U.S. DEPARTMENT OF ENERGY | Office of Science

DOE
Exascale Computing Initiative (ECI) Update

William Harrod
DOE, Office of Science (SC)
Aug 15, 2012

The World is Flat

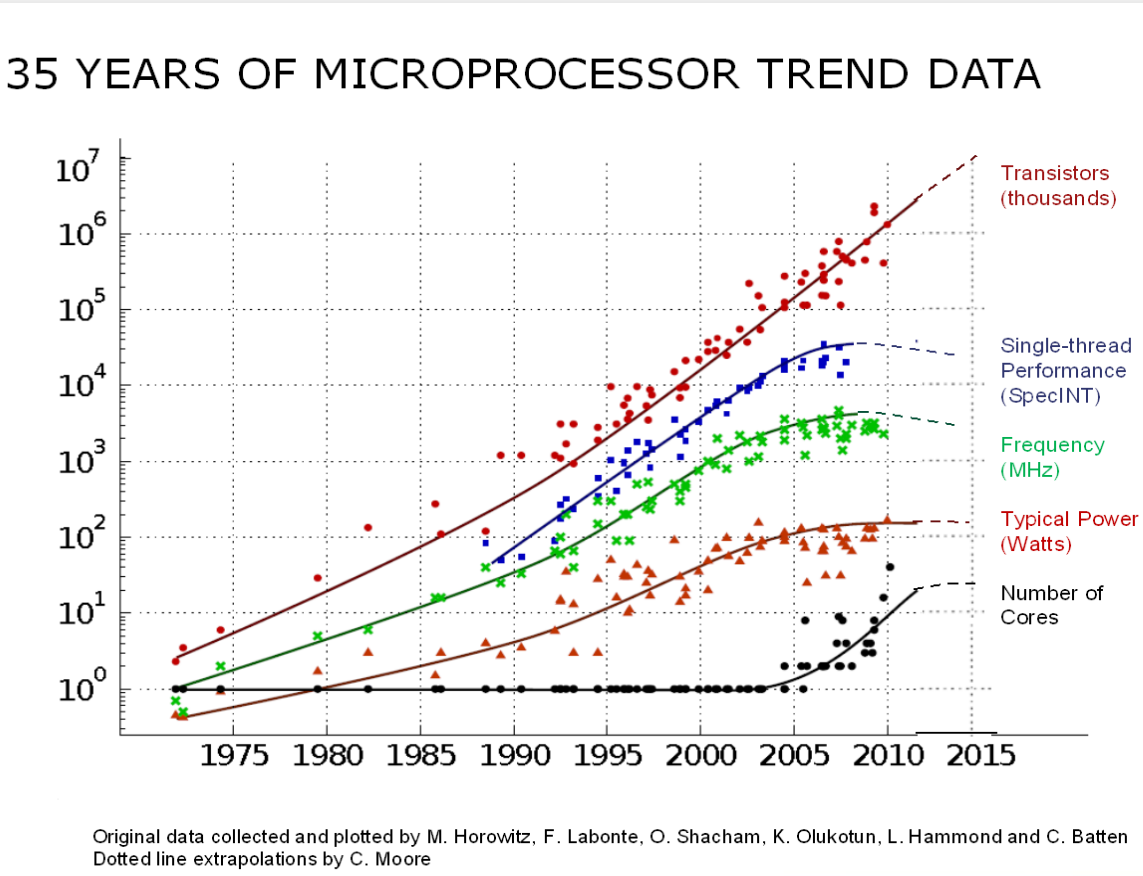

3 Data Processing in Exascale-class Computing Systems | April 27, 2011 | CRM

Moore's Law continues

- Transistor count still doubles every 24 months

Dennard scaling stalls – key parameters flatline:

- Voltage
- Clock Speed
- Power
- Performance/clock

HPC System Characteristics

(The Road to Exascale)

	ASCI Red	Road Runner	K Computer	Sequoia	Exascale
Peak (Tflops)	1.3	1,700	11,280	20,133	1,200,000
Linpack (Tflops)	1	1,000	10,510	16,325	1,000,000
Total Cores	9,298	130,464	705,024	1,572,864	1,000,000,000
Processors	9,298	12,960 +6,912	88,128	98,304	1,000,000
Cores/Processor	1	9, 2	8	16	1,000
Power	0.85 MW	2.35 MW	9.89 MW	7.9 MW	~20 MW
Year	2000	2008	2011	2012	2020+

Exascale Software Challenges

Parallelism

Data Movement

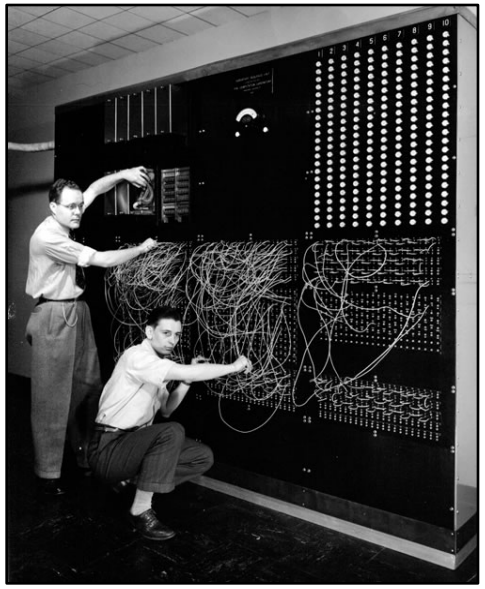

Programmability

Resiliency

ECI Program Strategy

- Conduct critical R&D efforts.
- Develop exascale software stacks.
- Fund computer technology vendors to move required technology from research to product space.
- Fund the design and development of exascale computer systems.
- Joint effort with NNSA.
- Collaboration with other United States government agencies and other countries.

FY12 Major Exascale Research Efforts

- **Runtime systems** to exploit unused/available resources through dynamic adaptive resource management and task scheduling (self aware)
- **Innovative operating system**, with a global perspective, for scalability, energy and reliability management, resource allocation and recovery, efficiency and protection
- **Architectures for cores and systems** to minimize latencies, preclude bottlenecks, reduce energy of data movement and control, IO systems, task instantiation and management, and address handling
- **Abstract machine models and performance models** to guide design and development of future Exascale machines

DOE/ASCR Progress Toward Exascale

- **FY2010:**
  - **Applied Math:** Uncertainty Quantification
  - **Computer Science:**
    - Advanced Architectures / X-Stack / Scientific Data Management and Analysis
- **FY2011:**
  - **Computational Partnerships:** 3 Exascale Co-Design Centers Funded
  - **Request for Information:** critical and platform technologies
- **FY2012:**
  - **Computer Science:** Programming Environments
  - **Applied Math:** Resilient, Extreme scale algorithms
  - **FastFoward**: Critical / Cross Cutting technologies (joint with NNSA)
- **FY2013**
  - **Exascale Strategy Plan to Congress**
  - **Exascale Software Plan**
  - **Computer Science**
    - Operating System Software / Storage System Software / ?
  - **Computational Partnerships:** Data Intensive Co-Design Centers

The OS Technical Council

Council: comprised of DOE lab researchers

Goal:

- Investigate issues, challenges and solutions for the DOE software plan for Exascale OS and runtime
- Develop conceptual OS & Runtime architectures and APIs

OS Software Preliminary Plan: November, 2012

Council Meetings:

- March 21, 2012: general council strategies, architectures, APIs
- May 14, 2012: vendor and facilities engagement strategies
- June 11-12, 2012: vendor engagement decisions, workshop planning, architectures, APIs
- July 19-21, 2012: meeting with vendors, workshop planning
- August 21-22, 2012: facilities and applications engagement planning
- September 12-14, 2012: vendor, facilities, and applications engagement final plan, architectures, APIs
- October 4-5, 2012: OS and Runtime Workshop

RX-Solvers

FOA-0000742: http://science.doe.gov/grants/pdf/SC_FOA_0000742.pdf

Resilient Extreme-Scale Solvers ("RX-Solvers")

- Scalable, resilient algorithms
- Support applications during the next 5-10 years
- Establish foundation for research of extreme-scale scientific computing."

Programmatic Details

- FOA Issued: 8 June 2012; closed: 13 August 2012
- Funding to start in FY13
- Total funds available: $4.5M per year for up to three years
- Proposals are being received; initial reviews underway

X-Stack: Programming Challenges, Runtime Systems, and Tools

FOA 0000619: http://science.doe.gov/grants/pdf/SC_FOA_0000619.pdf

Focus Areas

- Programming models, languages, runtime systems, and related technologies
- New energy-efficient and resilient programming techniques that are portable across future machine generations

Expected research investments

- That address fundamental Exascale challenges, while offering a transition path for existing scientific applications to fully explore the challenges and rewards of Exascale platforms.

Programmatic Details

- **Nov 22, 2011** **Solicitation Issued**
- Dec 20, 2011 Amendment FOA
- Dec 21, 2011 Pre-applications received
- Feb 6, 2012 Full proposals received
- **Apr 3-6, 2012** **Peer Review**
- May 29 - 30, 2012 Call Back Review
- **June 8, 2012** **Funding recommendations completed**

X-Stack Proposals

Proposals were expected to:

- Articulate complete solutions addressing multiple components of the system software stack and address Exascale challenges:
  - Scalability
  - Programmability
  - Performance Portability
  - Resilience
  - Energy Efficiency

Proposal were required to have:

- Description of plans for developing prototypes of the proposed solution;
- Description of the proposed path to integration and/or interoperation with existing programming environments;
- Evaluation plan using compact applications, mini-applications

FastForward

Project Goals & Objectives

- Initiate partnerships with multiple companies to accelerate the R&D of critical technologies needed for extreme-scale computing.
- DOE applications place extreme requirements on computations, data movement, and reliability.
- Fund innovative new and/or accelerated R&D of technologies targeted for productization in the 5–10 year timeframe.

Funded Projects

- AMD
- IBM
- Intel
- NVIDIA
- WhamCloud

For additional Details: Teri Quinn, LLNL, ASCAC presentation

ECI Timeline

(*actual lengths of phases could be longer*)

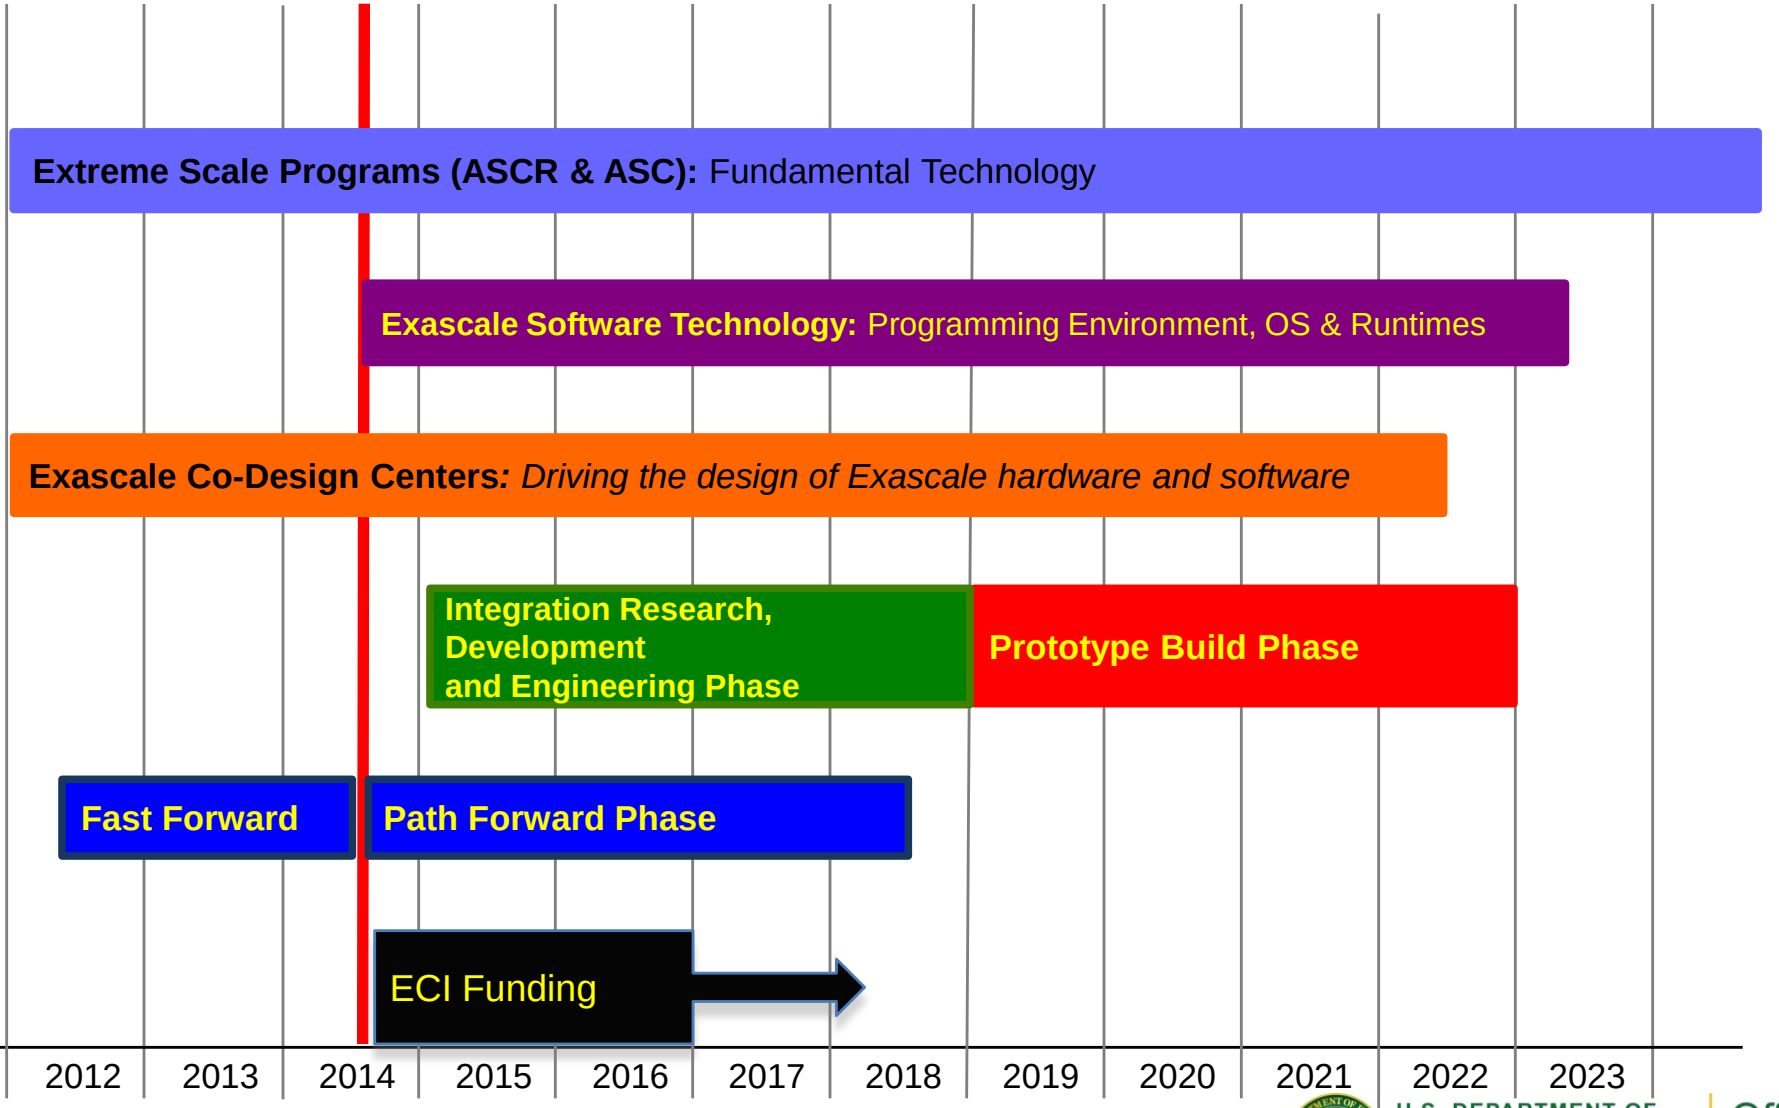

The Real Challenges

Avoiding mediocre solutions

- Evolving existing systems

Practicing Co-Design

Developing a new software stack for exascale systems

- Not treating it as an “after thought”

Need to explore radical concepts, but develop practical solutions

New computers designs based on a new execution model

- Must be based on COTS technology
- Exotic technology is not an option

Backup

ECI Timeline

(*actual lengths of phases could be longer*)