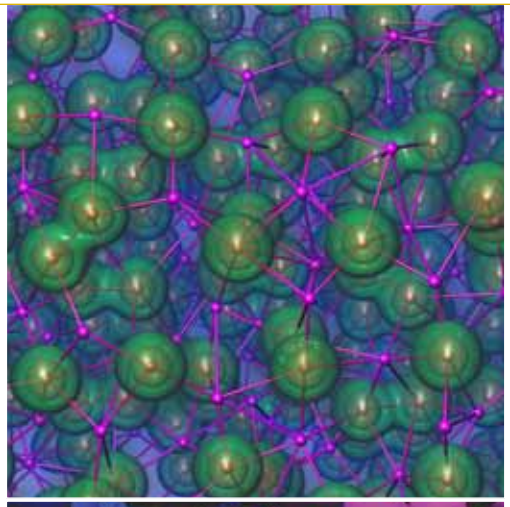

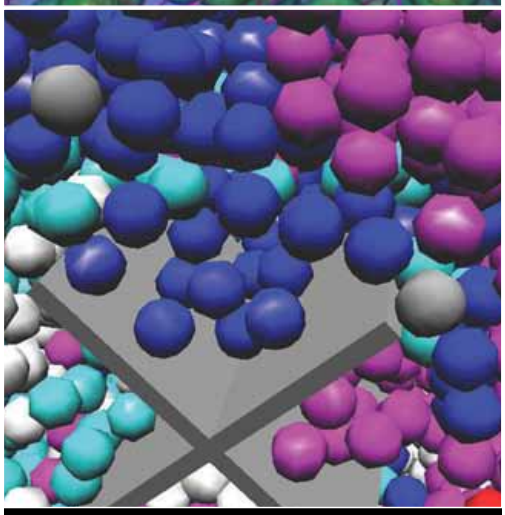

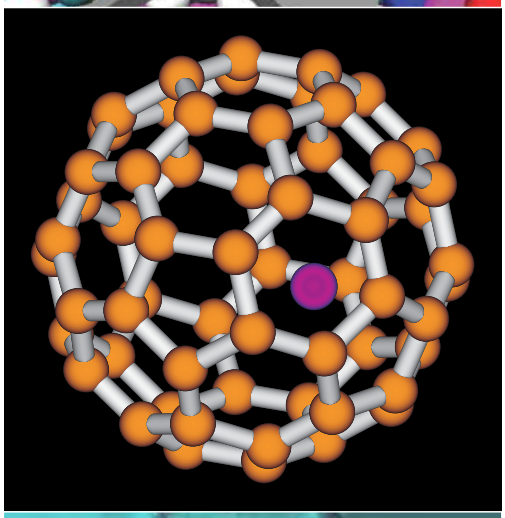

2011 science highlights

Argonne NATIONAL LABORATORY

Argonne Leadership Computing Facility

propelling innovation in science

www.alcf.anl.gov

November 2011

www.alcf.anl.gov

The Leadership Computing Facility Division operates the Argonne Leadership Computing Facility — the ALCF — as part of the U.S. Department of Energy's (DOE) effort to provide leadership-class computing resources to the scientific community.

About Argonne National Laboratory

Argonne is a U.S. Department of Energy laboratory managed by UChicago Argonne, LLC under contract DE-AC02-06CH11357. The Laboratory's main facility is outside Chicago, at 9700 South Cass Avenue, Argonne, Illinois 60439. For information about Argonne and its pioneering science and technology programs, see www.anl.gov.

Availability of This Report

This report is available, at no cost, at http://www.osti.gov/bridge. It is also available on paper to the U.S. Department of Energy and its contractors, for a processing fee, from:

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831-0062
phone (865) 576-8401
fax (865) 576-5728
reports@adonis.osti.gov

Disclaimer

contents

contents

ALCF mission and vision

Propelling Innovation in Science and Technology

Argonne operates the Argonne Leadership Computing Facility (ALCF) for the U.S. Department of Energy's (DOE) Office of Science as part of the larger DOE Leadership Computing Facility strategy. DOE leads the world in providing the most capable civilian supercomputers for science. Argonne researchers work closely with researchers from companies and universities, as well as federal, state, and municipal agencies to help them solve their specific problems, advance America's scientific leadership, and prepare the nation for a better future.

Mission

The Argonne Leadership Computing Facility's (ALCF) mission is to accelerate major scientific discoveries and engineering breakthroughs for humanity by designing and providing world-leading computing facilities in partnership with the computational science community.

Vision

The ALCF strives to be the forefront computational center for extending the frontiers of science by solving key problems for the nation that require innovative approaches and the largest-scale systems.

Supercomputing Resources

Researchers who are awarded allocations of computer time at the ALCF use one of the world's most powerful supercomputers--Intrepid, an IBM Blue Gene/P system. This 557-teraflops production machine features 40,960 quad-core compute nodes; 163,840 processors; and 80 terabytes of memory, yet is highly energy efficient. The ALCF also operates Surveyor, a 4,096-core system used for tool and application porting, software testing and optimization, and systems software development.

In addition, Eureka, a visualization supercomputer allows researchers to explore and visualize the flood of data they produce with Intrepid, and Gadzooks offers a test and development system for visualization.

User Services

ALCF catalysts, performance engineers, operations staff, and a data analytics and visualization team provide users with in-depth expertise and ongoing help in using the ALCF's computer systems. They provide users with information on ALCF services and resources, technical details on the IBM Blue Gene/P architecture, as well as hands-on assistance in porting and tuning applications.

Awards Honor ALCF Achievements

Argonne Leadership Computing Facility (ALCF) research and resources were honored with the following prestigious awards in 2011:

- Research at ALCF named Gordon Bell Prize finalist (multiscale brain blood flow simulations led by George Karniadakis of Brown University)—one of five finalists selected this year.
- Two DOE OASCRs awarded for Argonne computer visualizations of "Blood Flow: Multi-scale Modeling and Visualization" and "Modeling Early Galaxies Using Radiation Hydrodynamics" at the SciDAC Annual Conference in July.
- Intrepid, the IBM Blue Gene/P system at the ALCF, ranked No. 1 on the Graph 500 list for the second consecutive year.
- The IBM Blue Gene®/Q again topped the Green500 list. The BG/Q prototype at #1 on this list is different from the one that topped the 2010 list in that it delivers significantly better performance with the same number of cores and only a marginal increase in power consumption.

science director's message

Paul Messina
Director of Science
Argonne Leadership
Computing Facility

Partnering Yields Promising Opportunities for Scientific Innovation

Partnering with the computational science community is an essential strategy for carrying out the Argonne Leadership Computing Facility (ALCF) mission. Our users tackle complex science and engineering simulations that often benefit from interactions with our staff on a variety of issues. The topics range from detailed information about our leadership computing systems to algorithms that scale well on our systems, data mappings that are efficient for the application and the computer architecture, how to display the results of the simulations, and sometimes even how best to model the phenomena being investigated. The ALCF staff's computational science expertise is essential for being effective partners with the scientists and engineers who use our systems.

Our partnership activities extend to working with teams who are planning to use leadership-scale computers for the first time, or who are not familiar with the characteristics of our resources and how to use them efficiently. Although Intrepid, our IBM Blue Gene/P supercomputer, is now several years old, it continues to attract new projects and new research teams. These users typically gain access to ALCF systems by submitting a proposal for a Director's Discretionary allocation. They may attend our "Getting Started" workshops and webinars on how to apply for INCITE awards. When appropriate, our catalysts and performance engineers suggest different algorithms or implementations that will enable the applications to scale to the levels needed to have a shot at receiving an INCITE or ALCC award. The Operations staff may work with user teams to ensure that they understand the job scheduling and disk use policies.

Partnering with new project teams has been successful. Former recipients of Director's Discretionary awards make up a third of our current INCITE and ALCC projects. That is a high success record, given the stiff competition for those awards. And it is worth noting that some INCITE projects that started this year (CY2011) are using large fractions of Intrepid for production runs—even the entire system in some cases.

Propelling innovation in science and engineering is a continuous endeavor. Mira, the 10-petaflops IBM Blue Gene/Q that will be delivered in 2012, presents new opportunities for innovation, given its speed, memory size, and disk storage capacity. Through our Early Science Program we are partnering with 16 projects on enhancing their applications to take advantage of Mira's features and adding the functionality required to achieve their ambitious science goals.

As we move towards the next major challenge in high-performance computing—exascale levels of computation—no doubt the benefits of partnering will grow in tandem with the size of these extraordinary systems.

world-class resources

World-Class Resources for World-Changing Research

The Argonne Leadership Computing Facility (ALCF) provides leading scientists with next-generation, high-performance computing resources for breakthrough research to address global challenges.

ALCF resources are accessible through the following programs.

Innovative and Novel Computational Impact on Theory and Experiment (INCITE) Program

Provides allocations to computationally intensive, large-scale research projects to significantly advance key areas in science and engineering. The program encourages proposals from universities, other research institutions, and industry.

ASCR Leadership Computing Challenge Program (ALCC)

Allocates resources to projects with an emphasis on high-risk, high-payoff simulations in areas directly related to the U.S. Department of Energy's energy mission, national emergencies, or for broadening the community of researchers capable of using leadership computing resources. The program is open to scientists from the research community in academia and industry.

Early Science Program (ESP)

Offers researchers preproduction hours on Mira, the ALCF's next-generation, 10-petaflops IBM Blue Gene/Q system.

Director's Discretionary

Gives "start up" awards to potential future INCITE projects so that they can achieve computational readiness. Projects must demonstrate a need for leadership-class resources. Awards may be made year round to industry, academia, laboratories, and others.

For further details about these programs, contact info@alcf.anl.gov.

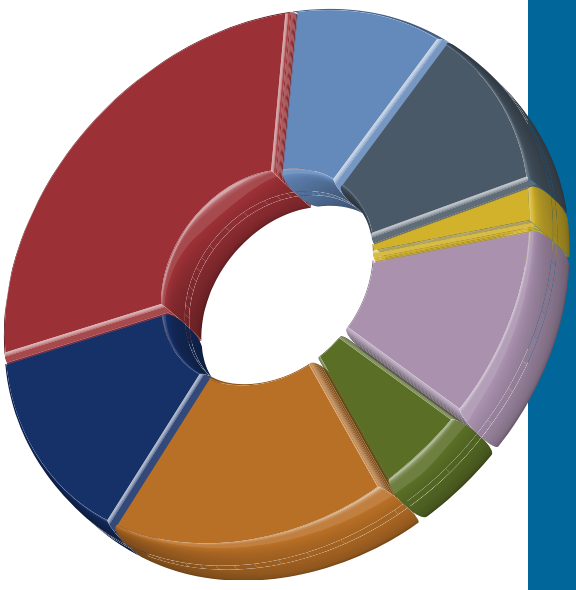

Scientific domains represented in INCITE allocations awarded for January 1, 2011 through December 31, 2011.

ALCF High-Performance Computing

Intrepid

Production scientific and engineering computing

- 40,960 quad-core compute nodes (163,840 processors)
- Memory: 80 TB
- 6.5 PB disk storage
- Peak Performance: 557 TF

Eureka

Converts data output to visual representations

- 100 dual quad-core servers
- 200 Quadro FX5600 GPUs in the S4s
- Memory: 3.2 TB+ RAM
- Peak Performance: More than 111 TF

Data Storage

Blue Gene/P data systems

- 640 I/O nodes
- 16 storage area networks (SANS)
- 7,680 disk drives
- 7.6 PB raw storage
- 88 GB/second transfer speed
- 2 parallel file systems—PVFS and GPFS

Networking

Internally, Intrepid uses five different networks. The main one is a 3D torus with an aggregate bandwidth of 5.1 Gbps to each node. Connectivity to outside institutions is provided by 20 Gbps of bandwidth to commercial and research networks.

biological sciences

Breakthroughs in Protein Structure Calculation and Design

Protein structure calculation and design challenges have direct relevance for biomedicine and provide stringent and objective tests of the understanding of the fundamental underpinnings of molecular biology. Structure prediction is key to understanding the function and interactions of biomolecules and provides the insights necessary to design new molecules with new and useful functions. Faster and more accurate predictions of structure can yield significant cost and time savings.

In collaboration with other researchers, David Baker from the University of Washington has achieved a number of exciting breakthroughs in protein structure calculation and design.

The researchers developed a new approach for computational analysis of Nuclear Magnetic Resonance (NMR) data that pushes the limits of protein size that can be structurally solved from NMR spectroscopy data. This is a very significant step forward since many of these larger proteins are not amenable to analysis by X-ray crystallography, and thus, NMR remains the only way to obtain their structure. ALCF staff assisted the project team with developing a system for automating ensemble runs. This system was critical to accomplishing the ensemble runs.

Furthermore, based on their previous work showing that accurate structures could be obtained from very sparse NMR data sets, the researchers reasoned that structure prediction methods guided by even very noisy density maps might be able to improve a poor molecular replacement model before applying crystallographic model-building techniques. They developed an approach in which electron density maps generated from molecular replacement solutions for each of a series of starting models are used to guide energy optimization by structure rebuilding, combinatorial side chain packing, and torsion space minimization. New maps are generated using phase information from the energy-optimized models most consistent with the diffraction data, and success is monitored through the free *R* factor. Using this approach, researchers were able to solve 8 of 13 challenging cases.

In another effort addressing the challenges of influenza, researchers have devised a computational method for designing protein-protein interactions de novo and used the method to design high-affinity binders to the conserved stem region on influenza hemagglutinin (HA). The results suggest that de novo computational design of antiviral proteins is feasible. HA is a prime candidate for drug development, as it is the major player in viral invasion of cells lining the respiratory tract.

INCITE Allocation:

30 Million Hours

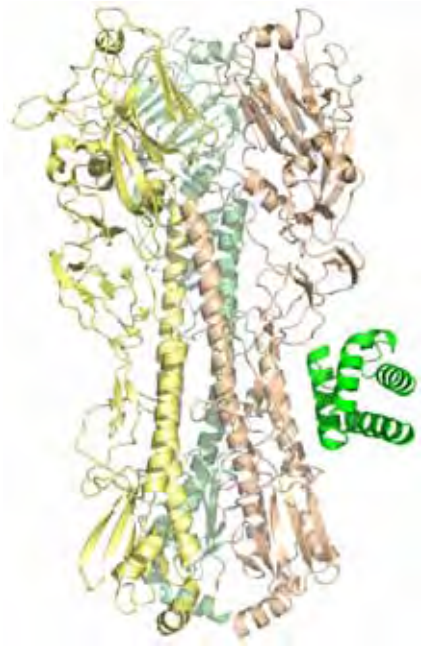

A crystal structure of Spanish influenza hemagglutinin (trimer) bound to a computationally designed binder (green). The computationally designed protein binds hemagglutinin with high affinity and blocks the conformational changes in hemagglutinin that underlie its function, thereby neutralizing influenza. The crystal structure is virtually identical to the designed model, validating the design approach. [Fleishman et al., (2011) Science 322: 816.] Image credit: Sarel Fleishman, Weizmann Institute of Science, Israel.

Contact ▸ David Baker
University of Washington | dabaker@u.washington.edu

Comparing Neuronal Spike Exchange Methods on the ALCF's Blue Gene/P

In simulations of large-scale spiking neural networks, the computational units—neurons—generate logical events—spikes—that are sent to thousands of other neurons with constant propagation delay that can be different for different connections. During time intervals between input events, the neuron is typically defined by a system of continuous ordinary differential equations, along with a threshold detector, which watches one of the states and determines when the output event is generated.

For neural network simulations on parallel machines, interprocessor spike communication can be a significant portion of the total simulation time. The performance of several spike exchange methods using Intrepid, the Blue Gene/P supercomputer at the Argonne Leadership Computing Facility, has been tested with 8 K to 128 K cores using randomly connected networks of up to 32 M cells with 1 k connections per cell and 4 M cells with 10 k connections per cell, i.e., on the order of $4 \cdot 10^{10}$ connections. (K is 1024, M is 1024^2, and k is 1000.)

A research team led by Michael Hines of Yale University found that the Deep Computing Messaging Framework (DCMF) gave three-fold better performance than the best MPI implementation. DCMF spike communication time is negligible since transmission overlaps with computation and is handled by a direct memory access controller. All simulations were carried out using the NEURON Version 7.2 simulation program.

The research team discovered that ideal performance scaling will be ultimately limited by imbalance between incoming processor spikes between synchronization intervals. Thus, counter-intuitively, maximization of load balance requires that the distribution of cells on processors should not reflect neural net architecture but be randomly distributed so that sets of cells which are burst firing together should be on different processors with their targets on as large a set of processors as possible.

Director's Discretionary Allocation:

770,000 Hours

Cores	Cells	Conn	MPI_Allgather		MPI_ISend		DCMF_Multicast		Record-Replay	
			comp	run	comp	run	comp	run	comp	run
8K	2M	1k	6.76	7.84	22.3	24.9	7.30	7.74	6.30	6.99
16K	2M	1k	4.01	5.44	11.5	13.5	3.67	4.00	3.17	3.62
32K	2M	1k	3.00	4.69	5.90	7.40	1.84	2.07	1.58	1.91
64K	2M	1k	2.98	5.05	2.87	4.20	0.957	1.25	0.808	1.03
128K	2M	1k	---	---	---	---	0.497	0.64	0.417	0.608
8K	1/4M	10k	5.23	5.64	18.5	21.5	6.50	6.90	5.81	6.58
16K	1/4M	10k	3.16	3.78	11.1	14.6	3.76	4.18	3.35	4.07
32K	1/4M	10k	2.04	3.11	6.37	9.37	2.04	2.39	1.79	2.64
64K	1/4M	10k	1.58	3.47	6.19	6.99	1.36	1.42	0.948	1.77
128K	1/4M	10k	---	---	---	---	0.573	0.799	0.494	1.39

Strong scaling performance of NEURON (seconds) on BG/P in Quad mode for simulation runs lasting 200 ms. Comp and run columns are in seconds. The Conn column refers to the average number of connections per cell. The MPI_Allgather method used a buffer size of 10 compressed spikes. The MPI_Isend and DCMF_Multicast methods utilize the two phase and two subinterval variants. The Record-Replay method utilizes the one subinterval variant. Image credit: Michael Hines, Yale University.

Contact ▸ Michael Hines
Yale University | michael.hines@yale.edu

biological sciences

Protein-Ligand Interaction Simulations and Analysis

To rationally design drugs to fight off diseases, scientists must first understand how a small molecule interacts with a given protein target. This process is the basis of how drugs work therapeutically in our bodies, by altering or disrupting essential enzymes required for pathogens to proliferate.

A fundamental goal of research in this arena is to computationally predict whether a given molecule will bind to a target and, if so, with what affinity. While a well-defined paradigm exists for predicting the ability of small molecules to interact with a protein surface, current methods have both theoretical and computational limitations. Access to leadership-class computing systems, like the resources at the Argonne Leadership Computing Facility, allows scientists to conduct studies using highly advanced, physics-based methodologies, such as MM-GBSA and FEP/MD-GCMC, whose prohibitive runtimes have limited their widespread use. Implementing these methods and making them more accessible to researchers will allow for a fuller application of the biological and pharmaceutical promise that molecular simulation holds.

With a multi-year INCITE award, Andrew Binkowski, with the Center for Structural Genomics of Infectious Diseases, Argonne National Laboratory, leads a team of researchers using the multi-stage Protein Ligand Interaction Atlas computational pipeline to conduct a comprehensive analysis of protein binding domain and small molecule interactions including receptor analysis, protein-ligand docking, and binding free energy calculations.

In collaboration with experimental researchers around the world, Binkowski's team conducts computer-aided drug discovery of high-value biomedical targets, including human pathogens, bioterrorism agents, and human disease-related proteins. The predicted computational results are validated experimentally. Using X-ray crystallography, they provide an important feedback mechanism necessary to evaluate the predictive power of biomolecular simulations.

INCITE Allocation:

50 Million Hours

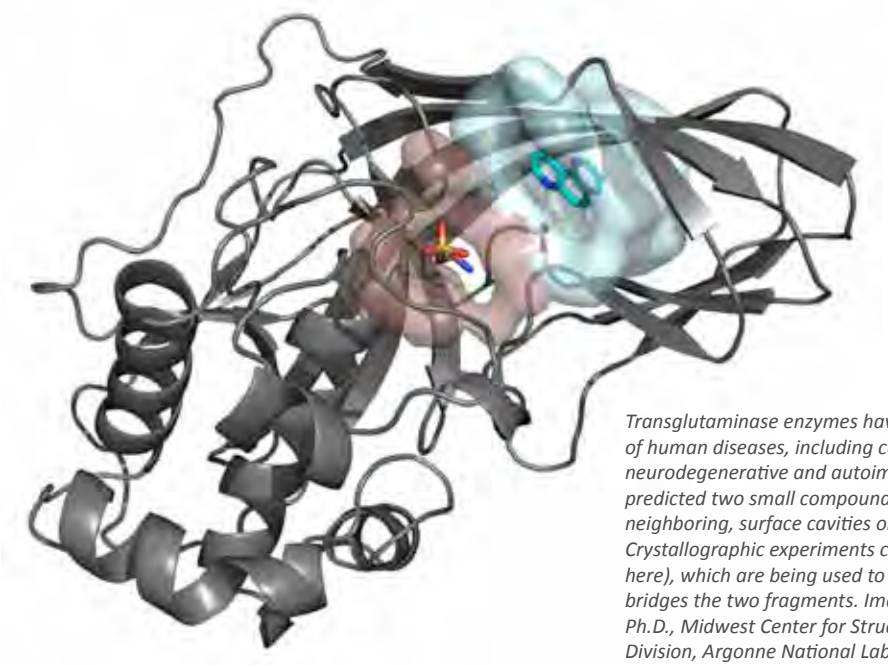

Transglutaminase enzymes have been linked to a number of human diseases, including cancers, tissue fibrosis, and neurodegenerative and autoimmune diseases. Virtual screening predicted two small compounds to bind to distinct, but neighboring, surface cavities on a transglutaminase enzyme. Crystallographic experiments confirmed the predictions (shown here), which are being used to develop a larger compound that bridges the two fragments. Image credit: T. Andrew Binkowski, Ph.D., Midwest Center for Structural Genomics & Biosciences Division, Argonne National Laboratory; Center for Structural Genomics of Infectious Diseases & Computation Institute, The University of Chicago.

Contact ▸ Andrew Binkowski
Argonne National Laboratory | abinkowski@anl.gov

Research on Multiscale Blood Flow Models Named Gordon Bell Prize Finalist

A research endeavor focused on multiscale brain blood flow simulations led by George Karniadakis from Brown University has been named a Gordon Bell Prize finalist—one of five finalists selected this year. Administered by the Association of Computing Machinery (ACM), the Prize is awarded annually to recognize outstanding achievement in high-performance computing applications.

To treat diseases involving disruptions of blood flow to the brain, doctors must first understand how multiple scales of blood vessel networks work within the brain, both alone and together. Currently, the researchers are using the supercomputing resources at the Argonne Leadership Computing Facility (ALCF) to conduct an INCITE project focused on creating "multiscale" models that show the interconnected workings of multiple scales of the brain's blood vessels. The simulations use MRI data of actual patients for generating the geometry and meshes of the vasculature. The research is an extension of an earlier ALCF Discretionary project, in which the runs were conducted for Karniadakis's Gordon Bell submission. Multiscale models provide doctors with a more realistic picture of blood flow, and the greatest hope for the development of new lifesaving treatments for such blood-related diseases as brain aneurisms, sickle-cell anemia, and cerebral malaria.

Karniadakis's goal is to enhance the ability to chart the behavior of individual red blood cells and improve predictive capabilities in medical procedures. The team is also developing software and algorithms pertaining to blood flow for petascale supercomputers.

Main research accomplishments to date are: (i) simulation of initial stages of clot formation; (ii) simulation of blood flow, and healthy and diseased red blood cells (RBCs) at different stages of cerebral malaria and sickle cell anemia; (iii) modeling of the glycocalyx layer, which plays an important role in protecting the arterial wall; and (iv) modeling of the microcirculation and distribution of RBCs in Y-shaped bifurcating arteries.

INCITE Allocation:

50 Million Hours

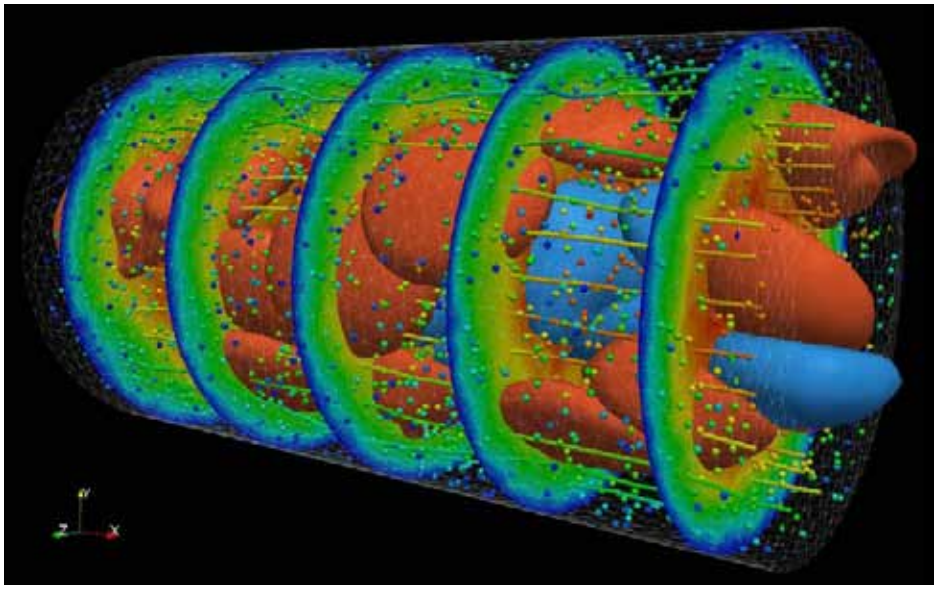

A flow of healthy (red) and diseased (blue) blood cells with a Dissipative Particle Dynamics method.

Contact ▸ George Karniadakis
Brown University | george_karniadakis@brown.edu

biological sciences

Simulation and Modeling of Membranes' Interactions with Unstructured Proteins

Researchers led by Igor Tsigelny from the University of California–San Diego (UCSD) are seeking to understand why some proteins aggregate and what structural or sequence features may increase their propensity to do so, how proteins penetrate membranes to form pore-like structures, and what role protective molecules play in influencing protein aggregation. The findings of this study offer important implications for identifying the molecular basis of Parkinson's disease and effective ways to treat it.

Through modeling on Intrepid, the IBM Blue Gene/P supercomputer at the Argonne Leadership Computing Facility, the team has produced three drug candidates with high potential for treating the disease by stopping or preventing pore formation. Each theoretical design was tested in the laboratory and found to be effective. Patents for all three drug candidates are pending.

Simulations conducted by the UCSD scientists also have shown the clear creation of annular protein aggregates that can penetrate to the bacterial membranes and create the pores in them. Such pores are open for absorption of radionuclides.

The team's exploration of pore formation in bacterial membranes may also yield answers to problems associated with environmental waste by reducing radionuclides—atoms that emit harmful radioactivity—and eliminating environmental waste through absorption. Working with scientists in Ukraine, the team is organizing an initiative in which it will conduct experimental validation of efforts to clean radionuclides from water in Chernobyl.

To this end, the researchers are building a program package for the study of proteins' aggregation and interaction with membranes. This set of programs will help equip future research teams to study other biomedical and environmental problems in which the membrane contact of proteins is involved.

INCITE Allocation:

4 Million Hours

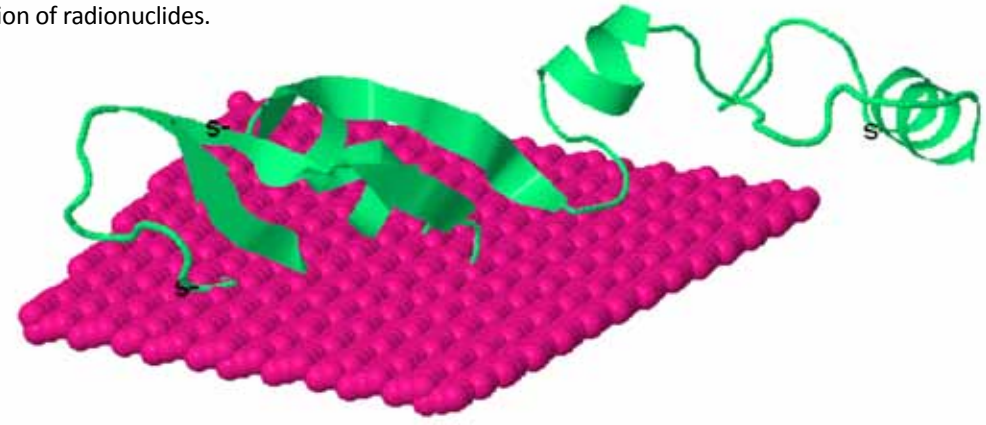

One of the selected MD conformers of MccE492 microcin is analyzed by the program MAPAS for the membrane-contacting surface location. The predicted orientation on the membrane surface is shown. This position is used for modeling of possible oligomerization of this microcin with further pore creation in bacterial membranes.. Image credit: Yuriy Sharikov and Igor Tsigelny, Department of Neurosciences, UCSD.

Contact ▸ Igor Tsigelny
University of California–San Diego | itsigeln@ucsd.edu

NAMD – The Engine for Large-Scale Classical MD Simulations of Biomolecular Systems Based on a Polarizable Force Field

Classical simulations of biomolecular systems are critical in life sciences. NAMD is a free, high-performance code used to generate these simulations. As petascale resources become available, researchers aim to enable the next level of simulation methods and scientific questions, and to ensure NAMD performance across platforms. To enhance the sampling efficiency beyond that of brute-force molecular dynamics (MD) simulations, researchers will implement several advanced strategies based on multiple copies, such as replica-exchange MD (REMD) and/or Hamiltonian-tempering (H-REMD. Because the quality and accuracy of the potential function is critical for meaningful MD simulations, researchers will implement a new all-atom polarizable force field for proteins, nucleic acids, and lipids, representing the state-of-the-art for next-generation simulation studies in computational biophysics and chemistry. The force field will be generally supported by NAMD, including on the IBM Blue Gene/Q. As such, it will be a tool of unprecedented accuracy and one of great use to the entire scientific community.

The project will include simulations covering a wide range of canonical and non-canonical DNA and RNA molecules for which a wealth of experimental data exists. Researchers will also examine the performance of the new force field for a suite of key problems about ions, peptides, proteins, and lipid membranes, where induced polarization is anticipated to be critical.

Standard (non-threaded) NAMD has been ported to Blue Gene/Q by researchers at IBM for initial tests. A threaded version has been developed by the Theoretical and Computational Biophysics Group of the Beckman Institute for Advanced Science and Technology at the University of Illinois, Urbana-Champaign. As a stepping stone to Blue Gene/Q, performance analysis of the threaded NAMD is under way on Blue Gene/P. Development will continue on Blue Gene/Q hardware as it becomes available.

Early Science Program Allocation:

80 Million Hours

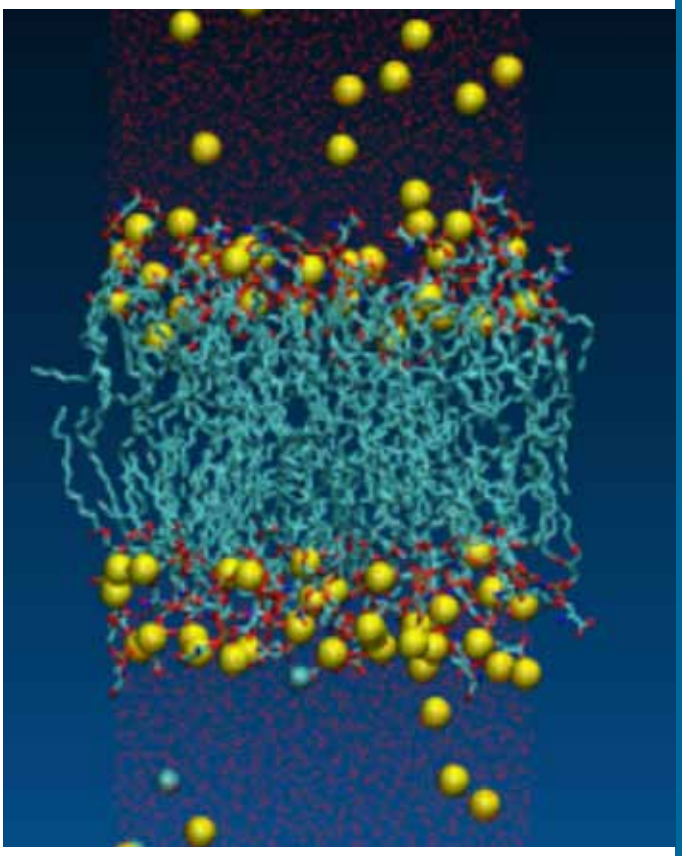

Typical configuration of a lipid bilayer membrane solvated by a salt solution. A polarizable force field is essential for meaningful simulations of such a system.

Contact ▸ Benoit Roux
The University of Chicago | roux@uchicago.edu

chemistry

Performing the Largest Unstructured Large-Eddy Simulation of a Real, Full Combustion Chamber

The increase of computer power has allowed science to make important strides in a variety of domains such as plasma studies, biomechanics, and molecular dynamics. With access to the INCITE program, researchers from CERFACS (the European Centre for Research and Advanced Training in Scientific Computation) have been able to perform top-of-the-line quality simulations on highly complex cases in their goal towards the fully numerical modeling of a real combustor.

This research is focused on Large Eddy Simulation (LES) of gas turbine engines with the inclusion of liquid phase phenomena. CERFACS has performed simulations and validation of two-phase flow experiments. In parallel, taking advantage of the leadership-class computer available at the Argonne Leadership Computing Facility, the researchers have performed the largest unstructured LES done to date of a real, full combustion chamber (330 million elements) on more than 16K cores. This simulation contributes to the validation of the LES approach when dealing with combustion instabilities. In these cases, the effects of mesh refinement are a highly critical point that was validated during the Stanford Center for Turbulence Research (CTR) 2010 Summer Program. A second mesh independency validation was performed, but this time it used a simpler, two-phase-flow single burner with three levels of refinement (4-, 8-, and 16-million elements). Research results are featured in a recent paper, "Using LES to Study Reacting Flows and Instabilities in Annular Combustion Chambers," published in *Flow Turbulence & Combustion,* Springer Science & Business Media, September 2011. Evaluation of the unbalance observed in Lagrangian simulations remains to be performed.

INCITE Allocation:

10 Million Hours

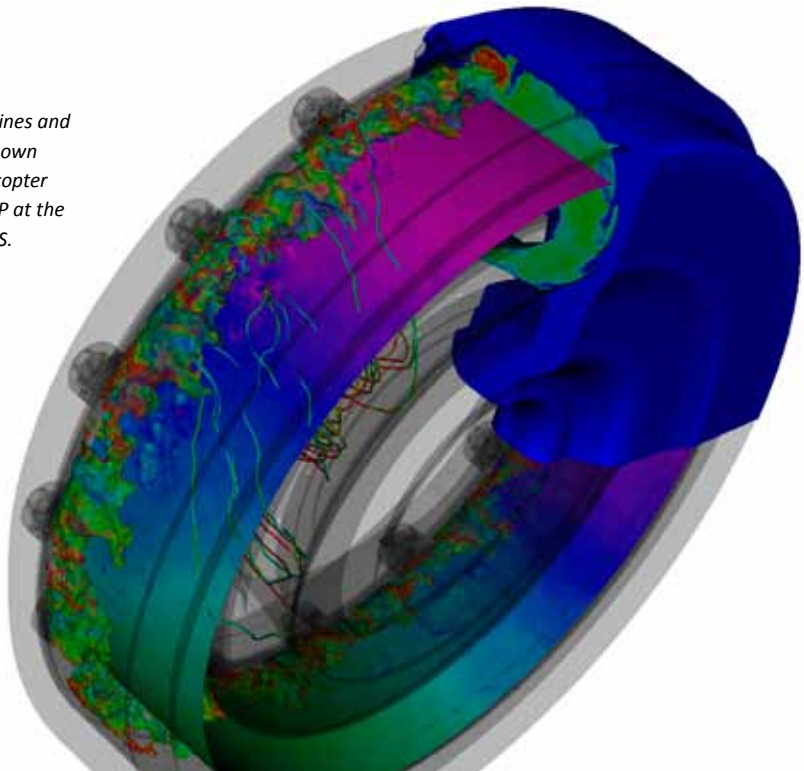

Fields of temperature and pressure, along with stream lines and a temperature isosurface representing the flame, are shown in a 330-million-elements simulation of a complete helicopter combustion chamber performed on the IBM Blue Gene/P at the ALCF. Image credit: Pierre Wolf, Turbomeca and CERFACS.

Contact ▸ Thierry Poinsot
CERFACS | Thierry.poinsot@cerfacs.fr

Potential Energy Surfaces for Simulating Complex Chemical Processes

Large-scale electronic structure theory can provide potential energy surfaces and force fields for simulating complex chemical processes important for technology and biological chemistry. This theory can reveal the details of these processes. An especially important challenge is to obtain accurate energies and stationary points for systems whose electronic structure is complex.

During this past year, University of Minnesota researchers extended their work to larger systems—metallofullerenes. Basically, they are conducting calculations to determine what structural factors make the metallofullerenes work as molecular electronic switches. From both a fundamental, theoretical point of view and a practical one, it is essential to find the energy minima and saddle points and to map the topography of the seams of conical intersections in these fascinating systems.

Staff at the Argonne Leadership Computing Facility (ALCF) worked with the project team to improve the performance of the key algorithm MCSCF (multi-configurational self-consistent field) in the GAMESS code. This improvement enabled the researchers to perform large-scaled calculations very efficiently.

For example, a state-averaged complete active space self-consistent field (CASSCF) single-point energy calculation of the Ca@C60 system (the "buckyball" fullerene interacting with a single Ca atom) running on 8,192 cores on Intrepid, the Blue Gene/P supercomputer at the ALCF, takes only ~40 minutes. This time is twice as fast as the same calculation using the standard version of the GAMESS code. This difference is significant considering that typically researchers need to run hundreds of such calculations. The researchers expect an even greater improvement in the performance for bigger problems (e.g., a larger wave function and/or larger molecular system).

INCITE Allocation:

15 Million Hours

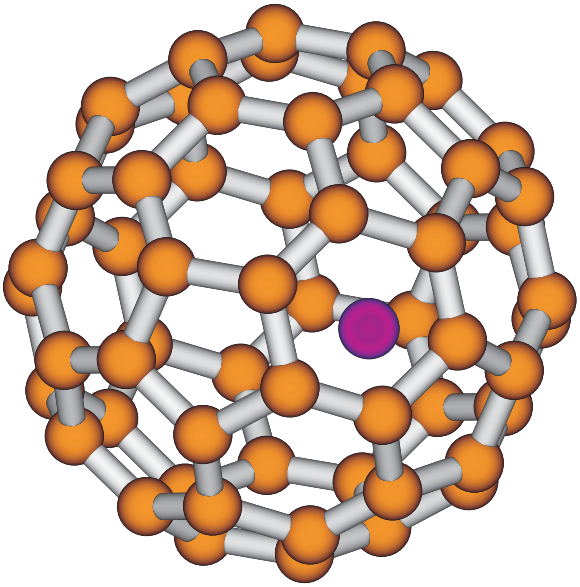

Researchers performed SA-CASSCF calculations of B@C_{60}—boron atom inside the fullerene cage shown here. The boron atom occupies a shallow energy-minimum "off center" due to charge transfer from the boron atom to the case wall. For these calculations, ALCF catalyst on the project, Graham Fletcher, developed a version of the GAMESS code for Intrepid that runs twice as fast as the standard version of GAMESS.
Image credit: Osanna Tischenko, University of Minnesota.

Contact ▸ Don Truhlar
University of Minnesota | truhlar@umn.edu

chemistry

Predicting Bulk Properties of Water Systems

INCITE Allocation:
10 Million Hours

Ames Laboratory researchers are using high-quality electronic structure theory, statistical mechanical methods, and massively parallel computers to predict bulk properties of water systems that require high fidelity. This molecular-scale challenge is of critical importance to national scientific issues such as global warming and the environment.

The research under way is targeting very large systems whose sizes can only be addressed with massive computational resources. The researchers are using GAMESS and NWChem codes in simulations on Intrepid, the IBM Blue Gene/P housed at the Argonne Leadership Computing Facility (ALCF). Both codes are without question the two most broadly distributed scalable electronic structure systems. The research is focused on understanding the molecular-level dynamics of water, the formation of aerosols important in cloud formation, and the interactions of dendrimers with ligands of environmental importance. In each of these applications, the underlying research is setting a new standard for the predictive computation of bulk properties.

Large clusters of water are now amenable to molecular dynamics simulations with the large improvements that have been made in the Fragment Molecular Orbital (FMO) method in GAMESS. Graham Fletcher at the ALCF has been instrumental in making these improvements. As the results below show, 32 racks of the BG/P system can be very effectively used for these simulations.

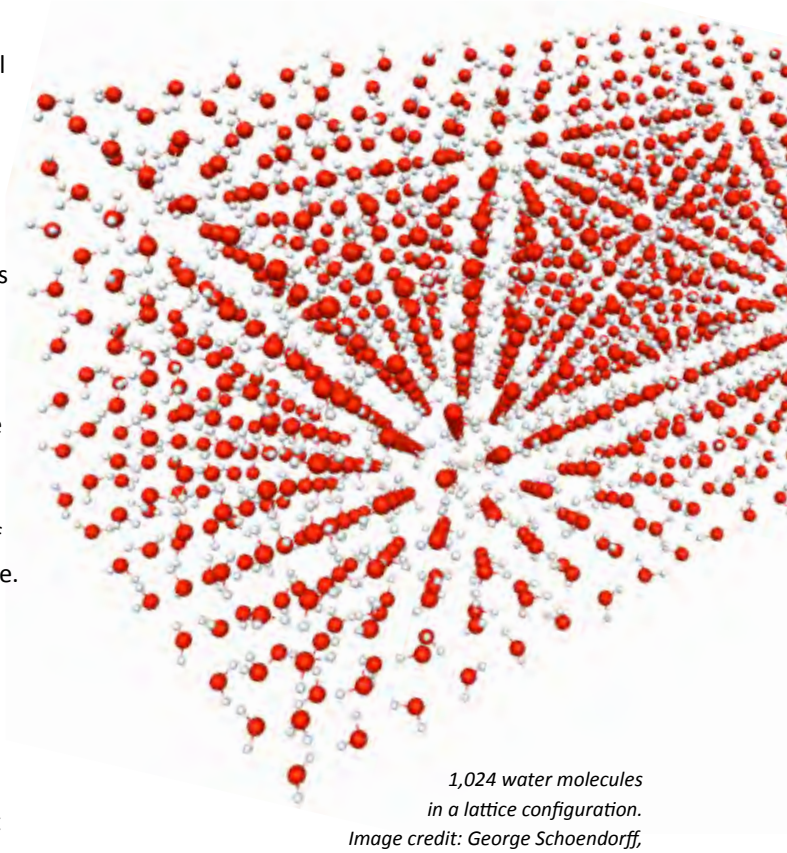

1,024 water molecules in a lattice configuration. Image credit: George Schoendorff, Iowa State University.

		Racks:	1	2	4	8	16	32
		Cores:	4,096	8,192	16,384	32,768	65,536	131,072
Waters	Atoms	Basis Functions	Wall Time (minutes)					
128	384	5,504	8.6	4.8	2.7	1.8		
256	768	11,008	19.8	10.5	5.8	3.4	2.2	
512	1536	22,016		28.9	15.4	8.6	4.9	3.2
1024	3072	44,032			41.1	22.0	12.2	7.1

The performance of FMO2-MP2 force calculations on Intrepid, the Blue Gene/P system at the ALCF. The atomic basis set is aug-cc-pVDZ.

Contact ▸ Theresa Windus
Ames Laboratory | theresa@fi.ameslab.gov

Simulations of Deflagration-to-Detonation Transition in Reactive Gases

Hydrogen is the most abundant element in the universe. It is an environmentally friendly, clean fuel that has the potential to reduce the nation' s dependence on foreign oil, improve the environment, and boost our economy. The challenge with hydrogen fuel is bringing it safely into our everyday lives. This fuel is very energetic and prone to accidents. In certain conditions, a hydrogen-oxygen mixture can react violently and detonate. This is a serious problem in nuclear reactor safety as well.

The process of transition from slow burning to a detonation is called a deflagration-to-detonation transition or DDT. Predicting DDT in various combustion settings remains an outstanding combustion theory problem. Led by Alexei Khokhlov with The University of Chicago, the High Speed Combustion and Detonation (HSCD) project uses Argonne Leadership Computing Facility (ALCF) resources to perform first-principles, reactive flow Navier-Stokes fluid dynamic simulations of DDT. These extremely detailed computer models allow researchers to safely study how hydrogen burns. Ultimately, this knowledge may help make hydrogen a viable fuel alternative for powering vehicles and other industrial applications.

INCITE Allocation:
18 Million Hours

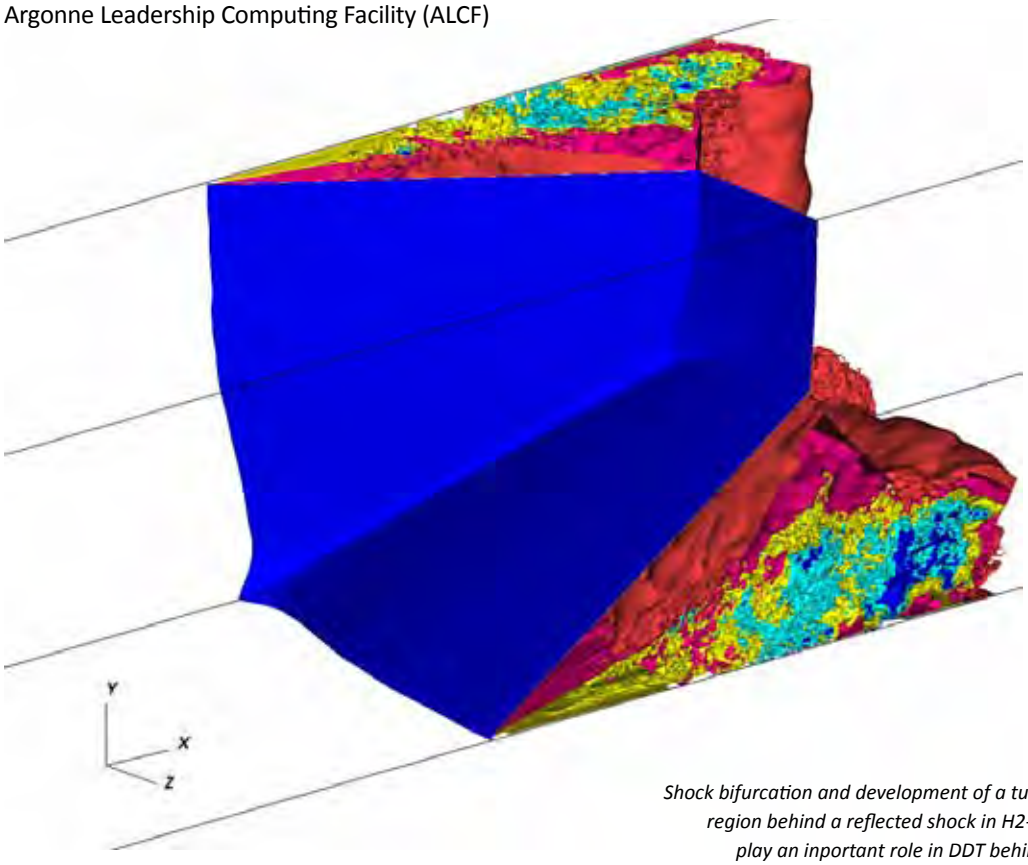

Shock bifurcation and development of a turbulent recirculation region behind a reflected shock in H2-O2. This region may play an inportant role in DDT behind a reflected shock. Image credits: A. Khokhlov and B. Clifford (UChicago), J. Austin and A. Knisely (UIUC), C. Bacon and S. Aithal (ANL).

Contact ▸ Alexei Khokhlov
The University of Chicago | ajk@oddjob.uchicago.edu

computer science

ExM: System Support for Extreme-scale, Many-Task Applications

Exascale computers will enable and demand new problem solving methods that involve many concurrent, interacting tasks. Running such "many-task" applications efficiently, reliably, and easily on extreme-scale computers is challenging. System software designed for today's mainstream, single program multiple data (SPMD) computations does not typically scale to the demands of many-task applications.

Michael Wilde, Mathematics and Computer Science Division, Argonne National Laboratory, is leading the ExM project to produce advances in computer science and usable middleware that enable the efficient and reliable use of exascale computers for new classes of applications. The project will both accelerate access to exascale computers for important existing applications and facilitate the broader use of large-scale parallel computing by new application communities for which it is currently out of reach.

The goals of the ExM project are to achieve the technical advances required to execute many-task applications efficiently, reliably, and easily on petascale and exascale facilities. ExM researchers are developing middleware that will enable new problem solving methods and application classes on these extreme-scale systems. To address these demands, the ExM project is designing, developing, applying, and evaluating two new system software components: the ExM data store and ExM parallel evaluator. The data store utilizes both memory-based filesystems and distributed hash tables to facilitate data exchange between tasks. The parallel evaluator breaks down scalability barriers by permitting an unprecedented rate of task creation and management operations in support of the Swift functional data-flow-based programming model (http://www.ci.uchicago.edu/swift). This implicitly parallel programming model is well suited for the upper-level logic of many prospective exascale applications ranging from climate model analysis to molecular biology to uncertainty quantification and extreme-scale ensemble studies.

Key accomplishments to date include the following: The *Jets* prototype extended the many-task model to a hybrid *many-parallel* (MPI) task model. The *Turbine* prototype, using ADLB, showed encouraging scalability and confirmed that exascale performance goals are reachable. The *AME* "anyscale" many-task engine and store measured Blue Gene/P scaling and data exchange to the 16 K-core level. *MosaStore* on the Blue Gene/P and other clusters is creating a model of a virtual data store. Researchers are already evaluating ExM tools on three science applications (earthquake simulation, image processing, and protein/RNA interaction). For more details about the project, visit the ExM project website (http://www.mcs.anl.gov/exm).

Director's Discretionary Allocation:

3 Million Hours

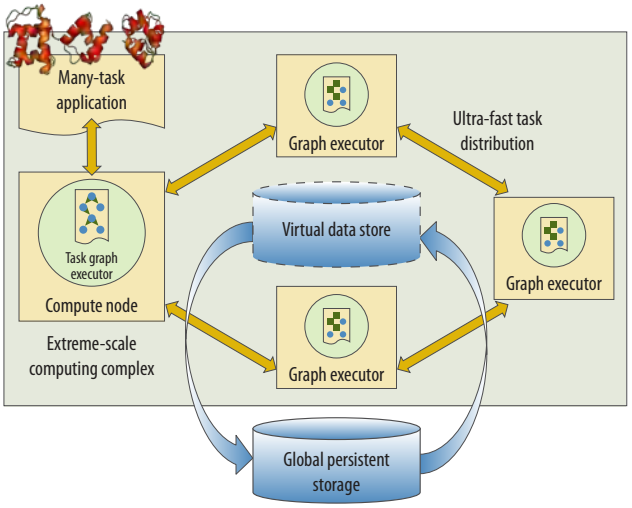

The ExM parallel data-flow program evaluator and virtual distributed data store are creating the task management and data exchange scalability needed to support exascale applications. Image credit: ExM project team.

Contact ▸ Michael Wilde
Argonne National Laboratory | wilde@mcs.anl.gov

How Can More Intricate Climate Models Help Curb Global Warming?

The effects of climate change are apparent in degrading air quality, intensified tropical storms, and the resulting destruction of coral reefs that protect the coasts from erosion and destructive waves. Global warming increases the occurrence of droughts, heat waves, wildfires, and floods. Scientists must improve the understanding of the impact of global warming so that society can optimally address climate adaptation considerations.

Advanced computation, like that possible on the Blue Gene/P at the Argonne Leadership Computing Facility (ALCF), allows researchers at the DOE laboratories and National Center for Atmospheric Research (NCAR) to develop more complex and intricate climate models. The vital information these improved models provide will help guide environmental policy.

The Department of Energy awards allocations of computing resources for climate studies across multiple laboratories through the INCITE program. In turn, the Climate Science Computational End Station (CCES) organizes and coordinates these computational efforts.

Using ALCF resources, CCES is advancing climate science through both aggressive model development activity and an extensive suite of climate simulations to correctly simulate the global carbon cycle and its feedback to the climate system, including its variability and modulation by ocean and land ecosystems.

Researchers are testing a new, highly scalable method for solving the fluid dynamics of the atmosphere for use in future climate simulations. This model, called HOMME, has run with a resolution as high as 1/8th of a degree of latitude on more than 80,000 cores.

Next, researchers will use HOMME to perform standard climate model benchmark simulations for comparisons with other models. They will also test the new version of the Community Earth System Model on the ALCF's Blue Gene/P.

INCITE Allocation:
40 Million Hours

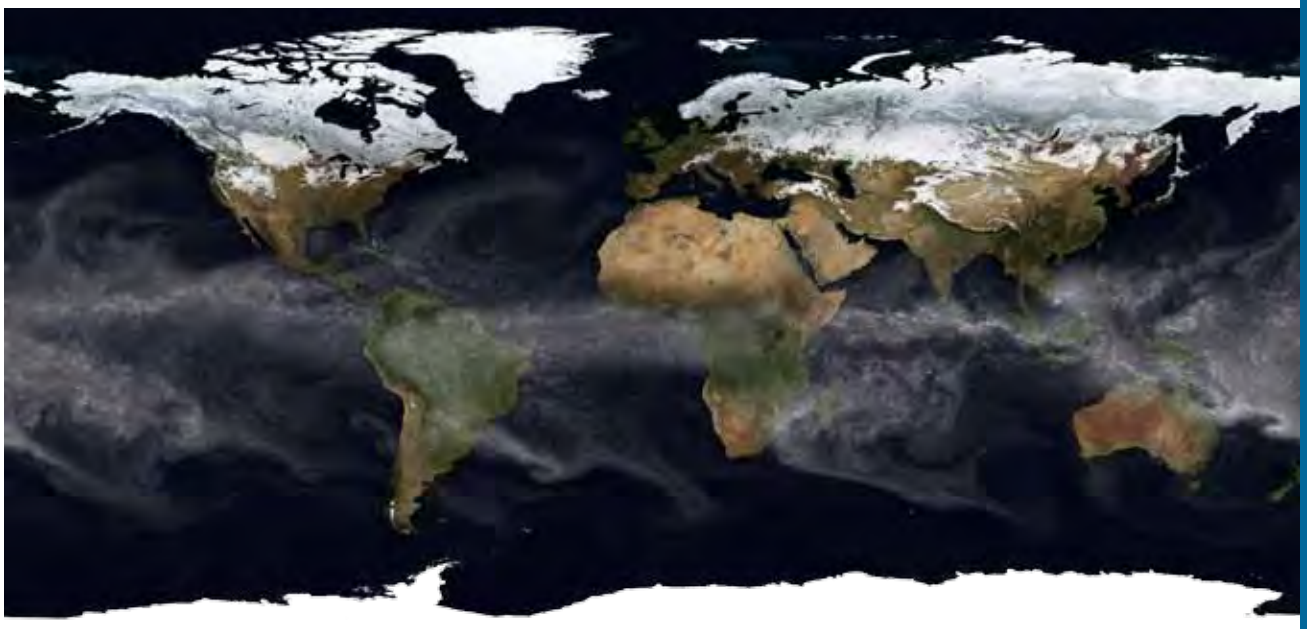

Total precipitable water, a measure of how much moisture is in the air from a single moment in time in the global simulation of the atmosphere at a resolution of half a degree of latitude. Image credit: Mark Taylor, Sandia National Laboratories.

Contact ▸ Warren Washington
National Center for Atmospheric Research | wmw@ucar.edu

earth science

Improving Models Used to Predict Global Climate Change

Standing on the shores of a raging river, it's easy to forget how such a powerful force of nature begins humbly upriver where several small creeks and streams flow together. But to accurately predict how the mighty river's course will change over time, details about the sources that feed it, no matter how small, become important.

The same is true in studying global climate change—small details matter. Scientists use computer models to simulate and study the factors that affect the balance between oceans and atmosphere. The more powerful the computer, the more detailed these climate models become, and the better their predictive capability. With access to a supercomputer, these models can begin to capture "small-scale" physics that ultimately have a big impact on climate change.

A team of researchers led by Susan Kurien from Los Alamos National Laboratory is using the supercomputer at the Argonne Leadership Computing Facility to improve climate models. Knowledge of small-scale physics will aid future studies in regional models, such as those of the Gulf of Mexico or the Arctic basin, and increase the accuracy of global climate predictions over long periods of time.

INCITE Allocation:

35 Million Hours

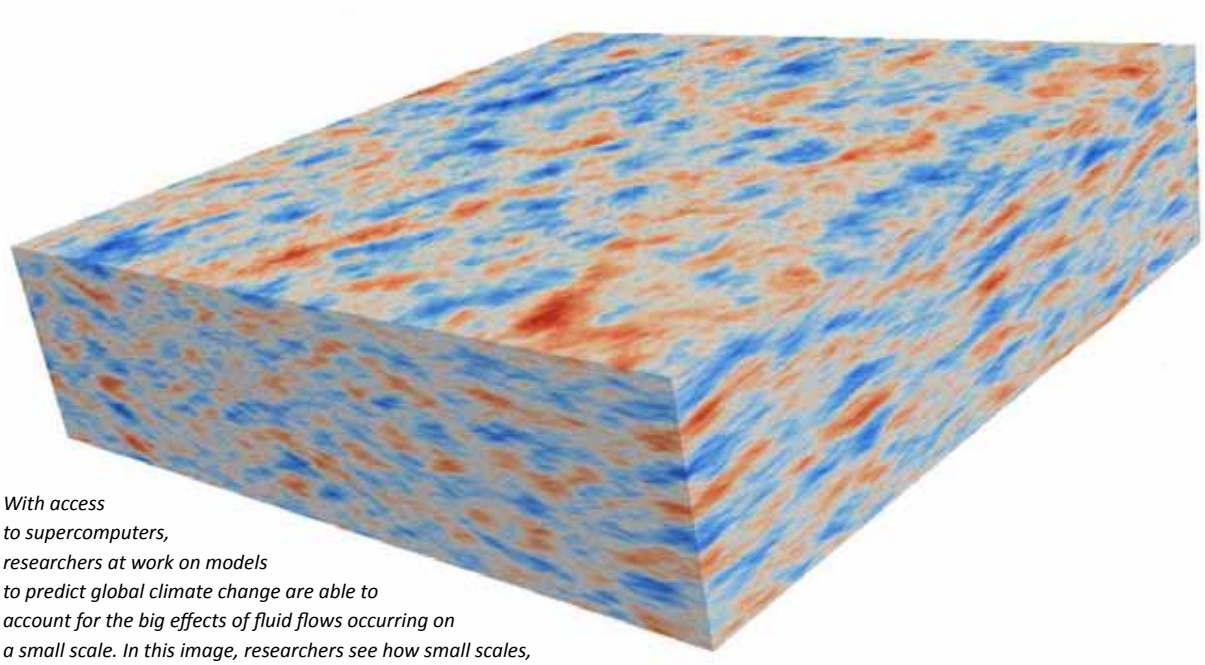

With access to supercomputers, researchers at work on models to predict global climate change are able to account for the big effects of fluid flows occurring on a small scale. In this image, researchers see how small scales, in fact, develop an interesting layered structure, indicated by fluid sloshing back and forth (red moving to the right and blue moving to the left). Image credit: The image was produced by Susan Kurien of Los Alamos National Laboratory using ALCF's Eureka data analytics cluster and ParaView.

Contact ▸ Susan Kurien
Los Alamos National Laboratory | skurien@lanl.gov

Oil Plume Behavior in a Stratified and Rotating Ocean

A key mystery for scientists is understanding the flow patterns, mixing behavior, surface presence, and distribution of buoyant plumes in a stratified and rotating environment such as the ocean. Knowledge of this behavior is directly applicable to understanding how the Deepwater Horizon oil spill distributed within the Gulf of Mexico. Studying the behavior of buoyant plumes, like those of the spill, could correlate the observations of oil on the ocean surface with the amount of oil that actually spilled.

Tamay Özgökmen from the University of Miami, in collaboration with Argonne National Laboratory researchers, is doing the first-ever simulations of the buoyant plumes in a stratified and rotating environment. This project uses Nek5000, an application to which the Argonne Leadership Computing Facility has contributed significantly, including achieving 19% of peak on 262 K cores. The research will improve the ability of evaluating oil release at underwater drilling sites and perhaps mitigate consequences at the surface.

Initial results showed that it can be difficult for the plume to cross stratification barriers caused by pressure and temperature variations. Additional results showed behavior of the plume after the spill was shut off. Further studies will quantitatively explore the dispersion behavior and rotational

The research was featured at a National Science Foundation (NSF) Hazards Research Expo hosted by the NSF on September 6, 2011 in Washington, DC. The exhibit highlighted results from the modeling effort of the oil-gas plume and its dispersion and featured stunning visualizations of the numerical simulations.

This research was critical in obtaining major funding to further examine the oil spill. The Gulf of Mexico Research Initiative is funding eight research consortia for $112.5 million over three years to support research that will investigate the fate of petroleum in the environment from the Deepwater Horizon, the impacts of the spill, and the development of new tools and technology for responding to future spills and improving mitigation and restoration. The University of Miami is the lead institution on the Consortium for Advanced Research on the Transport of Hydrocarbons in the Environment (CARTHE); Tamay Özgökmen is serving as the lead investigator.

Director's Discretionary Allocation:

2.3 Million Hours

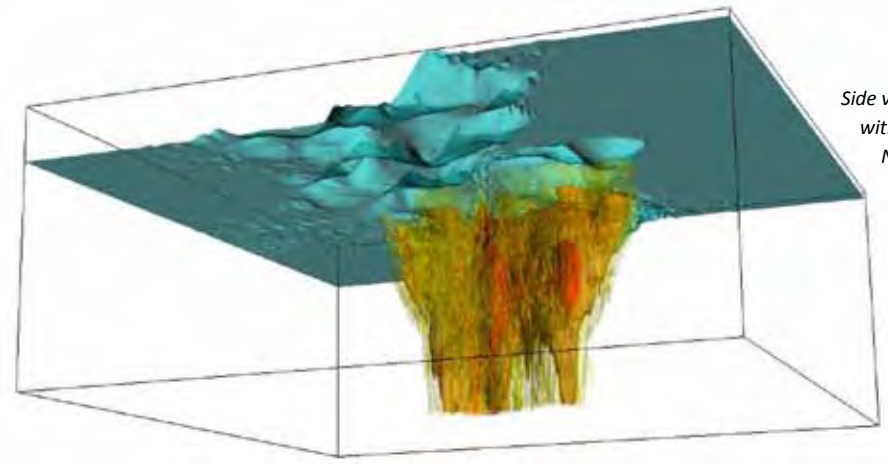

Side view of a buoyant plume interacting with ocean flows near the surface. Note the vertical coherence of the plume and inability to cross the base of the oceanic surface mixed layer. Image credit: Tamay Özgökmen, University of Miami; Paul Fischer and Aleks Obabko, Argonne National Laboratory; and Hank Childs, Lawrence Livermore National Laboratory.

Contact ▸ Tamay Özgökmen
University of Miami | tozgokmen@rsmas.miami.edu

earth science

Simulating Regional Climate at Convection-Permitting Resolution

Researchers at the National Center for Atmospheric Research (NCAR) and the Argonne Leadership Computing Facility (ALCF) are using the Nested Regional Climate Model (NRCM) on the ALCF's Blue Gene/P to demonstrate the opportunities and challenges of simulating regional climate at sufficient resolution to resolve individual thunderstorm circulations.

The work is supported by the National Science Foundation and the Willis Research Network and provides a thorough test of convection-permitting resolution on climate timescales with particular focus on high-impact weather and climate. Analysis will focus on phenomena that have high sensitivity to model resolution, including water-snowpack assessments for the mid- and western U.S., wind energy assessments, and high-impact events such as winter storms and hurricanes.

Thirteen million core-hours are being used on Intrepid, the ALCF Blue Gene/P system, to simulate both the record-breaking 2005 hurricane season and winter storms of 2005/2006, using 4 km grid spacing for a region covering the entire North Atlantic Basin and most of North America.

The experience gained at a resolution that is set to be the standard for the next generation of regional climate models represents an essential step towards our overarching goal to better quantify high-impact weather under climate variability and change. This will advance our understanding of Earth's climate for national emergencies and broaden the community of researchers capable of using leadership computing resources.

Director's Discretionary Allocation:

13 Million Hours

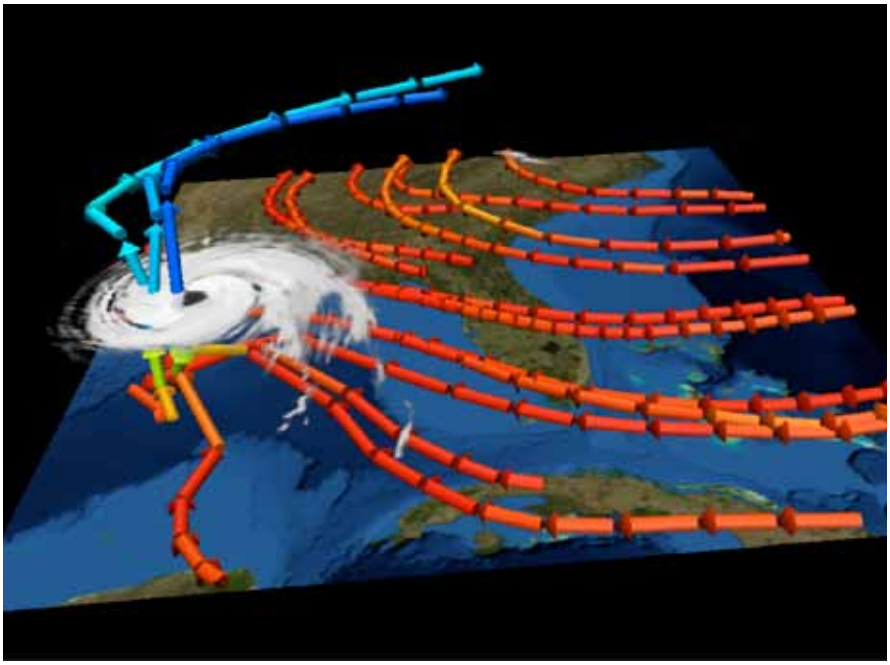

Snapshot of a simulated hurricane at 4 km grid spacing using the Nested Regional Climate Model at 4 km. White shading is simulated satellite imagery, and arrows indicate streamlines of evenly spaced seeds starting at 500 meters above the ground, colored by altitude. Image credit: James Done, NCAR Earth System Laboratory. Imagery produced by VAPOR (www.vapor.ucar.edu), a product of the Computational Information Systems Laboratory at the National Center for Atmospheric Research.

Contact ▸ James Done
National Center for Atmospheric Research | done@ucar.edu

Using Supercomputers to Improve Seismic Hazard Maps

The U.S. Geological Survey produces color-coded National Seismic Hazard maps that estimate how seismic hazards vary geographically. Because the seismic hazard information on these maps is of vital importance for businesses, corporations, governments, and other public policy organizations, regular improvements and updates are critical.

An interdisciplinary team led by Thomas Jordan at the Southern California Earthquake Center (SCEC) is using millions of supercomputer hours at the Argonne Leadership Computing Facility (ALCF) to simulate earthquake ground motions for use in constructing better seismic hazard maps. Already, this group has used detailed three-dimensional models of earth's crust to obtain simulation-based seismic hazard maps for Southern California. Jordan's research group has also simulated "worst-case" earthquake scenarios for California that allow scientists to estimate seismic hazards for rare, but possible, earthquakes in California of magnitudes of 8.0 or larger.

Improved seismic hazard maps have broad national impacts. Access to leadership-class supercomputing facilities, like the one here at Argonne, makes this work possible.

INCITE Allocation:

10 Million Hours

This U.S. Geological Survey seismic hazard map shows the maximum level of shaking expected to occur once in about 500 years, with warmer red regions expected to experience the strongest shaking. SCEC's earthquake system science research program uses high-performance computing to improve probabilistic seismic hazard maps like these. Image credit: U.S. Geological Survey Seismic Hazard Map (2008) showing Peak Ground Accelerations (PGA) with Probability of Exceedance of 2% in 50 years.

Contact ▸ Thomas Jordan
Southern California Earthquake Center | tjordan@usc.edu

energy technologies

Advanced Reactor Thermal Hydraulic Modeling

Advanced nuclear reactors are a key technology for providing power at a reasonable cost and with a low carbon footprint. This INCITE project is focused on analysis of thermal-hydraulics (heat transfer and coolant flow) in next-generation sodium- and gas-cooled nuclear reactors. The project simulations span a broad range of modeling scales. An important role for the compute-intensive simulations is to provide validation data for lower-cost reactor design simulations based on reduced-order models or subchannel codes. In the past, experiments provided the only accessible source for detailed validation of design codes. With petascale resources, it is now possible to efficiently and significantly augment the experimental parameter space through simulation.

As part of a worldwide validation effort for nuclear simulation codes, the Nuclear Energy Agency/ Organization for Economic Co-operation and Development (NEA/OECD) conducted a blind-benchmark study in 2010. Participants submitted computational results for velocity and temperature distributions in a T-junction (where streams of differing temperatures merge with limited mixing) that are important to understanding thermal-mechanical fatigue in light water reactors. Results were compared against experiments conducted for the benchmark.

Staff worked with the project team from Argonne National Laboratory and Lawrence Berkeley National Laboratory to enable the Nek5000 code to scale up to 163,840 cores on Intrepid, the IBM Blue Gene/P system at the Argonne Leadership Computing Facility (ALCF). This was leveraged in a workshop using the Juelich BG/P, where 19% of peak was realized on 262,000 cores. Runs from this project ranked first of 29 in temperature and sixth in velocity. This was a major verification and validation effort.

As a result of this success, the project is able to move forth with its hierarchical modeling strategy in which petascale simulations, backed up by detailed experimental data, are used to validate design scoping simulations performed with RANS (Reynolds-averaged Navier Stokes) and subchannel codes.

INCITE Allocation:

25 Million Hours

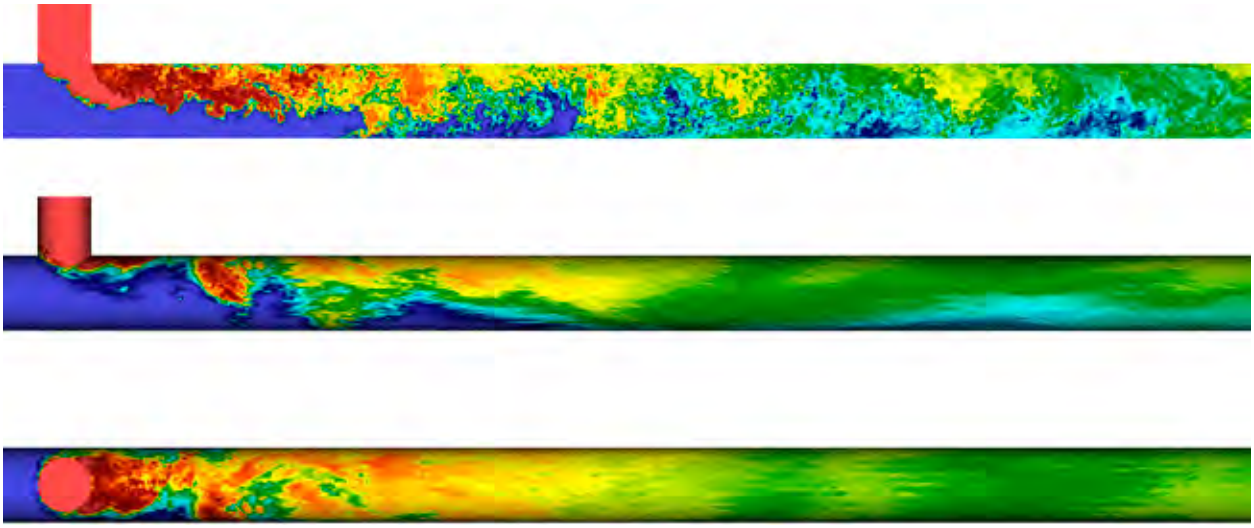

Temperature distribution illustrating thermal striping in a T-junction. Computed on Intrepid with Nek5000 and visualized on Eureka with VisIt at the ALCF. Image credit: Paul Fischer (ANL), Aleks Obabko (ANL), and Hank Childs (LBNL).

Contact ▸ Paul Fischer
Argonne National Laboratory | fischer@mcs.anl.gov

Large-Eddy Simulation for Green Energy and Propulsion Systems

An understanding of the complex turbulent mixing noise sources for wind turbine airfoils and jet exhaust nozzles is critical to delivering the next generation of "green," low-noise wind turbines and jet engines. Scientists at GE Global Research (GEGR) are leveraging the Argonne Leadership Computing Facility (ALCF) and Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program to develop/prove hi-fidelity direct-from-first-principles predictions of noise to characterize these hard-to-measure acoustic sources.

A scalable, compressible Large Eddy Simulation (LES)–based Computational Aeroacoustics (CAA) solver is being used to study free-shear layer noise from jet exhaust nozzles and boundary layer noise sources from airfoils. GE's LES strategy pushes application/validation to realistic conditions and scale, addresses fundamental physics and source characterization challenges, and extends capability to handle complex system interactions. Powered by scalability improvements at the ALCF, earlier INCITE work demonstrated how this first-principles-based LES capability can transform product development.

The research under way for jet engines has a two-pronged focus: (a) demonstrate a numerical jet-noise rig that provides faster, cheaper, and additional details of the turbulent flow-fields than possible via physical experimentation, helping to explain how various chevron designs affect noise; and (b) demonstrate the ability to capture propulsion-airframe integration effects such as the effect of pylon and jet-flap interaction.

For wind turbines, simulations are in progress to demonstrate LES readiness in guiding low-noise design, validation of the complex scaling of turbulent self-noise, and the use of effective wall-models to enable large-span blade computations. The latter are needed to drive noise-reduction concepts (e.g., serrations, advanced blade tips) beyond low-noise airfoil design.

With proof-of concept aero and acoustic LES calculations completed, the team is pursuing the use of LES+HPC to:

- Design noise-reduction features,
- Demonstrate numerical wind tunnel capability,
- Improve efficiency of numerical algorithms and parallel scalability beyond 32 K cores.

INCITE Allocation:

20 Million Hours

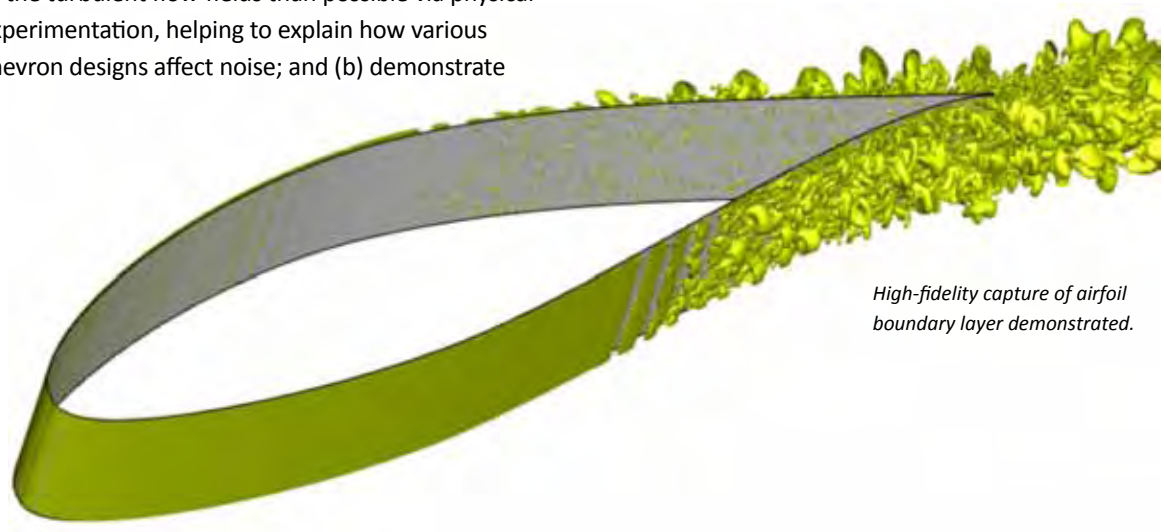

High-fidelity capture of airfoil boundary layer demonstrated.

Contact ▸ Umesh Paliath
GE Global Research | paliath@ge.com

energy technologies

Scalable, Explicit Geometry, Whole-Core Nuclear Reactor Simulations

With a primary focus on nuclear energy issues, the Department of Energy has developed a multi-physics simulation framework at Argonne for the analysis and design of nuclear reactors.

Making use of this framework, Micheal Smith and a team of researchers run PROTEUS—a reactor physics toolset focused on improving the existing methodology for nuclear reactor analysis and design, including fast reactor analysis of facilities like the MONJU reactor in Japan and the TerraPower reactor concept funded by Bill Gates.

With access to leadership-class computing systems like those at the Argonne Leadership Computing Facility, many of the approximations in legacy modeling methodologies are no longer necessary. PROTEUS provides the power distribution on the deformed geometry, allowing researchers to predict the change in power distribution derived from thermal expansion and improve the design and safe operation of nuclear reactors.

Two PROTEUS neutronics components were used in this study: MC^2 —which processes the continuous cross-section data into equivalent, multi-group data (material properties), and SN2ND—which solves the multi-group neutron transport equation on petascale architectures. Zero Power Reactor (ZPR) experiments provide clean benchmark cases against which PROTEUS modeling can be compared. With high-performance computing, PROTEUS improves the modeling accuracy in important reactions, effectively reducing errors from 5-10 percent to less than 4 percent.

SN2ND calculations routinely use 65,536 cores on Intrepid to solve problems with hundreds of billions to trillions of degrees of freedom. The group's primary efforts focused on detailed models of ZPR-6 experiments carried out at Argonne and the MONJU restart activities carried out in Japan.

A steady-state, drawer-homogenized SN2ND calculation of a ZPR-6/6A model was used for validation purposes. SN2ND was a finalist for the 2009 Gordon Bell prize, where the primary physics calculation was a detailed plate-by-plate model of ZPR-6/6A.

ALCC Allocation:

38 Million Hours

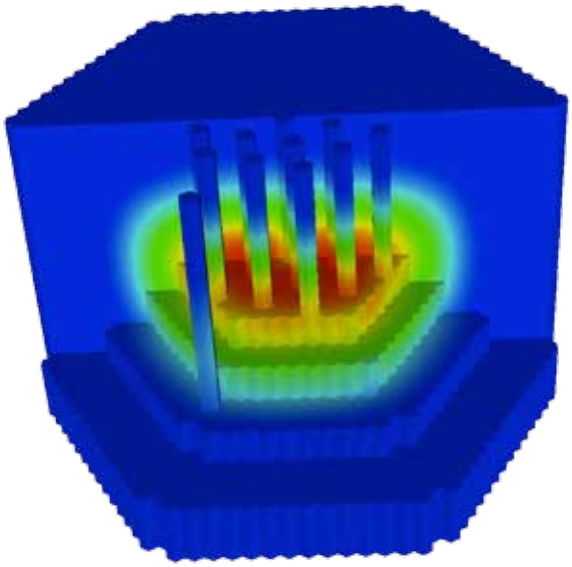

A heterogeneous model of MONJU constructed for PROTEUS, where the geometrical details of the driver, blanket, shield, and control rod assemblies are clearly visible. In these heterogeneous solutions, an outer control rod was partially inserted, leading to a radial offset in the power distribution. This was done to investigate the impacts on the solution accuracy.

Contact ▸ Micheal Smith
Argonne National Laboratory | masmith@anl.gov

Understanding the Ultimate Battery Chemistry: Rechargeable Lithium/Air

A rechargeable Lithium/Air battery can potentially store ten times the energy of a Lithium/Ion battery of the same weight, making practical widespread use of fully electric cars. But realizing this enormous potential is a very challenging scientific problem. Therefore, an interdisciplinary effort from IBM Research, Vanderbilt University, Oak Ridge National Laboratory (ORNL), and Argonne National Laboratory (ANL) is focusing on this problem.

Currently, the main challenges are to realize a high percentage of the theoretical energy density, improve electrical efficiency of recharging, increase the number of times the battery can be cycled, and improve the power density. As the key reaction product in a Lithium/Air battery, various properties of Li_2O_2 remain elusive and needed to be explored. Thus, at Argonne, the current primary focus is on the basic study of lithium oxides (Li_2O), the peroxides (Li_2O_2) system (e.g., bulk crystals, surfaces, and nanoparticles) that resides on the cathode interfaces as the Li/Air cells' main reaction products.

So far, the fundamental understanding of these systems that are governed by electronic, structural, and thermodynamic properties have been obtained by using the well-parallelized Density Functional code installed on Intrepid, the IBM Blue Gene/P system at the Argonne Leadership Computing Facility. These results will provide useful insights for the design of Li/Air cells in solving the discharge/recharge reactions at the electrode-electrolyte interface in the future.

INCITE Allocation:
25 Million Hours

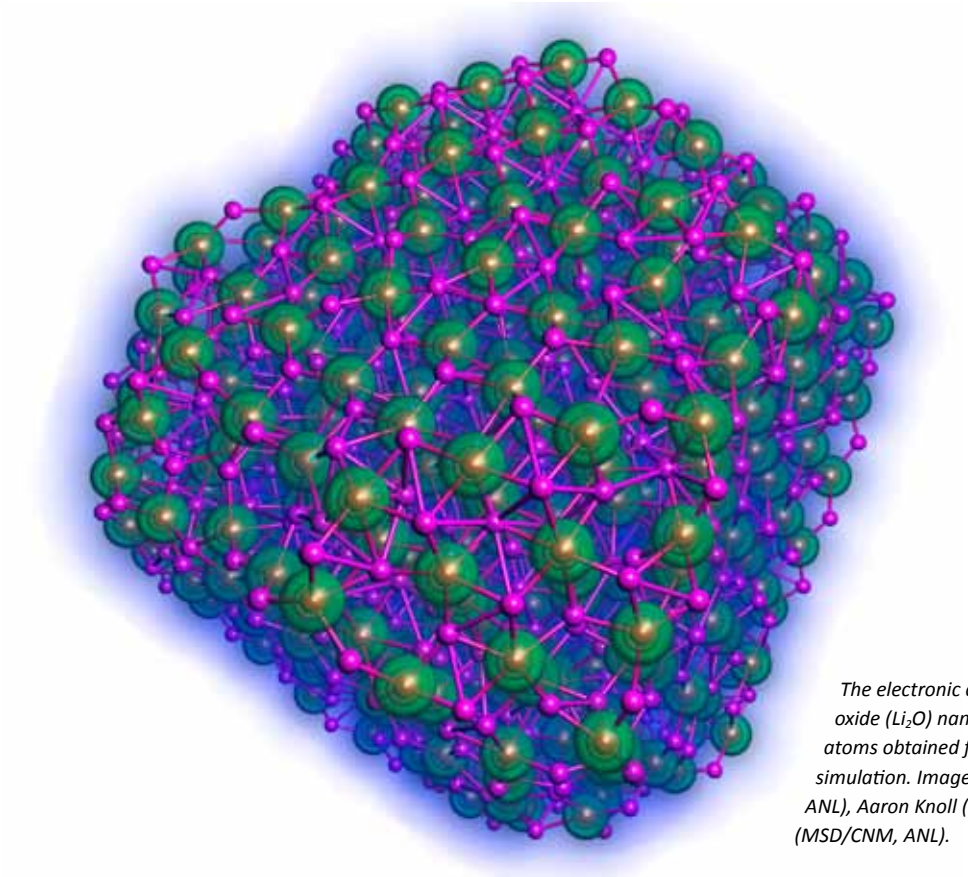

The electronic charge density of a lithium oxide (Li_2O) nanoparticle consists of 1500 atoms obtained from Density Functional Theory simulation. Image credit: Kah Chun Lau (MSD, ANL), Aaron Knoll (MCS, ANL), Larry A. Curtiss (MSD/CNM, ANL).

Contact ▸ Larry A. Curtiss
Argonne National Laboratory | curtiss@anl.gov

engineering

Boosting Fuel Economy through Cutting-edge Computational Physics

United Airlines, the world's largest airline, consumes one of every 350 gallons of global oil production. In 2010, the airline's fuel costs reached $9.6 billion. Combating the economic impact of soaring fuel costs requires major innovation in aircraft engineering. To improve jet fuel efficiency, scientists need a thorough understanding of the physics at work in turbulent flows, including drag.

Before it can become airborne, an aircraft must power through the effects of drag, burning costly fuel in the process. A team of researchers led by Robert Moser from the University of Texas is using the resources at the Argonne Leadership Computing Facility (ALCF) to shed light on the physics that impact fuel economy in jets and other systems affected by drag in turbulent flows. With commercial aircraft flying at near supersonic speeds, the calculations in these studies require leadership-class supercomputing resources, available at the ALCF.

Extreme-scale exploration of basic science, like the physics of drag, enables scientific and engineering breakthroughs with far-reaching industrial, societal and economic impact. With access to the ALCF's resources, Moser's work will facilitate design innovations in next-generation jets that translate into significant fuel and dollar savings.

INCITE Allocation:
40 Million Hours

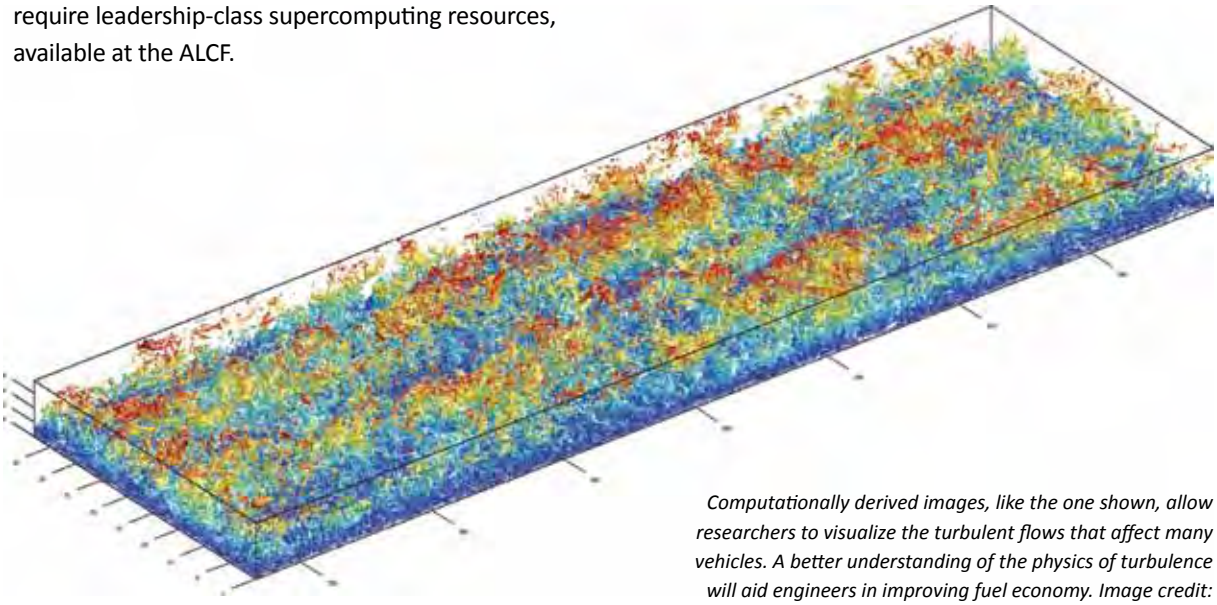

Computationally derived images, like the one shown, allow researchers to visualize the turbulent flows that affect many vehicles. A better understanding of the physics of turbulence will aid engineers in improving fuel economy. Image credit: Juan Sillero, Universidad Politécnica de Madrid.

Contact ▸ Robert Moser
University of Texas at Austin | rmoser@ices.utexas.edu

Detached-Eddy Simulations and Noise Predictions for Tandem Cylinders

Researchers led by Philippe Spalart at Boeing are conducting a thorough numerical investigation of the turbulence and noise generated by Tandem Cylinders, a prime test case for detailed comparisons among computational fluid dynamics (CFD) approaches and with NASA experiments. It involves massive separation, the impingement of turbulence on a solid body, and the resulting noise. These features are typical of various applications (aircraft landing gear, wind turbines, bridges, heat exchangers, and buildings).

The simulations reflect the state of the art in turbulence-resolving CFD approaches to massively separated flows. The IBM Blue Gene/P at the Argonne Leadership Computing Facility allows sufficiently fine grids (up to 60 million cells) and wide domains (up to 16-cylinder diameters), with either periodic or free-slip boundary conditions in the lateral direction. The objective is to gain insight into the fundamental problem of low-Mach-number aero-acoustic computations. Very low Mach numbers are notorious for affecting CFD convergence; furthermore, the acoustic energy is minute.

The accuracy of the simulations is confirmed by the Power Spectral Density of velocity at a point in the wake: the inertial range of frequencies with exponent -5/3 is a decade long, revealing the turbulent energy cascade. The mean and unsteady wall pressures also compare well with experiments.

A central issue of aero-acoustics for low-Mach-number flows is the accuracy of noise predictions made using only the solid-surface terms in the Ffowcs-Williams-Hawkings (FWH) integral. It is widely assumed that this convenient simplification suggested by Curle is adequate for airframe noise. The researchers' work on a landing gear indicated, however, that it is not adequate even at a Mach number as low as 0.115; this is confirmed in the present study at Mach 0.1285 for the higher frequencies. *This unexpected finding needs to be explored, including the level of numerical errors, first inside the turbulence and then in the radiated noise, via the FWH equation.*

INCITE Allocation:

45 Million Hours

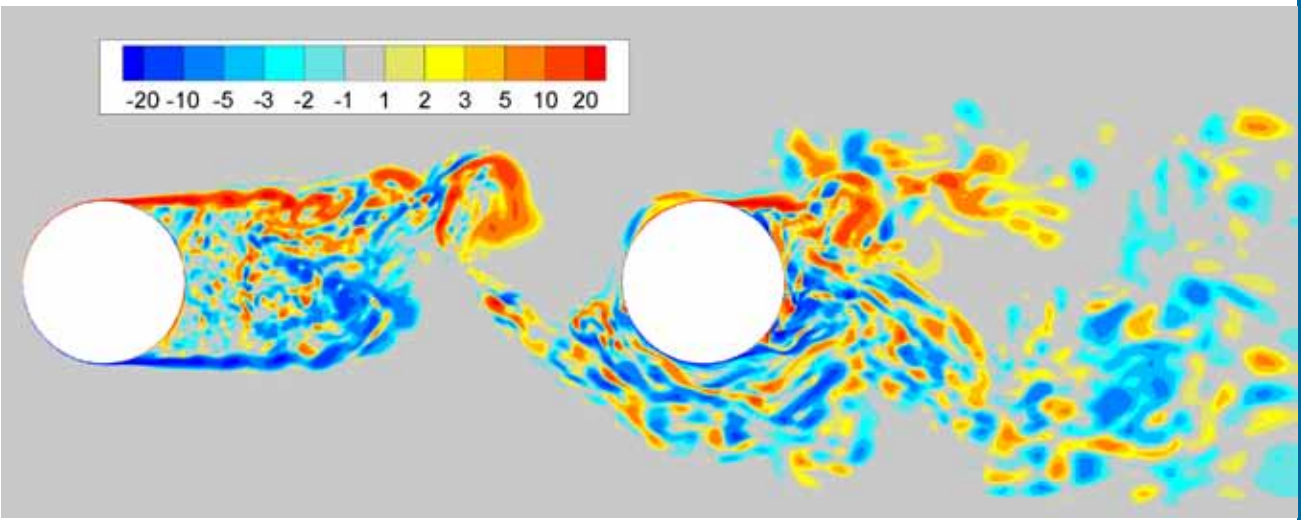

Flow visualization (snapshot of vorticity magnitude from DDES of the Tandem Cylinders). Image credit: Dr. Mikhail Strelets, Dr. Mikhail Shur, and Dr. Andrey Travin, New Technologies and Services, Ltd. and Saint-Petersburg State Polytechnic University.

Contact ▸ Philippe Spalart
Boeing | philippe.r.spalart@boeing.com

engineering

Prediction of Supersonic Jet Noise Using Large-Eddy Simulation

A noise generated by engine exhaust jets is one of the greatest barriers to the deployment of high-speed commercial aircraft. Very high jet velocities, typical of supersonic aircraft, heighten noise pollution levels in airport communities and accelerate hearing loss for crew on aircraft carrier decks.

Modifying the nozzle geometry by adding chevrons (serrated geometric edges) to the nozzle lip is experimentally known to reduce jet noise, but the physical reasons for this are not well understood. Small-scale turbulence and shocks generated by chevrons strongly influence the development of the large turbulent eddies far downstream that control the production of noise. To capture these effects, compressible large-eddy simulation using unstructured grids, which minimizes numerical dissipation and dispersion, is needed. For this purpose, the CharLES computational infrastructure, developed at the Center for Turbulence Research, is being used for these massively parallel simulations at the Argonne Leadership Computing Facility (ALCF).

Using the resources of the ALCF, researchers will work to fully resolve the effects of chevrons and heating on the noise produced in supersonic turbulent flow from a rectangular nozzle. These simulations are expected to reveal new insight into how enhanced shear-layer mixing due to chevrons sustains itself and to address the question of the optimal chevron penetration angle from a jet-noise standpoint. This research may further efforts in noise-mitigation strategies that are critical to the design and viability of next-generation supersonic aircraft.

ALCC Allocation:
60 Million Hours

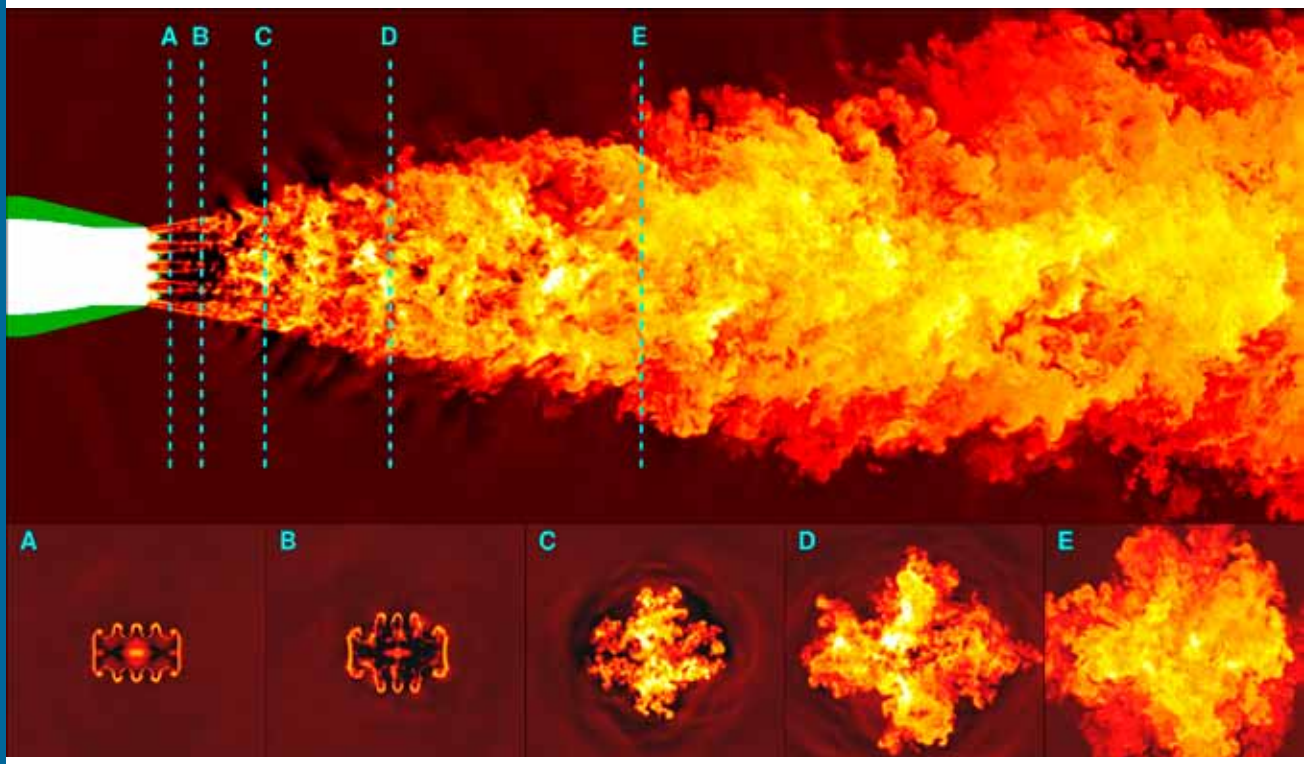

Supersonic jet exhaust from a 2:1 rectangular nozzle trimmed with 8 mechanical chevrons along the nozzle lip. Temperature contours shown.

Contact ▸ Parviz Moin
Center for Turbulence Research, Stanford University | moin@stanford.edu

Simulations of Turbulent Flows with Strong Shocks and Density Variations

This research project aims to develop a new capability based on high-order and high-resolution schemes to simulate shock-turbulence interactions and multi-material mixing in planar and spherical geometries, and to study Rayleigh-Taylor and Richtmyer-Meshkov turbulent mixing. These fundamental problems have direct application in high-speed engineering flows, such as inertial confinement fusion capsule implosions and scramjet combustion, and also in the natural occurrence of supernovae explosions. This is achieved by leveraging the capabilities of Intrepid, the Blue Gene/P supercomputer at the Argonne Leadership Computing Facility, which enables the study of fundamental physics governing these phenomena.

High-resolution simulations of this phenomena capture the scales at which viscous dissipation and molecular mixing occur while representing the nonlinear dynamics of the energy-containing scales. The resulting databases will enable a fundamental study of the mechanisms at play in turbulent multi-material mixing in high-speed accelerated flows and help develop improved models for engineering calculations.

Results from the canonical, nominally planar shock-turbulence interaction simulations suggest that the amplification of turbulence kinetic energy is predominantly a linear, large-scale process, well predicted by existing linear theories, *whereas the shock-induced change in the turbulence structure is a non-linear process that both theory and previous (under-resolved) simulations fail to capture.*

Simulations of the spherical converging/reflecting shock-turbulence interaction show that vorticity is amplified following the expanding/reflected shock, generating smaller-scale structures. The shock itself is distorted by the turbulence, deviating from Guderley's theoretical scaling in its increased asphericity and decreased compression, especially for weaker shocks.

Simulations of the planar shock-curtain interaction show that the growth of the primary instability is highly sensitive to the initial concentration profile of species present in the flow, posing a challenge for comparison with experiments. A double-peak time evolution is observed for RMI in spherical geometry, with peaks that correspond to the linear growth phase and the post-reshock turbulent mixing layer growth.

INCITE Allocation:
12 Million Hours

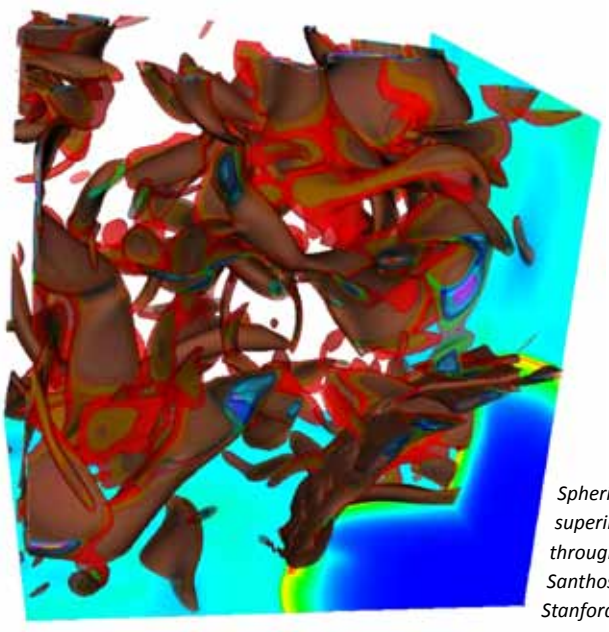

Spherical Shock-Turbulence Interaction. Iso-surfaces of vorticity superimposed on contours of density for a strong blast wave propagating through a region of isotropic turbulence. Image credit: Ankit Bhagatwala, Santhosh Kumar, Johan Larsson, Ivan Bermejo-Moreno, and Sanjiva Lele, Stanford University.

Contact ▸ Sanjiva Lele and Ivan Bermejo-Moreno
Stanford University | lele@stanford.edu | ibermejo@stanford.edu

engineering

Uncertainty Quantification for Turbulent Mixing

Rayleigh-Taylor (RT) unstable mixing flow challenges conventional ideas of computational science and sheds light on differing views for the computation of turbulent mixing flows. The problem is encountered in a variety of contexts, such as combustion, inertial confinement fusion, supernovae explosions, and geophysics.

Led by James Glimm of Stony Brook University, researchers are verifying and validating parameterized subgrid models for turbulent mass, momentum, and thermal diffusion that capture unresolved phenomena as it impacts coarse grid scales in a Large Eddy Simulation (LES) having steep numerical gradients. Refined mesh simulations based on the front tracking algorithm, with INCITE allocations on Intrepid, the Blue Gene/P high-performance computer at the Argonne Leadership Computing Facility, resolve key length scales and minimize the role of models.

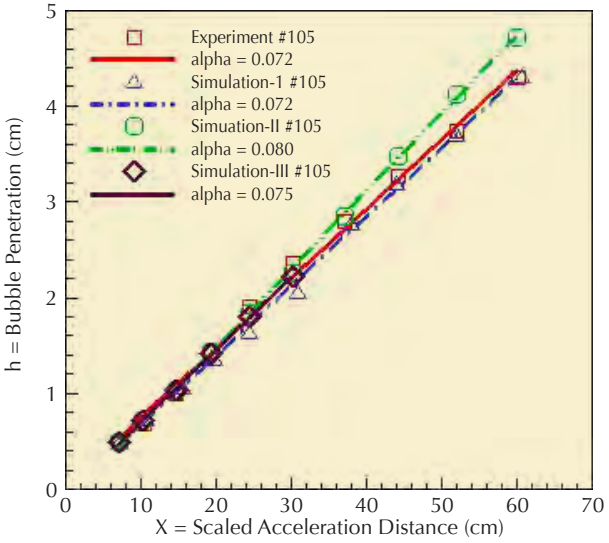

Plot of the bubble penetration distance h_b vs. a scaled acceleration distance, Agt^2. The slope of this curve is the mixing growth rate α_b. Researchers plotted the experimental data points and three simulation results, which have (I) 0 x, (II) 2 x the best reconstruction of the initial long wavelength perturbations, as extrapolated by backward propagation in time from the early time experimental plates. (III) They added inferred initial conditions for long wavelength perturbations to their previously successful simulation-experiment validation study. The main point of their current simulation is to resolve the Weber scale; $\Delta x = 111\mu m < l_{We} = 130\mu m$, where l_{We} is the critical bubble size (predicted by Weber number theory). Image credit: Tulin Kaman, Stony Brook University.

The front tracking method is an adaptive computational method that provides sharp resolution of a wave front by tracking the interfaces between distinct materials. It represents interfaces explicitly as lower dimensional meshes moving through a rectangular grid. The method has been implemented in the code FronTier, which solves the equations with the following steps:

1. Interface propagation,
2. Interpolation reconstruction,
3. Interior states update.

The problem is to predict the growth rate α, defined in terms of the penetration distance h of the light fluid into the heavy fluid via the equation $h = \alpha Agt^2$, where A is the Atwood ratio, a dimensionless measure of the density contrast, and g is the acceleration, applied normal to an interface separating fluids of density ρ^1 and ρ^2.

The error estimates in the $t = t_0$ reconstructed initial data are used to establish an uncertainty quantification interval for numerical simulations of the RT growth rate α.

- Maximum ±5% effect on the growth rate
- Agreement with experimental data of Smeeton & Youngs, 1987
- Resolution of critical length scales, to diminish and verify the role of models

Director's Discretionary Allocation:

1.5 Million Hours

Contact ▸ **James Glimm and Tulin Kaman**
Stony Brook University | glimm@ams.sunysb.edu | tkaman@ams.sunysb.edu

Global Simulation of Plasma Microturbulence at the Petascale and Beyond

As scientists look for alternatives to fossil fuels to meet the world's energy needs, there is increasing interest in nuclear fusion, the power source of the sun. Of utmost importance in the design and operation of future fusion power sources is an understanding of turbulent transport losses. Acquiring this knowledge requires computational efforts at the extreme scale.

Researchers will study the influence of plasma size on confinement properties in advanced tokamak systems, like ITER. This requires a systematic analysis of the underlying nonlinear turbulence characteristics in magnetically confined tokamak plasmas that span the range from current scale experiments, which exhibit an unfavorable scaling of confinement as the plasma radius increases, to ITER-scale plasmas, expected to be insensitive to size variations.

Present-day tokamaks are not even one-third the radial dimension of ITER, making high-fidelity predictive simulations even more critical, since improvements in ITER-sized devices can only be validated after they are constructed and operational.

In dealing with this challenge, researchers will deploy the GTC-P and GTS codes, which are highly scalable particle-in-cell gyrokinetic codes used for simulating microturbulence-driven transport in tokamaks.

To date, efforts to scale the codes to the petascale power of the next-generation Blue Gene/Q have yielded significant results:

- GTC parallelism: MPI plus OpenMP
  - Best Blue Gene/P configuration: 1 MPI rank per node, with 4 OpenMP threads
  - Best Blue Gene/Q configuration (on early access hardware): 1 MPI rank per node with 64 OpenMP threads
    - BG/Q has 16 cores per node, 4 hardware threads per core (Ruud Haring, IBM, Hot Chips Meeting, July 2011)
- Long-duration simulations of ITER plasmas will demand O(10^8) grid points and O(10^{10}) particles

Early Science Program Allocation:

50 Million Hours

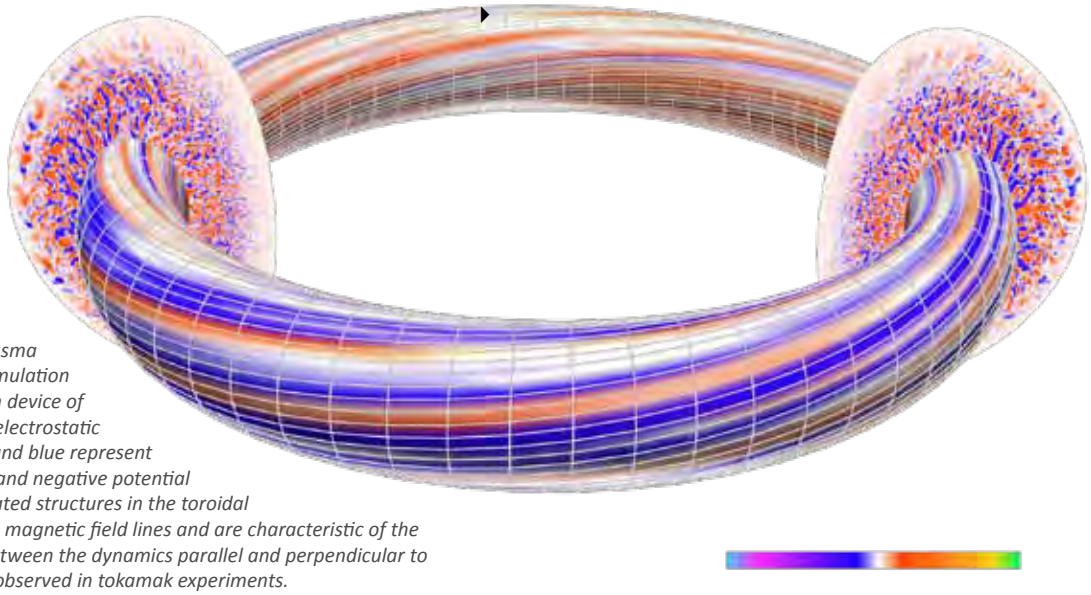

Fully kinetic 3-D plasma microturbulence simulation in a tokamak fusion device of the self-consistent electrostatic potential. The red and blue represent regions of positive and negative potential respectively. Elongated structures in the toroidal direction follow the magnetic field lines and are characteristic of the large anisotropy between the dynamics parallel and perpendicular to the magnetic field observed in tokamak experiments.

Contact ▸ William Tang
Princeton Plasma Physics Laboratory | tang@pppl.gov

materials science

ALCF's Blue Gene/P Enables New Insights into Concrete's Flow Properties

Flow simulations of thousands of irregularly shaped particles on the Argonne Leadership Computing Facility's (ALCF) Blue Gene/P supercomputer are enabling new insights into how to measure and control flow properties of large-particle dense suspensions like concrete that can't be accurately measured in industrial settings. This study of the flow of matter—known as rheology—will provide a better understanding of concrete's flow properties to help ensure its optimum performance and eliminate cost overruns.

Measuring the true rheological properties of fresh concrete is quite a challenge. Laboratory devices that measure the way a suspension flows in response to applied forces—called rheometers—can just approximate rheological properties of suspensions because these devices gauge only torque and angular velocity. Without detailed knowledge of the flow in the rheometer, scientists can't obtain fundamental rheological parameters such as viscosity and yield stress (resistance to initiating flow). Computing the flow allows researchers to correctly interpret empirical measurements in terms of fundamental units.

Through their simulations on the Blue Gene/P at the ALCF, researchers led by William George from the National Institute of Standards and Technology (NIST) have gained fundamental new insights into the yield stress of dense suspensions. Studies by the group indicate that particle contacts are an important factor in controlling the onset of flow in dense suspensions. Further, such interactions can be strongly influenced by the shape of the aggregates and lead to jamming effects that restrict the easy placement of concrete in forms. The researchers also discovered that for suspensions with a non-Newtonian fluid matrix, the local shear rates between aggregates strongly determine their rheological properties. These results have been validated against physical experiments with excellent agreement.

Knowledge gained through this research has technological application in the building, coatings, water-treatment, food-processing, and pharmaceutical industries and furthers DOE's mission to develop revolutionary new materials and processes.

INCITE Allocation

25 Million Hours

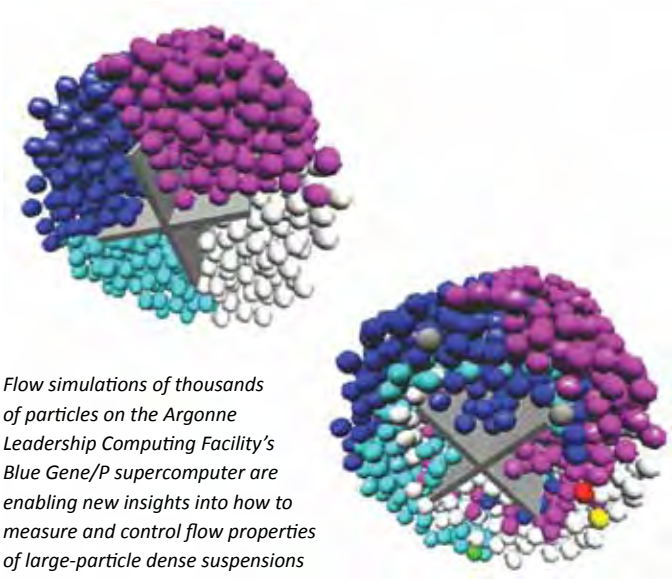

Flow simulations of thousands of particles on the Argonne Leadership Computing Facility's Blue Gene/P supercomputer are enabling new insights into how to measure and control flow properties of large-particle dense suspensions like concrete that cannot currently be accurately measured in industrial settings. A better understanding of these properties will help ensure the optimum performance of concrete and eliminate cost overruns. Image credit: Images and the software used to produce them were developed by Steven Satterfield, John Hagedorn, and John Kelso of NIST and Marc Olano of NIST and the University of Maryland–Baltimore County.

Contact ▸ William George
National Institute of Standards and Technology | william.george@nist.gov

Better Catalytic System Designs through Nanoscale Research

In life, sometimes to get the ball rolling, a little nudge is needed. In a chemical reaction, that nudge often comes in the form of a catalyst. A catalyst is an agent that speeds a chemical reaction along, or causes a reaction that otherwise would not have occurred. Platinum, a common catalyst, is used in catalytic converters to remove toxins from exhaust. Improved emissions control requires an understanding of how catalysts behave at their most fundamental atomic level—the nanoscale.

Jeff Greeley of Argonne National Laboratory leads a team, including researchers from the Technical University of Denmark and Stanford University, that uses the supercomputing resources at the Argonne Leadership Computing Facility (ALCF) to study catalytic nanoparticles. Calculating catalysis on particles with as few as one thousand atoms takes several days on the world's fastest supercomputers. The process is so time intensive and the calculations are so complex, the research would be impossible without a leadership-class system like the ALCF's Blue Gene/P.

At the ALCF, large-scale, basic science exploration yields significant industrial, societal, and economic impact. With access to the world-class computing resources needed to explore the behavior of catalysts at the nanoscale, Greeley and his team are paving the way for improved catalytic system designs with wide-ranging industrial applications.

INCITE Allocation:
15 Million Hours

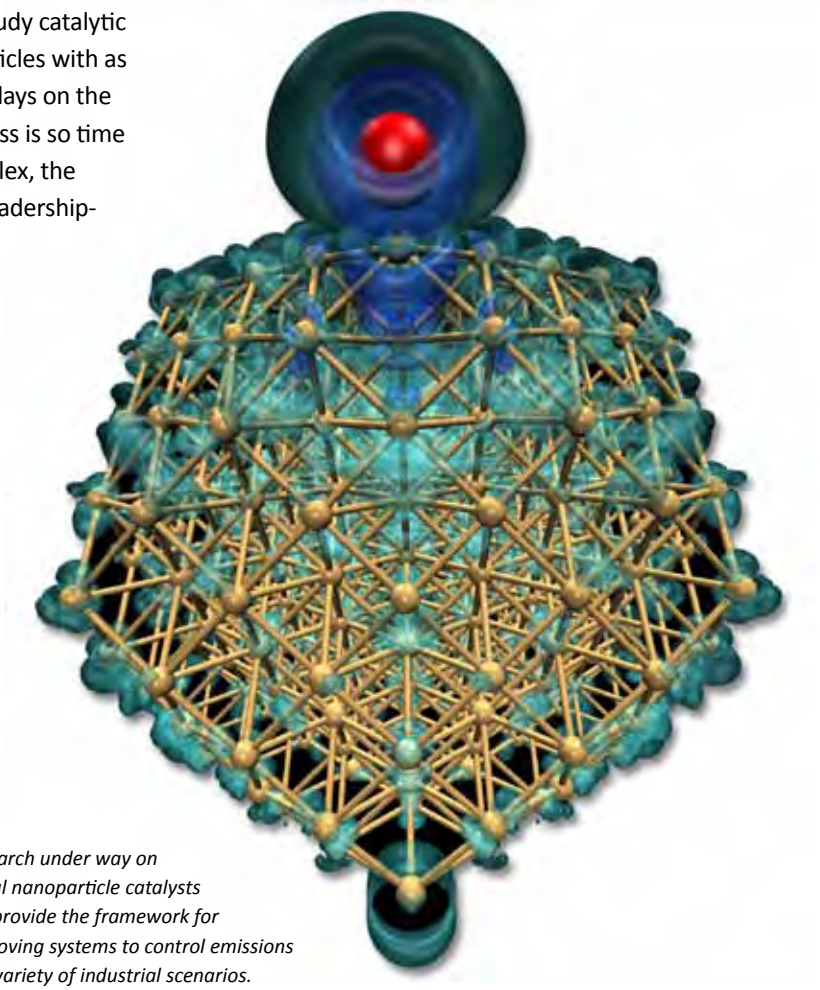

Research under way on metal nanoparticle catalysts will provide the framework for improving systems to control emissions in a variety of industrial scenarios.
Image credit: Jeff Greeley, Argonne National Laboratory.

Contact ▸ Jeff Greeley
Argonne National Laboratory | jgreeley@anl.gov

materials science

First Principles Simulations of the Infrared Spectrum of Liquid Water

Understanding pure water is an essential prerequisite for understanding the behavior of the liquid interacting with naturally occurring compounds such as carbon dioxide or methane. Several accurate flavors of density functional theory (PBE0 and van der Waals functionals), that include computationally demanding exact exchange terms, have been used in Professor Giulia Galli's laboratory at the University of California-Davis (UC-Davis) to investigate water and ions in water with atomistic resolution. Specifically, *ab initio* molecular dynamics calculations of liquid water were carried out using the massively parallelized density functional theory code Qbox, developed in Prof. Gygi's group at UC-Davis, within an INCITE allocation on the Blue Gene/P high-performance computer at the Argonne Leadership Computing Facility. Structural and vibrational properties of liquid water were determined for several temperatures. Rather conventional generalized gradient (PBE)-based calculations, currently considered to be state-of-the-art, were also performed.

The PBE0-based predictions of electronic, structural, and vibrational properties of liquid water substantially improve the prediction when compared to PBE, using experimental results as a reference. In particular, there is a dramatic improvement in the description of infrared spectra. These important findings suggest that through the use of high-performance computing, we can improve our predictive power of aqueous environments. Prof. Galli and her collaborators' research was published in three papers, two in the *American Chemical Society's Journal of Chemical Theory and Computation* and one in the *Journal of Physical Chemistry B*. Using PBE0, the authors are currently extending their investigations towards solvation properties of important ions, such as chloride, or mixtures of water with methane or carbon dioxide.

INCITE Allocation:

15 Million Hours

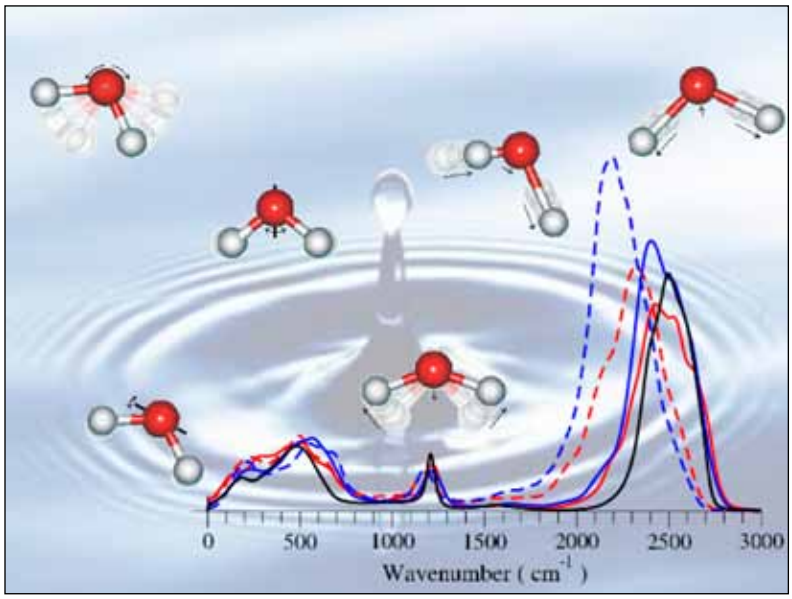

Infrared spectra of liquid water computed using hybrid (solid lines) and semi-local (dotted line) functionals, computed by ab-initio *molecular dynamics with the Qbox code. Image credit: Dr. Cui Zhang, UC-Davis.*

Contact ▸ Giulia Galli
University of California-Davis | gagalli@ucdavis.edu

Materials Design and Discovery: Catalysis and Energy Storage

The energy future hinges on the design and discovery of new materials—like materials to replace the oils currently used to make plastics, and materials to power electric vehicles.

Scientists at Argonne's Center for Nanoscale Materials (CNM) and the Materials Science Division (MSD) are pairing the power of the IBM Blue Gene/Q with newly available electronic structure codes to conduct massively parallel quantum mechanical calculations for use in the design of breakthrough materials that may have energy-related applications.

A team of researchers, led by Larry Curtiss with Argonne's MSD and CNM, is focusing research efforts on catalytic materials and on materials used for electric energy storage. Catalytic materials are used for bond-specific activation for efficient chemical transformations. This research could yield new strategies for more energy-efficient, environmentally friendly chemical synthesis to help reduce greenhouse gases, or new methods for replacing petrochemicals with inexpensive, abundant small alkanes.

Creating new materials for electrical energy storage, (specifically, for the interface between electrolyte and electrode) could lead to safer, longer-range batteries for electric vehicles.

Using the extreme compute power of the Blue Gene/Q at the Argonne Leadership Computing Facility, researchers will employ high-accuracy quantum mechanical calculations using density functional theory (DFT) and quantum monte carlo (QMC). In preparation for the new architecture of the Blue Gene/Q, key computational kernels will be re-written to employ OpenMP nested parallelism.

Early Science Program Allocation:

50 Million Hours

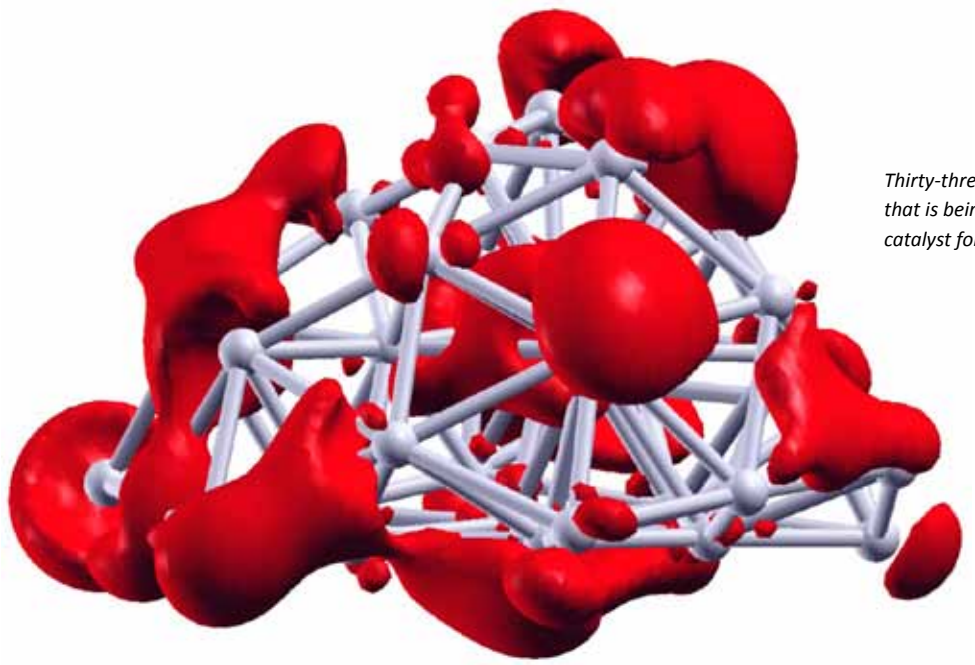

Thirty-three atom silver cluster that is being studied as a new catalyst for propylene epoxidation.

Contact ▸ Larry A. Curtiss
Argonne National Laboratory | curtiss@anl.gov

materials science

Reactive MD Simulation of Shock-Induced Cavitation Damage

Maintaining the soundness of nuclear reactors is a major concern for scientists, engineers, and the general public. Among many factors, "cavitation erosion" of cooling system components is a significant mechanism for long-term degradation in nuclear power plants. Cavitation occurs when a liquid experiences rapid change in pressure that creates low-pressure cavities within the liquid. These cavities, or cavitation bubbles, cause stress when they collapse and hit a solid surface, and therefore cause deterioration of the surfaces of materials.

However, cavitation bubbles also provide benefits. Nanobubbles are used to prevent Stress Corrosion Cracking (SCC)—the biggest reason the lifetime of nuclear reactors is shortened. When nanobubbles form, they create low-pressure regions, but when they collapse near a solid surface, the result is the creation of high-pressure areas that relieve the tensile stresses that cause SCC in the material.

To get a molecular-level understanding of nanobubble collapse near a solid surface, Priya Vashishta and his colleagues, Rajiv Kalia and Aiichiro Nakano, at the University of Southern California (USC) are using Intrepid, the IBM Blue Gene/P system at the Argonne Leadership Computing Facility (ALCF), to simulate and unravel the complex mechanochemistry problem. The 1-billion-atom simulation is feasible because it runs efficiently on 163,840 cores, the full Intrepid system. The goal of this nanobubble collapse simulation is to understand molecular processes to improve both the safety and longevity of nuclear reactors. The efficiency with which these simulations run on Intrepid is the result of successful work conducted by the USC group using an Argonne Director's Discretionary allocation in 2010.

INCITE Allocation:

45 Million Hours

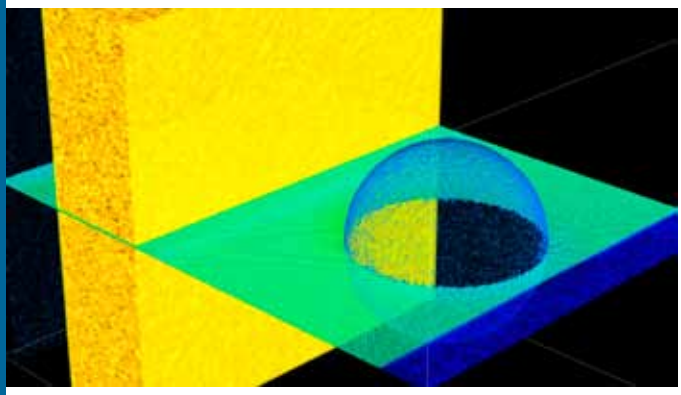

Billion-atom reactive molecular dynamics simulation of nanobubble collapse in water near a ceramic surface under shock compression. A 2km/sec shock wave compresses the nanobubble and creates high compressive stress and novel chemical reactions (production of hydronium ions) not found under normal conditions. The simulations reveal that high pressure in the shock wave deforms the ceramic surface and also accelerates water molecules from the bubble periphery towards the center of the bubble. These high-velocity water molecules bunch up to form a nanojet. The nanojet impact creates damage on the ceramic surface. The simulation results reveal atomistic mechanisms of mechanically induced chemistry, which is the key to understanding the safety-threatening damage in nuclear reactors. The simulations were carried out using 163,840 cores on Intrepid, the BlueGene/P system, at the Argonne Leadership Computing Facility. Image credit: Billion-atom reactive simulations by Adarsh Shekhar, Ken-ichi Nomura, Rajiv K. Kalia, Aiichiro Nakano, and Priya Vashishta, University of Southern California. Visualization by Joseph Insley at the ALCF.

Contact ▸ Priya Vashishta
University of Southern California | priyav@usc.edu

Ab Initio Reaction Calculations for Carbon-12

This project aims to calculate several fundamental properties of the ^{12}C nucleus: the imaginary-time response, the one-body density matrix, and transition matrix elements between isospin-0 and -1 states. Using Green's Function Monte Carlo (GFMC), these calculations will allow researchers to reliably compute neutrino-^{12}C scattering, quasi-elastic electron scattering, and the results of older reactions, such as ($\pi i,\pi i'$) on ^{12}C.

Such properties are critical to a real understanding of the physics of nucleonic matter. Electron scattering experiments in the quasi-elastic regime, where the dominant process is knocking a single nucleon out of the nucleus, are under way for a range of nuclei. The separation into longitudinal and transverse response allows study of the propagation of charges and currents, respectively, in the nucleus. The nontrivial changes from studies of the nucleon and deuteron to larger nuclei require processes beyond one-nucleon knockout. Researchers will compute the transition density matrices on a two-dimensional grid of the magnitudes of the initial and final positions and will make a partial wave expansion of the angle between the two vectors. These matrices will be used as inputs for a variety of reaction calculations.

To address Blue Gene/Q-related issues for GFMC, researchers changed the format for some matrices from dense to compressed, realizing a universal improvement on Blue Gene architecture. Secondly, researchers identified a set of important kernel routines and separated the threading part from the serial vectorizable part, allowing the use of SIMD instructions within threads. They have also tuned several Q-related parameters for data sizes. Currently running on a midplane of Blue Gene/Q, the code shows one of the highest per-node performance improvements from Blue Gene/P to Blue Gene/Q.

Early Science Program Allocation:

110 Million Hours

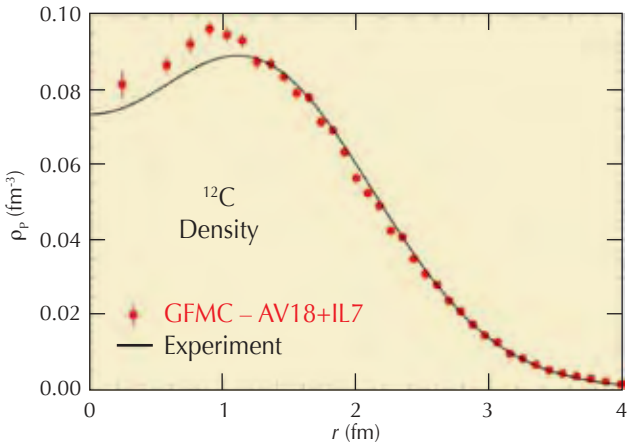

The calculated density (red dots) is in excellent agreement with experimental data (solid curve).

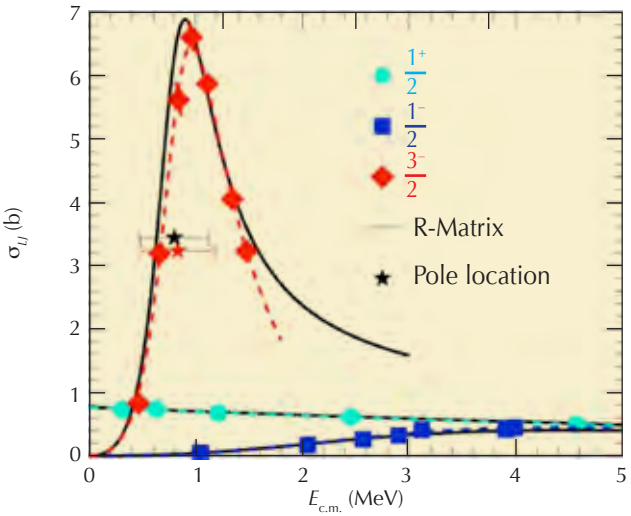

Contact ▸ Steven C. Pieper
Argonne National Laboratory | spieper@anl.gov

physics

Advancing the Understanding of Nuclear Structure

Researchers from Iowa State University, Oak Ridge and Argonne national laboratories are using complementary techniques, including Green's Function Monte Carlo, the No Core Shell Model, and Coupled-Cluster methods to perform *ab initio* calculations of both structural and reaction properties of light- and medium-mass nuclei. The calculations use realistic models of nuclear interactions, including both two- and three-nucleon forces. Their work could advance understanding of the triple-alpha burning reaction, which is essential to life on earth.

They also are exploring the role of the three-nucleon force in substantially heavier nuclei. Furthermore, researchers are using Density Functional Theory (DFT) to calculate properties of nuclei across the entire range of nuclear masses. These DFT studies will help predict nuclear properties relevant to nuclear reactions such as neutron-nucleus reaction cross-sections and fission. This new understanding of nuclei has far-reaching implications, impacting the fields of energy and astrophysics. The researchers are conducting their calculations on the IBM Blue Gene/P (BG/P) at the Argonne Leadership Computing Facility and the Cray XT at Oak Ridge National Laboratory.

The BG/P research team has completed ground-state C calculations—a key milestone. The ground state represents the best converged *ab initio* calculations of C ever. The researchers have continued developing and testing various formulations of starting wave functions for the first excited 0 state of ^{12}C (the Hoyle or triple-alpha burning state). With these they are achieving converged energies for the Hoyle state that are in reasonable agreement with the experiment.

The BG/P team has used the ^{12}C GFMC calculations to develop the Automatic Dynamic Load Balancing (ADLB) library, a general purpose work-sharing library. By using OpenMP on each node, the team is now getting essentially perfect scaling up to 65,536 cores (32,768 are used for their production runs). The team will be testing a new multi-threaded version of Automatic Dynamic Load Balancing (ADLB) on Argonne's BG/Q when that machine is available.

INCITE Allocation:

15 Million Hours

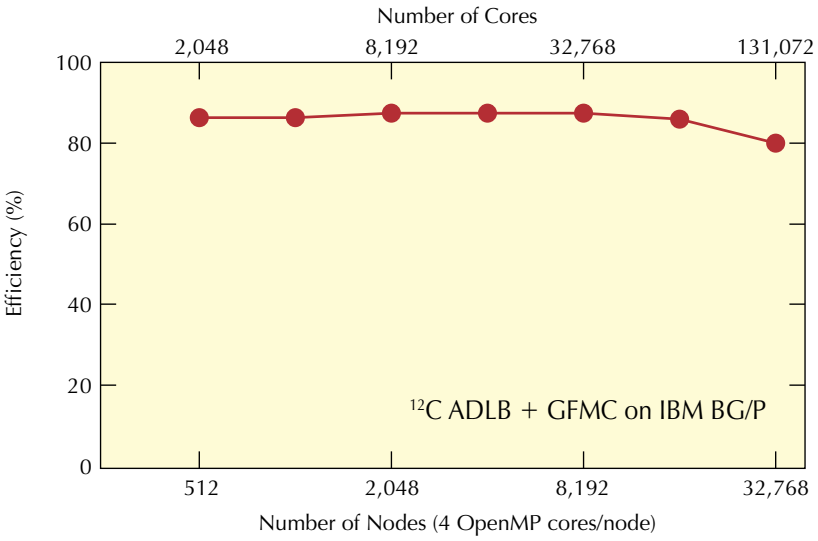

Excellent scaling is achieved by the production Automatic Dynamic Load Balancing (ADLB) library on the BG/P. Image credit: Steve Pieper, Argonne National Laboratory.

Contact ▸ James Vary
Iowa State University | jvary@iastate.edu

Calculating Laser-induced Ultrafast Magnetism

Researchers led by Guoping Zhang from Indiana State University are focusing on the emerging field of laser-induced femtosecond demagnetization in ferromagnets with potential applications in magnetic recording. Their goal is to reveal the mechanism of ultrafast demagnetization. This work is significant in targeting a new field with a new concept (femtomagnetism) and has broader applications in magnetic recording (faster switching). The results will have important applications in high-temperature superconductors and transition metal oxides.

To achieve their research goal, they are using nickel thin films to investigate the change in magnetization induced by a femtosecond laser. The initial test run was very successful. With the Director's Discretionary allocation from Argonne, they have shown for the first time that correlation between optical and magnetic response indeed exists for longer laser pulses. There is a clear dependence on the laser pulse duration and the excitation energy. The researchers completely mapped out how the magnetization is changed in the entire Brillouin zone. Their results have been published in *Nature Physics* and *Physical Reviews.* This paves the way to the future manipulating orbital momentum in ferromagnetic materials on a several- hundred-femtosecond time scale.

One key challenge of this project is the huge number of crystal momentum points to converge the results, but Nature has been very kind. These points are intrinsically parallelable. This is where the massively parallel computing power at the Argonne Leadership Computing Facility plays a big role. One team member, Dr. Yihua Bai, Center for Instruction, Research and Technology (CIRT) at Indiana State University, developed the parallel FEMTOMAG code. This code fully exploits the parallel computing power. This finally enables the researchers to pursue two new frontiers: (1) orbital momentum excitation, and (2) coherent control of spin dynamics.

Director's Discretionary Allocation:

3.5 Million Hours

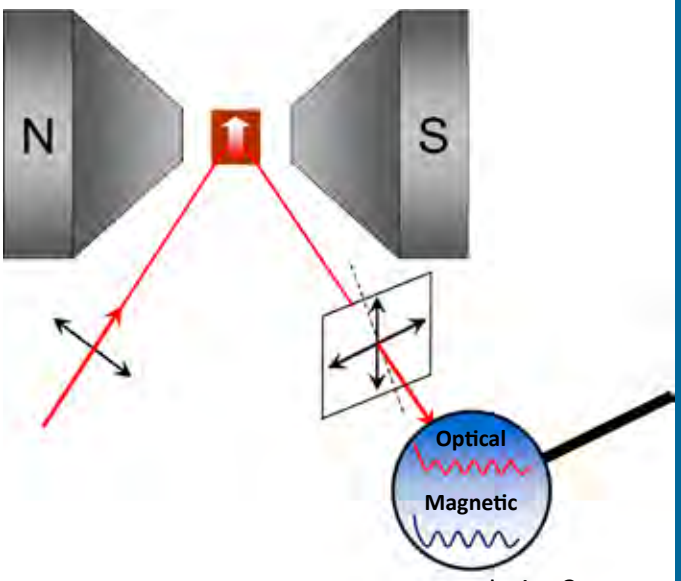

Schematic geometry of the time-resolved magneto-optical Kerr effect (TRMOKE). Only the optical signal (Kerr ellipticity and rotation) is measured, and how well the signal reflects the magnetization change on the time scale is key to the success of TRMOKE. Image credit: Dr. Yihua Bai, Ms. Sarah Edwardson-Pigg (CIRT at Indiana State University).

Contact ▸ Guoping Zhang
Indiana State University | gpzhang@indstate.edu

physics

Cosmic Structure Probes of the Dark Universe

Dark energy and dark matter are the dominant components of the Universe. Their ultimate nature, however, remains a mystery. Ambitious ground- and space-based missions investigating different aspects of the "Dark Universe" constitute a major national and international initiative. The discovery potential of these missions relies on theoretical modeling of the formation and evolution of the large-scale structure of the Universe. Remarkably, observational error estimates for various cosmological statistics will soon be at the percent level, setting a very high bar for simulation capabilities required to match the details captured in the data and to properly interpret them.

A team of researchers led by Salman Habib, Los Alamos National Laboratory, is using the Argonne Leadership Computing Facility to carry out some of the largest high-resolution simulations of the distribution of matter in the Universe, resolving galaxy-scale mass concentrations over observational volumes representative of state-of-the-art sky surveys. A key aspect of the project will be a major simulation suite covering approximately a hundred different cosmologies -- an essential resource for interpreting next-generation observations. This effort targets an approximately two- to three-orders-of-magnitude improvement over currently available resources. The database created from this project will be an essential component of Dark Universe science for years to come.

The simulation program is based around the new HACC (Hardware/Hybrid Accelerated Cosmology Code) framework aimed at exploiting emerging supercomputer architectures such as the IBM BG/Q arriving at Argonne in 2012. HACC's medium-resolution infrastructure has been tested for scalability on the current BG/P across the full machine and is now running on a BG/Q prototype system. Porting the high-resolution modules is currently in progress; the scaling properties of HACC do not depend on the short-range force calculations.

Early Science Program Allocation:

12 Million Hours

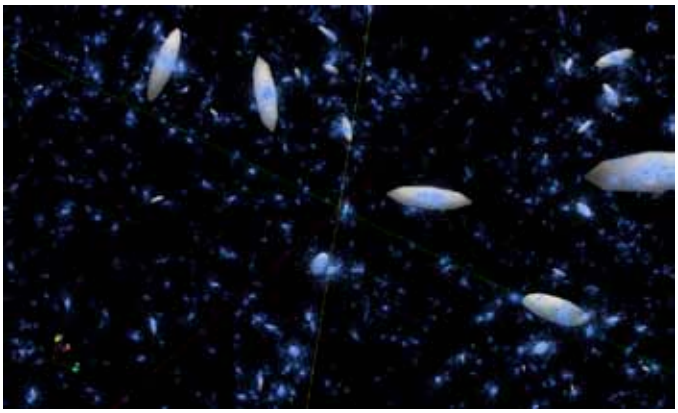

Parallel visualization (using ParaView) of the distribution of mass as traced by particles in a sub-volume of a large cosmic simulation. Mass clumps into compact structures called halos. The ellipsoids are oriented along the velocity vector of the halo.

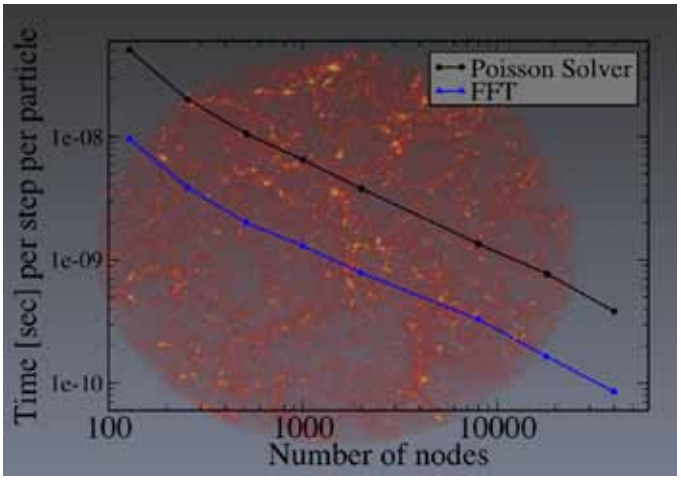

Scaling of the HACC Poisson solver across the full Argonne IBM Blue Gene/P, Intrepid. The background shows the density field from an 8-billion-particle HACC simulation run.

Contact ▸ **Salman Habib**
Los Alamos National Laboratory | habib@lanl.gov

"Cosmic Yardsticks" Extend Understanding of Dark Energy

Leadership-class computing provides an unparalleled ability to model and simulate Type Ia (thermonuclear-powered) supernovas. The ability to do 3-D, large-scale simulations of these explosions led to the discovery of an entirely new and unexpected explosion mechanism, termed the gravitationally confined detonation (GCD) model, by researchers at the Flash Center for Computational Science at The University of Chicago.

Don Lamb, the center's director, and his team at the Flash Center have done extensive simulations of Type Ia supernovas, as well as key physical processes important to the explosions, with computing time at the Argonne Leadership Computing Facility awarded through DOE's INCITE program. Type Ia supernovas are among the brightest explosions in the universe, making them important tools for measuring cosmological distances. The goal of the center's Type Ia supernova project is to better understand these explosions, and by doing so, help observers use them to determine the properties of dark energy, a hypothetical form of energy that permeates all of space. Understanding dark energy ranks among the most compelling problems in all of physical science.

The FLASH Center team is using information gained from the simulations to develop a better flame model for simulating Type Ia supernovas. The simulations have also produced new insights into how these stars explode and led to the discovery of robust signatures for the different supernova models, holding the promise that observations can discriminate among them.

Currently, the team is working with the Sloan Digital Sky Survey II (SDSS-II) Supernova Survey Team and its collaborators to confront the light curves and spectra predicted by the simulations with high-quality data.

INCITE Allocation:

80 Million Hours

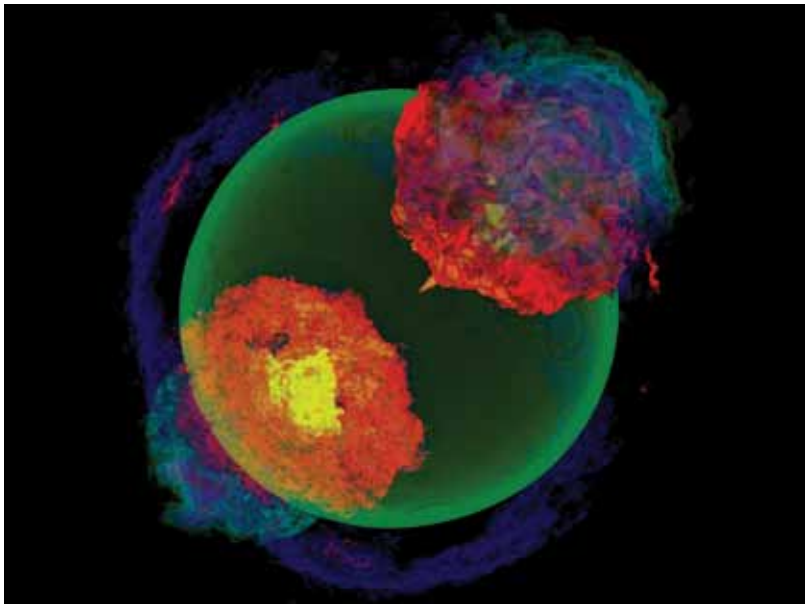

Type Ia supernovas, among the brightest exploding stars in the universe, are "cosmic yardsticks" whose application led to the discovery of dark energy. Astrophysicists are creating leadership computing simulations to better understand these explosions, and by doing so, help observers use them to determine the properties of dark energy. Shown is a 3-D, whole-star simulation of the GCD model of a Type Ia supernova at the moment at which a detonation wave begins to sweep through the white dwarf star, incinerating it. Image courtesy Flash Center for Computational Science, The University of Chicago.

Contact ▸ Don Lamb
The University of Chicago | d-lamb@uchicago.edu

Deepening the Understanding of Interactions Between Quarks and Gluons

While the behavior of protons and neutrons is well understood, less is known about the interactions of quarks and gluons, which compose them. In the field of Quantum Chromodynamics (QCD), scientists use a four-dimensional lattice representation of space-time and the assistance of supercomputers to study these interactions. This research plays a key role in ongoing efforts to develop a unified theory of the fundamental forces of nature.

Logging over 800 million core hours on Argonne's Blue Gene/P, scientists have generated gauge configurations with up, down, and strange quarks on lattices that are sufficiently fine-grained and have sufficiently small up and down quark masses to enable the extrapolation of key quantities to their physical values found in nature. The gauge configurations are being used to determine a wide range of physical quantities of importance in high energy and nuclear physics.

Using the Blue Gene/P, the generation of gauge configurations has been accelerated in many cases by a factor of 5 to 10 over what was previously possible with other machines.

Domain-wall configuration ensembles of lattice spacings 0.114 femtometers (fm) and 0.086 fm were completed and analyzed on lattices of sizes 24^3x64 and 32^3x64, respectively. These are the largest domain-wall lattices ever attempted. The most challenging staggered ensembles generated to date (lattice spacing of 0.06 and 0.045 fm) also have been completed and analyzed. These results are enabling precise tests of the Standard Model, aiding in a deeper understanding of fundamental physics.

Improved versions of both methods for lattice fermions are under way. For domain-wall fermions, the newly developed "AuxDet" method will permit a closer approach to the physical, light quark limit. For staggered fermions, an improved discretization method ("hisq" fermions) substantially reduces discretization errors. New ensembles with the improved methods are expected soon.

INCITE Allocation:

50 Million Hours

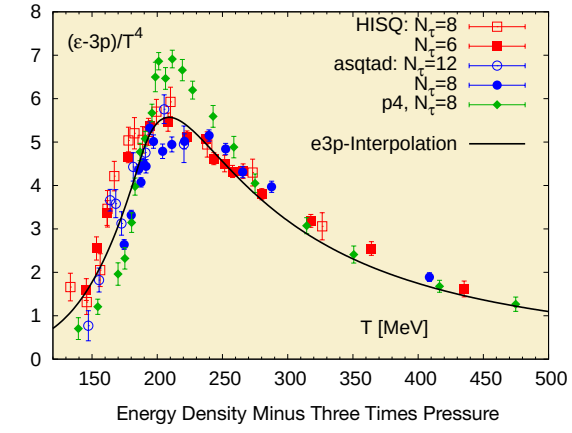

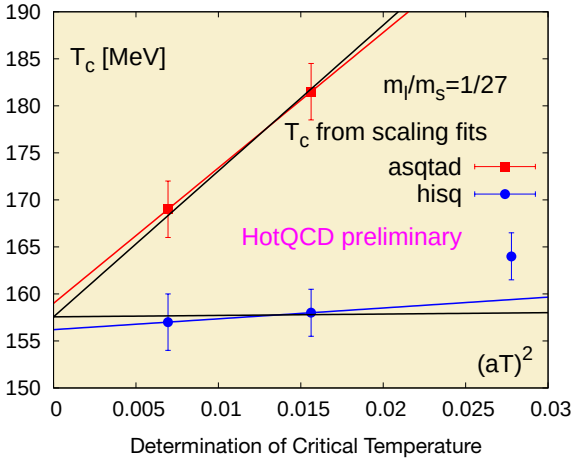

The lattice QCD calculations perfomed at the ALCF of the thermodynamics of strongly interacting matter aim at an accurate determination of the transition temperature of strongly interacting elementary particle matter and its equation of state. This provides input to the description of expanding hot and dense matter in the early universe as well as those created in heavy ion collisions at the Relativistic Heavy Ion Collider at Brookhaven National Laboratory. Image credit: Paul Mackenzie, Fermilab.

Contact ▸ Paul Mackenzie
Fermilab | mackenzie@fnal.gov

Direct Multiobjective Optimization of Storage Ring Lattices for the APS Upgrade and Beyond

The brightest storage ring-generated x-ray beams in the Western Hemisphere are created by Argonne's Advanced Photon Source (APS) and are used by more than 5,000 scientists worldwide. A planned upgrade to the APS will reconfigure the facility's magnets (its "lattice") to enhance this world-class resource. The addition of long superconducting devices will increase brightness by an order of magnitude for x-rays above 20 keV compared to standard APS devices. The upgrade will also accommodate systems for dramatically reducing the x-ray pulse length, giving the APS a unique position for enabling time-resolved science with hard x-rays. Without disruption to current operating modes, the upgrade will result in an improved source of high-energy, high-brightness, tunable x-rays for scientific research.

Scientists at work on the APS upgrade are challenged with optimizing the nonlinear dynamics to provide both sufficient dynamic aperture (to ensure high-injection efficiency) and momentum aperture (to ensure sufficient beam lifetime). To tackle this challenge, researchers will pair the extreme computing power of the Blue Gene/P at the Argonne Leadership Computing Facility with the APS-developed code "elegant" to create complex particle-tracking simulations.

The vast compute power of the Blue Gene/P gives scientists at work on the APS upgrade the ability to resolve more challenging problems faster. To keep pace with the community's seemingly insatiable appetite for increased brightness, researchers will use a portion of their ALCC allocation to advance important concept work for next-generation "ultimate" storage rings.

ALCC Allocation:

36 Million Hours

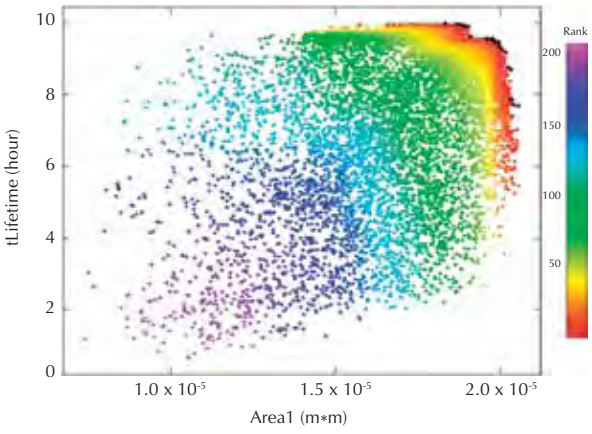

Map of solutions found by multi-objective evolutionary optimization for APS-U lattice with the Short Pulse X-ray (SPX) optics insertion for a chromaticity of 5 in both planes, obtained on Intrepid. The first-rank solutions, exhibiting the best lifetime and injection area, are marked with black symbols.

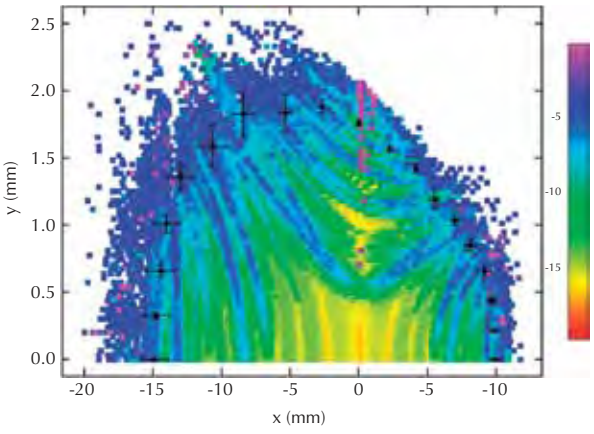

Dynamic aperture for 50 error ensembles overlaid on frequency map for APS-U lattice with SPX and Reduced Horizontal Beamsize insertions. The optimization process successfully excludes strong resonances from the interior region.

Contact ▸ Michael Borland
Argonne National Laboratory | borland@aps.anl.gov

Simulating Laser-Plasma Interactions in Targets for the National Ignition Facility and Beyond

Lawrence Livermore National Laboratory (LLNL) has been tasked with achieving ignition at the National Ignition Facility (NIF). An important aspect of the ignition campaign involves quantitative prediction of the level of laser backscatter in these targets. Mitigation of laser backscatter is important, as backscatter reduces the amount of input energy available for driving the fusion process. It can also alter implosion symmetry as well as preheat the ignition capsule via generation of hot electrons.

Recent experimental results from the National Ignition Campaign at NIF show that backscatter occurs in the laser target where quads of laser beams are overlapped. The goal of these simulations is to quantify how overlapping beam quads impact laser backscatter. Shown here are two different laser beam inputs to pF3D simulations performed on Intrepid, the IBM Blue Gene/P supercomputer at the Argonne Leadership Computing Facility (ALCF). These simulations used 80% of the machine for approximately three weeks. The first simulation (top image) determined that two overlapping quads in the region of backscatter can share a reflected light wave, which acts to increase reflectivity. The second simulation (bottom image) determined that spatial non-uniformity across the quads, as is imposed by crossed-beam transfer from other quads of NIF beams, also acts to increase reflectivity. Both effects are important to include when simulating beam propagation and backscatter in National Ignition Campaign targets.

The simulations being conducted by LLNL researchers will generate scientific results that will have a major impact on the national ignition campaign—inertial fusion—as well as on the fundamental science of LPI. These state-of-the-art simulations are only possible because of the INCITE award allocated to Intrepid, the Blue Gene/P system at the ALCF.

INCITE Allocation:

50 Million Hours

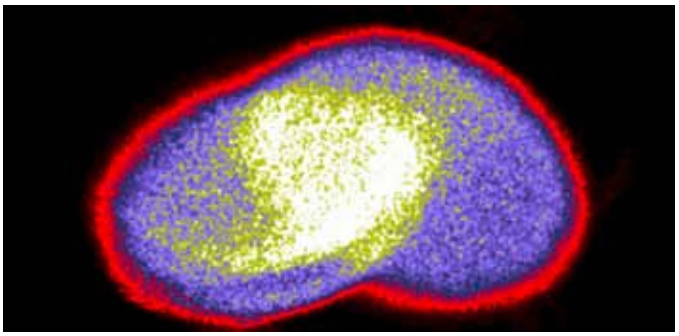

Laser input to a pF3D simulation of two NIF laser quads propagating through an ignition target. Here, power transferred from other quads of laser beams is distributed uniformly across the laser beams. Two quads overlap in the simulated region. This enhances reflectivity through a shared reflected light wave.

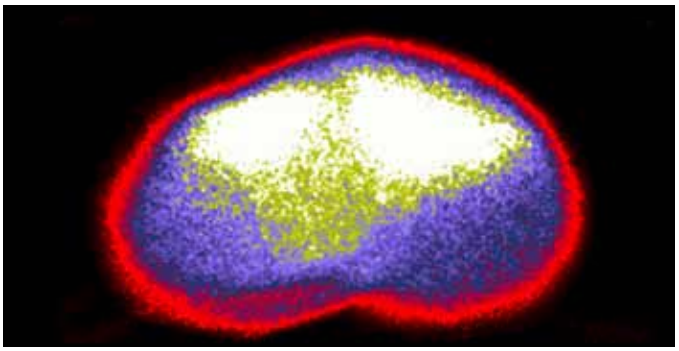

Laser input to a pF3D simulation of two NIF laser quads propagating through an ignition target. Here, power transferred from other quads provides a spatially non-uniform distribution of power across the beams. The bright "triangle" in the upper region of each laser quad drives higher levels of reflectivity within each quad. The overlap of the two quads drives reflectivity through a shared reflected light wave.

Contact ▸ Denise Hinkel
Lawrence Livermore National Laboratory | hinkel1@llnl.gov

next-generation HPC

Coming Soon to the ALCF: Mira

Watch for Mira, IBM's next-generation Blue Gene/Q supercomputer, to be delivered to the Argonne Leadership Computing Facility (ALCF) in 2012. Computational scientists and engineers look to the speed, memory size, and disk storage capacity of this leadership-class system to propel innovation in science and technology.

The 10-petaflops supercomputer will feature 48 K 16-way compute nodes (768 K processors), and 768 terabytes of memory. Like the ALCF's Intrepid, the IBM Blue Gene/P system, Mira will be made available to scientists from industry, academia, and government research facilities around the world.

Argonne anticipates that Mira will be one of the fastest and most energy-efficient supercomputers in the world after its construction and installation are complete, due to a combination of innovative new chip designs and extremely efficient water cooling.

Argonne also envisions Mira as a stepping-stone to exascale-class computers. Exascale computing has the potential to address a class of highly complex work loads that formerly have been beyond scientists' reach—not just due to their sheer size, but because of their inherent uncertainties and unpredictability—challenges like understanding the impacts of regional climate change and the design of safe nuclear reactors.

Mira will offer an opportunity for scientists to become familiar with characteristics of exascale systems and the programming changes they will require. For example, it will provide a platform for scaling current computer codes to more than 750,000 individual computing cores and exploring thread-oriented programming. This capability will give researchers preliminary experience on how scalability might be achieved on an exascale-class system with hundreds of millions of cores.

Once Mira is in production, computer time will be awarded through the Department of Energy's INCITE and ALCC programs, as well as through Argonne's Director's Discretionary projects. Watch for the remarkable scientific advances made possible by Mira in the blink of an eye.

research projects

INCITE

In 2011, 30 projects in diverse scientific disciplines were awarded approximately 732 million supercomputer core-hours on the Blue Gene/P at the ALCF through the Department of Energy's INCITE (Innovative and Novel Computational Impact on Theory and Experiment) Program. Of these, 12 projects were renewed from 2010, and 18 are new projects.

biological sciences

Multiscale Blood Flow Simulations
George Karniadakis, Brown University
Allocation: 50 Million Hours

Protein-Ligand Interaction Simulations and Analysis
T. Andrew Binkowski, Argonne National Laboratory
Allocation: 20 Million Hours

Simulation and Modeling of Membranes Interactions with Unstructured Proteins and Computational Design of Membrane Channels for Absorption of Specified Ions
Igor Tsigelny, University of California–San Diego
Allocation: 4 Million Hours

Towards Breakthroughs in Protein Structure Calculation and Design
David Baker, University of Washington
Allocation: 30 Million Hours

chemistry

Ab Initio Dynamical Simulations for the Prediction of Bulk Properties
Theresa Windus, Iowa State University
Allocation: 10 Million Hours

Large Eddy Simulation of Two-phase Flow Combustion in Gas Turbines
Thierry Poinsot, European Centre for Research and Advanced Training in Scientific Computation
Allocation: 10 Million Hours

Potential Energy Surfaces for Simulating Complex Chemical Processes
Donald Truhlar, University of Minnesota
Allocation: 15 Million Hours

Simulations of Deflagration-to-Detonation Transition in Reactive Gases
Alexei Khokhlov, The University of Chicago
Allocation: 18 Million Hours

computer science

Performance Evaluation and Analysis Consortium End Station
Patrick Worley, Oak Ridge National Laboratory
Allocation: 10 Million Hours

Scalable System Software for Performance and Productivity
Ewing Lusk, Argonne National Laboratory
Allocation: 5 Million Hours

Trace Collection for Simulation-driven Co-design of Exascale Platforms and Codes
David Evensky, Sandia National Laboratory
Allocation: 5 Million Hours

earth science

Climate-Science Computational Development Team: The Climate End Station II
Warren Washington, National Center for Atmospheric Research
Allocation: 40 Million Hours

Deterministic Simulations of Large Regional Earthquakes at Frequencies up to 4Hz
Thomas Jordan, University of Southern California
Allocation: 10 Million Hours

Numerical Study of Multiscale Coupling in Low-Aspect Ratio Rotating Stratified Turbulence
Susan Kurien, Los Alamos National Laboratory
Allocation: 35 Million Hours

energy technologies

Advanced Reactor Thermal Hydraulic Modeling
Paul Fischer, Argonne National Laboratory
Allocation: 25 Million Hours

Large Eddy Simulation for Green Energy and Propulsion Systems
Umesh Paliath, GE Global Research
Allocation: 20 Million Hours

Understanding the Ultimate Battery Chemistry: Rechargeable Lithium/Air
Jack Wells, Oak Ridge National Laboratory
Allocation: 15 Million Hours

research projects

engineering

Detached-Eddy Simulations and Noise Predictions for Tandem Cylinders
Philippe Spalart, Boeing
Allocation: 45 Million Hours

Simulation of High Reynolds Number Turbulent Boundary Layers
Robert Moser, University of Texas at Austin
Allocation: 40 Million Hours

Turbulent Multi-material Mixing in the Richtmyer-Meshkov Instability
Sanjiva Lele, Stanford University
Allocation: 12 Million Hours

Uncertainty Quantification for Turbulent Mixing
James Glimm, Stony Brook University
Allocation: 10 Million Hours

materials science

High-Fidelity Simulation of Complex Suspension Flow for Practical Rheometry
William George, National Institute of Standards and Technology
Allocation: 25 Million Hours

Petascale Simulations of Stress Corrosion Cracking
Priya Vashishta, University of Southern California
Allocation: 45 Million Hours

Probing the Non-scalable Nano Regime in Catalytic Nanoparticles with Electronic Structure Calculations
Jeffrey Greeley, Argonne National Laboratory
Allocation: 15 Million Hours

Vibrational Spectroscopy of Liquid Mixtures and Solid/Liquid Interfaces
Giulia Galli, University of California–Davis
Allocation: 15 Million Hours

physics

Advanced Simulations of Plasma Microturbulence at the Petascale and Beyond
William Tang, Princeton Plasma Physics Laboratory
Allocation: 8 Million Hours

Lattice QCD
Paul Mackenzie, Fermi National Accelerator Laboratory
Allocation: 50 Million Hours

Nuclear Structure and Nuclear Reactions
James Vary, Iowa State University
Allocation: 15 Million Hours

Simulations of Laser-plasma Interactions in Targets for the National Ignition Facility and Beyond
Denise Hinkel, Lawrence Livermore National Laboratory
Allocation: 50 Million Hours

Study of Buoyancy-driven Turbulent Nuclear Burning and Validation of Type Ia Supernova Models
Donald Lamb, The University of Chicago
Allocation: 80 Million Hours

research projects

ASCR Leadership Computing Challenge

Open to scientists from the research community in academia and industry, the ASCR Leadership Computing Challenge (ALCC) program allocates up to 30 percent of the computational resources at the Argonne Leadership Computing Facility, NERSC and Oak Ridge. Projects in the program are of special interest to the Department, with an emphasis on high-risk, high-payoff simulations in areas directly related to the Department's energy mission, national emergencies, or for broadening the community of researchers capable of using leadership computing resources. Proposals are awarded an ALCC allocation based on a peer review for scientific merit and computational readiness.

chemistry

Electrocatalyst Durability from First Principles Calculations
PI: Jeffrey Greeley, Argonne National Laboratory
Allocation: 20 Million Hours

earth science

Sensitivity and Uncertainty of Precipitation of a Climate Model
PI: Laura Zamboni, Argonne National Laboratory
Allocation: 40 Million Hours

Simulating Regional Climate at Convection-permitting Resolution
PI: Greg Holland, National Center for Atmospheric Research (NCAR) Earth System Laboratory
Allocation: 13 Million Hours

energy technologies

First Principles Calculations of Interfaces in Electrical Energy Storage Systems
PI: Larry Curtiss, Argonne National Laboratory
Allocation: 30 Million Hours

engineering

Prediction of Supersonic Jet Noise Using Large Eddy Simulation
PI: Parviz Moin, Stanford University
Allocation: 60 Million Hours

The Interactions Between Vaporizing Liquid Droplets and a Turbulent Flow: Fully Resolved Direct Numerical Simulation
PI: Said Elghobashi, University of California, Irvine
Allocation: 10 Million Hours

materials science

Toward Crystal Engineering from First Principles
PI: James Chelikowsky, University of Texas at Austin
Allocation: 12 Million Hours

research projects

Early Science Program (ESP)

Allocations through the Early Science Program (ESP) provide researchers with preproduction hours (between system installation and full production) on the ALCF's next-generation, Mira, the 10-petaflops IBM Blue Gene/Q system. This early science period provides projects with a significant head start for adapting to the new machine and access to substantial computational time. During this shakedown period, users assist in identifying the root causes of any system instabilities, and work with ALCF staff to help develop solutions. More than two billion core hours are allocated through ESP.

biological sciences

Multiscale Molecular Simulations at the Petascale
PI: Gregory Voth, The University of Chicago
Allocation: 150 Million Hours Materials Science

NAMD - The Engine for Large-Scale Classical MD Simulations of Biomolecular Systems Based on a Polarizable Force Field
PI: Benoit Roux, The University of Chicago
Allocation: 80 Million Hours

chemistry

Accurate Numerical Simulations of Chemical Phenomena Involved in Energy Production and Storage with MADNESS and MPQC
PI: Robert Harrison, Oak Ridge National Laboratory
Allocation: 150 Million Hours

High-Accuracy Predictions of the Bulk Properties of Water
PI: Mark Gordon, Iowa State University
Allocation: 150 Million Hours

High-Speed Combustion and Detonation (HSCD)
PI: Alexei Khokhlov, The University of Chicago
Allocation: 150 Million Hours

earth science

Climate-Weather Modeling Studies Using a Prototype Global Cloud-System Resolving Model
PI: Venkatramani Balaji, Geophysical Fluid Dynamics Laboratory
Allocation: 150 Million Hours

engineering

Direct Numerical Simulation of Autoignition in a Jet in a Cross-Flow
PI: Christos Frouzakis, Swiss Federal Institute of Technology
Allocation: 150 Million Hours

Petascale, Adaptive CFD
PI: Kenneth Jansen, University of Colorado–Boulder
Allocation: 150 Million Hours

Petascale Direct Numerical Simulations of Turbulent Channel Flow
PI: Robert Moser, University of Texas
Allocation: 60 Million Hours

geophysics

Using Multi-scale Dynamic Rupture Models to Improve Ground Motion Estimates
PI: Thomas Jordan, University of Southern California
Allocation: 150 Million Hours

materials science

Materials Design and Discovery: Catalysis and Energy Storage
PI: Larry Curtiss, Argonne National Laboratory
Allocation: 50 Million Hours

physics

Ab-initio Reaction Calculations for Carbon-12
PI: Steven C. Pieper, Argonne National Laboratory
Allocation: 110 Million Hours

Cosmic Structure Probes of the Dark Universe
PI: Salman Habib, Los Alamos National Laboratory
Allocation: 150 Million Hours

Global Simulation of Plasma Microturbulence at the Petascale and Beyond
PI: William Tang, Princeton Plasma Physics Laboratory
Allocation: 50 Million Hours

Lattice Quantum Chromodynamics
PI: Paul Mackenzie, Fermilab
Allocation: 150 Million Hours

Petascale Simulations of Turbulent Nuclear Combustion
PI: Don Lamb, The University of Chicago
Allocation: 150 Million Hours

research projects

Director's Discretionary

Discretionary allocations are "start up" awards made to potential future INCITE projects. Projects must demonstrate a need for leadership-class resources. Awards may be made year round to industry, academia, laboratories and others, and are usually between three and six months in duration. The size of the award varies based on the application and its readiness/ability to scale; awards are generally from the low tens of thousands to the low millions of hours. A sampling of 2011 Director's Discretionary projects appears below.

Biological Sciences

Computer-Aided Design of Vaccine Nanoparticles
PI: Peter Ortoleva, Indiana University–Bloomington
Allocation: 2.25 Million Hours

Load-Balanced Spike Communication in Large-Scale Neural Networks
PI: Michael Hines, Yale University
Allocation: 770,000 Hours

Computer Science

I/O Forwarding Scalability Layer
PI: Rob Ross, Argonne National Laboratory
Allocation: 1.5 Million Hours

Many Task Computing Science Applications
PI: Michael Wilde, Argonne National Laboratory
Allocation: 3 Million Hours

Earth Science

Large-Scale Hurricane Simulations
PI: James Done, National Center for Atmospheric Research
Allocation: 13 Million Hours

Oil Plume Behavior
PI: Tamay Özgökmen, University of Miami
Allocation: 2.3 Million Hours

Engineering

Interactions between Vaporizing Liquid Droplets and a Turbulent Flow: Fully Resolved DNS
PI: Said Elghobashi, University of California–Irvine
Allocation: 1.2 Million Hours

Large Eddy Simulation of Turbulent Mixing by Rayleigh-Taylor Instability
PI: James Glimm, Stony Brook University
Allocation: 1.5 Million Hours

Physics

First-Principles Calculations of Laser-Induced Ultrafast Magnetism
PI: Guoping Zhang, Indiana State University
Allocation: 3.5 Million Hours

On the cover

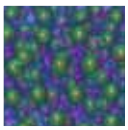

The electronic charge density of a lithium oxide (Li_2O) nanoparticle consists of 1500 atoms obtained from Density Functional Theory simulation. *Image credit: Kah Chun Lau (MSD, ANL), Aaron Knoll (MCS, ANL), Larry A. Curtiss (MSD/CNM, ANL).*

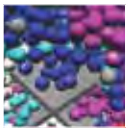

Flow simulations of thousands of particles on the Argonne Leadership Computing Facility's Blue Gene/P supercomputer are enabling new insights into how to measure and control flow properties of large-particle dense suspensions like concrete that cannot currently be accurately measured in industrial settings. A better understanding of these properties will help ensure the optimum performance of concrete and eliminate cost overruns. *Image credit: Images and the software used to produce them were developed by Steven Satterfield, John Hagedorn, and John Kelso of NIST and Marc Olano of NIST and the University of Maryland–Baltimore County.*

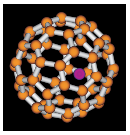

Researchers performed SA-CASSCF calculations of $B@C_{60}$—boron atom inside the fullerene cage shown here. The boron atom occupies a shallow energy-minimum "off center" due to charge transfer from the boron atom to the case wall. For these calculations, ALCF catalyst on the project, Graham Fletcher, developed a version of the GAMESS code for Intrepid that runs twice as fast as the standard version of GAMESS. *Image credit: Osanna Tischenko, University of Minnesota.*

Side view of a buoyant plume interacting with ocean flows near the surface. Note the vertical coherence of the plume and inability to cross the base of the oceanic surface mixed layer. *Image credit: Tamay Özgökmen, University of Miami; Paul Fischer and Aleks Obabko, Argonne National Laboratory; and Hank Childs, Lawrence Livermore National Laboratory.*

Argonne NATIONAL LABORATORY | Argonne Leadership Computing Facility

Argonne Leadership Computing Facility
Argonne National Laboratory
9700 South Cass Avenue, Bldg. 240
Argonne, IL 60439-4847

www.alcf.anl.gov

CONTACT ▸ info@alcf.anl.gov | (877) 737-8615

alcf-scibro-1111