U.S. DEPARTMENT OF ENERGY | Office of Science

Exascale Update

William J. Harrod
Division Director
Advanced Scientific Computing Research
Office of Science

November 18, 2013

Significant Events

- Meetings with the Secretary of Energy to discuss proposed Exascale Computing Initiative
- Meetings with the Secretary of Energy Advisory Board (SEAB) to discuss proposed Exascale Computing Initiative
- Developing an updated exascale project plan
- ASCAC Exascale Top 10 Study
- Programs: Design Forward, OS/R, UQ, RX-Solvers, Math/Stat for Data

Mission: Extreme Scale Science

Next Generation of Scientific Innovation

- **DOE's mission is to push the frontiers of science and technology to:**
  - Discovery science
  - Mission-focused basic science in energy
  - Provide state-of-the-art scientific tools
  - Plan, implement, and operate user facilities
- **The next generation of advancements will require Extreme Scale Computing**
  - 1,000X capabilities of today's computers with a similar size and power footprint
- **Extreme Scale Computing, however, cannot be achieved by a "business-as-usual" evolutionary approach**

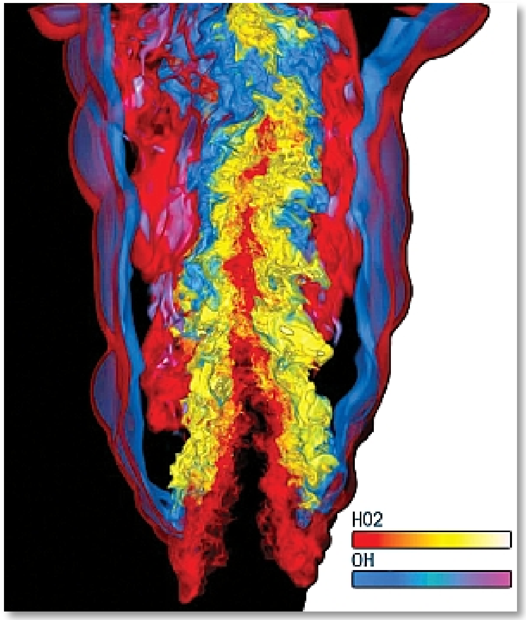

- **Extreme Scale Computing will require major novel advances in computing technology – Exascale Computing**

Exascale Computing Will Underpin Future Scientific Innovations

Mission: Extreme Scale Science
Data Explosion

Genomics

Data Volume increases to 10 PB in FY21

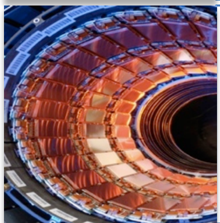

High Energy Physics (Large Hadron Collider)

15 PB of data/year

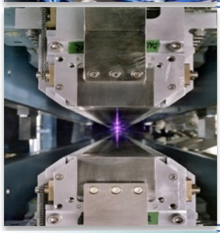

Light Sources

Approximately 300 TB/day

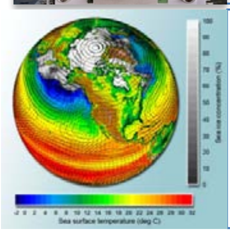

Climate

Data expected to be hundreds of 100 EB

Driven by exponential technology advances

Data sources

- Scientific Instruments
- Simulation Results
- Observational data

Big Data and Big Compute

- Analyzing Big Data requires processing (e.g., search, transform, analyze, …)
- Extreme scale computing will enable timely and more complex processing of increasingly large Big Data sets

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

Strategy for Moving Forward

- **Exascale Research and Development**
  - Enable extreme scale science
  - Develop energy efficient system designs
  - Highly productive and accessible for a wide range of users
- **Extreme Scale Applications**
  - Exploit the full potential of exascale systems
  - Drive future system requirements
  - Enable new scientific discoveries
- **Facilities**
  - Designed and optimized for the full application workflow
  - Data management infrastructure
  - Collaborative environments
  - “On ramp” to exascale systems

Exascale Computing
The Vision

- **Exascale computing**
  - Achieve order 10^{18} operations per second and order 10^{18} bytes of storage, computers with a similar size and power footprint for today's petascale systems
  - Address the next generation of scientific, engineering, and large-data problems
  - Set the US on a new trajectory of progress – towards a broad spectrum of computing capabilities over the succeeding decade
- **Productive system**
  - Usable by a wide variety of scientists and engineers
  - "Easier" to develop software & management of the system
- **Based on marketable technology**
  - Not a "one off" system
  - Scalable, sustainable technology, exploiting economies of scale and trickle-bounce effect
- **Deployed in early 2020s**

Exascale Computing
R&D Challenges and Issues

- **Four primary challenges must be overcome**
  - Parallelism / concurrency
  - Reliability / resiliency
  - Energy efficiency
  - Memory / Storage
- **System design issues**
  - Scalability
  - Efficiency
  - Time to solution
  - Reliability
- **Productivity issues**
  - Managing system complexity
  - Portability
  - Generality
- **Co-Design**

Exascale Computing
Proposed Timeline

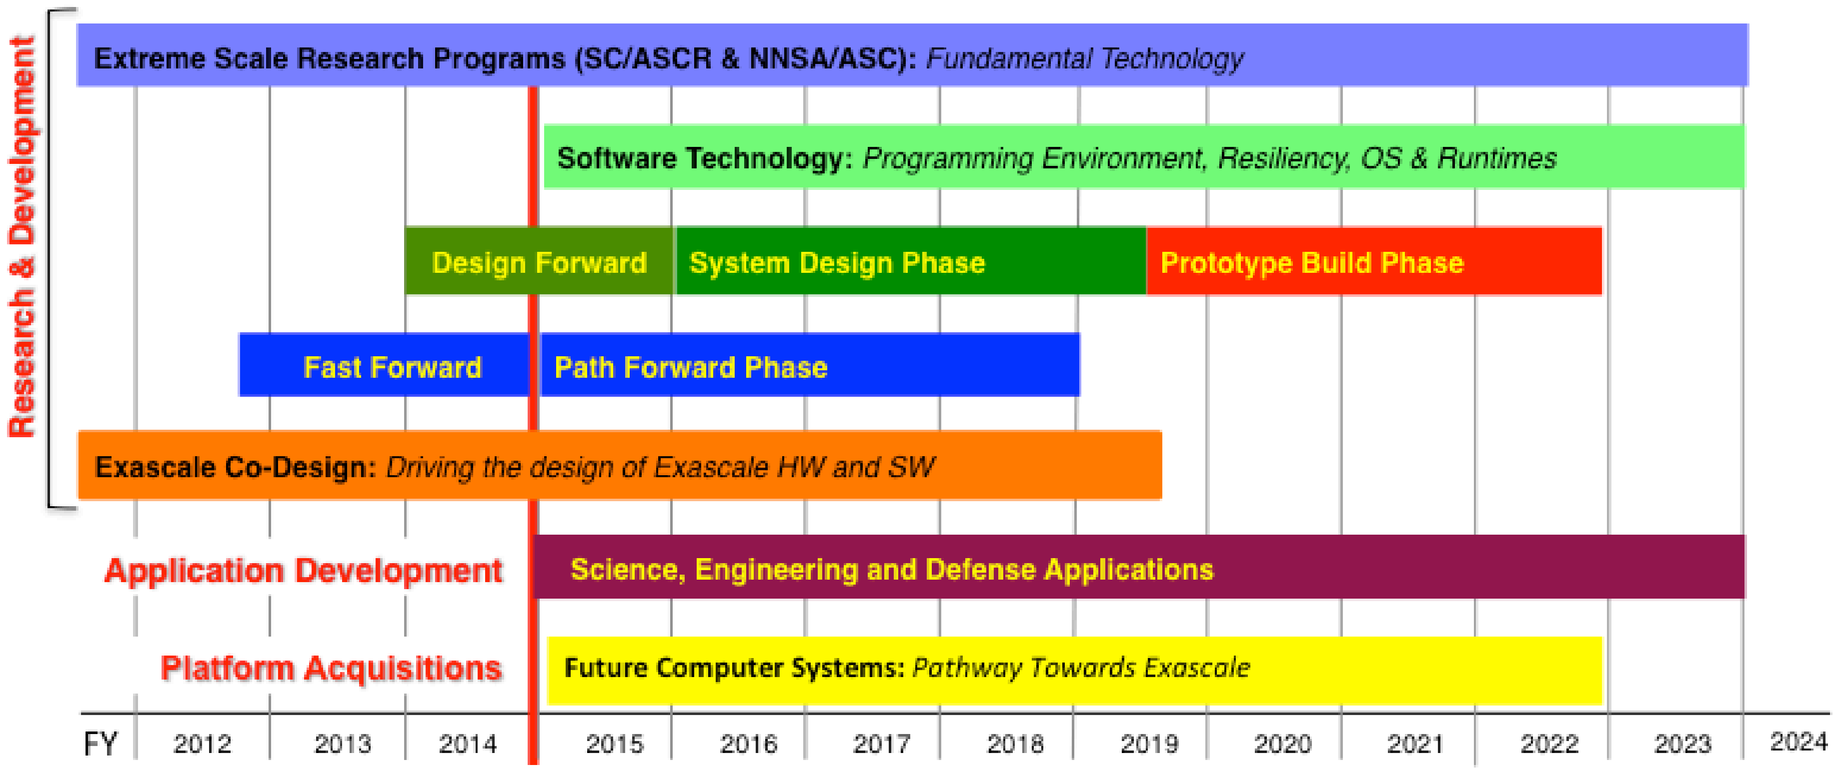

Exascale Co-Design

- **Application-driven co-design is the process by which:**
  - Scientific problem requirements guide computer architecture and system software design
  - Computer technology capabilities and constraints inform formulation and design of algorithms and software

- **Need shared global perspective across the design-space - to establish conceptual framework for co-design and interoperability**
  - Parallelism
  - Latency
  - Overhead
  - Dependability

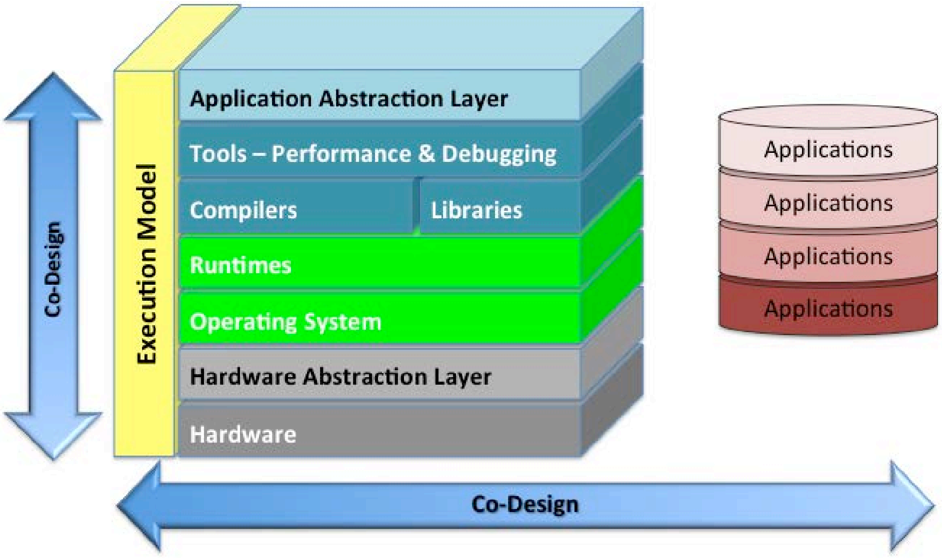

Exascale Co-Design Progress

- **Co-Design Centers have made (nontrivial) progress toward influencing hardware architectures:**
  - Extensive interactions with Fast Forward vendors, including through regular teleconferencing and numerous "Hack-a-thons" with Nvidia, AMD, and Intel
- **Co-Design Centers provide realistic case studies for the Exascale Ecosystem**
  - Collaborations include X-Stack, Design Space Explorations, Execution Models, and Performance Modeling and Simulation.

Exascale Co-Design Center for Materials in Extreme Environments (ExMatEx) http://exmatex.lanl.gov

Center for Exascale Simulation of Combustion in Turbulence (ExaCT) http://exactcodesign.org

Center for Exascale Simulation of Advanced Reactors (CESAR) https://cesar.mcs.anl.gov

Fast Forward Projects

Fast Forward

- Jointly funded by SC & NNSA
- Two year contracts, started July 1, 2012

Project Goals & Objectives

- Initiate partnerships with multiple companies to accelerate the R&D of critical technologies needed for extreme-scale computing.
- Fund technologies targeted for productization in the 5–10 year timeframe.

Vendor	SCOPE	Value
AMD	Processor / Memory	$12,600,000
IBM	Memory	$10,476,714
Intel	Processor / Memory	$18,963,437
NVIDIA	Processor	$12,398,893
WhamCloud (Intel)	Storage & I/O	$7,996,053
Total		**$62,435,097**

Design Forward Projects

- **Design Forward**
  - Jointly funded by SC & NNSA
  - Two year contracts, started Fall 2013
  - $25.4 Million in Contracts
- **Project Goals & Objectives**
  - Initiate partnerships with multiple companies to accelerate the R&D of interconnect architectures for future extreme-scale computers.
  - Fund technologies targeted for productization in the 5–10 year timeframe.
- **Projects Funded**
  - **AMD:** interconnect architectures and associated execution models
  - **Cray:** open network protocol standards
  - **IBM:** energy-efficient interconnect architectures and messaging models
  - **Intel:** interconnect architectures and implementation approaches
  - **NVIDIA:** interconnect architectures for massively threaded processors.

Extreme Scale Software Projects

2012 X-Stack

- D-TEC: LLNL and MIT
- Traleika Glacier:Intel
- DEGAS: LBNL
- XPRESS: Sandia
- DAX (ETI):
- Autotunig: U. Utah
- GVR: U. Chicago
- SLEEC: Purdue
- CORVETTE: UCB

Co-Design Centers

Misc. Software

- Exascale MPI

Mod / Sim

- CoDEX
- DMD
- Blackcomb
- Thrifty
- BSM
- Execution Models

2013 OS/R

- Hobbes: SNL
- Argo: ANL

Resilient Extreme-Scale Solvers

("RX-Solvers")

- **Program Goals**
  - Support basic research in scalable, resilient, extreme-scale solvers, targeted for supercomputers in the next 5-10 years
  - Establish the foundation for research in numerical algorithms for extreme-scale scientific computing
- **Program Requirements: projects must address**
  - Advances in solvers
  - Fault tolerance and resilience at the algorithmic level
  - Demonstration of performance of proposed algorithms
- **FOA Issued: 8 June 2012; closed: 13 August 2012**
- **18 proposed projects (48 proposals total) received**
- **Four (4) three-year-long projects started in June 2013**
- **Total funding: $4.5M per year for up to three years**

Mathematical and Statistical Methodologies for DOE Data-Centric Science at Scale

- **Basic research in novel mathematical and statistical methods, models, and tools for the representation, analysis, and understanding of DOE data-centric science at scale.**
- **Areas of interest:**
  - Data and dimension reduction
  - Automated analysis
  - Integration of observational data, experimental data, simulation and models
- **Funding Opportunity released May 2013 totaling $3M/year for 3 years**
- **Six projects selected for funding; projects to begin January 2014.**
- **Awarded projects address machine learning, scalable statistics, optimization, solution of statistical inverse problem, and Gaussian processes.**

UQ Methodologies for Enabling Extreme-Scale Science

- **Applied Mathematics basic research that significantly advances uncertainty quantification (UQ) methodologies as an enabling technology in extreme-scale scientific computing**
- **Six FY13 awarded-projects at $5M/year for 3 years**
- **Portfolio: DOE-mission science impact, rigorous UQ methodologies, and applied mathematics or statistics basic research advances**
  1. Extreme-scale Bayesian inference for UQ of complex simulations
     George Biros (UT Austin) & ORNL
  2. Probabilistic approach to enable extreme-scale simulations under uncertainty & system faults
     Bert Debusschere (Sandia National Labs) & Duke
  3. Scalable multilevel UQ concepts for extreme-scale multiscale problems
     Yalchin Efendiev (Texas A&M) & LLNL
  4. Scalable multi-chain MCMC methods for high-dimensional statistical inverse problems
     Jaideep Ray (Sandia National Labs) & PNNL
  5. Mathematical foundations for UQ in materials design
     Petr Plechac (Univ of Delaware) & Brown, UMass-Amherst
  6. Mathematical environment for quantifying uncertainty: Integrated & optimized at the extreme-scale
     Clayton Webster (Oak Ridge National Lab) & Sandia, Georgia Tech, Florida State

Conclusions and Summary

- **Unique opportunity to create the future epoch of Exascale computing**
  - Empower an extraordinary suite of extreme scale science and engineering applications
- **We are setting a new direction for future generations of computing**
  - Substantial research is required and has begun
- **We have made an exciting beginning**
  - We need to work together as a cohesive community to achieve the shared goal

BACKUP

Exascale Computing
Target System and Execution Strategy

- **Target System Characteristics**
  - 500 to 1,000 times more **performance** than a Petaflops system
  - **1 Billion** degrees of concurrency
  - **20 MW** Power requirement
  - **200** cabinets
  - **Development and execution time productivity improvements**
- **Perform research, development and integration required to deploy exascale computers in 2020+**
- **Partnership involving:**
  - Government
  - Computer industry
  - DOE laboratories
  - Academia
  - International researchers

Exascale Computing
We Need to Reinvent Computing

Traditional path of 2x performance improvement every 18 months has ended

- For decades, Moore’s Law plus Dennard scaling provided more, faster transistors in each new process technology
- This is no longer true – we have hit a power wall!
- The result is unacceptable power requirements for increased performance

We cannot procure an exascale system based on today’s or projected future commodity technology

- Existing HPC solutions cannot be usefully scaled up to exascale
- Energy consumption would be prohibitive (~300MW)

Exascale will require partnering with the computing industry to chart the future

- Industry at a crossroads and is open to new paths
- Time is right to push energy efficiency into the marketplace

Limits in Device Physics Abet Increased Parallelism

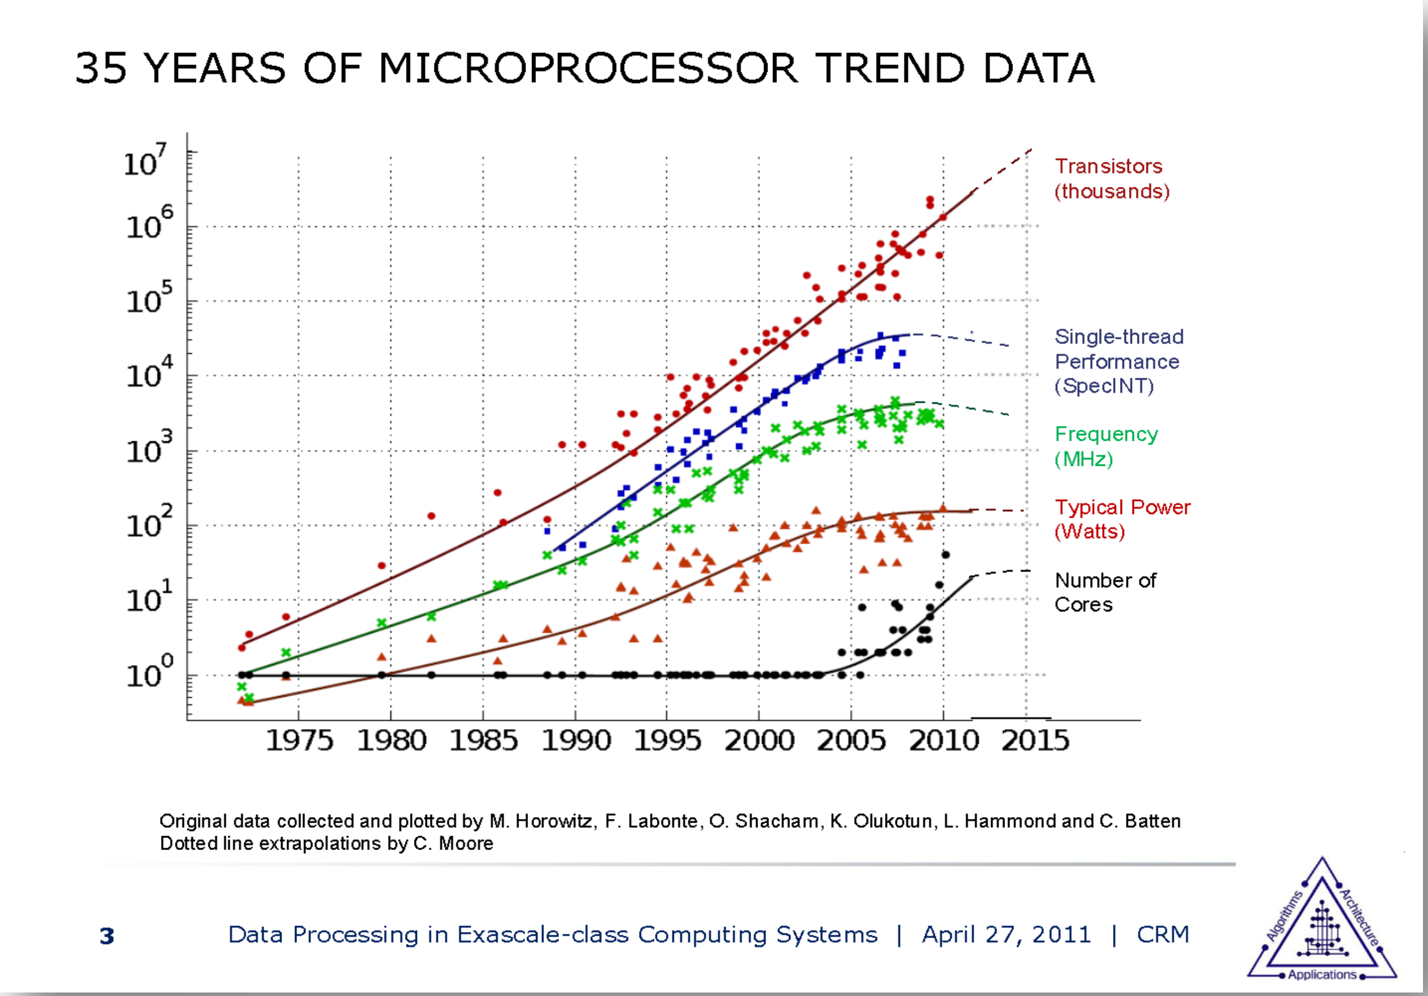

Moore's Law continues

- Transistor count still doubles every 24 months

Dennard scaling stalls – key parameters flatline:

- Voltage
- Clock Speed
- Power
- Performance/clock

Exascale Technology
Impact on Computing World

Terascale Embedded Systems

Field Deployable Systems

Cloud Computing Data Centers

Terascale Desktop Systems

Petascale Department Systems

Exascale Data Center Systems

Exascale Technology Will Have Significant Impact Across the Computer Industry and will Expand Use of HPC in S&T Organizations