Advanced Scientific Computing Research

Michael Strayer

Associate Director, Advanced Scientific Computing Research

Office of Science

Department of Energy

Department of Energy Organization

Dr. Steven Chu, Secretary of Energy
Winner of Nobel Prize for Physics in 1997

Dr. Steven Koonin, Undersecretary for Science
Formerly, Chief Scientist for BP; Provost of CalTech (1995-2004)

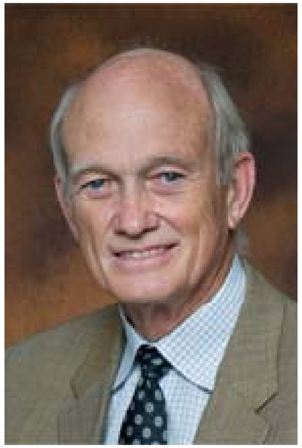

Dr. William Brinkman, Director, Office of Science
Former Vice President of Bell Laboratories; Vice President of Research at DOE’s Sandia National Laboratories

Dr. Michael Strayer, Associate Director, Office of Advanced Scientific Computing Research;
Formerly Distinguished R&D Staff at ORNL

U.S. DEPARTMENT OF
ENERGY

ASCR Vision

First in Computational Science

Recognized as "Best in Class" advancing science and technology through modeling and simulation

- **Excellence in applied mathematics and computer science research**
  - **Research Programs started in 1950's**
- **Cross-disciplinary partnerships**
  - **In 2000, initiated Scientific Discovery through Advanced Computing**
- **Forefront Computing for Science Applications**
  - **Top Computing Facilities in the World for Open Science**

U.S. DEPARTMENT OF ENERGY

DOE - Extreme scale computing today: 15 years of world leadership

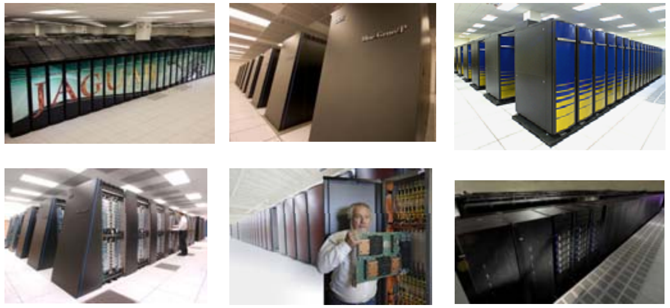

The Office Of Advanced Scientific Computing Research

NNSA
Advanced Simulation and Computing

Top 500 list, November 2009

Machine	Place	Speed (max)	On list Since
Jaguar	ORNL	1.75 PF	2009 (1)
Roadrunner	LANL	1.04 PF	2009 (3)
Dawn	LLNL	0.478 PF	2007 (7)
BG/P	ANL	0.458 PF	2007 (8)
NERSC	LBL	0.266 PF	2008 (15)
Red Storm	SNL	0.204 PF	2009 (17)

INCITE: 2.5x oversubscribed

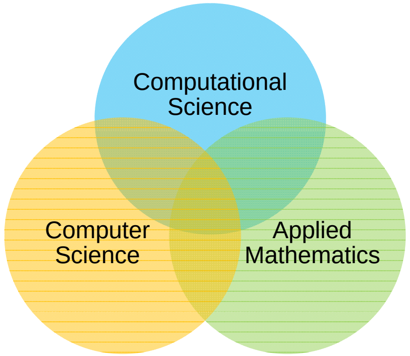

ASC and ASCR *provide much more than machines*:

• Applications	(Computational Science)
• Algorithms	(Applied Mathematics)
• Systems	(Computer Science)
• Integration	(SciDAC, Campaigns)

ASCR Facilities Roadmap

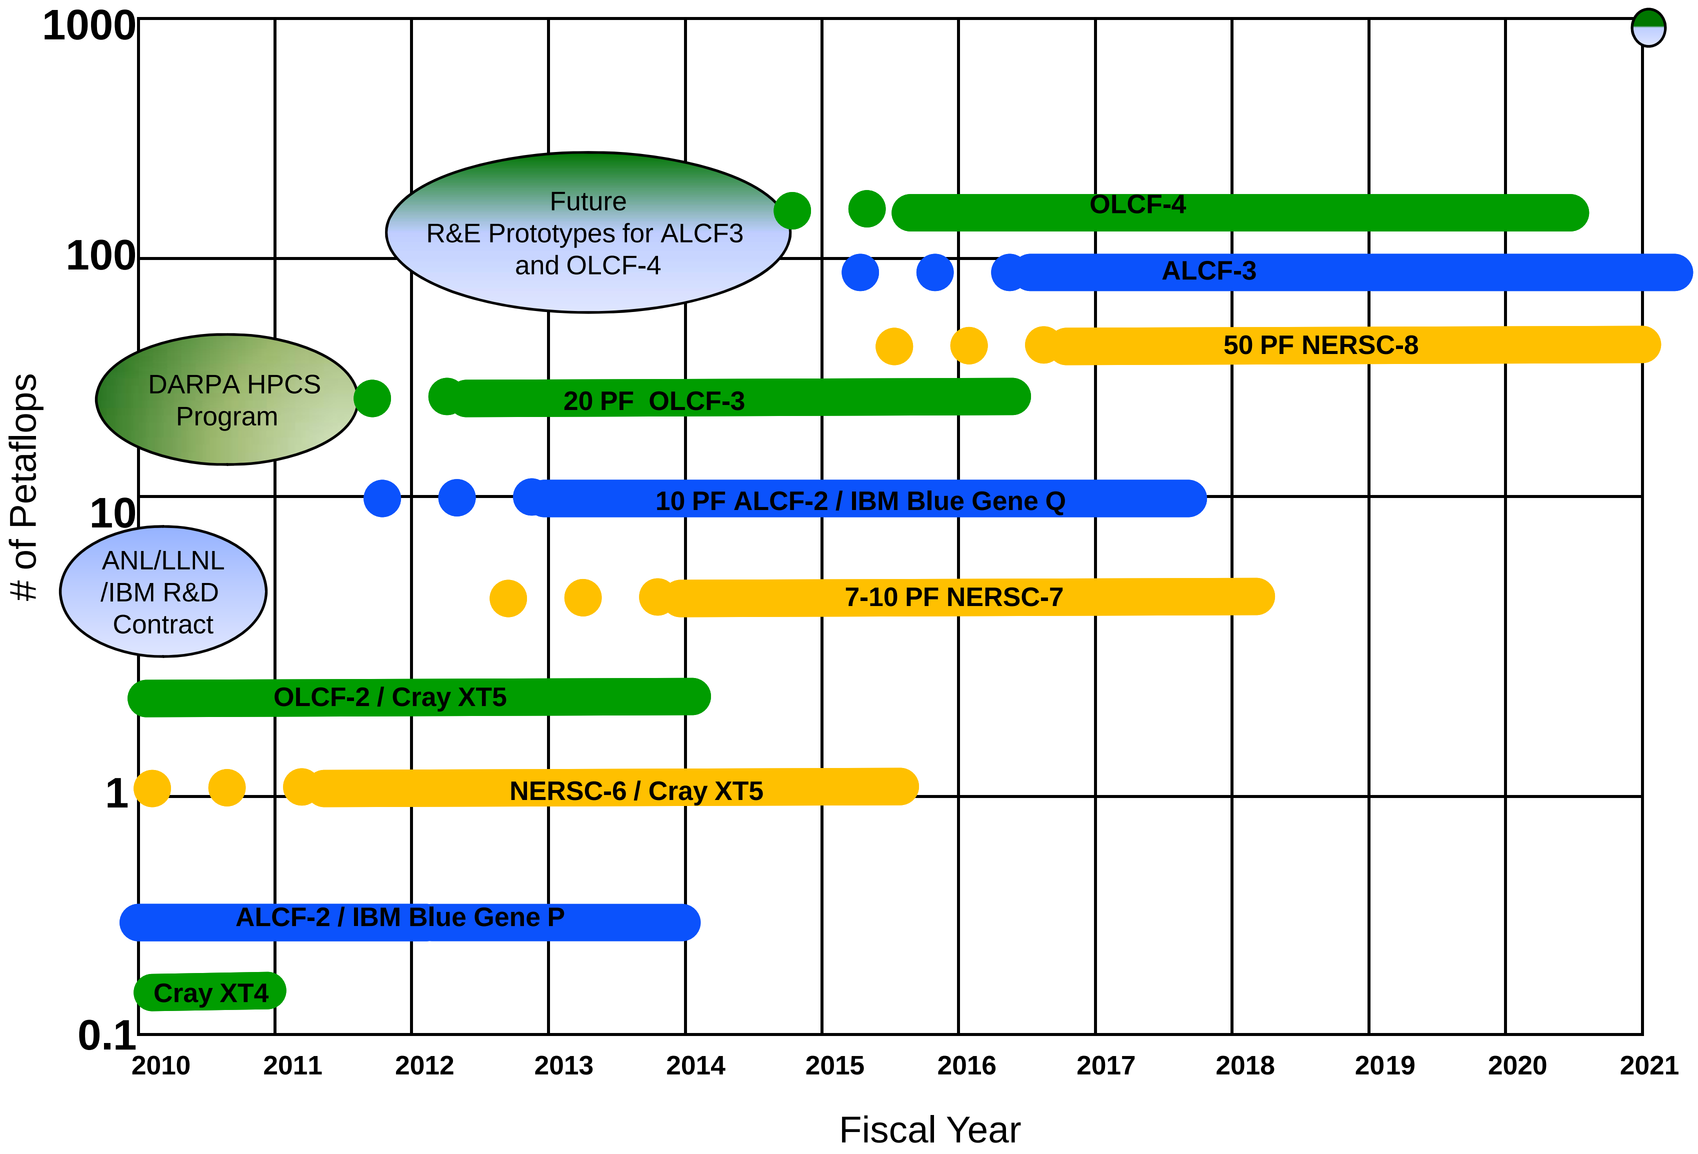

Energy Sciences Net

Innovative and Novel Computational Impact on Theory and Experiment

- **Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program started in 2004.**
- **INCITE**
  - **Provides Office of Science computing resources to a**
    - **computationally intensive research**
    - **high-impact scientific advances**
  - **Open to national and international researchers, including industry**
  - **No requirement of DOE or Office of Science funding or topic area**
  - **Peer and computational reviews**

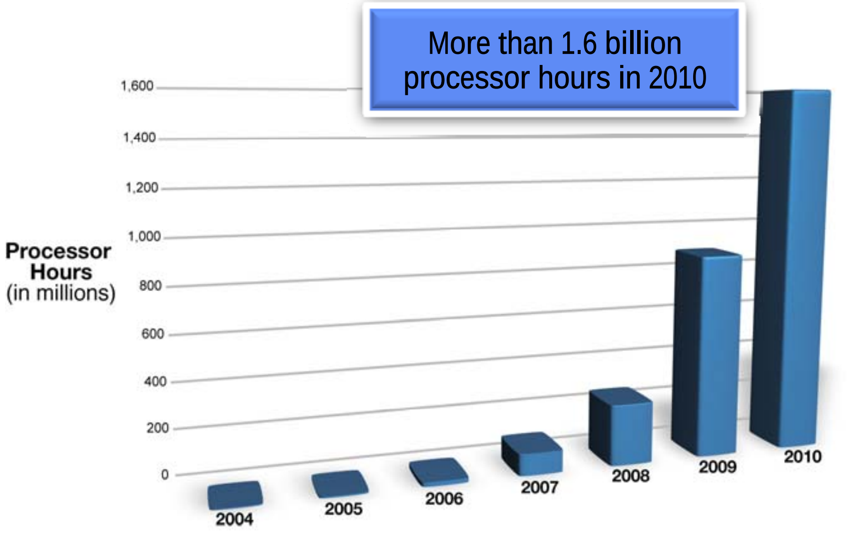
More than 1.6 billion processor hours in 2010

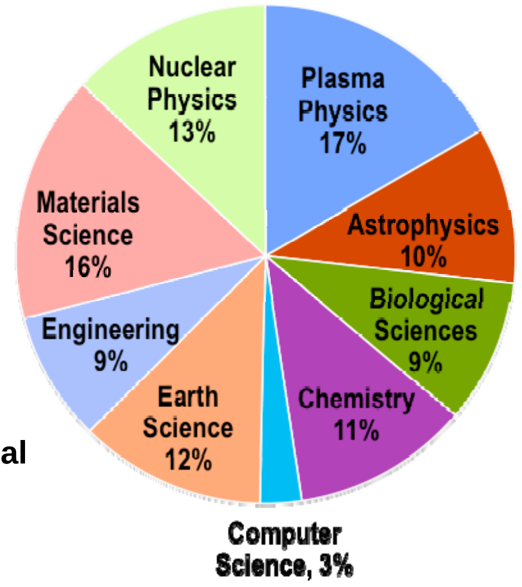

INCITE 2010 awards:
35 new projects and 34 renewal projects

International INCITE Users

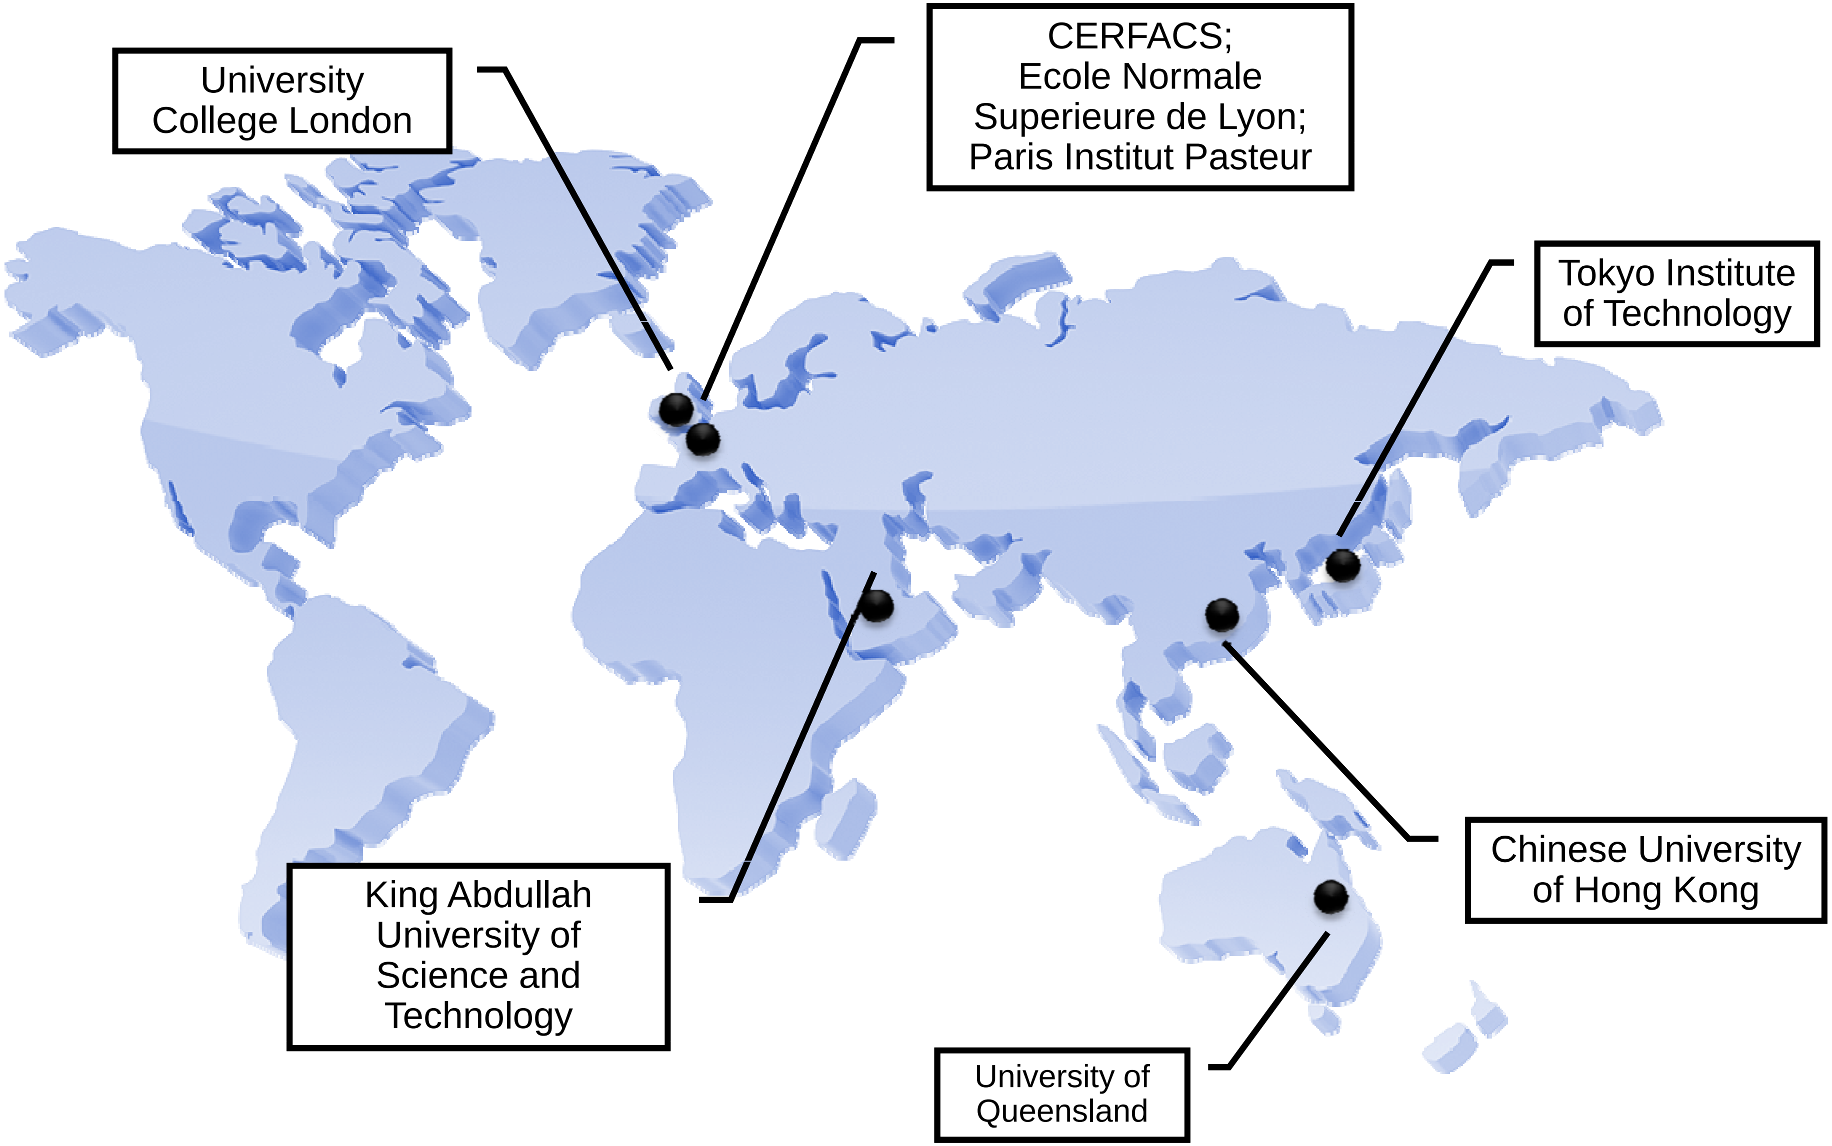
University College London

CERFACS;
Ecole Normale Superieure de Lyon;
Paris Institut Pasteur

Tokyo Institute of Technology

King Abdullah University of Science and Technology

Chinese University of Hong Kong

University of Queensland

Large Eddy Simulation of combustion instabilities in a gas turbine combustion chamber

Thierry Poinsot, CNRS; Gabriel Staffelbach, Marta Garcia, Olivier Vermorel, CERFACS

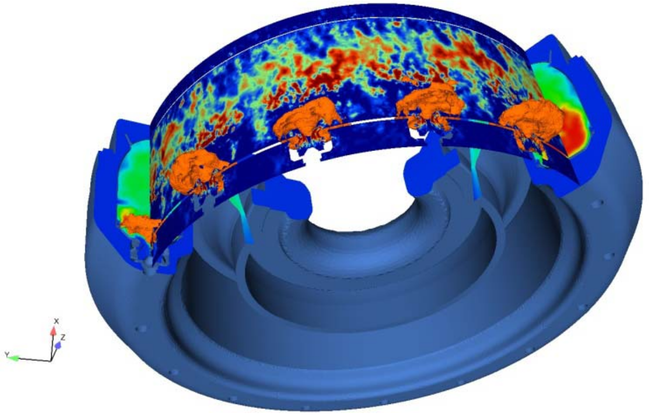

- First simulation of a complete helicopter gas turbine combustion chamber using Large Eddy Simulation of the turbulent reacting flow within the chamber
- Allowed to reproduce and suppress acoustic modes observed in real engines
- Feature paper of *Combustion and Flame* , July 2008.
- Scalability verified up to 12,000 cores

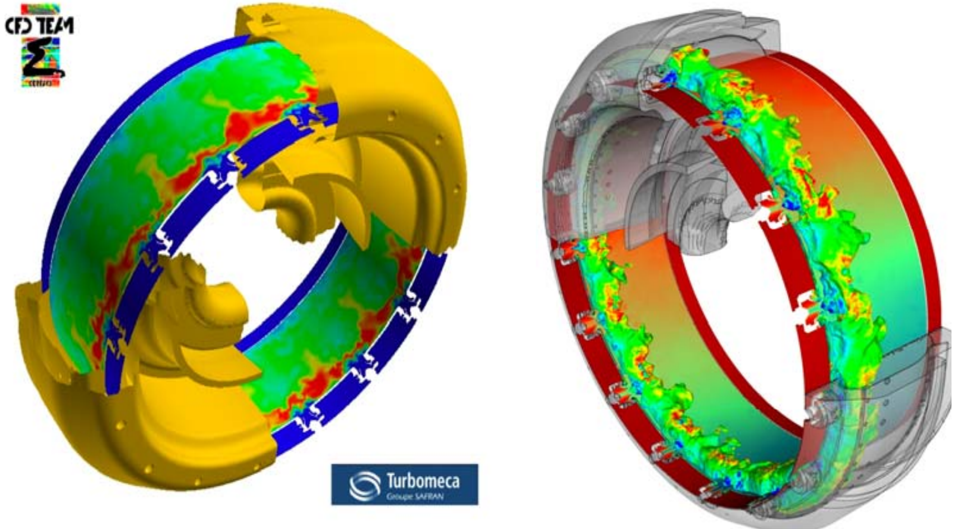

Figure 1: (left) instantaneous temperature field on a cylinder view plane passing through the injectors of the helicopter chamber. Red zones correspond the hot gases and blue zones to cold air coming from the compressor. (right) instantaneous pressure field in the same view plane with isosurfaces of temperature.

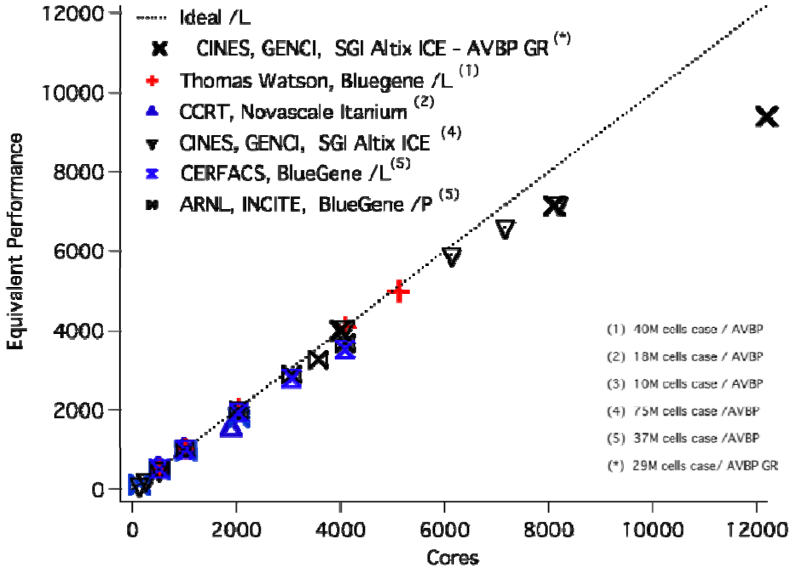

ASCR
Mathematical, Computational, and Computer Science Subprogram

Substantial innovation is needed to provide essential system and application functionality in a timeframe consistent with the anticipated availability of hardware

Providing forefront research knowledge and foundational tools since the early 1950's :

- Applied Mathematics
- Computer Science
- SciDAC
- Next Generation Networking for Science

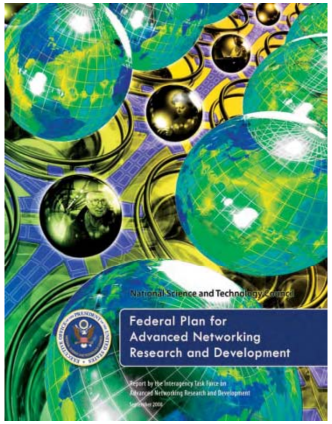

National Science and Technology Council

Federal Plan for
Advanced Networking
Research and Development

Report by the Interagency Task Force on
Advanced Networking Research and Development

September 2006

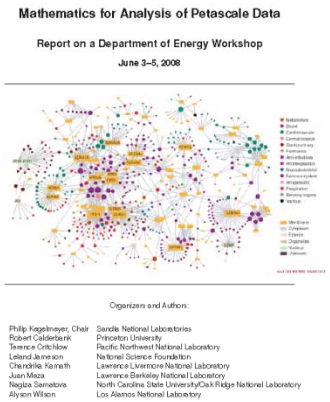

Mathematics for Analysis of Petascale Data

Report on a Department of Energy Workshop

June 3–5, 2008

Organizers and Authors:

Philip Kegelmeyer, Chair	Sandia National Laboratories
Robert Calderbank	Princeton University
Terence Critchlow	Pacific Northwest National Laboratory
Leland Jameson	National Science Foundation
Chandrika Kamath	Lawrence Livermore National Laboratory
Juan Meza	Lawrence Berkeley National Laboratory
Nagiza Samatova	North Carolina State University/Oak Ridge National Laboratory
Alyson Wilson	Los Alamos National Laboratory

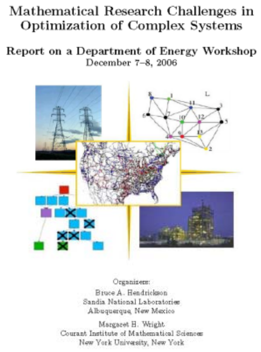

Mathematical Research Challenges in
Optimization of Complex Systems

Report on a Department of Energy Workshop
December 7–8, 2006

Organizers:
Bruce A. Hendrickson
Sandia National Laboratories
Albuquerque, New Mexico

Margaret H. Wright
Courant Institute of Mathematical Sciences
New York University, New York

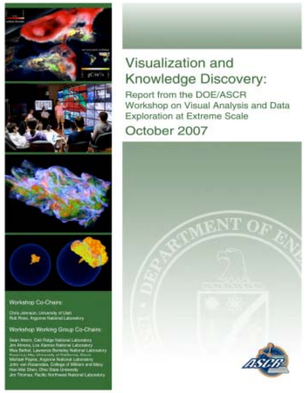

Visualization and
Knowledge Discovery:
Report from the DOE/ASCR
Workshop on Visual Analysis and Data
Exploration at Extreme Scale
October 2007

Workshop Co-Chairs:

Chris Johnson, University of Utah
Rob Ross, Argonne National Laboratory

Workshop Working Group Co-Chairs:

Sean Ahern, Oak Ridge National Laboratory
Jim Ahrens, Los Alamos National Laboratory
Wes Bethel, Lawrence Berkeley National Laboratory
Kwan-Liu Ma, University of California, Davis
Michael Papka, Argonne National Laboratory
John van Rosendale, College of William and Mary
Han-Wei Shen, Ohio State University
Jim Thomas, Pacific Northwest National Laboratory

US Department of Energy
Office of Science

Workshop Report on
Advanced Networking for
Distributed Petascale Science:
R&D Challenges and Opportunities

April 8–9, 2008

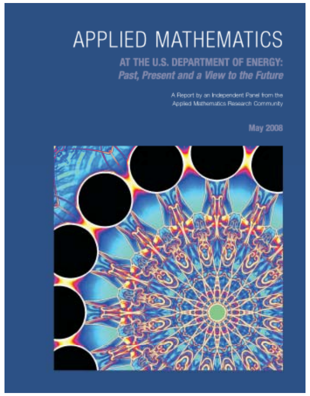

APPLIED MATHEMATICS

AT THE U.S. DEPARTMENT OF ENERGY:
Past, Present and a View to the Future

A Report by an Independent Panel from the
Applied Mathematics Research Community

May 2008

Applied Math Research

Fast Multipole Methods (FMM)

- Adaptive**,** multi-scale methods that reduce the cost of many-body interactions from $O(N^2)$ to $O(N \log N)$ with user-controlled precision.
- Applicable to electromagnetics, gravitation, CFD, acoustics, elasticity, heat transfer, chemistry
- Provide orders of magnitude speedup over direct methods
- Integrated into numerous commercial codes in aerospace, automotive, semiconductor and chemical industries

Impact

- Greengard: “Fast Algorithms for Classical Physics”, *Science,* Vol 265, Aug 1994. Applicable to astrophysics, fluid dynamics and computational chemistry
- FMM dramatically reduced the computational requirements of the electronic quantum Coulomb problem enabling ab initio quantum chemistry calculations of large molecules (*Science*, Vol 271, Jan & Feb 1996)
- Widespread and active research in massively parallel adaptive fast-multipole method on heterogeneous architectures (2009)

APPLICATIONS:

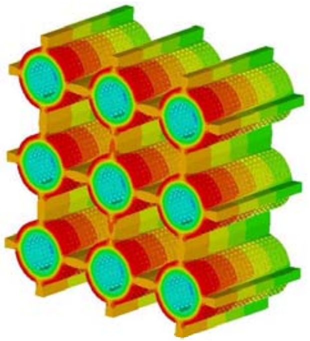

Thermal analysis of fuel cells

(Y. Liu)

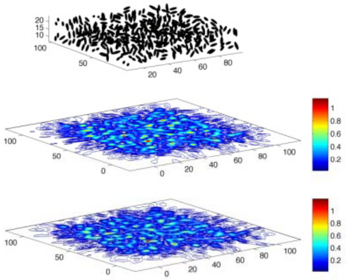

Electromagnetics of microstructured materials

(Gimbutas, Greengard)

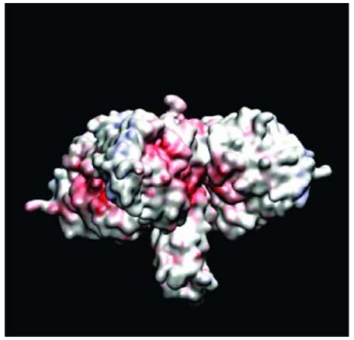

Molecular electrostatics

(Lu, Cheng, Huang, McCammon)

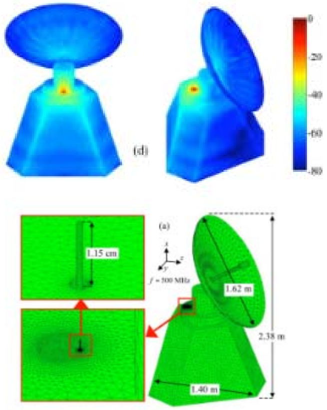

Radiation from multiscale structure

(E. Michielssen et al.)

Applied Math Research

Adaptive Mesh Refinement (AMR):

- Set of algorithms using recursively nested block-structured refinement to concentrate computational resources where needed
- Applicable to wide variety of time-dependent physical phenomena
- Mesh length and time scales automatically adapted to local features
- Allows orders of magnitude increase in “effective” resolution at a fraction of cost of uniform refinement
- Robust software: open source codes and frameworks freely available

Impact

- Berger and Colella, “Local Adaptive Mesh Refinement for Shock Hydrodynamics”, *Journal of Computational Physics,* May,1989.
- Used extensively in fluid dynamics, combustion, magnetohydrodynamics, astrophysics, and other science and engineering applications
- 2004 Sidney Fernbach Award Recipient, Marsha Berger. AMR is “*considered to be one of the seminal ideas in numerical partial differential equations”*
- Very active research area supporting massively parallel computations (2009)

APPLICATIONS:

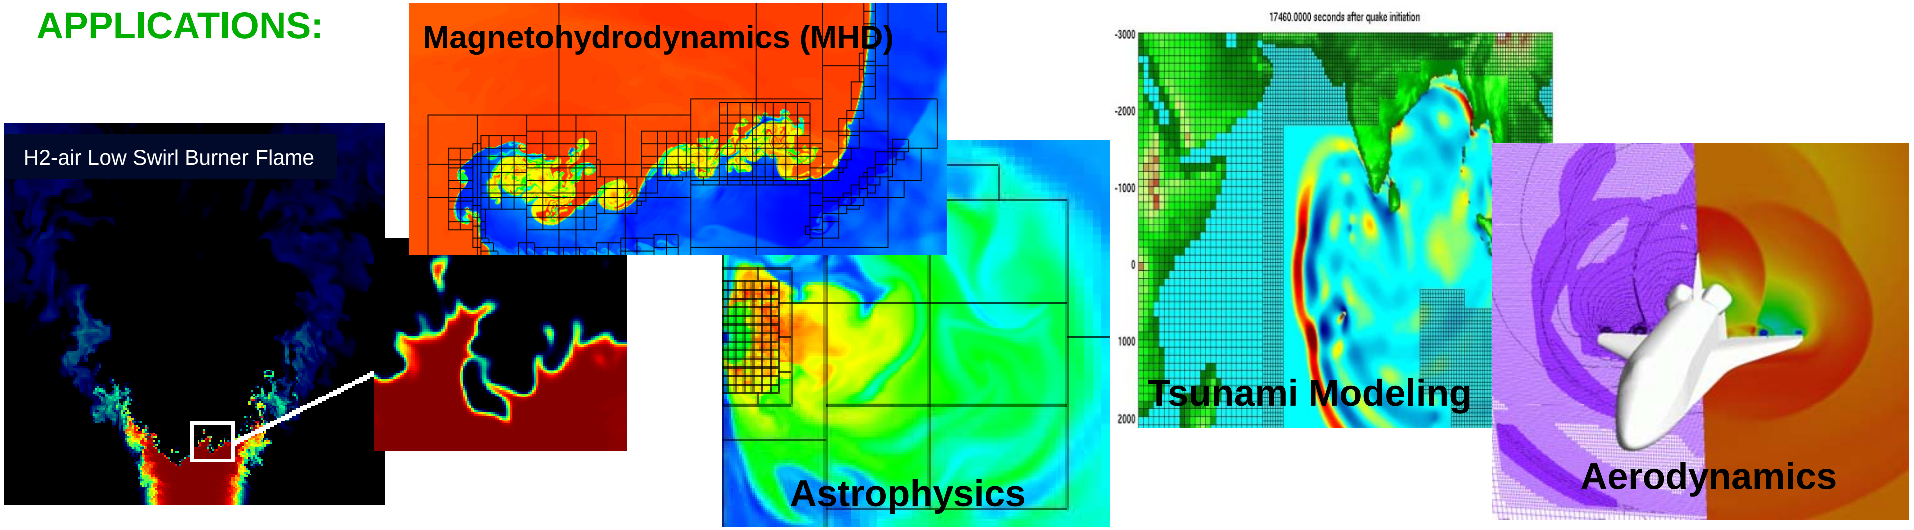

Marsha Berger, Courant Institute, NYU & Philip Colella, Lawrence Berkeley National Laboratory

Applied Math Research

Level Set Methods and Fast Marching Methods

- General mathematical /computational framework for tracking moving interfaces separating different regions
- Ideal for interfaces driven by complex physics and chemistry
- Particularly powerful for moving boundaries that break, merge, and distort
- Key idea: Recast moving interface as higher-dimensional "Hamilton-Jacobi" equation, solve using advanced PDE-schemes

Scientific, engineering and industrial applications:

Molecular Pathways in Chemical Systems

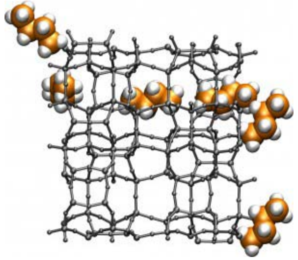

Robotic navigation to detect accessible pathways

Semiconductor Manufacturing

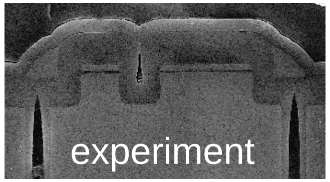

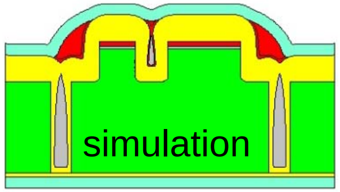

Evolving metal deposition interface

Seismic Imaging Oil Recovery

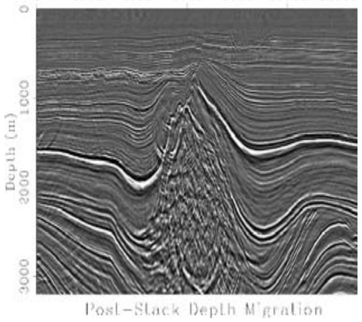

Deep earth imaging using moving interface velocity reconstruction

Biological Imaging

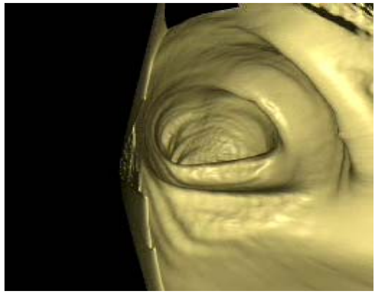

(movie)

Automatic reconstruction from CT scans: Virtual colonoscopy

J.A. Sethian (LBNL, UC Berkeley: www.math.lbl.gov/~sethian)

Portable Programming With MPI and MPICH

- Problem
  - Before MPI, development of parallel programs was stalled; application writers could not commit to a moving target approach to programming.
- Solution
  - Computer scientists worked with parallel computer vendors and application developers defined a standard programming interface: MPI (Message Passing Interface).
  - Argonne computer scientists developed the first complete implementation, MPICH, helping to promote adoption of the standard.
  - DOE support over the last 15 years has enabled MPICH to scale to larger and larger machines, allowing applications to scale as well.

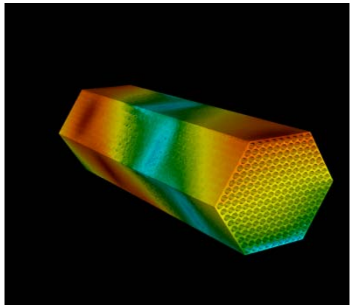

- Impact
  - Nearly all large-scale parallel scientific applications, in all areas of computational science, are written either for MPI directly or for a library in turn implemented in MPI.
  - 14 of the 15 largest machines in the world run MPICH

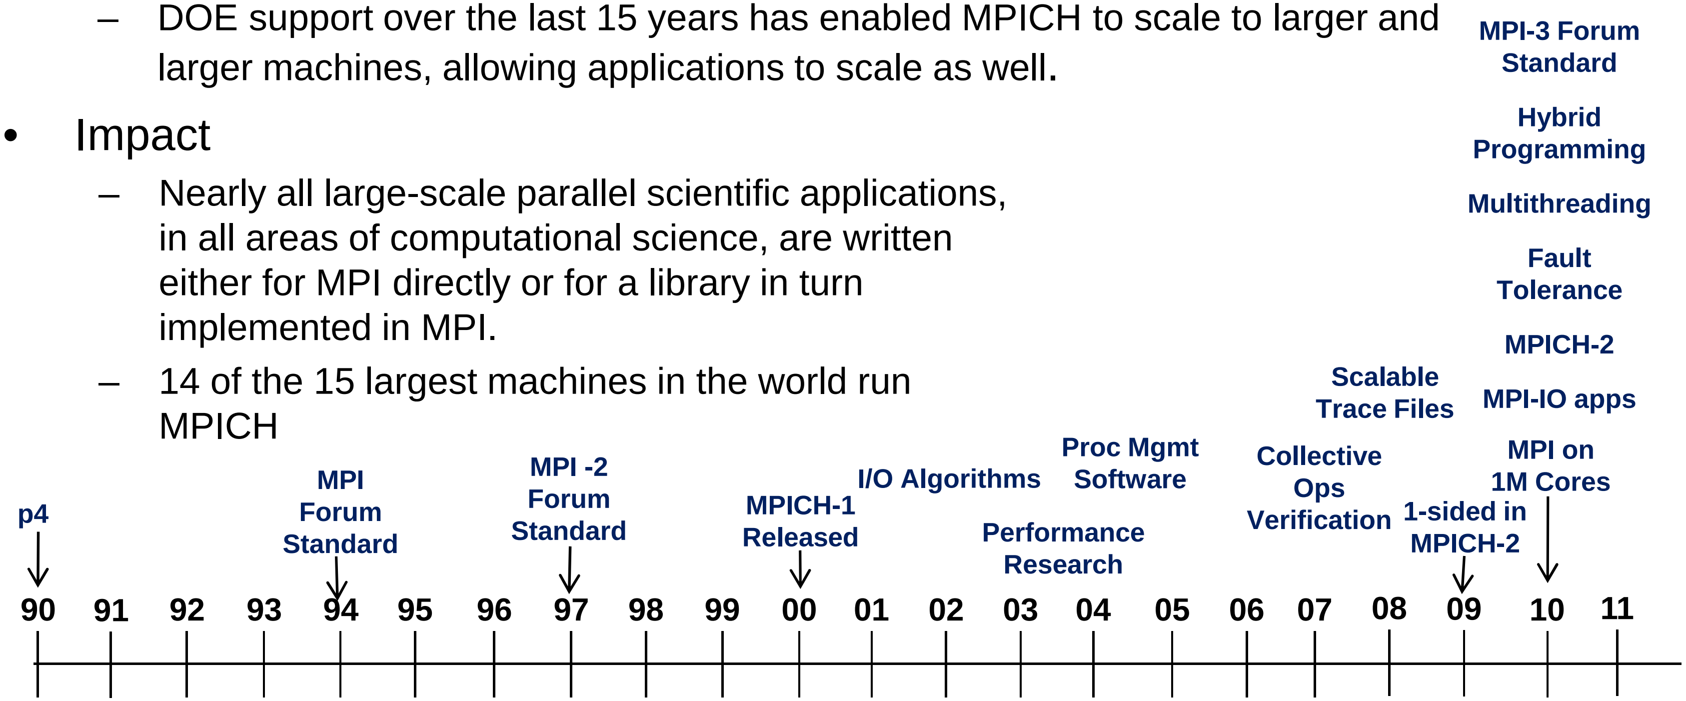

FastBit - Efficient Search Technology for Data Driven Science

- **Problem**
  - Quickly find records satisfying a set of user-specified conditions in a large, complex data set
  - Example: High-energy physics data –find a few thousand events based on conditions on energy level and number of particles in billions of collision events, with hundreds of variables,
- **Solution**
  - Developed new indexing techniques and a new compression method for the indexes, achieved 10-100 fold speedup compared with existing methods
  - Efficient software implementation: available open source from http://sdm.lbl.gov/fastbit/ (1000s of downloads), received a R&D 100 Award
- **Impact**
  - Laser Wakefield Particle Accelerator data analysis: FastBit acts as an efficient back-end for a visual analytics system, providing information for identifying and tracking particles
  - Combustion data analysis: FastBit identifies ignition kernels based on user specified conditions and tracks evolution of the regions
  - Testimonial “FastBit is at least 10x, in many situations 100x, faster than current commercial database technologies” – Senior Software Engineer, Yahoo! Inc

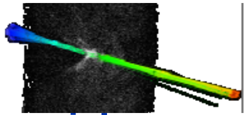

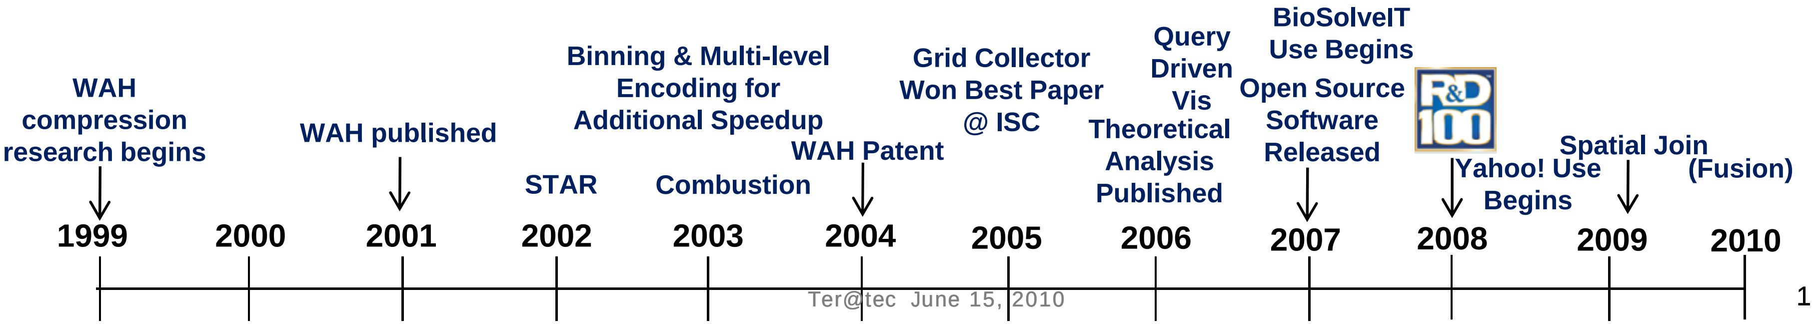

WAH compression research begins

WAH published

STAR

Binning & Multi-level Encoding for Additional Speedup

Combustion

WAH Patent

Grid Collector Won Best Paper @ ISC

Query Driven Vis

Theoretical Analysis Published

BioSolveIT Use Begins

Open Source Software Released

Yahoo! Use Begins

Spatial Join

(Fusion)

1999 2000 2001 2002 2003 2004 2005 2006 2007 2008 2009 2010

U.S. DEPARTMENT OF ENERGY

Scientific Discovery through Advanced Computing

SciDAC – Building Community of Computational Scientists

- **Started in 2001; Redesigned and re-competed in 2006**
  - **Centers for Enabling Technologies (CETs)** – Research and development in mathematical and computing systems software to support SciDAC application.
  - **SciDAC Institutes** are university-led centers of excellence intended to complement CETs with concentrated efforts and a focus on outreach and training.
  - **Science Application Partnerships (SAPs)** provide support for multidisciplinary interactions among application domains, computer science and applied mathematicians.
  - **The SciDAC Outreach Center** provides a collection point for SciDAC technologies and capabilities intended for industry, academia and other scientific communities.
  - **SciDAC Conference** provides an opportunity of computational science community to come together.
  - **SciDAC Review** highlights accomplishments of computational science community especially SciDAC and INCITE.

AMR Simulations of Pellet Injection in Tokamaks

Ravi Samtaney, PPPL; Phil Colella, LBNL

- Injection of frozen hydrogen pellets is a viable method of fueling a tokamak (e.g. ITER)
- Developed an AMR MHD code to simulate pellet injection
- Mass deposition during pellet injection: large scale MHD driven which qualitatively reproduces experimentally observed results that HFS (high-field side) pellet launches are more efficient that LFS (low field side) injection.

AMR meshes during pellet injection

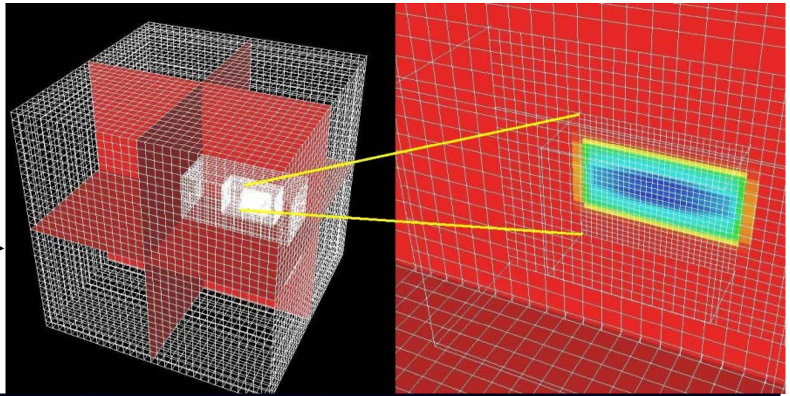

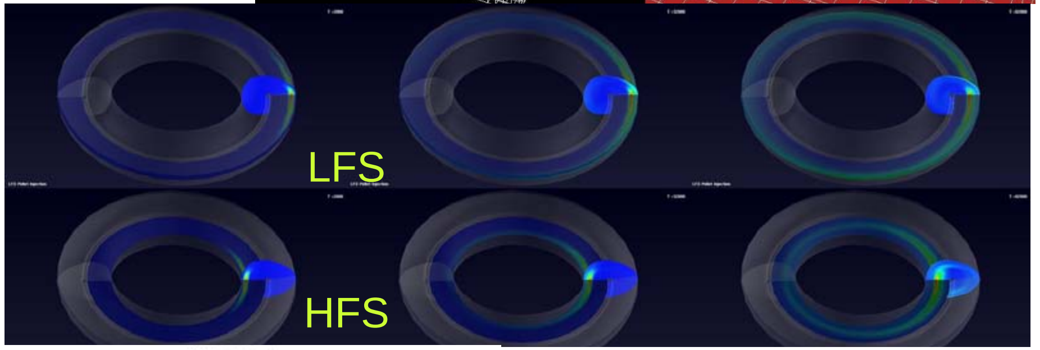

Average density profiles during HFS and LFS pellet injection (Arrow indicated pellet location)

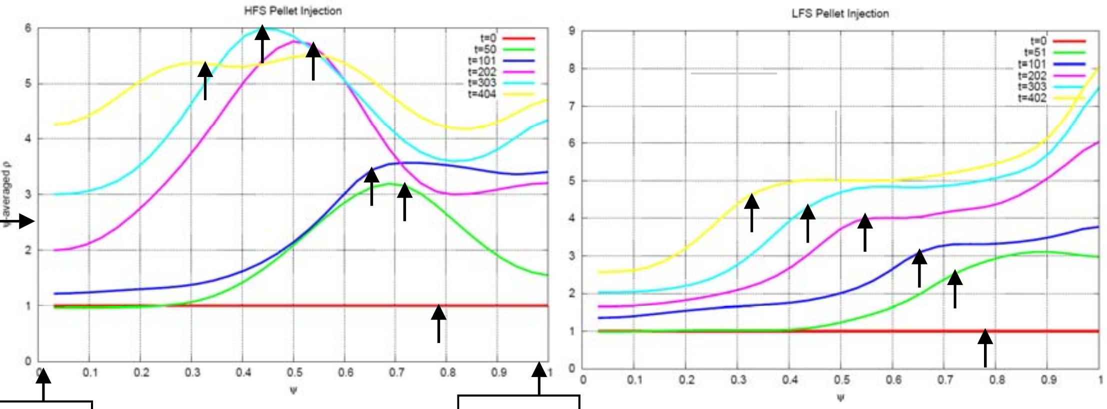

Core

Edge

U.S. DEPARTMENT OF ENERGY

Hybrid quantum simulations of biomolecules

M. Hodak, R. Chisnell, S. Wang, W. Lu, J. Bernholc, NCSU

- **New method for quantum simulations of biomolecules in solution was developed by combining Kohn-Sham density functional theory (DFT) for the biomolecule and its first solvation shells with orbital-free DFT for distant water molecules.**
- **This method is particularly suitable for studying transition-metal (TM) interactions with proteins, where a full quantum treatment is necessary. TMs govern crucial biological processes and are key in anti-cancer drugs.**
- **The numerical methods utilize multigrid approach, domain decomposition, and a linear-scaling orbital-free implementation. This results in highly-scalable parallel code, able to perform quantum calculations of unprecedented sizes, and to include thousands of water molecules at less than 10% of the total simulation cost.**
- **The first application, just published in PNAS, investigated a link between copper and the normal function of prion protein. It was shown that:**
  - **Normal prion protein can function as a copper buffer, protecting human tissue from damage.**
  - **As copper ions bind to the PrP, its structure changes, becomes more stable and more resistant to misfolding, which could prevent the Cruetzfeldt-Jakob disease in humans or “mad cow” disease in cattle.**

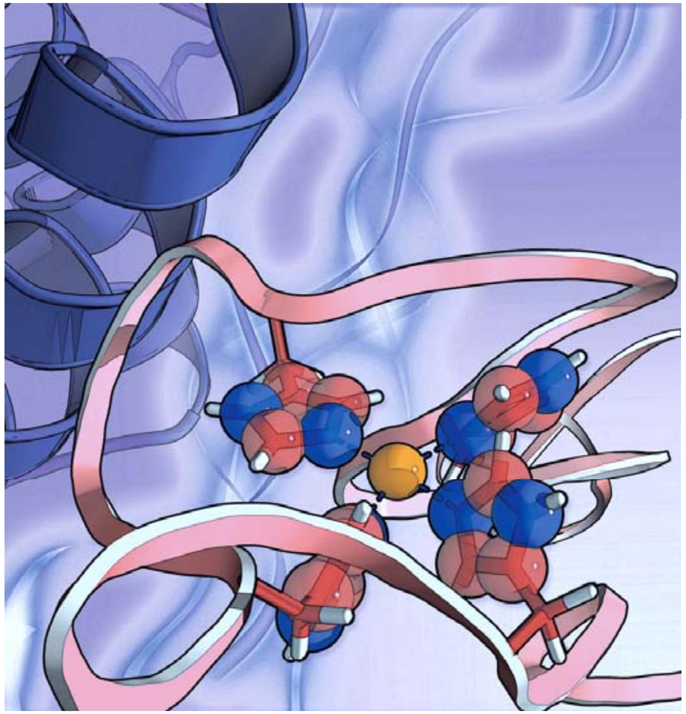

The prion protein, with copper bound to its unstructured part in the foreground, and its folded portion containing alpha helices in the background.

SciDAC funding and INCITE award

Solving CEBAF BBU Using Shape Uncertainty Quantification Method

V. Akcelik, K. Ko, L. Lee, Z. Li, C. Ng, L. Xiao, SLAC, F. Marhauser, C. Reece, R. Rimmer
and H. Wang, TJLAB; E. NG, LBNL

SciDAC Success as a Collaboration between Accelerator Simulation, Computational Science and Experiment – Beam Breakup (BBU) instabilities at well below the designed beam current were observed in the CEBAF12 GeV upgrade of the Jefferson Lab (TJNAF) in which Higher Order Modes (HOM) with exceptionally *high* quality factor (Q) were measured. Using the shape uncertainty quantification tool developed under SciDAC, the problem was found to be a deformation of the cavity shape due to fabrication errors. This discovery was achieved as a team effort between SLAC, TOPS, and JLab which underscores the importance of the SciDAC multidisciplinary approach in tacking challenging applications

Method of Solution - Using the measured cavity parameters as inputs, the deformed cavity shape was recovered by solving the ***inverse*** problem through an optimization method. The calculations showed that the cavity was 8 mm shorter than designed, which was subsequently confirmed by measurements. The result explains why the troublesome modes have high Qs because in the deformed cavity, the fields shift away from the HOM coupler where they can be damped. This shows that quality control in cavity fabrication can play an important role in accelerator performance. .

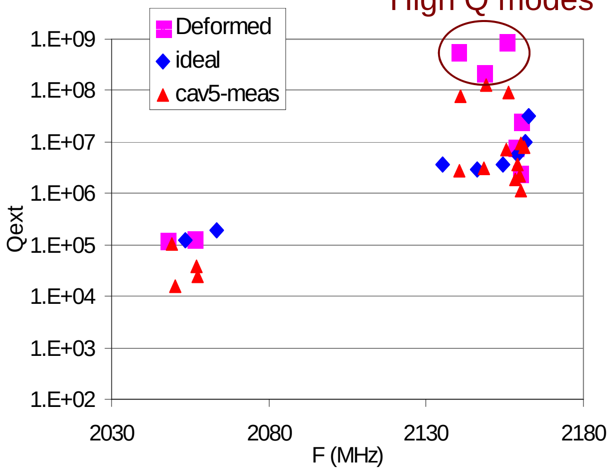

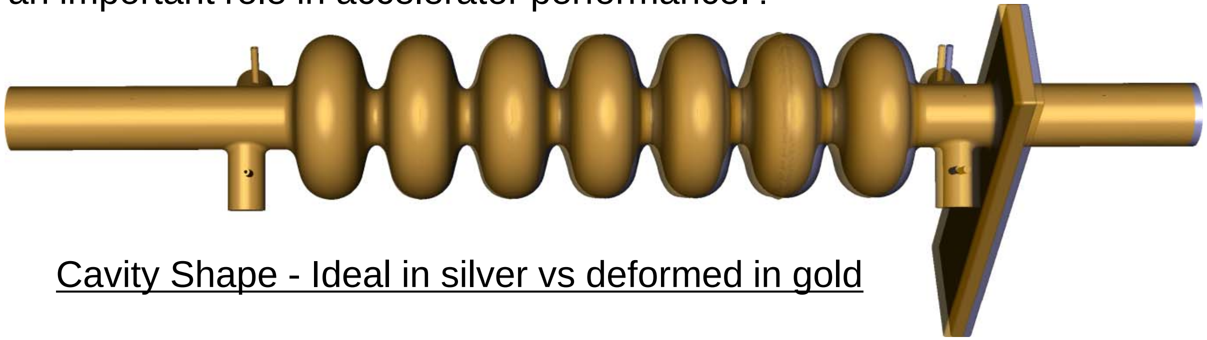

Cavity Shape - Ideal in silver vs deformed in gold

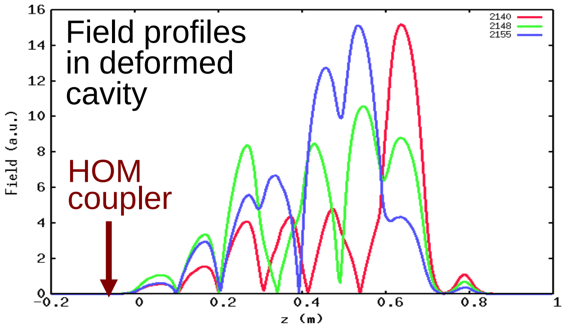

Leadership Computing: Scientific Progress at the Petascale

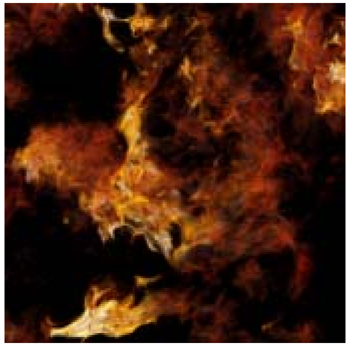

Turbulence

Understanding the statistical geometry of turbulent dispersion of pollutants in the environment.

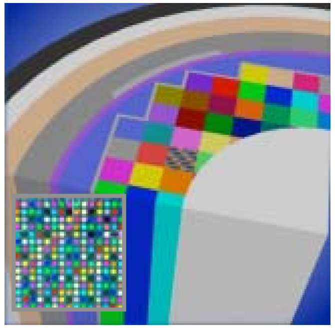

Nuclear Energy

High-fidelity predictive simulation tools for the design of next-generation nuclear reactors to safely increase operating margins.

Energy Storage

Understanding the storage and flow of energy in next-generation nanostructured carbon tube supercapacitors

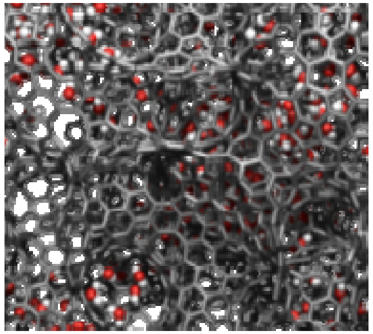

Fusion Energy

Substantial progress in the understanding of anomalous electron energy loss in the National Spherical Torus Experiment (NSTX).

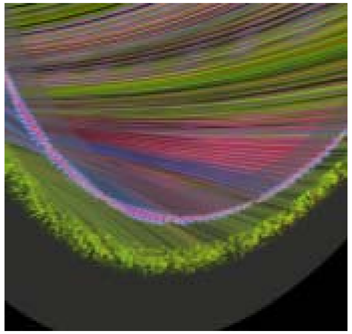

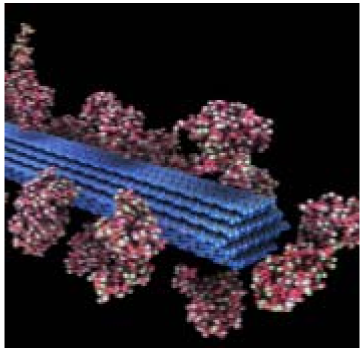

Biofuels

A comprehensive simulation model of lignocellulosic biomass to understand the bottleneck to sustainable and economical ethanol production.

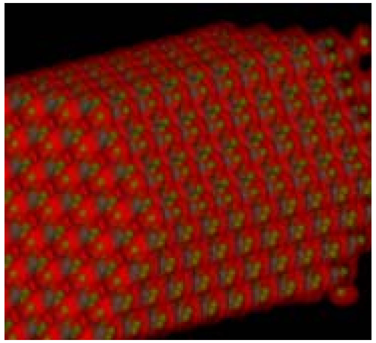

Nano Science

Understanding the atomic and electronic properties of nanostructures in next-generation photovoltaic solar cell materials.

All known sustained petascale science applications to date have been run on OLCF system

DOE Exascale

Goal: Provide the next generation of extreme scale computing capability to solve problems of importance in Energy, the Environment, Security, and Science

Exascale Applications and Technology

- **Town Hall Meetings April-June 2007**
- **Scientific Grand Challenges Workshops November 2008 – October 2009**
  - **Climate Science (11/08),**
  - **High Energy Physics (12/08),**
  - **Nuclear Physics (1/09),**
  - **Fusion Energy (3/09),**
  - **Nuclear Energy (5/09) (with NE)**
  - **Biology (8/09),**
  - **Material Science and Chemistry (8/09),**
  - **National Security (10/09) (with NNSA)**
- **Cross-cutting workshops**
  - **Architecture and Technology (12/09)**
  - **Architecture, Applied Mathematics and Computer Science (2/10)**
- **Meetings with industry (8/09, 11/09)**
- **External Panels**
  - **ASCAC Exascale Charge (FACA)**
  - **Trivelpiece Panel**

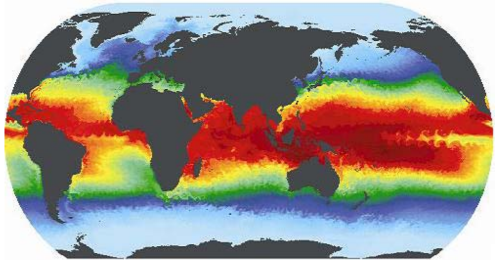

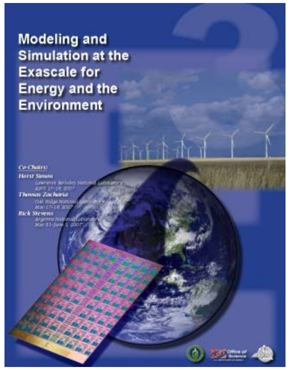

Modeling and Simulation at the Exascale for Energy and the Environment

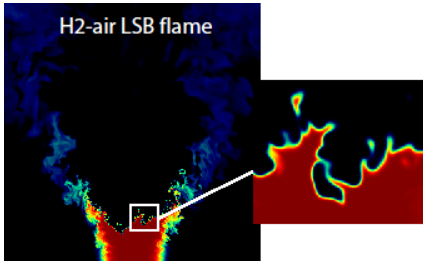

MISSION IMPERATIVES

Scientific Grand Challenges
FOREFRONT QUESTIONS IN NUCLEAR SCIENCE AND THE ROLE OF COMPUTING AT THE EXTREME SCALE

Scientific Grand Challenges
CHALLENGES FOR UNDERSTANDING THE QUANTUM UNIVERSE AND THE ROLE OF COMPUTING AT THE EXTREME SCALE

Scientific Grand Challenges
CHALLENGES IN CLIMATE CHANGE SCIENCE AND THE ROLE OF COMPUTING AT THE EXTREME SCALE
November 6-7, 2008 • Washington D.C.

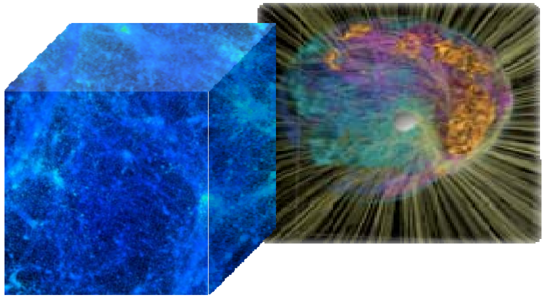

FUNDAMENTAL SCIENCE

Exascale: Major Components

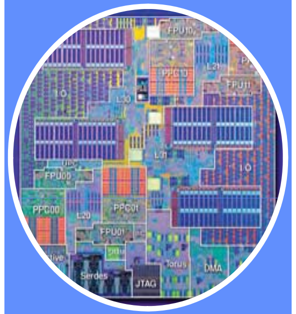

Platform R&D

2 Vendor Tracks

- Power
- Integration
- Risk Mitigation

Critical Technologies

(everyone benefits)

- Memory
- Nonvolatile storage
- Optics

Software and Environments

- Operating environment
- Systems Software
- System reliability
- Programming model

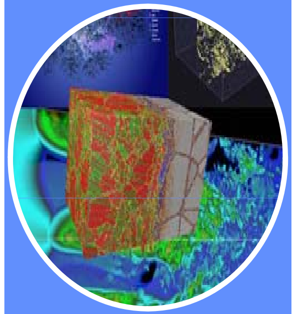

Co-design

- Applied Mathematics
- Performance models
- Simulators
- Applications integration with vendors

Platforms

- Early prototypes to ensure component integration and usefulness
- Risk mitigation for vendors – Non recoverable engineering cost

Exascale Initiative

Potential System Architectures

Systems	2009	2021 +1/-0	2021 +1/-0
System peak	**2 Petaflop**	**1 Exaflop**	**1 Petaflop**
Power	**6 MW**	**~20 MW**	**~20 kW**
System memory	0.3 PB	64 PB (+)	.1PB
Node performance	125 GF	1-2 or 10TF	1-2 or 10TF
Node memory BW	25 GB/s	2-4TB/s	2-4TB/s
Node concurrency	12	O(1k) or 10k	O(1k) or 10k
Total Node Interconnect BW	3.5 GB/s	200-400GB/s (1:4 or 1:8 from memory BW)	~1 GB/s (1:4 or 1:8 from memory BW)
System size (nodes)	18,700	O(1M) or O(100,000)	O(1,000) or O(100)
Total concurrency	225,000	O(billion) * O(10) to O(100) for latency hiding	O(million) * O(10) to O(100) for latency hiding
Storage	15 PB	500-1000 PB (>10x system memory is min)	1-15 PB (>10x system memory is min)
IO	0.2 TB/s	60 TB/s (how long to drain the machine)	.1 TB/s (how long to drain the machine)
MTTI	days	O(1 day)	days

Exascale for the Environment: Climate

"The climate problem has an insatiable appetite for computing. The realistic representation of the Earth system is paced by computing." -- James Hack, ORNL

- **Validate Models**
  - **Increasing sensitivity for comparison to observation**
- **Push limits of Predictive Skill**
  - **Resolve time and distance scales relevant for decision makers**
- **Build Bio-geochemistry into simulations**
  - **Include ecological interactions between nitrogen and carbon**
- **Analyze new processes**
  - **Global cloud system resolving models**
- **Enable decadal forecasts**
  - **Decision support for mitigation or adaptation**
- **Explore tipping points**
  - **Abrupt climate change due to feedbacks**
- **Assess geoengineering**
  - **Deliberate manipulation of climate system to counteract consequences of GHG emissions**

The road to exascale will enable this important policy area to have a robust scientific underpinning

Exascale for Energy: Combustion

Develop a validated, predictive, multiscale, combustion modeling capability that can optimize the design and operation of evolving fuels in advanced engines and power plants.

Thousands of design iterations – each corresponding to a high-fidelity multiscale simulation – would accelerate optimization and implementation of new technologies.

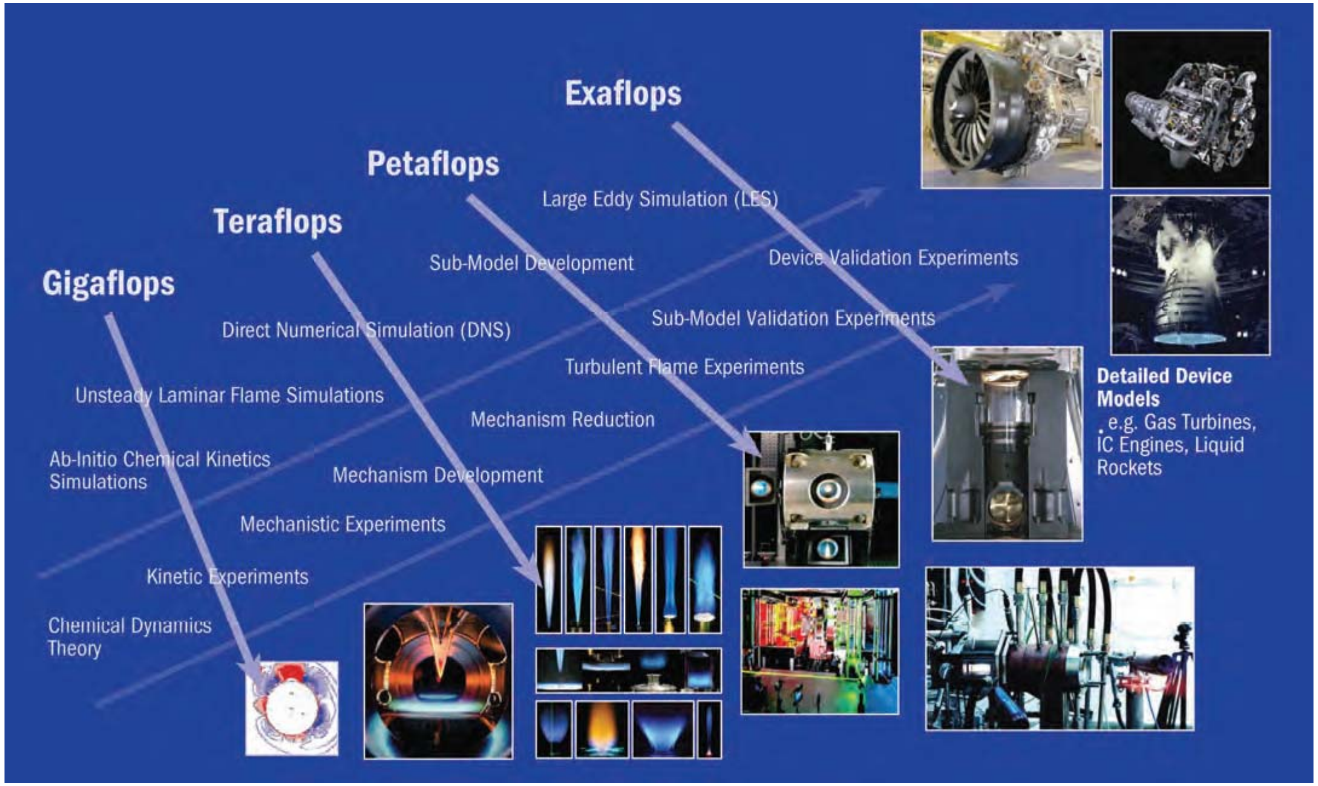

The road to exascale will enable robust simulations in this important area and will directly impact US competitiveness, jobs, the economy, and the policy landscape.

Exascale for Energy: Nuclear

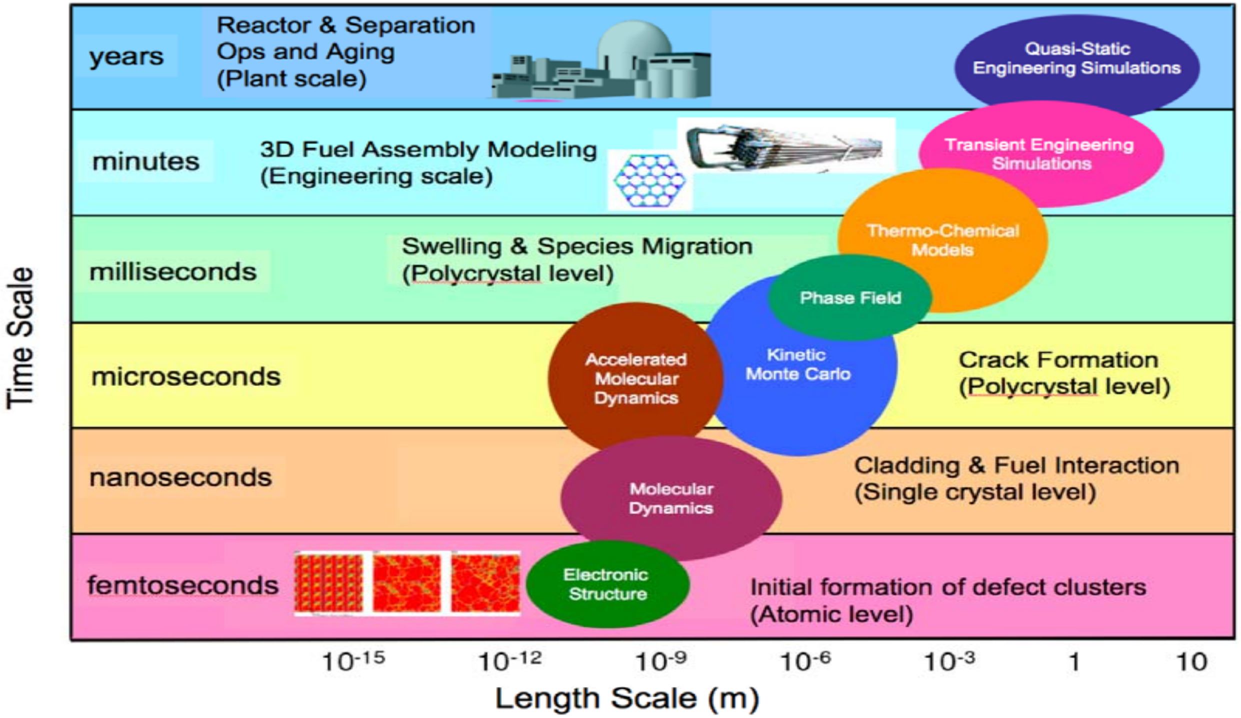

Bridging length and time scales to resolve scientific unknowns [in nuclear energy] will require 3D simulations 100x standard resolution = A 10 Exaflop problem.

Science-Based, Nuclear Energy Systems Enabled by Advanced Modeling and Simulation at the Extreme Scale

The road to exascale will:

- Accelerate the iteration cycle of technology evaluation, design, engineering, and testing to optimize existing and new nuclear energy applications
- Shorten the licensing process by providing reliably predictive integrated performance models that reduce uncertainties
- Reduce construction and operations costs while also reducing uncertainty and risk

Different approaches to on-chip clustering

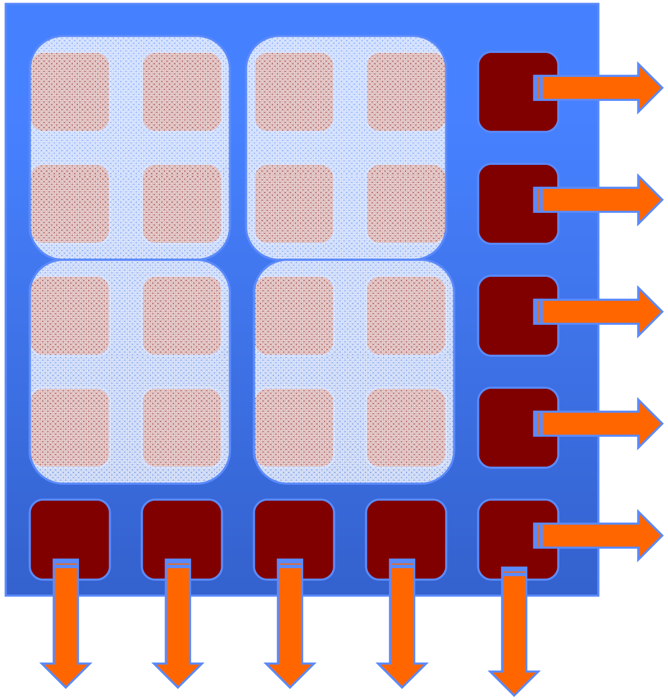

- Cost of moving long-distances on chip motivates clustering on-chip
  - 1mm costs ~6pj (today & 2018)
  - 20mm costs ~120 pj (today & 2018)
  - FLOP costs ~100pj today
  - FLOP costs ~25pj in 2018

- Different Architectural Directions
  - GPU: WARPs of hardware threads clustered around shared register file
  - CMP: limited area cache-coherence
  - CMT: hardware multithreading clusters

Power Consumption

- **Barriers**
  - **Power is leading design constraint for computing technology**
  - **Target ~20MW, estimated > 100MW required for Exascale systems (DARPA, DOE)**
  - **Efficiency is industry-wide problem (IT technology >2% of US energy consumption and growing)**
- **Technical Focus Areas**
  - **Energy efficient hardware building blocks (CPU, memory, interconnect)**
  - **Novel cooling and packaging**
  - **Si-Photonic Communication**
  - **Power Aware Runtime Software and Algorithms**
- **Technical Gap**
  - **Need 5X improvement in power efficiency over projections that include technological advancements**

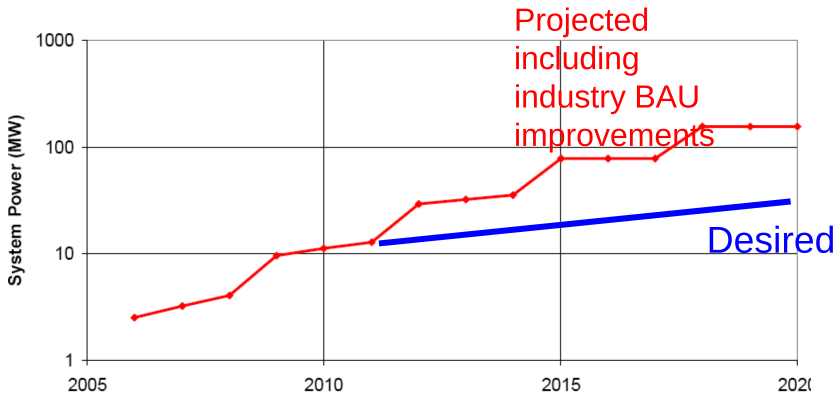

Possible Leadership class power requirements
From Peter Kogge (on behalf of Exascale Working Group), “Architectural *Challenges* at the Exascale Frontier”, June 20, 2008

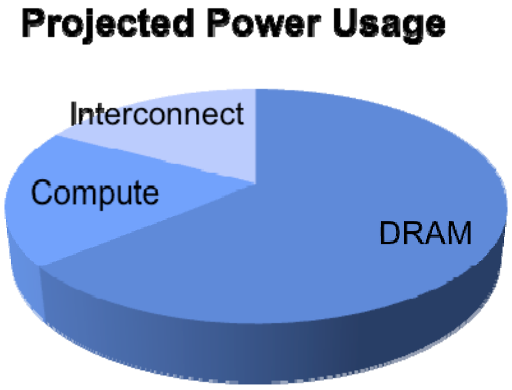

System memory dominates energy budget

Memory and Storage Bandwidth

- Barriers
  - **Per-disk performance, failure rates, and energy efficiency no longer improving**
  - **Linear extrapolation of DRAM vs. Multi-core performance means the height of the memory wall is accelerating**
  - **Off-chip bandwidth, latency throttling delivered performance**
- Technical Focus Areas
  - ***Efficient Data Movement***
    - **Photonic DRAM interfaces**
    - **Optical interconnects / routers**
    - **Communications optimal algorithms**
  - ***New Storage Approaches***
    - **Non-volatile memory gap fillers**
    - **Advanced packaging (chip stacking)**
    - **Storage efficient programming models (Global Address Space)**
- Technical Gap
  - **Need 5X improvement in memory access speeds to keep current balance with computation.**

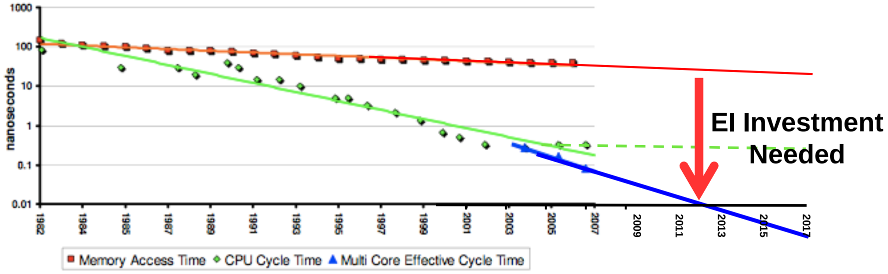

Figure 6.12: CPU and memory cycle time trends.

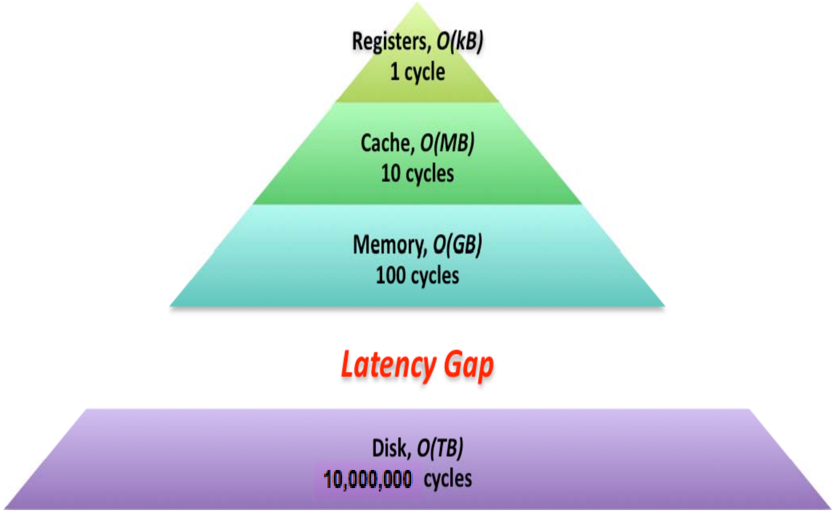

Reliability and Resilience

- **Barriers**
  - Number of system components increasing faster than component reliability
  - Mean time between failures of minutes or seconds for exascale
  - Silent error rates increasing
  - No job progress due to fault recovery if we use existing checkpoint/restart
- **Technical Focus Areas**
  - Improved hardware and software reliability
    - Better Reliability, Availability and Serviceability (RAS) collection and analysis (root cause)
    - Greater integration
  - Fault resilient algorithms and applications
  - Local recovery and migration
- **Technical Gap**
  - Need 1000X improvement in MTTI so that applications can run for many hours. Goal is 10X improvement in hardware reliability. Local recovery may and migration may yield another 10X. However, for exascale, applications will need to be fault resilient.

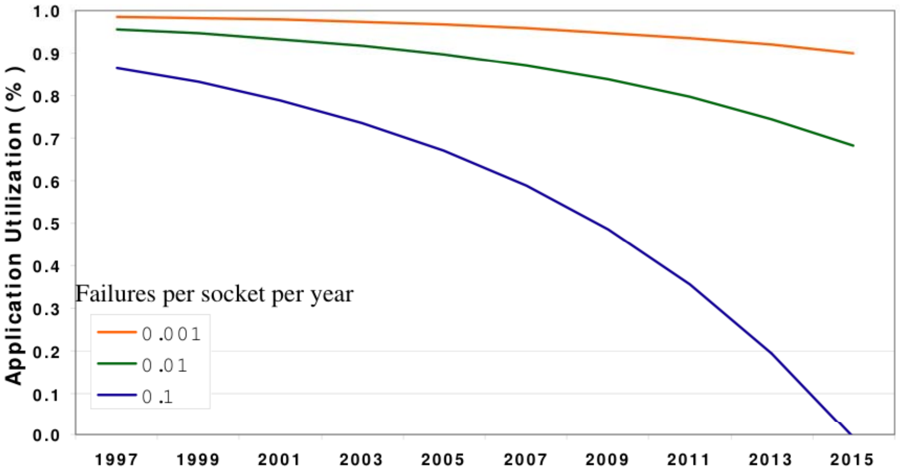

Effective application utilization (including checkpoint overhead) at 3 rates of hardware failure

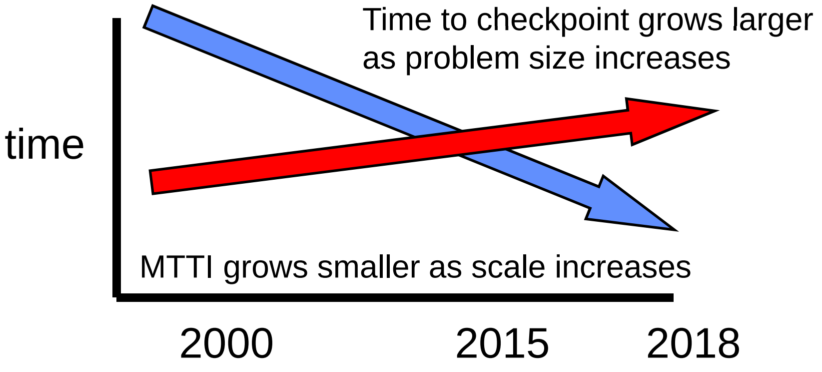

Time to checkpoint grows larger as problem size increases

MTTI grows smaller as scale increases

By exascale checkpoint/restart no longer viable

U.S. DEPARTMENT OF ENERGY

System Software Scalability

- Barriers
  - Fundamental assumptions of system software architecture did not anticipate exponential growth in parallelism
  - Requirements for resilience at scale
  - IO wall reducing effectiveness of simulation environment
- Technical Focus Areas
  - System Hardware Manageability
  - System Software Scalability
  - Applications Scalability
  - Supporting investments in infrastructure to support systems
  - Initial deliveries to validate software and operations path
- Technical Gap
  - 1000x improvement in system software scaling
  - 100x improvement in system software reliability
  - Need application hooks into Reliability, Availability and Serviceability (RAS) system

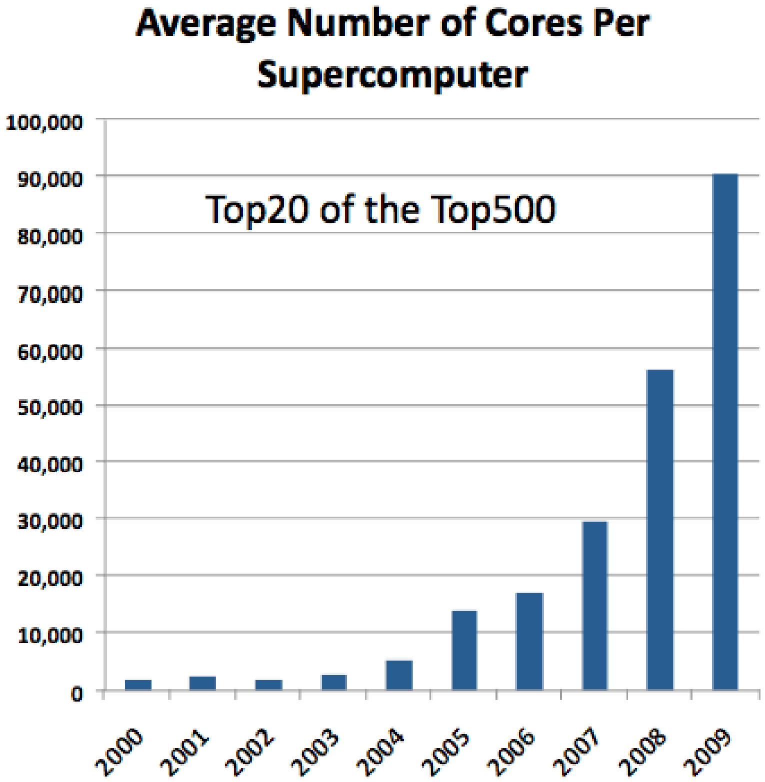

Programming Models and Environments

- **Barriers:** Delivering a complex large-scale scientific instrument that is productive and fast.
  - O(1B) way parallelism in Exascale system
    - Massive lightweight cores for low power
    - Some “full-feature” cores lead to heterogeneity
  - O(1K) way parallelism in a processor
    - Data and independent thread parallelism
  - Effective management of locality
    - Software-managed memory (local store)
    - Effective abstractions for explicitly managed memory hierarchies
    - Communication avoiding algorithms
    - Communication optimized for architecture
  - Complexity of scientific applications
  - Programming for resilience
  - Science goals require complex codes

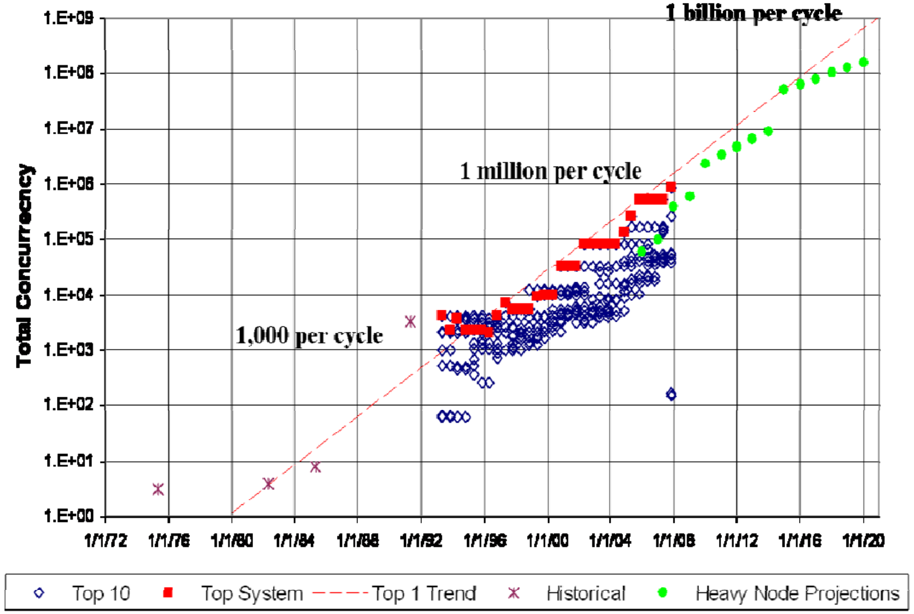

How much parallelism must be handled by the program?

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

- Technology Investments
  - Evolutionary: extend existing models used in science for scalability and to hide system complexity, e.g., heterogeneity and failures
  - Moderate: leverage emerging models in scientific computing
  - Revolutionary: develop a new paradigm for high usability at extreme scales
- **Technical Gap:** Productivity, Performance and Correctness for 1000x more parallelism while increasing programming productivity of scientists by 10x

U.S. DEPARTMENT OF ENERGY

Road to Exascale

“somewhere I have never traveled gladly beyond any experience.”

e.e. cummings