# 2012 SCIENCE HIGHLIGHTS

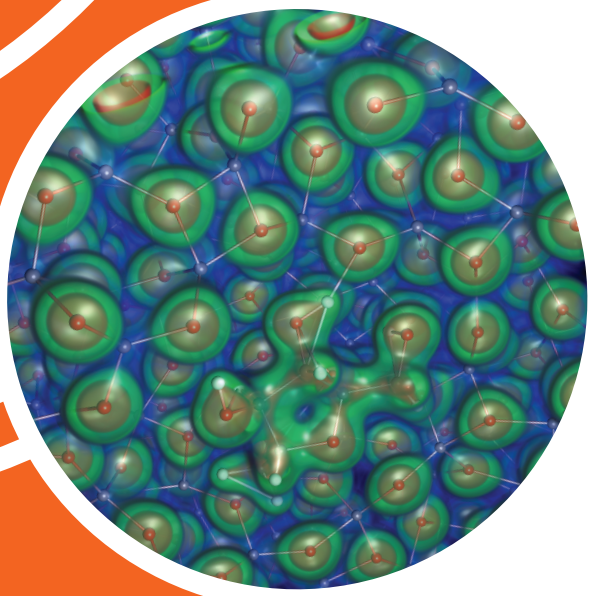

Argonne NATIONAL LABORATORY

Argonne Leadership Computing Facility

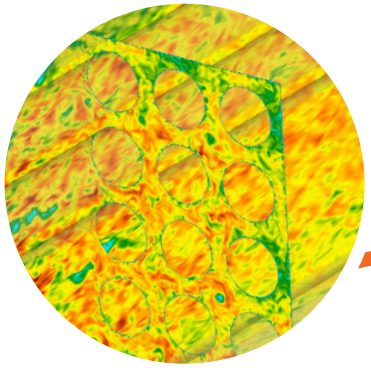

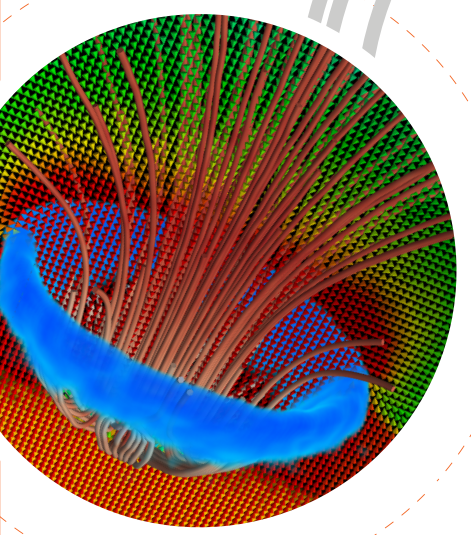

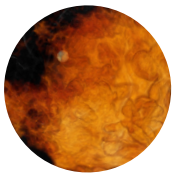

Argonne **Leadership Computing** Facility

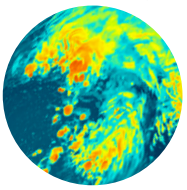

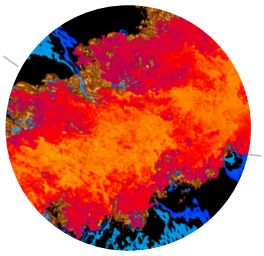

CONTENTS

# CONTENTS

# LEADERSHIP-CLASS COMPUTING EXPANDS SCIENTIFIC HORIZONS

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. Argonne researchers work closely with researchers from industry 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 ALCF's 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

Mira, the ALCF's magnificent 10-petaflops system, will enable scientific discoveries only imagined in the past. 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 next-generation IBM Blue Gene/Q supercomputer features 48K 16-way compute nodes (768K processors), and 768 terabytes of memory.

Tukey will be twice as fast as its predecessor, Eureka, enabling unprecedented levels of productivity and capability in visualization and data analysis. Tukey shares the Mira network and parallel filesystem, enabling direct access to data generated by Mira simulations.

## SERVICES FOR USERS

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.

# ALCF ACHIEVEMENTS

ALCF's research and resources were honored with the following prestigious awards in 2012:

- A paper on cosmology research led by Argonne's Salman Habib and co-PI Katrin Heitmann, "The Universe at Extreme Scale: Multi-Petaflop Sky Simulation on the BG/Q," earned a finalist designation for the ACM Gordon Bell Prize and will be included in the Gordon Bell Prize sessions at SC12.
- Mira is the third fastest supercomputer in the world, according to the June 2012 TOP500 list. It achieved 8.1 petaflops per second on the LINPACK benchmark, using 786,432 processing cores on 48 racks.
- Argonne's Mira tied for first place on the Graph 500 list with Sequoia at Lawrence Livermore National Laboratory. Each supercomputer achieved a score of more than 3,500 GTEPS (giga traversed edges per second).
- Blue Gene/Q systems hold the top five spots on the Green 500 list that ranks the top 500 supercomputers in the world by energy efficiency. Due to innovative new chip designs and extremely efficient water cooling, Mira will be five times more energy efficient than Intrepid, its Blue Gene/P predecessor at the ALCF.

**Paul Messina**
Director of Science
Argonne Leadership Computing Facility

# EARLY SCIENCE PROGRAM PREPARES MIRA FOR NEW OPPORTUNITIES IN INNOVATION

Mira, the 10-petaflops IBM Blue Gene/Q, was installed in 2012 and–as planned–is starting to be used by the projects that participate in our Early Science Program (ESP). The ESP projects play essential roles in the initial stages of our new system. Leadership systems enable computational science and engineering research that promises to have high impact and advance state-of-the-art modeling and simulation and that cannot be done practically on lesser machines. Enhancing applications software so that it can tackle new research goals often requires substantial changes—new algorithms, mathematical models that simulate new or additional phenomena, different mesh generation, etc. The ESP projects proposed new simulation campaigns that include new models, methods, and algorithms, as well as computational problem scales demanding the size and speed of Mira—mostly those not possible to run on Intrepid.

In addition, the architecture of new systems is always somewhat different than the previous generation. Although codes that ran on Intrepid, our Blue Gene/P system, usually run on Mira with no modifications, achieving optimal performance can take months of tuning. Given that the useful lifetime of supercomputers is typically four to five years, it is important to be poised to use them for leadership-level research as soon as possible after installation. I am pleased to report that all sixteen Early Science projects are running on Mira, and some are beginning to make runs aimed at new results. As expected, most of the codes required few or no modifications to run and for initial tuning.

A second role that ESP projects serve is to tease out bugs that leading-edge systems inevitably have and to characterize the behavior of new hardware and software features. Scientific applications utilize system features in different ways than standard benchmarks suites or the tests provided by the computer designers. Therefore, throughout the history of advanced computers, what we often call "real apps" are the ultimate tools for shaking out quirks in new systems. And yes, ESP project participants, in collaboration with our catalysts and performance engineers, have found bugs, many of which have already been fixed by IBM's highly responsive and dedicated software engineers.

Our ESP has a third vital component: the BG/Q Performance Tools, Debuggers, Programming Models, and Libraries project. Through this partnership, developers, twelve performance tools, five programming models, and two vendor debuggers have been ported to Mira, and fifteen libraries have been installed. The availability of these software tools at the beginning of the system's life will contribute to the success of all the research projects.

Mira presents new opportunities for innovation, given its speed, memory size, and disk storage capacity. Through the Early Science Program, those opportunities will be pursued starting now.

# MIRA USHERS IN A NEW ERA OF SCIENTIFIC SUPERCOMPUTING

Mira, the new petascale IBM Blue Gene/Q system being installed at the ALCF, ushers in a new era of scientific supercomputing. An engineering marvel, the 10-petaflops machine is capable of carrying out 10 quadrillion calculations per second. To put this speed into perspective: If every single person on earth performed one calculation per second, around the clock, it would take them more than two weeks to do the work that Mira will accomplish in one second!


As computers get faster and more powerful, they must also become more energy efficient to be practical. Mira is among the 20 IBM Blue Gene/Q systems rated by the Green500.org to be the most energy-efficient supercomputers in the world and will be five times more energy efficient than Intrepid, its Blue Gene/P predecessor. To gain this power efficiency, copper tubes are used to pipe cold water directly alongside the Mira's chips, which saves power by eliminating an extra cooling step. Mira also fits sixteen cores in each "system on a chip" that is the building block of the system, compared to four cores on Intrepid's. The system functions on each chip include two memory controllers, caches, internal crossbar switch, chip-to-chip network router and ports, and external I/O logic. By integrating system functions and by placing multiple cores per chip, energy-expensive off-chip data transfers are significantly reduced and system reliability is improved thanks to a lower chip count for the entire system.

Mira will be in full production in 2013. This remarkable next-generation supercomputer represents a milestone in the ongoing effort to develop exascale systems equipped with hundreds of millions of processors by the end of the decade.

## MIRA PROVIDES PETAFLOPS POWER FOR SCIENCE

Mira, the ALCF's next-generation Blue Gene/Q system, will consist of:

- 48 racks
- 1,024 nodes per rack
- 1.6 GHz 16-core processor and 16 GB RAM per node
- 384 I/O nodes
- 240 GB/s, 35PB storage
- For a total of 768K cores, 768 terabytes of RAM, and a peak performance of 10 petaflops. The system is capable of carrying out 10 quadrillion floating-point operations per second.

# THE ALCF'S EARLY SCIENCE PROGRAM— STEPPING STONES FOR SCIENCE ON MIRA

The multi-year effort of standing up a resource the size of Mira is far from idle time for the computational scientists and teams of researchers at the ALCF. The ALCF's Early Science Program (ESP) is using this critical pre-production time period, plus access to portions of the resource as it becomes available, to prepare key applications for the architecture and scale of Mira and to solidify necessary libraries and infrastructure.

The ESP program was launched in 2009. At that time, sixteen projects were selected to participate in the program and were allocated a combined two billion core-hours on Mira. In addition to their promise of delivering exciting new science, the projects were chosen based on their state-of-the-art, petascale applications—applications especially well suited to exploiting the unique characteristics of Blue Gene/Q architecture. The ESP projects form a representative sample of the science arenas present in the current ALCF workload, including science to simulate advanced materials, explore the universe, model biological organisms, and further the design of new, safe, and reliable sources of energy. Together, these early projects will provide a diversity and wealth of experience for future production applications.

Long before the first pieces of Mira hardware began arriving at Argonne, ESP project teams, in collaboration with ALCF and IBM staff, embarked on intensive efforts to adapt their software to take full advantage of Mira's Blue Gene/Q architecture. To aid efforts, the ALCF appointed a postdoctoral appointee to each team, provided modest allocations on Intrepid (Mira's predecessor) and granted access to prototype systems.

Broadly summarized, ESP preparations for Mira include:

- Introduction of multithreading and optimization.
- Implementation, validation, benchmarking, and optimization of new physics models required for proposed simulations.
- Implementation, testing, and optimization of new algorithms required for proposed simulations.
- Setup of proposed simulations on Mira, including preliminary and lower-resolution runs on Intrepid to establish application parameters.

These preparations have yielded significant, sometimes groundbreaking successes in code modifications for better performance and algorithmic developments to enable even greater scientific discovery once Mira is fully operational. A sampling of these achievements follows.

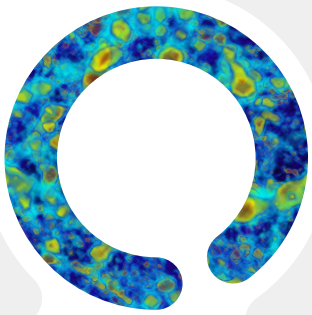

# Accurate Numerical Simulations of Chemical Phenomena Involved in Energy Production and Storage with MADNESS and MPQC

**Robert Harrison | Stony Brook University**

Award: 150 Million Hours

*A team of researchers led by Robert Harrison from Stony Brook University is using ALCF resources to create computational tools to study catalysis of chemical processes on metal-oxide surfaces, and heavy-element chemistry for nuclear fuel reprocessing.*

This computational chemistry project uses the MADNESS and MPQC packages, which use modern software development techniques including generic and object-oriented programming as well as hybrid parallelism (MPI+Pthreads). Team members ported MADNESS to the Blue Gene/Q architecture and then optimized the most computationally intensive routines using automatic code generation. New C++ math libraries such as Elemental and Eigen were incorporated to improve parallelism and robustness. As a result of all these developments, MADNESS is running efficiently on more than 1,000 nodes of Mira.

The team has incorporated a new library for exchange-correlation potentials for density-functional theory, implemented molecular geometry optimizers, and developed a prototype for linear response calculations. Together, these features provide a robust environment to simulate complicated chemical processes in surface catalysis, energy storage and condensed phase systems.

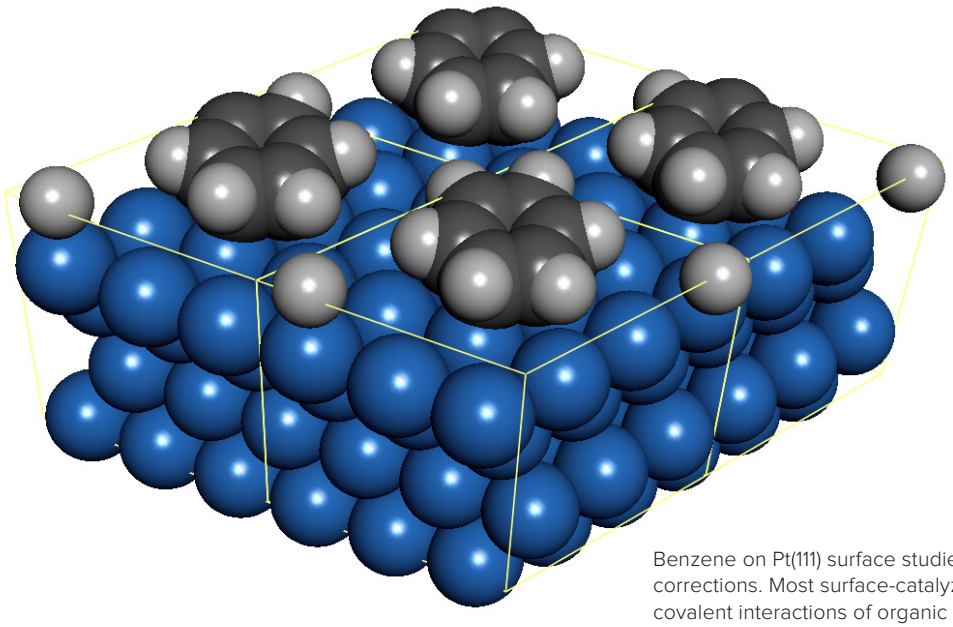

Benzene on Pt(111) surface studied with DFT with van der Waals corrections. Most surface-catalyzed reactions are driven by non-covalent interactions of organic molecules and active metallic faces. The deposition organic molecules, such as benzene, on platinum surfaces are still challenging simulations. The interaction energies and structural arrangements estimated with the methodologies developed at the ALCF allow for accurate comparisons of simulations with experiments.

**Image Credit**

Alavaro Vazquez-Mayagoitia, Argonne National Laboratory

## High-Accuracy Predictions of the Bulk Properties of Water

**Mark Gordon | Iowa State University**

Award: 150 Million Hours

*Researchers from Iowa State University led by Mark Gordon will be the first to simulate large-scale chemical systems with high accuracy and high precision, employing new, advanced techniques that allow fully* ab initio *quantum mechanical methods to scale up to problems with tens of thousands of atoms.*

This project uses the highly scalable GAMESS quantum chemistry package to execute dynamical simulations of liquid water. To maximum performance of GAMESS on Blue Gene /Q, the group is multi-threading its integral kernels via OpenMP. Multi-threading improves efficiency through enhanced instruction pipelining and memory latency hiding. The integral kernels are first modified to render them amenable to multi-threading, then OpenMP directives are inserted to distribute the kernel arithmetic over multiple execution threads.

| | | | Number of Nodes | | | | |
|---|---|---|---|---|---|---|---|
| System | | | Blue Gene/P | Blue Gene/Q* | | | |
| Waters | Atoms | Basis Functions | 2048 | 128 | 256 | 512 | 1024 |
| 128 | 384 | 5504 | 4.8 | 12.0 | 6.2 | 3.3 | 1.9 |
| 256 | 768 | 11008 | 10.5 | 31.1 | 15.8 | 8.2 | 4.4 |
| 512 | 1536 | 22016 | **28.9** | 94.2 | 48.8 | **24.9** | 12.9 |

*Using -O2 IBM XL compiler option and 4 ranks per node

A comparison of the performance of GAMESS' FMO-MP2 force calculations for several water clusters on Blue Gene/P and Blue Gene/Q. All timings are in minutes. The Blue Gene/Q calculations employed 4 ranks per node. All calculations used aug-cc-pvdz basis set. A Q/P speed-up per node of about 4.6 to 6.4 times is observed even without tuning and optimizing the GAMESS program.

## High-Speed Combustion and Detonation

**Alexei Khokhlov | The University of Chicago**

Award: 150 Million Hours

*Alexei Khokhlov and his research team from The University of Chicago is exploring the physical mechanisms of the burning and detonation of hydrogen-oxygen mixtures with the goal of aiding the design of safe systems for future use of hydrogen fuel.*

The group has capitalized on the Early Science period to refine the threading model and improve single core performance. Tuning certain data exchange and balance loops has optimized the code for both the Blue Gene/P and Blue Gene/Q systems. They team has also successfully eliminated a significant bottleneck in the mesh refinement algorithm which was influencing the code's overall scalability at and above 16,000 cores. Through their efforts, the previously non-scaling portion of code was sped up four orders of magnitude, reducing the overall runtime on 32K cores. In addition to preparing the code for scaling to Mira's larger number of processors and threads, researchers have used this time to explore the weak ignition regime of hydrogen combustion and detonation.

# Petascale Simulations of Turbulent Nuclear Combustion

**Don Lamb | The University of Chicago**

Award: 150 Million Hours

*Researchers from The University of Chicago led by Don Lamb will conduct the largest-ever hydrodynamical simulations of thermonuclear-powered Type Ia supernovae.*

The FLASH early science applications will make use of a new unsplit hydrodynamics solver. This solver is considerably more efficient and more accurate than the split solver previously used for FLASH simulations of this type. The nuclear burning flame model used for both Raleigh-Taylor Flame (RTFlame) box simulations and Deflagration-to-Detonation Transition (DDT) Type Ia Supernova (SNIa) simulations features a much-improved turbulent flame model. This model supplies a more accurate phenomenological treatment of sub-grid turbulent effects on the nuclear flame speed. Accurate modeling of these effects is essential for obtaining accurate estimates of nuclear energy release and general hydrodynamic development of SNIa explosions.

The team has multithreaded several key units in FLASH to make better use of the increased on-node parallelism of Blue Gene/Q relative to Blue Gene/P. This includes the hydrodynamics solver, flame model and Equation of State (EOS) that are key components of the FLASH early science applications. The team has improved the memory access locality in the hydrodynamics solver and EOS thereby boosting serial performance by 30%. Finally, researchers are able to avoid an expensive phase of communication which happens during FLASH initialization: In runs from time=0.0s the communication now happens between fewer MPI ranks in a reduced MPI communicator, and in restart runs the communication is avoided entirely by reading additional metadata from the checkpoint file.

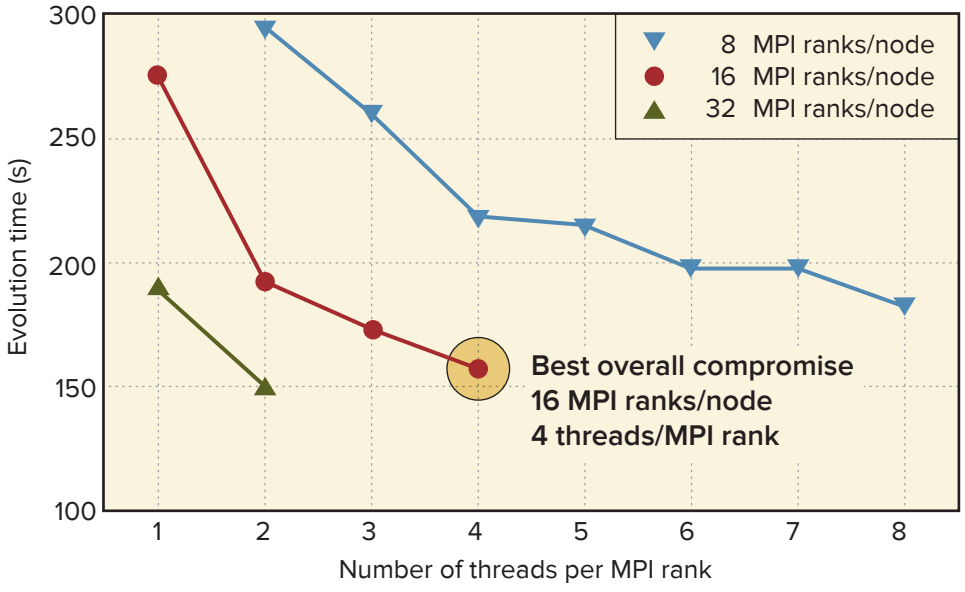


Time to complete 30 evolution steps of FLASH as a function of the number of MPI ranks and OpenMP threads per BG/Q node. Measurements obtained by running the same FLASH simulation on 128 nodes of Vesta Blue Gene/Q. The black circle shows the configuration with the best compromise between time to solution and memory per MPI rank. The high memory overhead per MPI rank for this application necessitates the compromise.

# ALLOCATION PROGRAMS FOR WORLD-CHANGING RESEARCH

The Argonne Leadership Computing Facility (ALCF) is one of two U.S. Department of Energy (DOE) Leadership Computing facilities in the nation, and provides leading scientists with next-generation, high-performance computing resources for breakthrough research to address global challenges. The ALCF invites researchers from universities, government agencies, and industry who are prepared to accelerate scientific discovery in climate, materials, alternative energy sources and energy storage, chemistry, nuclear physics, astrophysics, quantum mechanics, and the other areas of scientific research.

ALCF resources are accessible through the following programs:

## INNOVATIVE & NOVEL COMPUTATIONAL IMPACT ON THEORY AND EXPERIMENT (INCITE)

The DOE's INCITE program provides allocations to computationally intensive, large-scale research projects that aim to address "grand challenges" in science and engineering. The program encourages proposals from academia, industry, and other research institutions. DOE sponsorship is not required.

The INCITE program conducts a two-part review of all proposals: a peer review by an international panel of experts and a computational-readiness review. The annual call for proposals is issued in April and the allocations are awarded in millions of core hours for 1 to 3 years. Sixty percent of ALCF resources are allocated through the INCITE program.

## ASCR LEADERSHIP COMPUTING CHALLENGE (ALCC)

The DOE's ALCC program allocates resources to projects with an emphasis on high-risk, high-payoff simulations in areas directly related to the DOE's energy mission, national emergencies, or for broadening the community of researchers capable of using leadership computing resources. The program is available to researchers in academia, industry, and other research institutions. DOE sponsorship is not required.

The DOE conducts a peer review of all proposals based on scientific and technical merit of the project; appropriateness of the proposed method or approach; competency and adequacy of personnel and proposed resources; and the reasonableness and appropriateness of the proposed allocation request. The 2013 deadline for submitting a proposal is February 1, 2013, and the yearlong allocation cycle runs from July 1 to June 30. The average ALCC award size is low to high millions of core hours and represents up to 30% of ALCF resource allocations.

## DIRECTOR'S DISCRETIONARY (DD)

The Director's Discretionary program provides "start up" awards to researchers working toward an INCITE or ALCC allocation so that they can achieve computational readiness. The program is available to researchers in academia, industry, and other research institutions. DOE sponsorship is not required.

Projects must demonstrate a need for leadership-class resources and are reviewed by the ALCF. The award size is low tens of thousands to the low millions of core hours and the renewable awards typically last 3 to 6 months. Ten percent of ALCF resources are allocated through this program.

Please visit **alcf.anl.gov/programs** for more information on how to get an allocation at the ALCF.

# SCIENCE HIGHLIGHTS

**Benoît Roux**
The University of Chicago | roux@uchicago.edu

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

*Biology, at the atomic and molecular level, is governed by complex interactions involving a large number of key constituents, including water, ions, proteins, nucleic acids, and lipid membranes. The goal of this project is to develop new technologies to simulate biomolecular systems with unprecedented accuracy.*

Molecular dynamics (MD) simulations provide a new perspective to understand the function of biologically important molecular systems. This computational approach consists of first constructing a detailed atomic model of a macromolecular system which describes microscopic forces with a potential function (force field), and then using Newton's classical equation, $F = ma$, to literally "simulate" the dynamical motions of all the atoms as a function of time.

The success of an MD simulation is determined by the accuracy of the potential function and on the efficiency of the dynamic algorithm enabling the adequate sampling of motions. NAMD, a high-performance code under development at the University of Illinois, is one of the most optimal programs used to carry out classical simulations of biomolecular systems.

University of Chicago biophysicist Benoît Roux and his colleagues are enhancing the sampling efficiency of NAMD beyond that of brute-force MD simulations, and implementing several advanced strategies based on multiple copies such as replica-exchange MD (REMD) and Hamiltonian tempering (H-REMD). The team is also implementing a new force field that incorporates the effect of induced polarization.

The team has already carried out initial scaling tests with NAMD and the polarizable force field. The application of NAMD on the Argonne Leadership Computing Facility's BG/Q consists of first porting the code then carrying out performance/validation tests in multi-threaded mode. As NAMD achieves petascale performance the team will start to explore the next level of simulation methods and scientific questions while continuing to ensure that NAMD performance remains portable across platforms.

In addition, an all-atom polarizable force field for proteins, nucleic acids, and lipids will be available to the computational biophysics and chemistry community. The new force field will then represent the state-of-the-art for the next generation of simulation studies of proteins and lipid systems. This will provide an investigative tool of unprecedented accuracy that will be of great use for the entire scientific community.

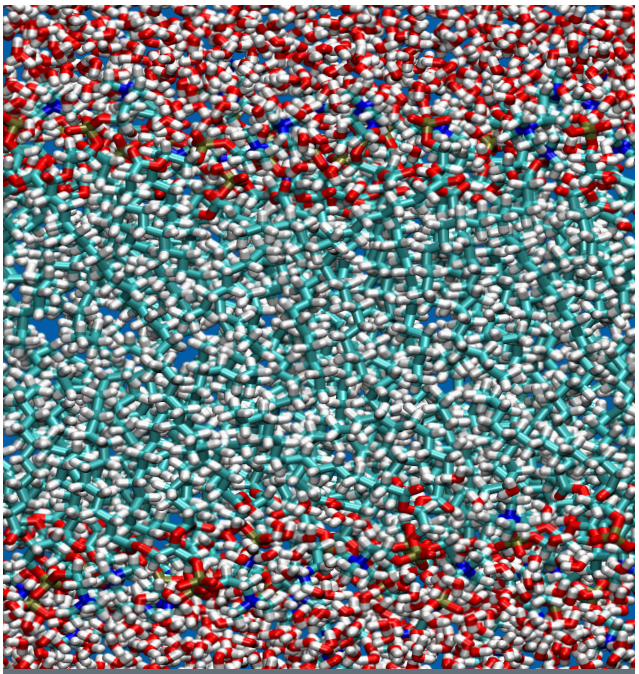

Typical configuration of a zwitterionic lipid bilayer membrane. A polarizable force field is essential for meaningful simulations of such a system.

**Image Credit**
Janamejaya Chowdhary and Benoît Roux, The University of Chicago; Wei Jiang and Yun Luo, Argonne National Laboratory; Alex MacKerell, University of Maryland

**Andrew Binkowski**
Argonne National Laboratory | abinkowski@anl.gov

# Protein-Ligand Interaction Simulations and Analysis

*A fundamental goal of research in rational drug design is to predict computationally whether a given molecule will bind to a protein 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. ALCF resources allow researchers to greatly extend their studies of molecular docking problems for applications in structural genomics and drug discovery.*

To design effective disease-fighting pharmaceuticals, scientists must first understand how a small molecule interacts with a given protein target. This process is the basis of how disease-fighting drugs work; drugs are essentially small molecules that bind to a pathogen and alter or disrupt the enzymes required for it to proliferate. However, the success rate for simulating protein-ligand docking has been historically mixed. Many of the shortcomings point to the scoring functions and the approximations and heuristics necessary to make the computer run-time feasible. Powerful Argonne Leadership Computing Facility resources have remove some of these constraints, allowing more advanced physics-based methods to be applied.

A University of Chicago and Argonne National Laboratory research team led by Andrew Binkowski is conducting a comprehensive analysis of protein binding domain and small molecule interactions through an automated system, including receptor analysis, protein-ligand docking, and binding free energy calculations. In addition, the team is performing the first large-scale study of the computationally intensive free energy–perturbation molecular dynamics (FEP/MD) grand canonical Monte Carlo (GCMC) methodology, or FEP/MD-GCMC, for estimating free binding energy. Implementing these methods and making them more accessible to researchers helps to realize more of the promise that molecular simulation holds.

In collaboration with the Center for Structural Genomics of Infectious Diseases, the team is also conducting computer-aided drug discovery of high-value biomedical targets, including human pathogens, bioterrorism agents, and human disease-related proteins. The predicted computational results are experimentally tested in binding assays and X-ray crystallography experiments, allowing for the important validation step necessary to evaluate the predictive power of biomolecular simulations.

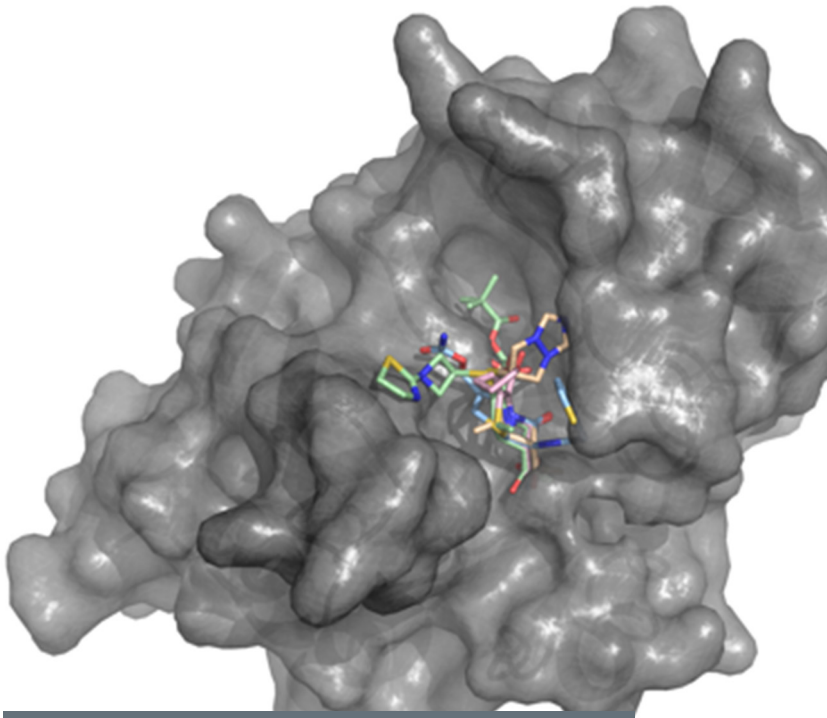

Advanced docking simulations of existing antibiotics (shown in stick form) provide insight on the antibiotic resistance of the NDM-1 enzyme (gray).

**Image Credit**
T. Andrew Binkowski, Argonne National Laboratory/The University of Chicago

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

# Towards Breakthroughs in Protein Structure Calculation and Design

*Through a series of INCITE allocations at the ALCF, University of Washington biophysicist David Baker has extended and validated the capabilities of his landmark structural prediction and design computer program, Rosetta. Now transitioning to modeling the structures of naturally occurring biomolecules and interactions, Baker and his team are also attempting to design new macromolecules with new and useful functions in biomedicine and biotechnology.*

Biomedicine is a modern branch of medical science informed by natural science principles and advanced through studies of biological and physiological processes. Biotechnology encompasses a broader range of applications than medicine, but shares many of the same *de novo* design problems, such as the creation of novel protein folds, binding interfaces, or enzymatic activities. Computation has become an indispensable tool to this discipline's investigations.

Protein structure prediction is key to understanding the function and interactions of biomolecules—the realm where the foundational advances lie. Biophysicist David Baker uses the Argonne Leadership Computing Facility to develop high-resolution protein structure prediction tools to build models of proteins with *atomic-level accuracy* and to computationally engineer both proteins and enzymes with new functions for applications ranging from basic research to therapeutics to bioremediation.

Rosetta, the computational methodology that Baker and his team have been validating through a series of INCITE allocations at the ALCF, will help experimentalists solve structures of biologically important proteins for which experimental X-ray phases are not available or are hard to obtain. In 2012, the team continued to develop the parallelized version of Rosetta, which will allow them to get multifold improvement in performance on ALCF machines, and transition its applications to the new algorithm.

The researchers also intend to test how general the newly developed approach is for the design of protein binders. To that end, they have projects involved in targeting specific patches on proteins useful for diagnostics for infectious diseases, as purification aids for human antibodies, as new diagnostics for malignant transformations, and for laboratory reagents that can target and disrupt oncogenes thought to be crucially important for cancer malignancy.

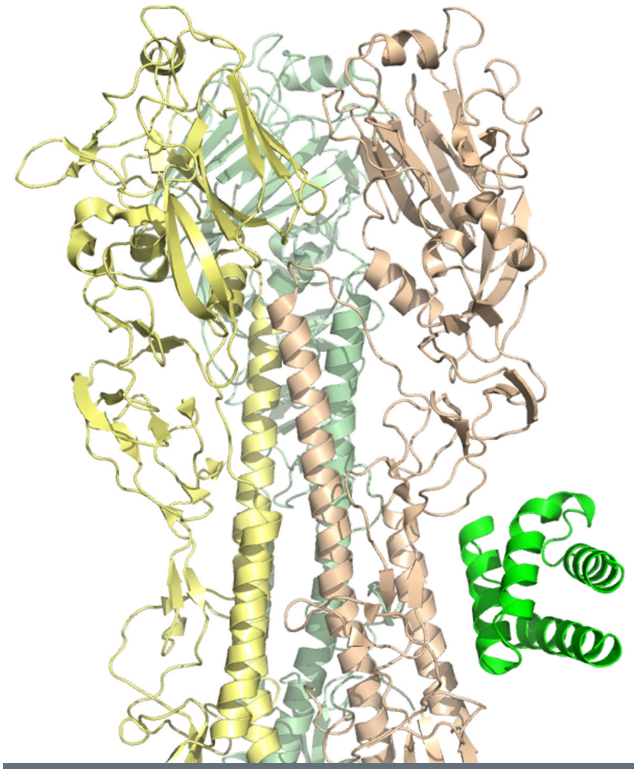

A crystal structure of Spanish influenza hemagglutinin (trimer) bound to a computationally designed binder (green). The computationally designed protein binds hemagglutinin with high affinity and blocks the conformational changes in hemagglutinin that underlie its function, thereby neutralizing influenza. The crystal structure is virtually identical to the designed model, validating the design approach.

**Image Credit**
Sarel Fleishman, Weizmann Institute of Science

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

# Simulations of Deflagration-to-Detonation Transition in Reactive Gases

*Hydrogen is an abundant, environmentally friendly fuel with the potential to reduce our dependence on foreign oil, improve the environment, and boost our economy. Researchers led by Alexei Khokhlov of the University of Chicago are using ALCF supercomputing resources to understand how hydrogen transitions from burning to detonation, thereby furthering efforts to bring hydrogen fuel safely into our everyday lives.*

In reactive gases, the process of transition from slow burning to detonation is called 'deflagration-to-detonation transition,' or DDT. Predicting DDT in various combustion settings is an outstanding combustion theory problem. Led by Alexei Khokhlov, the High Speed Combustion and Detonation (HSCD) project is performing first-principles, reactive flow Navier-Stokes fluid dynamic simulations of DDT at the Argonne Leadership Computing Facility. 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 for use in other industrial applications.

With an allocation of resources made available through the DOE's INCITE program, HSCD has conducted simulations that have reproduced a transition from strong to weak ignition in a hydrogen-oxygen mixture with decreasing Mach numbers. In these simulations, the location of the initiation spot moves away from the end wall of the shock tube. At the same time, the ignition time delay decreases compared to an adiabatic induction time in ideal reflected shock conditions whereas for a strong ignition regime, the ignition time coincides with the adiabatic predictions. These results are fully consistent with experimental observations and represent a significant advancement for the project. The next steps will be to analyze and understand the physics underlying flame acceleration and DDT transition in long tubes.

Shock bifurcation In $CO_2$. Calculations including both viscosity and heat conduction. Left - experiment. Right - simulation. Shock is moving to the left. 1 and 2 - angles in Table 2.
Angle 1: experiment, 36, simulated 37
Angle 2: experiment 38, simulation 34

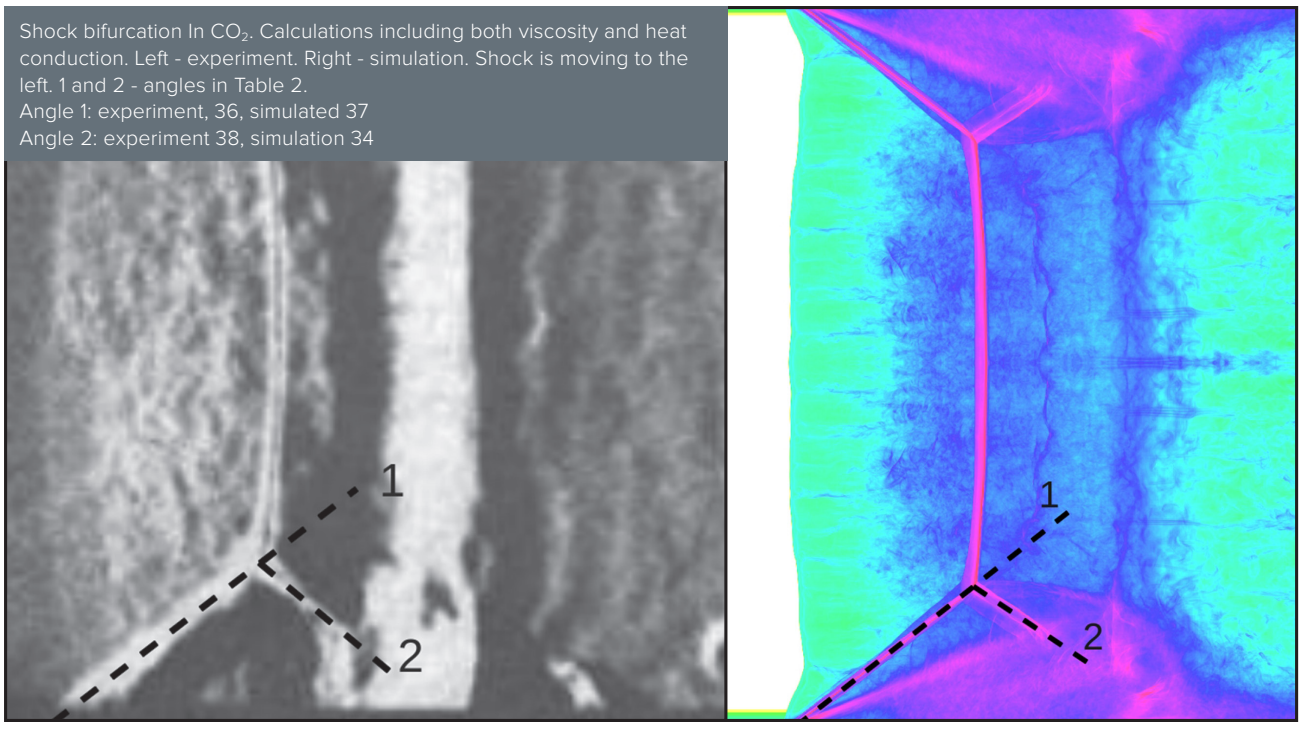


**Image Credit**
Alexei Khokhlov and Ben Clifford, The University of Chicago; Shashi Aithal and Charles Bacon, Argonne National Laboratory; Joanna Austin and Andrew Knisely, University of Illinois at Urbana-Champaign

**Dario Alfè**
University College London | d.alfe@ucl.ac.uk

# Water Systems from Highly Accurate Quantum Monte Carlo Calculations

*Essential to life on earth, water is one of the most familiar substances still not fully understood by modern science. Using the highly accurate electronic structure theory method, Quantum Monte Carlo (QMC), University College London researchers led by Dario Alfè are leveraging ALCF resources to model liquid water with unprecedented accuracy. Their efforts will yield fundamental insights applicable in aqueous chemistry across scientific disciplines.*

Currently, QMC is the only highly accurate method of computational electronic structure theory that can be used to solve the Schrödinger equation to numerical precision for systems bigger than smallish molecules. Its quadratic scaling of computer time with system size is favorable to the $N^7$ scaling of the standard CCSD(T) coupled cluster technique that is widely regarded as quantum chemistry's 'gold standard.' Furthermore, QMC may also be trivially applied to periodic systems with hundreds of atoms per unit-cell.

Dario Alfè's team at University College London has recently added the capability of coupling QMC with Density Functional Theory (DFT) based molecular dynamics within their CASINO program (developed in Cambridge by R.J. Needs, M.D. Towler, N. Drummond and P. Lopez-Rios), allowing them to calculate the accurate QMC energy of a system on a dynamical trajectory generated using the less expensive DFT calculations. This constitutes the team's first step towards the challenging goal of simulating liquid water with QMC, a material notoriously difficult to account for with DFT. QMC will yield highly accurate results for all the occurring interactions, including the ubiquitous hydrogen bonds.

Making use of a Director's Discretionary allocation at the ALCF, the team is currently using Intrepid to conduct benchmarking and testing of various new routines and scripts to implement novel functionality that will also allow for calculation of QMC forces, eventually enabling the first direct QMC simulation of liquid water.

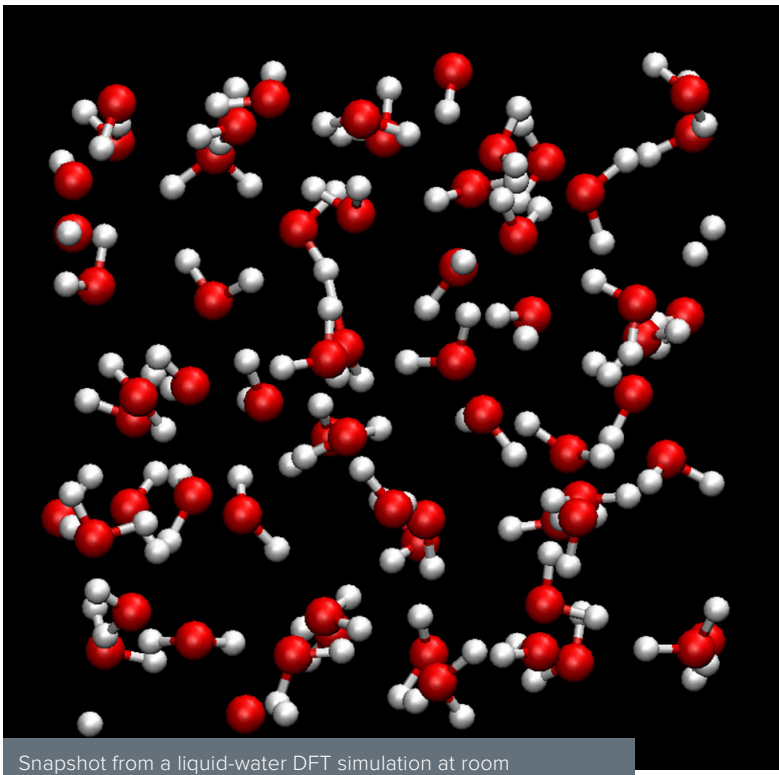

Snapshot from a liquid-water DFT simulation at room temperature. QMC energies on such a system have only become possible recently on leadership-class machines like Intrepid and Mira. This image generated via a density functional theory simulation on the Cray XT6 HECToR in the U.K.

**Image Credit**
Dario Alfè, Michael Gillan, and Michael Towler, University College London

**Ewing Lusk**
Argonne National Laboratory | lusk@mcs.anl.gov

# Scalable System Software for Performance and Productivity

*System software is a critical component of any computing system and forms the infrastructure on which all applications depend. Using supercomputing resources at the ALCF, an Argonne-led research team is improving and extending the capabilities of existing system software to allow applications to benefit from current leadership-class systems.*

As hardware complexity skyrockets in leadership-class systems, it is not easy for applications to take complete advantage of the available system resources and to avoid bottlenecks. This INCITE project aims to improve the performance and productivity of key system software components on these platforms. The research team is studying four different classes of system software: message passing libraries, parallel I/O, data visualization, and operating system.

The collaborators are using Intrepid, the IBM Blue Gene/P platform at the Argonne Leadership Computing Facility (ALCF), to understand and solve problems that occur at scale. Through rigorous experimentation, analysis, and design cycles, the team is improving the capabilities of not only systems being deployed in the near term, but of all systems pushing scalability limits in the near future.

In 2012, the researchers:

- Began developing a framework where users can annotate applications with relevant computation/communication information, allowing an automated program transformation system to modify such code to use different architecture-sensitive communication mechanisms.
- Started working on a compiler support for refactoring MPI applications.
- Evaluated latest I/O optimizations at scale and in real science applications. They used a new technique in their Parallel-NetCDF I/O library to provide double the observed I/O bandwidth to the FLASH astrophysics code without making any changes to the output file format.
- Studied and improved the I/O behavior of the Parallel IDX (PIDX) library at significant fractions of the machine. This work resulted in an SC12 publication.
- Ported the CODES event-driven simulator, developed at Rensselaer Polytechnic Institute, to Intrepid and used CODES to carry out storage hardware analysis. Previously, the CODES simulator ran exclusively on Linux clusters. Using techniques such as “time warping,” the researchers can explore the design space of machines years before they are available to researchers.

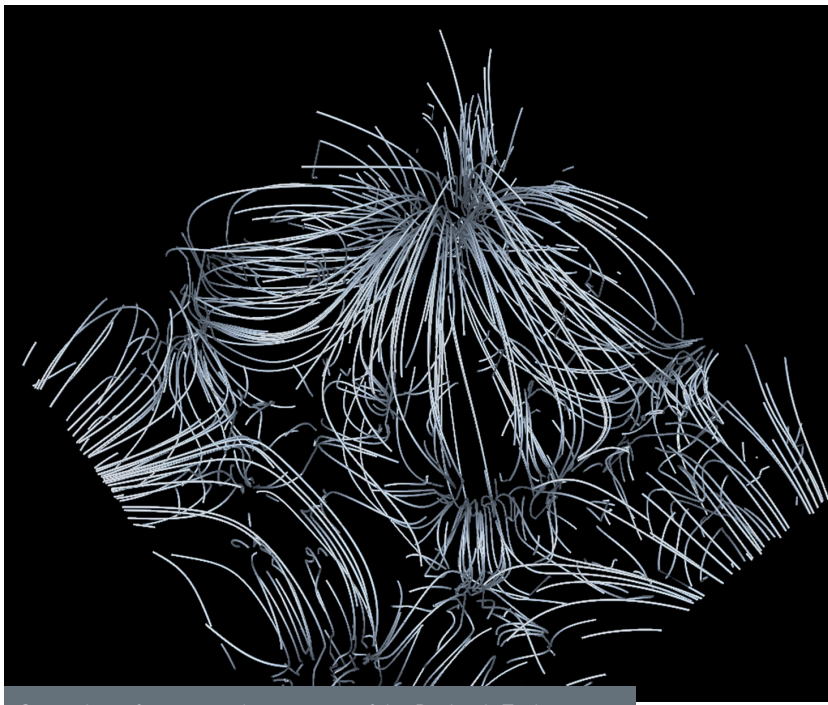

Streamlines from an early time step of the Rayleigh-Taylor instability depend on scalable storage, communication, and data analysis algorithms developed at extreme scale using INCITE resources.

**Image Credit**
Tom Peterka, Argonne National Laboratory

**Greg Holland**
National Center for Atmospheric Research | gholland@ucar.edu

# Simulating Regional Climate at Convection-Permitting Resolution

*Using the world-class supercomputing resources at the ALCF, researchers from the National Center for Atmospheric Research (NCAR) will complete the longest, highest-resolution run of the Nested Regional Climate Model (NRCM) over a large domain. Their findings will be used to create next-generation climate modeling systems that may ultimately aid in national planning for extreme weather events.*

NRCM is a state-of-the-art climate modeling code that combines the strengths of NCAR's Weather Research and Forecasting model and Community Climate System Model to allow for fundamental progress on the understanding and prediction of regional climate variability and change.

Using NRCM, with support from the National Science Foundation, the Research Partnership to Secure Energy for America and the Willis Research Network, this project provides a thorough test and exploratory examination of the advantages of convection-permitting resolution on climate timescales with an emphasis on high-impact weather and climate. In a breakthrough for climate simulations, this project has run NRCM at a horizontal resolution of 4 kilometers for a large domain extending over North America and the Atlantic Ocean basin at the Argonne Leadership Computing Facility. Analysis is now underway on phenomena with high sensitivity to model resolution, including water-snowpack assessments for the mid- and western U.S., and high-impact events such as winter storms and hurricanes. In addition to providing an improved understanding of phenomena, the knowledge gained from this research is already being used in the design of the next-generation regional climate modeling system at NCAR. All data obtained are being made available to the community so that other research and development activities may benefit. These latest successes represent essential steps towards the community's overarching goal to better quantify high-impact weather under climate variability and change. Ongoing efforts will continue to advance our understanding of Earth's climate for national emergency preparedness and will broaden the community of researchers capable of using leadership computing resources.

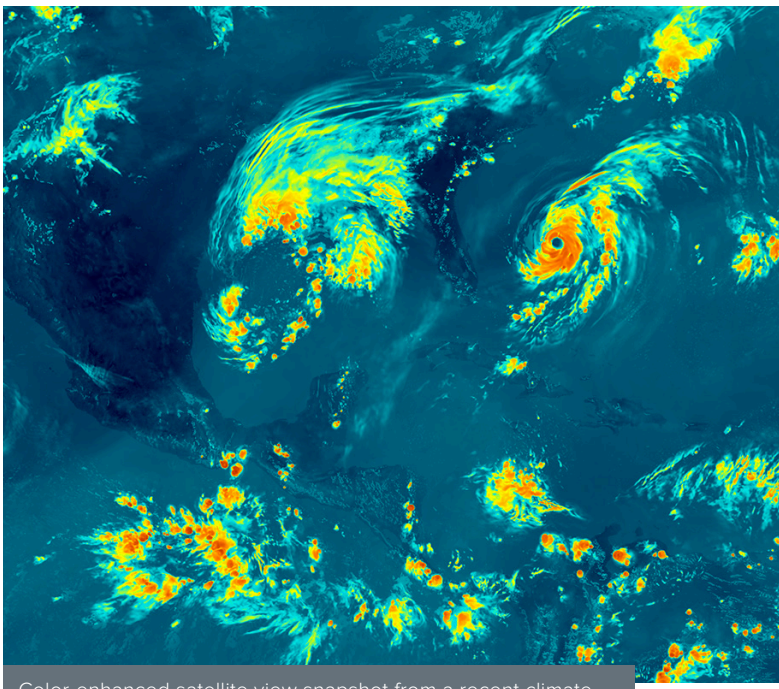

Color-enhanced satellite view snapshot from a recent climate simulation using the Nested Regional Climate Model at 4 kilometers horizontal resolution.

**Image Credit**
James Done, National Center for Atmospheric Research (NCAR) Earth System Laboratory

**Thomas Jordan**
University of Southern California | tjordan@usc.edu

# Using Supercomputers to Improve Seismic Hazard Maps

*Supercomputer simulations from the ALCF help scientists model what goes on inside an earthquake and better understand seismic hazard. Their findings are expected to inform the work of building designers and emergency planners worldwide.*

An interdisciplinary research team led by Thomas Jordan at the Southern California Earthquake Center (SCEC) is designing, conducting, and analyzing high-resolution simulations of earthquakes at the Argonne Leadership Computing Facility. SCEC's scenario simulations of past and expected future quakes use realistic 3-D structural models of Southern California to study the dynamic rupture processes inside an earthquake.

Many physical processes take place during an earthquake. In addition to the static and dynamic friction between fault surfaces, the roughness of the fault generates powerful high-frequency seismic waves. Leveraging the capabilities of the IBM Blue Gene/P supercomputer at the ALCF, the team has created high-resolution dynamic rupture models that are more physically accurate than previous models. These new models allow researchers to study how a quake's rupture processes produce seismic waves and how those waves affect ground motion. The findings have enabled scientists to develop more accurate estimates of ground motion in Southern California and to apply their learnings to earthquakes in other parts of the world. Their research is expected to help emergency management groups, civil engineering groups, tall building engineers, and governmental organizations worldwide gain a better understanding of earthquake hazard and risk.

SCEC researchers have also constructed worst-case scenarios for California that enable them to estimate seismic hazards for rare but possible earthquakes in magnitudes of 8.0 and larger. Their findings promise to help communities prepare for future earthquakes and reduce earthquake risk.

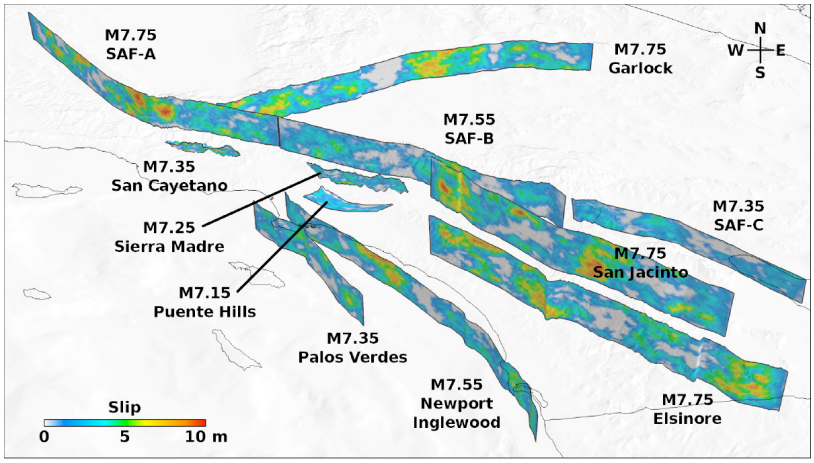


A 3-D view showing potential source faults for Southern California's next "big one." Dynamic rupture and wave propagation simulations produce a model of ground motion at the earth's surface. Colors indicate possible distributions of displacement across the faults during rupture.

*"Using ALCF supercomputers, we are able to simulate earthquakes in Southern California across a much broader range of space-time scales, and these simulations are improving our understanding of earthquake processes worldwide."*
*— Thomas Jordan, Southern California Earthquake Center*

**Image Credit**
Geoffrey Ely, Southern California Earthquake Center

**Umesh Paliath**

GE Global Research | paliath@ge.com

# Delivering "Green" Jet Engines and Wind Turbines

*A GE Global Research team is studying the complex flow of air in jet exhaust nozzles and wind turbine airfoils. The researchers are conducting simulations on the ALCF's supercomputer to understand and predict flow in jet engines and wind turbines. Such information is key to developing quieter, more fuel-efficient wind turbines and jet engines and to improving engine life cycles in an extremely competitive global market.*

GE Global Research's (GE's) recent work on Large Eddy Simulations (LES) has leveraged the petascale computing enabled by the INCITE program to break barriers related to the aero-acoustics of wind turbine and aircraft engines. The GE LES strategy has been to accelerate its industrial impact by pushing application/validation to realistic conditions, addressing fundamental physics and source characterization challenges, and extending capability to handle complex system interactions. Powered by scalability improvements at the Argonne Leadership Computing Facility (ALCF), the previous INCITE work has demonstrated how this first-principles-based LES capability can transform product development. This includes proof-of-concept for airfoil and jet noise computation at realistic Reynolds numbers, design guidance for novel jet noise reduction features, and turnaround capability to replace a physical design/test campaign.

For jet engines, current research includes examining the experimentally hard-to-quantify effect of jet-flap interaction to characterize installation effects on jet noise, and aid in developing low-noise exhaust systems. Simulations are under way to understand the noise generation and radiation in the presence of installation geometry, such as pylon and wing. This will demonstrate how the validated LES solver, coupled with high-performance computing, can be used as a numerical rig to gather high-resolution flow and acoustic data to better understand, effect of nozzle geometry variations on noise. Driven by a recent push into ultra-high, bypass-ratio-ratios engines for increased propulsive efficiency, fundamental work in fan broadband noise has also been initiated.

For wind turbines, GE is working towards demonstrating readiness in guiding low-noise design, validating the complex scaling of turbulent self-noise, and using wall-models to enable large-span blade computations. Simulations are in progress to assess the effect of angle of attack on airfoil trailing edge noise, an important step towards enabling improved trailing edge designs for noise reduction. The research team's goal is to improve the design process for both individual turbines and large turbine farms.

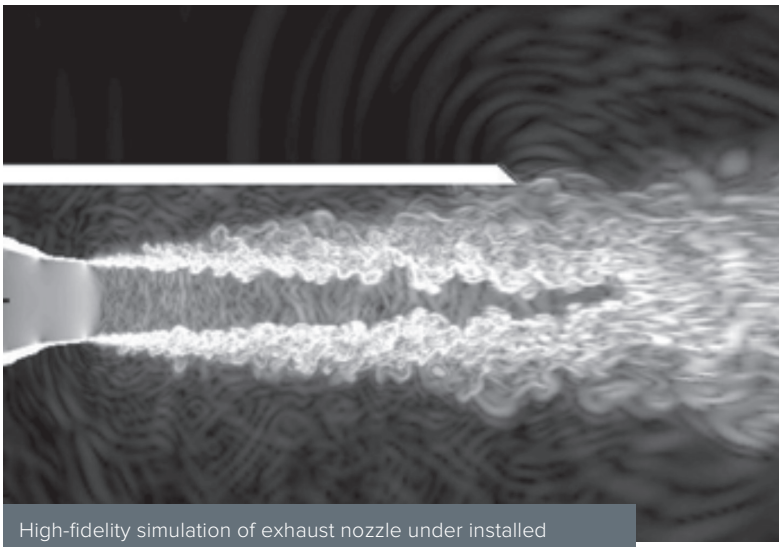

High-fidelity simulation of exhaust nozzle under installed configuration.

*"High-performance computing at the ALCF has enabled us, for the first time, to visualize the complex phenomena of turbulent mixing in jet exhaust flow and flow over wind turbine airfoils in ways not possible through experiments."*
*— Umesh Paliath, GE Global Research*

**Image Credit**

Umesh Paliath, GE Global Research; Joseph Insley, Argonne National Laboratory

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

# Global Simulation of Plasma Microturbulence at the Petascale and Beyond

*As scientists look for alternatives to fossil fuels to meet our energy needs, there is increasing interest in nuclear fusion. Of utmost importance in the design and operation of future fusion power sources is an understanding of turbulent transport losses. Led by William Tang, researchers from Princeton Plasma Physics Laboratory are tapping into the extreme-scale resources of the ALCF to acquire this knowledge.*

Researchers are studying 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 code, a 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 Mira, the Argonne Leadership Computing Facility's next-generation Blue Gene/Q, have yielded significant results:

- GTC parallelism: MPI plus OpenMP
  - On Mira, running with two different combinations: (i) 16 MPI rank per node and 4 threads per MPI rank (this is equivalent to 1 MPI rank per core and 4 threads per core); and (ii) 4 MPI rank per node and 16 threads per MPI rank
  - On Mira, running with 1 MPI rank per core and 4 threads per core
- Scaled efficiently to 524,288 cores—two-thirds of Mira
  - New load-balance method allows running ITER-scale problems using the capability of Mira
- Long-duration simulations of ITER plasmas will demand O($10^8$) grid points and O($10^{10}$) particles

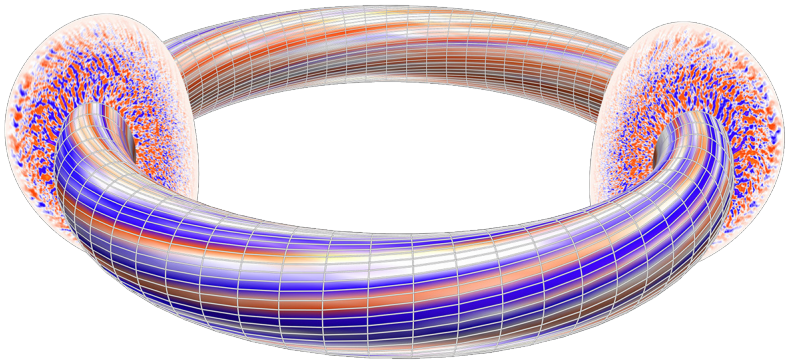

Fully kinetic 3-D plasma microturbulence simulation in a tokamak fusion device of the self-consistent electrostatic potential. The red and blue represent regions of positive and negative potential respectively. Elongated structures in the toroidal direction follow the magnetic field lines and are characteristic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

**Image Credit**
Kwan-Liu Ma, University of California, Davis

**Paul F. Fischer**

Argonne National Laboratory | fischer@mcs.anl.gov

# Making Safe, Clean Nuclear Energy Available Locally

*The United States is committed to new technologies that expand the availability of safe, clean nuclear energy to help meet growing global demand. Using the power of the ALCF's Blue Gene/P supercomputer, researchers are performing large-scale numerical simulations in the quest for an economical source of power.*

Liquid-metal-cooled fast reactors may hold the key to producing safe, clean nuclear energy. These reactors permit the recycling of nuclear fuel and show promise as an economical source of power. A research team led by Paul Fischer and including scientists from Argonne National Laboratory and the University of Illinois carried out large-scale numerical simulations of turbulent thermal transport in sodium-cooled reactor cores. Using the Argonne Leadership Computing Facility's IBM Blue Gene/P supercomputer, researchers simulated wire-wrapped fuel rods with 7-, 19-, 37-, and 217-pin bundles. This simplified geometry allowed them to resolve most of the scales of turbulent motion with minimal modeling assumptions involved in large-eddy simulations (LES).

Their current computations are some of the largest to date with the spectral element code Nek5000, and involve several hundred million grid points on unstructured meshes. Validated high-fidelity Nek5000 results are being used to benchmark steady-state and unsteady low-fidelity Navier-Stokes codes that employ turbulence models and to provide input to reactor design codes that require only coarse (mean flow) data. The team's findings are expected to inform the design of next-generation nuclear reactors capable of providing sustainable energy with a low carbon footprint.

As part of a worldwide validation effort for nuclear simulation codes, the Nuclear Energy Agency of the Organization for Economic Co-operation and Development conducts blind-benchmark studies. In 2010, the team's computational results for velocity and temperature distribution in a T-junction—flow geometry important to understanding thermal-mechanical fatigue in light water reactors—ranked number one in temperature and sixth in velocity predictions. This year, in a Matis benchmark study, the team's submission results ranked sixth in predictions of average velocity and first in rms velocity that characterizes a level of turbulence. Also, the team's submission was the only LES calculation that managed to simulate the full experimental geometry including the outflow plenum.

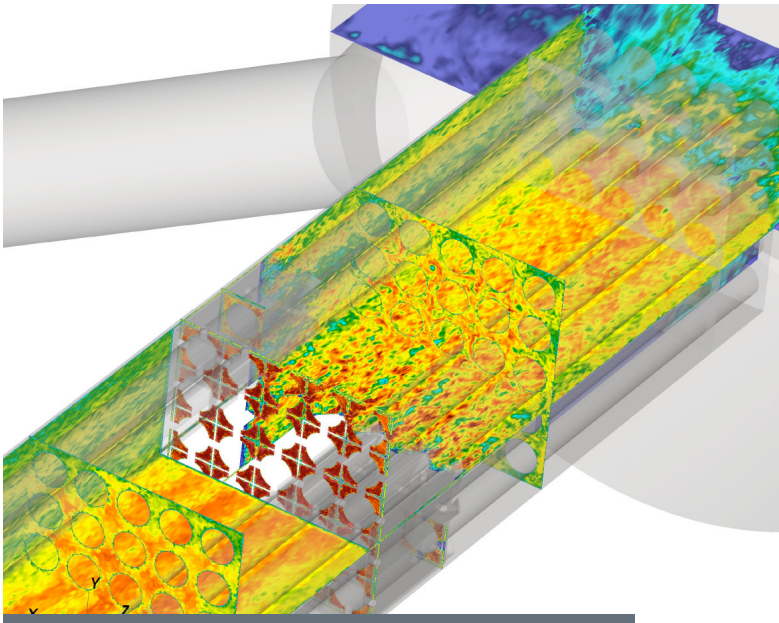

Velocity magnitude distribution in a flow through the 25-pin swirl-vane spacer grid of Matis benchmark. Computed on Intrepid with Nek5000 and visualized on Eureka with VisIt at the ALCF.

*"Advanced simulation is considered essential in bringing new reactor technology to fruition in an economic and timely manner."*
*— Paul Fischer, Argonne National Laboratory*

**Image Credit**

Aleks Obabko, Paul Fischer, and Tim Tautges, Argonne National Laboratory

**Micheal Smith**
Argonne National Laboratory | masmith@anl.gov

# Scalable, Explicit Geometry, Whole-Core Nuclear Reactor Simulations

*Nuclear power is a safe, abundant, clean energy resource capable of meeting our nation's growing needs. Argonne's Micheal Smith is using the resources at the ALCF to accomplish breakthrough enhancements in computer models used to simulate new nuclear reactor designs that will allow for dramatic improvements in the safety, sustainability and economics of nuclear energy systems in the long term.*

The Simulation-Based High-efficiency Advanced Reactor Prototyping (SHARP) project at Argonne aims to improve the ability of nuclear reactor simulation codes to solve real-world design problems. SHARP is supported primarily by DOE's Nuclear Energy Advanced Modeling and Simulation (NEAMS) program with additional support from the Center for Exascale Simulation for Advanced Reactors (CESAR). NEAMS endeavors to develop new tools that push the envelope of nuclear reactor analysis, while CESAR prepares for the future of nuclear reactor research through the co-design of fundamental algorithms and new supercomputers.

Before researchers can justify confidence in a code, it must first undergo substantial verification and validation. Previously, the majority of the validation work on Argonne's neutron transport code, UNIC, was done using well-understood methodologies involving spatial homogenization. This project will investigate the challenges in cross section generation in heterogeneously modeled systems, and will address a request from Idaho National Laboratory to demonstrate the ability of the UNIC code to model the geometrically complex Advanced Test Reactor (ATR). The team will continue refining the cross section generation technique for homogeneous, mixed homogeneous-heterogeneous, and fully heterogeneous geometry modeling in steady-state as well as kinetics/dynamics modeling of fast reactors. Additional efforts at the Argonne Leadership Computing Facility will include initial validation of the heterogeneous modeling capability of UNIC on fast- and thermal-reactor problems using a wider set of benchmark problems, including the Jezebel and Godiva reactors.

Knowledge gained through simulations of nuclear reactors can:

- Extend the life of the nation's current thermal reactor fleet,
- Boost confidence in safety margins, allowing reactors to run at even higher powers with greater efficiencies,
- Reduce nuclear waste through improved recycling technology and through the advancement of fast reactor technology (a highly fuel-efficient reactor type that generates substantially less nuclear waste).

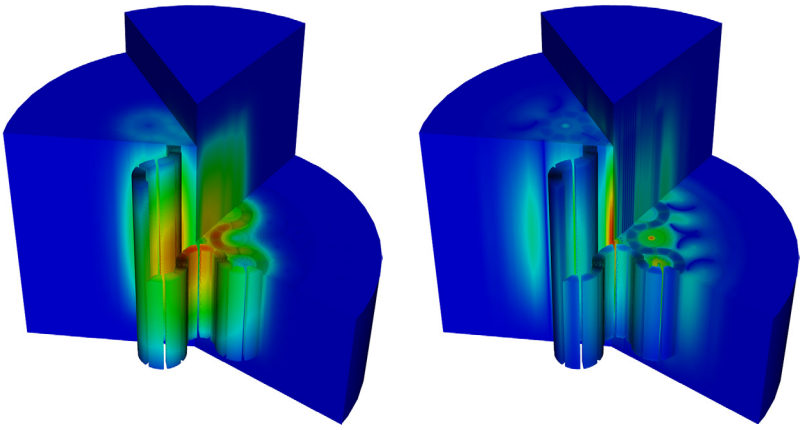

Images depicting the 3-D distribution of neutron flux in the ATR for both fast (left) and thermal (right) energies calculated by the UNC-SN2ND code on Intrepid, the ALCF's Blue Gene/P.

**Image Credit**
Micheal A. Smith and Emily R. Wolters, Argonne National Laboratory

**Parviz Moin**
Center for Turbulence Research | moin@stanford.edu

# Prediction of Multiscale, Multiphysics Turbulent Flow Phenomena Using Unstructured Large Eddy Simulation

*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. Using ALCF resources, Parviz Moin continues his research for supersonic jet noise crackle in complex geometry nozzles and has expanded his research to include supersonic combustion in mixing layers.*

One of the major problems confronting the aircraft industry is the aerodynamic noise generated by engine exhaust jets and airframes, particularly during take-off and landing approaches.

One way to make jets less noisy is to modify the shapes of the nozzles. Pointed cutouts, called chevrons, are known to reduce noise when added to the engine exhaust nozzles, as seen on Boeing's new 787. The physical reasons for this are not well understood. For high exhaust jet velocities, typical of supersonic aircraft, jet mixing, shocks and crackle account for most noise sources. Past Argonne Leadership Computing Facility allocations enabled a comprehensive study of the physical mechanisms by which chevrons affect the jet mixing and shock-associated noise. New simulations are underway to determine how chevrons may reduce crackle noise, which involves highly intermittent, non-linear pressure waves that are very loud.

Another way to make an aircraft less noisy is to manage or control the airframe noise during landing. As slats and flaps on aircraft wings adjust for a landing approach, they create extra noise. As with engine chevrons, different shapes are added to the wing to make the flow "more smooth" and reduce the noise. The best position where these control elements should be placed is, however, not always obvious. Moin's team is investigating the optimal sensor placement strategies to control the noise and drag for multi-element airfoils using ALCF's resources.

New simulations of combustion instabilities are also being examined this year. To capture and ultimately control these complex geometry, multi-scale and multi physics turbulent flows, scalable computational tools for large-eddy simulation with minimal numerical dissipation and dispersion on unstructured grids are 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 ALCF.

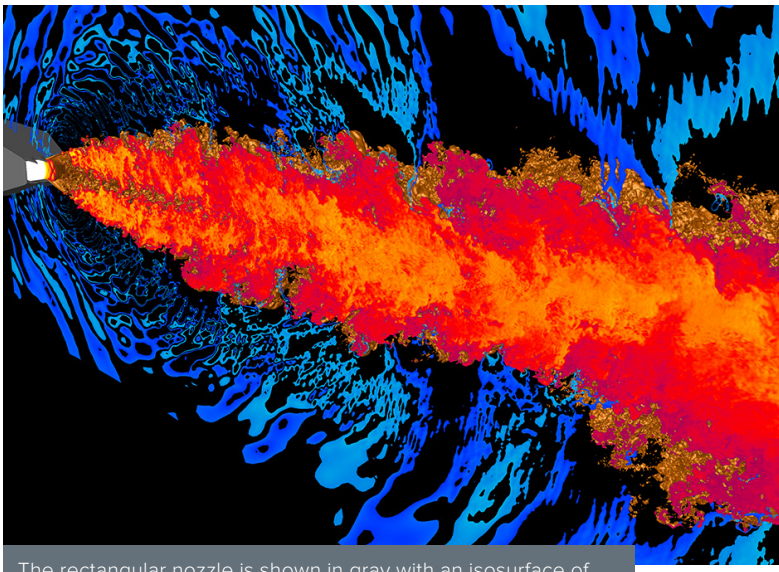

The rectangular nozzle is shown in gray with an isosurface of temperature (gold) cut along the center plane of the nozzle showing temperature contours (red/yellow). The acoustic field is visualized by (blue/cyan) contours of the pressure field taken along the same plane. The chevrons enhance turbulent mixing just downstream of the nozzle exit shaping how the jet spreads downstream. This significantly reduces the noise produced by the supersonic jet compared to a rectangular nozzle without chevrons (not shown).

**Image Credit**
Joseph Nichols, Center for Turbulence Research

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

# Reducing Drag and Energy Dissipation in Transportation

*Transportation accounts for 28% of U.S. energy consumption. Researchers from the University of Texas at Austin are performing a direct numerical simulation of turbulent channel flow at high Reynolds numbers on an ALCF supercomputer that will lead to advances in understanding the physics of these flows and contribute to the improved efficiency of transportation systems.*

Central to energy losses inherent in transportation is turbulence in the boundary layer formed when a vehicle moves through air or water or when fluid moves through a pipe. The energy expenditure results from the interaction between the solid surfaces and the fluid flowing past them, leading to drag and the dissipation of energy by turbulence. This drag is caused by friction and pressure at the vehicle's surface that are enhanced when the boundary layers separate, slowing the vehicle and consuming energy.

Engineering developments to reduce drag and energy consumption are greatly impeded by the lack of accurate models of the turbulence phenomena involved. Direct numerical simulation (DNS) at high Reynolds numbers and the subsequent analysis of the resulting data can provide the insights needed to develop such models, as well as new concepts for manipulating wall-bounded turbulence.

A research team from the University of Texas at Austin is focusing on reaching sufficiently high Reynolds numbers to explore the physics that arise in the overlap region. The overlap region is where the viscous near-wall turbulence interacts with the outer-layer turbulences. This interaction is key to understanding high Reynolds number turbulent wall layers. To investigate this interaction, the Reynolds number needs to be sufficiently high so that there is a substantial disparity in scale between the inner and outer layers. The results can then be extrapolated to arbitrary Reynolds numbers. This simulation is being conducted using the supercomputing software that the team has developed and benchmarked on Intrepid, the IBM Blue Gene/P supercomputer at the Argonne Leadership Computing Facility (ALCF), and on Mira, the next-generation Blue Gene/Q system.

To date, the researchers have completed the highest-ever Reynolds number simulations for spatially evolving incompressible turbulent boundary layers using DNS. The simulations show excellent agreement with experimental data sets in literature and other simulations. The researchers found that turbulent boundary layer fluctuations are higher than those in channel flows.

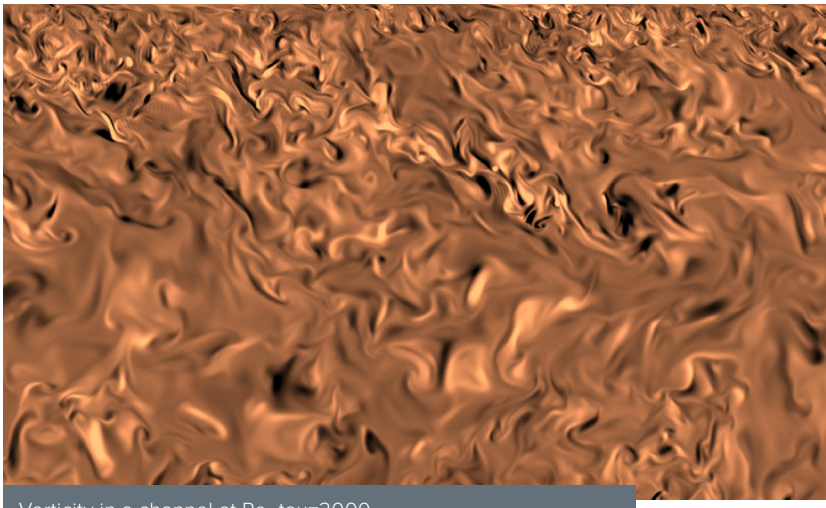

Vorticity in a channel at Re_tau=2000.

*"Simulations on the most advanced supercomputers enable turbulence research by exposing the structure of turbulence in unprecedented detail. Supercomputer simulations of turbulence like those conducted at the ALCF are a game changer for scientists. With them, progress is not limited by what can be measured, only by the insight of the investigator."*
*— Robert Moser, University of Texas at Austin*

**Image Credit**
Sergio Hoyas and Oscar Flores, Universidad Politécnica de Madrid

**Sanjiva K. Lele**
Stanford University | lele@stanford.edu

# Studying the Interaction between Shock and Turbulence

*How shock waves and turbulence interact is important in studying explosions in supernovae as well as in designing propulsion systems for aircraft. In an unprecedented study, researchers are leveraging the computational power of the ALCF's Blue Gene/P to simulate those phenomena, expanding our understanding of their relevant physics and developing better models for engineering applications that include these flow features.*

Richtmyer-Meshkov instability (RMI) occurs when a shock wave interacts with a perturbed interface, separating fluids of different densities. After the shock refracts through the interface, perturbations grow and the instability can evolve into a turbulent mixing region. While RMI occurs in a wide range of flows, the turbulent mixing is not well understood. Until recently, little research had been conducted on RMI, chiefly because of the inadequacy of numerical algorithms and high computational costs.

A team led by Sanjiva Lele of Stanford University is using the IBM Blue Gene/P supercomputer at the Argonne Leadership Computing Facility to create a set of canonical shock/turbulence interactions at unprecedented fidelity with sufficient resolution to capture the smallest turbulent eddies. Researchers analyzed the datasets to explain how the shock wave and turbulence affect each other. They found that post-shock turbulence rapidly becomes isotropic at the smallest scales, but the largest eddies remain anisotropic for a long distance. This finding contradicts linear theory and shows that nonlinear effects are important in the post-shock evolution of the turbulence.

In contrast, the amplification of turbulence kinetic energy during the interaction is consistent with linear theory. The team also found that the instantaneous shape of the shock wave is highly distorted for strong incoming turbulence, and that turbulence can even "punch holes" in the shock wave where compression is smooth.

This research enables direct numerical simulations of the turbulent multi-material mixing generated after RMI. Capturing these turbulence statistics, researchers will be able to answer fundamental questions about fluid mechanics.

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

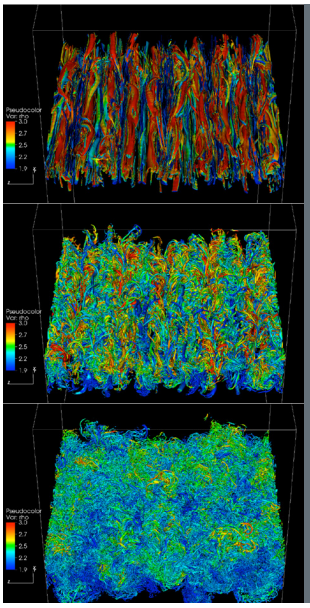

Simulation of Richtmyer-Meshkov instability between air and a curtain of SF6: density iso-contours at three time instants after re-shock impact show how the two different density fluids (red and blue) mix rapidly with each other in the turbulent regime after the re-shock.

*"Using the power of the Blue Gene/P, we're able to simulate interactions between shock waves and turbulence, and find answers to some fundamental physics questions."*
*— Sanjiva Lele, Stanford University*

**Image Credit**
Santhosh K. Shankar and Sanjiva K. Lele, Stanford University

**Jeff Greeley**
Argonne National Laboratory | jgreeley@anl.gov

# Better Catalytic System Designs through Nanoscale Research

*Carbon monoxide is a toxic gas produced in car exhaust and in various industrial processes. Catalytic agents, including precious metals, can aid the conversion of carbon monoxide to carbon dioxide. To improve emissions control technology, Jeffrey Greeley of Argonne National Laboratory is using ALCF's world-class resources to better understand how catalysts behave at their most fundamental level—the nanoscale.*

Nanoparticles of precious metals such as gold, platinum, and rhodium are catalysts of interest for the conversion of carbon monoxide to carbon dioxide. Most catalysts used in industry consist of nanoparticles, typically 3-20 nm in diameter, of the catalytic material.

Quantum mechanical calculations of adsorption energies and charge transfer effects are indicators of catalytic properties. With computational resources made available through the DOE's INCITE program, Jeff Greeley of Argonne National Laboratory leads a team, including researchers from the Technical University of Denmark and Stanford University at the Argonne Leadership Computing Facility in the study of the catalytic properties of nanoparticles with over 10,000 valence electrons at the quantum mechanical level of theory.

Nanoparticles interact differently from bulk materials. Determining the size of the particles—the number of atoms—to use in catalysts for optimal results is difficult to do with physical experiments. Large-scale computer simulation enables the study of the interactions, and therefore the catalytic properties, for many configurations. A number of factors were analyzed in this study; with charge density plots being just one of many. The adsorption energies are the clearest indicator of finite-size effects.

This project has calculated the properties of carbon monoxide adsorbed on gold nanoparticles for sizes up to 1,415 atoms. A key finding is that nanoparticles smaller than 561 atoms show clear finite-size (as opposed to bulk) effects. This information determines the extent to which the carbon monoxide molecules can be catalyzed by gold nanoparticles of various sizes.

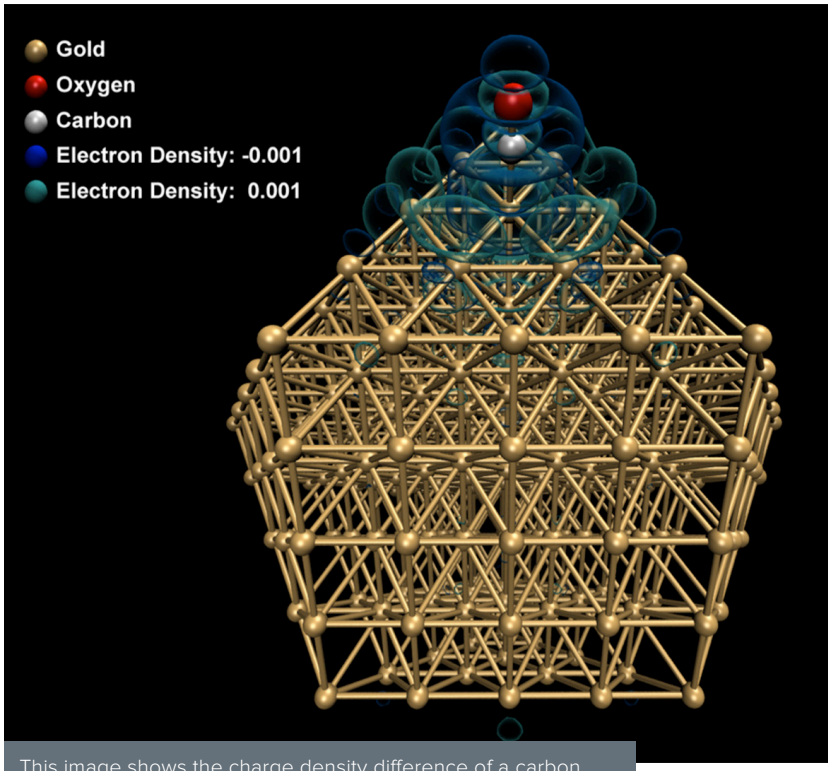


This image shows the charge density difference of a carbon monoxide (CO) molecule adsorbed on a gold nanoparticle of 309 atoms. Quantum mechanical calculations of the adsorption of CO on a series of gold nanoparticles ranging from 13 to 1,415 atoms revealed that nanoparticles larger than 561 atoms (2.7 nm) behave like bulk materials instead of large molecules. These results help to determine the extent to which CO molecules can be catalyzed by gold nanoparticles of various sizes. The blue translucent isosurface depicts charge density depletion. The green translucent isosurface depicts charge density accumulation.

**Image Credit**
Jeffrey Greeley, Vitali Morozov, Nichols A. Romero, and Joseph Insley, Argonne National Laboratory; Jesper Kleis, Ask Larsen, Jens Jørgen Mortensen, Christian Glinsvaad, and Karsten Jacobsen, Technical University of Denmark; Frank Abild-Pederson, Lin Li, and Jens K. Norskov, Stanford University

**William George**

National Institute of Standards and Technology | william.george@nist.gov

# High-Fidelity Simulation of Complex Suspension Flow for Practical Rheometry

*Highly versatile and stone-like in strength, concrete is the most widely used fabricated building material. Accessing the resources of the ALCF through an INCITE award, William George with the National Institute of Standards and Technology (NIST) is researching the flow of concrete to improve its workability and to reduce the impact of concrete production on the environment.*

Over 7.5 billion cubic meters of concrete are made each year, powering a $100 billion U.S. industry. Given its predominance, there is broad interest in making concrete a more sustainable material by finding new ways to recycle it, and by reducing the amount of greenhouse gas created during its production. Architects also continue to push for enhanced workability that will allow concrete to flow easily into increasingly intricate forms. These and other improvements require that scientists first find a way to accurately and reliably measure the viscosity of concrete.

Modeling the flow of concrete presents many challenges. Concrete is a dense suspension of aggregates (e.g., sand, gravel) in a non-Newtonian fluid matrix with aggregate size ranges of several orders of magnitude. Significant effects from wide variations in the shape of the aggregates cannot be accounted for by modeling them as idealized spheres. As concrete undergoes shear flow, microstructures formed by the aggregates dynamically change as they are built and destroyed over time. The geometries of rheometers and the characteristics of the suspensions do not allow analytical solutions to relate torque and angular velocity to fundamental rheological parameters.

Results from this study at the Argonne Leadership Computing Facility will advance the science of dense suspensions and enable the measurement science needed for rheometer design with applicability to industries ranging from concrete, food processing and water treatment to coatings and pharmaceuticals.

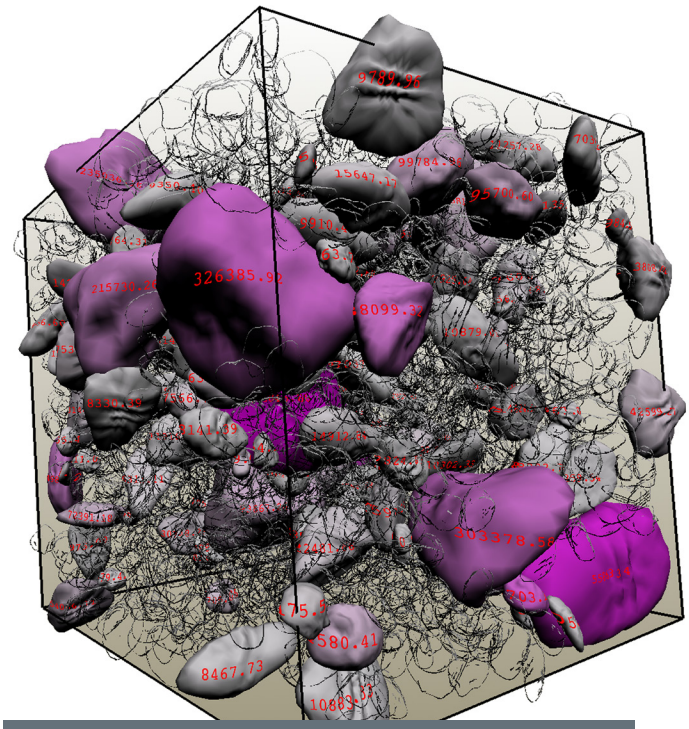


A visualization of the flow of concrete, a complex suspension. In this snapshot of the simulation, the stress on each suspended particle is shown color-coded with its specific value drawn on its surface. Suspended particles that have a stress value below a specific threshold value are shown in outline form in order to better view those particles that are carrying the majority of the stress in the system.

*"You need massive amounts of memory for these simulations, and a machine that's powerful enough to compute the forces at work in the mix, which is why we run at the ALCF where there are resources to match our needs." — William George, National Institute of Standards and Technology*

**Image Credit**

Steven Satterfield, John Hagedorn, and John Kelso, National Institute of Standards and Technology; Marc Olano, National Institute of Standards and Technology/University of Maryland-Baltimore County

**Larry Curtiss**
Argonne National Laboratory | curtiss@anl.gov

# Materials Design and Discovery: Catalysis and Energy Storage

*Pairing the power of the ALCF's new supercomputer, Mira, with newly available electronic structure codes, Larry Curtiss of Argonne's Center for Nanoscale Materials and the Materials Science Division is leading breakthrough research in novel materials with energy-related applications both to reduce our nation's dependence on fossil fuels and to mitigate their environmental impact.*

Our nation's energy future hinges on the design and discovery of new materials—materials that could one day be used to power electric vehicles and to replace the oils currently used in the production of plastics.

To further national efforts to develop sustainable alternative energy solutions, Curtiss's team is conducting field-leading research at the Argonne Leadership Computing Facility (ALCF) into catalytic materials 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. Findings may also lead to new approaches for replacing petrochemicals with inexpensive, abundant small alkanes.

In their efforts, researchers will employ high-accuracy, massively parallel quantum mechanical calculations using Quantum Monte Carlo. Based on its potential to deliver breakthrough science using petascale applications, this project was selected to be one of the ALCF's 16 Early Science Program projects given first access to Mira's resources. In preparation for Mira's availability, the group is collaborating with ALCF staff and IBM in intensive efforts to adapt their software to exploit the new Blue Gene/Q architecture. Computational scientists at the ALCF have successfully ported QMCPACK to Mira, and have already doubled its speed in early optimization efforts. Next steps include developing a mixed-precision version of the main algorithm to make it possible to simulate catalytic nanoparticles with twice as many electrons up to four times faster. Curtiss's team will also use CPMD code, from IBM Zurich, to perform Density Functional Theory calculations. IBM has optimized CPMD for Mira where it has shown to scale to millions of threads.

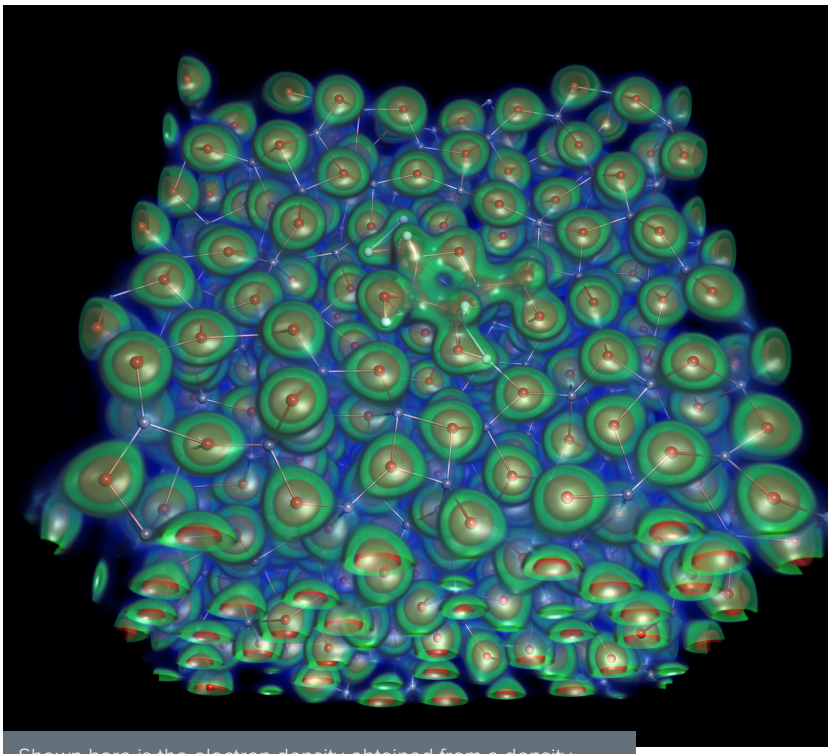

Shown here is the electron density obtained from a density functional theory (DFT) calculation of lithium oxide ($Li_2O$) performed with the GPAW code. This visualization was the result of a simulation run on Intrepid, a supercomputer at the ALCF.

**Image Credit**
Kah Chun Lau, Aaron Knoll, and Larry Curtiss, Argonne National Laboratory

**Priya Vashishta and Aiichiro Nakano**

University of Southern California | priyav@usc.edu and anakano@usc.edu

# Petascale Simulations of Stress Corrosion Cracking

*The performance and lifetime of materials widely used in energy and nuclear technologies are often severely limited by corrosion under stress loads. Simulations performed at the ALCF are revealing the atomistic mechanisms that control stress-induced corrosion within nuclear reactors—which is key to understanding the phenomenon, and ultimately, to developing new technologies to prevent it.*

Understanding the fundamental mechanisms behind stress corrosion cracking—the formation of tiny cracks in highly stressed regions of otherwise corrosion resistant structures—will aid in the design of safe, next-generation nuclear reactors and nuclear waste containers.

A University of Southern California research team that includes Priya Vashishta, Rajiv Kalia, and Aiichiro Nakano are using Argonne Leadership Computing Facility's (ALCF) Intrepid supercomputer to study the complex electrochemical processes leading to the degradation of nickel-based alloys and silicate glass. The team runs molecular dynamics (MD) simulations to learn how a small amount of impurities within these materials can cause fractures and potentially catastrophic failures.

The team has been investigating nanobubble collapse in water under shock and the damage it causes to a silica wall. This research requires huge computational resources to simulate the atoms and molecules interacting over a time span that matches the chemical kinetics of the natural damage mechanisms. Furthermore, the extent of damage depends on the shock speed, so the team will simulate several particle velocities to establish the lowest velocities at which the damage is minimal or nonexistent, and the upward boundary at which the damage is catastrophic. In 2011, the team used the full Intrepid machine to run a one-billion-atom reactive MD simulation at a particle velocity of 2.0 km/s.

In 2012, the team examined the ionic flow characteristics of water inside silica nanochannels. It ran small-scale calculations on a baseline water system of approximately 230,000 atoms, and on a twenty angstrom-wide silica nanochannel with water and sodium and chlorine ions.

The team aims to exploit the capabilities of the ALCF's Mira supercomputer to attempt simulations that encompass several billion atoms for nanosecond timescales, and several million atoms for microseconds.

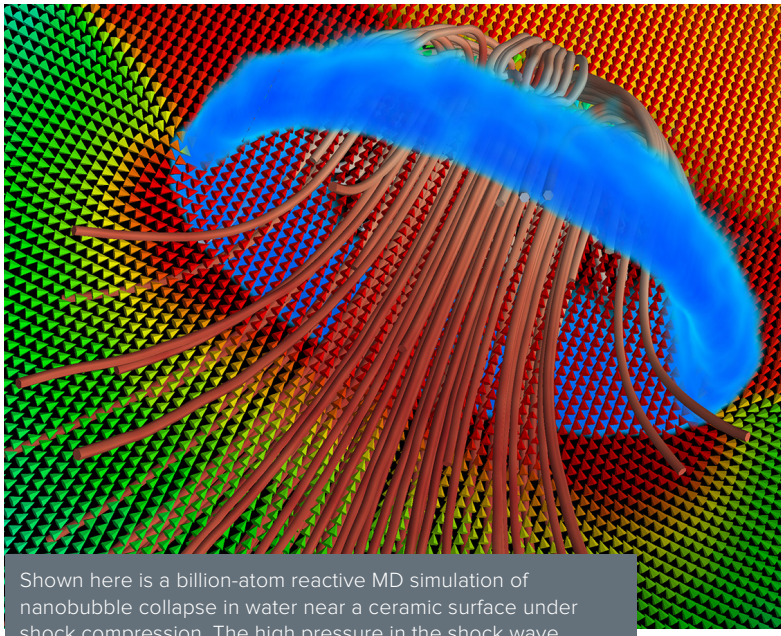

Shown here is a billion-atom reactive MD simulation of nanobubble collapse in water near a ceramic surface under shock compression. The high pressure in the shock wave deforms the ceramic surface and also accelerates water molecules from the bubble's periphery inward. These high velocity water molecules bunch up to form a nanojet, whose impact damages the ceramic surface. These simulations 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 Intrepid, the ALCF's IBM Blue Gene/P.

**Image Credit**

Adarsh Shekhar, Ken-ichi Nomura, Rajiv K. Kalia, Aiichiro Nakano, and Priya Vashishta, University of Southern California

**James R. Chelikowsky**
University of Texas at Austin | jrc@ices.utexas.edu

# Toward Crystal Engineering from First Principles

*This crystal engineering project aims at a bottom-up approach to designing new crystalline materials from molecular building blocks with applications ranging from drug design to hydrogen storage. By leveraging the ALCF's computational resources, the project team developed a highly accurate and efficient method for describing the structure and energetics of molecular crystals and applied it to several case studies.*

A deeper understanding of the dispersive interactions that govern the structure and properties of molecular crystals will enable computational crystal engineering from first principles. For this purpose, University of Texas at Austin researchers Noa Marom and James Chelikowsky, in collaboration with Alexandre Tkatchenko from the Fritz Haber Institute of the Max Planck Society in Berlin, are employing density functional theory (DFT) in conjunction with newly developed methods for describing dispersion interactions. Their research focuses primarily on demonstrating the capability of their approach for a series of case studies reflecting the wide variety of applications of crystal engineering from biological systems, such as amino acids and antimalarial drugs, to technological applications, such as organic semiconductors for organic electronics and dye-sensitized $TiO_2$ clusters for solar cells. The systems being studied comprise several hundred atoms, pushing the size limits of fully quantum mechanical calculations. Potential energy surface (PES) exploration for such systems requires massively parallel computing.

To date, the accurate and efficient many-body dispersion (MBD) method has been developed and applied to the prediction of the structure, energetics, and electronic properties of molecular crystals. The MBD method is a twofold improvement over the Tkatchenko-Scheffler (TS) dispersion method. First, depolarization effects arising from the long-range electrostatic screening in the crystalline medium are accounted for via the self-consistent screening equation of classical electrodynamics. Second, the MBD energy is calculated to infinite order by solving the Schrodinger equation for a system of coupled fluctuating dipoles. The MBD contribution is crucial for locating the PES minima for molecular crystals. In the figure this is illustrated for PES projections on to the a-b plane of the γ-Glycine crystal.

The Argonne Leadership Computing Facility team helped to write scripts for exploring the PES of molecular crystals using the MBD method, and in adapting evolutionary optimization codes to the IBM Blue Gene/P platform.

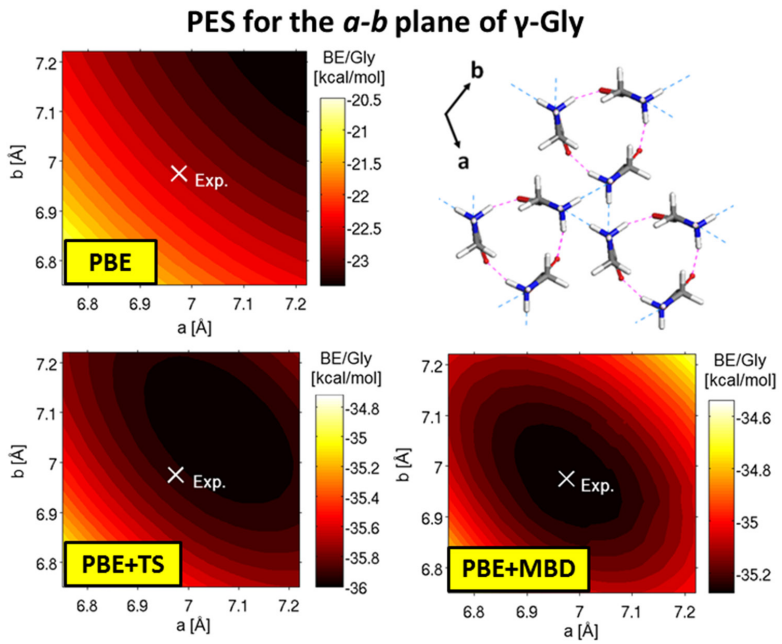


The potential energy surface projected on to the a-b plane of the gamma-glycine crystal (illustrated on the top right) calculated using DFT without any dispersion method (top left), DFT with the TS dispersion method (bottom left), and DFT with the MBD dispersion method (bottom right). Only the latter is in agreement with experiment.

**Image Credit**
Noa Marom and James R. Chelikowsky, University of Texas at Austin; Viktor Atalla, Sergey Levchenko, and Alexandre Tkatchenko, Fritz Haber Institute of the Max Planck Society; Leslie Leiserowitz, Weizmann Institute of Science; Robert A. DiStasio, Jr., Princeton University

**Giulia Galli**

University of California, Davis | gagalli@ucdavis.edu

# Vibrational Spectroscopy of Liquid Mixtures and Solid/Liquid Interfaces: Improving the Predictive Power of Aqueous Environments

*Understanding pure water is an essential prerequisite for grasping the behavior of the liquid interacting with surfaces and solutes. Findings from this INCITE project conducted on ALCF supercomputing resources suggest that, through the use of high-performance computing, scientists can significantly improve their predictive power of aqueous environments.*

A research team led by Giulia Galli from the University of California-Davis is using first-principles molecular dynamics (MD) simulations conducted on Intrepid, the supercomputer at the Argonne Leadership Computing Facility, to compute vibrational properties of liquid mixtures and liquid/solid interfaces. The team's goal is to compare the properties directly with experimental data on specific systems (e.g., water/salt mixtures and water in contact with oxide and graphitic-like surfaces). These calculations are playing a key role in understanding the structure and dissociation of water in contact with solid hydrophobic and hydrophilic surfaces. The research is enabling the team to separate the effects responsible for observed phenomena that can't be easily probed experimentally, such as interface structure dependence on vibrational properties of liquids and liquid mixtures.

The researchers carried out *ab initio* MD calculations using the Qbox code. They determined the structural and vibrational properties of liquid water for several density functionals, e.g., non-local van der Waals functionals and the PBE0 functional (including a portion of exact exchange). Work is in progress to compare their computed vibrational properties with ongoing neutron scattering experiments at the Spallation Neutron Source at Oak Ridge National Laboratory. The researchers also studied dissolution of salts and sulfuric acid in water and unraveled the structural properties of the respective solvation shells. In addition, they analyzed the electronic properties of the solutions. Work is under way to study the structural and electronic properties of the interface between these solutions and oxide surfaces.

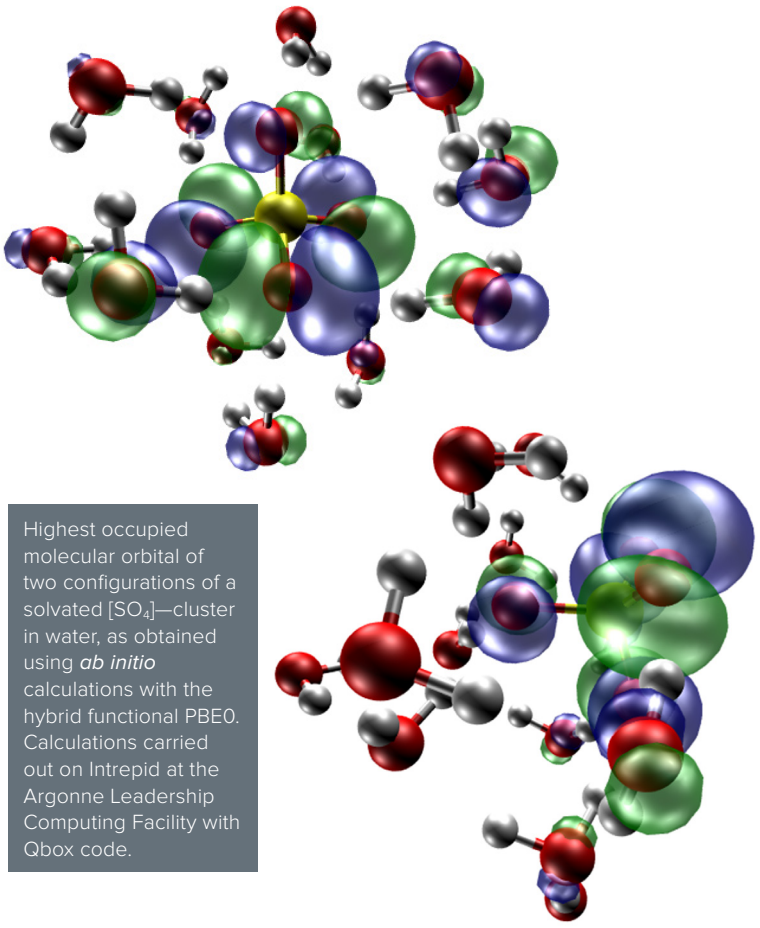

Highest occupied molecular orbital of two configurations of a solvated $[SO_4]$—cluster in water, as obtained using *ab initio* calculations with the hybrid functional PBE0. Calculations carried out on Intrepid at the Argonne Leadership Computing Facility with Qbox code.

**Image Credit**

Giulia Galli, Francois Gygi, Leonardo Spanu, and Quan Wan, University of California, Davis

**Salman Habib**
Argonne National Laboratory | habib@anl.gov

# Argonne Scientists Probe the Cosmic Structure of the Dark Universe

*In the standard model of cosmology, dark energy and dark matter together account for 95% of the mass energy of the universe; however, their ultimate origin remains a mystery. The ALCF will allocate significant supercomputing resources towards unraveling one of the key puzzles—the nature of the dark energy causing the universe to accelerate its current expansion rate.*

Scientists have long been fascinated by the structure of the Universe, hoping ultimately to be able to reveal its fundamental origins. As part of this quest, researchers are now planning the most detailed, largest-scale simulation of structure formation in the universe to date. The simulation targets an improved understanding of the nature of the "Dark Universe."

The new computational cosmology project is led by physicists Salman Habib and Katrin Heitmann of Argonne National Laboratory's High Energy Physics (HEP) and Mathematics and Computer Science (MCS) divisions. The project team will conduct some of the largest high-resolution simulations of the distribution of matter in the universe. They will resolve galaxy-scale mass concentrations over observational volumes representative of state-of-the-art sky surveys by using Mira, a petascale supercomputer at the Argonne Leadership Computing Facility. A key aspect of the project involves developing a major simulation suite covering approximately 100 different cosmologies—an essential resource for interpreting next-generation observations. This initiative targets an improvement of approximately two to three orders of magnitude over currently available resources.

The Argonne researchers will use a new framework they have developed called HACC, short for Hardware/Hybrid Accelerated Cosmology Codes. HACC is designed for extreme performance, but with great flexibility in mind, combining a variety of algorithms and programming models that make it easily adaptable to different platforms.

A paper on the team's research, "The Universe at Extreme Scale: Multi-Petaflop Sky Simulation on the BG/Q," earned a finalist designation for the ACM Gordon Bell Prize and will be included in the Gordon Bell Prize sessions at SC12.

These powerful computational tools will enable the team's ultimate goal to compare the best telescope observations of structure in the universe to the structure found in the numerical simulations, thereby testing various cosmological models and possible explanations for dark energy. In turn, insights gained from the simulations will suggest new theoretical approaches, as well as new methods for designing and cross-correlating observational probes of structure formation.

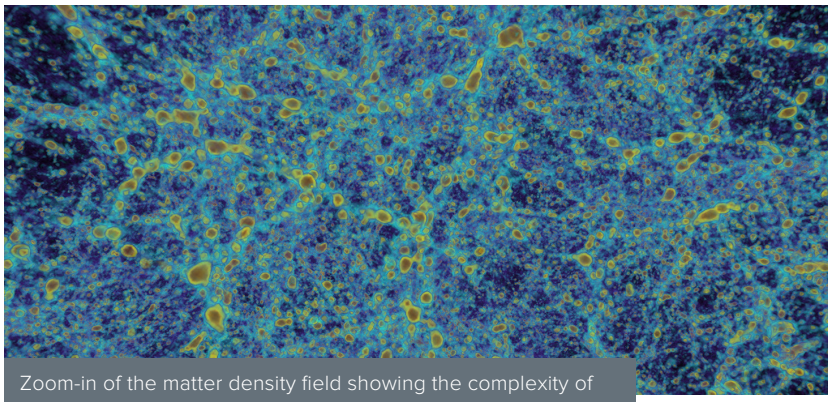

Zoom-in of the matter density field showing the complexity of cosmological structure formation as resolved in a 68-billion-particle simulation carried out on the early access Blue Gene/Q system at the ALCF. The simulation is based around the new HACC framework aimed at exploiting emerging supercomputer architectures such as Mira.

**Image Credit**

Hal Finkel, Salman Habib, Katrin Heitmann, Kalyan Kumaran, Vitali Morozov, Tom Peterka, Adrian Pope, Tim Williams, Mark Hereld, Joseph A. Insley, Aaron Knoll, Michael E. Papka, and Venkatram Vishwanath, Argonne National Laboratory; David Daniel, Patricia Fasel, and Nicholas Frontiere, Los Alamos National Laboratory; Zarija Lukic, Lawrence Berkeley National Laboratory; Eric C. Olson, The University of Chicago

**Paul Mackenzie**
Fermi National Accelerator Laboratory | mackenzie@fnal.gov

# Deepening Our Understanding of Quarks and Gluons Guiding Experiments

*The United States Quantum Chromodynamics (USQCD) Collaboration carries out extensive lattice gauge theory calculations at ALCF. This theory describes how quarks and gluons, the fundamental entities of protons and neutrons, help form atomic nuclei. These efforts are the heart of DOE's large experimental programs in high energy and nuclear physics.*

Quantum chromodynamics (QCD) research plays a key role in the ongoing efforts to develop a unified theory of the fundamental forces of nature. While the behavior of atomic particles such as protons and neutrons is well understood, less is known about the interactions of subatomic particles like quarks and gluons, which compose the atomic particles. Paul Mackenzie from Fermi National Accelerator Laboratory (Fermi) leads the USQCD INCITE project that studies these interactions. Not only does this work directly impact the understanding of physics, but also the research done on resources at the Argonne Leadership Computing Facility provides crucial high-precision lattice QCD calculations needed for new or in-progress experiments and for analyzing results from completed experiments.

Brookhaven National Laboratory's Relativistic Heavy Ion Collider used results from USQCD computation to firmly constrain heavy-ion collision models for the first time. The papers that detail these results are the first- and second-most cited papers of high-temperature QCD with LQCD. Fermilab has used calculations by USQCD members combined with experimental results, allowing many of the fundamental parameters of the Standard Model to be determined more accurately than ever before.

This year, the RBC (Riken-BNL-Columbia) sub-group of USQCD performed the first ab initio analysis of the decay of a K meson into two pions. This process provided the first evidence of charge-parity (CP) violation, the type of physics responsible for the fact that the universe is composed of matter rather than antimatter. The calculation was a landmark because it was the first calculation from first principles of the decay of a hadron (a particle containing quarks) into two other hadrons.

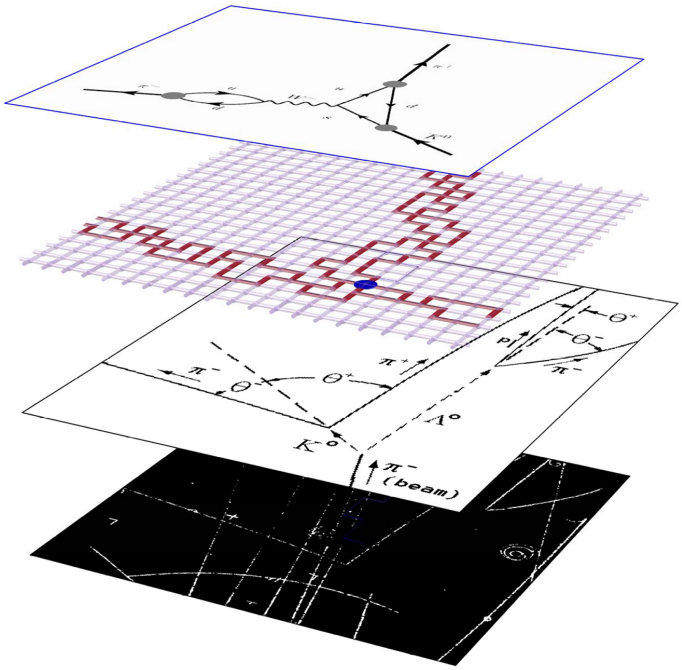


This diagram illustrates the wide range of distance scales that must be understood before the calculations of the decay of a kaon into two pions can be performed.

*"The expanding resources of the leadership computing centers is making possible the large volumes required for processes like the new lattice QCD calculation of the decay of a kaon into two pions, while preserving the needed small lattice spacings."*
*— Paul Mackenzie, Fermi National Accelerator Laboratory*

**Image Credit**
Norman Christ, Columbia University; RBC/USQCD collaboration

**Jean Perez**
University of New Hampshire | jeanc.perez@unh.edu

# Furthering the Understanding of Coronal Heating and Solar Wind Origin

*Two of the most important unsolved problems in heliospheric physics are the origin of the solar wind and the heating of the solar corona. Some observations suggest that Alfvén wave (AW) turbulence in the extended solar atmosphere may be key to solving these unknowns. University of New Hampshire scientists are tackling these challenges with simulations performed on the ALCF's Intrepid, a IBM Blue Gene/P supercomputer.*

Researchers from the Space Science Center at the University of New Hampshire, led by co-principal investigators Jean Perez and Benjamin Chandran, expect to arrive at new theoretical understandings in this area through their INCITE research. The team is conducting the first direct numerical simulations of Alfvén wave (AW) turbulence in the extended solar atmosphere that account for the inhomogeneities in the density, flow speed, and background magnetic field within a narrow magnetic flux tube extending from roughly one solar radius to eleven solar radii. They are comparing numerical simulations conducted on Argonne Leadership Computing Facility (ALCF) supercomputers with remote observations.

To carry out the research, the team is using the Inhomogeneous Reduced Magnetohydrodynamics Code (e.g., IRMHD Code) developed by Perez. The code has been extensively tested and benchmarked on Intrepid, the ALCF's Blue Gene/P supercomputer.

Using the TAU Performance System, a portable profiling and tracing toolkit for performance analysis of parallel programs, performance improvements have been achieved in the asynchronous model of communications being used.

TAU's profile visualization tool, paraprof, allowed the research team to quickly identify sources of performance bottlenecks in the application using the graphical interface. By choosing the best communication strategy for Intrepid, the time in some key routines was reduced by nearly half. This performance improvement translates into big savings—millions of CPU hours. This will allow a more complete and well-resolved physics study with the project's INCITE hours.

The simulations being conducted for the project are of added interest to the space physics community, given the preparations now under way for NASA's Solar Probe Plus mission. With a planned launch date in 2018, this mission will send a spacecraft in to a distance of closest approach to the Sun of Rsun, where Rsun = 695,990 kilometers is one solar radius inside the region that the researchers are simulating numerically. The simulations will enable detailed comparisons with future spacecraft measurements.

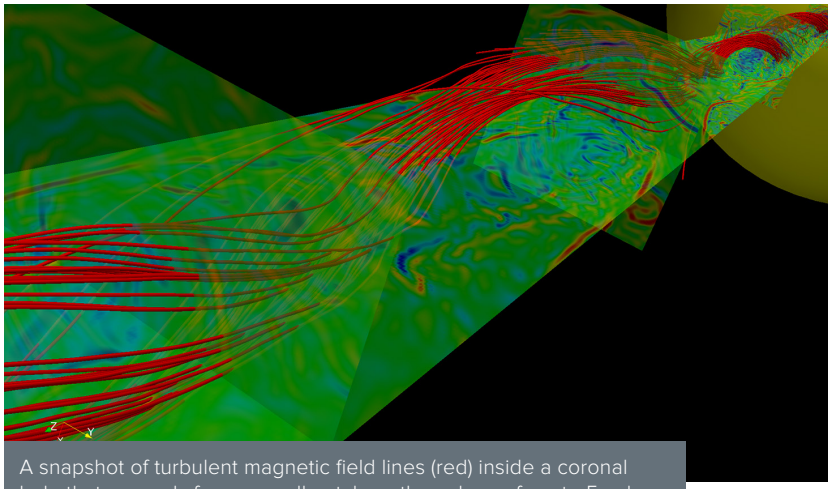

A snapshot of turbulent magnetic field lines (red) inside a coronal hole that expands from a small patch on the solar surface to 5 solar radii. AW, launched by convective motions on the photosphere, propagate in the inhomogeneous Solar atmosphere, producing primary reflected waves that interact non-linearly with the outward waves, driving a turbulent cascade. This cascade continues with secondary reflections in a very complex interplay between wave reflections and nonlinear interactions. Selected slices across the simulation domain show contours of plasma current, indicating the generation of small-scale structures where the turbulent energy ultimately dissipates.

**Image Credit**
Jean Perez, University of New Hampshire

**Sean Couch**
The University of Chicago | smc@flash.uchicago.edu

# Petascale Simulation of Magnetorotational Core-Collapse Supernovae

*A team of researchers led by Sean Couch, Hubble Fellow at the University of Chicago, is using the ALCF's cutting-edge resources to conduct first-ever simulations for exploring the influence of magnetic fields and rotation on the mechanisms that reverse stellar core collapse and drive a supernova. These groundbreaking 3-D studies may profoundly expand our understanding of stellar death and the creation of heavy elements throughout the universe.*

Core-collapse supernovae (CCSNe) are the luminous explosions that herald the death of massive stars. Neutron stars, pulsars, magnetars, and stellar-mass black holes are all born out of these explosions. While their importance in galactic evolution cannot be underestimated, the mechanism of CCSN is not fully understood and remains one of the most important challenges for modern computational astrophysics.

In their studies, the team is conducting a series of 3-D magnetohydrodynamics simulations of the collapse of rotating, magnetic stellar cores using the FLASH multi-physics, adaptive mesh refinement simulation framework. Efforts at the Argonne Leadership Computing Facility include approximate treatments for neutrino physics, realistic progenitor rotation and magnetic fields. These simulations allow researchers to predict the spins, kicks, magnetic field strengths and alignments of newly formed neutron stars, pulsars, and magnetars, as well as the dependence of these parameters on originating conditions. Researchers also hope the simulations will reveal the relative importance of neutrino-driven convection and the standing accretion shock instability in driving the stalled supernova shock outward in realistic, rotating and magnetic progenitors.

Using *in situ* analysis, the team is computing the gravitational wave signal from these simulations, providing invaluable theoretical data for the international physics and astrophysics communities, and aiding in experiments such as the National Science Foundation's Laser Interferometer Gravitational-Wave Observatory. Finally, the research team is calculating nucleosynthetic yields from derived simulations using Lagrangian tracer particles and a large, sophisticated nuclear network.

Using ALCF's Intrepid, Couch conducted a large parameter study of multidimensional core-collapse supernovae exploring the impact of variations in the equation of state on explosion dynamics resulting in a paper that will appear in the *Astrophysical Journal.* This was the most comprehensive study of the dependence of multidimensional core-collapse on the equation of state yet conducted, achievable only with ALCF-scale resources.

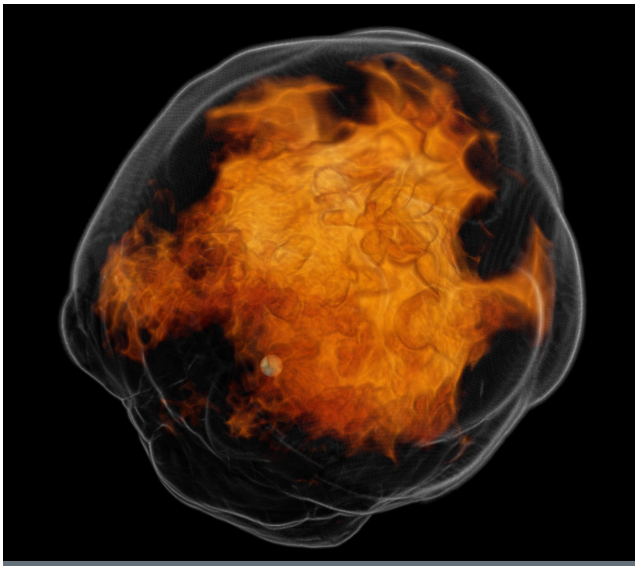

Volume rendering of the highest entropy plumes in a 3-D core-collapse supernova simulation. The edge of the nascent neutron star is shown as the faint blue sphere near the center. Entropy, a measure of the thermodynamic disorder, shows the buoyant convective plumes and turbulence that play a critical part in the core-collapse supernova mechanism. The surface of the outgoing shock wave is also shown in faint gray.

**Image Credit**
Sean Couch, The University of Chicago

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

# Simulations of Laser-Plasma Interactions in Targets for the National Ignition Facility and Beyond

*Providing for the world's energy demands is one of the most urgent and difficult challenges facing our society. Inertial fusion energy provides an attractive solution to the demands for safe, secure, environmentally sustainable energy. Principal Investigator Denise Hinkel is using supercomputing resources at the ALCF to perform simulations that will have a major impact on inertial fusion energy and obtaining ignition at the National Ignition Facility (NIF), a large, laser-based research facility located at Lawrence Livermore National Laboratory (LLNL).*

Scientists have been working to achieve self-sustaining nuclear fusion and energy gain in the laboratory for more than half a century. Lawrence Livermore National Laboratory (LLNL) has been tasked with achieving ignition at the NIF using the inertial confinement fusion process to bring the goal of fusion energy closer to realization. Ignition at NIF will resolve the question of whether through fusion energy is possible through—an attractive solution to the demands for safe, secure, environmentally sustainable energy.

NIF uses 192 powerful lasers to heat and compress a small amount of hydrogen fuel in a laser target to the point where nuclear fusion reactions take place. An important aspect of coupling NIF's laser energy to the fusion target is minimizing laser backscatter, where the incident laser energy scatters off self-generated plasma density fluctuations. This can be a direct loss of input energy. Moreover, backscatter can alter implosion symmetry as well as preheat the ignition capsule via generation of hot electrons. Recent experimental results from NIF show that backscatter occurs in the laser target where quads of laser beams are overlapped.

LLNL's Denise Hinkel is running state-of-the-art simulations at the Argonne Leadership Computing Facility to quantify how overlapping beam quads impact backscatter. These simulations show that overlapping quads can share a reflected light wave, thereby enhancing reflectivity. This is important to understanding the energetics of the experiment, as well as to the improvements in the target design.

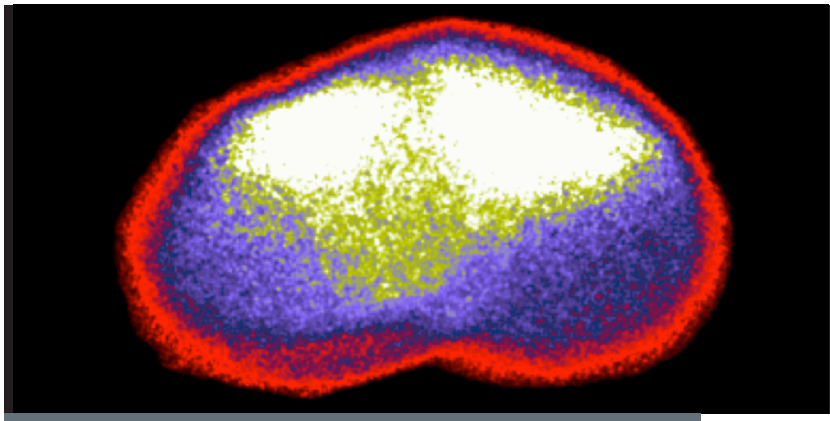

Spatially Non-Uniform Laser Input: Laser input to a pF3D simulation of two NIF laser quads propagating through an ignition target. Here, power transferred from other quads provides a spatially non-uniform distribution of power across the beams. The bright "triangle" in the upper region of each laser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

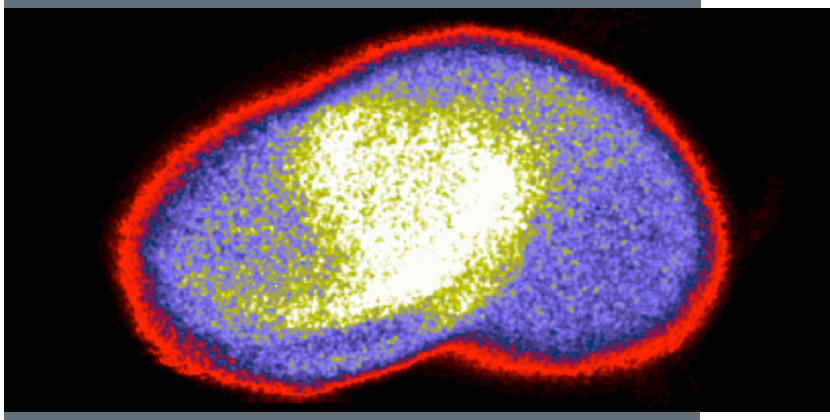

Spatially Uniform Laser Input: Laser input to a pF3D simulation of two NIF laser quads propagating through an ignition target. Here, power transferred from other quads of laser beams is distributed uniformly across the laser beams. Two quads overlap in the simulated region. This enhances reflectivity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

**Image Credit**
Denise Hinkel, Bruce Langdon, Edward Williams, Steven Langer, and Charles Still, Lawrence Livermore National Laboratory

**Jack Wells**

Oak Ridge National Laboratory | wellsjc@ornl.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. If perfected, Lithium/Air batteries could be key to replacing gasoline vehicles with electric vehicles. In a collaborative effort, IBM Research, Vanderbilt University, Oak Ridge National Laboratory, and Argonne National Laboratory are conducting research that may have breakthrough practical applications for electric vehicles, including new technologies for longer-range batteries.*

This multidisciplinary team studied the use of polypropylene carbonate (PC) as a potential electrolyte for a rechargeable Lithium/Air battery, and set a two-year goal for establishing the practicality of PC as the basis for Lithium/Air technology. The fundamental understanding of these systems 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.

As the traditionally preferred choice of gel, problems with PC were overlooked out of the belief that it was stable, based on the behavior of similar systems. Simulations revealed, however, that during discharge, oxygen reduces to peroxide, forming layers of lithium peroxide ($Li_2O_2$) that immediately degrade PC. Moreover, the special properties of the surface of $Li_2O_2$ act to enhance its reactivity. Among the team's other findings, they explained the recent failure to identify $Li_2O_2$ in the discharge products on the cathode surface during experimental research into this technology, and determined the major chemical mechanisms behind the charge/discharge processes.

Making use of leadership-class resources like those available at the ALCF, researchers can employ sophisticated models for both the surface and the electrolyte in sufficient detail to permit studies of the full complexity of PC degradation by lithium/air discharge products. Using these realistic models, the team demonstrated for the first time that PC is not a suitable choice of electrolyte for Lithium/Air-based batteries. Such insights are critical to ultimately realizing goal of a Lithium/Air battery.

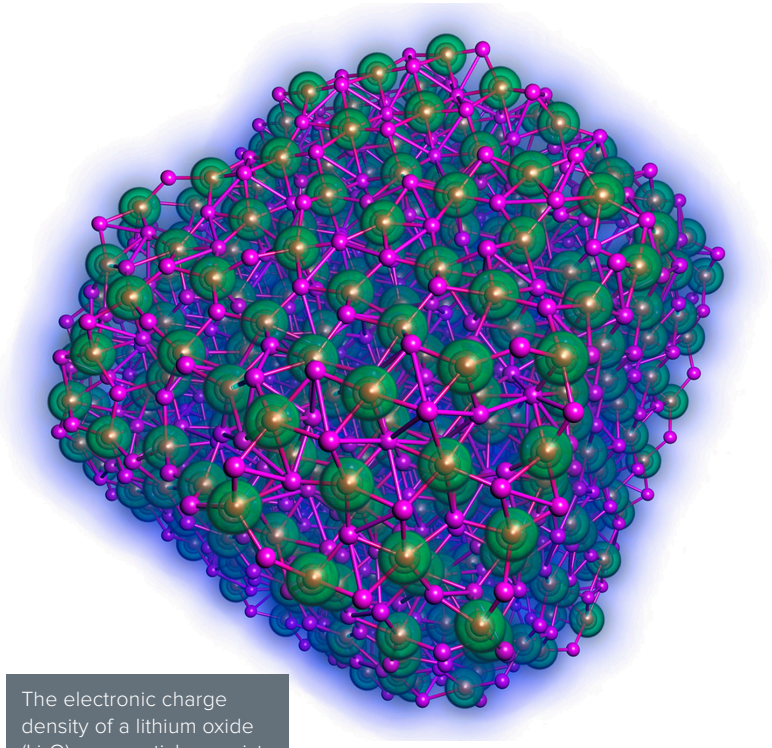

The electronic charge density of a lithium oxide ($Li_2O$) nanoparticle consists of 1500 atoms obtained from Density Functional Theory simulation.

**Image credit:**
Kah Chun Lau, Aaron Knoll, and Larry A. Curtiss, Argonne National Laboratory

**Philippe Spalart**
The Boeing Company | philippe.r.spalart@boeing.com

# Detached-Eddy Simulations and Noise Predictions for Tandem Cylinders

*Tandem Cylinders were introduced by NASA as an example of turbulence and noise generated by flow separation, creating turbulence that collides with another body. These cylinders are relevant to aerospace, ground transportation, energy, civil engineering, and other fields. In an INCITE project conducted on Intrepid, the ALCF's IBM Blue Gene/P, allowed simulations in a wide domain with 84 million grid points and 60,000 time steps.*

Led jointly by Philippe Spalart at Boeing and Ramesh Balakrishnan at the Argonne Leadership Computing Facility, a research team from NTS funded by Boeing achieved results in the turbulent region, which confirmed that Detached-Eddy Simulation (DES) is accurately resolving the instabilities and coherent structures. There is excellent agreement with experiment for the time-average pressures, and more importantly, for the pressure fluctuations on the second cylinder, which reflect turbulence impingement and generate the bulk of the noise. The effect of varying the numerics and, in particular, the numerical dissipation is very moderate by these measures.

The calculation of far-field noise is the true objective of the simulations and remains far more delicate. A key question concerns the validity of simplifying assumptions, based on the low value of the Mach number, proposed by Curle in 1955. In associated work, a mathematical error was found in that article, and misunderstandings in its application were shown to be common in the literature. Two methods were compared: one using the Ffowcs-Williams-Hawkings (FWH) equation with the solid surface of the cylinders as suggested by Curle, and another one, in principle more correct, with a permeable surface surrounding the turbulence. The Tandem-Cylinder noise consists of tonal noise from the vortex shedding and broad-band turbulence noise. The shedding noise is reproduced well by both approaches. The turbulence noise brings out differences. The solid-surface approach agrees with measurements for a wider range of frequencies than the permeable-surface approach. This also happens with landing-gear noise. However, it must be somewhat fortuitous, because a similar effect is obtained by increasing the numerical dissipation, which lowers the accuracy of the simulation. There is a complex interplay between numerics and acoustics.

Progress in the prediction of noise from first principles is real for theory and quality of the simulations, but limited. The community needs more rigor, and if possible, a breakthrough in FWH and similar approaches to extract noise from simulations.

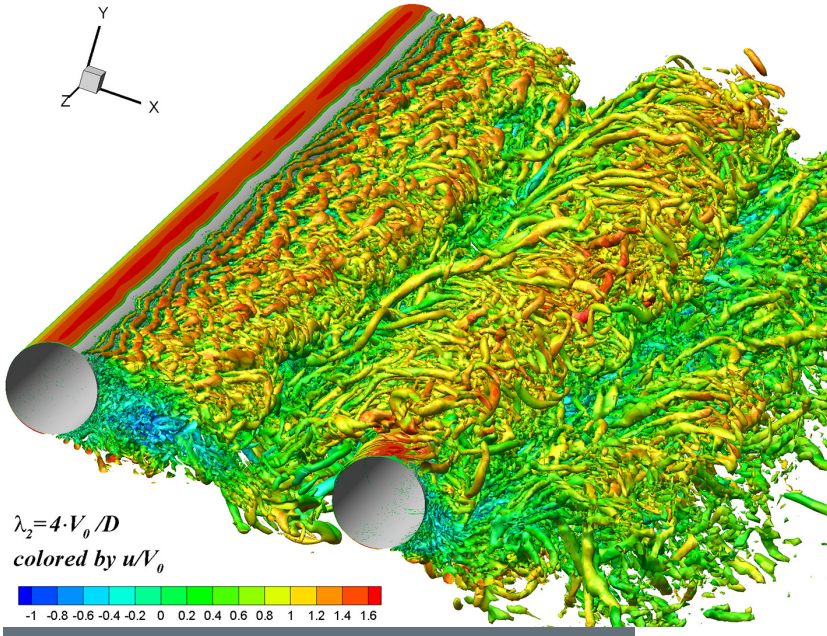


Results in the turbulent region confirmed that Detached-Eddy Simulation (DES) is accurately resolving the instabilities and coherent structures.

**Image credit:**
Philippe Spalart, The Boeing Company

# INCITE

*Open to researchers from academia, government labs, and industry, the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program is the major means by which the scientific community gains access to some of the fastest supercomputers. The program aims to accelerate scientific discoveries and technological innovations by awarding, on a competitive basis, time on supercomputers to researchers with large-scale, computationally intensive projects that address "grand challenges" in science and engineering.*

## Biological Science

### Protein-Ligand Interaction Simulations and Analysis
T. Andrew Binkowski, Argonne National Laboratory
*Allocation: 10 Million Hours*

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

## Chemistry

### Towards Breakthroughs in Protein Structure Calculation and Design
David Baker, University of Washington
*Allocation: 33 Million Hours*

### Simulations of Deflagration-to-Detonation Transition in Reactive Gases
Alexei Khokhlov, The University of Chicago
*Allocation: 20 Million Hours*

### Energetic Aspects of CO2 Absorption by Ionic Liquids from Quantum Monte Carlo
William Lester, UC Berkeley
*Allocation: 4 Million Hours*

### Large-Eddy Simulation of Two-Phase Flow Combustion in Gas Turbines
Thierry Poinsot, European Center 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*

## Computer Science

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

### Fault-Oblivious Exascale Computing Environment
Ronald Minnich, Sandia National Laboratories
*Allocation: 10 Million Hours*

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

## Earth Science

### CyberShake 3.0: Physics-based Probabilistic Seismic Hazard Analysis
Thomas Jordan, Southern California Earthquake Center
*Allocation: 2 Million Hours*

### Large-Eddy Simulations of Contrail-to-Cirrus Transition
Roberto Paoli, CERFACS
*Allocation: 20 Million Hours*

### Climate-Science Computational Development Team: The Climate End Station II
Warren Washington, National Center for Atmospheric Research
*Allocation: 30 Million Hours*

## Energy Technologies

### Optimization of Complex Energy System Under Uncertainty
Mihai Anitescu, Argonne National Laboratory
*Allocation: 10 Million Hours*

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

### Atomistic Adaptive Ensemble Calculations of Eutectics of Molten Salt Mixtures
Saivenkataraman Jayaraman, Sandia National Laboratories
*Allocation: 10 Million Hours*

Enabling Green Energy and Propulsion Systems via Direct Noise Computation
Umesh Paliath, GE Global Research
*Allocation: 45 Million Hours*

## Engineering

Direct Simulation of Fully Resolved Vaporizing Droplets in a Turbulent Flow
Said Elghobashi, University of California—Irvine
*Allocation: 20 Million Hours*

Stochastic (w*) Convergence for Turbulent Combustion
James Glimm, Stanford University
*Allocation: 35 Million Hours*

Adaptive Detached Eddy Simulation of a Vertical Tail with Active Flow Control
Kenneth Jansen, University of Colorado—Boulder
*Allocation: 40 Million Hours*

Turbulent Multi-Material Mixing in the Richtmyer-Meshkov Instability
Sanjiva Lele, Stanford University
*Allocation: 20 Million Hours*

## Materials Science

Vibrational Spectroscopy of Liquid Mixtures and Solid/Liquid Interfaces
Giulia Galli, University of California—Davis
*Allocation: 25 Million Hours*

High-Fidelity Simulation of Complex Suspension Flow for Practical Rheometry
William George, National Institute of Standards
*Allocation: 22 Million Hours*

Probing the Non-Scalable Nano Regime in Catalytic Nanoparticles with Electronic Structure Calculations
Jeffrey Greeley, Argonne National Laboratory
*Allocation: 10 Million Hours*

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

Multiscale Modeling of Energy Storage Materials
Gregory Voth, The University of Chicago
*Allocation: 25 Million Hours*

## Physics

Simulations of Laser-Plasma Interactions in Targets for the National Ignition Facility and Beyond
Denise Hinkel, Lawrence Livermore National Laboratory
*Allocation: 63 Million Hours*

Toward Exascale Computing of Type Ia and Ib, c Supernovae: V&V of Current Models
Donald Lamb, The University of Chicago
*Allocation: 40 Million Hours*

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

Petascale Simulations of Inhomogeneous Alfvén Turbulence in the Solar Wind
Jean Perez, University of New Hampshire
*Allocation: 10 Million Hours*

Nuclear Structure and Nuclear Reactions
James Vary, Iowa State University
*Allocation: 18 Million Hours*

# ASCR Leadership Computing Challenge (ALCC)

*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 DOE, 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. This list reflects the 2012-2013 ALCC awards.*

## Computer Science

HPC Colony: Adaptive System Software for Improved Resiliency and Performance
Terry Jones, Oak Ridge National Laboratory
*Allocation: 3 Million Hours*

## Energy Technologies

Validation Work for Heterogeneous Nuclear Reactor Calculations
Micheal Smith, Argonne National Laboratory
*Allocation: 30 Million Hours*

## Engineering

Petascale Thermal-Hydraulic Simulations in support of CESAR
Elia Merzari, Argonne National Laboratory
*Allocation: 30 Million Hours*

Prediction of Multiscale, Multiphysics Turbulent Flow Phenomena Using Unstructured Large-Eddy Simulation
Parviz Moin, Stanford University
*Allocation: 60 Million Hours*

U.S.-Russia Collaboration on Verification and Validation in Thermal Hydraulics: Nek5000 and Conv3D Simulation of "SIBERIA" Experiment
Aleksandr Obabko, The University of Chicago
*Allocation: 30 Million Hours*

## Materials Science

Dynamics of Conformational Transition in Thermo-Sensitive Polymers and Hydrogels
Subramanian Sankaranarayanan, Argonne National Laboratory
*Allocation: 50 Million Hours*

Liquid-Solid Interfaces in Electrocatalysis from First Principles
Jeffrey Greeley, Argonne National Laboratory
*Allocation: 20 Million Hours*

## Physics

*Ab Initio* Hyper-Nuclear Physics
Kostas Orginos, College of William & Mary
*Allocation: 20 Million Hours*

ALCC: Exploring the Nature of the Lightest Massive Particles in the Universe
Katrin Heitmann, Argonne National Laboratory
*Allocation: 4 Million Hours*

# 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 Science

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

NAMD – The Engine for Large-Scale Classical MD Simulations of Biomolecular Systems Based on a Polarizable Force Field
Benoit Roux, Argonne National Laboratory and 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
Robert Harrison, Oak Ridge National Laboratory
*Allocation: 150 Million Hours*

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

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

## Earth Science

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

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

## Energy Technologies

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

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

## Engineering

Direct Numerical Simulation of Autoignition in a Jet in a Cross-Flow
Christos Frouzakis, Swiss Fed. Inst. Tech.
*Allocation: 150 Million Hours*

Petascale, Adaptive CFD
Kenneth Jansen, University of Colorado-Boulder
*Allocation: 150 Million Hours*

## Physics

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

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

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

LatticeQCD - Early Science
Paul Mackenzie, Fermilab
*Allocation: 150 Million Hours*

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

## 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 2012 Director's Discretionary projects appears below.*

### Biological Science

Multi-Scale Simulations of Deformable Blood Vessels
Leopold Grinberg, Brown University
*Allocation: 500,000 Hours*

High-Performance Neuron Simulations on the Blue Gene/Q
Michael Hines, Yale University
*Allocation: 1 Million Hours*

### Chemistry

Water Systems from Highly Accurate Quantum Monte Carlo Calculations
Dario Alfè, University College London
*Allocation: 1 Million Hours*

Multi-scale Modeling of Catalytic Interfaces Based on 2-D Sub-Nano Surface-Deposited Clusters
Anastassia N. Alexandrova,
University of California, Los Angeles
*Allocation: 1 Million Hours*

Port ACES III and SIAL
Erik Deumens, University of Florida
*Allocation: 500,000 Hours*

Machine Learning for the Exploration of Chemical Compound Space
O. Anatole von Lilienfeld, Argonne National Laboratory
*Allocation: 1 Million Hours*

### Computer Science

Parallel Run-Time Systems
Jeff Hammond, Argonne National Laboratory
*Allocation: 4 Million Hours*

Charm++ and its Applications
Laxmikant V. Kale, University of Illinois at Urbana-Champaign
*Allocation: 1.5 Million Hours*

Parallel Boost Graph Library
Andrew Lumsdaine, Indiana University
*Allocation: 100,000 Hours*

SciDAC Scalable Data Management Analysis and Visualization
Michael Papka, Argonne National Laboratory
*Allocation: 900,000 Hours*

Visualization and Analysis Research and Development for Argonne Leadership Computing Facility
Michael E. Papka, Argonne National Laboratory
*Allocation: 500,000 Hours*

Distributed File Systems for Exascale Computing
Ioan Raicu, Illinois Institute of Technology
*Allocation: 350,000 Hours*

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

Parallel Performance Evaluation Using the TAU Performance System
Sameer Shende, ParaTools, Inc.
*Allocation: 100,000 Hours*

TotalView Debugger on Blue Gene/P
Peter Thompson, TotalView Technologies
*Allocation: 250,000 Hours*

### Earth Science

Dynamic Downscaling of Climate Models
V. Rao Kotamarthi, Argonne National Laboratory
*Allocation: 1.25 Million Hours*

Large-Eddy Simulations of Atmospheric Convection
David Romps, Lawrence Berkeley National Laboratory
*Allocation: 700,000 Hours*

Sensitivity and Uncertainty of Precipitation of the GFDL High-Resolution Model
Laura Zamboni, Argonne National Laboratory
*Allocation: 100,000 Hours*

## Energy Technologies

Molecular Modeling of Nanoscale Transport Pertinent to Energy Conversion and Storage
Soumik Banerjee, Washington State University
*Allocation: 250,000 Hours*

Oxygen-Sensitivity Problem of Hydrogenases
Martin Stiebritz, ETH Zurich
*Allocation: 1 Million Hours*

## Engineering

A Computational Study of a Proprietary Parallel, Arbitrary Cell Type Code Used by G.E. Global Research
Ramesh Balakrishnan, Argonne National Laboratory
*Allocation: 1 Million Hours*

Porting ExaFMM to Blue Gene/P Architecture and Turbulence Validation Studies
Lorena A. Barba, Boston University
*Allocation: 1 Million Hours*

Quantum Lattice Algorithm for Quantum Turbulence
George Vahala, College of William & Mary
*Allocation: 300,000 Hours*

## Fusion Energy

Global Gyrokinetic Particle-in-Cell Investigations of ITER-relevant Fusion Plasmas
William M. Tang, Princeton University, PICSciE and Princeton Plasma Physics Laboratory
*Allocation: 5 Million Hours*

## Materials Science

Quantum Monte Carlo Methods for Solids and Liquids
Dario Alfè, University College London
*Allocation: 500,000 Hours*

SPaSM Molecular Dynamics Simulations of Material Dynamics
Timothy C. Germann, Los Alamos National Laboratory
*Allocation: 500,000 Hours*

First-principle Investigations of Oxygen Defects in Metal-Oxide-Metal Heterostructures
Olle Heinonen, Argonne National Laboratory
*Allocation: 500,000 Hours*

Modeling Oil Properties with Molecular Dynamics
Detlef Hohl, Shell International E&P, Inc.
*Allocation: 150,000 Hours*

Phonon Parallelization in Quantum ESPRESSO
William Parker, Argonne National Laboratory
*Allocation: 500,000 Hours*

Wavelet Basis Set in Density Functional Theory Methods for Photoelectric Materials
Alvaro Vazquez-Mayagoitia, Argonne National Laboratory
*Allocation: 900,000 Hours*

## Nuclear Energy

Modeling of Defects in Materials for Energy Applications
John J. Low and Marius Stan, Argonne National Laboratory
*Allocation: 500,000 Hours*

## Physics

3-D Simulations of Magnetorotational Core-Collapse Supernovae
Sean M. Couch, University of Chicago and Fermilab
*Allocation: 5 Million Hours*

NEK5000
Paul Fischer, Argonne National Laboratory
*Allocation: 1 Million Hours*

Electromagnetics
Misun Min, Argonne National Laboratory
*Allocation: 500,000 Hours*

First-Principles Calculation of Laser Induced Ultrafast Magnetism
Guoping Zhang, Indiana State University
*Allocation: 1.5 Million Hours*

## 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.

The Argonne Leadership Computing Facility is supported by DOE's Office of Science, Advanced Scientific Computing Research (ASCR) program.

## 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

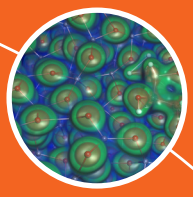

Electron density obtained from a density functional theory (DFT) calculation of lithium oxide ($Li_2O$) performed with the GPAW code. This visualization was the result of a simulation run on Intrepid, a supercomputer at the ALCF.

A billion-atom reactive molecular dynamics simulation of nanobubble collapse in water near a ceramic surface under shock compression. The high pressure in the shock wave deforms the ceramic surface and also accelerates water molecules from the bubble's periphery inward. These high velocity water molecules bunch up to form a nanojet, whose impact damages the ceramic surface. These simulations 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 Intrepid, an IBM Blue Gene/P at the ALCF.

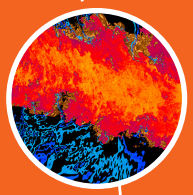

The rectangular nozzle is shown in gray with an isosurface of temperature (gold) cut along the center plane of the nozzle showing temperature contours (red/yellow). The acoustic field is visualized by (blue/cyan) contours of the pressure field taken along the same plane. The chevrons enhance turbulent mixing just downstream of the nozzle exit shaping how the jet spreads downstream. This significantly reduces the noise produced by the supersonic jet compared to a rectangular nozzle without chevrons (not shown).

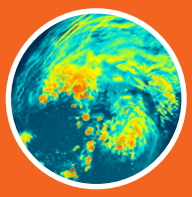

Color-enhanced satellite view snapshot from a recent climate simulation using the Nested Regional Climate Model at 4 kilometers horizontal resolution.

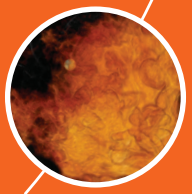

Volume rendering of the highest entropy plumes in a 3D core-collapse supernova simulation. The edge of the nascent neutron star is shown as the faint blue sphere near the center. Entropy, a measure of the thermodynamic disorder, shows the buoyant convective plumes and turbulence that play a critical part in the core-collapse supernova mechanism. The surface of the outgoing shock wave is also shown in faint gray.

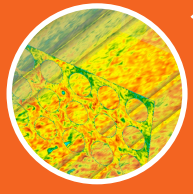

Velocity magnitude distribution in a flow through the 25-pin swirl-vane spacer grid of Matis benchmark. Computed on Intrepid with Nek5000 and visualized on Eureka with VisIt at the ALCF.

www.alcf.anl.gov

## CONTACT

media@alcf.anl.gov

1-877-252-8615

U.S. DEPARTMENT OF ENERGY

Argonne National Laboratory is a U.S. Department of Energy laboratory managed by UChicago Argonne, LLC.

alcf_indsbro_1112_ss_mn