A Path to Capable Exascale Computing

July 31, 2016

Paul Messina, Argonne National Laboratory, ECP Director

EXASCALE COMPUTING PROJECT

National Nuclear Security Administration

U.S. DEPARTMENT OF ENERGY | Office of Science

Why Exascale instead of Exaflops?

- Peak speed of computers has traditionally been measured by the maximum number of floating point operations that can be performed in a second (FLOPS), or peak flops
  - Exaflops is 10^18 flops = 1,000 Petaflops
  - But on current computer architectures most applications can achieve only a modest fraction of the peak speed
- Therefore we are transitioning to a less precise but more useful measurement: Exascale, which is based on the ratio of execution time of full applications (not simple benchmarks) on today's petaflops systems (e.g., Mira 10 PF) and future systems
- So if a system can run a real app 100X faster than it runs on Mira, we would say it is an exascale system

The National Strategic Computing Initiative* and DOE's role

Through Executive Order 13702 he signed July 29, 2015, President Obama established the National Strategic Computing Initiative (NSCI) to maximize the benefits of HPC for US economic competitiveness and scientific discovery.

DOE is a lead agency within NSCI with the responsibility that the DOE Office of Science and DOE National Nuclear Security Administration will execute a joint program focused on advanced simulation through a **capable** exascale computing program emphasizing sustained performance on relevant applications.

* https://www.whitehouse.gov/the-press-office/2015/07/29/executive-order-creating-national-strategic-computing-initiative

Approach to executing that DOE role in NSCI

- Starting this year, the Exascale Computing Project (ECP) was initiated as a DOE-SC/NNSA-ASC partnership, using DOE's formal project management processes
- The ECP is a ten-year project led by DOE laboratories and executed in collaboration with academia and industry
- The ECP leadership team has staff from six U.S. DOE labs
  – Staff from most of the 17 DOE national laboratories will take part in the project
- The ECP will collaborate with the facilities that operate DOE's most powerful computers

Why is a project needed?

- We have reached technology and physics limits that mean it is not business as usual
  - In fact, it has not been business as usual since 2004
- The holistic approach of advancing all aspects of the HPC ecosystem requires the structure of a project
  - Co-design
  - Integrated planning and execution
  - Milestones
  - Joint milestones

Progress in CMOS CPU Technology

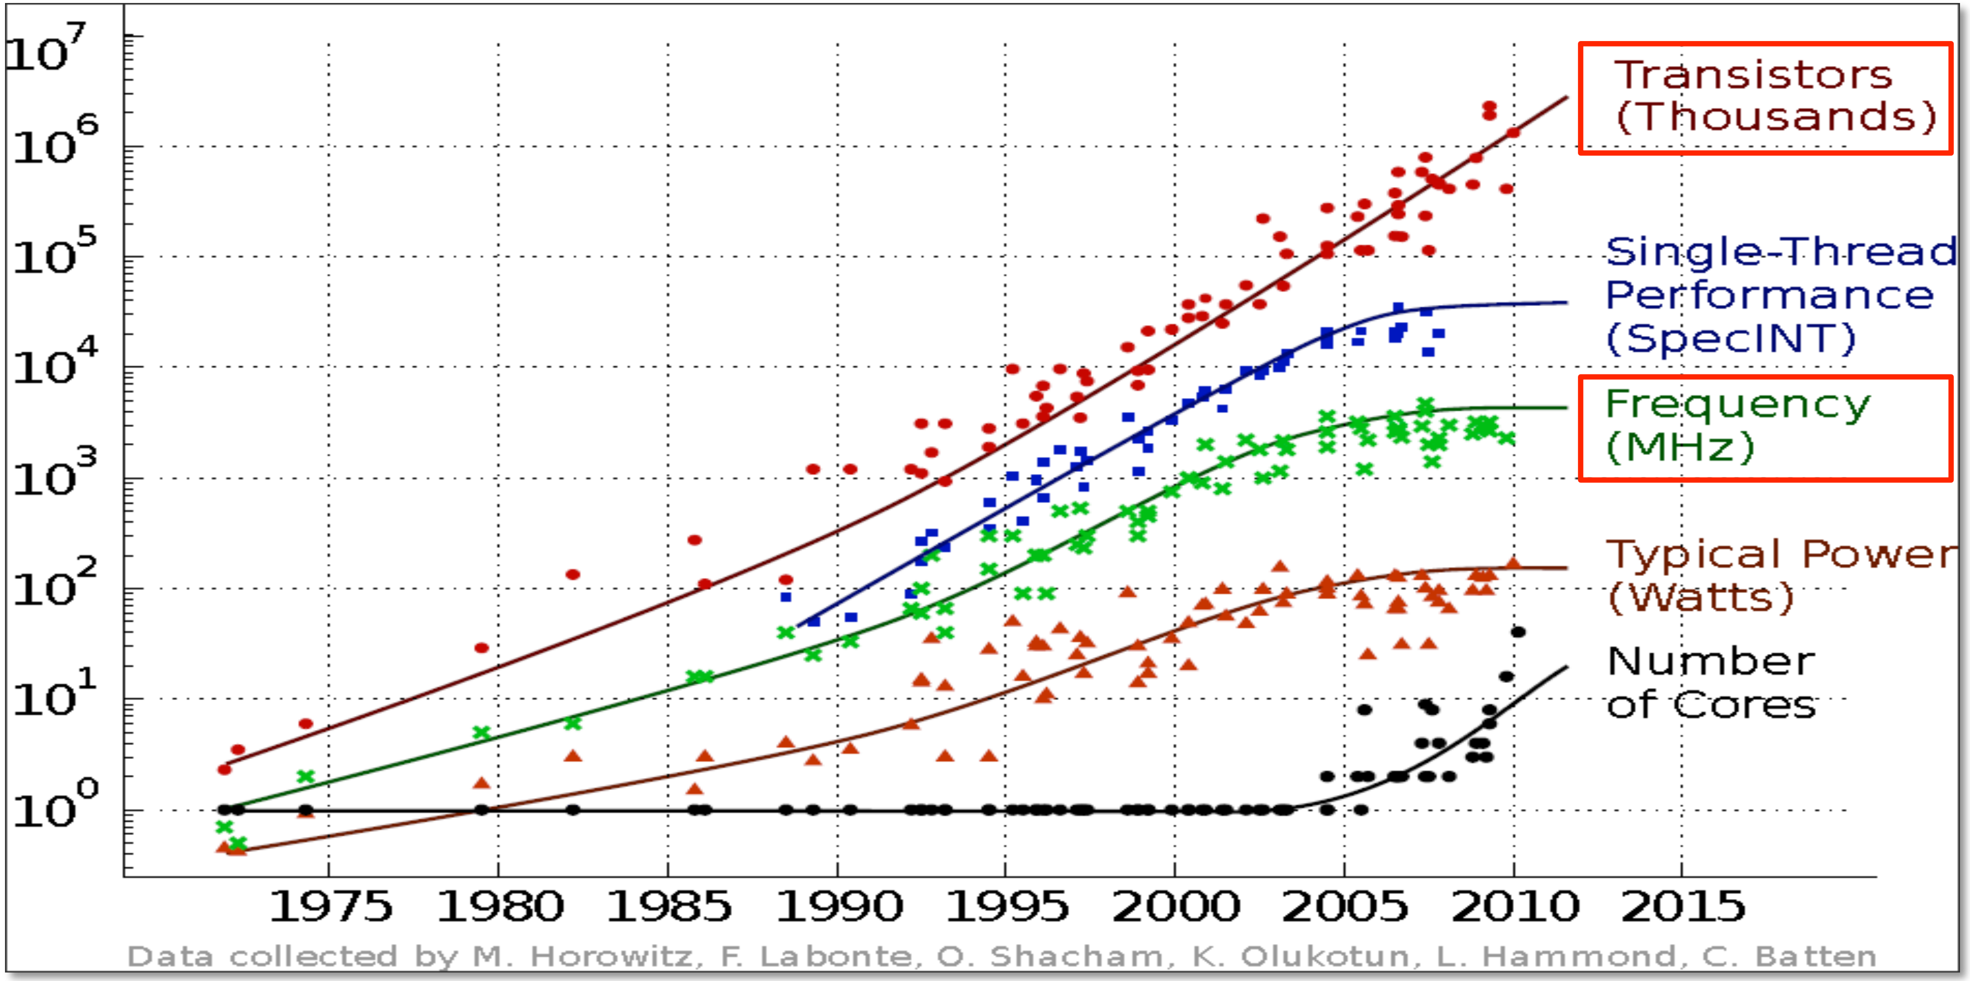

Moore's Law continues

- Transistor count still doubles every ~24 months

Dennard scaling stalls

- Voltage
- Clock Speed
- Power
- Performance/clock

From Giga to Exa, via Tera & Peta*

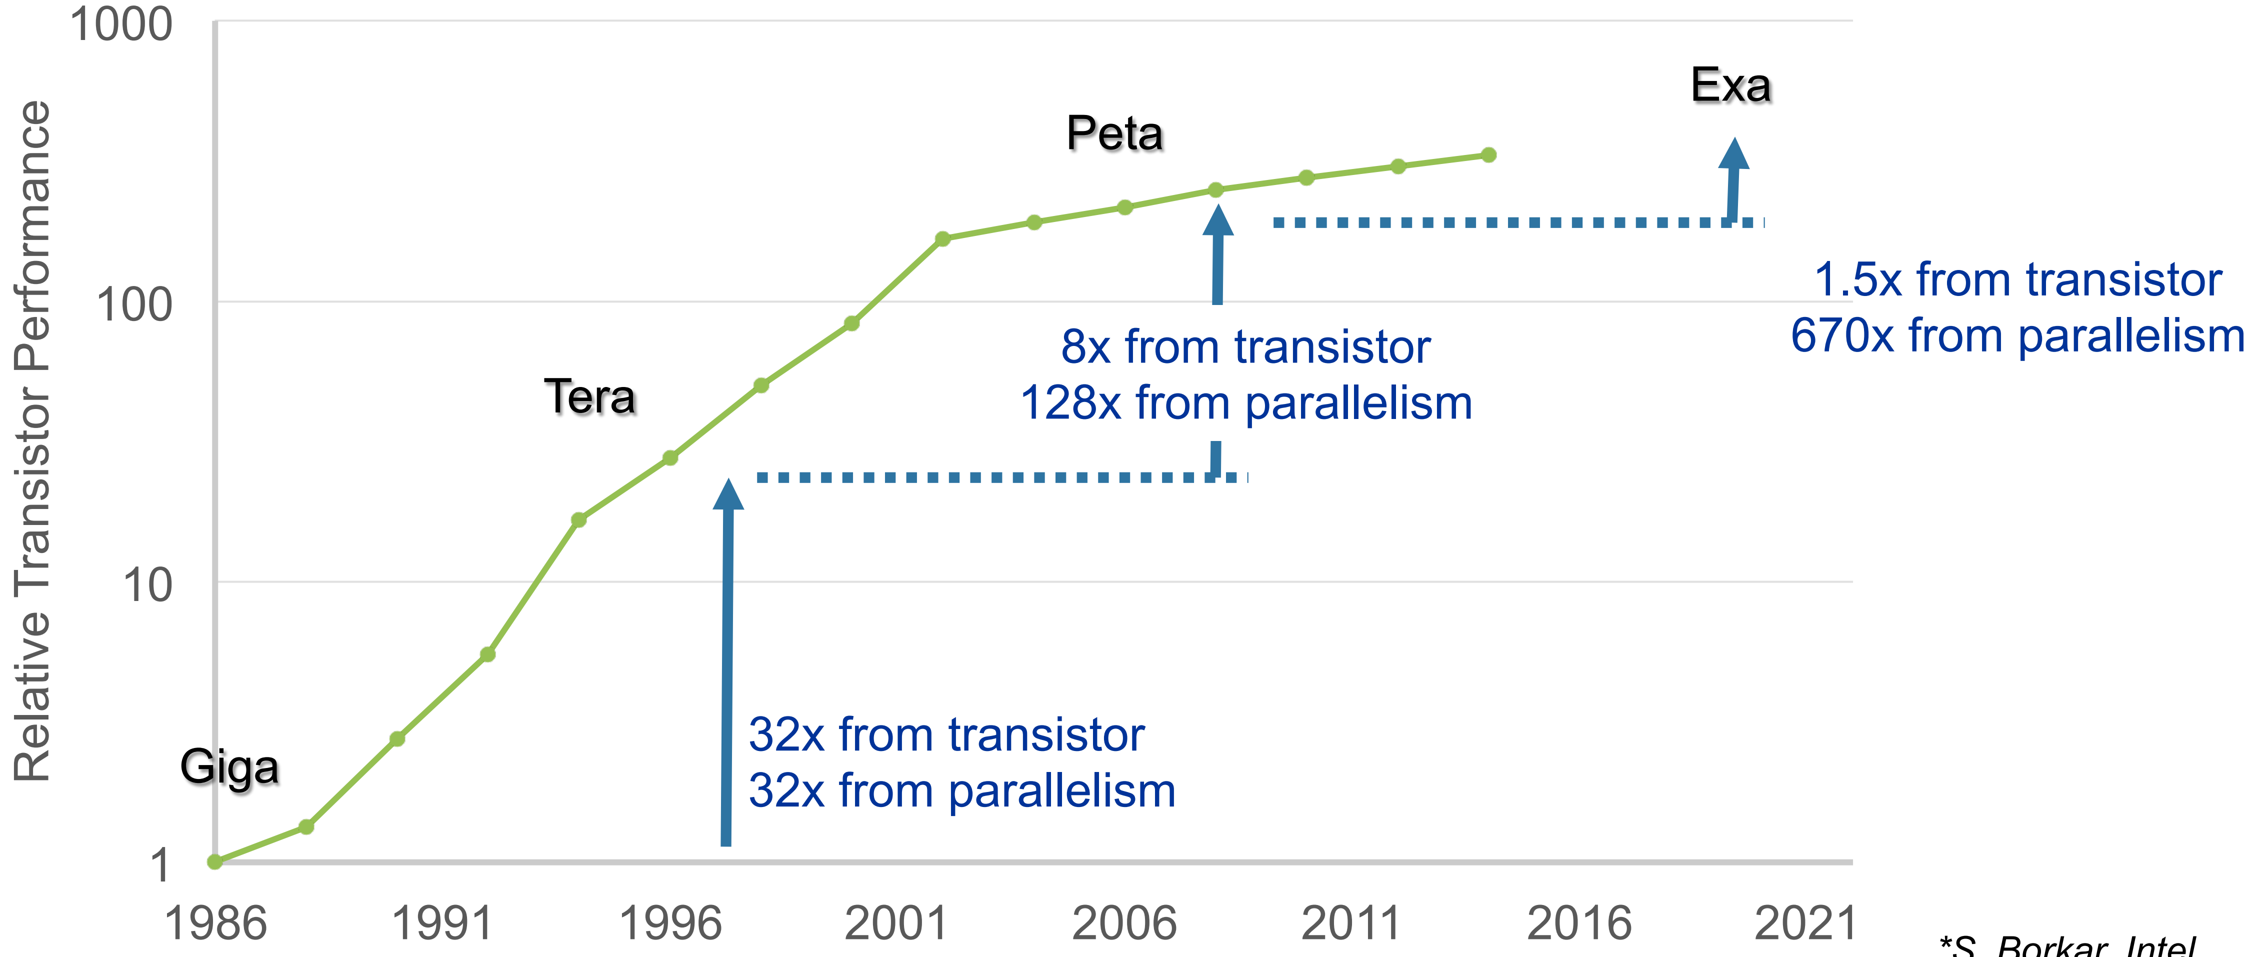

*S. Borkar, Intel

Performance from parallelism

Top Ten Technical Challenges for Exascale
(Advanced Scientific Computing Advisory Committee Findings)

1. **Energy efficiency**: Creating more energy efficient circuit, power, and cooling technologies.
2. **Interconnect technology**: Increasing the performance and energy efficiency of data movement.
3. **Memory technology**: Integrating advanced memory technologies to improve both capacity and bandwidth.
4. **Scalable System Software**: Developing scalable system software that is power and resilience aware.
5. **Programming systems**: Inventing new programming environments that express massive parallelism, data locality, and resilience
6. **Data management**: Creating data management software that can handle the volume, velocity and diversity of data that is anticipated.
7. **Exascale algorithms**: Reformulating science problems and refactoring their solution algorithms for exascale systems.
8. **Algorithms for discovery, design, and decision**: Facilitating mathematical optimization and uncertainty quantification for exascale discovery, design, and decision making.
9. **Resilience and correctness**: Ensuring correct scientific computation in face of faults, reproducibility, and algorithm verification challenges.
10. **Scientific productivity**: Increasing the productivity of computational scientists with new software engineering tools and environments.

http://science.energy.gov/~/media/ascr/ascac/pdf/reports/2013/report.pdf

Our definition of "*capable* exascale computer"

- A capable exascale system is defined as a supercomputer that
  - can solve science problems 50X faster (or more complex) than on the 20PF systems (Titan, Sequoia) of today or 100X faster than on Mira
  - in a power envelope of 20-30 MW
  - is sufficiently resilient that user intervention due to hardware or system faults is on the order of a week on average, and
  - has a software stack that meets the needs of a broad spectrum of applications and workloads.

To achieve capable exascale requires the integration of multiple complex elements working together

Application Development	Software Technology	Hardware Technology	Exascale Systems
Science and mission applications	Scalable software stack	Hardware technology elements	Integrated exascale supercomputers

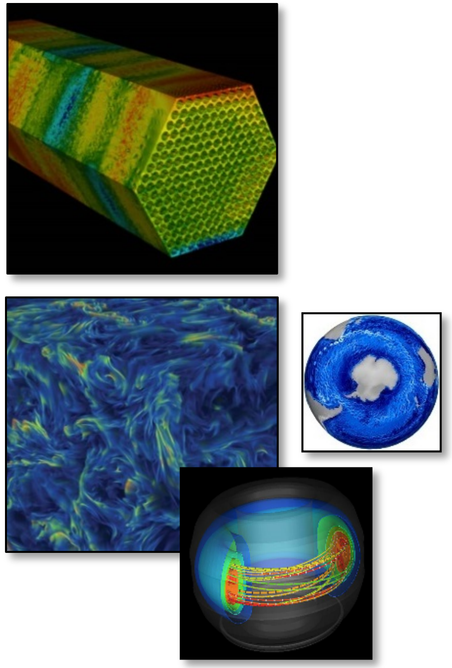

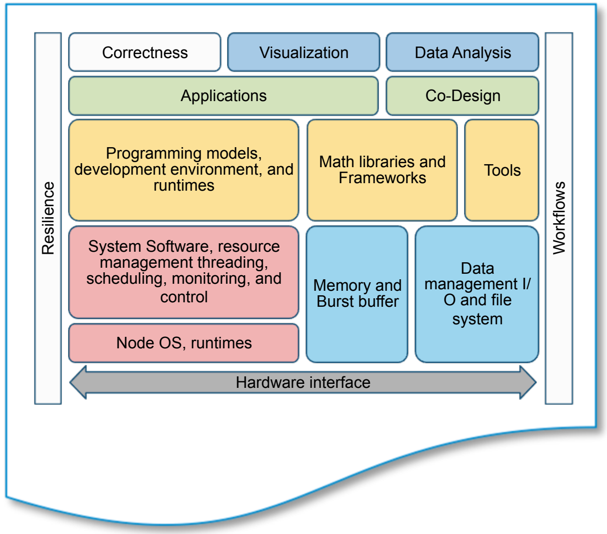

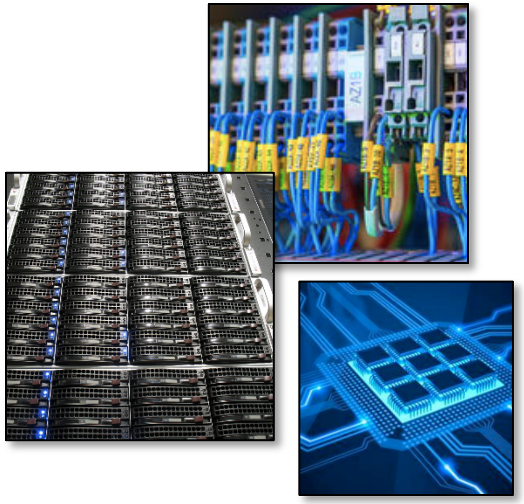

ECP EXASCALE COMPUTING PROJECT

U.S. Strategic Advantage: NSCI Critical Applications

NSCI Vision: Supported by a national tech base, integrated capabilities, R&D foundation

National Security	Economic Competitiveness	Scientific Discovery
• Stockpile Stewardship • Decision Support • Battlefield Command • Counter-Terrorism • Secure Communication • Cyber Defense • Signals Intelligence	• Energy Production • Advanced Manufacturing • Digital Engineering • Drug Design • Personalized Medicine • Health Care • Business Analytics • Financial Services • E-Commerce • Social Networking	• Climate Science • Fusion Science • Materials Genome • Particle Physics • Neuroscience • Weather Prediction • Genomic Discovery

Applications Development activities

- Fund applications development teams
  - Each aiming at capability and specific challenge problems
  - Following software engineering practices
  - Tasked to provide software and hardware requirements
  - Execute milestones jointly with software activities
- Establish co-design centers for commonly used methods
    - E.g., Adaptive Mesh Refinement, Particle-in-Cell
- Developer training

The central role of software

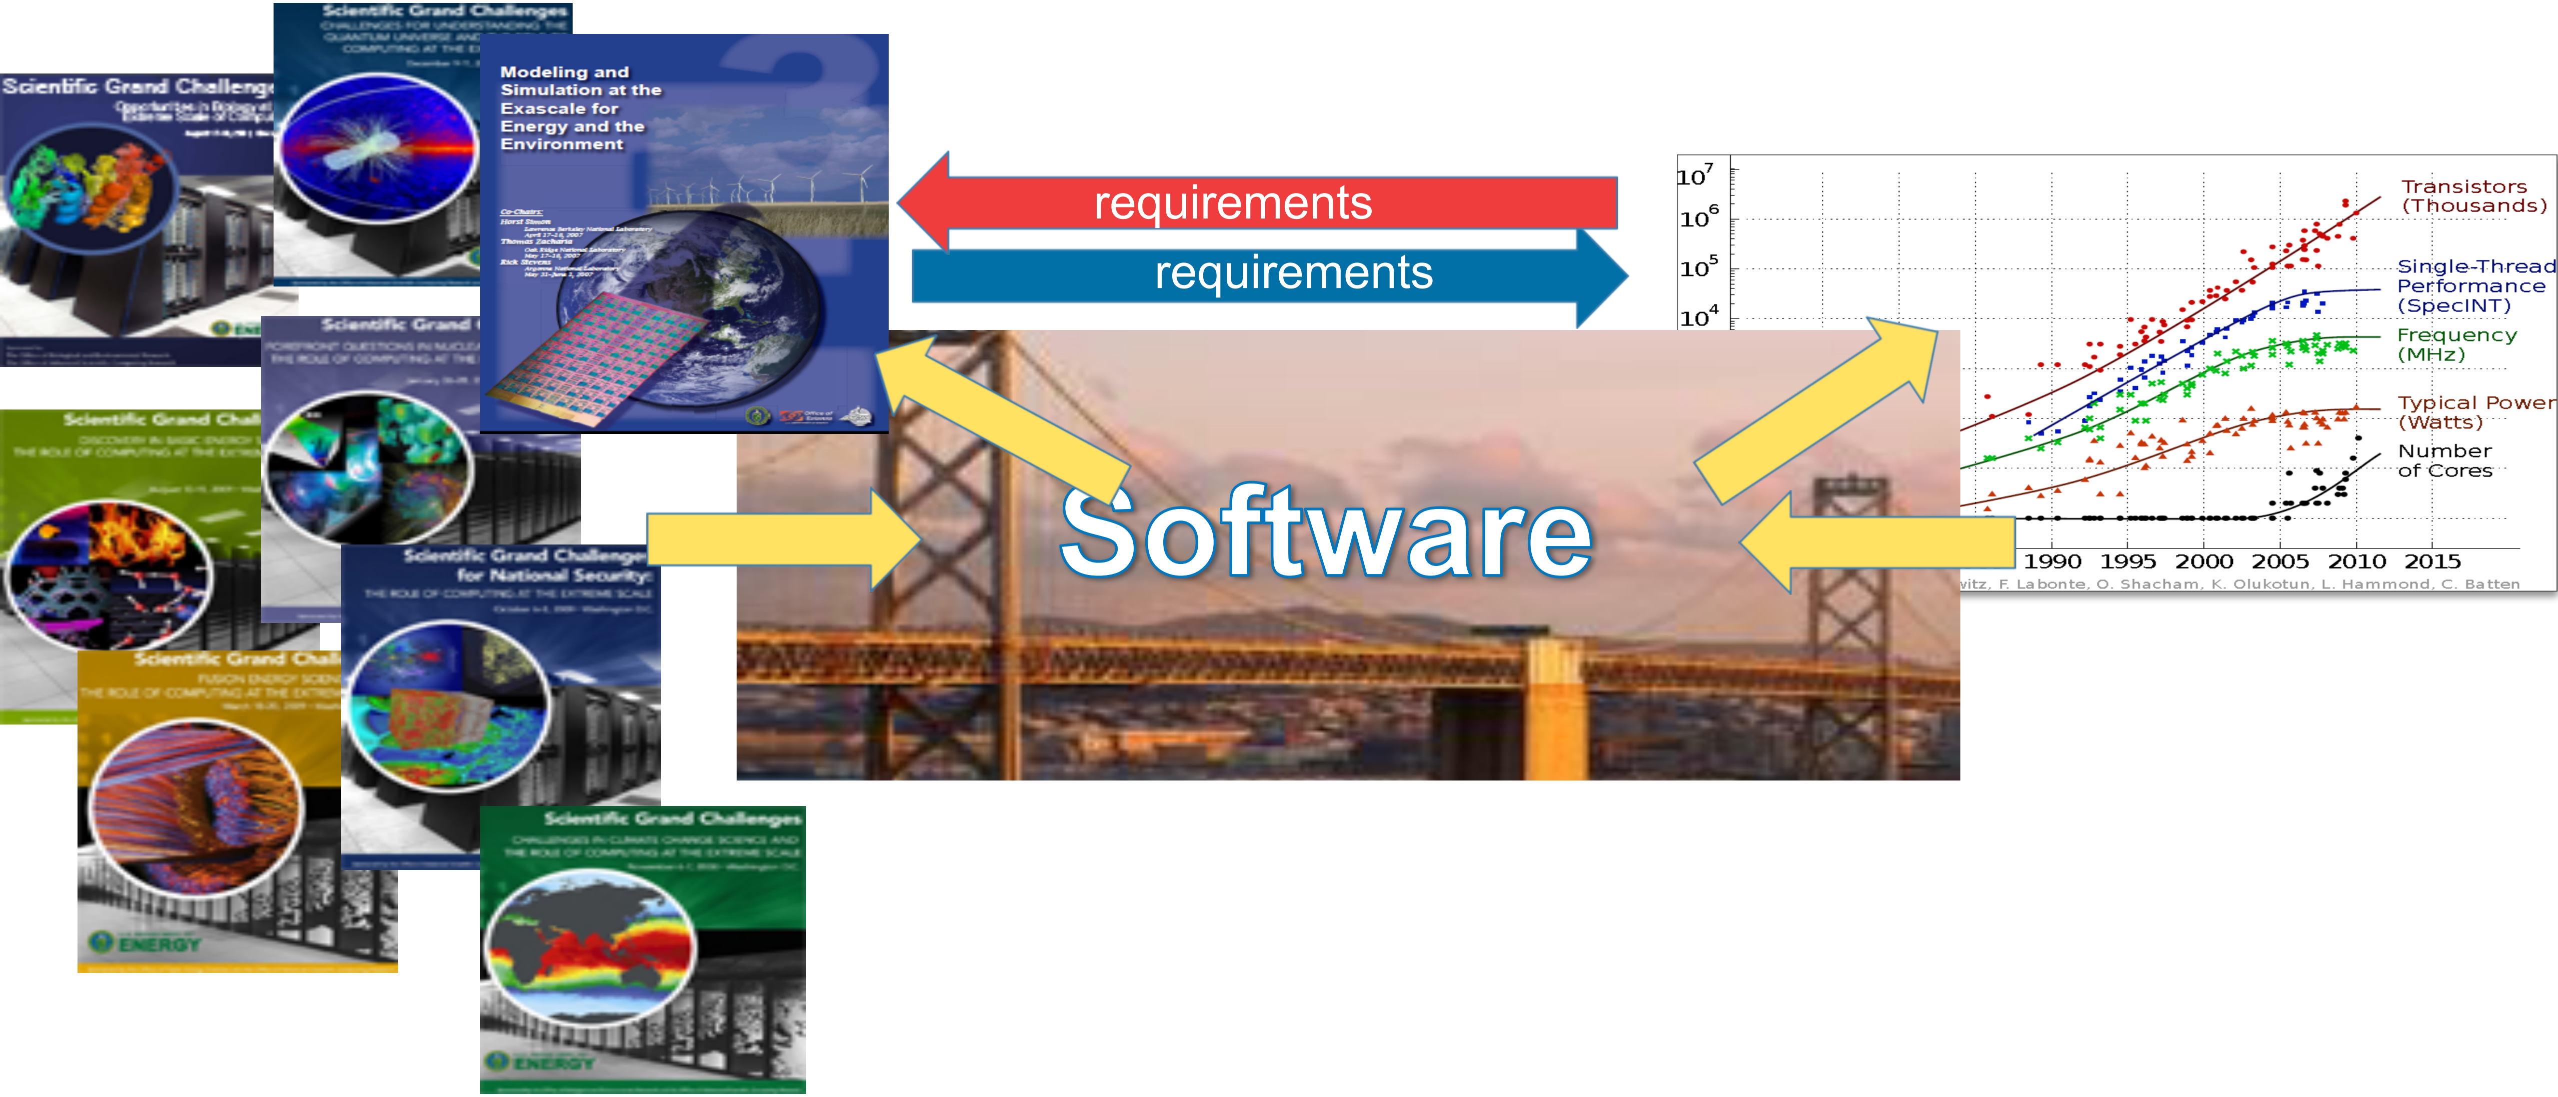

EXASCALE COMPUTING PROJECT

Conceptual ECP Software Stack

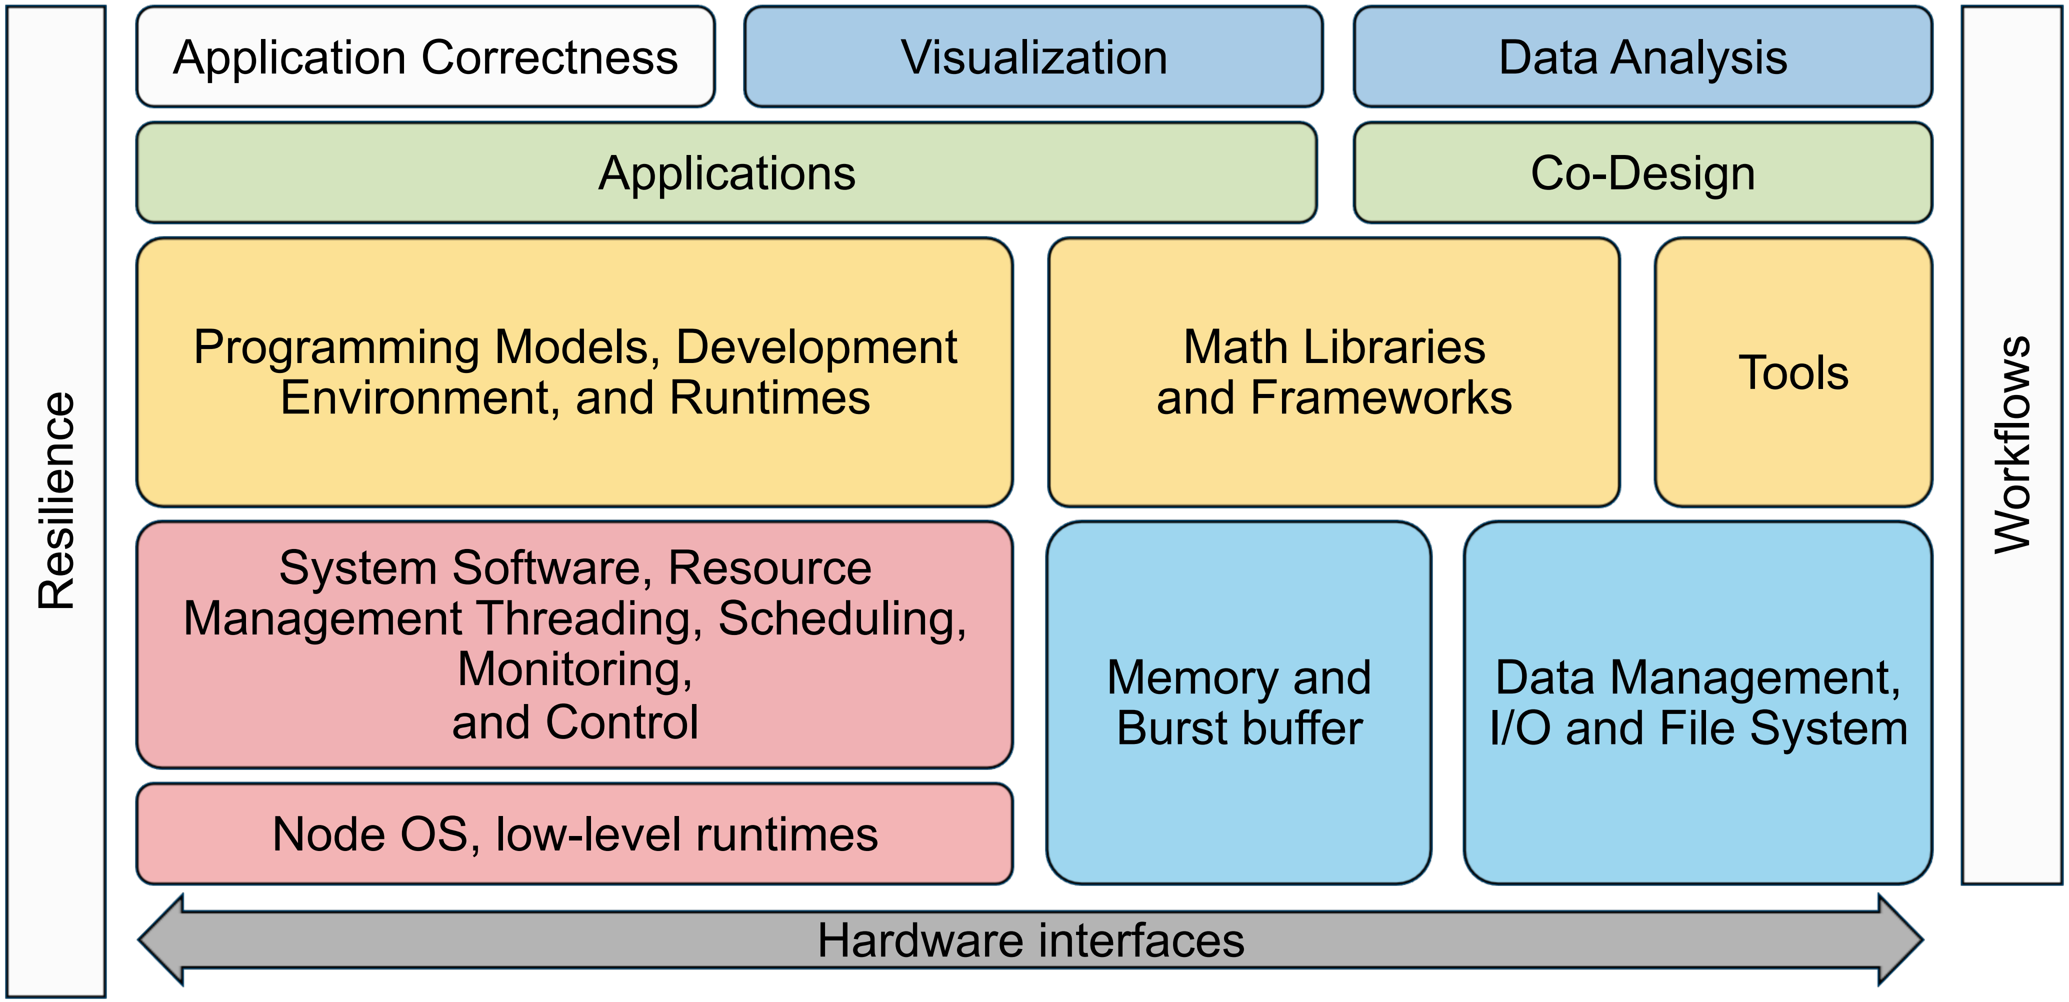

Addresses Top Ten Challenges 4, 5, 6, 7, 8, and 10

EXASCALE COMPUTING PROJECT

Bay Bridge, San Francisco Bay

Old span on left, new on right

Hardware Technology Activities

- PathForward: support DOE-vendor collaborative R&D activities required to develop exascale systems with at least two diverse architectural features; quote from RFP:
  - PathForward seeks solutions that will improve application performance and developer productivity while maximizing energy efficiency and reliability of an exascale system.
  - Aimed at addressing Top Ten Challenges 1, 2, 3, and 9
- Design Space Evaluation
  - Apply laboratory architectural analysis capabilities and Abstract Machine Models to PathForward designs to support ECP co-design interactions

ECP WBS

- **Exascale Computing Project** 1.
  - Project Management 1.1
    - Project Planning and Management 1.1.1
    - Project Controls & Risk Management 1.1.2
    - Business Management 1.1.3
    - Procurement Management 1.1.4
    - Information Technology and Quality Management 1.1.5
    - Communications & Outreach 1.1.6
    - Integration 1.1.7
  - Application Development 1.2
    - DOE Science and Energy Apps 1.2.1
    - DOE NNSA Applications 1.2.2
    - Other Agency Applications 1.2.3
    - Developer Training and Productivity 1.2.4
    - Co-Design and Integration 1.2.5
  - Software Technology 1.3
    - Programming Models and Runtimes 1.3.1
    - Tools 1.3.2
    - Mathematical and Scientific Libraries and Frameworks 1.3.3
    - Data Management and Workflows 1.3.4
    - Data Analytics and Visualization 1.3.5
    - System Software 1.3.6
    - Resilience and Integrity 1.3.7
    - Co-Design and Integration 1.3.8
  - Hardware Technology 1.4
    - PathForward Vendor Node and System Design 1.4.1
    - Design Space Evaluation 1.4.2
    - Co-Design and Integration 1.4.3
  - Exascale Systems 1.5
    - Site Preparation 1.5.1
    - NRE 1.5.2
    - Testbeds 1.5.3
    - Base System Expansion 1.5.4
    - Co-Design and Integration 1.5.5

ATPESC

EXASCALE COMPUTING PROJECT

Summary

- The DOE Exascale Computing Project aims to advance to exascale
  - a broad class of applications,
  - the software ecosystem,
  - hardware design, and
  - the facilities that will operate the initial exascale systems

using a holistic approach with a strong focus on co-design and integration

- The work will be carried out by the DOE national laboratories, industry, and academia
  - With involvement by other U.S. government agencies

Closing remarks

- Risk 2001: Unable to recruit and/or retain qualified staff needed to execute R&D
- You might someday work for the ECP project
  - It is an exciting project that I believe will have positive impacts on many aspects of society
  - Its broad scope makes it likely you will find topics you would like to work on
- Tomorrow Doug Kothe will give a talk on the ultimate payoff of ECP: applications
  - That is why I did not give much detail on applications

Questions?