U.S. DEPARTMENT OF ENERGY | Office of Science

ASCR Update

August 24, 2010

Michael Strayer

Associate Director, Advanced Scientific Computing Research

Office of Science (SC)

FY 2011 Budget Request to Congress

(B/A in thousands)

	FY 2009		FY 2010	FY 2011		
	Current Base Approp.	Current Recovery Act	Current Approp.	Request to Congress	Request to Congress vs. FY 2010 Approp.	
Advanced Scientific Computing Research	358,772	161,795	394,000	426,000	+32,000	+8.1%
Basic Energy Sciences	1,535,765	555,406	1,636,500	1,835,000	+198,500	+12.1%
Biological & Environmental Research	585,176	165,653	604,182	626,900	+22,718	+3.8%
Fusion Energy Sciences	394,518	91,023	426,000	380,000	-46,000	-10.8%
High Energy Physics	775,868	232,390	810,483	829,000	+18,517	+2.3%
Nuclear Physics	500,307	154,800	535,000	562,000	+27,000	+5.0%
Workforce Development for Teachers & Scientists	13,583	12,500	20,678	35,600	+14,922	+72.2%
Science Laboratories Infrastructure	145,380	198,114	127,600	126,000	-1,600	-1.3%
Safeguards & Security	80,603	——	83,000	86,500	+3,500	+4.2%
Science Program Direction	186,695	5,600	189,377	214,437	+25,060	+13.2%
Small Business Innovation Research/Technology Transfer (SC)	104,905	18,719	——	——	——	——
Subtotal, Science	**4,681,572**	**1,596,000**	**4,826,820**	**5,121,437**	**+294,617**	**+6.1%**
Congressionally-directed projects	91,064	——	76,890	——	-76,890	-100.0%
Small Business Innovation Research/ Technology Transfer (DOE)	49,534	36,918	——	——	——	——
Use of prior year balances	-15,000	——	——	——	——	——
Total, Office of Science	**4,807,170**	**1,632,918**	**4,903,710**	**5,121,437**	**+217,727**	**+4.4%**

Congressional FY11 Budget Action

- **Senate:**
  - SC: $5.012B , $109.4M under the request
  - ASCR: $418M, $8M under the request
- **House:**
  - SC: $4.9B, $221.0M under the request, $4M under the FY 2010 enacted.

ASCR FY11 Budget Request
Highlights

- Maintain FY10 increases in Applied Mathematics and Computer Science to prepare for challenges of future architectures, huge datasets, and multi-disciplinary science.
- Continue Exascale research programs started in FY2010
- Continue focus on Computational Partnership teams on transforming critical DOE Applications to be ready for running at scale on multicore computers
- Fulfill obligations at computing facilities (leases) and to DARPA HPCS program.
- Begin site preparation for ALCF-2 upgrade to be installed in 2012-2013.
- Support NERSC-6 operations.
- Support ESnet deployment of 100Gbps technologies to meet growing science requirements.

Exascale Progress since May, 2010

- **Proposals processed in Exascale related topic areas:**
  - **Applied Math: Uncertainty Quantification**
  - **Computer Science: Advanced Architectures**
  - **Computer Science: X-Stack**
  - **Computational Partnerships: Co-Design (21 Proposals requesting ~ $160M/year)**
- **Exascale Coordination meeting with DOD and DARPA – June 15 with decisions to --**
  - **Have yearly coordination/planning meetings in accordance with current MOU**
  - **Follow-on meeting of Program Managers to identify gaps and overlaps in computer science and applied math areas currently funded**
  - **Address critical technologies issues, initially in memory**
- **Project documentation started**
  - **Mission Need**
  - **High level Acquisition Strategy that outlines program, siting of machines; roles and responsibilities, etc.**
- **Ongoing meetings with NNSA**

Advanced Architectures

(28 proposals requesting ~$28M/yr)

- **Blackcomb: Hardware-Software Co-design for Non-Volatile Memory in Exascale Systems**
  - ORNL, Hewlett Packard Labs, University of Michigan, Penn State University
- **Enabling Co-Design of Multi-Layer Exascale Storage Architectures**
  - ANL, RPI
- **NoLoSS: Enabling Exascale Science through Node Local Storage Systems**
  - ANL, LLNL
- **CoDEx: A Hardware/Software Co-Design Environment for the Exascale Era**
  - LBNL, SNL, LLNL, GATECH
- **Data Movement Dominates: Advanced Memory Technology to Address the Real Exascale Power Problem**
  - SNL, LBNL, Micron, Columbia, UMD
- **Thrifty: An Exascale Architecture for Energy-Proportional Computing**
  - UIUC, LLNL, UCSD, Intel

Six projects funded at $5M/yr

X-Stack Software

(55 proposals requesting ~$40M/yr)

- **A Fault-oblivious Extreme-scale Execution Environment**
  - SNL, PNNL, LLNL,OSU, IBM, Alcatel-Lucent Bell Labs, BU
- **Auto-tuning for Performance and Productivity on Extreme-scale Computations**
  - LBNL, UCSB, PPPL
- **Software Synthesis for High Productivity ExaScale Computing**
  - MIT, UC Berkeley
- **COMPOSE-HPC: Software Composition for Extreme Scale Computational Science and Engineering**
  - ORNL, Galois, Inc., LLNL, PNNL, SNL
- **Damsel: A Data Model Storage Library for Exascale Science**
  - Northwestern University, ANL), NCSU, The HDF Group
- **Vancouver: Designing a Next-Generation Software Infrastructure for Productive Heterogeneous Exascale Computing**
  - ORNL, UIUC, University of Oregon, GATECH

X-Stack Software (cont)

(55 proposals requesting ~$40M/yr)

- **Enabling Exascale Hardware and Software Design through Scalable System Virtualization**
  - UNM, Northwestern University, SNL, ORNL
- **ZettaBricks: A Language, Compiler and Runtime Environment for Anyscale Computing**
  - MIT
- **Compiled MPI: Cost-Effective Exascale Application Development**
  - UIUC, LLNL, Indiana
- **ExM: System support for extreme-scale, many-task applications**
  - ANL, Chicago, University of British Columbia
- **An Open Integrated Software Stack for Extreme Scale Computing**
  - ANL, LBNL, ORNL, SNL, UTK

Eleven projects funded at $8.5M/yr

Summary of UQ Projects

(90 projects received requesting ~ $45M/yr)

- **Modeling and Simulation of High-Dimensional Stochastic Multiscale PDE Systems at the Exascale**
  - Guang Lin (PNNL), Nicholas Zabaras (Cornell), and Ioannis Kevrekidis, (Princeton)
- **Advanced Dynamically Adaptive Algorithms for Stochastic Simulations on Extreme Scales**
  - Dongbin Xiu (Purdue), Richard Archibald, Ralf Deiterding, and Cory Hauck (ORNL)
- **A High-Performance Embedded Hybrid Methodology for Uncertainty Quantification with Applications**
  - Charles Tong (LLNL), Barry Lee (PNNL), Gianluca Iaccarino (Stanford)
- **Enabling Predictive Simulation and UQ of Complex Multiphysics PDE Systems by the Development of Goal-Oriented Variational Sensitivity Analysis and a-Posteriori Error Estimation Methods**
  - John Shadid (SNL), Don Estep (CSU), Victor Ginting (UWyoming)
- **Bayesian Uncertainty Quantification in Predictions of Flows in Highly Heterogeneous Media and its Application to CO2 Sequestration**
  - Yalchin Efendiev (Texas A&M), Panayot Vassilevski (LLNL)
- **Large-Scale Uncertainty and Error Analysis for Time-Dependent Fluid/Structure interactions in Wind Turbine Applications**
  - Juan Alonso (Stanford) and Michael Eldred, et al (SNL)

Six projects funded at $3M/yr

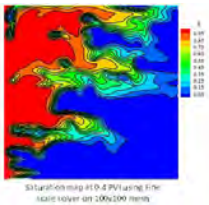

Fine mesh

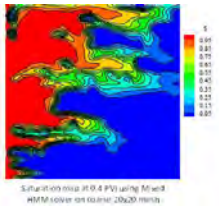

Coarse-graining

Saturation profile of oil-water porous media flow using fine and multiscale coarse-graining solver

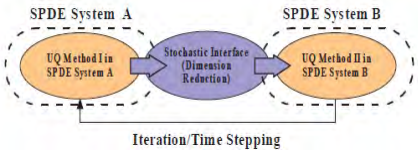

Sketch of hybrid UQ method between multi-physics stochastic PDE systems

Applied Mathematics

Hybrid Optimization Parallel Search Package

Software for derivative-free optimization: Culmination of 10 years research resulting in mature software package

- Addresses general optimization problems
- Easy interface to application codes; multiple solvers can collaborate
- Asynchronous parallel evaluations; MPI and multithreading

Previous approach: Hand-tuning by experts. A time-consuming, iterative process that directly involved the scientist. Cannot keep up with expected future demands.

New approach: HOPSPACK used to automatically tune 50 parameters in direct analysis code in one hour. Will enable more sophisticated and accurate models to be developed.

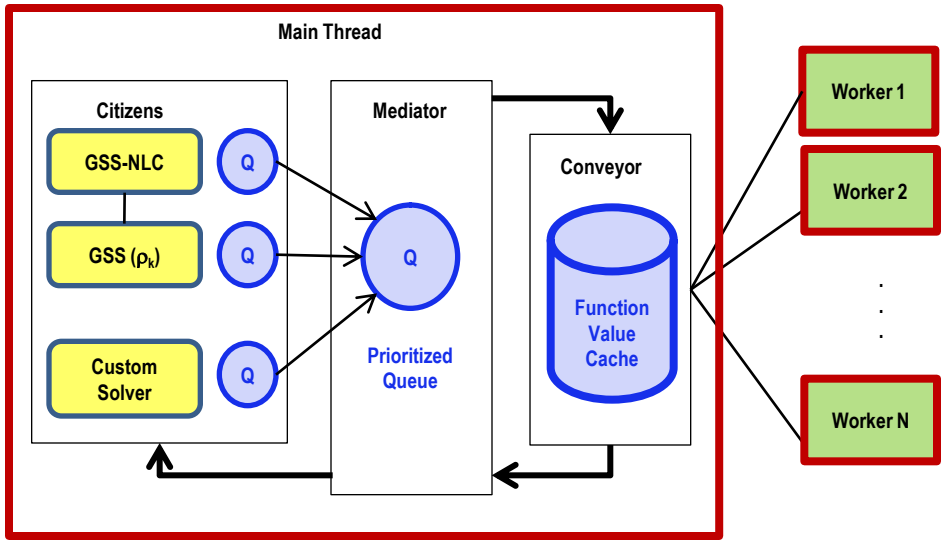

IMPACT

- Over 200 downloads since 10/2009; 65% for Windows, 30% for Linux, and 5% for Mac.
- National laboratory, academic, and commercial users

POCs:
Tammy Kolda and Todd Plantenga, SNL

SMARTMAP – New OS Capability

SMARTMAP Improves Communication Performance on Multi-Core Processors

- SMARTMAP team: Ron Brightwell, Sandia

SMARTMAP can significantly increase MPI communication performance on multi-core processors over traditional shared memory

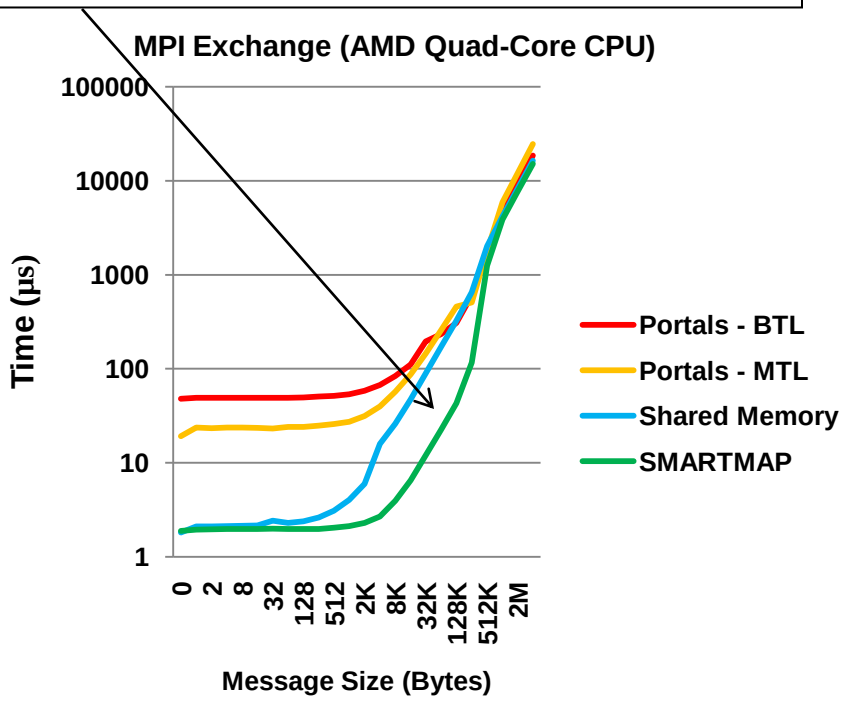

2009 R&D 100

Accomplishments

- Provides direct access shared memory between processes on a multi-core processor
- No change to programming model
- Eliminates all extra memory copies for MPI
- Enables several other MPI optimizations
- Not just for MPI
- Can be used by applications directly to optimize computational libraries
- R&D100 Award-winning technology

Joint Math/CS Institute

EASI: Multicore Solver

EASI Addresses Critical Bottleneck in Iterative Solvers

- **EASI project team: ORNL, SNL, UC-Berkeley, UIUC, UT-Knoxville**
- Iterative Solvers are Dominant cost of many apps
- Exascale challenges:
  - Collectives, synchronization.
  - Memory latency/BW.

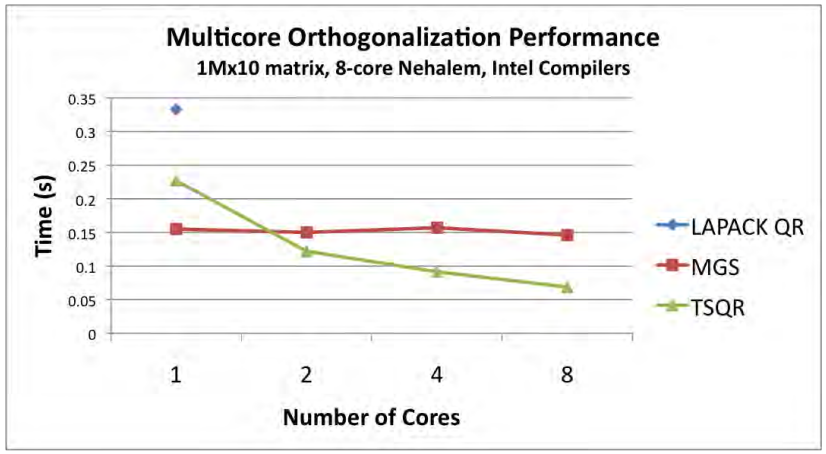

LAPACK – Serial, MGS –Threaded modified Gram-Schmidt

TSQR capability:

- **Critical for exascale solvers.**
- **Part of the Trilinos scalable multicore capabilities.**
- **Helps all iterative solvers in Trilinos (available to external libraries, too).**
- **Lead developer: Mark Hoemmen (post-doc, UC-Berkeley)**
- **Part of Trilinos 10.6 release, Sep 2010.**

Accomplishment

- **Communication-avoiding (s-step) iterative solvers:**
  - s times fewer synchronizations & data transfers.
- **New orthogonalization algorithm:**
  - Tall Skinny QR factorization (TSQR).
- **TSQR Implementation:**
  - 2-level parallelism (Inter and intra node).
  - Memory hierarchy optimizations.
  - Flexible node-level scheduling via Intel Threading Building Blocks.

VACET Automatic Feature Detection

The Visualization and Analytics Center for Enabling Technologies (VACET) delivers first-ever ability to see relationship between simulation parameters (e.g., level of turbulence) and scalar dissipation rate.

- **VACET project team: LBNL, LLNL, ORNL, UTAH**
- New family of robust combinatorial methods for identifying and tracking in time topological features related to local mixing rates and to a scalar representative of autoignition.

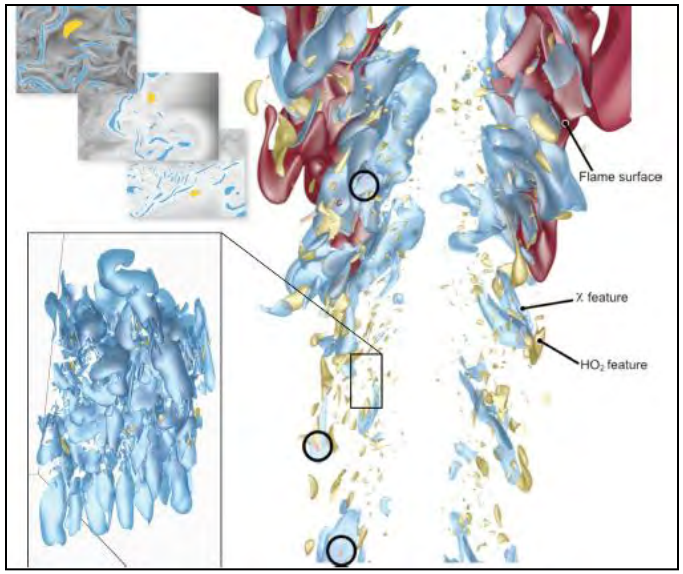

Accomplishments

"These methods have enabled us, for the first time, to robustly track over time large numbers of extinction or ignition regions, which is essential to the development of a better understanding of the flame extinction/ reignition and stabilization dynamics."

- Jackie Chen, SNL-CA.

CEDPS Data Movement Service

The Center for Enabling Distributed Petascale Science (CEDPS) provides secure, high-bandwidth "fire-and-forget" file transfer services

- **CEDPS project team: ANL, LBNL, USC/ISI, FNAL, UW-Madison**
- Serving scientists and resource providers from: LBNL / NERSC, ORNL, ANL, GFDL, TeraGrid, Europe, Australia

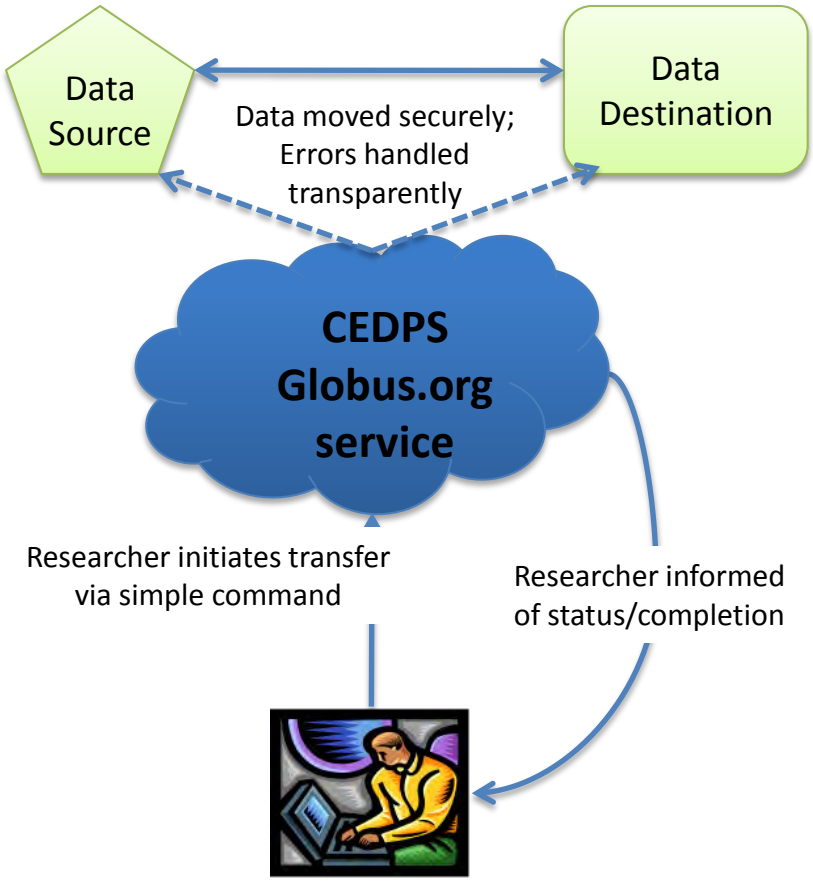

Accomplishments

Transferred without human intervention, automatically recovering from many network and systems failures:

- **STAR: 300GB of data between BNL and NERSC/PDSF**
- **SCEC: 600,000 4MB files from TeraGrid resources to OSG**
- **CEDPS: 20TB among 12 sites (including DOE labs, TeraGrid, Australia) in under 15 hours**

Powered by:

NERSC

Hopper installation planned for completion in late 2010

- 1.25 PFlop/s peak performance Cray XE6
- Over a billion core-hours to science per year
- Gemini high performance resilient interconnect
- Two 12-core AMD Magny-Cours chips per node
- Collaboration with NNSA ACES on testing

NERSC/Cray Center of Excellence

Programming Models

- Ensures effective use of new 24-core nodes

Hopper installation, August 2010

Scientific Accomplishments at NERSC

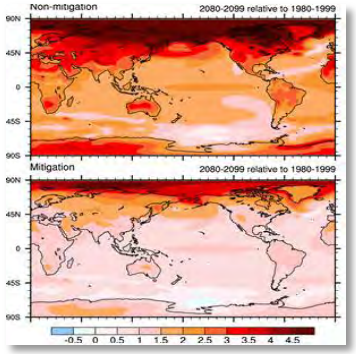

Climate

Studies show that global warming can still be diminished if society cuts emissions of greenhouse gases.
(Warren Washington, NCAR)

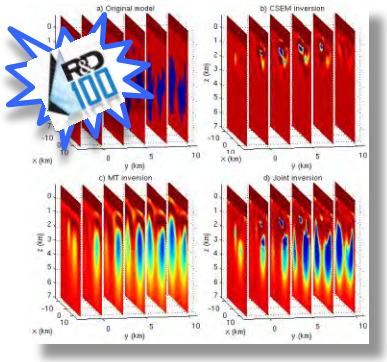

Energy Resources

Award-winning software uses massively-parallel supercomputing to map hydrocarbon reservoirs at unprecedented levels of detail.
(Greg Newman, LBNL)

.

Fusion Energy

A new class of non-linear plasma instability has been discovered that may constrain design of the ITER device.
(Linda Sugiyama, MIT)

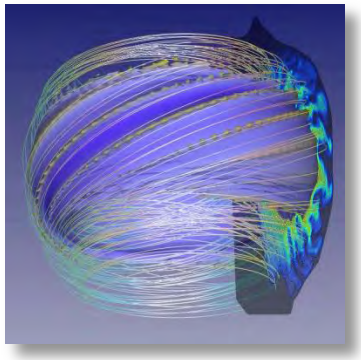

Combustion

Adaptive Mesh Refinement allows simulation of a fuel-flexible low-swirl burner that is orders of magnitude larger & more detailed than traditional reacting flow simulations allow.
(John Bell, LBNL)

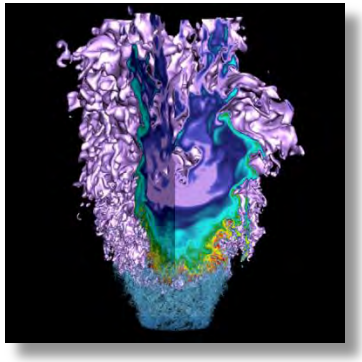

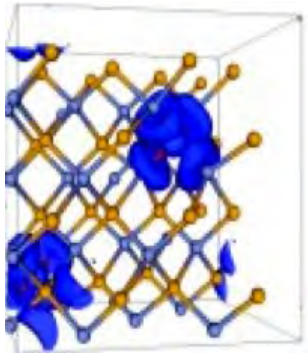

Materials

Electronic structure calculations suggest a range of inexpensive, abundant, non-toxic materials that can produce electricity from heat.
(Jeffrey Grossman, MIT)

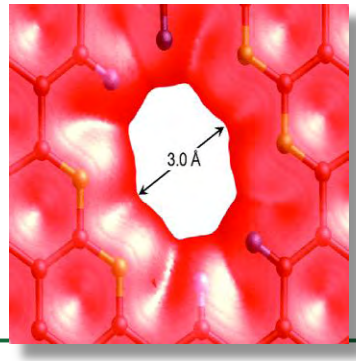

Nano Science

Using a NERSC NISE grant researchers discovered that Graphene may be the ultimate gas membrane, allowing inexpensive industrial gas production.
(De-en Jiang, ORNL)

Argonne Leadership Computing Facility

ALCF - now

- more CPUs/node (x 4)
- faster CPUs (x 60)
- faster storage (x 2.7)
- more storage (x 3.6)
- more RAM/CPU (x 2)

ALCF - 2013

Intrepid

Blue Gene/P—peak **557 TF**
40 racks
40,960 nodes (quad core)
163,840 processors (3.4 GF Peak)
80 TB RAM
8 PB storage capacity
88 GB/s storage rate
Air-cooled

Mira

Blue Gene/Q—peak **10 PF**
48 racks -- *just one row more than Intrepid*
49,152 nodes (16 core)
786,432 processors (205 GF peak)
786 TB RAM
28.8 PB storage capacity
240 GB/s storage rate
Water-cooled

~20 times the computing power in only 20% more space

ALCF Scientific Accomplishments

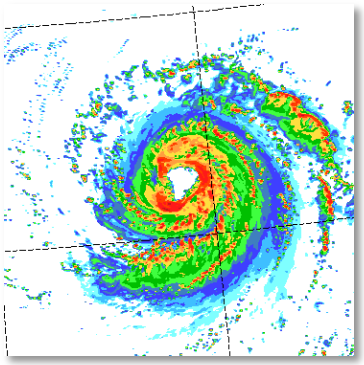

Climate

Used leadership class, vortex-following calculation to more accurately predict hurricane track, to better mitigate risks. **(John Michalakes, NCAR)**

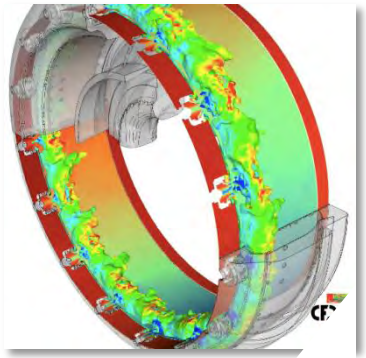

Gas Turbines

Two-phase flow and combustion modeling identified instability mechanisms that reduce efficiency, leading to design of more efficient aircraft engines. **(Thierry Poinsot, CERFACS)**

Nuclear Energy

High-fidelity fluid flow and heat transfer simulation of next-generation reactor designs, aiming to reduce the need for costly experimental facilities. **(Paul Fischer, ANL)**

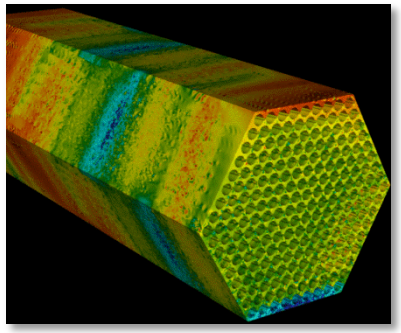

Nano Catalysts

Mapped out properties of a wide range of gold nanoparticles to design catalysts for fuel cells and methane conversion. **(Jeff Greeley, ANL)**

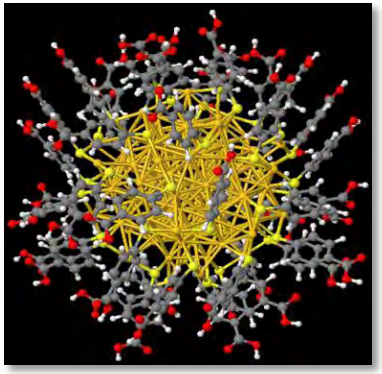

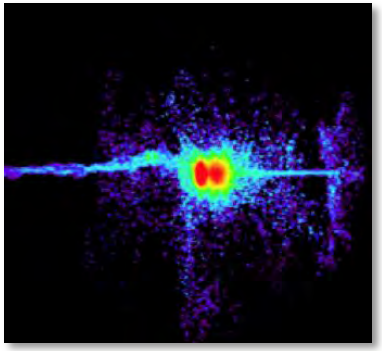

Fusion Energy

New hybrid algorithm allowed study of physics in Fast Ignition inertial confinement fusion over a much greater density range than planned. **(John Tonge, UCLA)**

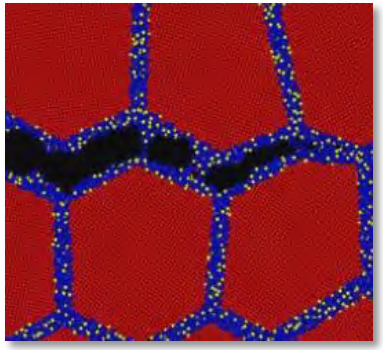

Materials Science

Molecular dynamics simulation explained how a minute sulfur impurity embrittles nickel—relevant to next-generation nuclear reactor design. **(P. Vashishta, USC)**

Oak Ridge Leadership Computing Facility

- **ALCF and OLCF Mission Need (CD-0) signed January, 2009 for 20-40 Petaflop upgrade across facilities**
- **July, 2009 Lehman Review recommended**
  - approval of acquisition strategy (CD-1) for heterogeneous processor machine
  - Application Readiness review to
    - Demonstrate the viability of system design for science
    - Facilitate getting science day 1
- **August 17-18, 2010, Application Readiness review**
  - Focused on six applications: S3D, combustion; CAM-HOMME, climate; WL-LSMS, nanoscience; LAMMPS, biology/materials; Denovo, nuclear energy; PFLOTRAN, geoscience
  - Panel recommended approval of OLCF application readiness plan
- **Lehman Review for CD-2 scheduled for December, 2010**

OLCF: Advancing Scientific Discovery

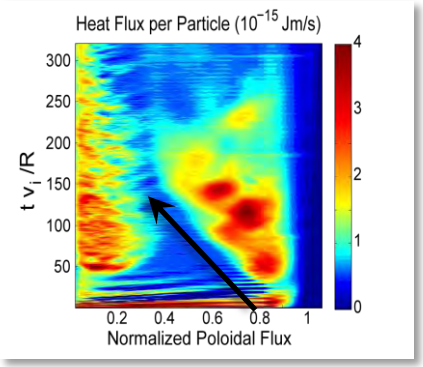

Fusion Efficiency in ITER
First-principles plasma heating simulation using a realistic device geometry confirms plasma heating by edge confinement. Deeper understanding will immediately impact ITER projections. **(C. S. Chang, NYU**

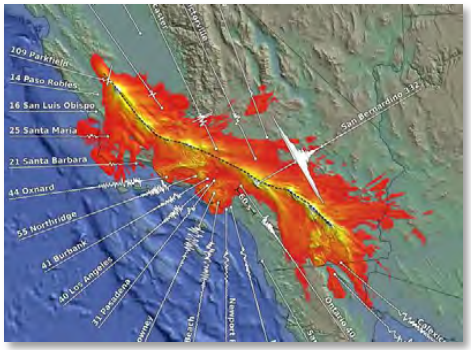

Emergency Response in Earthquakes
Largest simulation of an earthquake ever performed shows a magnitude-8 quake and its impact on the region. Accurate forecasts lead directly to improved planning for emergency management. **(Tom Jordan, USC)**

PEDOT to Aid Future Electronics
Simulations of an organic polymer-widely used as a conductive material in devices like light-emitting diodes and televisions-could mean more efficient and cheaper electronics. **(Bobby Sumpter, ORNL)**

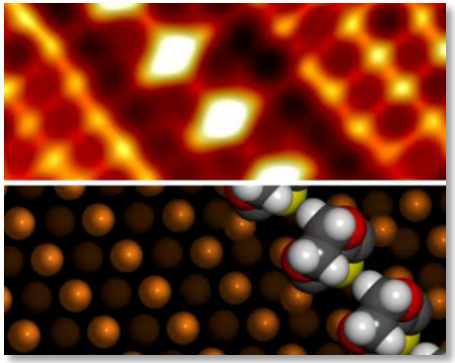

The Origins of Cosmic Lighthouses
MHD simulations of supernovae provide a mechanism to generate highly-magnetized neutron stars, the hearts of pulsars and magnetars and the possible progenitors of gamma-ray burtsts. **(Tony Mezzacappa, ORNL)**

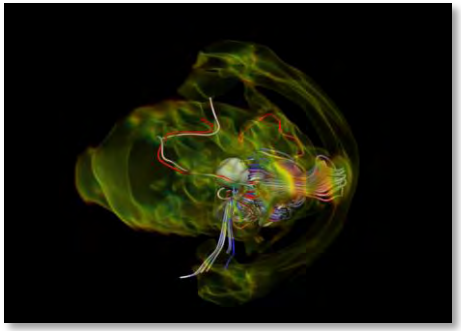

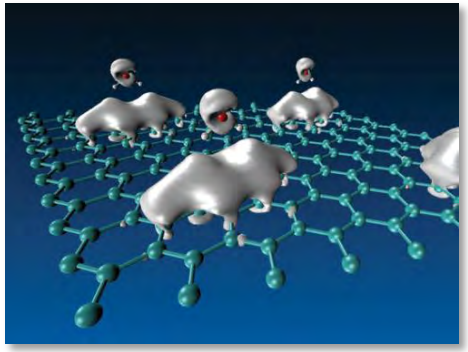

Exploring Carbon-Water Union
Researchers obtain a full-binding energy curve between water and graphene, research that could lead to better industrial lubricants, sreamlined catalysis, and more efficient hydrogen batteries. **(Dario Alfè, UCL)**

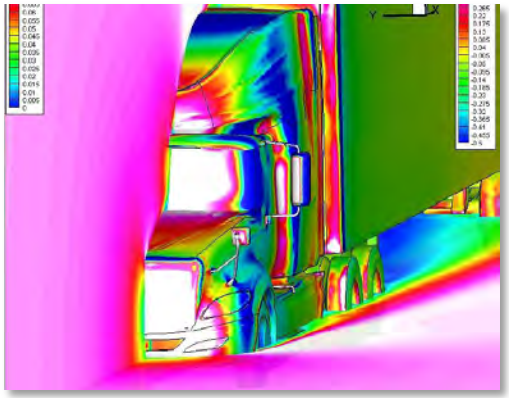

Fuel Efficiency from SmartTrucks
The most accurate and detailed numerical models of 18-wheelers ever created lead to the design of add-on parts that can save more than 3100 gallons of fuel per truck each year. **(Branden Moore, BMI)**

ASCR Facilities

4 of 6 ACM 2010 Gordon Bell Prize Finalists

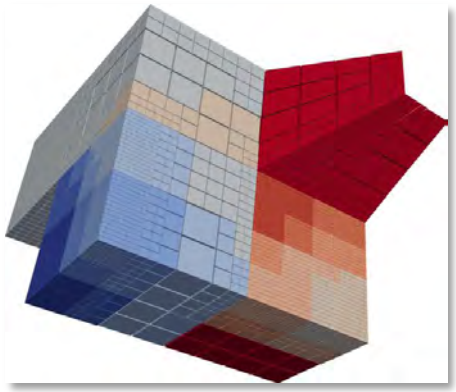

Extreme-Scale AMR
Scaling difficult Adaptive Mesh Refinement techniques to over 224,000 cores on **Jaguar** demonstrating excellent scaling.

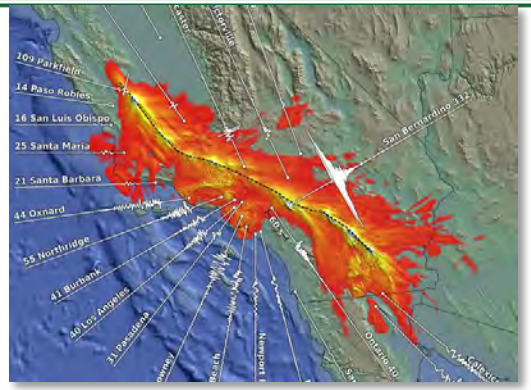

Scalable Earthquake Simulation
Largest simulation of an earthquake ever performed shows a magnitude-8 quake and its impact on the region. Run on 223,074 cores on **Jaguar** opening new territory for earthquake science.

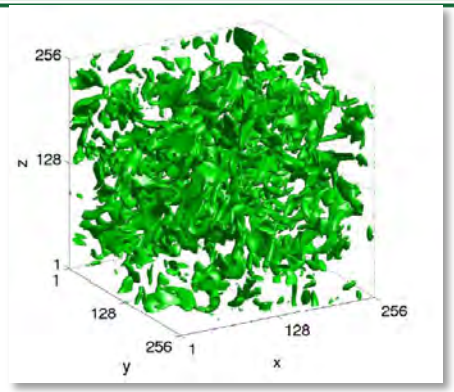

Astrophysical N-body Simulation
An astrophysical N-body simulation with 3,278,982,596 particles using a treecode algorithm shows a sustained performance of 190.5 Teraflops on DEGIMA, a 144 node GPU cluster at Nagasaki U.

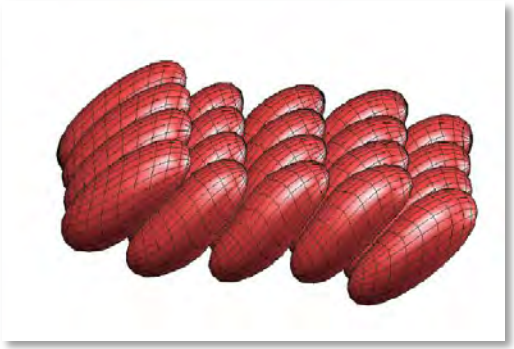

DNS of Blood Flow on 200K Cores
The first high-fidelity petascale direct numerical simulation of blood flow that directly resolves the interactions of 200 million deformable red blood cells in plasma. Runs on GPUs, and achieves 700 Teraflops on 200k cores of **Jaguar**.

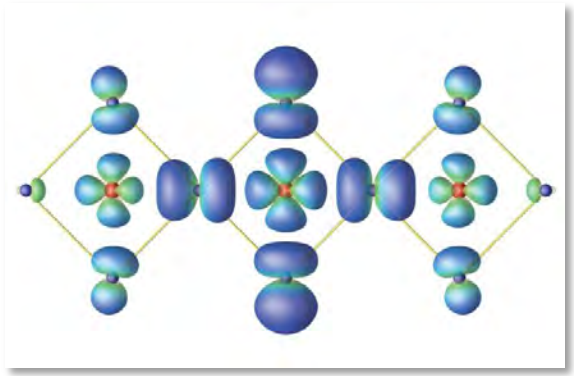

Simulation of Excited Electron States
First principle simulations of excited state properties and strong electron correlations that are based on many-body Green's function approach run on **Jaguar** at over 1.3 Petaflops. .

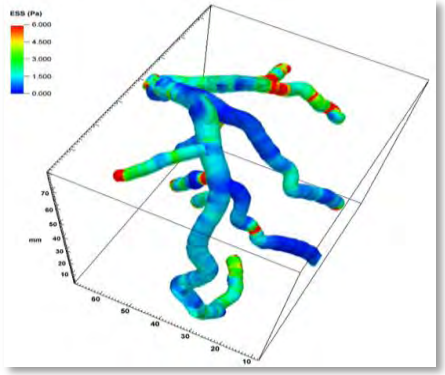

Full Heart-Circulation System
The first large-scale simulation of blood flow in the coronary artieries, with a realistic description of human arterial geometry at spatial resolutions from centimeters down to 10 microns run on Jugene at Jülich.

ESnet Accomplishments

- **Network:**
  - Easily handled an almost 300% increase in network traffic from July 2009 to June 2010
  - 12-month reliability now 99.985% (up from 99.954% in 2007)
  - Purchased dark fiber to provide scalable and cost effective bandwidth to BNL; the dedicated DOE network created will be operational Nov/Dec
- **OSCARS virtual circuits:**
  - Created cloud service between NERSC and JGI, a feat repeatable for other labs
  - Interagency proof-of-concept virtual circuit set up between NASA Ames supercomputing facility and USGS for TOPS climate project http://ecocast.arc.nasa.gov/
  - In past year, ~50% of ESnet traffic used OSCARS circuits
  - Deployed as production infrastructure in the Internet2 and NORDUnet (Scandinavian) networks

Staffing
In Progress

- **To Be posted**
  - Director, Computational Science and Research Partnerships (SciDAC) Division
  - Computer Scientist – GS 14/15

- **Vacancy announcements Closed:**
  - Mathematician/Physical Scientist – GS15
  - Mathematician/Physical Scientist – GS 12/13/14
  - Computer Scientist – GS 15