SAND2016-10467C

Exceptional service in the national interest

Sandia National Laboratories

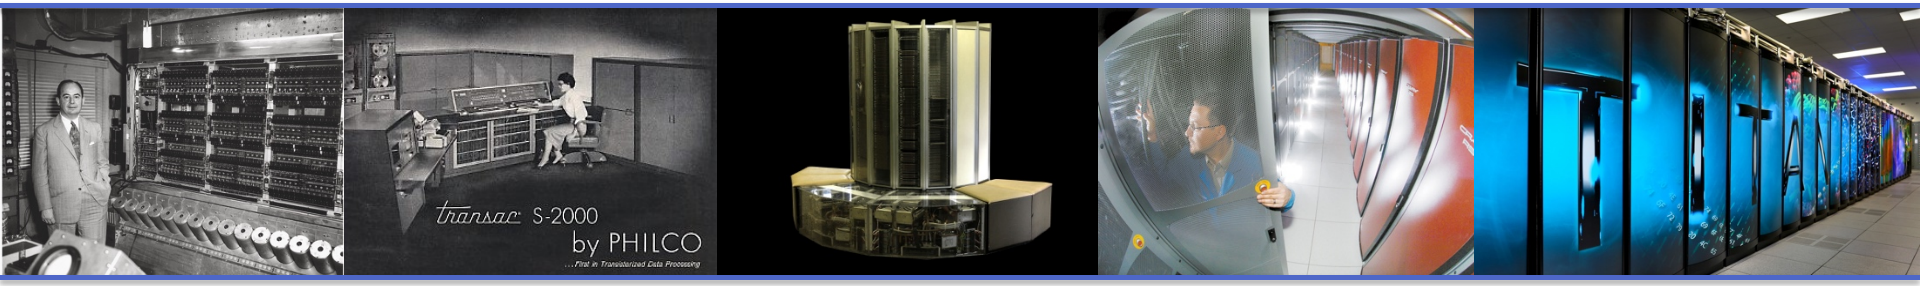

Historical Impact of Government Investment in High Performance Computing

IEEE International Conference on Rebooting Computing
San Diego, CA
October 17, 2016

Rob Leland
Vice President, Science & Technology
Chief Technology Officer
Sandia National Laboratories

U.S. DEPARTMENT OF ENERGY

NNSA National Nuclear Security Administration

Sandia National Laboratories is a multi-mission laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Outline

- Some context
- My thesis
- Historical review of major transitions in computing
- Laying out the pattern
- Looking forward

Some context

- **Why might this be important?**
  - Many indicators suggest we are in need of a major advance in computing
  - If we better understand how such advances have come about in the past, we are more likely to generate similar advances in the future

- **What are we really talking about?**
  - High Performance Computing (HPC) means …
    - *Computing systems operating at or near the maximum performance achievable in a given technological era.*
  - My thesis derives from a subjective historical pattern

- **How can we make use of this thesis?**
  - What are the indicators of a coming transition?
  - What approach/mindset is suggested by history?

My thesis regarding evolution of HPC

- **Periodic shifts in architecture and programming models**
  - Response to national security context in which …
  - Prevailing capability seen as insufficient for future defense or intelligence needs
- **Overlaying shifts with government investments shows pattern**
  - Government funded pilot/prototype typically *preceded* shift by 5-7 years
- **Suggests government investment has played a key role in catalyzing shifts**
  - Resonates for many in field
  - Particularly if they lived through an example
- **Acknowledge**
  - Correlation NOT= causation
  - Other contributing factors
  - Relative contribution often not clear
  - Some major government investments did not lead to a shift
  - Viewed from predominantly U.S. government perspective

Early electronic computing era

- **Electronic computing developed to meet military needs in WWII**
  - Colossus, Bletchley Park, 1943 … code breaking
  - ENIAC, U. Pennsylvania, 1945 … ballistics tables
  - Von Neumann
    - Initiated 1946 at Institute for Advanced Studies (IAS)
    - Hydrog simulations
    - IAS, Princeton, 1951
    - MANIAC, LANL, 1952
    - ORACLE, ORNL, 1953

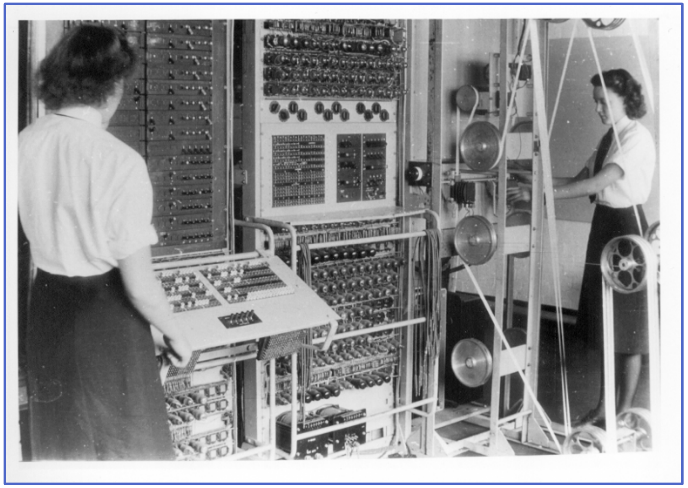

Women's Royal Naval Service operating Colossus during World War II

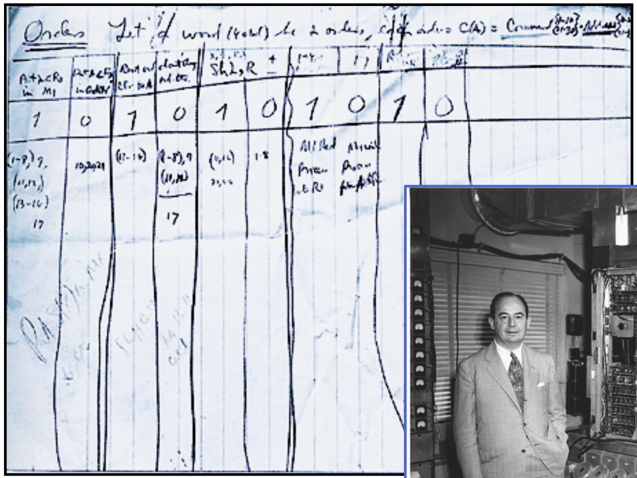

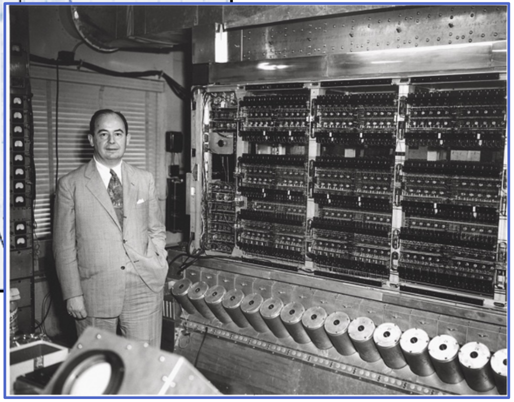

First line of code written for the von Neumann Digital Computing Project

Mainframe era

- **Large govt. investments for military and intelligence applications led to civilian “mainframe” market:**
  - Data processing
  - Weather, climate modeling
  - Scientific R&D
- **IBM701, 1953, close copy of IAS machine**
  - Vacuum tube-based system
- **Solid state computing**
  - Philco’s transistorized computer, 1955, NSA funded
  - Transac S-2000, 1957, first commercial version
- **NSA and AEC (DOE predecessor) are primary sponsors of HPC mainframe development throughout 1960’s**
- **CDC 6600, 1966, federal sponsorship**
  - Fast logic, RISC, functional parallelism
  - Order of magnitude performance improvement
  - Inspired famous TJ Watson memo

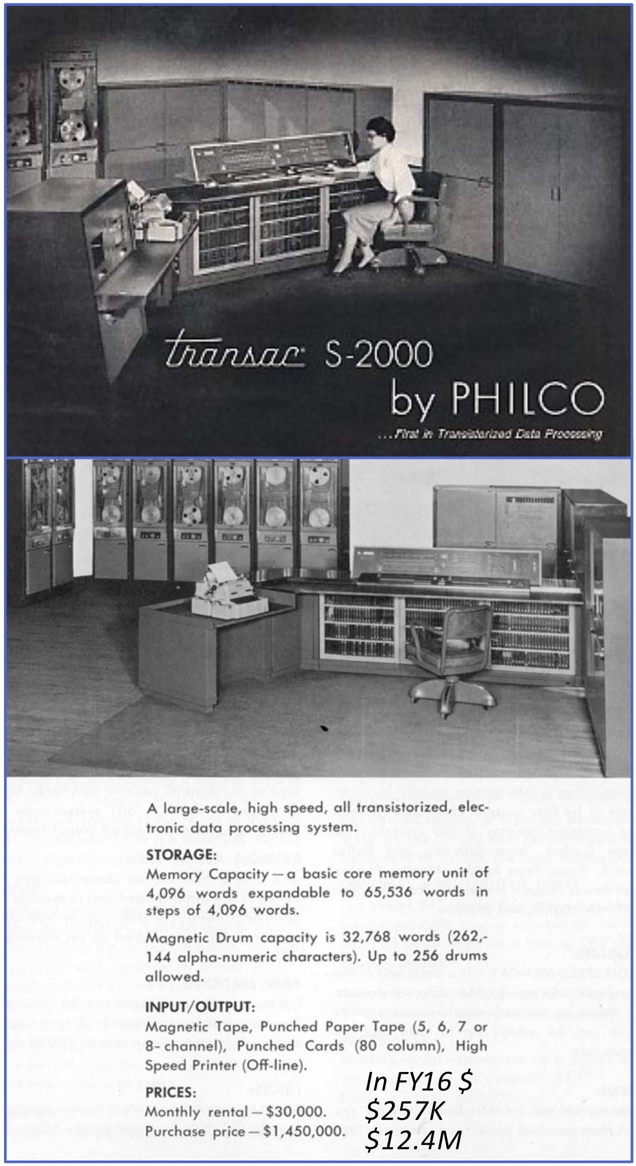

In FY16 $
$257K
$12.4M

Vector era

- **In the 1970s, National Lab needs outpaced performance growth of mainframes, driving development of vector-based (data parallel) systems**

- **STAR-100, LLNL, first IC system**
  - Million word memory
  - Legacy codes a poor fit
  - Lab overhauls codes to take advantage of vector machine architecture

- **Seymour Cray leaves CDC to form Cray Research**
  - Cray 1 builds on STAR-100, adds optimized scalar processing, memory size

- **Critical partnerships between Labs and computing suppliers**
  - New machines and software for vector systems
  - Aids expansion of supercomputing into new areas (oil and gas, aerospace, automotive)

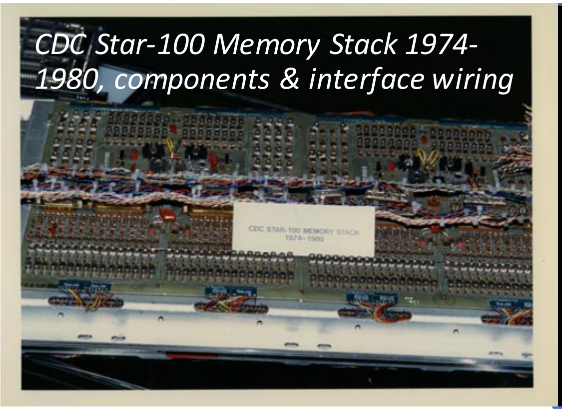
CDC Star-100 Memory Stack 1974-1980, components & interface wiring

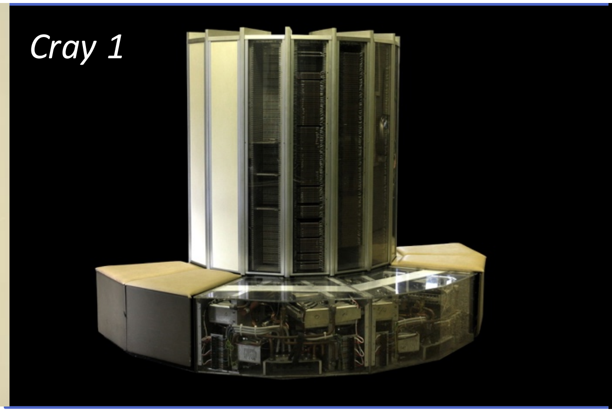
Cray 1

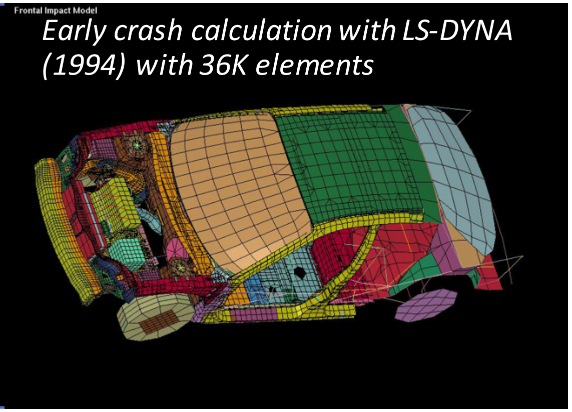

Early crash calculation with LS-DYNA (1994) with 36K elements

Distributed memory MPP era

- **Attack of the killer micros, late 1980's**
  - Personal computer market drives key components to commodities
  - Used as building blocks in innovative architectures

- **High Performance Computing and Communications Initiative, 1991**
  - Multi-agency federal program to invest in leading edge technology development
    - DoD, DOE, NASA, NSF, etc.
  - Early MPP machines, algorithms, tools, applications, education
  - Fiber optical network, web browser, ...

- **DOE/NNSA Accelerated Strategic Computing Initiative, 1995**
  - Accelerated shift from vector systems to commodity microprocessor-based, distributed memory MPP systems
  - ASCI Red, 1996, Intel/SNL – first terascale system
  - Blue Gene, 2005, IBM/LLNL – advanced power efficiency
  - Red Storm, 2005, Cray/SNL – high bandwidth interconnect
  - Linux cluster commoditization of MPP technology

Cray Red Storm, SNL

IBM Blue Gene, LLNL

SALINAS weapon system simulation; ASCI Red, 3000 processors (2002)

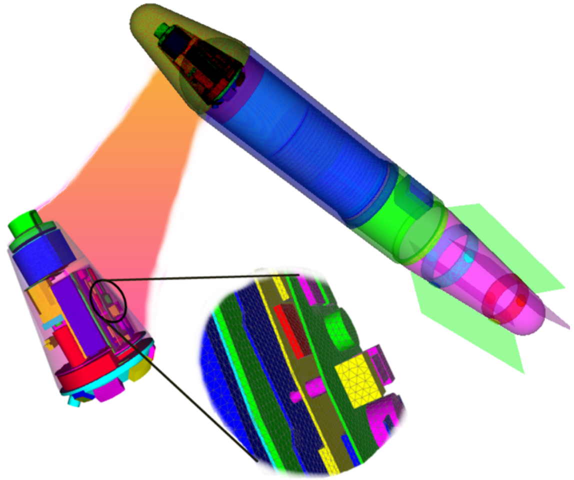

Heterogeneous node/many core era

- **Shift from scaling up to scaling out**
  - Add processing cores to a single chip to extend gains from existing technology
  - Add accelerator chips to node to leverage commodity technology from other markets
  - Software development focus on associated parallel processing challenges

- **DOE/NNSA ASCI is formalized as the ASC Program, 2004**
  - Roadrunner, 2008, IBM/LANL – hybrid CPU processor architecture
  - Commodity Technology Systems – 7 petaflop capacity
    - "Scalable Unit" Linux clusters (2015+) replace the Tri-Lab Linux Capacity Cluster (2011)
  - Advanced Technology Systems – goal of 40+ petaflop capability
    - Trinity, 2015, Cray/LANL/SNL, 11 petaflops, "burst buffer" storage, advanced power management

- **DOE/SC ASC Research Program (ASCR) to supply HPC**
  - National Leadership Computing Facilities at ORNL/ANL, 2004
    - Titan, 2012, Cray/ORNL – hybrid CPU/GPU processor architecture
    - Mira, 2012, IBM/ARNL – third generation Blue Gene

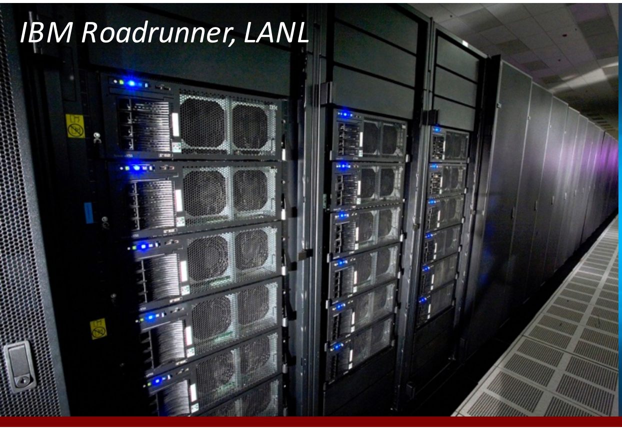
IBM Roadrunner, LANL

Cray Titan, ORNL


First successful numerical simulation of ground-based acoustic tests of weapons systems (2014)

Historical investment patterns and eras of HPC

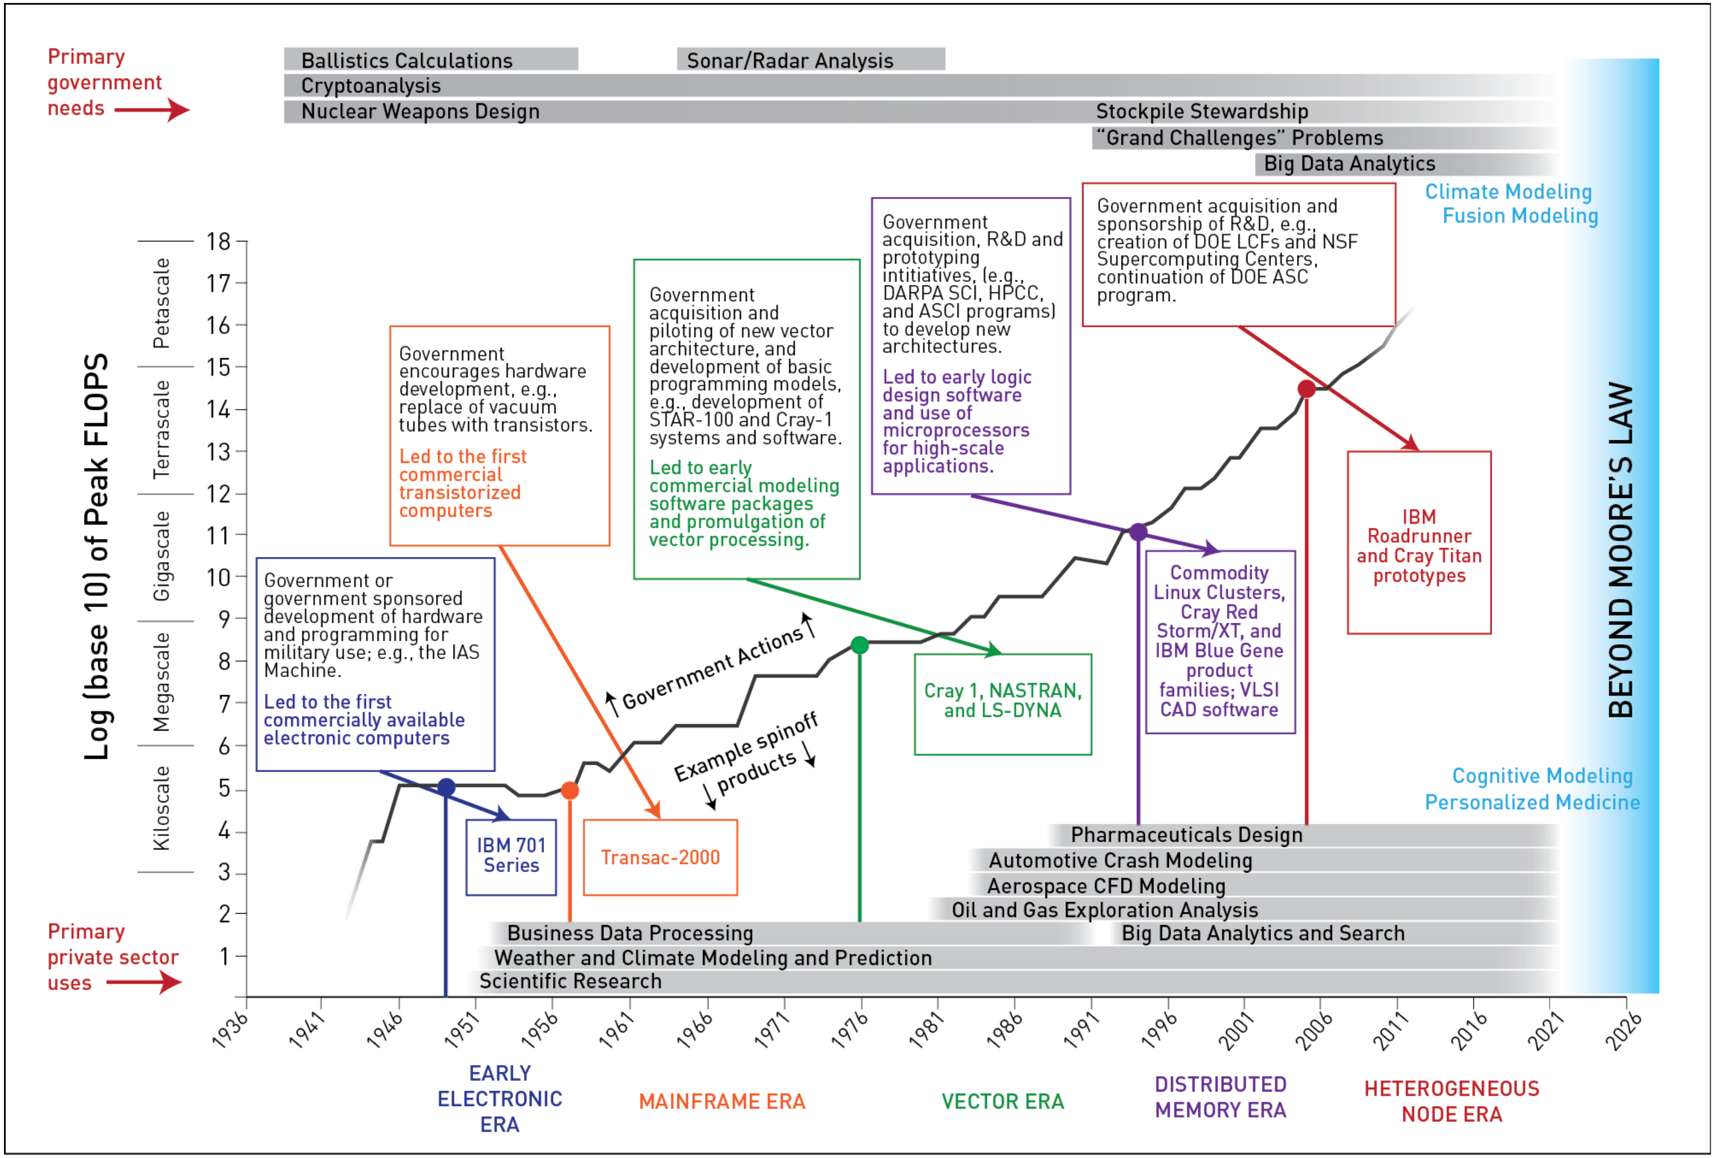

Looking forward

- **Lessons from the past**
  - Urgency
  - Technical opportunity
  - Conceptual breakthrough
  - Collaboration
- **Strategic drivers for change**
  - Increasing international competition at the high end
  - Erosion of Moore’s Law
  - Rise of Big Data
  - Coming to an end of the Massively Parallel Processing era
- **National Strategic Computing Initiative (NSCI)**
  - Executive Order issued July 2015
  - Ten agency joint vision, strategy, roadmap
  - Five strategic objectives …

NSCI Strategic Objectives

Sandia National Laboratories

- (1) Accelerating delivery of a capable exascale computing system that integrates hardware and software capability to deliver approximately 100 times the performance of current 10 petaflop systems across a range of applications representing government needs.
- (2) Increasing coherence between the technology base used for modeling and simulation and that used for data analytic computing.
- (3) Establishing, over the next 15 years, a viable path forward for future HPC systems even after the limits of current semiconductor technology are reached (the "post- Moore's Law era").
- (4) Increasing the capacity and capability of an enduring national HPC ecosystem by employing a holistic approach that addresses relevant factors such as networking technology, workflow, downward scaling, foundational algorithms and software, accessibility, and workforce development.
- (5) Developing an enduring public-private collaboration to ensure that the benefits of the research and development advances are, to the greatest extent, shared between the United States Government and industrial and academic sectors.

Exascale Computing Project

- **DOE Exascale Computing Project**
  - Partnership between SC and NNSA
  - Full HPC stack
  - 2016 PathForward RFP for a new exascale design meeting the following:
    - 50x improvement in application figure of merit
    - 20-30 MW peak power
    - 1 week on average between human intervention due to hardware or system faults

DDN/ORNL

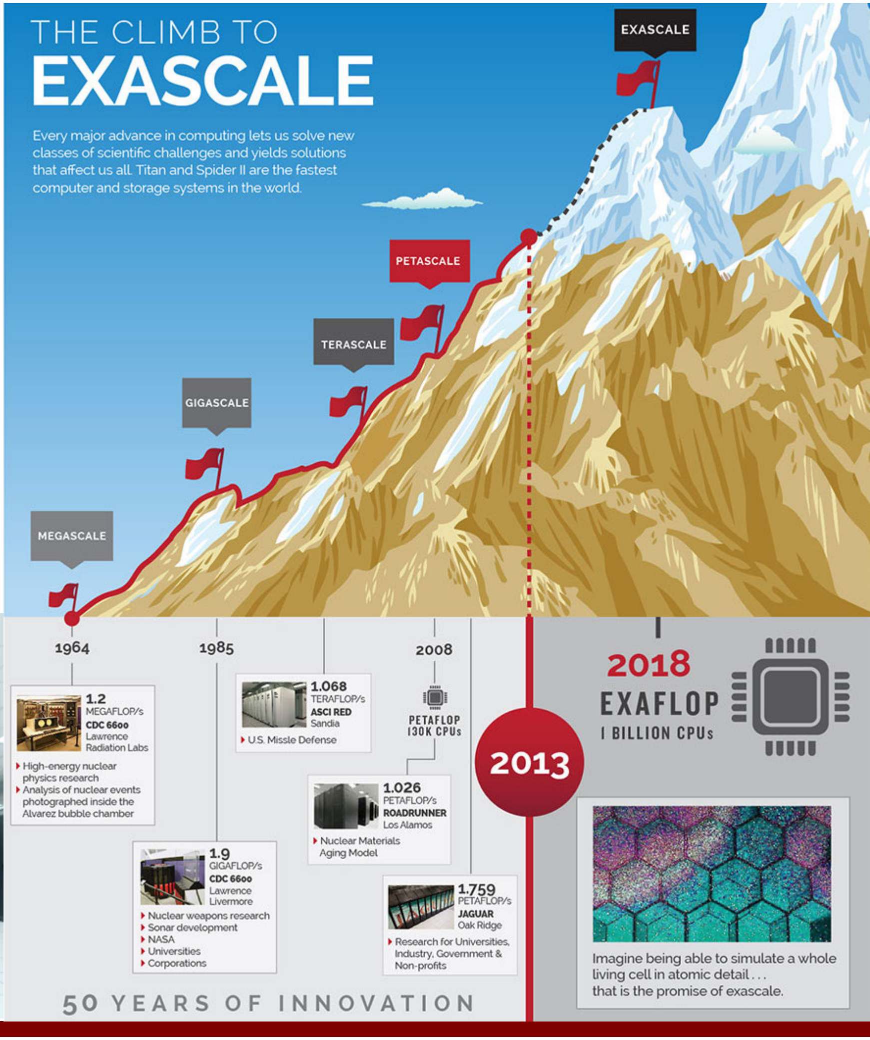

ORNL parallel coordinate visualization

TaihuLight, 10.6 milion cores, 33 petaflop/s

Beyond Moore's Law era

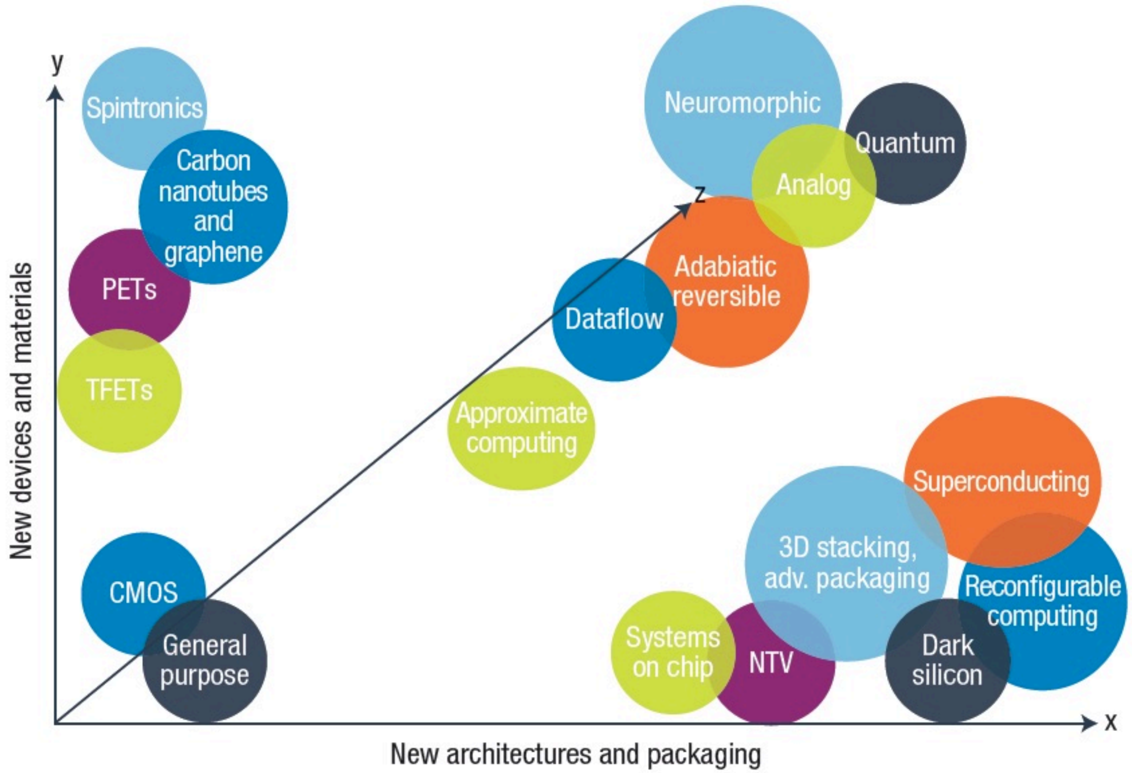

FIGURE 1. Technology scaling options along three dimensions. The graph's origin represents current general-purpose CMOS technology, from which scaling must continue. All the dimensions, which are not mutually exclusive, aim to squeeze out more computing performance. PETs: piezo-electric transistors; TFETs: tunneling field-effect transistors; NTV: near-threshold voltage.

Shalf and Leland, Computing Beyond Moore's Law, Computer 48:12, 2015.

IEEE Rebooting Computing Summit
Remarks, October, 2015

Some history first
Thought it would be good to review last really major "Reboot"

- *The modern digital era really began in the late 1940's with the development of the first electronic, stored program computer developed by Von Neumann's team at the Princeton Institute for Advanced Study and Los Alamos.*

Built out of vacuum tubes and oscilloscopes, programmed in assembler language

Plagued by technical challenges – as you can imagine

But momentous.

Motivated by and applied to key national security problems of the era
Design of the hydrogen fusion bomb
Cryptanalysis
Computing ballistics tables

Proved useful for early computational science efforts
Genetic evolution
Weather and climate

Later systems inspired by this design and rendered in solid-state became the foundation for the modern computing industry.

- *Ever since, leading edge computers – "High Performance Computers" as we now call them – have been crucial tools for national security, economic prosperity and scientific discovery.*

(Recommend Turing's Cathedral by George Dyson.)

What lessons to draw?
Conceptual advance
Combination of hardware and stored (flexible) program
Previous theoretical basis (Turing)

Technology opportunity
Experimental feedback – Feynman's approach not scalable

Mature base technology (vacuum tubes)
Mother of invention – tube memory, physical layout, reliability engineering
Engineering expertise

Urgency
Lessons of World War 2 (cryptography, ballistics, fission bomb)
Cold War driver

So here is one reading: conceptual advance, technology opportunity, sense of urgency

Welcome your thoughts on others

What are the key strategic issues today?

1. Increasing international competition at the high end.

As you all know, China now has the Top 1 Linpack system

Just one measure and of course much debated

Real challenge is in putting these big systems to good use
Programming them effectively to solve a variety of hard problems
Training the workforce to enable this more broadly
Designing the next generation of systems with more effective architectures.

Most HPC experts believe the US still leads here

But international competitors are working to close the gap

Significant efforts reported in Europe, Japan and Russia as well

2. Foundational technology base in microelectronics is set to go through a major transition as Moore's Law plateaus.

Thank Bob Colwell and Peter Highnam for their excellent presentations

Just will add emphasis to one key point: Already seeing the significant effects of this transition

Will drive a major change in the underlying algorithmic and software base as
Single processor paradigm fades away
Whole software stack will need to work on chips which have become parallel computers

More difficult to sustain historical rates of technical progress on which many mission and commercial capabilities now depend.

Require a broad and deep R&D effort to adapt the existing industrial code base to parallel processing

Impressive educational effort to train a work force to maintain and optimize this software base

(Recommend 2011 NRC report: *The Future of Computing Performance – Game over or Next Level?)*

3. Rise of Big Data

Somewhere between Neal Armstrong's moonwalk and Google/Twitter/BuzzFeed, we crossed over from the "space age" to the "information" or "digital" age.

Increasingly substitute information for materials/energy to improve societal efficiency (Toffler)

One estimate:

World information "content" currently growing 2+ zetabytes/yr.
Doubles every two years
85% unstructured

Improved sensing and processing drives ever faster

Internet, business, finance of course

But also science, engineering, security – and much more

Examples

Google Search	embedded in modern life
Human genome sequencing	transformational for medicine
Climate data analysis	better understanding fate of the planet
Discovery of Higgs particle	why there is a planet in the universe

Number of talks now about this as a "fourth" paradigm of science

Experiment
Theory
Simulation
` Data exploration

Certainly is highly significant

What is the relative growth of data and processing? Anyone know of an authoritative estimate?

Moore's Law growth in both and constant investment implies parity

Need to look more closely

People/institutions now looking at lot more data
Look at Order N data sets rather than Order 1 data set

Also a progression in sophistication of analysis

Arithmetic
Algebra
Calculus

So: More people, looking a more data, more closely

Aggregate complexity could certainly rise dramatically faster than Moore's Law

What are limitations – processing, networking, theory?

Etc.

How do these trends and the early "Rebooting" lessons correlate, if at all?

Urgency

Growing international competition (issue from US point of view)
Moore's Law erosion and necessary lead time to change course

Technical Opportunity

Peter Highnam's list of alternative technologies
New commodity curve: mobile, graphics, …
New architecture: data centric, processor in memory, global high BW interconnect, …

Conceptual breakthrough

What would be comparable to the idea of flexible programming?
Networked world – makes big data paradigm possible
Radically non Von Neumann architecture e.g. learning system (very little programming)
New mathematics?

Last seems to be where we are most challenged, but also perhaps where there is most leverage.

Welcome your ideas on the best path forward