U.S. DEPARTMENT OF ENERGY

ASCR Update

November 3, 2009

Michael Strayer

Associate Director, Advanced Scientific Computing Research

New Positions
Postings Closed

- Research Program Managers
  - – Computer Science
  - – Computer Science
  - – Physical Scientist/SciDAC – SC programs
  - – Physical Scientist/SciDAC – Applied programs

- Facilities Program Manager
  - – Leadership Computing

ASCR Budget Details

	FY 2009 Appropriation	FY 2010 Conference	Change from FY09 to FY 10
Advanced Scientific Computing Research			
Applied Mathematics	40,164	44,850	+4,686
Computer Science	33,618	46,800	+13,182
Computational Partnerships (includes SciDAC)	52,064	53,235	+1,171
Next Generation Networking for Science	14,321	14,321	0
SBIR/STTR	4,038	4,586	+548
Total, Mathematical, Computational, and Computer Sciences Research	***144,205***	***163,792***	***+19,587***
High Performance Production Computing (NERSC)	54,790	55,000	+210
Leadership Computing Facilities	115,000	123,000	+8,000
Research and Evaluation Prototypes	23,900	16,182	-7,718
High Performance Network Facilities and Testbeds (ESnet)	25,000	30,000	+5,000
SBIR/STTR	5,925	6,026	521
High Performance Computing and Network Facilities	***224,615***	***229,776***	***+5,593***
Total, Advanced Scientific Computing Research	**368,820**	**394,000**	**+25,180**

U.S. DEPARTMENT OF ENERGY

President's FY2010 Budget Request
Highlights

- Applied Math
  - Cyber Security research moved from Next Generation Networking
  - Proposed new fellowship program in Applied Math and High performance computer science
- Computer Science
  - New effort in Advanced Computer Architecture design for science
  - Bridges efforts in advanced computer architecture design with ongoing efforts in computer science and applied mathematics to address needs of DOE science applications
- Computational Partnerships
  - Support for interdisciplinary teams focused on transforming critical DOE applications for extreme scale computing
- Facilities
  - Increases support lease payments and site preparation at ANL for proposed upgrade
  - ESnet will begin to deliver 100-400 Gbps to SC laboratories

2009 New Research Initiatives

ISICLES: *Ice Sheet Initiative for CLimate ExtremeS* :

Enable scientific and computational breakthroughs in accurate dynamical representation – with uncertainty quantification -- of ice sheets in climate models

- Eight proposals received and reviewed
- Six projects awarded, $3M/yr for three years

Mathematics for the Analysis of Petascale Data

Research in the mathematics of extracting features from extremely large datasets and understanding these features

- 81 proposals received and reviewed
- 11 projects awarded, $4M/yr for three years

Mathematics for Complex, Distributed Interconnected Systems (Lab only)

Research into the behavior of large-scale complex, distributed, interconnected systems not as decentralized component models but as integrated entities.

- 38 proposals received and reviewed
- 5-7 awards, $3.5M/yr for three years

Joint Math/CS Institute

Collaborative research in applied mathematics and computer science to bridge the gap between large complex scientific applications software and next-generation hardware

- 25 proposals received; 24 proposals reviewed
- Three projects awards, $4M/yr for three years

ISICLES: Ice Sheet Initiative for CLimate ExtremeS

Overarching GOAL: Enable scientific and computational breakthroughs in representation of ice sheet dynamics in HPC climate models to improve reliability of predictions

Research Topics:

- Ice sheets fracture modeling
- Uncertainty Quantification (UQ) for large-scale ice sheet modeling and simulation.
- Scalable solution framework for the Community Ice Sheet Model
- Scalable unstructured ice-sheet solvers and simulation infrastructure
- Adaptive algorithms for ice sheet modeling
- Ice sheet dynamics modeling

Expected Outcome:

- Accurate simulation of ice sheets at ~1km resolution for near-term IPCC report
- Refined UQ techniques to improve credibility of climate projections
- Development of solution frameworks for extreme-scale simulations of ice sheets

"*Dynamics of the slow-moving ice and of ice shelves are reasonably well understood and can be modeled adequately, but this is not so for fast-moving ice streams and outlet glaciers… recent changes in ice sheet margins and ice streams cannot be simulated accurately with…* [current] *models…*"

Joint Math/CS Institute

Overarching Goal: Foster collaborative research in applied mathematics and computer science to bridge the gap between large complex scientific applications software and next-generation hardware

Research Topics:

- Research in PDE solvers that better acknowledge memory hierarchy
- Develop multi-precision and architecture-aware algorithms and libraries
- Develop necessary software stack (including runtime environment) to support scalable, resilient performance on multicore architecture
- Mathematical modeling and optimization of the use of I/O resources through the coordination of “productive I/O” (running the application) and “defensive I/O” (checkpoint restart)

Expected Outcome:

- Improved time to solution on DOE applications
- Reduction of the time PDE solvers spend in accessing memory
- Optimized software portable across different platforms
- Better understanding in the I/O usage for fault-tolerance in applications
- Open distribution of software developed

Mathematics for the Analysis of Petascale Data

Overarching Goal: Address the mathematical challenges involved in extracting insights from extremely large datasets ("petascale data") and investigating fundamental issues in finding key features and understanding the relationships between those features.

Research Topics:

- Data analysis with uncertainty
- Rare-event detection
- Machine learning for massive data sets
- Scalable graph decompositions
- Geometric analysis for data reduction
- Scalable statistical analysis

Expected Outcomes:

Mathematical models, tools, and methods for the representation and analysis of petascale data sets, including:

- real-time anomaly identification in streaming and evolving data;
- dimension reduction of data for extracting subsets, features of interest, or low-dimensional patterns;
- rigorous mathematical methodology for combining data of different types and quality

Mathematics for Complex, Distributed Systems

Overarching Goal: Research in mathematical models, methods and tools for the modeling, simulation and analysis of complex, distributed, interconnected systems.

Research Topics:

- Scalable methods for representing, characterizing, and generating large graphs
- Agent-based integrated model for complex networks
- Intrusion detection for high-performance computing
- Stochastic modeling of complex networks
- Risk management and planning of complex networks

Expected Outcomes:

- Analysis of data generated from complex, distributed, interconnected systems for situational awareness;
- Modeling and simulation of the key properties and emergent behavior on large-scale complex distributed interconnected systems;
- Mathematical methods for modeling and analysis of the dynamics and evolution of large-scale, complex, distributed interconnected systems

Facilities Update

- NERSC
  - Quad core upgrade to Franklin accepted June 17, 2009
  - NERSC-6 contract awarded to Cray for at least 1 petaflop Cray XT5
- LCFs Next Generation
  - Mission needed approved January, 2009: “
    - “The upgrade of the Leadership Computing Facilities to tens of petaflops by the 2011-2013 timeframe is vital to the U.S. playing a leading role in several important international programs including: climate science (International Panel on Climate Change), fusion energy research (ITER) and the Nuclear Energy Advanced Modeling and Simulation (NEAMS) program”
  - Follow-on Lehman Reviews held
    - OLCF -- July 7-8, 2009
    - ALCF -- July 28-29, 2009
- Operational Assessment Reviews completed

Leadership Computing Upgrade

ASCR Deploys World's Most Powerful Computer for Open Science at ORNL

- ORNL's Cray XT5 was upgraded from 2.3 GHz quad-core processors to 2.6 GHz 6-core processors with ARRA funds.
- Increases system peak performance to 2.3 Petaflops
- Increases allocatable hours by 50% (from 1 billion to 1.5 billion hours)
- Upgrade was done in steps, keeping part of the system available
- System undergoing acceptance testing now

U.S. DEPARTMENT OF ENERGY

New NERSC-6 System "Hopper"

- Cray system selected competitively:
  - Used application benchmarks from climate, chemistry, fusion, accelerator, astrophysics, QCD, and materials
  - Best application performance per dollar based
  - Best sustained application performance per MW
  - External Services for increased functionality and availability

Grace Murray Hopper
(1906-1992)

Phase 1: Cray XT5

- 668 nodes, 5,344 cores
- 2.4 GHz AMD Opteron
- 2 PB disk, 25 GB/s
- Air cooled

Phase 2: Cray system

- > 1 Pflop/s peak
- ~ 150K cores, 12 per chip
- 2 PB disk, 80 GB/s
- Liquid cooled

3Q09	4Q09	1Q10	2Q10	3Q10	4Q10

ALCF-2 Upgrade System “Mira”

- Enables key science impact:
  - Predict abrupt regional climate change
  - Design safer, cost-effective nuclear power reactors
  - Enhancement of the extraction of biofuels from biomass
  - In silico design of nano-structured storage systems
- Builds on ASCR/NNSA investment and LLNL BG/Q Sequoia competitive bid procurement

Science Enablement Program

- 15 Science Teams
- Call: 1Q10, 1st Selection: 3Q10
- Port, Optimize, New Modeling Approaches
- Early Access to Pre-Production hardware

Mira Blue Gene/Q System

- 10 Pflop/s peak
- ~800K cores, 16 per chip
- ~70PB disk, ~470 GB/s I/O bandwidth
- Power efficient, water cooled

4Q10	1Q11	2Q11	3Q11	4Q11	1Q12	2Q12	3Q12	4Q12

Applied Math Accomplishments

- Committee: David L. Brown, Phillip Colella, Donald Estep, Paul Fischer, Omar Ghattas, Leslie Greengard, Bruce Hendrickson, Michael Holst , Sallie Keller-McNulty , Randall Leveque, Tom Manteuffel , Dianne O'Leary, Linda Petzold, James Sethian, Margaret Wright, and David Keyes (chair)

- Identifying major ASCR-funded accomplishments in Applied Mathematics in past 10 years

Applied Math Accomplishments
Areas under consideration

- Categories
  - Recently developed topics that contains rather new ideas all around (e.g., analysis of networks, adaptive algebraic multigrid)
  - Mature areas that have had substantial recent activity and achievement (e.g., a posteriori error estimate and automated adaptivity)

- high-order discretizations
- stability conserving discretizations (conservative/mimetic, monotone, etc.)
- harmonic analysis, including sublinear complexity low uncertainty wavelet transforms
- clustering/machine learning methods
- adaptive algebraic multigrid solvers
- kernel-free multipole solvers
- new optimization methods: semidefinite programs, conic optimization, semi-infinite optimization, combinatorial optimization
- sensitivity analysis and uncertainty quantification
- statistical and probabilistic methods for inverse problems, including data assimilation and forecasting

- reduced-order modeling
- network modeling
- tensor decompositions
- sparse grid methods [not "sparse matrices", but "sparse grids"]
- latency-tolerant parallel algorithms
- graph partitioners
- robust unstructured geometric mesh generation
- *a posteriori* error estimation and automated solution-based adaptivity
- front-tracking methods
- integrators: stiff, DAE, symplectic, geometric

Computer Science Accomplishments

- **Committee:** Ian Foster, Brian Tierney, Garth Gibson, Pete Beckman, Mary Hall, Rusty Lusk, Arie Shoshani, Rob Ross, Jeffrey Vetter, Leonid Oliker, Jack Dongarra, Michael Heroux, Barton Miller, Wes Bethel, Jon Bashor, and **Kathy Yelick (chair)**
- Identified major ASCR-funded accomplishments in Computer Science in past 8 years
- Criteria: Impact and difficulty
- Observations on kinds of accomplishments:
  - Individual projects that produced a software artifact
  - Creation and support of research communities, e.g., performance tools, lightweight OS kernels, benchmarking

Computer Science Accomplishments (not finalized)

Scientific Libraries
- LAPACK/ScaLAPACK
- Trilinos
- SuperLU
- PETSc

Programming Models
- MPI
- PGAS Languages
- Global Arrays

Autotuning
- Atlas

Architecture and Performance
- Performance modeling
- Architecture Benchmarking
- Petascale Architecture Partnerships

Performance tools
- Dyninst

LightweightOS
- Zeptos
- Catamount

Networking and Security
- OSCARS: Hybrid networks
- Bro at 10+ Gbps

Analyzing Large Data Sets
- Petascale production Visualization
- Fastbit
- Topologically based data analysis

Storing and Accessing Large Data Sets
- Object Store / Parallel Filesystems
- HL I/O Libraries

Moving Large Data Sets
- TCP Autotuning
- Grid Security Infrastructure
- GridFTP
- Storage Resource Management (SRM)

Computational Science Accomplishments

- **Committee:** Jackie Chen, Giulia Galli, Jim Hack, Doug Kothe, Paul Messina, Juan Meza, Chris Mundy, Claudio Rebbi, Nagiza Samatova, Panagiotis Spentzouris, Bill Tang, and Tony Mezzacappa (chair)
- Continuation 2008 Breakthroughs with 15-20 additional breakthroughs;
- Not finalized

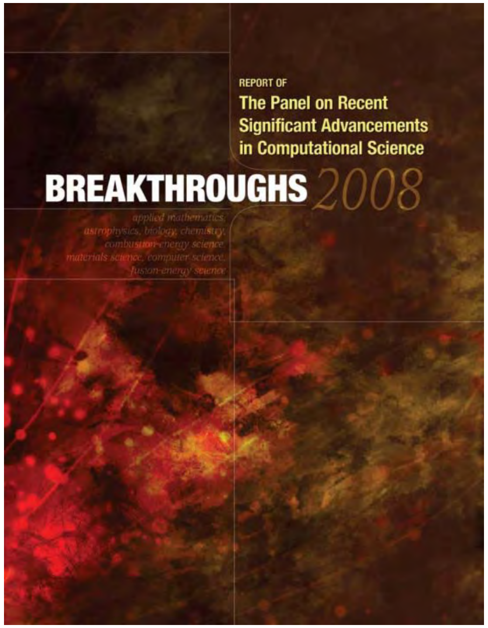

First Self-consistent Simulation of Baryon Acoustic Oscillations in the Intergalactic Medium

Michael Norman, Robert Harkness, Pascal Paschos, UCSD

2048^3 cell/particle hybrid simulation, 330 Mpc volume, ENZO code
2048 procs, 1.2 million CPU-hrs, 6 TB RAM, 200 TB output, 6 month job , NERSC Seaborg
2006 INCITE award; 2006 Joule Metric application

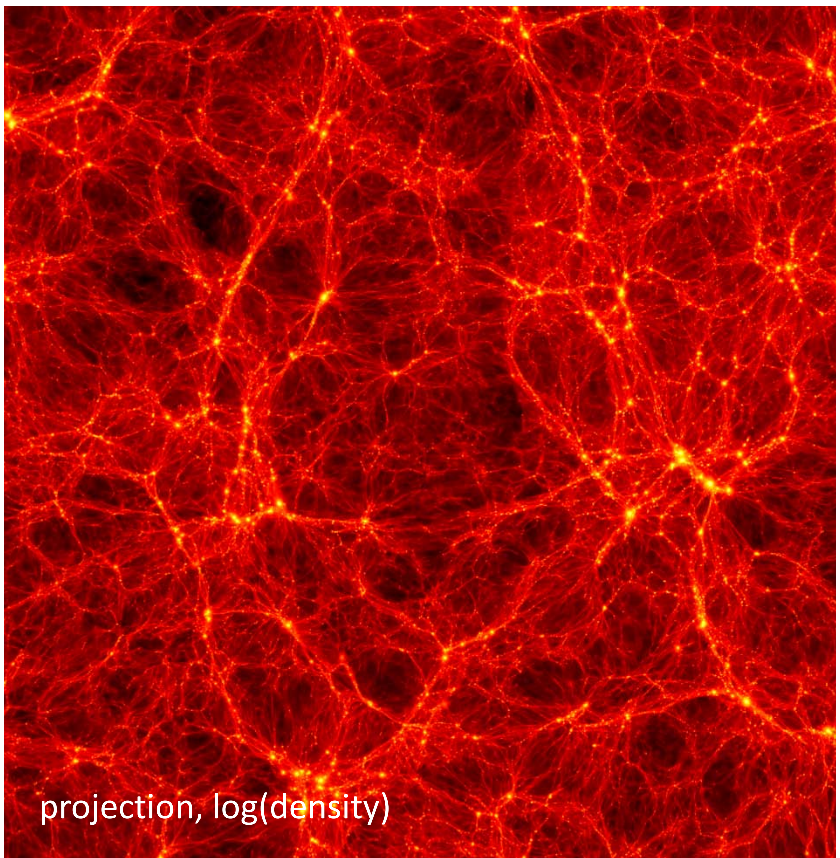

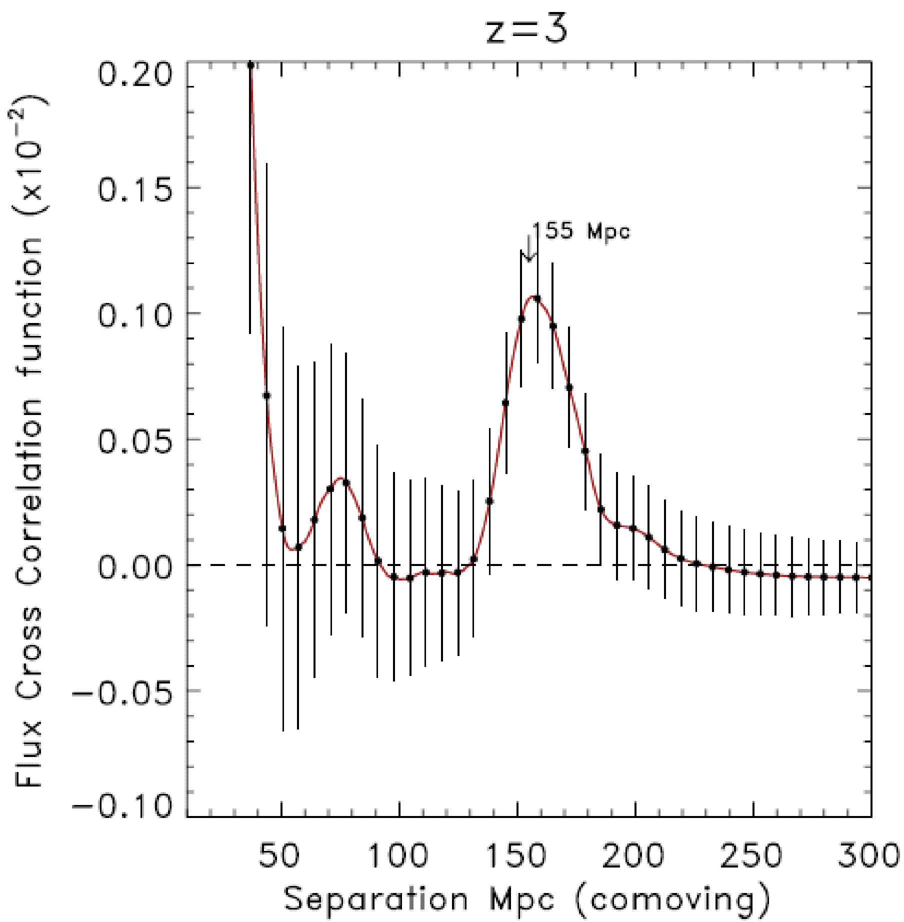

Large Eddy Simulation of combustion instabilities in a gas turbine combustion chamber

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

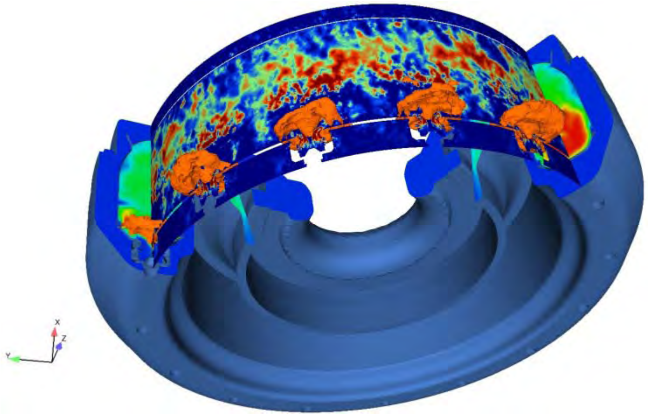

- First simulation of a complete helicopter gas turbine combustion chamber using Large Eddy Simulation of the turbulent reacting flow within the chamber
- Allowed to reproduce and suppress acoustic modes observed in real engines
- Feature paper of *Combustion and Flame* , July 2008.
- Scalability verified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

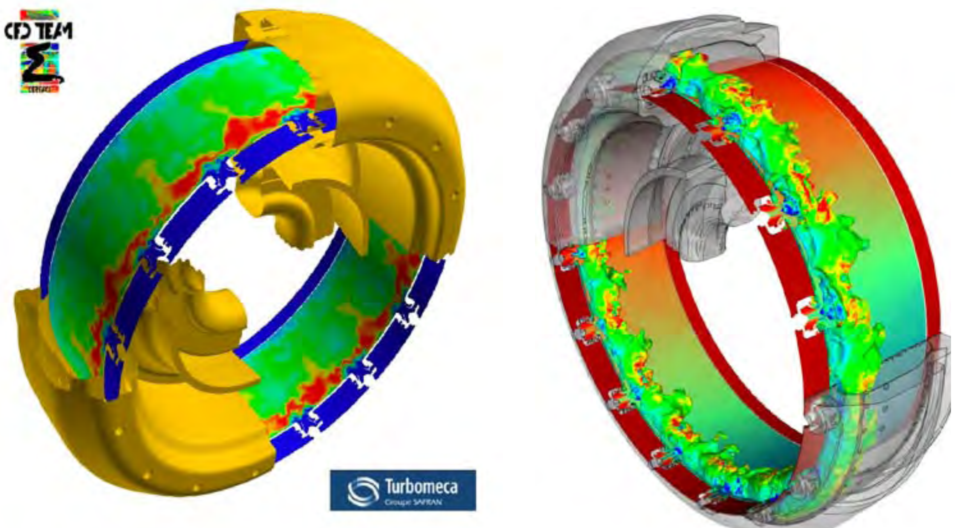

Figure 1: (left) instantaneous temperature field on a cylinder view plane passing through the injectors of the helicopter chamber. Red zones correspond the hot gases and blue zones to cold air coming from the compressor. (right) instantaneous pressure field in the same view plane with isosurfaces of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

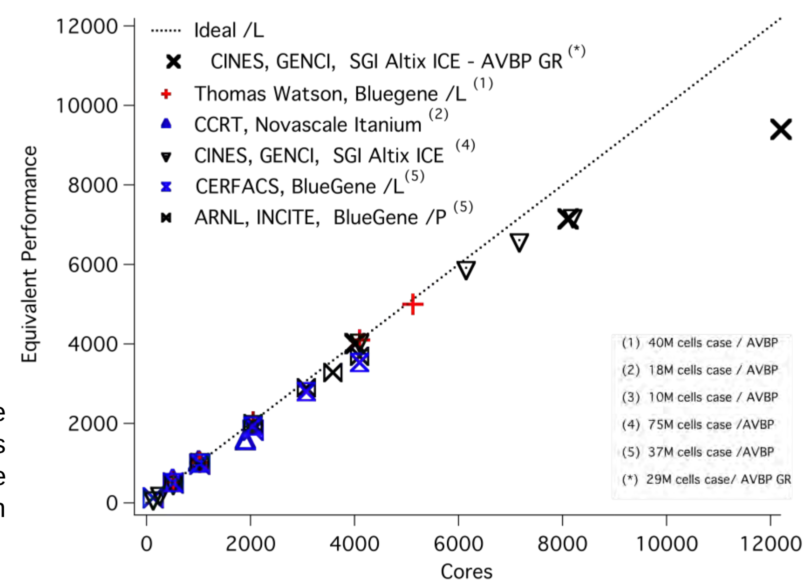

AMR Simulations of Pellet Injection in Tokamaks

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

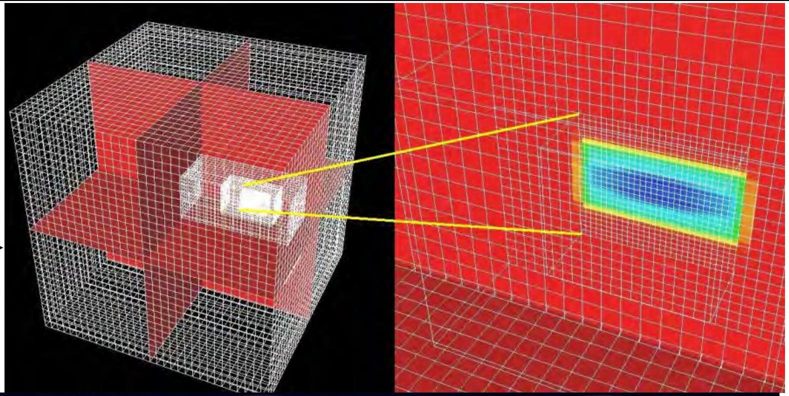

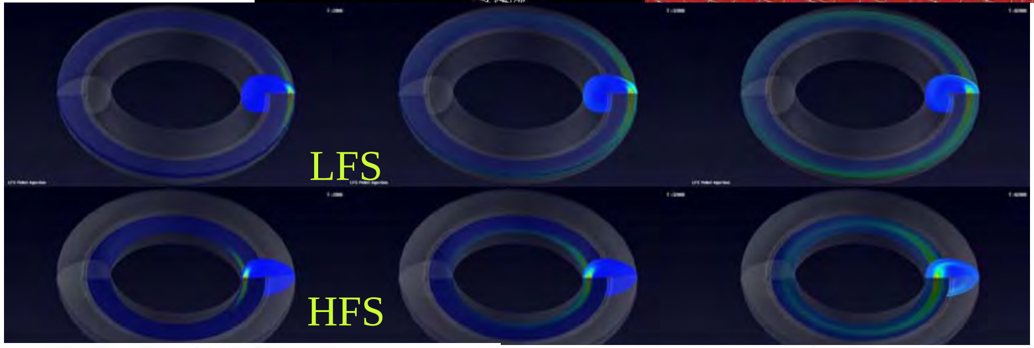

Average density profiles during HFS and LFS pellet injection (Arrow indicated pellet location)

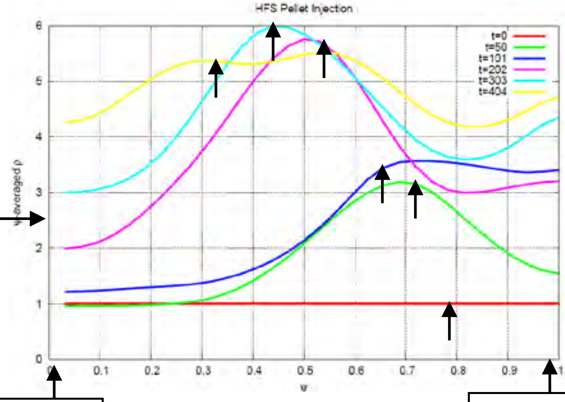

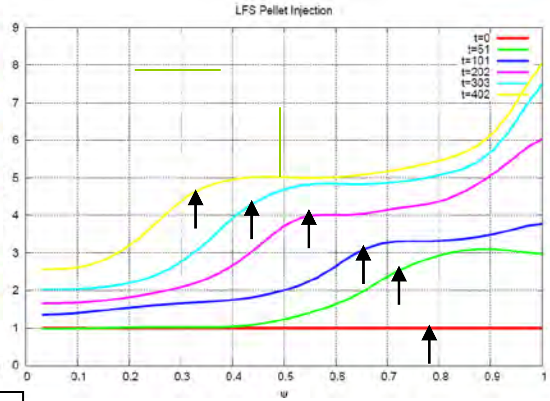

Core

Edge

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.

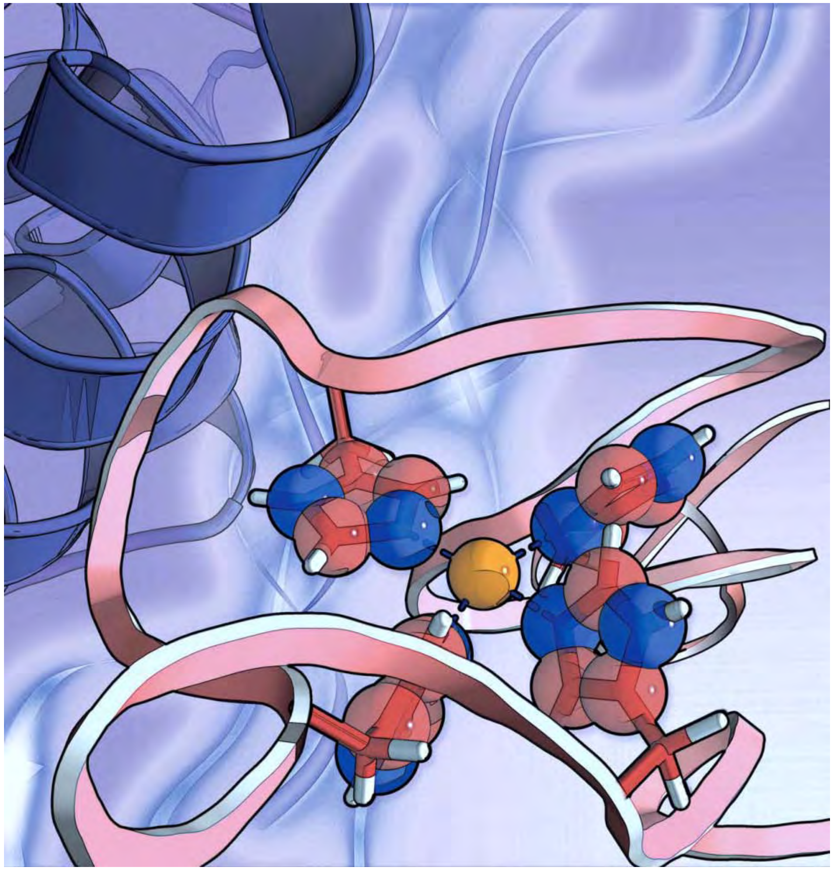

The prion protein, with copper bound to its unstructured part in the foreground, and its folded portion containing alpha helices in the background.

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

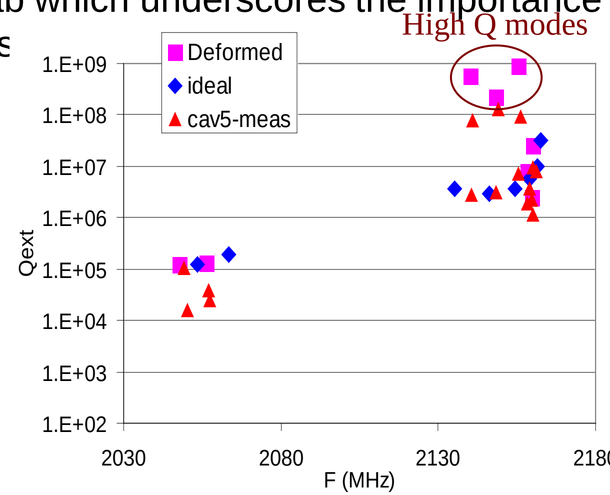

Method of Solution - Using the measured cavity parameters as inputs, the deformed cavity shape was recovered by solving the ***inverse*** problem through an optimization method. The calculations showed that the cavity was 8 mm shorter than designed, which was subsequently confirmed by measurements. The 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. .

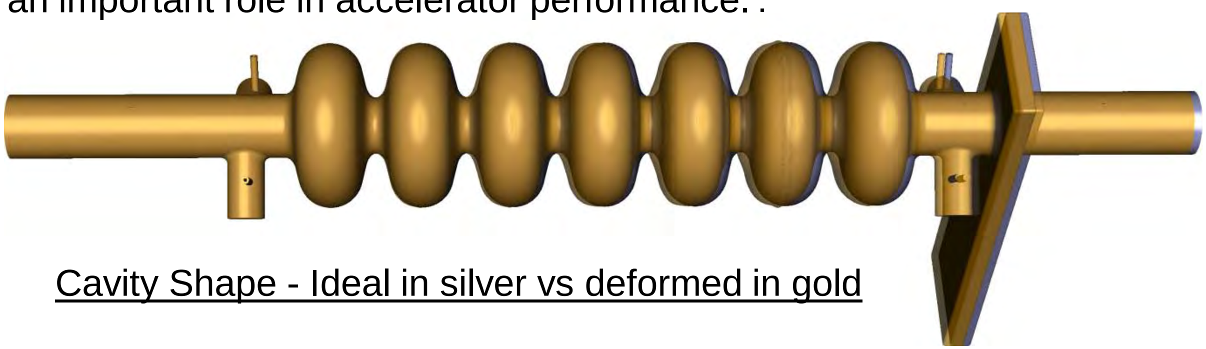

Cavity Shape - Ideal in silver vs deformed in gold

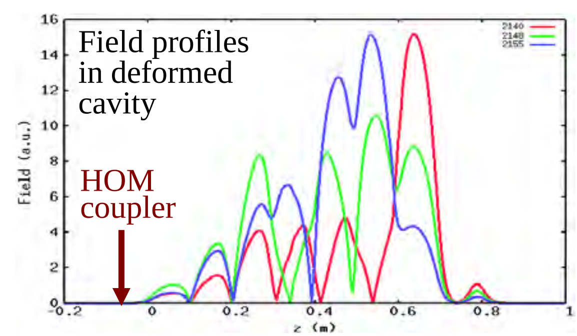

Simulation of lean premixed hydrogen flames

J. Bell, M. Day, R. Cheng, M. Lijewski, V. Beckner, LBNL; S. Tachibana, LBNL and JAXA, Tokyo

- **Methodology**
  - Low Mach number formulation that exploits mathematical structure of the problem
  - Advanced numerical methodology
    - Projection methodology
    - Adaptive mesh refinement
    - Parallel implementation using BoxLib
- **Simulation of lean-premixed hydrogen flame stabilized on a low-swirl burner**
    - Detailed chemistry and transport
    - No explicit models for turbulence or turbulence / chemistry interaction
  - 25 cm x 25 cm x 25 cm
  - Combined methodology enables simulation at effective resolution of 2048^3
- **Simulation captures cellular structure of thermodiffusively unstable lean hydrogen flames**
  - Quantify enhanced burning from local enrichment of the fuel resulting for high H_2 diffusion
  - Provide insight into the analysis of experimental diagnostics

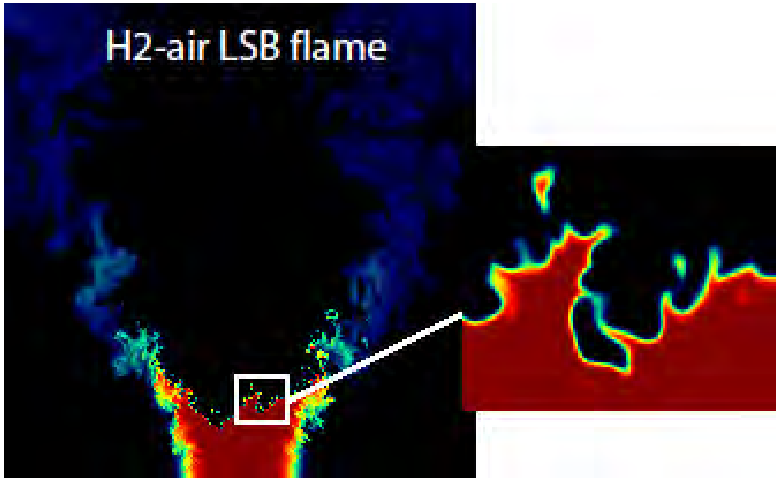

Fuel

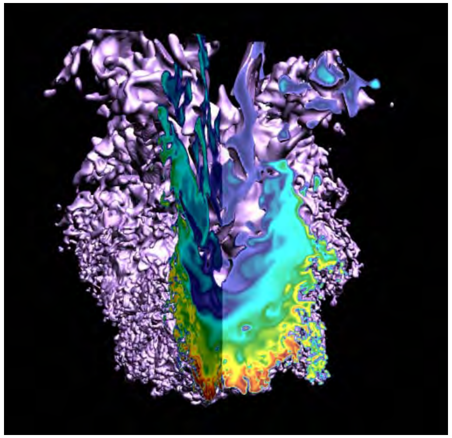

OH Radical

Simulations performed at NERSC under an INCITE grant

Unique high-fidelity calculations of direct-injection processes for IC-engine applications

J. Oefelein, SNL; R. Sankaran, ORNL

2008-2009 INCITE Award

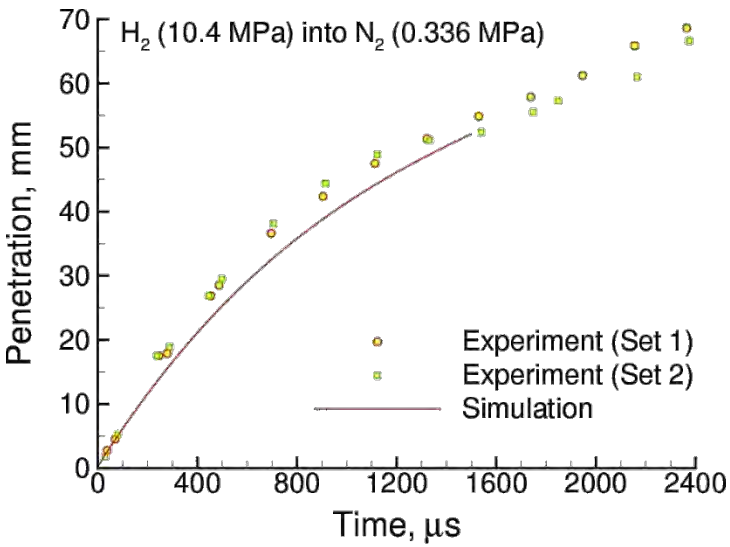

Representative comparison of LES with penetration measurements

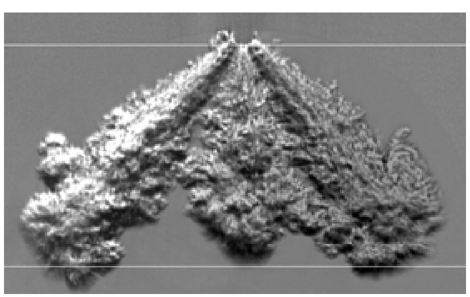

Shadowgraph (U. Wisconsin)

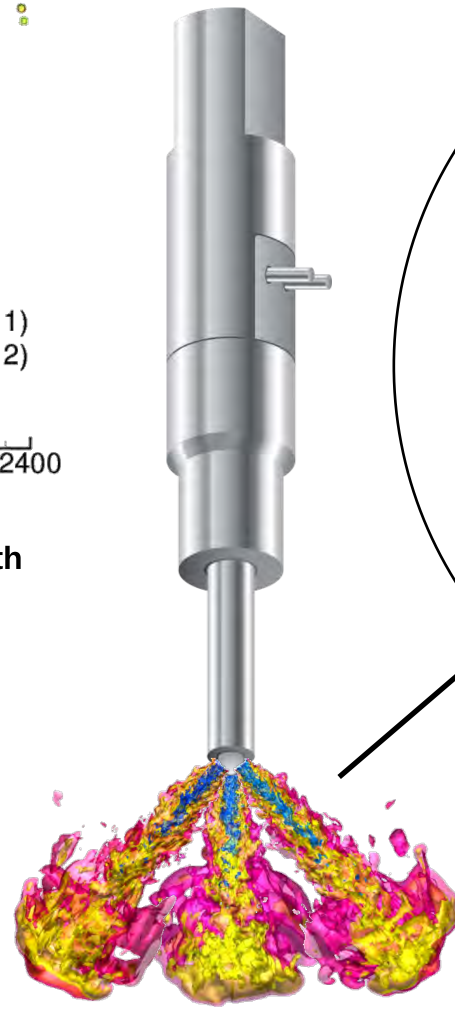

Large Eddy Simulation

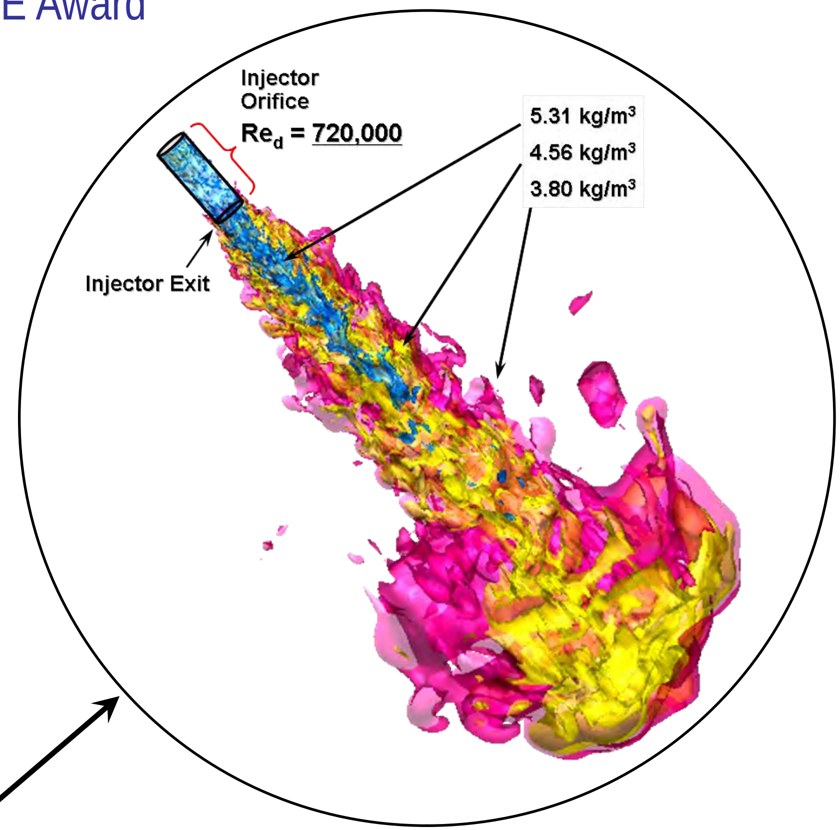

Iso-Contours of Density ($H_2 - N_2$)	
Orifice Diameter	0.8 *mm*
Injection Pressure	10.4 *MPa*
Injection Temperature	298 *K*
Chamber Pressure	0.336 *MPa*
Chamber Temperature	298 *K*

Reactor Core Hydrodynamics

P. Fischer, D. Pointer, A. Siegel, ANL; C. Pantano, U. Illinois

- Large-scale simulations of turbulent thermal transport in sodium-cooled reactor cores are increasing the understanding of fundamental thermal mixing phenomena within advanced recycling reactor (ARR) cores.
  - *Insights gained from these simulations will enable higher power output while increasing safety*

- Potential benefits of ARRs:
  - *Improved safety and economy*
  - *Economical power sources*
  - *Used to recycle spent nuclear fuel*
  - *Reduced loading demands (nearly 100-fold) in geological repositories*

- Work was conducted on the Argonne Leadership Computing Facility BG/P through INCITE awards
  - *Simulation sizes have increased ten-fold in the past 12 months and sustain 80% efficiency on 131,000 cores.*

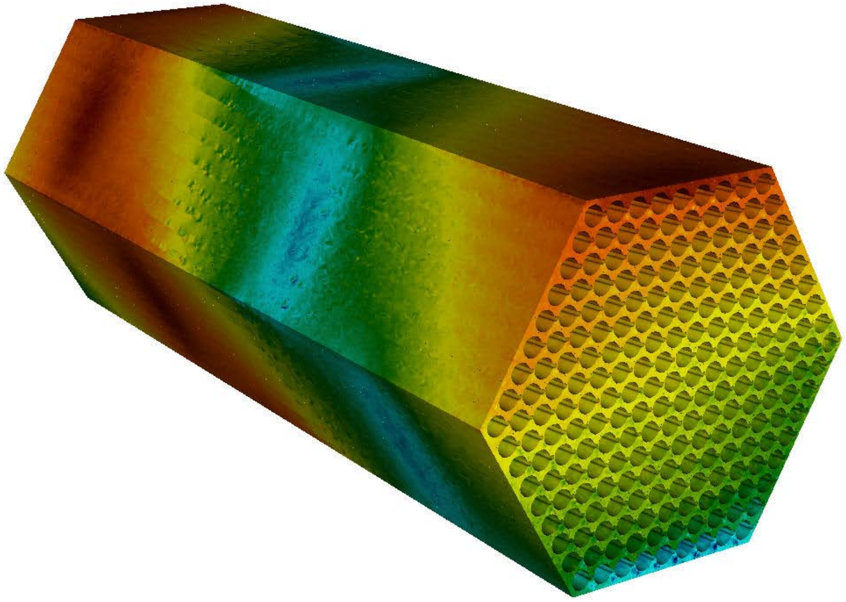

Pressure distribution in a 217-pin fuel assembly with wire wrap