The Bump in the Road to Exaflops and Rethinking LINPACK

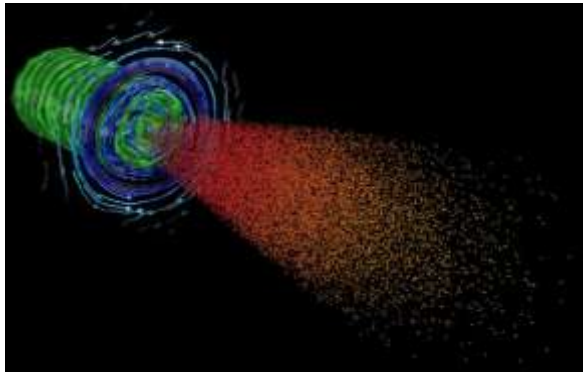

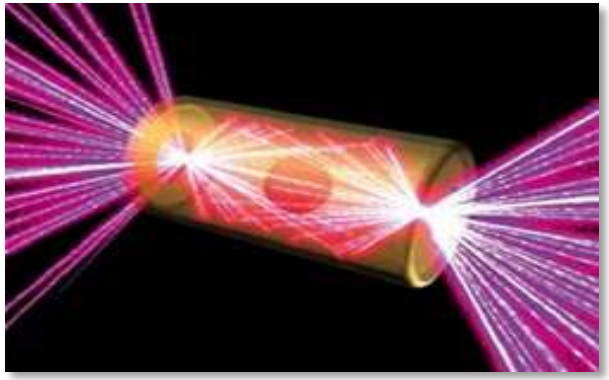

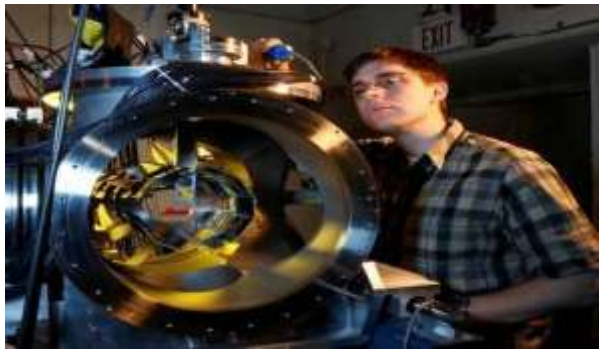

Bob Meisner, Director

Office of Advanced Simulation and Computing

Theme

- Actively preparing for imminent profound shift in computing architectures by making computing investments NOW in H/W and S/W
- We can't wait for an initiative to save us
- BUT, an initiative would enable a comprehensive approach to building exascale system
- Exascale isn't about an exaflop, but about how effectively we transition to a new era of computing
- A better corrolated benchmark would help

Episodic disruption defines high-end computing

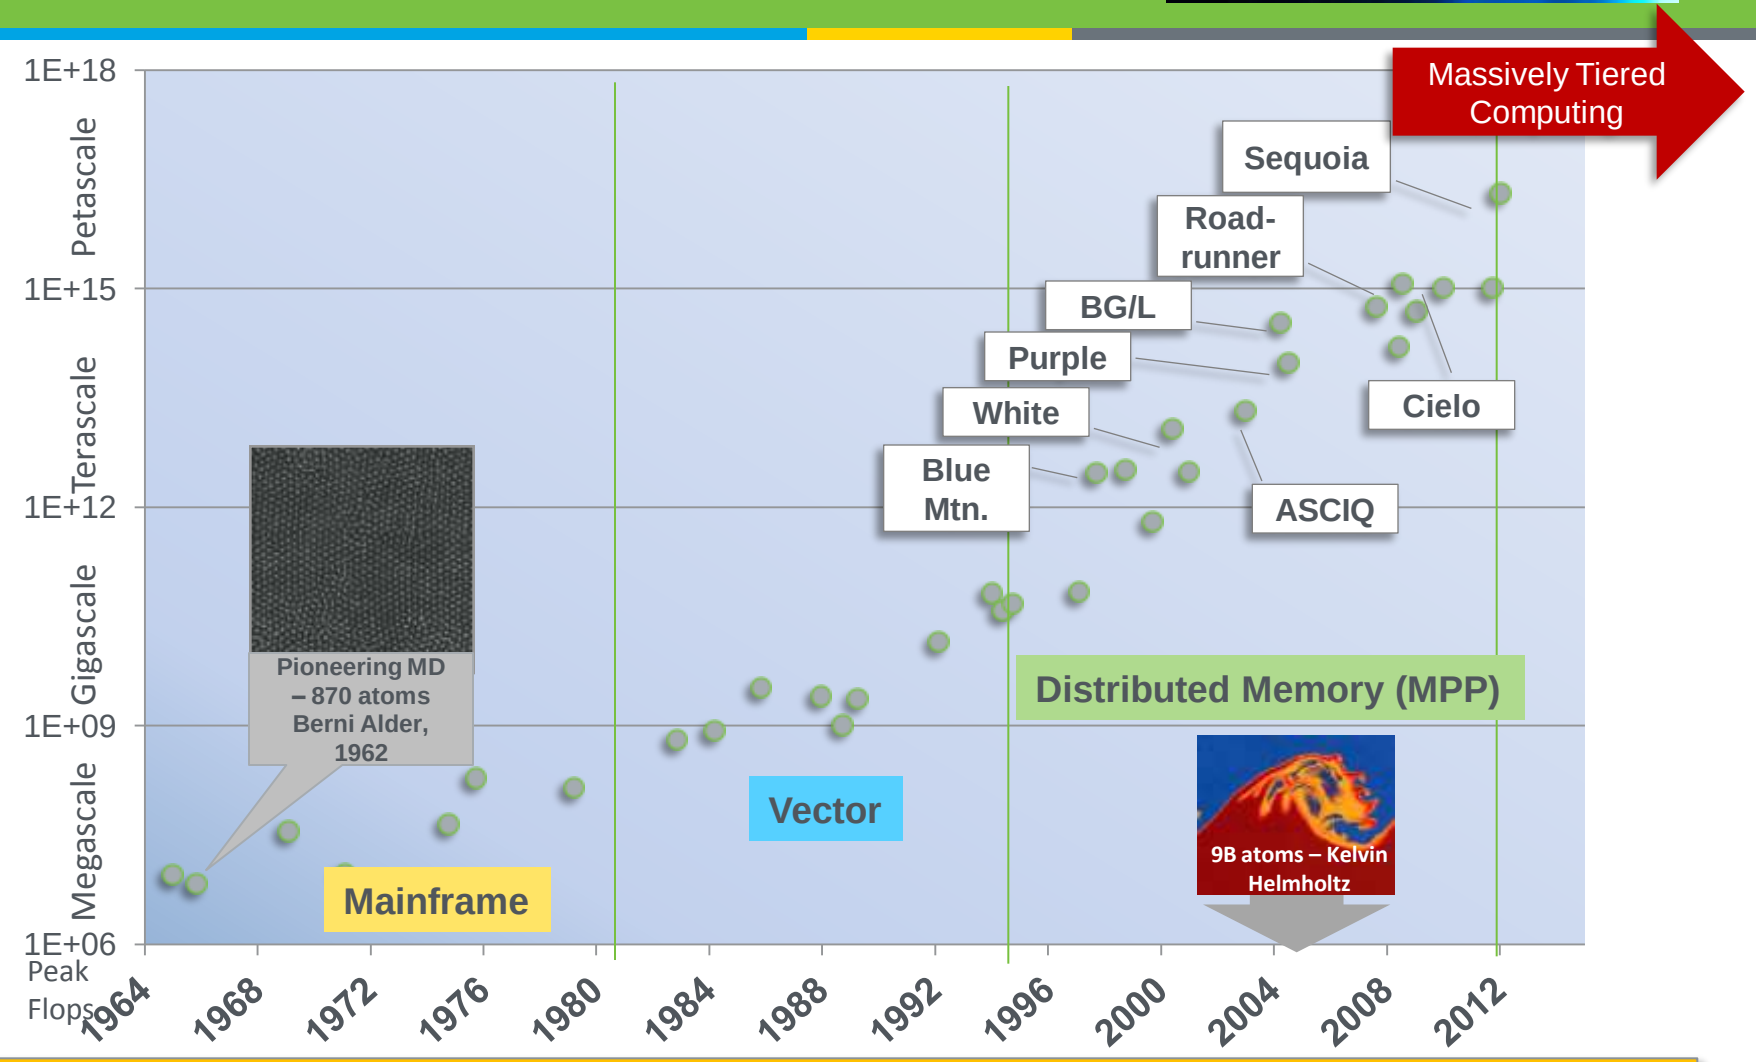

Architectural stability has made possible remarkable advances in science. But, programming model transitions are tough…and we are approaching one now….

The next disruption is NOW

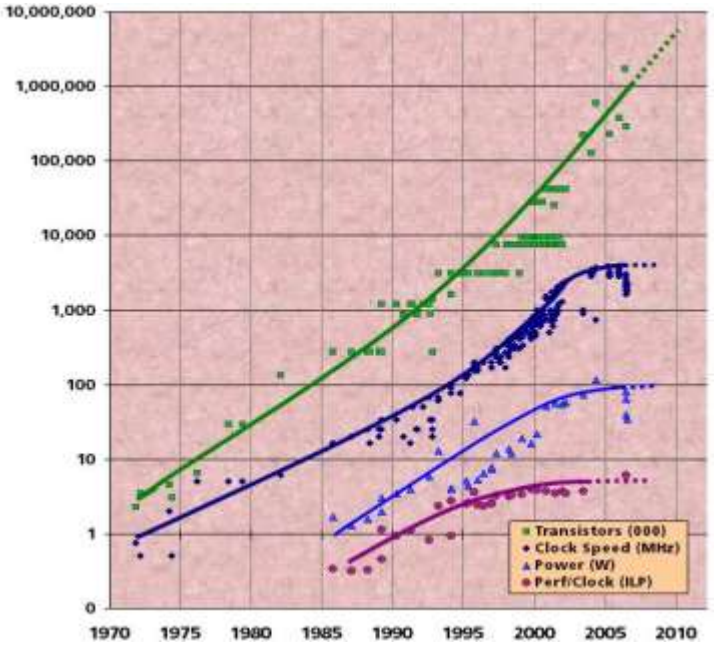

Figure courtesy of Kunle Olukotun, Lance Hammond, Herb Sutter, and Burton Smith, 2004

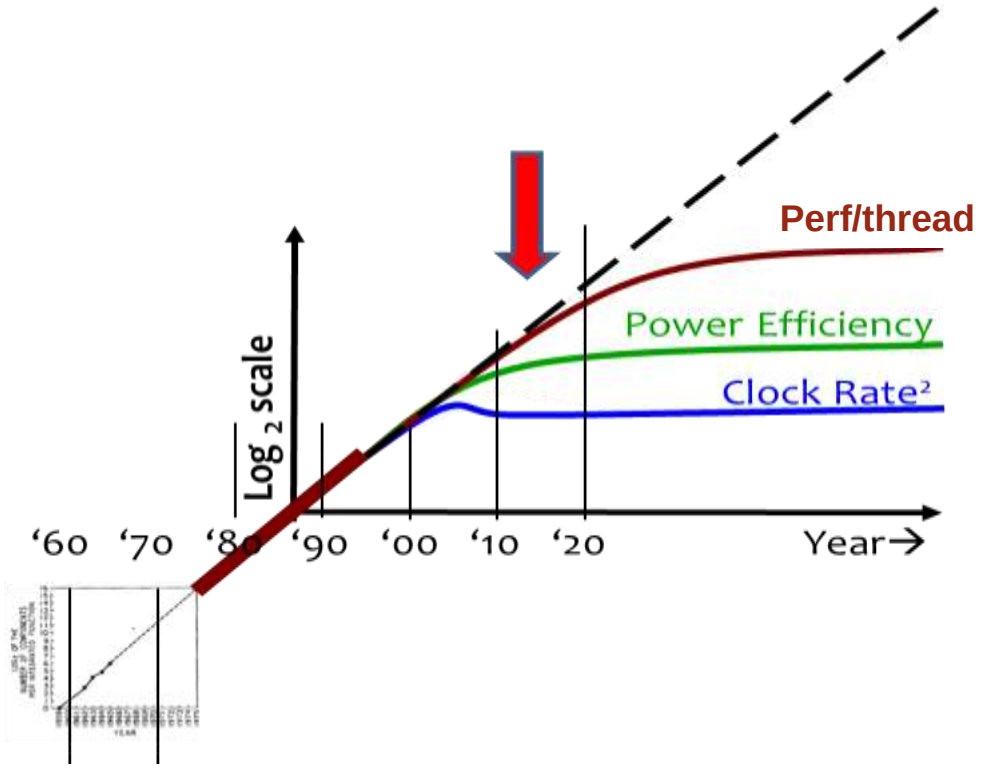

New Epoch is forcing us to address issues in several broad areas:

- Exponentially growing parallelism
- Data movement management
- System complexity
- Application code evolution

Continuing to advance computational science will require mastering architectural complexity

- Resolution increases have led to critical scientific insights and further increases are necessary for continued progress
- Science requirements drive the need for higher performance computers, while computational progress depends on successfully transitioning to complex architectures

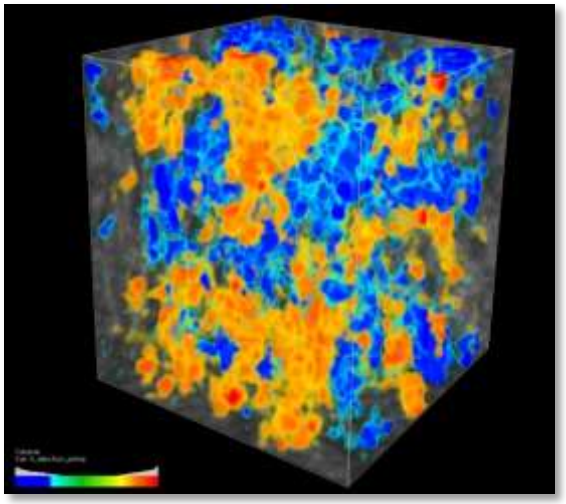

MD plasma simulation

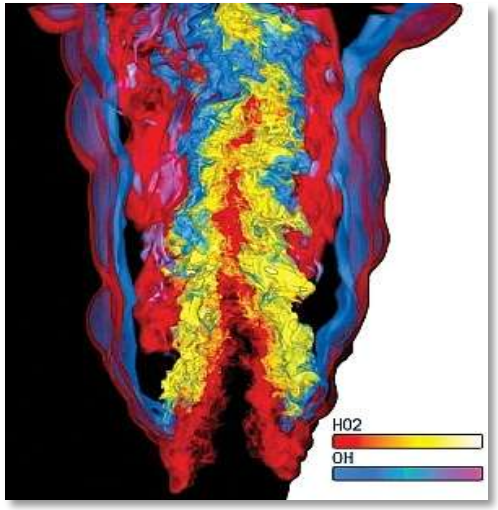

Combustion simulation

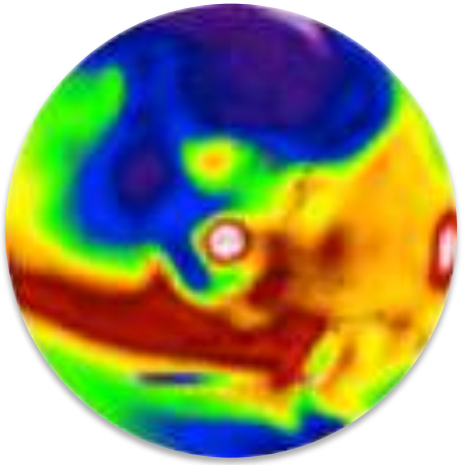

Global climate simulation

We are in a new era of computing and need to quickly adapt our codes

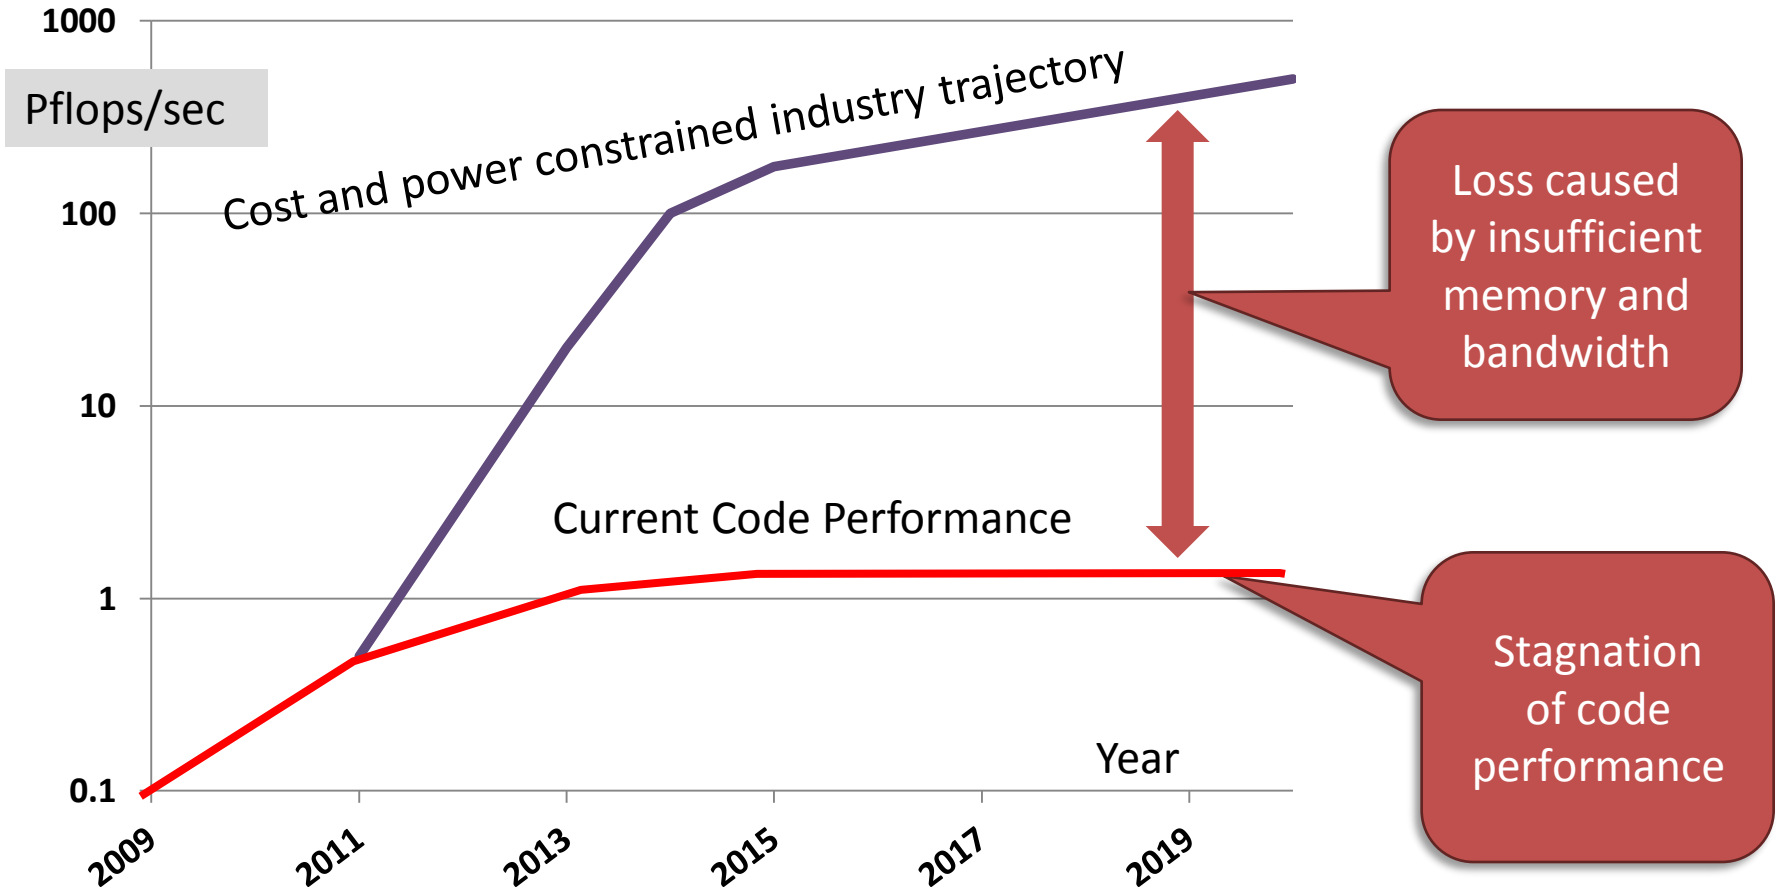

Unless we take action, our future will be keeping performance from deteriorating rather than improving

An exascale initiative may be our long term salvation, but we need a short term life jacket as well

- Scientific simulations must be ready for new architectures
- To prepare for the dramatic, impending changes we are pursuing :
  - Partnerships with industry to develop advanced processor, memory and interconnect technologies
  - Investments in software environments and application codes
  - Non-Recurring Engineering (NRE) investments
- We are investigating a new metric to confirm the performance of high-end computers

Promoting industry innovation through codesign

- Formed partnerships with multiple companies to accelerate the R&D of critical technologies needed for extreme-scale computing

- Targeted innovative new and/or accelerated R&D of technologies for productization in the 5–10 year timeframe
  - $25.4 million focusing on interconnect architectures and implementation approaches
  - $62.5M focusing on processor/memory and storage

- Future investments planned

- NRE also critical to move vendors in suitable directions

Entering a new episode in HPC-- Rethinking the community benchmark

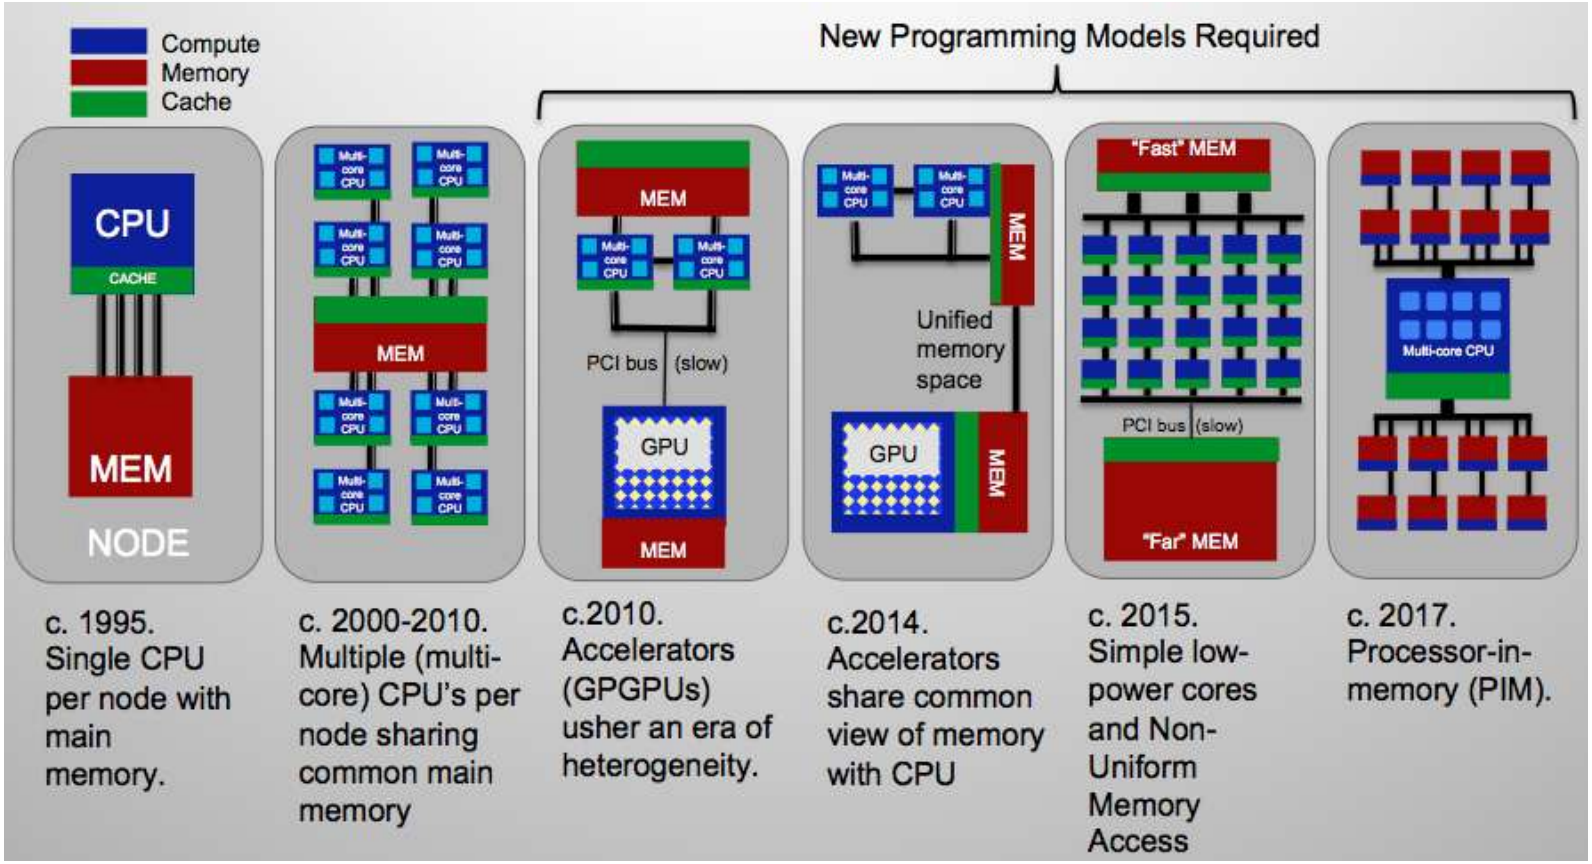

We have entered a new era in HPC architectural complexity and need to move beyond High Performance LINPACK (HPL) as a metric

HPL: Pros

- Easy to run
- Easy to understand
- Easy to check results
- Good tool for community outreach
- “Understandable” to the outside world
- Historical database of performance information

HPL: Cons

- Has poor balance of floating point and data movement compared to modern codes
- Overall usability of a system is not measured
- Used as a marketing tool
- Can require long run times—wasting valuable resources
- Not sensitive to new architectural features
- Does not have sufficient fidelity for procurements

Promote the pros, fix the cons-- Evolving the community benchmark

- Develop a new metric that correlates with important scientific and technical apps not well represented by HPL
  - Replicate the good (enduring) features of HPL
  - Replace the outdated features

- Accurately predict rankings for a target suite of scientific applications

- Encourage vendors to focus on architectural features needed for high performance on important scientific and technical apps

- Not intended to define procurements

- PLUS--Support a historical record of performance information on existing and future systems

Proposal: HPCG for ranking scientific systems

- High Performance Conjugate Gradient (HPCG)
  - Solve *Ax=b, A* large, sparse, *b* known, *x* computed
  - Physics-based A matrix
- Contains communication patterns that are prevalent in a variety of methods for discretization and numerical solution of PDEs
- More relevant patterns of computation:
  - Dense and sparse computations
  - Dense and sparse collectives
  - Data-driven parallelism

HPC Technical Reports

- *HPCG Technical Specification*
  - Jack Dongarra, Michael Heroux, Piotr Luszczek
- *Toward a New Metric for Ranking High Performance Computing Systems*
  - Jack Dongarra and Michael Heroux

HPCG results presented at ISC 2014

June 2014 Top 500	Site	Computer	Cores	Peak (Pflops)	HPL RMAX (Pflops)	HPCG (Pflops)	HPCG/RMAX
1	National Super Computer Center in Guangzhou	Tianhe-2 NUDT,Xeon 12C 2.2GHz + IntelXeon Phi (57c) + Custom	3,120,000	54.90	33.9	0.58	1.71%
2	DOE / OS Oak Ridge Nat Lab	Titan, Cray XK7 (16C) + Nvidia Kepler GPU (14c) + Custom	560,640	27.10	17.6	0.322	1.83%
4	RIKEN Advanced Inst for Comp Sci	K computer Fujitsu SPARC64 VIIIfx (8c) + Custom	705,024	11.30	10.5	0.426	4.06%
5	MIRA DOE / OS Argonne Nat Lab	BlueGene/Q, Power BQC 16C 1.60GHz, Custom	786,432	10.10	8.59	0.101	1.18%
6	Swiss CSCS	Piz Daint, Cray XC30, Xeon 8C + Nvidia Kepler (14c) + Custom	115,984	7.80	6.27	0.099	1.58%
11	HPC2	Intel Xeon 10C 2.8 GHz + Nvidia Kepler (14c) + IB	62,640	4.00	3	0.0489	1.63%

- HPCG is real and has been run on several systems
- Performance is consistent with our expectations and experience

Comments on early HPCG benchmark results

- The disparity between HPL and HPCG is not a surprise, it's a fact of life
- The results reflect the intrinsic nature of many challenging scientific applications: climate, combustion, turbulence, etc…
- These are typical of the currently available systems for mission-critical applications
- Not all vendors have developed optimized versions

In Summary…

- The transition to the next era in high-end computing is going to affect all scientific computer users long before an exaflop system is available
- We need to take a comprehensive approach to next-gen platforms
- We are preparing for the inevitable and significant changes through
  - Hardware and software codesign efforts
  - Funded collaborations with industry to ensure that exascale architecture computers will meet our scientific computing needs
  - Application code redesign to address expected processor, memory and storage changes
- We are investigating new, more informative ways to measure performance

Thank You

Robert E. Meisner

Office of Advanced Simulation and Computing

U.S. DEPARTMENT OF ENERGY