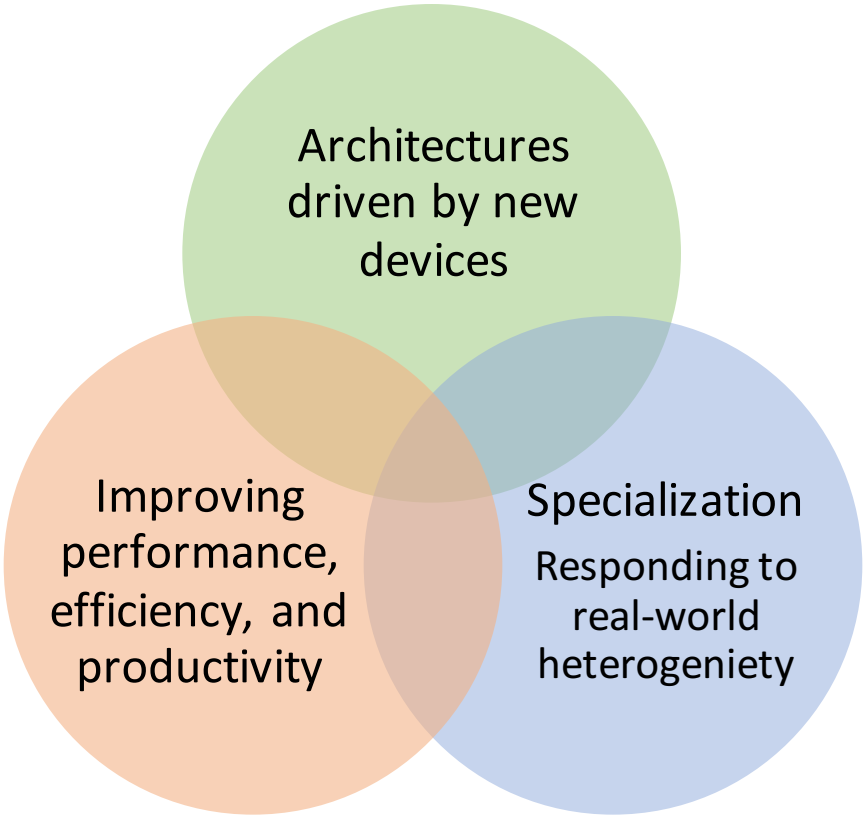

Architectures driven by new devices

ARPA-e: SWITCHES

SWITCHES projects aim to find innovative semiconductor materials, device architectures, and device fabrication processes that will enable increased switching frequency

IBM Research Initiative

Research programs are aimed at "7 nanometer and beyond" silicon technology and developing alternative technologies for post-silicon-era chips using entirely different approaches
($3 Billion investment)

Stanford N3XT Project

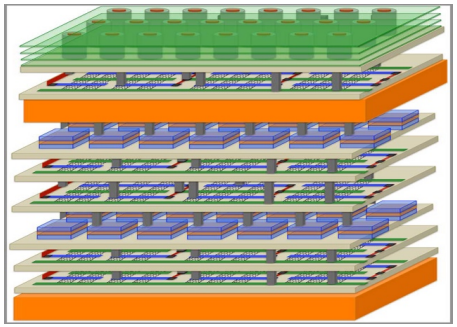

Stanford-led skyscraper-style chip design boosts electronic performance by factor of a thousand

Carbon nanotube transistors

Improving performance, efficiency, and productivity

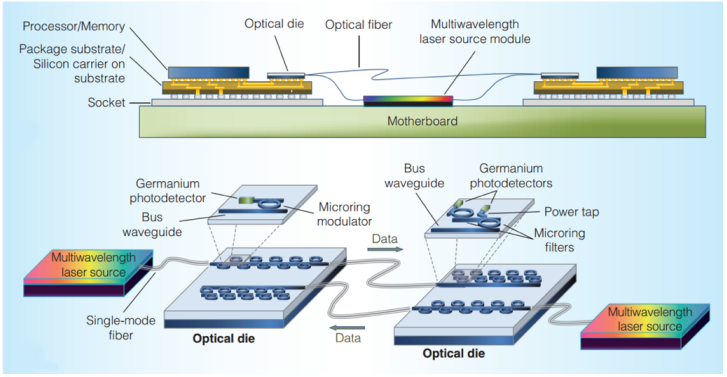

Silicon photonics (SiP)

Bandwidth

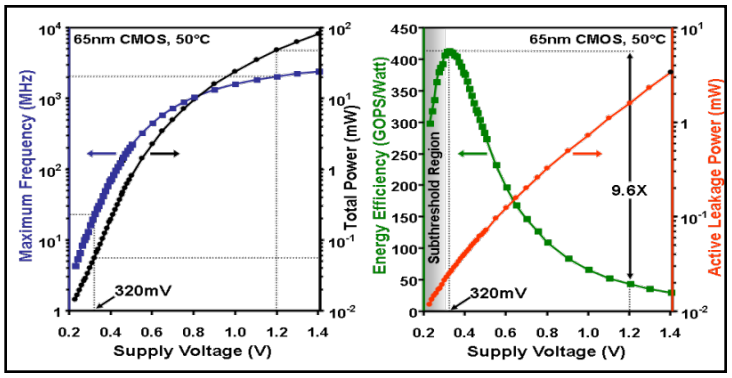

Near Threshold Voltage (NTV)

Energy Efficiency

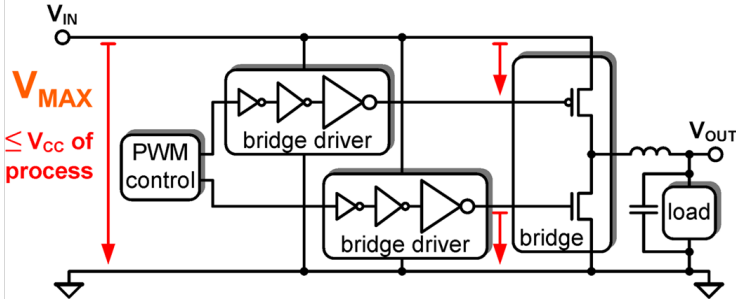

Circuits

Power Reduction

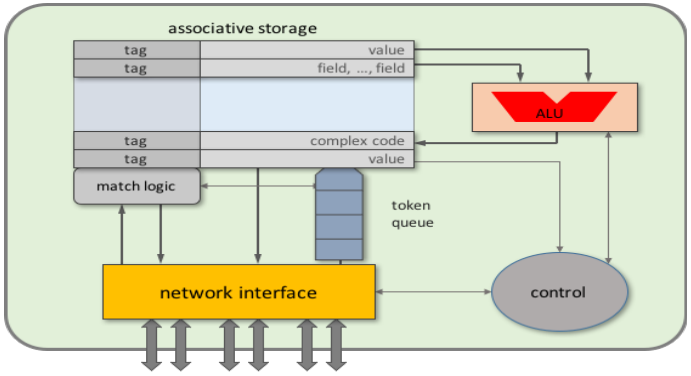

Specialization

Integrating current and future processing technology into the computing fabric – heterogeneous processing

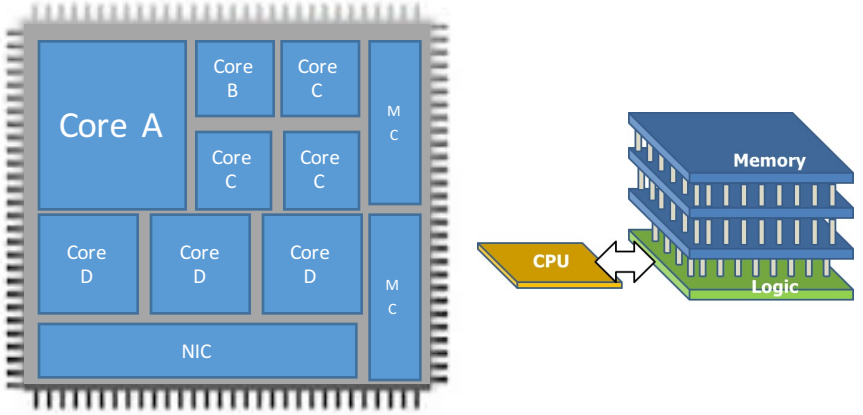

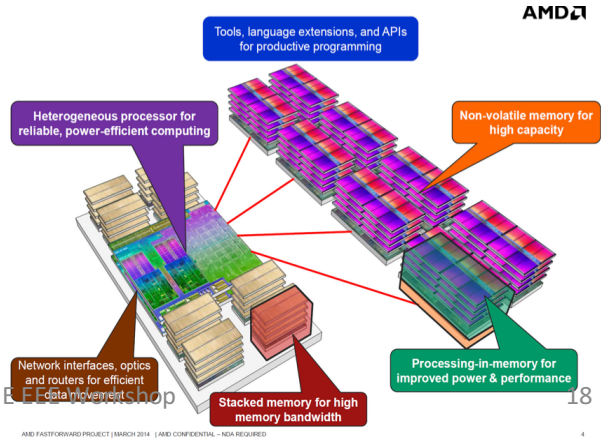

Final Words

- The semiconductor industry is crucial to the U.S. economy – *drives a > $2T/yr IT market*
- What's after CMOS? – *efficient CMOS & substantially improved HW/SW architectures*
- There is an exponentially increasing demand for information technology – *new technologies are limited by cost to develop and manufacture, not innovations*
- It isn't clear what is the enabling technology for the "Post CMOS" | "Post Moore's Law" epoch – *requires involvement from many different organizations*
- We need a significant investment in HW / SW architecture that utilizes the co-design process – *$$$*
- We need to stop trying to make tomorrow's computers look and operate like yesterday's computers – *don't be constrained by the von Neumann paradigm*
- The ultimate challenge isn't finding the technical solutions – *it's accepting that a change in how we think about computers is required to enable the next major advances*