

ASCAC Subcommittee on Exascale Modeling & Simulation

Robert Rosner
For the Subcommittee

March 30, 2010
Washington, DC

Steve Ashby
Pete Beckman
Jackie Chen
Phil Colella
Bill Collins
Dona Crawford

Jack Dongarra
Doug Kothe
Rusty Lusk
Paul Messina
Tony Mezzacappa
Parviz Moin

Mike Norman
Vivek Sarkar
Andrew Siegel
Fred Streit
Andy White
Margaret Wright

The Charge(s)



Department of Energy
Office of Science
Washington, DC 20585
October 29, 2009

Office of the Director

Professor Roscoe Giles, Chair
Department of Electrical & Computer Engineering
Boston University
8 St. Mary's Street
Boston, MA 02215

Dear Professor Giles:

Over the last few years, several workshops and subcommittee reports have identified and described the scientific opportunities for high performance computing. By this letter I am charging the Advanced Scientific Computing Advisory Committee (ASCAC) to assemble a subcommittee to look at the results of these activities and to analyze the opportunities and challenges for the Office of Advanced Scientific Computing Research (ASCR) and the Office of Science associated with exascale computing. Specifically, I would like the sub-committee to deliver a report that:


- Assesses the opportunities and challenges of exascale computing for the advancement of science, technology, and Office of Science missions.
- Identifies strategies that ASCR can use to address the challenges and deliver on such opportunities.

We would appreciate the committee's preliminary comments by March 30, 2010, and a final report by August 15, 2010. I appreciate ASCAC's willingness to undertake this important activity.

If you have any questions regarding this matter, please contact either Michael Strayer, the Associate Director of the Office of Science for ASCR, or Christine Chalk, the Designated Federal Official for the ASCAC.

Sincerely,

W. F. Brinkman
Director, Office of Science



Printed with soy ink on recycled paper



Under Secretary for Science
Washington, DC 20585

November 2, 2009

Professor Roscoe Giles
Chair, Department of Electrical & Computer Engineering
Boston University
8 St. Mary's Street
Boston, MA 02215

Dear Prof. Giles,

As your subcommittee begins to work on Dr. Brinkman's new charge, I would like for you to consider the following.

Modeling and simulation at the extreme scale have the potential to span the entire Department and to forge lasting interlocks between applied and basic science, technology, and engineering. I would therefore ask the sub-committee to acknowledge and comment upon the broader issue of opportunities and challenges for exascale computing to advance *Department of Energy* missions. I can imagine that will entail including among the subcommittee members familiar with both NNSA and various applied programs that are amenable extreme scale computing.

We live in times when the Nation is faced with important issues involving energy, environment, and national security, yet resources are constrained. It is therefore very important to have clear justification for how new endeavors in extreme scale computing (ultimately reaching the exascale) will affect science, technology, and society.

I look forward with great interest to your report.

Sincerely,

Steven E. Koonin
Under Secretary for Science



Printed with soy ink on recycled paper

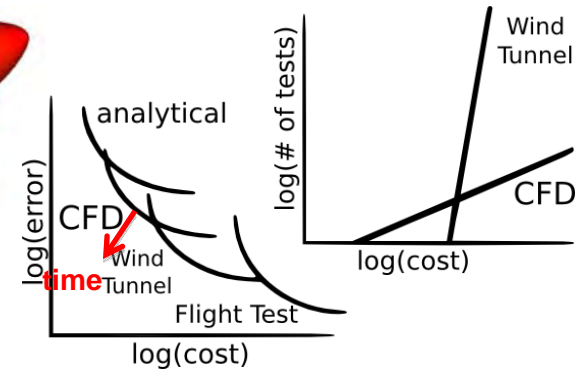
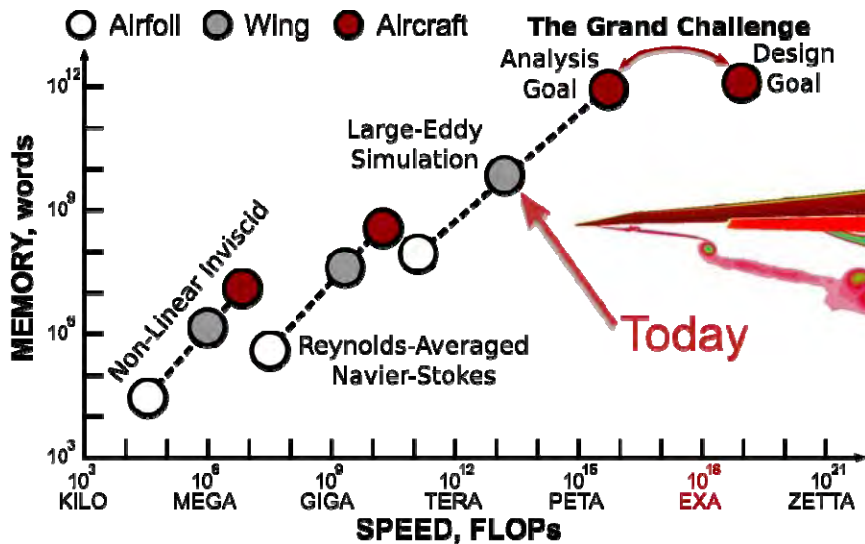
Our Committee Strategy and 'ground rules' for answering the Charge(s)

- Don't go discipline-by-discipline – that is what previous Exascale Workshop reports have done already
 - **Use Workshop reports as inputs to our Report**
- Focus on transformational calculations: What can be done that will change a discipline, change the way work is done, transform an industry ...
 - **Bigger may not be better (enough) ...**
- Identify the roadblocks common to all apps: from programming models to scalable algorithms to unstable architectures to
 - **What things need to be worked on in order to succeed in getting to the exascale on the software and hardware end of things?**
 - **How does what is now known about the (hardware) roadmap to the exascale affect our prospects?**
- Think broadly – not only about Office of Science applications, but also about more general applications, including applied science and engineering disciplines, NNSA, and industry

The structure of our (interim) Report


- Introduction: What's so exciting about exascale computing and beyond?
- What are the 'big questions' that can be answered by exascale computing?
- What are the 'big obstacles', and what needs to be done to resolve them?
- ... and then there are some things we did not get to cover **as yet**, in particular
 - Biology
 - Health sciences, and more generally
 - Disciplines that focus on extremely large data sets, or deal with data in extremely high-dimensional 'state spaces'

In concrete terms, what does 'transformational' really mean? Consider the impact of HPC & CFD on Boeing's Design Cycle




“High Performance Computing has fundamentally changed the way that Boeing designs flight vehicles.”

-Director, Boeing Commercial Airplanes



Consider predictive simulation by Cummins, leading to a new diesel engine design, and by Goodyear, leading to a new tire design





Cummins achieved a milestone in engine design by bringing a diesel engine, the 2007 ISB 6.7 liter, to market solely with computer modeling and analysis tools. The only testing was after-the-fact to confirm performance. Cummins achieved a reduction in development time and cost (estimated to be about 10 to 15% for this first effort). As important, they realized a more robust design, improved mileage, and met all environmental and customer constraints.



Goodyear's Assurance® Triple Tred all-weather tire was its first product designed using predictive modeling simulation tools developed in conjunction with Sandia National Laboratories. This tire and the subsequent products utilizing advanced modeling capabilities resulted in a factor of three reduction in product development time and led to record profits for Goodyear.

Consider the discovery of a possible new mechanism for Type Ia supernova explosions





The gravitationally-confined detonation of white dwarfs as a possible mechanism underlying Type Ia supernovae was discovered using the HPC *Flash* code.

Prior to this discovery, it was not known how the transition between nuclear deflagration and detonation (DDT) occurred; and it was simply assumed that it did occur. The *Flash* simulation provided the first concrete physical explanation how such a DDT could arise in a system without hard boundaries (e.g., container walls).

The *Flash* code – funded by the DOE/ASC Academic Alliance program at the Univ. of Chicago - was designed from its conception to be an open-source compressible computation CFD code capable of running on any of the largest computers in the world.

The computations shown here were supported by both the NNSA/ASC and the SC/ASCR INCITE programs

Consider the winner of the 2007 Gordon Bell prize: first-of-a-kind micron-scale simulation of Kelvin-Helmholtz instability in molten metals



While Kelvin-Helmholtz instability has been thoroughly studied for years and its behavior is well understood at the macro-scale, scientists did not clearly understand how it evolves at the atomic scale until this LLNL simulation by J.N. Glosli et al. was carried out.

Understanding how matter transitions from a continuous medium at macroscopic length scales to a discrete atomistic medium at the nanoscale has important implications for applications ranging from spray formation at the air-ocean interface to laser fusion experiments.

What did these examples all have in common?

- ▶ HPC was essential
 - The problems could not be solved any other way
- ▶ Quantitative prediction
 - It was not just all about the pictures
 - It was all about understanding
- ▶ Game-changing
 - Scientific and/or technical discovery
 - Exploration of phenomena where physical experiments are difficult or impossible to carry out
 - Faster (non-Edisonian) 'design' and sub-system/system-level optimization, which allows for
 - Systematic, rather than random, exploration of phenomena,
 - Completion within the effective attention span of expert investigators

Can we expect similar advances by going beyond the petascale, to the exascale?

The exascale and related workshops have identified specific exascale applications and their benefits

- ▶ They have examined the role of *extreme scale* computing in many areas
- ▶ A real (= credible) case for value-added by going to the exascale has been established in a handful of areas (including those at right)
- ▶ Compelling societal, economic, and scientific benefits are associated with a few areas

But: we need to remember that it is in the nature of this subject that some of the ultimately most interesting applications will be ones we have not yet thought of, or about ...

- ▶ Climate
- ▶ Combustion
- ▶ Nuclear Reactors
- ▶ Catalysis
- ▶ Electric Grid
- ▶ CCS
- ▶ Fusion
- ▶ Nuclear Stockpile
- ▶ Supernovae
- ▶ Materials
- ▶ Accelerators

One can make the case that exascale simulation will deliver real benefits to Society, Industry, and Science

▶ Societal benefits

- National Security (stockpile stewardship/weapons certification in the absence of testing)
- Climate Understanding (planetary stewardship)
- Clean Energy Sources (quality of life maintenance)

▶ Industrial benefits

- Electricity production
- Alternative fuels
- Combustion engine efficiency
- Reduce design cycle time & expensive physical testing

▶ Scientific Discovery

- Understanding the world around us
- Manipulating matter creatively

Exascale simulation could enable these advances within the next decade

- ▶ Climate
- ▶ Combustion
- ▶ Nuclear Reactors
- ▶ Catalysis
- ▶ Electric Grid
- ▶ CCS
- ▶ Fusion
- ▶ Stockpile
- ▶ Supernovae
- ▶ Materials
- ▶ Accelerators
- ▶ Understand water and carbon cycles
- ▶ 30% efficiency gain in combustion engines
- ▶ High-burn (low-waste), proliferation-resistant designs
- ▶ Infrastructure-ready bio-fuels
- ▶ Robust, adaptive smart-grid
- ▶ Emissions sequestration with confidence
- ▶ Understand plasma behavior in fusion reactors
- ▶ Understand weapons aging and performance
- ▶ Understand source of nuclei in periodic table
- ▶ Fundamental nature of materials
- ▶ Predict beam loss and activation

The following exemplify – and illustrate how – exascale computing can enhance USA leadership in key technology areas

▶ Information technologies

- Processors, computer systems
- Programming paradigms
- Algorithms
- Application codes

▶ Materials science and engineering

- Fundamental understanding of materials: their efficient and ‘process-aware’ production and their ultimate use
- Myriad applications, from storage (e.g., batteries) and sensors


▶ Systems-level engineering

- Effective sub-system to system-level optimization of processes
- Rapid prototyping brought to the systems-level
- Myriad applications, from smart grids to nuclear reactor design and operation

I will briefly provide some specific examples ...

- ▶ There are strong arguments – and a technical consensus – for benefits derived from exascale computing in a number of select areas
 - Climate
 - Combustion
 - Reactor Design
 - Weapons Science
- ▶ However: In none of these cases is simulation technology now ready for going to exascale – obstacles discussed ‘below’ currently stand in the way

Simulating Climate Change



Drivers:

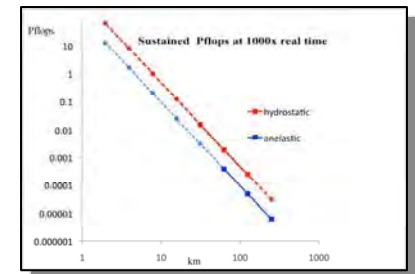
- ▶ What are the critical cloud controls on climate and the hydrologic cycle?
- ▶ What is the strength of the global carbon sink, and how will it change?

Exascale Impact:

- ▶ Reliable predictions of water and carbon cycles in warmer climates



Current state:

- ▶ Our best petascale climate simulations are quite uncertain due to parameterization of clouds, convection, and ocean eddies.



Need for exascale capability now:

- ▶ Simulating clouds at their native scale for global climate is inherently an exascale problem.
- ▶ Modeling fully turbulent exchange of heat and gases between the atmosphere and ocean is an exascale problem.



Transformational exascale capability:





- ▶ Robust climate models for early warning, adaptation, and mitigation



For further reading: *Challenges in Climate Change Science and the Role of Computing at the Extreme Scale* (DOE BER, 2009).

Predictive Simulation of Combustion for Transportation and Power Generation in an Evolving Fuel Environment

- ▶ **Drivers:** Clean Energy & Security Act 2009
- ▶ **What:** Validated predictive simulation that industry will use in the design cycle
- ▶ **Transformational technology:** 30% increase in efficiency while meeting stringent emissions standards (2/3 of Nation's energy consumption and 1/4 of greenhouse gas emissions is due to transportation)
- ▶ **Needed science and tools:** Combustion science (including high-pressure chemistry and mixed-mode multi-phase combustion of renewable bio-fuels); coupled exascale simulations and experimentation to provide benchmark data for physical insights and validation of models
- ▶ **Requirements for exascale:** Incorporating complex chemistry of biofuels ($\sim O(100)$ transported species, 1000's reactions) in CFD; higher spatial dynamic range driven by pressure and Reynolds # scaling; multi-phase CFD



Next-generation nuclear reactor modeling

- ▶ **Driver:** The need for a broad mix of energy technologies that both avoid further contributions to global warming and serve as reliable energy sources for base electric power has led to a renewal of interest in nuclear power
- ▶ **What:** Validated predictive (science-based) modeling and simulation tools for the design, construction & operation of new-generation nuclear reactors, reactors that
 - Are cost-effective both in construction and operations
 - Are safe and secure to operate
 - Are proliferation-resistant by design
- ▶ **Transformational science/engineering**
 - High fidelity, robust and well-validated thermal hydraulics, neutronics and structure modeling tools
 - Fully coupled thermal hydraulics, neutronics and structure analyses
 - Predictive materials analysis tools for fuels, cladding, reactor vessel welds, ..., all needed for increased fuel utilization, power uprates, and reactor life extensions
- ▶ **Requirements for exascale:** Fully coupled thermal hydraulics, neutronics & structure modeling, w/ iterations required by systems-level optimization, will call for exascale capabilities

Simulating Core Collapse Supernovae



- ▶ **Driver:** Understanding the physics of Type II supernovae is a necessary ingredient in establishing a scientifically secure understanding of cosmic nucleosynthesis

- ▶ **What:** Understand origin of a large fraction of the elements in Periodic Table, especially the heavy nuclei

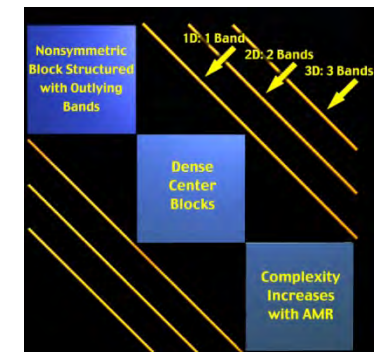
Core Collapse SNe dominate element production between O and Fe and are likely responsible for half of the elements heavier than Fe (the “heavy” elements).

- ▶ **Transformational Science:** Physically realistic 3-D supernova models

3D supernova simulation will involve linear systems with > 1 trillion unknowns.

- ▶ **Requirements for exascale:**

- Solution of exascale linear systems underpinning solution of the neutrino transport equations will require exascale platforms
- Each complete multi-physics supernova simulation will require days to weeks to complete on such platforms




Stockpile Stewardship – and ‘Going to Zero’

- ▶ **Driver:** Assuring the security, safety and operability of our nuclear weapons stockpile – and doing so in an internationally believable fashion – as the total # of weapons is sharply reduced
- ▶ **What:** For political and geopolitical reasons, our confidence in the national nuclear weapons stockpile must remain high in the face of no future nuclear weapons tests (e.g., the CTBT and NPT), sharply reduced numbers of on-duty weapons, and the certainty that there is no known alternative to stockpile certification via simulations
- ▶ **Exascale simulation**
 - Greater physics fidelity (fewer ‘knobs’)
 - Advanced V&V using NIF, DARHT, ...
 - Increased confidence in life extension programs ...

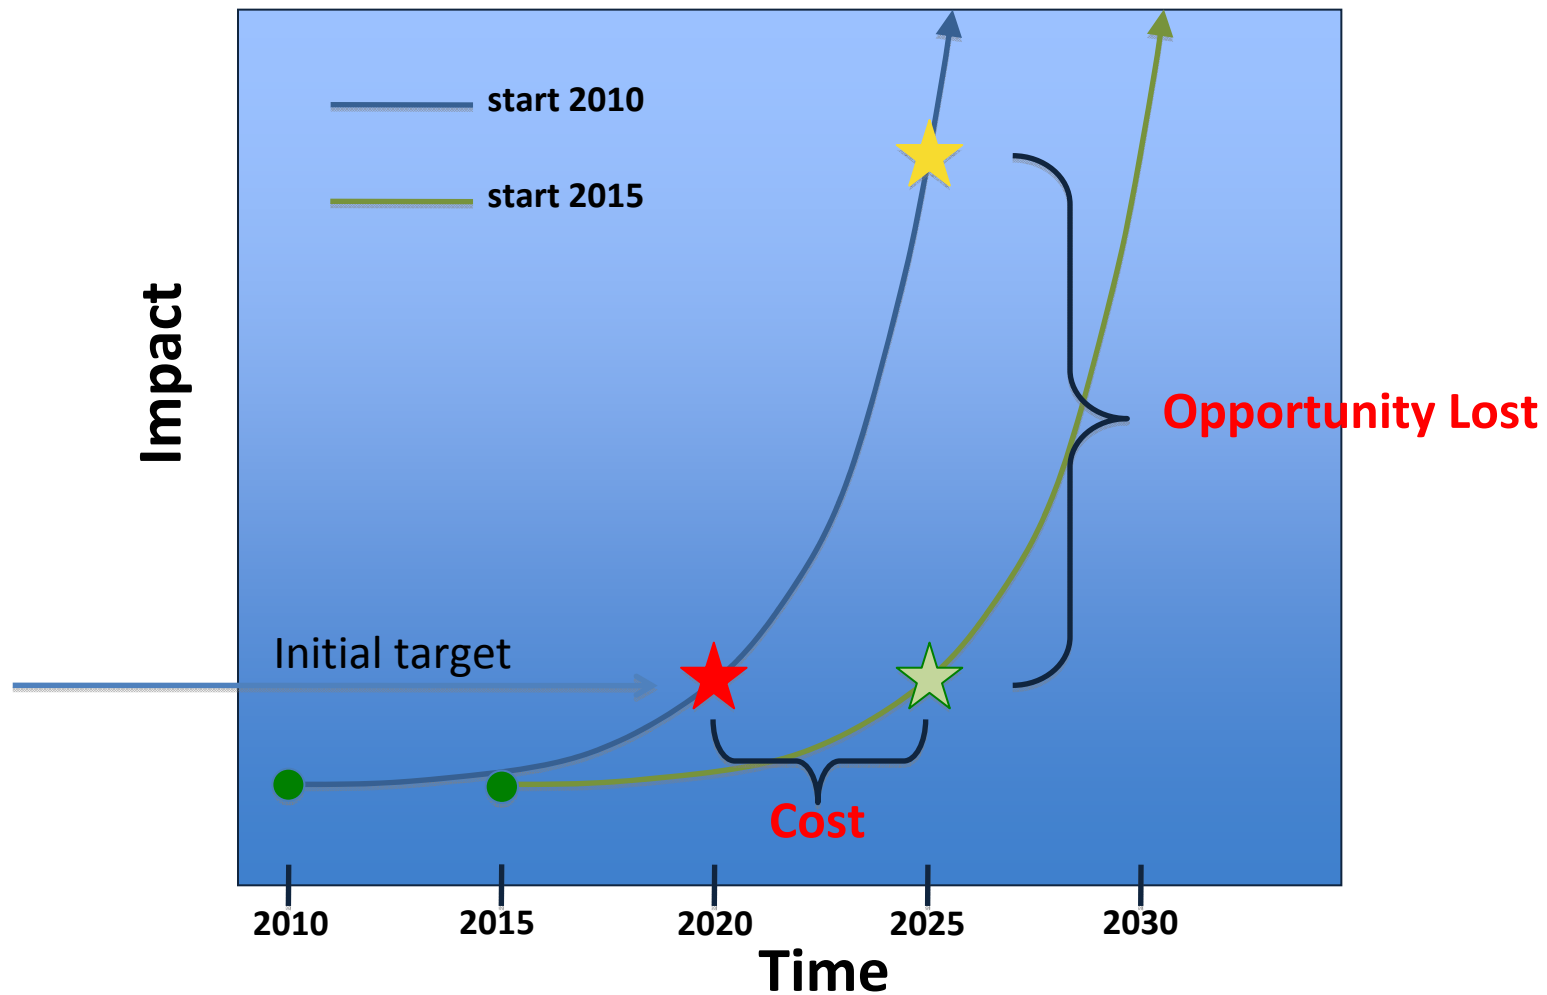
Exascale Now! There *is* considerable time urgency

- ▶ Some applications will need this capability
 - Global security: 'Going to Zero'
 - Meet near-term CTBT obligations (including reducing the stockpile) while maintaining a credible deterrence
 - Clean Energy and Security Act of 2009
 - 80% reduction CO2 from 2005 levels by 2050
 - Energy security
 - ...
- ▶ American competitiveness and domestic jobs
 - Position USA to dominate tomorrow's industries: owning the IP, producing ...
 - Create a broad spectrum of high-value jobs
- ▶ Enabling and enhancing USA Science
 - Long-term industrial competitiveness depends on the strength of our underlying science and technology base
 - Going to exascale will support existing and planned leading edge science experiments & facilities
- ▶ We are not alone ...
 - Others are competing head-on with us

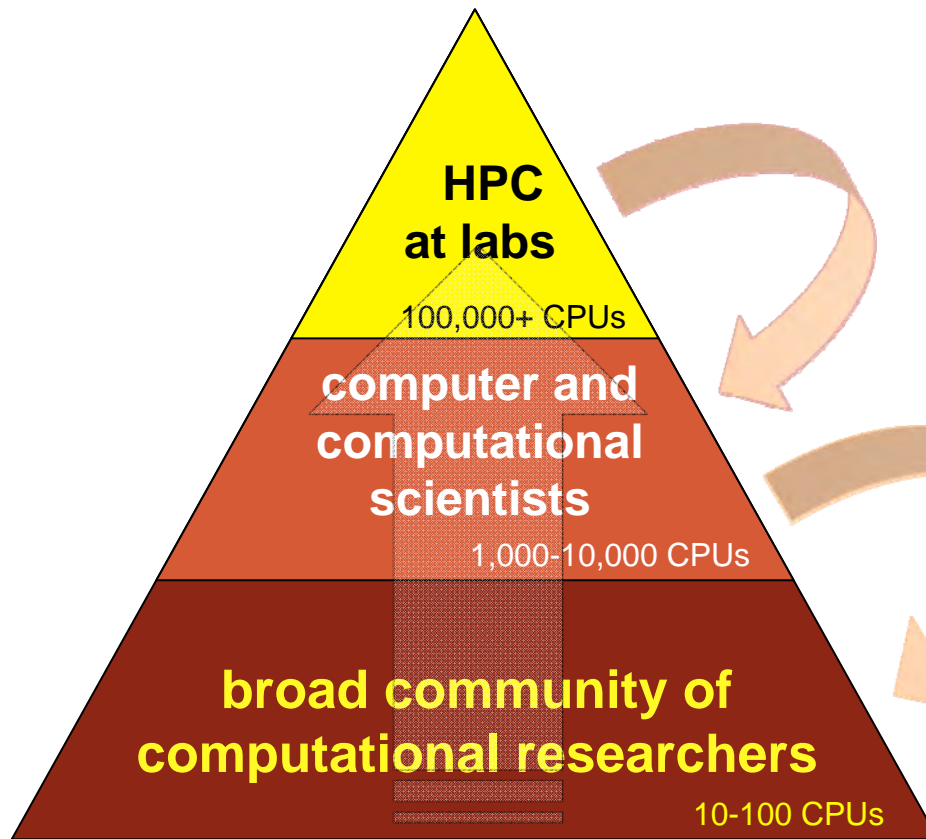
For example, China is aggressively investing in HPC ...



Failure to invest in emerging exascale technology will have a profound impact on US competitiveness



Each generation of architecture exposes new challenges and opportunities



- Pioneering simulations on bleeding-edge technology leads to improvements in algorithms and methods
- Collaborations hasten adoption to lower levels of scalability
- Wider community leads to greater levels of innovation

Achieving exascale computing will require the engagement of every level of the HPC community

... and, unsurprisingly, there are obstacles/barriers to getting to the exascale


Overall, what seems to be needed is a “Value Proposition” for HPC Applications and Technology

- ▶ Optimize simulation/modeling software, algorithms, software, technology choices, based on applications outcomes
 - Modeling and algorithmic approaches, tied to (and integrated with) architecture choices
 - Error tolerance
 - Checkpoint requirements
 - CPU/memory balance
- ▶ Standards for preserving and archiving computational results
 - Published results should be reproducible ...

Consensus programming models and environments do not yet exist, but require early investment



- ▶ **Barriers:** Delivering a large-scale scientific instrument that is productive and fast.
 - O(1B) way parallelism in exascale system
 - O(1K) way parallelism in a processor chip
 - Massive lightweight cores for low power
 - Some “full-feature” cores lead to heterogeneity
 - Data movement costs power and time
 - Software-managed memory (local store)
 - Programming for resilience
 - Science goals require complex codes



How much parallelism must be handled by the program?

From Peter Kogge (on behalf of Exascale Working Group), “Architectural Challenges at the [Exascale Frontier](#)”, June 20, 2008

▶ Technology Investments

- Extend inter-node models for scalability and resilience, e.g., MPI, PGAS (includes HPCS)
- Develop intra-node models for concurrency, hierarchy, and heterogeneity by adapting current scientific ones (e.g., OpenMP) or leveraging from other domains (e.g., CUDA, OpenCL)
- Develop common low level runtime for portability and to enable higher level models

▶ Technical Gap:

- No portable model for variety of on-chip parallelism methods or new memory hierarchies
- Goal: Hundreds of applications on the Exascale architecture; tens running at scale



Candidate programming models are available ...

- ▶ At least in the near term, hybrid approaches:
 - Well-known MPI among address spaces
 - MPI-3 addressing scalability, interfaces for multiple hybrid approaches, reference implementations needed
 - Multiple approaches for parallelism within an address space
 - MPI + {OpenMP, PGAS, Cuda, OpenCL, ...}
 - Reference implementations needed
- ▶ Longer term, unified approaches may find use
 - PGAS (UPC, CAF)
 - HPCS languages (Chapel, X10)
 - High-performance reference implementations needed
- ▶ Even low-level models are useful for implementing higher-level approaches, such as domain-specific languages.



Algorithms, architecture and “co-design”

- ▶ Model and algorithm development is needed to both support new architectures, **and** take advantage of them
 - Applied math/computational math (= ‘Applicable Mathematics’) must develop methods that both take advantage of, and be exploited in, new architectures
 - Algorithms must reflect impact of memory and communication constraints, rather than flops
 - Locality more important than ever
 - Out-of-core algorithms to avoid memory size constraints
 - New paradigms: streams, UQ inside the loop, reduced precision
 - Move away from bulk-synchronous programs to multi-task approaches

- ▶ Modeling and Simulation for Co-Design
 - Need for simulation tools that emulate node architectures in order to test new algorithms (PDE)
 - Need for “mini-applications” that provide info to hardware and system (PDE)

- ▶ Address question of in-situ vs. off-mainframe data intensive (e.g. analysis) processing
 - Diverse data structures for codes are confounding issue



System and programming model support

- ▶ Communications & concurrency
 - small light-weight messaging; light-weight fine-grained synchronization; high degree of threading, fast collectives; communication patterns; overlapping of computation, communication and I/O
- ▶ Memory access
 - Global arrays; reconfigurable memory hierarchies; data provenance; cache coherence?
- ▶ Performance / Resource measurement and analysis tools
- ▶ Fault management/system resilience
 - Taxonomy of faults; relationship of fault-tolerance & UQ: tools for fault tolerance management; API for checkpointing; fault tolerant MPI collectives
- ▶ Dynamic resource allocation
 - Scheduling; load balancing; UQ ensemble calculations; adaptive runtime systems; power allocation;
- ▶ New exascale programming paradigm: MPI+X?
 - Need for MPI+X (hybrid) programming model with effective abstractions and abstract machine model
 - Since PM is changing anyway, it's a good time to design new functionality: support for multiple networks; "uncertainty-carrying" variables; ...
- ▶ I/O ...
- ▶ And – don't forget legacy interoperability!
 - Interoperability with old PMs/languages needed to promote adoption of the new PMs/paradigms (PDE)





The “Concurrency challenge”

- ▶ Goal: expose at least 1000x more concurrency relative to current systems, subject to the following constraints
 - 10x – 100x lower bytes/ops ratios than current systems
 - Serial bottlenecks arising from software stack (Amdahl’s Law)
 - Additional concurrency for latency hiding (Little’s Law)
 - Data movement constraints implied by energy constraints

- ▶ Implications
 - Exploit strong scaling or “new era” weak scaling in applications
 - Redesign software stack to reduce serial bottlenecks
 - Redesign software stack to enable more efficient data movement



Projected Concurrency in High End Computing





Implications for Exascale Application Development

Challenge: increased levels of concurrency in applications

- ▶ Strong scaling: apply more resources to the same problem size to get results faster
 - In practice, many applications are not amenable to strong scaling
 - May need to rewrite application from 1st principles to obtain strong scaling, e.g., Qbox
- ▶ Traditional weak scaling: use more resources by increasing problem size
 - Reduce grid size and time-step interval
 - Memory requirement increases in $\frac{3}{4}$ power for 3-D problems and $\frac{2}{3}$ power for 2-D problems
 - 1000x increase in computational work requires ~ 180x and 100x in memory for 3-D and 2-D problems respectively
- ▶ “New-era” weak scaling
 - Application trends in which additional work is done per datum e.g., multi-scale, multi-physics, interaction analysis, data mining.
 - No increase in memory requirement



Performance Nondeterminism

Challenge: Exascale systems will exhibit performance nondeterminism at multiple levels

- ▶ **Node performance:** multithreading, heterogeneous cores, variable clock frequencies, variable memory latencies
 - Even regular computations may appear to be imbalanced across nodes (jitter magnification)
- ▶ **Communication performance:** multithreaded nodes, interconnect structure, messaging runtime
 - Communication performance (latency & bandwidth) will vary with source node, destination node, message size, as well as degree of multithreading on source and destination nodes
- ▶ **Resiliency overhead:** checkpointing, recovery
 - Performance impact of resiliency overhead will depend on TTI distribution, I/O links, and checkpointing & recovery overheads

Need solutions for decreased reliability and a new model for resiliency



► Barriers

- System components, complexity increasing
- Silent error rates increasing
- Reduced job progress due to fault recovery if we use existing checkpoint/restart

► Technical Focus Areas

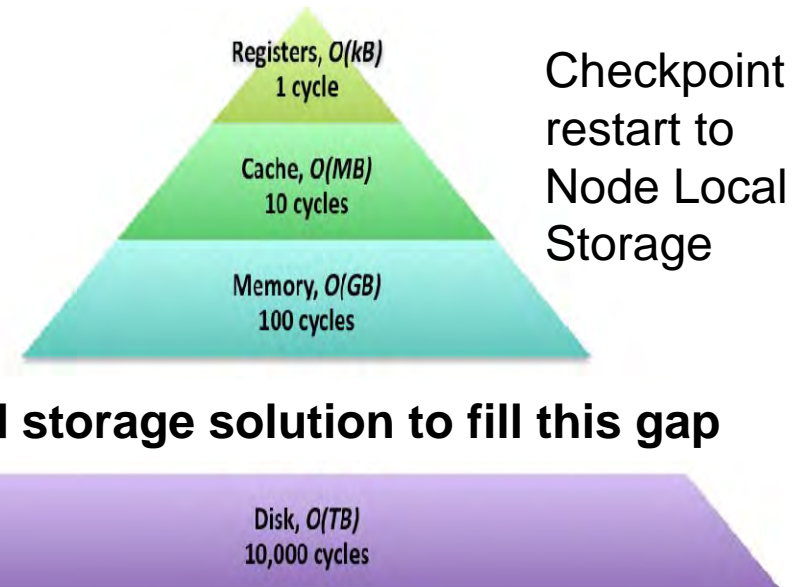
- Local recovery and migration
- Development of a standard fault model and better understanding of types/rates of faults
- Improved hardware and software reliability
 - Greater integration across entire stack
- Fault resilient algorithms and applications

► Technical Gap

- Maintaining today's MTTI given 10x – 100x increase in sockets will require:
- 10X improvement in hardware reliability
- 10X in system software reliability, and
- 10X improvement due to local recovery and migration, as well as research in fault resilient applications

Taxonomy of errors (h/w or s/w)

- **Hard errors:** permanent errors which cause system to hang or crash
- **Soft errors:** transient errors, either correctable or short term failure
- **Silent errors:** undetected errors either permanent or transient. *Concern is that simulation data or calculation have been corrupted and no error reported.*



Export Control and International Collaboration

- ▶ Advances needed for extreme scale computing require innovation and persistent development efforts by top-level research groups around the world
- ▶ Coordination and open sharing of results among these efforts is critical to engender benefits to DOE
- ▶ DOE must facilitate these interactions by appropriate interpretation of export control laws

Intra-DOE & Extra-DOE (=Inter-Agency) interactions

- ▶ SC “owns” math, computer science, applications, and leadership class computers
- ▶ NNSA “owns” math, computer science, applications, and leadership class computers
- ▶ Applied Energy Offices “own” applications
- ▶ Working across DOE to gain the benefit of extreme scale computing is essential to advance the goals of both DOE and each individual office
- ▶ Coordination across federal agencies and with industry will help maximize investments and ensure the bases are covered

Education and Training

▶ The problem:

- Going to the exascale will prove at least as challenging as the transition from vector to massively-parallel computing
- The workforce experienced in the requisite software underpinnings does not exist at the scale needed for this transition; expertise is needed to
 - Develop codes
 - Use codes

▶ The needed solution:

- Start ongoing workforce development ***now***
- Outreach from “exascale community” to industrial users necessary
- University programs in high-performance computing

... which brings us to

▶ Questions & Discussion