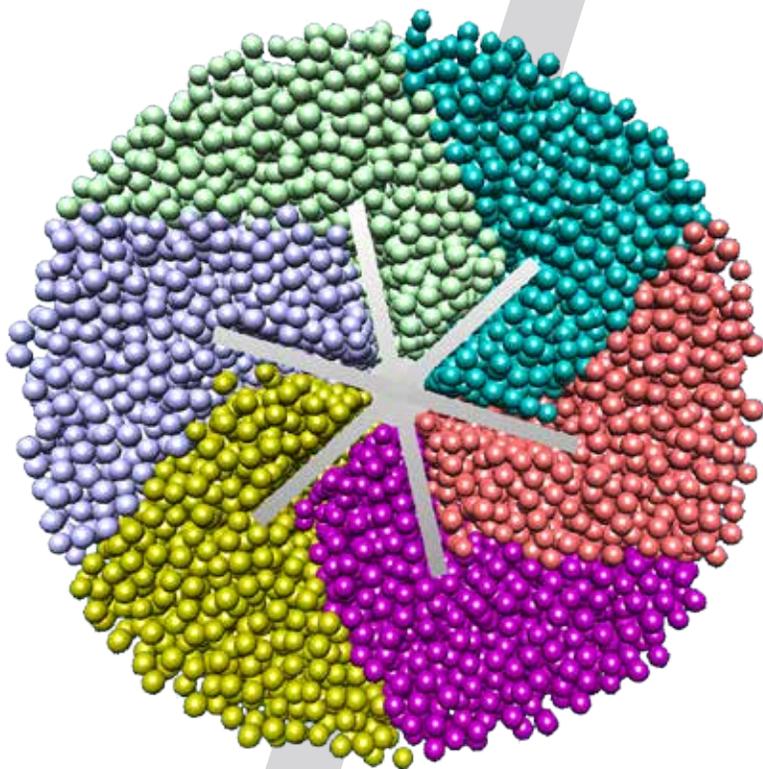


ARGONNE **LEADERSHIP COMPUTING** FACILITY

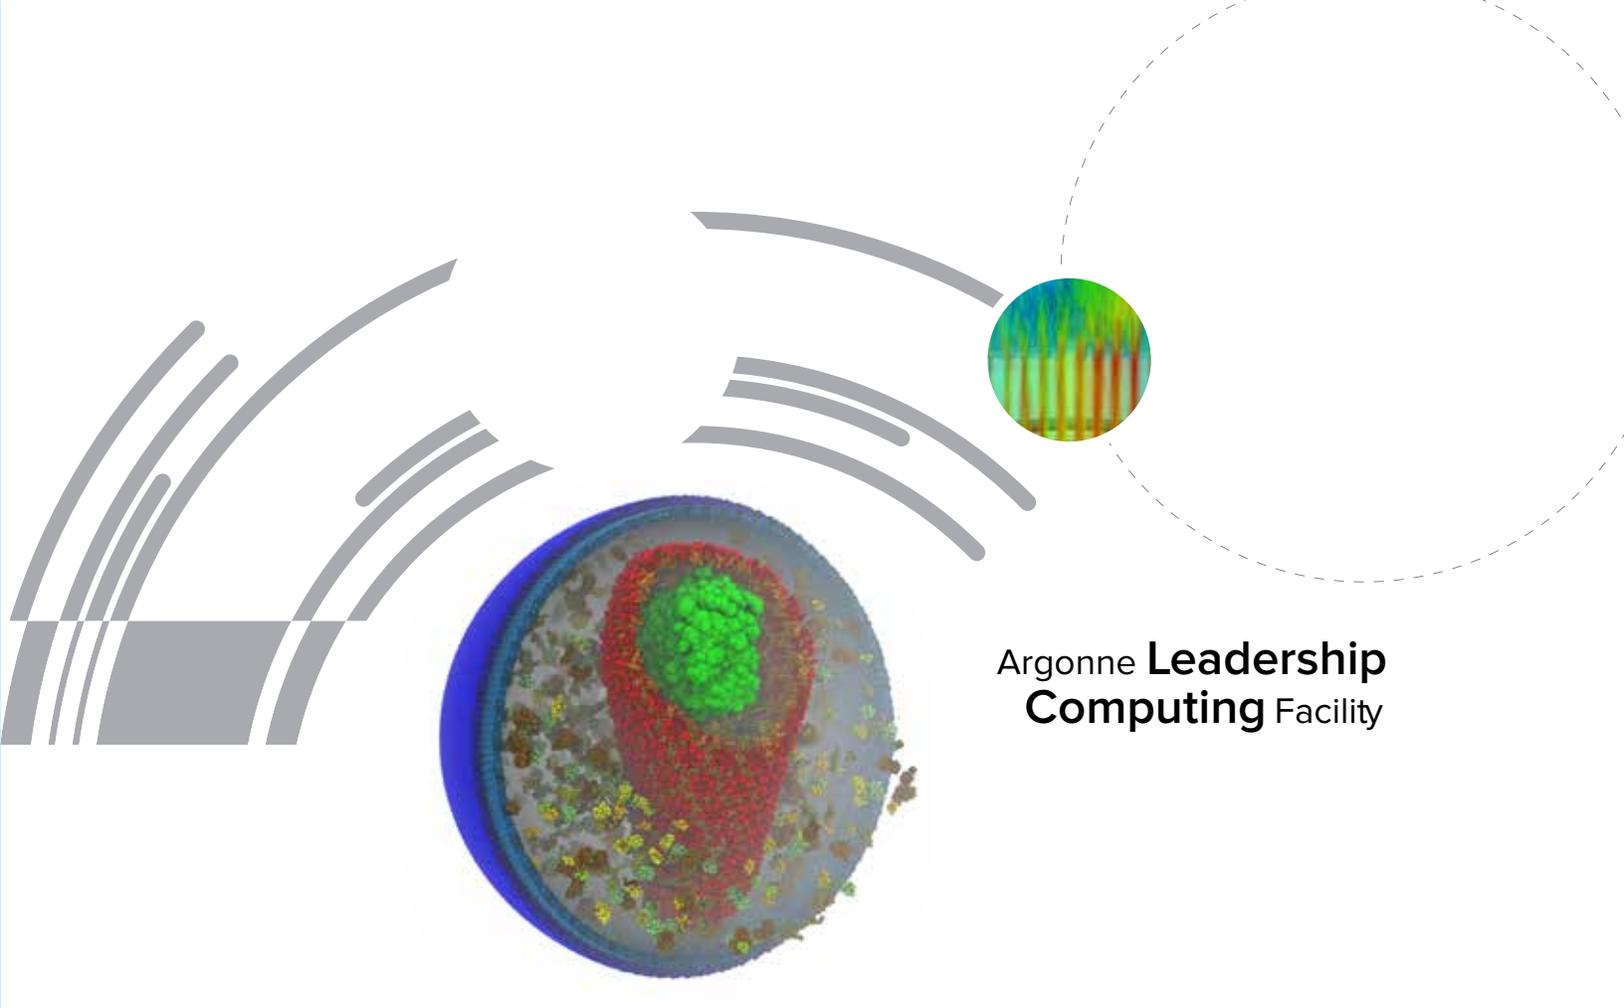


**2013**

SCIENCE  
HIGHLIGHTS

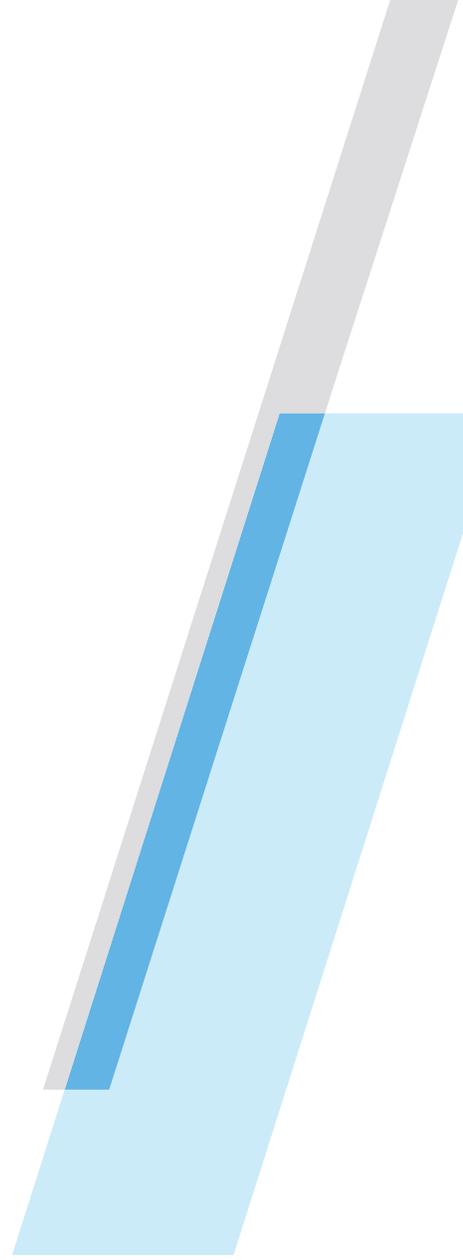
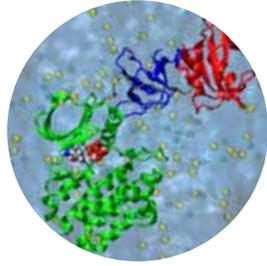






Argonne **Leadership**  
**Computing** Facility





# CONTENTS

|   |    |
|---|----|
| ABOUT ALCF.....   | 2  |
| MIRA.....   | 3  |
| SCIENCE DIRECTOR'S MESSAGE.....   | 4  |
| ALLOCATION PROGRAMS.....  | 5  |
| EARLY SCIENCE PROGRAM.....  | 7  |
| Ab Initio Reaction Calculations for Carbon-12 .....   | 8  |
| Accurate Numerical Simulations of Chemical Phenomena Involved in Energy Production and Storage with MADNESS and MPQC..... | 9  |
| Climate-Weather Modeling Studies Using a Prototype Global Cloud-System Resolving Model.....                               | 10 |
| Cosmic Structure Probes of the Dark Universe .....  | 11 |
| Direct Numerical Simulation of Autoignition in a Jet in a Cross-Flow .....  | 12 |
| Global Simulation of Plasma Microturbulence at the Petascale and Beyond .....   | 13 |
| High-Accuracy Predictions of the Bulk Properties of Water.....  | 14 |
| High-Speed Combustion and Detonation.....   | 15 |
| Lattice QCD - Early Science.....  | 16 |

|   |           |
|---|-----------|
| Materials Design and Discovery: Catalysis and Energy Storage .....  | 16        |
| Multiscale Molecular Simulations at the Petascale.....  | 17        |
| NAMD - The Engine for Large-Scale Classical MD Simulations of Biomolecular Systems Based<br>on a Polarizable Force Field..... | 18        |
| Petascale, Adaptive CFD .....   | 19        |
| Petascale Direct Numerical Simulations of Turbulent Channel Flow.....   | 20        |
| Petascale Simulations of Turbulent Nuclear Combustion .....   | 21        |
| Using Multiscale Dynamic Rupture Models to Improve Ground Motion Estimates.....   | 22        |
| <b>SCIENCE HIGHLIGHTS .....</b>   | <b>23</b> |
| <b>BIOLOGICAL SCIENCES</b>  |           |
| Computational Studies of Nucleosome Stability .....   | 24        |
| Protein Binding and Interaction Analysis of Human Pathogen Protein Targets.....   | 25        |
| Studies of Large Conformational Changes in Biomolecular Machines .....  | 26        |
| Towards Breakthroughs in Protein Structure Calculation and Design .....   | 27        |
| <b>CHEMISTRY</b>  |           |
| Simulations of Deflagration-to-Detonation Transition in Reactive Gases.....   | 28        |
| <b>COMPUTER SCIENCE</b>   |           |
| Developing and Testing Future Applications and Operating Systems for Exascale.....  | 29        |
| <b>EARTH SCIENCE</b>  |           |
| Attributing Changes in the Risk of Extreme Weather and Climate.....   | 30        |
| CyberShake 3.0: Physics-Based Probabilistic Seismic Hazard Analysis .....   | 31        |
| Evaluation of Mesoscale Atmospheric Model for Contrail Cirrus Simulations .....   | 32        |
| <b>ENERGY TECHNOLOGIES</b>  |           |
| Direct Numerical Simulations of High Reynolds Number Turbulent Channel Flow.....  | 33        |
| Optimization of Complex Energy System Under Uncertainty.....  | 34        |
| Predictive Modeling for Complex Phenomena in Electromagnetics and Fluid Systems.....  | 35        |
| Thermodynamics of Binding Biomass to Cellulases for Renewable Fuel .....  | 36        |

**ENGINEERING**

|  |    |
|--|----|
| Combustion Stability in Complex Engineering Flows .....  | 37 |
| Direct Numerical Simulation of Autoignition in a Jet in a Cross-Flow .....                                       | 38 |
| Enabling Green Energy and Propulsion Systems via Direct Noise Computation .....                                  | 39 |
| Petascale Thermal-Hydraulic Simulations in Support of CESAR .....  | 40 |
| Prediction of Multiscale, Multiphysics Turbulent Flow Phenomena Using<br>Unstructured Large Eddy Simulation..... | 41 |
| U.S.-Russia Collaboration on Verification and Validation in Thermal Hydraulics.....                              | 42 |

**MATERIALS SCIENCE**

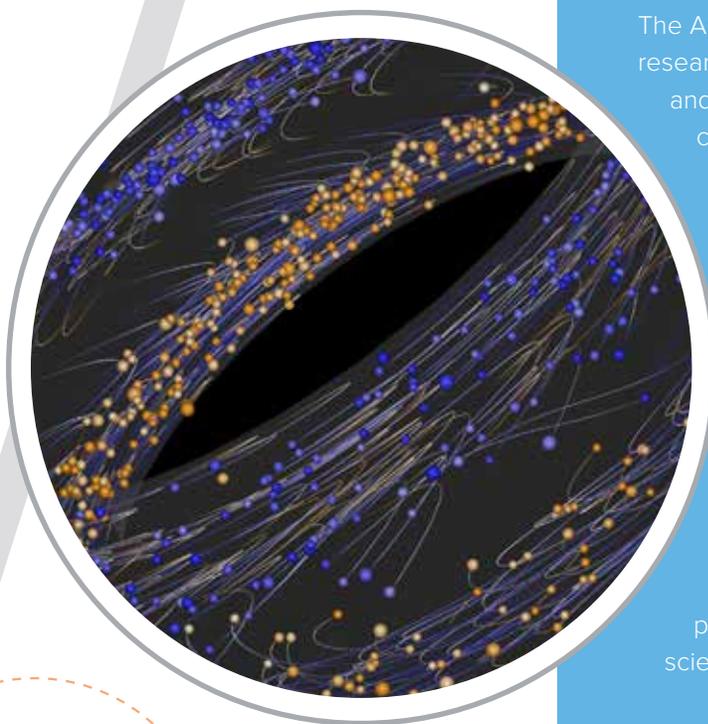
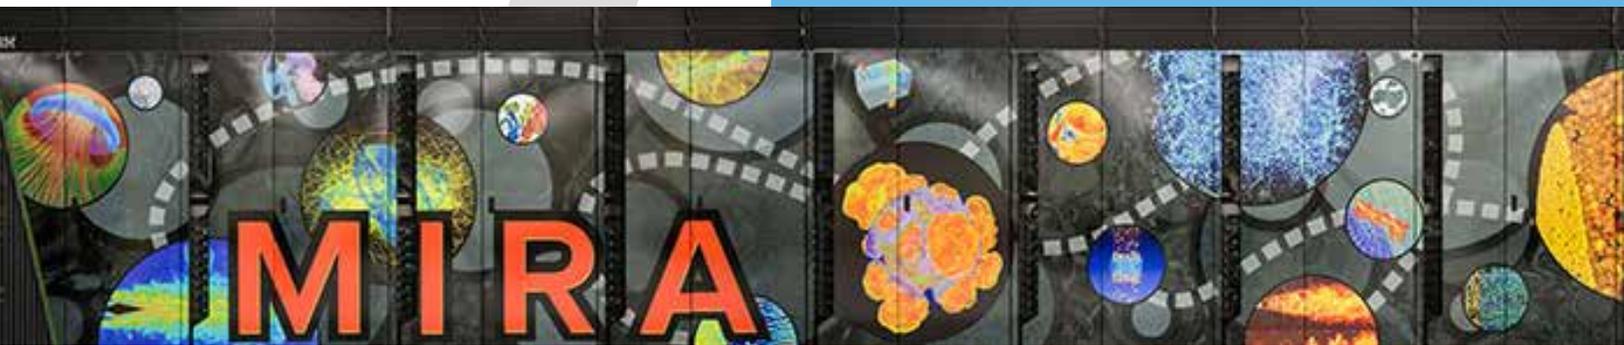
|   |    |
|---|----|
| Dynamics of Conformational Transition in Thermo-Sensitive Polymers and Hydrogels..... | 43 |
| High-Fidelity Simulation of Complex Suspension Flow for Practical Rheometry.....      | 44 |
| Multibillion-Atom MD Studies of the Mechanical Response of Nanocrystalline Ta .....   | 45 |
| Multiscale Modeling of Energy Storage Materials .....                                 | 46 |
| Non-Covalent Bonding in Complex Molecular Systems with Quantum Monte Carlo .....      | 47 |
| Petascale Simulations of Stress Corrosion Cracking .....                              | 48 |
| Vibrational Spectroscopy of Liquid Mixtures and Solid/Liquid Interfaces.....          | 49 |

**PHYSICS**

|  |    |
|--|----|
| Ab Initio Quantum Liquid Water and Aqueous Ionic Solutions.....                        | 50 |
| Accelerator Simulations for the Intensity Frontier of Particle Physics.....            | 51 |
| Computing the Dark Universe .....  | 52 |
| Kinetic Simulations of Fusion Energy Dynamics at the Extreme Scale .....               | 53 |
| Lattice QCD.....   | 54 |
| Next-Generation Petascale Simulations of Type Ia Supernovae.....                       | 55 |
| Petascale Simulation of Magnetorotational Core-Collapse Supernovae .....               | 56 |
| Simulation of Laser-Plasma Interaction in National Ignition Facility Experiments ..... | 57 |

**PROJECT AWARDS..... 59**

|                      |    |
|----------------------|----|
| INCITE 2013 .....    | 60 |
| ALCC 2012-2013 ..... | 62 |
| ESP.....             | 63 |
| DD 2013.....         | 64 |



## ABOUT ALCF

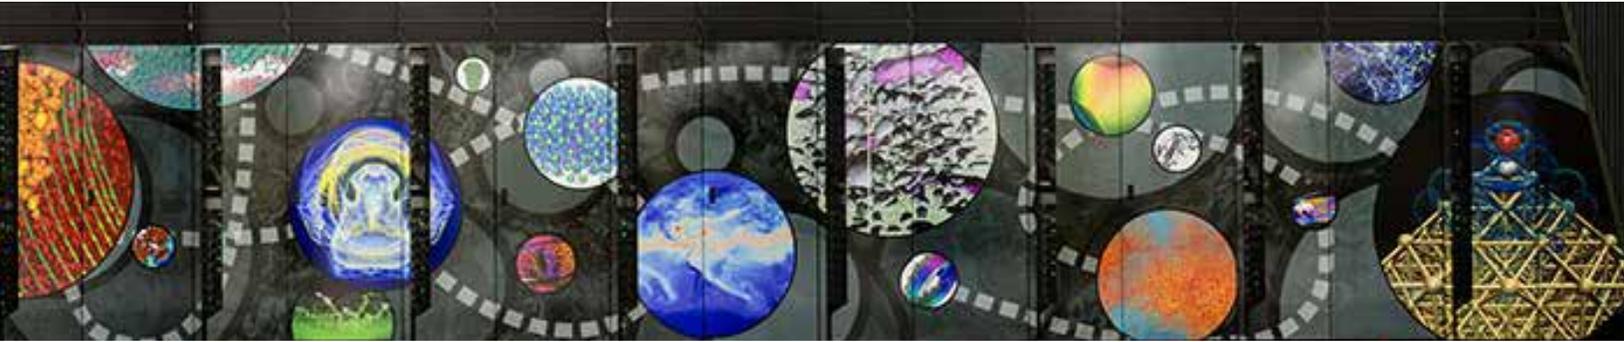
The Argonne Leadership Computing Facility provides researchers from national laboratories, academia, and industry with access to high-performance computing capabilities to enable breakthrough science and engineering. Supported by the U.S. Department of Energy Office of Science, the ALCF is one of two leadership computing facilities in the nation dedicated to open science.

## MISSION

The ALCF's mission is to accelerate major scientific discoveries and engineering breakthroughs by designing and providing world-leading computing facilities in partnership with the computational science community.



# {MIRA} AN ENGINE FOR DISCOVERY



The advent of petascale supercomputers allows researchers to conduct calculations and simulations on unprecedented scales, and the ALCF houses one of the fastest in the world. An IBM Blue Gene/Q supercomputer, Mira went into production on April 9, 2013. Its architecture is the result of a co-design collaboration between Argonne National Laboratory, Lawrence Livermore National Laboratory, and IBM Research.

Capable of carrying out 10 quadrillion calculations per second, Mira's enormous computational power can be attributed to a symbiotic architecture that combines novel designs for fine- and coarse-grain parallelism with corresponding network and I/O capabilities that redefine standards in supercomputer processing and networking.

With the ability to employ 10 times as much RAM and 20 times as much peak speed in floating-point calculations, Mira has far exceeded the capabilities of its immediate predecessor Intrepid, an IBM Blue Gene/P.

An escalation in Mira's node-level parallelism increases the number of tasks a node can perform. Each compute node has 16 processing cores supported by simultaneous multi-threading (SMT), which allows a core to run four simultaneous threads of execution—all adding up to 64 simultaneous tasks per node. This is further enhanced by quad processing extensions (QPX), which allow for single instruction multiple data (SIMD) operation on each thread. With this, Mira can execute eight floating-point operations with a single instruction.

This type of high-level computing would not be achievable without an equally highly structured

communications network, such as Blue Gene/Q's 5D torus interconnect. Designed to provide greater bandwidth and lower latency relative to a 3D torus, more information can be relayed simultaneously and faster.

Buttressing the main architecture are custom features that help the user access and modify data more efficiently.

- ▶ The L1 cache prefetcher greatly reduces the time spent waiting for memory access and allows applications to run better by predicting and delivering data that the user needs before it is actually requested.
- ▶ Atomic operations ensure that data is read and written correctly when accessed by multiple threads simultaneously. These are implemented in the L2 cache rather than at the processor core, which dramatically reduces their cost.
- ▶ Transactional memory manages the concurrent use of multiple threads, protecting against conflicting operations in the atomic region and correcting for it. It is the first such hardware in a commercial supercomputer.

The Blue Gene/Q system demonstrates the major benefits that can be gained from innovative architectural features that are developed with a deep understanding of scientific computing requirements. Such an approach is already leading researchers down the path toward exascale computing. Making it possible to perform computations 100 times faster than today's most powerful computers, these machines will truly prove themselves engines of discovery.

## ENTERING THE MIRA ERA

This has been an exciting year for the ALCF. We are finishing up an extremely productive five-year run on Intrepid, while also beginning the next era of supercomputing at the ALCF with the launch of Mira, our 10-petaflops IBM Blue Gene/Q supercomputer.



**Paul Messina**

Director of Science, ALCF

Since entering production mode in April, we've seen firsthand the value Mira is providing to our research community. Our new supercomputer is enabling simulations of unprecedented scale and accuracy, allowing scientists to tackle larger and more complex problems than ever before.

For example, we have one research team using Mira to study the behavior of enzymes in an effort to improve biofuel production. The power of Mira is allowing them for the first time to calculate accurate binding free energy values for enzymes with a computer. This information is difficult to obtain through conventional lab experiments, so these computations have the potential to accelerate biofuels research.

Turbulence research is also entering new territory, thanks to Mira. Scientists are studying wall-bounded turbulence at an unprecedented level of detail with the largest production direct numerical simulations performed to date. Wall-bounded turbulence is the phenomenon that results in drag and energy dissipation when a vehicle (or a pipe) interacts with the fluid flowing past it. This work is expected to lead to an improved understanding of the underlying physics and the creation of new tools for the reduction of drag.

Another project is taking advantage of Mira to conduct the largest and most detailed simulations of the universe ever performed. This has given researchers the first simulations that are accurate enough to be used in comparisons with the state-of-the-art sky surveys of today and tomorrow, including the images provided by the Large Synoptic Survey Telescope (LSST). The hope is that this effort will lead to a better understanding of one of the greatest mysteries in science today—the nature of dark energy and dark matter.

As a final example, Mira is also helping scientists to blaze new trails in particle physics with the most complex and detailed accelerator beam dynamics simulations ever performed. Results from these studies will help minimize the cost and risks of developing a new generation of particle accelerators that will, in turn, provide another scientific tool to advance both basic and applied sciences.

These are just a few of the pioneering projects currently using Mira. In the following pages, you'll read about how dozens of our users are tapping Mira (and Intrepid for one final year) to fast-forward their research efforts into the future.

## ALLOCATION PROGRAMS

Any researcher with a question that requires large-scale computing systems can submit a proposal for time on ALCF resources to run simulations for their experiments. Typically awarded in chunks of millions of core-hours, the following allocation programs add up to billions of hours of computing time per year.

### 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 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 one to three years.

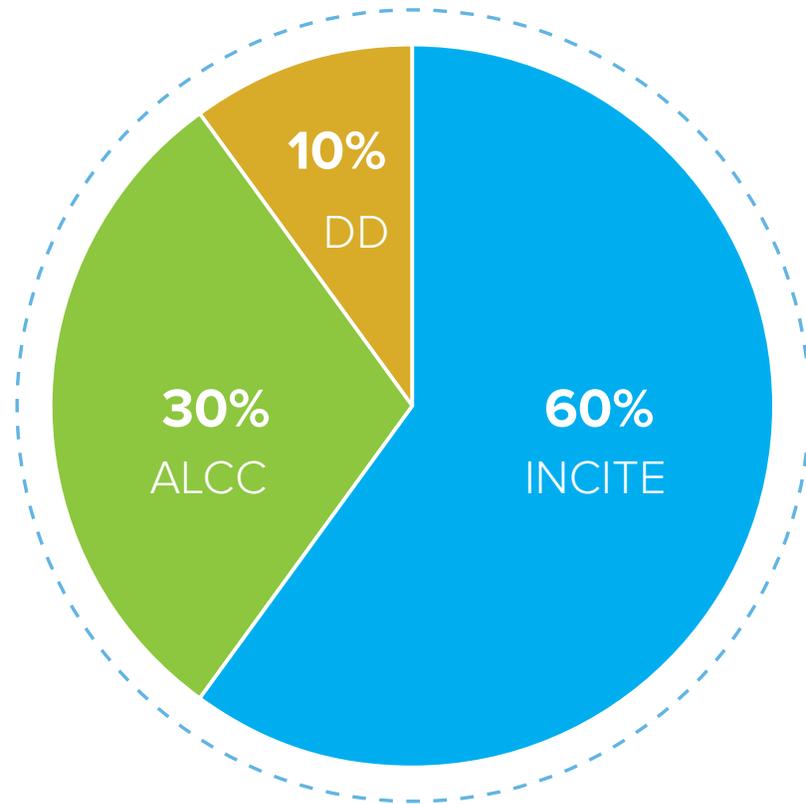
### ASCR LEADERSHIP COMPUTING CHALLENGE (ALCC)

The DOE's ALCC program allocates resources to projects directly related to the DOE's energy mission, national emergencies, or for broadening the community of researchers capable of using leadership computing resources. 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 yearlong allocation cycle runs from July 1 to June 30.

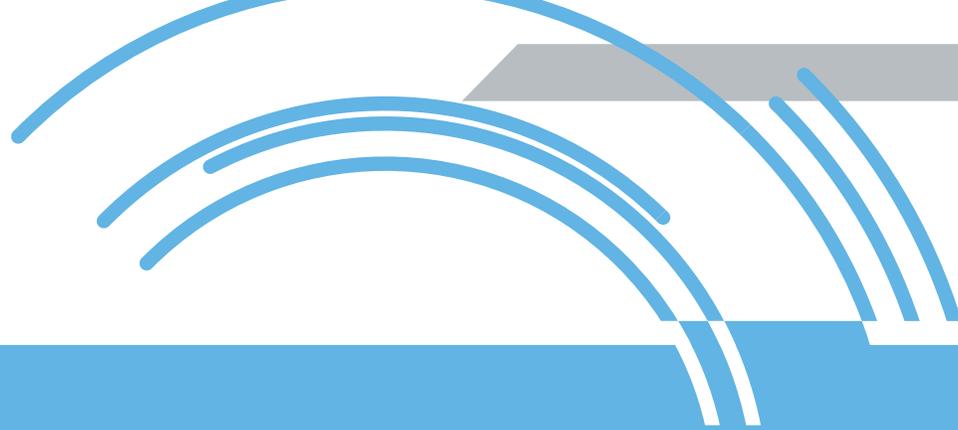
### DIRECTOR'S DISCRETIONARY (DD)

The ALCF's DD program provides "start up" awards to researchers working toward an INCITE or ALCC allocation to help them achieve computational readiness. 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 hundreds of thousands to the low millions of hours.

Please visit [alcf.anl.gov/programs](http://alcf.anl.gov/programs) for more information on how to get an allocation at the ALCF.



A breakdown of how ALCF computing time is allotted among the three allocation programs.



# ESP TECHNICAL ACHIEVEMENTS

## EARLY SCIENCE PROGRAM: TECHNICAL ACHIEVEMENTS

The ALCF's Early Science Program (ESP) was established to use the critical pre-production time period (between system installation and full production) on Mira to prepare key applications for the architecture and scale of the new supercomputer and to solidify the necessary libraries and infrastructure. When the program launched in 2010, 16 projects were selected to participate and awarded a combined two billion core-hours on Mira.

The ESP projects span a diverse range of scientific fields, numerical methods, programming models, and computational approaches, representing a large portion of the ALCF's current and projected computational workload. These projects took advantage of the Early Science period to pursue exciting new science, while characterizing the behavior of new hardware and software features.

Even 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. 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 became fully operational. A summary of these technical achievements follows.



## Ab Initio Reaction Calculations for Carbon-12

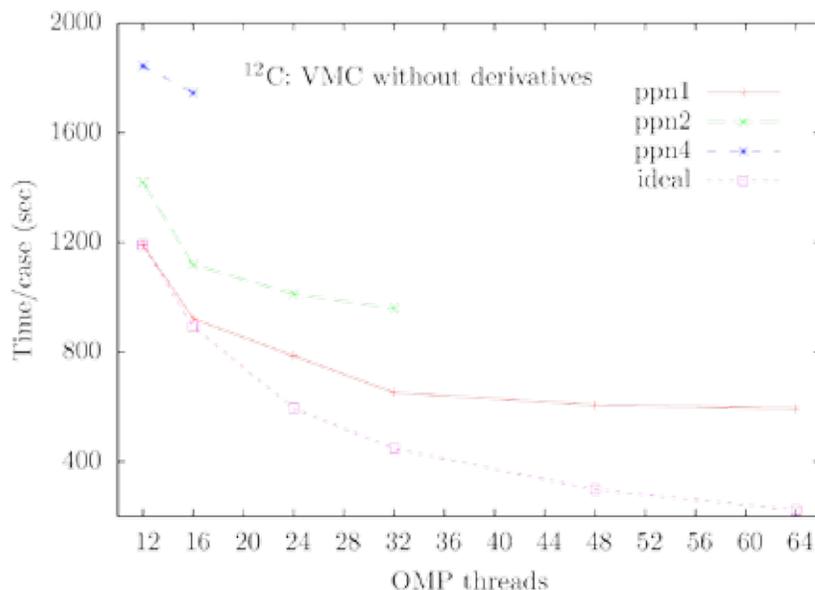
Steven Pieper | Argonne National Laboratory

Award: 110 Million Core-Hours

*Researchers from Argonne, Los Alamos, and Jefferson national laboratories, led by Steven Pieper, are calculating several fundamental properties of the  $^{12}\text{C}$  nucleus to describe quasi-elastic electron scattering data and to enable more reliable neutrino detector calibrations.*

For this ESP project, researchers set out to revise the Argonne Green's Function Monte Carlo (GFMC) code to take advantage of Mira's unique capabilities to compute sum rules and the response to weak probes of  $^{12}\text{C}$ . Both the direct calculation of neutrino scattering to the  $\Psi_{T=1}$  state and the evaluation of the sum rules and response would not have been possible on Intrepid due to the limited amount of RAM per node (2 GB). However, with 16 GB of RAM per node, Mira enabled the team to perform sufficiently large calculations.

The team employed the general purpose Automatic Dynamic Load Balancing (ADLB) library (originally developed on Intrepid) in the code and found it worked extremely well on Mira, with no modifications required. Furthermore, the team developed major new subroutines for the evaluation of the response functions and sum rules; these subroutines are multi-threaded using OpenMP and distribute their work over many nodes using ADLB. The conversion to Mira also consisted of improving the already existing OpenMP sections of the GFMC code to enable the efficient use of up to 64 threads per MPI process.



OpenMP strong scaling performance for sum rule calculation with two values of momentum transfer. The total wall time for 32 MPI ranks is shown.

**Image Credit**

Alessandro Lovato and Steven Pieper, Argonne National Laboratory

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

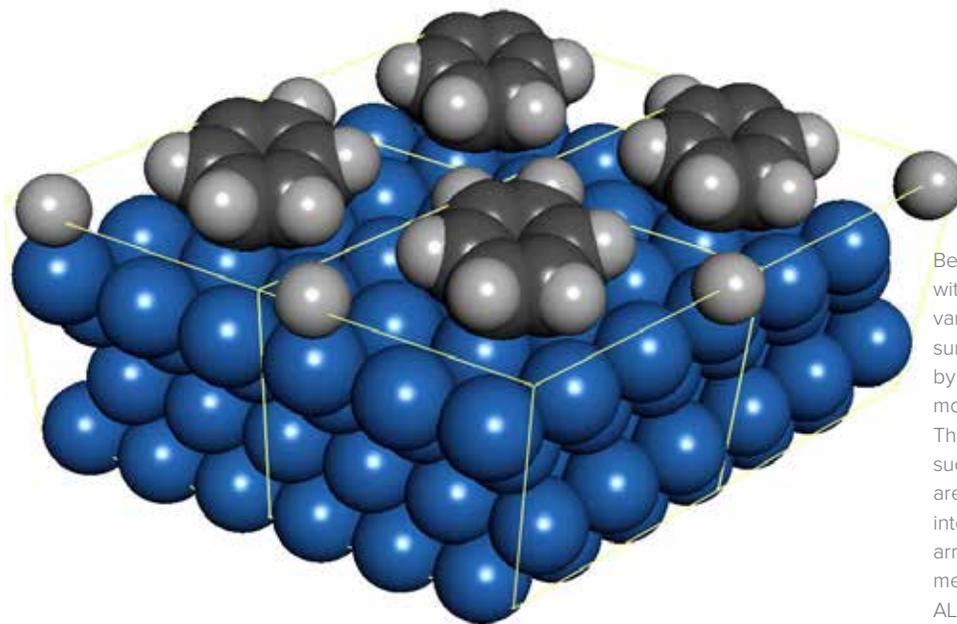
Robert Harrison | Stony Brook University

Award: 150 Million Core-Hours

*A team of researchers led by Robert Harrison from Stony Brook University is using ALCF resources to create computational tools to study complex quantum mechanical systems where systematic and robust numerical accuracy is required.*

This project uses the MADNESS simulation framework to study complex problems in chemistry, materials science, and nuclear physics by enabling multiresolution, adaptive numerical representations. Team members ported MADNESS to the Blue Gene/Q architecture by exploiting its portable programming model, which provides a dynamic and asynchronous runtime using MPI and Pthreads. The most computationally intensive routines were optimized using automatic code generation to achieve more than 50% of peak in mathematical kernels. New C++ linear algebra libraries such as Elemental and Eigen were incorporated to improve scalability. As a result of all these developments, MADNESS is running efficiently on thousands of nodes of Mira.

The team has also incorporated a new library for exchange-correlation potentials for density functional theory, implemented molecular geometry optimizers, and developed a prototype for linear response calculations, all of which are essential components in a quantum mechanical simulation code for use in atomistic simulation. The generality of this framework is demonstrated by the simultaneous application of these codes to quantum systems ranging from atomic nuclei to hundreds of water molecules.



Benzene on Pt(111) surface studied with density functional theory with van der Waals corrections. Most surface-catalyzed reactions are driven by noncovalent interactions of organic molecules and active metallic faces. The deposition of 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**

Alvaro Vazquez-Mayagoitia, Argonne National Laboratory

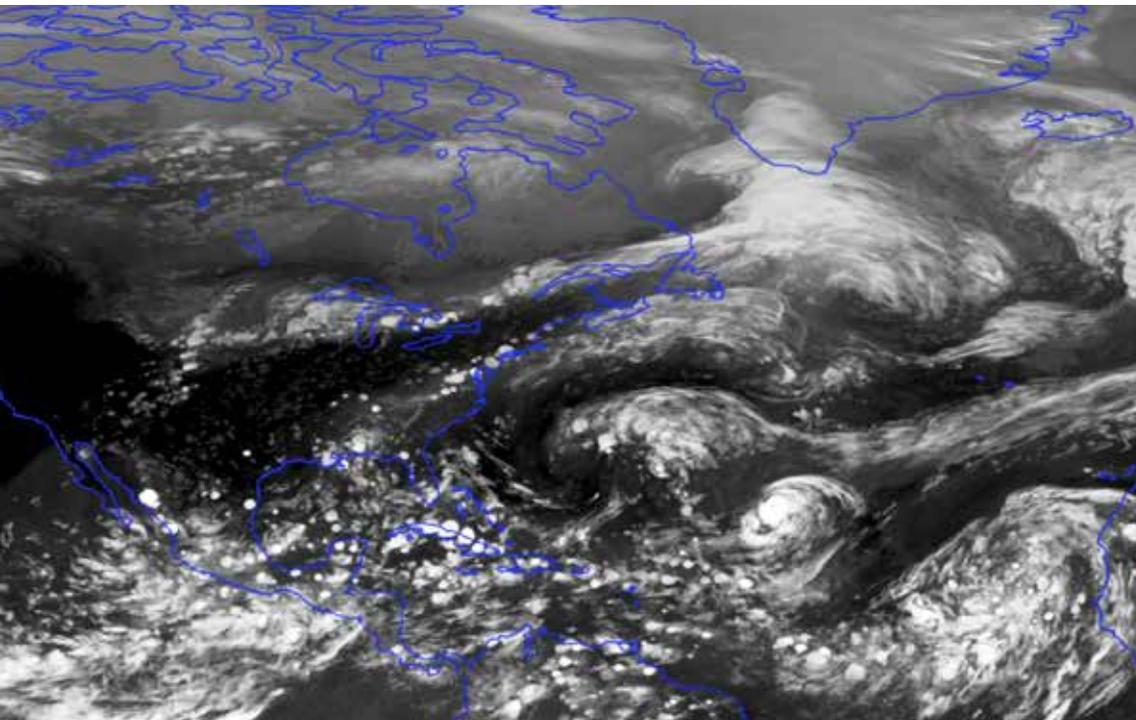
## Climate-Weather Modeling Studies Using a Prototype Global Cloud-System Resolving Model

Venkatramani Balaji | Geophysical Fluid Dynamics Laboratory

Award: 150 Million Core-Hours

*Researchers from the National Oceanic and Atmospheric Administration's Geophysical Fluid Dynamics Laboratory (GFDL) used the Early Science period to explore the frontier of weather prediction and climate modeling with their newly developed GFDL global cloud-resolving model.*

To tune the GFDL model, researchers performed experiments from the active 2005 hurricane season in the North Atlantic and Pacific basins and used those diagnostics to improve the model—a process that took six months and resulted in extensive changes to dynamics and physics schemes in the model. With early access to resources and expertise provided through the ESP, the team was able to refine the model resolution from their Atmospheric Model Intercomparison Project from 25 to 3.5 km. A significant amount of development time was dedicated to improving the computational and I/O performance of High-Resolution Atmospheric Model (HiRAM) and enhancing the post-processing diagnostic packages for the high-resolution experiments. Overall, the ESP work has resulted in an 8x improvement of performance per node from the Blue Gene/P to the Blue Gene/Q, with demonstrated strong scaling. In addition, the code now scales to 1 million hardware threads and is tuned to run on other leadership-class systems.



The figure shows the outgoing longwave radiation for August 14, 2005 from the global cloud-resolving HiRAM run at 3.5 km resolution.

**Image Credit**  
Venkatramani Balaji, Geophysical Fluid Dynamics Laboratory

## Cosmic Structure Probes of the Dark Universe

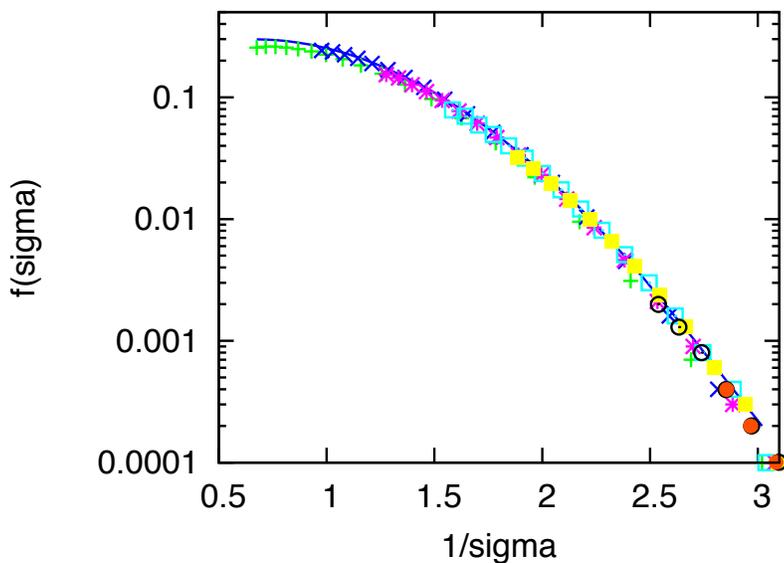
Salman Habib | Argonne National Laboratory

Award: 150 Million Core-Hours

*Led by Salman Habib of Argonne National Laboratory, this ESP project used Mira to perform some of the largest-scale simulations of the universe ever performed.*

To enable their groundbreaking studies of the universe, researchers developed a new framework called HACC (Hardware/Hybrid Accelerated Cosmology Code). 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. To get optimal performance on Mira, it was necessary to create as much data locality as possible, decompose the problem in a balanced way (both among nodes and among threads on each node), and then map the problem of interest onto what most modern machines do well: fused multiply-add operations.

For the most expensive part of the simulation, the short-range force calculation, the team created a customized hierarchical force approximation algorithm (the RCB tree) that maximized data locality and created balanced partitions of the overall problem. The researchers then used the force kernel, based on a polynomial fit to the short-range force formula. Polynomial evaluations use mostly fused multiply-add instructions, and so this mapped the problem onto what Mira does well. Work during the Early Science period also included integrating techniques from GLEAN into HACC and other applications, which achieved throughput of 170 GB/s. With all of the performance improvements, the team attained nearly 14 petaflops of sustained performance at 69% of peak and 90% parallel efficiency on more than 1.5 million cores, setting a benchmark for the largest cosmological simulation to date. The speedup achieved with HACC is enabling time-to-solution improvements for all science runs.



Differential mass function of dark matter halos in the universe measured from the “Outer Rim Simulation,” the largest cosmological simulation ever performed. The plot shows results starting at a redshift  $z=10$  when the universe was only  $\sim 0.5$  Gyr old until  $z=1$  when the universe was  $\sim 6$  Gyr (today the universe is 13.7 Gyr). Nearly 1.1 trillion particles are evolved in a cosmological volume of  $(4.225 \text{ Gpc})^3$ . A simulation of this size is required for researchers to extract detailed measurements of the structures in the universe over such a long time span. Enabled by work performed during the ESP period, these simulations continue to run on Mira under a 2013 INCITE allocation.

**Image Credit**  
Suman Bhattacharya and Katrin Heitmann, Argonne National Laboratory

Bhattacharya et al. fit ————  
 $z=1$  +  
 $z=5$  ■  
 $z=8$  ●  
 $z=10$  ▲

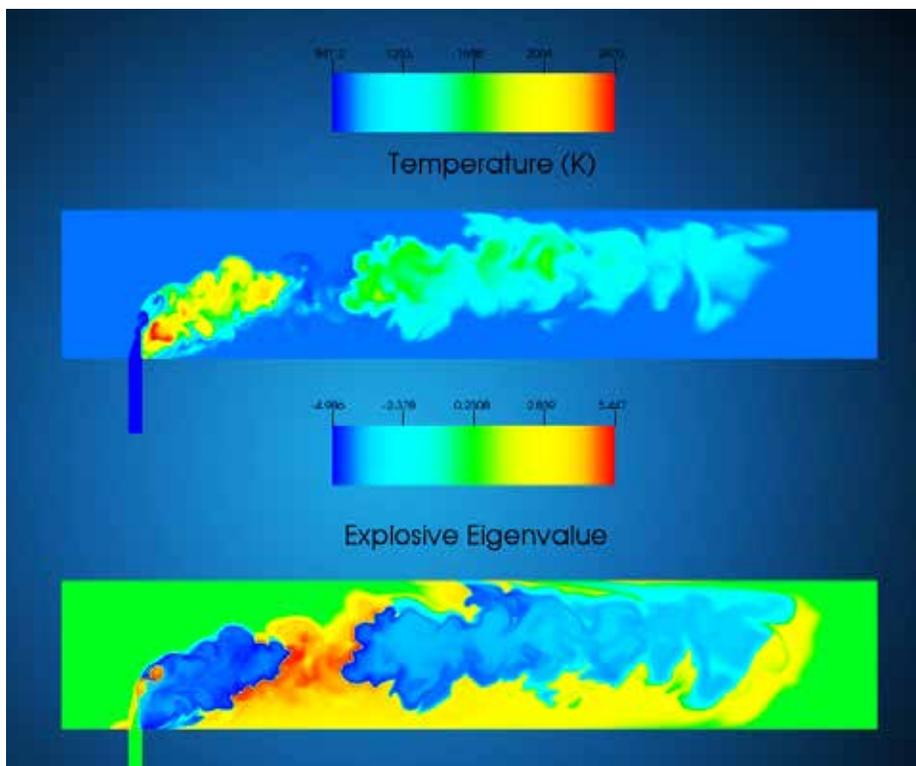
## Direct Numerical Simulation of Autoignition in a Jet in a Cross-Flow

Christos Frouzakis | Swiss Federal Institute of Technology Zurich

Award: 150 Million Core-Hours

*Led by Christos Frouzakis of the Swiss Federal Institute of Technology Zurich, this ESP project is aimed at studying the fundamental aspects of autoignition in a fuel-air mixing pattern directly applicable to mixing ducts in gas turbines.*

Researchers tapped Mira to perform exploratory numerical simulations on experimental computational grids as well as production runs during the Early Science period. They used the NEK5000-based code for low Mach number reactive flows to investigate autoignition in a laboratory-scale jet in a cross-flow configuration (JICF) using large-scale direct numerical simulation. The code relies heavily on two highly optimized mxm and thermochemistry kernels for evaluating chemical species production rates, transport, and thermodynamics properties. Scaling results on Mira proved to be successful for the current problem size, showing that ideal parallel efficiency of 100% could be sustained up to 130,000 MPI ranks and up to 60% efficiency on nearly 500,000 ranks. The team was able to run two full simulations at the lower friction Reynolds number of 180 (for two cross-flow temperatures: 930 K and 950 K). Further optimization of the algebraic multigrid solver initialization routines was found to be necessary for tackling high Reynolds number simulations on the proposed 2.5-million element grid. The team continues to investigate the high Reynolds number case using higher local order of polynomial approximation (p refinement) instead of increasing the number of elements (h refinement).



A 2D slice at the mid spanwise plane of a hydrogen JICF simulation for air cross-flow temperature of 950 K and a friction Reynolds number of 180. Two distinct burning regions can be identified from the top/temperature panel: a spatially coherent flame in the vicinity of the fuel nozzle and spatially localized random ignition spots further downstream. Significant insight into the reactivity of the mixture can be learned from the distribution and magnitude of the explosive eigenvalue obtained by Chemical Explosive Mode Analysis (CEMA).

**Image Credit**

Ammar M. Abdilghanie, Argonne National Laboratory

## Global Simulation of Plasma Microturbulence at the Petascale and Beyond

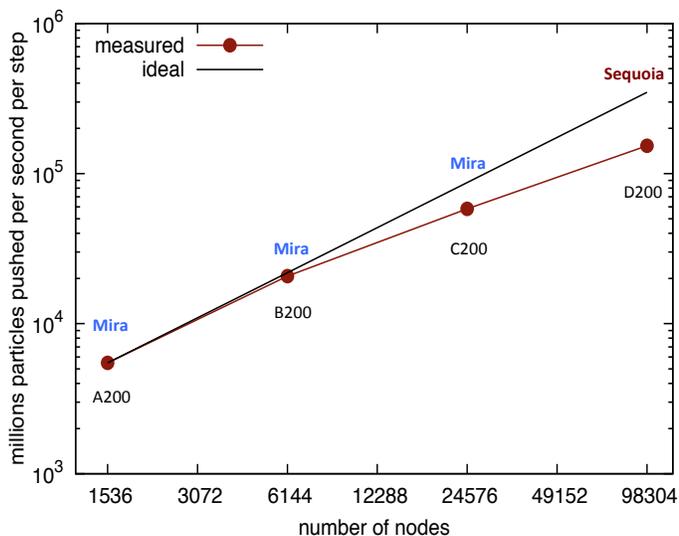
William Tang | Princeton Plasma Physics Laboratory

Award: 50 Million Core-Hours

*William Tang and his research group at the Princeton Plasma Physics Laboratory with collaborators at the Lawrence Berkeley and Argonne national laboratories set out to advance the understanding of plasma microturbulence to help accelerate the design and operation of future nuclear fusion energy devices.*

Fusion energy physics spans an extreme range of time and spatial scales, requiring associated simulations to exploit local concurrency to take full advantage of Mira and other powerful supercomputing systems. For this ESP project, researchers developed a C-version of the lead code GTC-P (originally written in Fortran-90) to enable better compatibility with forefront computer science advances. Consequently, GTC-P-C readily exploited, for example, new multi-threading methods for low memory-per-core systems to significantly improve performance on Mira.

In order to run efficiently on the Blue Gene/Q's large number of nodes, the team mapped MPI ranks onto the physical nodes and cores of the machine and optimized the code to use OpenMP efficiently to achieve major performance gain on the system's highly multi-threaded architecture. Together with an improved load-balance approach, the team was able to demonstrate excellent scalability to 524,288 cores (two-thirds of Mira) during the ESP development phase for this modern code. Overall, moving GTC-P from the Blue-Genie/P to the Blue-Genie/Q resulted in a greater-than-10 gain in time to solution, while further software improvements produced an additional 50% gain. In terms of practical impact, this improvement in carrying out high-resolution production runs on large plasma systems (e.g., ITER) represents a significant step forward in the fusion energy science application domain.



The GTC-P C code enabled researchers to take full advantage of the Blue Gene/Q's highly multi-threaded nodes and large scalable interconnect, allowing them to achieve excellent scaling on both Mira and Sequoia (the 20-petaflops Blue Gene/Q system at Lawrence Livermore National Laboratory). The graph shows the number of nodes required for increasing plasma system sizes, from current laboratory experiments to the largest device under construction (with D being ITER).

#### Image Credit

Stephane Ethier and William Tang, Princeton Plasma Physics Laboratory; Bei Wang, Princeton University

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

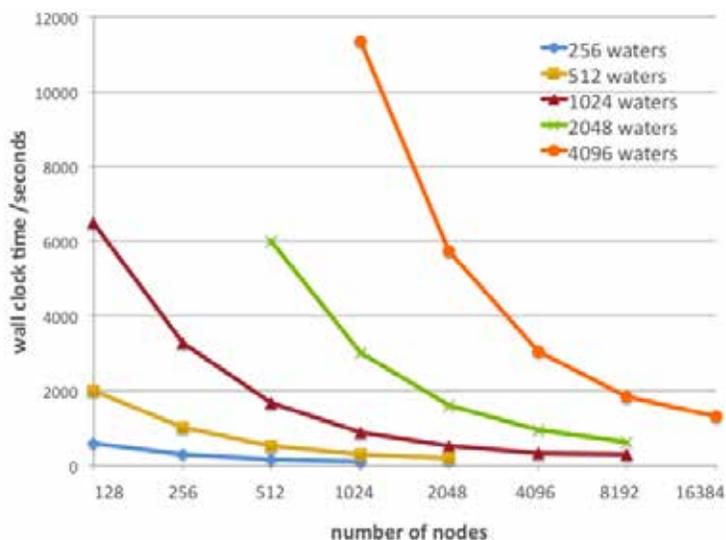
Mark Gordon | Iowa State University

Award: 150 Million Core-Hours

*Led by Mark Gordon of Iowa State University, this project broke new ground with the first high-accuracy, high-precision models to employ advanced techniques that allow ab initio quantum mechanical methods to scale up to simulations of thousands of water molecules.*

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

The team also employed the RATTLE method into GAMESS to significantly reduce the cost of the project to one-half or one-quarter of its original cost. RATTLE is a Lagrange-multiplier-based method developed especially for Velocity Verlet, an algorithm for integrating Newton's equations implemented in GAMESS. Systems with large numbers of hydrogen atoms (such as water) can waste many time steps when the trajectory is trapped within local minima corresponding to the vibrational modes of all the hydrogen atoms. Implementing RATTLE into the code applies constrained dynamics to help smooth out the trajectory and to improve the overall efficiency by removing any abrupt shifts in molecule positions that may hinder the solution of the Schrödinger equation for each water molecule. In doing so, RATTLE has resulted in improved performance on Mira by allowing fewer, larger, and more efficient time steps to be taken.



GAMESS performance on Mira showing scalability of FMO2-MP2 (forces)/aug-cc-pvdz calculations on 256 to 4096 water clusters in c16 mode (16 MPI ranks per node).

**Image Credit**

Yuri Alexeev, Graham Fletcher, and Maricris Mayes, Argonne National Laboratory

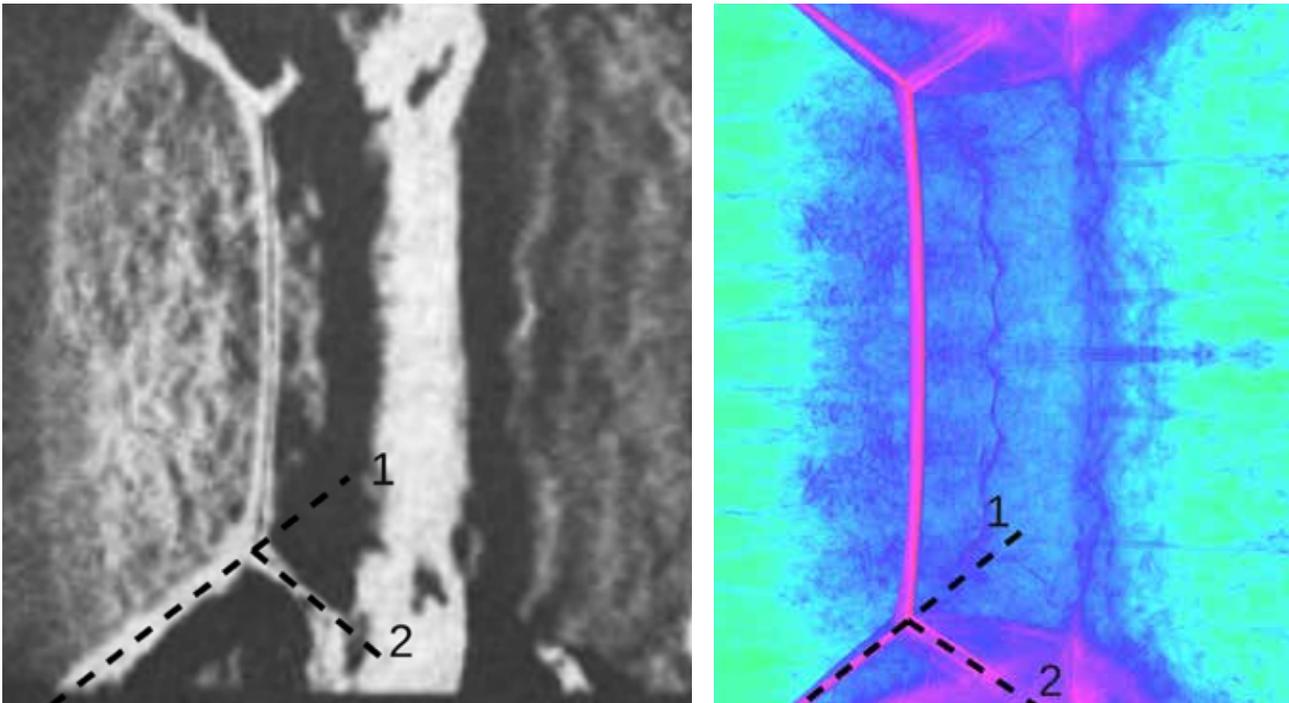
## High-Speed Combustion and Detonation

**Alexei Khokhlov | The University of Chicago**

Award: 150 Million Core-Hours

*Alexei Khokhlov and his research team from the University of Chicago explored 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 capitalized on the Early Science period to refine the threading model and improve single-core performance. Tuning certain data exchange and balance loops has helped to optimize the code for the Blue Gene/Q system. The 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 32,000 cores. Overall, this early access to Mira has enabled modeling with greater accuracy, and a 2.5x performance per core improvement over efforts on Intrepid. 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.



Shock bifurcation in CO<sub>2</sub>. Calculations including both viscosity and heat conduction. Left - experiment. Right - simulation. Shock is moving to the left.  
**Angle 1:** experiment 36, simulated 37    **Angle 2:** experiment 38, simulation 34

### Image Credit

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

## Lattice QCD - Early Science

**Paul Mackenzie | Fermilab**

Award: 150 Million Core-Hours

*Led by Paul Mackenzie of Fermilab, this ESP project sought to produce the high-precision lattice quantum chromodynamics (QCD) calculations that are urgently needed in the analysis of crucial experiments in high energy and nuclear physics.*

Lattice QCD is a numerical method used to simulate QCD including non-perturbative effects. It plays an important role in high energy particle physics and nuclear physics, but requires enormous computing resources. With Mira's unique capabilities, researchers can advance Lattice QCD code to generate configurations that would support calculations of the spectrum, decay properties and internal structure of strongly interacting particles, and tests of the standard model, of unprecedented precision.

To fully leverage the power of the Blue Gene/Q system, researchers used the Early Science period to make several improvements to the code. They developed a new communication library (qspi) using the Blue Gene/Q's low-level communications interface, resulting in approximately 20x less latency than achieved with normal MPI in 5D halo exchange tests. The team's optimal node mapping strategy helped to reduce the surface volume at the boundaries and the number of communication hops. In addition, researchers reduced the number of memory fragmentations in the communications by applying the optimal site ordering. Moreover, using QPX in QLA (linear algebra library) resulted in significant speedup. Overall, they achieved a 2.3x speedup compared to performance of the code on Mira before the ESP enhancements were made.

## Materials Design and Discovery

**Larry Curtiss | Argonne National Laboratory**

Award: 50 Million Core-Hours

*Researchers from Argonne National Laboratory paired the power of the Blue Gene/Q with newly available electronic structure codes to conduct massively parallel quantum mechanical calculations for use in the design of breakthrough materials.*

Initially aimed at performing fast-accurate density functional theory calculations on materials for energy storage and catalysis, the research team shifted the project's focus after determining intrinsic algorithmic limitations. Instead, the researchers decided to pursue Quantum Monte Carlo (QMC) as a complimentary method on Mira, optimizing a valuable tool for the scientific community.

Taking advantage of Mira's massive parallelism, QMC proved to be successful in using a large number of cores to enable the study of very complex systems at chemical accuracy. The ESP work focused on porting the code to Mira, optimizing the main kernel using QPX, and prefetching. These efforts resulted in a 2.67x speedup over the original algorithms, offering a substantial improvement in time to solution. Working with the QMCPACK simulation package provided a breakthrough that enabled researchers to study a larger spectrum of materials at chemical accuracy for many applications including materials design and biochemistry. Further work with QMCPACK on Mira is expected to continue in the future, including the implementation of nested OpenMP, which could potentially improve the time to solution by 4x.

## Multiscale Molecular Simulations at the Petascale

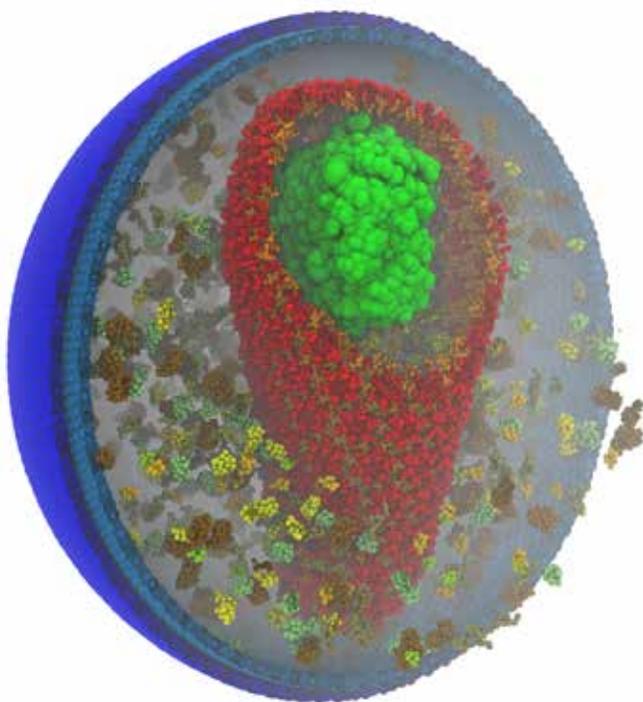
**Gregory Voth | The University of Chicago**

Award: 150 Million Core-Hours

*Gregory Voth and his team at the University of Chicago and Argonne National Laboratory aimed to further our understanding of cellular-scale biological processes via coupling of multiscale computer simulation methodologies with petascale computational algorithms and hardware.*

For this ESP project, researchers transformed a poorly scaling RAPTOR code (add-on package for reactive molecular simulations in LAMMPS software) into one that could harness half of Mira (24 racks) with appreciable parallel efficiency. This was accomplished by introducing multi-threading via OpenMP, developing new state decomposition parallel algorithms, and creating a novel implementation of replica exchange umbrella sampling with the LAMMPS Ensembles code. These new, highly scalable simulation codes are now being used on Mira to compute free energy surfaces of proton transport in complex biological systems at an unprecedented rate.

The algorithmic advances made possible with this ESP project are sufficiently general that they are also being used in an INCITE project to address challenging questions related to energy storage technologies. Targeted for larger-scale biological processes, a new 2D “Hot-Bead” Hamiltonian replica-exchange algorithm was implemented in NAMD. In this algorithm, domains of a protein are defined as “Hot-Beads” within which interactions are scaled down, facilitating local conformational changes. When “Hot-Bead” replicas are combined with temperature exchange, the resulting reduction in free energy barriers has enabled large-scale conformational changes not yet observed with standard brute-force techniques. To access larger scale phenomena, novel cellular-scale simulation software has been developed to specifically handle dynamic, sparse systems with an efficient load-balancing strategy and low memory footprint. This software is being used in studies of the retroviral maturation pathway.



Coarse-grained representation of encapsulated HIV virion (red) with RNA (green) used in cellular-scale molecular simulations.

**Image Credit**

John Grime, Argonne National Laboratory/The University of Chicago

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

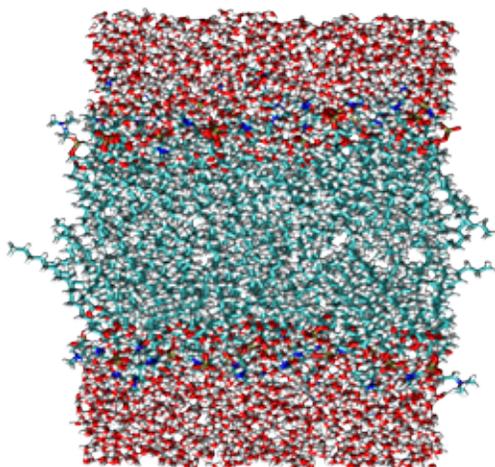
**Benoit Roux, The University of Chicago**

80 Million Core-Hours

*A research team led by University of Chicago biophysicist Benoît Roux applied advanced strategies based on large-scale molecular dynamics (MD) simulations to create virtual models of biomolecular systems with unprecedented accuracy.*

NAMD is currently one of the most optimal and efficient tools used to carry out MD simulations of biomolecular systems, but the program traditionally only supported brute-force MD simulations of a single trajectory. A new parallelization paradigm able to drive multiple concurrent MD trajectories with fine-grained inter-trajectory communication was needed. To address this issue, the team aimed to develop an extremely scalable generic replica exchange MD (REMD) method driven by the asynchronous programming system Charm++. This work allowed NAMD to take full advantage of Mira and other leadership-class systems.

Researchers successfully implemented REMD in NAMD with multiple MPI sub-communicators in Charm++. However, such MPI level implementation of Charm++ can lose some network topology information and therefore result in extra network contentions or latency. Alternatively, within the Charm++ Converse layer, each Charm++ processing element is logically mapped onto a designated local partition (instead of MPI\_COMM\_WORLD). When NAMD instances enter inter-copy communication phase, all localized processing elements are mapped back to global state. As the Converse layer is the interface with the underlying machine-specific communication layer, the new implementation can obtain further performance gains from Parallel Active Messaging Interface (PAMI), the low-level communication library on the Blue Gene/Q system. With the PAMI-version MCAs on Mira, researchers obtained better stability and more than 20% speedup compared to the previous MPI version. Already, several other ALCF research projects involving biological systems are benefitting from these improvements to NAMD.



Typical configuration of a zwitterionic lipid bilayer membrane using a novel polarizable force field based on the classical Drude oscillator model for molecular dynamics simulations [Chowdhary et al. (2013). *J. Phys. Chem. B.* 117(31):9142-60. doi: 10.1021/jp402860e].

#### Image Credit

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

## Petascale, Adaptive CFD

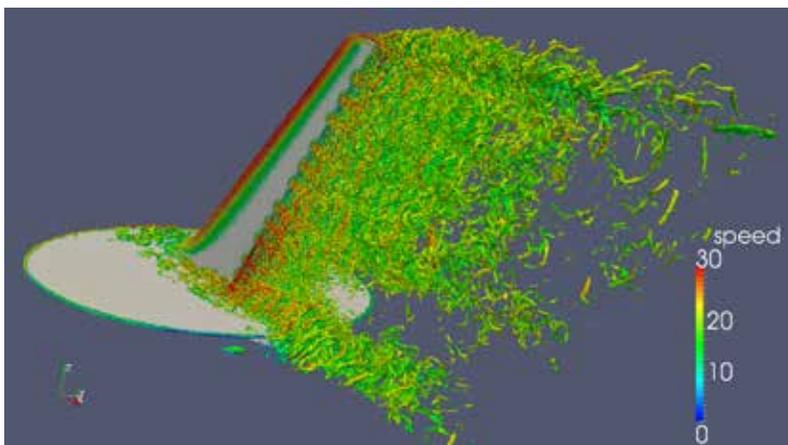
**Kenneth Jansen | University of Colorado Boulder**

Award: 150 Million Core-Hours

*Kenneth Jansen and his research team at the University of Colorado Boulder performed simulations of active flow control based on synthetic jet actuation at a high Reynolds number, approaching engineering application scales for the first time.*

Prior to Mira, researchers used Intrepid, the ALCF's Blue Gene/P supercomputer, to perform simulations with up to 500 million finite elements for this work. With access to Mira through the ESP, they had the ability to simulate flow control on a realistic aerodynamic configuration with billions of finite elements (typically around 5 billion). Because their models must resolve much of the flow's turbulence, a very fine resolution is needed. This requires the use of a very large number of elements, which leads to a significant amount of equations that must be solved. Mira was the first resource with enough processors and sufficient memory to make these computations feasible. In that context, the team also managed to push the limits of their workflow by running successful flow control simulations with a 92 billion unstructured finite element mesh on the full Mira system.

To take full advantage of the Blue Gene/Q architecture, the researchers used the Early Science period to optimize the PHASTA (Parallel Hierarchic Adaptive Stabilized Transient Analysis) code for Mira. Work included increasing from one MPI process per core to two or four MPI processes per core, which resulted in a significant performance improvement. Additionally, the team developed new algorithms to enable unstructured mesh partitioning associated with complex geometries at a large scale. They also developed the capacity to perform live, interactive data co-visualization, which has the potential to provide continuous and reconfigurable insight into massively parallel simulations, paving the way for interactive simulation and simulation steering.



Full span view of a vertical tail assembly of a commercial aircraft at  $Re=3.6 \cdot 10^5$ , highlighting the root and tip vortex along with the turbulent structures in the wake of the deflected rudder through isosurface of instantaneous Q criterion colored by speed on a 1.4 billion unstructured finite element mesh.

**Image Credit**

Michel Rasquin, Argonne National Laboratory

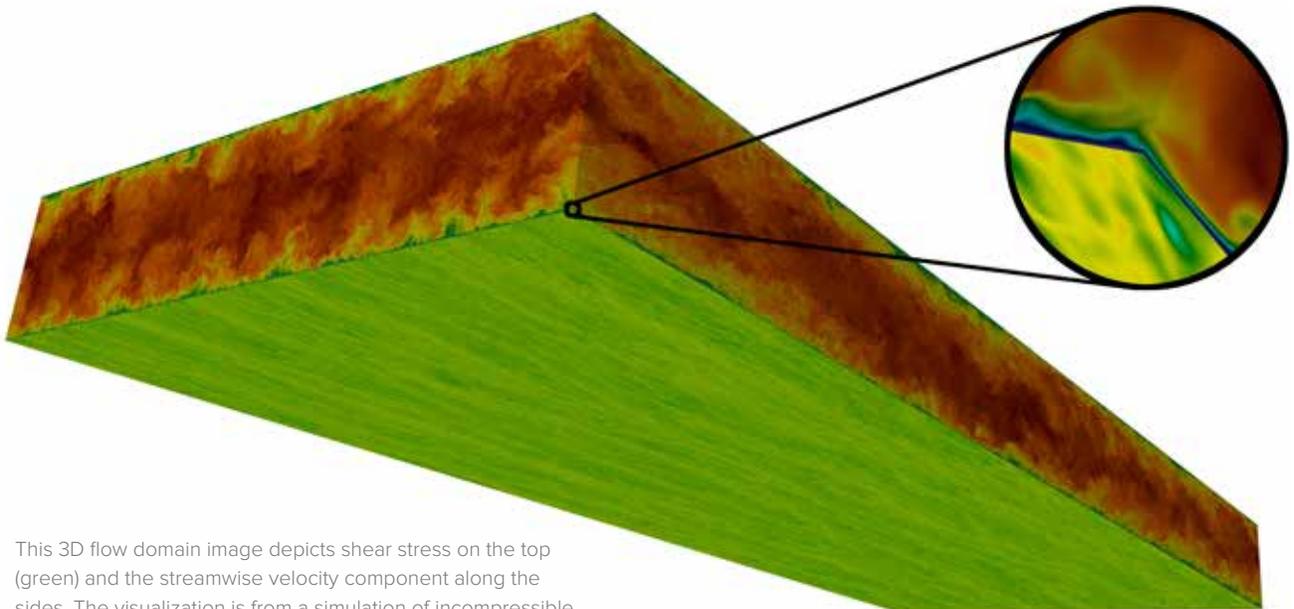
## Petascale Direct Numerical Simulations of Turbulent Channel Flow

Robert Moser | University of Texas at Austin

Award: 60 Million Core-Hours

*Researchers from the University of Texas at Austin led by Robert Moser have conducted the largest-ever production direct numerical simulation to investigate the physics of wall-bounded turbulence.*

During the Early Science period, researchers developed an entirely new codebase to fully exploit Mira's unique hardware features for their turbulence research. At the core level, the team achieved efficient memory access by hand unrolling and fusing loops to improve cache reuse. At the node level, implementing OpenMP threading with loop fusion techniques allowed them to maximize the size of threaded blocks. Mira's threading, cache, and memory characteristics enabled the researchers to employ a hybrid OpenMP/MPI model to take advantage of the natural concurrence in their algorithm. This provided flexibility in on-node execution resulting in efficient management of cache and execution threads. By minimizing the inter-memory access between OpenMP threads, they achieved near-perfect OpenMP scalability (99%). At the system level, they were able to take full advantage of the 5D torus network by replacing the existing library for 3D global Fast Fourier Transforms (P3DFFT) with a new library they developed using the FFTW 3.3 communication library. Together, all of this work resulted in a 2x performance increase compared to the old code. Without these performance improvements, the groundbreaking DNS would not have been possible.



This 3D flow domain image depicts shear stress on the top (green) and the streamwise velocity component along the sides. The visualization is from a simulation of incompressible turbulent flow ( $Re_\tau = 5200$ ) between two parallel smooth planes. Simulations of wall-bounded turbulence such as this are analogous to an experimental laboratory, permitting scientists to probe the interaction between a vehicle (or a pipe) and the fluid flowing past it, transferring momentum and determining the drag.

**Image Credit**

Myoungkyu Lee, University of Texas at Austin

## Petascale Simulations of Turbulent Nuclear Combustion

---

**Don Lamb | The University of Chicago**

Award: 150 Million Core-Hours

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

Researchers made many changes to the FLASH code to increase its robustness and performance on the Blue Gene/Q system for this ESP project. They multi-threaded the hydrodynamics solver, flame model, and Equation of State (EOS), which are expensive code sections in FLASH ESP and INCITE applications. This enabled FLASH to make better use of the increased on-node parallelism of Blue Gene/Q relative to Blue Gene/P. The team experimented with placing OpenMP directives at two different levels of granularity in the code and obtained best performance when using finer-grained multi-threading. The hybrid MPI+OpenMP configuration improved time to solution over an MPI-only version of FLASH and is now used in all production FLASH simulations on Blue Gene/Q. Researchers also improved serial performance by reordering arrays in the unsplit hydrodynamics solver and EOS. The new data structures improve access locality and avoid the creation of Fortran array temporaries in subroutine calls. These serial optimizations improved the performance of a fully optimized, hybrid MPI+OpenMP FLASH simulation by 1.5x.

Additionally, the ESP team created Fortran wrappers around mallinfo and Kernel\_GetMemorySize (part of the System Programming Interface on BG/Q) to monitor memory use during an application run. This capability helped identify two prohibitive memory leaks. The first memory leak happened within a communication-heavy initialization subroutine in FLASH. There was no leak at the application layer, but the team used the situation it uncovered as motivation to decrease the amount of communication in the subroutine, thereby reducing memory growth and improving scalability. The researchers modified the subroutine so that communication would occur between significantly fewer MPI ranks in a reduced communicator. This circumvented the communication layer memory growth and reduced initialization time from minutes to seconds. The team also added new metadata to FLASH checkpoint files to avoid the subroutine call entirely when restarting a run. The second memory leak correlated with the writing of FLASH checkpoint and plot files. The team worked with ALCF staff to identify the leak and then verified that a fix to the ROMIO Blue Gene driver stopped the leak. A forthcoming driver upgrade on Mira will enable the largest FLASH simulations of Type Ia supernovae to date to run for the maximum job runtime.

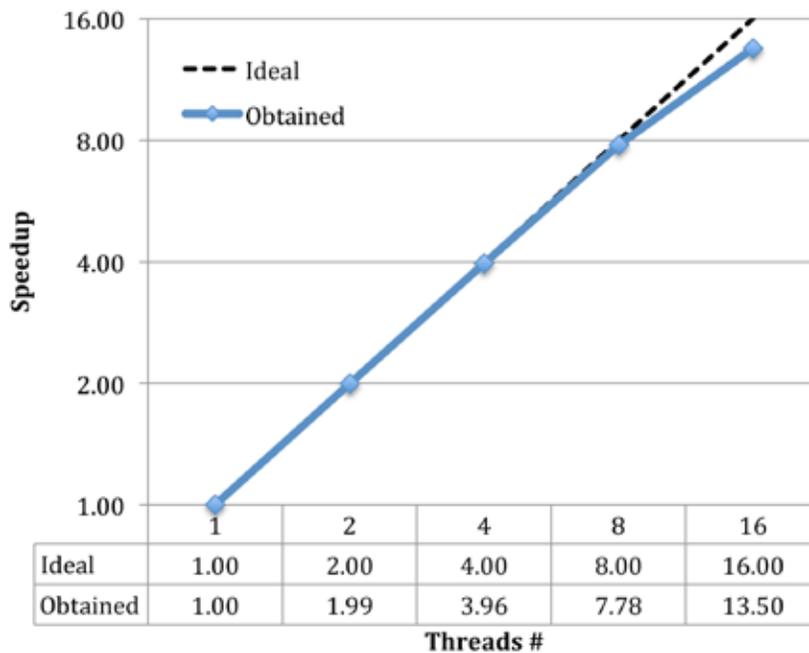
## Using Multiscale Dynamic Rupture Models to Improve Ground Motion Estimates

Thomas Jordan | University of Southern California

Award: 150 Million Core-Hours

*Researchers at the Southern California Earthquake Center at the University of Southern California used their ESP project to perform dynamic rupture simulations to investigate high-frequency seismic energy generation.*

To enable high-resolution simulations of the dynamic rupture processes inside of earthquakes, the research team uses the Support Operator Rupture Dynamics (SORD) code. During the ESP period, they achieved nearly ideal weak scaling of SORD on Mira in pure MPI mode. The team introduced OpenMP multi-threading for the kernels in SORD, and achieved nearly ideal strong scaling up to 16 threads per node. ESP work also included implementing a cache-tilting scheme to make better reuse of arrays in the cache, helping to overcome memory bandwidth bottlenecks. Taken with the increased core count and bus speed for the Blue Gene/Q system, their efforts resulted in a total of 20x per-node speedup compared to performance on the Blue Gene/P system.



SORD OpenMP strong scaling benchmark for single node Blue Gene/Q.

Image Credit  
Geoffrey Ely, Argonne National Laboratory



# SCIENCE HIGHLIGHTS

GEORGE SCHATZ

Northwestern University | schatz@chem.northwestern.edu

## Computational Studies of Nucleosome Stability

*Modeling the mechanical properties of DNA based on structures with 100 or more base pairs is now feasible through simulations conducted on ALCF's IBM Blue Gene supercomputers, Intrepid and Mira. By understanding factors involved in the formation and stability of nucleosomes, the primary genomic DNA packaging in mammalian cells, researchers can begin to identify abnormalities that can be used as biomarkers in the diagnosis of genetic diseases.*

This INCITE research project led by George Schatz of Northwestern University uses molecular dynamics and electronic structure methods to study nucleosomes, which are complexes of DNA and proteins in chromatin that account for 75-90% of genomic DNA packaging in mammalian cells. Nucleosomes consist of 147 base pairs of DNA wrapped around a complex of proteins known as histone octamers, though research has typically focused on an 18-base-pair model. Understanding base-pair sequence dependence and formation propensities has long been a goal in cell biology, as nucleosome stability is thought to be an important component of transcriptional regulation, the transferring of genetic data from DNA to protein-producing RNA molecules.

In particular, DNA methylation and histone acetylation, which play key roles in epigenetics—external genetic forces—and have been associated with cancer, are important components of nucleosome structures and will be extensively studied.

High-performance computing (HPC) provides a significant opportunity for taking a major step toward the understanding of nucleosome binding. Using ALCF's IBM Blue Gene systems Intrepid and Mira, researchers will fundamentally advance computational modeling of nucleosome stability. Classical molecular dynamics calculations are now possible for whole nucleosome structures. Conducting such calculations using NAMD on an HPC system for a time (~1 microsecond) long enough to equilibrate

structures that include explicit water and ions would help determine useful information about nucleosome binding, including both sequence dependence effects and the effect of epigenetic markers. Electronic structure studies of smaller nucleosome models using the fragment molecular orbital (FMO) method will also be performed using GAMESS.

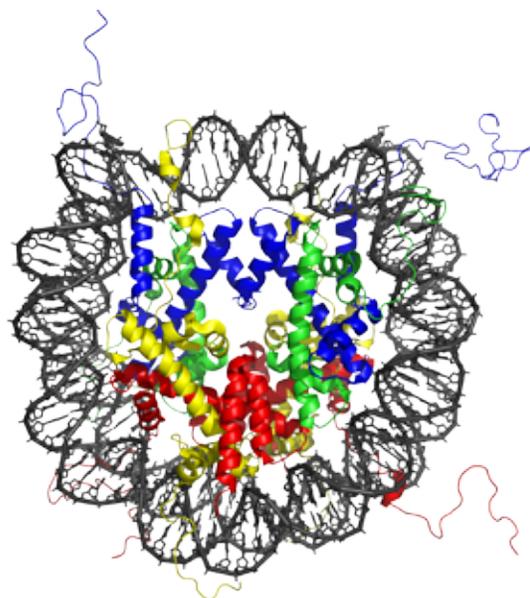


Image of the nucleosome 1KX5 from the Protein Data Bank. This DNA/protein complex will serve as the primary target of simulation studies to be performed by the Schatz group as part of the INCITE program.

Image Credit  
Xiao Zhu, University of Texas at Austin

**IMPACT** » This research will prove essential in developing coarse-grained models of nucleosome positioning in genomic DNA, and help elucidate the interplay and competition between nucleosome formation and transcription factor binding. It will also provide insight into the development of chromatin structure models, and the modifications in chromatin that arise both from natural epigenetic modifications and from epigenetic drugs that are becoming important in cancer therapy.

T. ANDREW BINKOWSKI

Argonne National Laboratory | abinkowski@anl.gov

## Protein Binding and Interaction Analysis of Human Pathogen Protein Targets

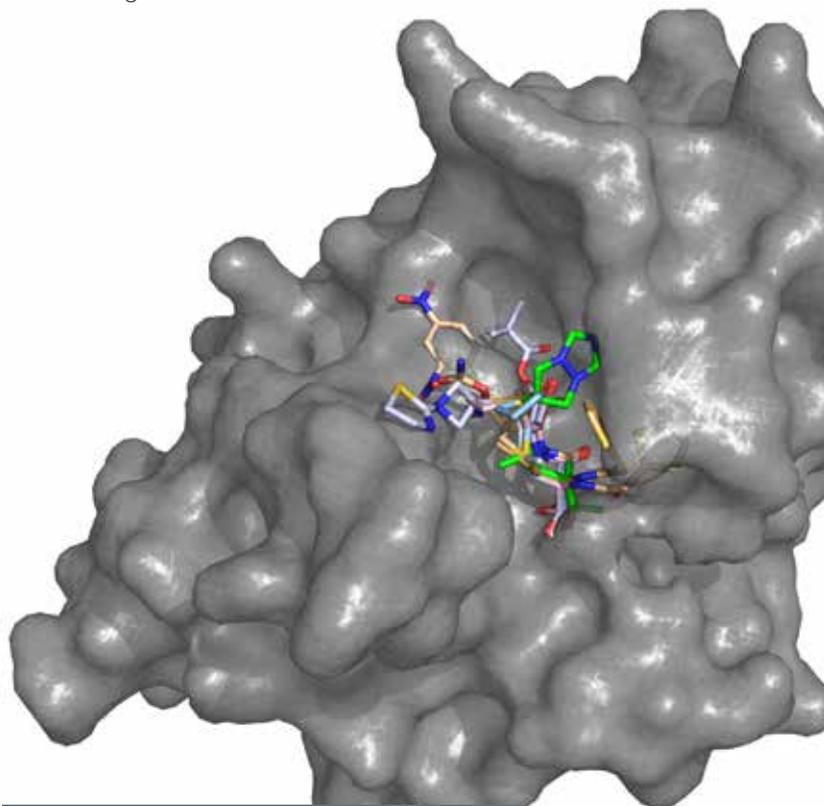
*Researchers have tried to keep pace with antibiotic resistant enzymes by continually producing larger and larger antibiotics, but some proteins are adapting faster than the pipeline for drug design, which can take as long as 20 years. Computer-aided drug design represents a substantial speedup over traditional experiments and is being significantly advanced through molecular simulation techniques being developed at the ALCF.*

For the past two years, Argonne structural biologist Andrew Binkowski has been using ALCF resources to develop an automated system to conduct a comprehensive analysis of protein binding domain and small molecule interactions. He and his team are also conducting computer-aided drug discovery of high-value biomedical targets, including human disease-related proteins. Once such protein, NDM-1, is a recently emerging threat that became known several years ago and has since been detected globally.

NDM-1 is a member of the beta-lactamase enzyme family characterized by their ability to hydrolyze or “cut” the lactam ring in common antibiotics. The antibiotics would normally kill bacteria by disrupting cell wall synthesis, but once cut they are unable to carry out this function. What makes NDM-1 so dangerous is that its cavity is particularly large and flexible so most antibiotics can enter the cavity and be rendered ineffective.

In 2012, researchers from Argonne’s Midwest Center for Structural Genomics (MCSG) and Structural Biology Center (SBC) began collaborating with researchers at Texas A&M who had been studying the enzyme, but had been unable to determine the 3D structure of the protein. Within a few months, the MCSG team solved the 3D structure of the enzyme at the Advanced Photon Source and deposited it in the Protein Data Bank. Binkowski, a member of MCSG, then used

10 million core-hours on ALCF resources over the course of one month to determine why NDM-1 made bacteria so drug resistant. By analyzing the structure and simulating how nine different antibiotic compounds bind to NDMs, he also discovered a flaw in the design of new antibiotics.



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

**IMPACT** » The design of an entirely new drug can take years, however, modifying an existing drug can reduce that time significantly. The computational pipeline at the ALCF is helping to better understand the mechanisms of target proteins to provide insights into developing strategies to combat public health threats.

BENOÎT ROUX

The University of Chicago | roux@uchicago.edu

## Studies of Large Conformational Changes in Biomolecular Machines

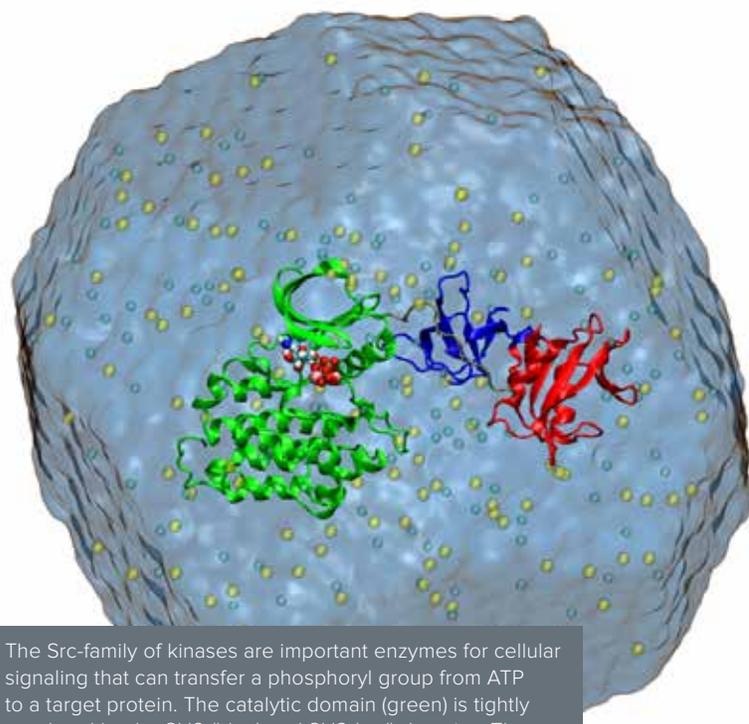
*Proteins are extremely complex “molecular machines.” Their concerted action triggers many of the critical activities occurring in living organisms. Membrane-associated proteins play an essential role in controlling the bidirectional flow of material and information. Malfunction of some vital proteins can lead to diseases such as cancer. This project is using ALCF supercomputers to advance our understanding of such protein functions.*

Molecular dynamics (MD) simulations provide a new perspective to understand the function of biologically important molecular systems. 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 developed by the University of Illinois, is one of the most optimal programs used to carry out classical simulations of biomolecular systems.

Through an Early Science Program project at the ALCF, University of Chicago biophysicist Benoît Roux and his colleagues enhanced the sampling efficiency of NAMD beyond that of brute-force MD simulations, and implemented several advanced strategies based on multiple copies such as temperature and Hamiltonian tempering replica-exchange MD (T-REMD and H-REMD). The team also implemented a new force field that incorporates the effect of induced polarization.

To further research using NAMD, Roux is using a 2013 INCITE award to conduct simulations at the ALCF to gain a deep mechanistic perspective of protein function, linking structure to dynamics by characterizing the free energy landscape that governs key functional motions. Src tyrosine kinases and the ATP-driven ion pumps will be studied within a unified computational perspective. As a benchmark for quantifying the accuracy of the approach, the researchers will also investigate the conformational

propensity of small peptides in solution. By studying experimentally well-characterized systems of increasing size and complexity within a unified theoretical framework based on free energy landscapes, this project will push the envelope and advance theory-modeling-simulation (TMS) technology. TMS offers a virtual route to address fundamental biological questions and help solve the problem of rational protein design.



The Src-family of kinases are important enzymes for cellular signaling that can transfer a phosphoryl group from ATP to a target protein. The catalytic domain (green) is tightly regulated by the SH2 (blue) and SH3 (red) domains. The active state is shown here with a large water box and a concentration of ion-mimicking biological conditions.

Image Credit

Avisek Das, Mikolai Fajer, and Benoît Roux, The University of Chicago

**IMPACT** » The enhancements to NAMD have provided an investigative tool of unprecedented accuracy for the entire scientific community, enabling studies of biological systems that were not possible in the past. Results from the INCITE project will serve as a roadmap for simulating, visualizing, and elucidating how biomolecular nano-machines work.

DAVID BAKER  
University of Washington | dabaker@u.washington.edu

## Towards Breakthroughs in Protein Structure Calculation and Design

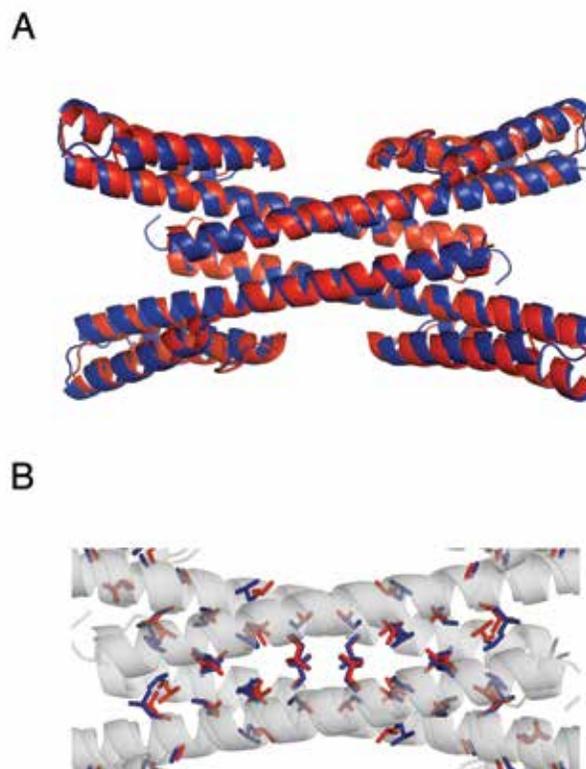
*The calculation of protein structure and the design of novel proteins are two of the most important challenges in structural biology today. Addressing these challenges will help researchers cure diseases and design proteins for a variety of medical and industrial uses. This 2013 INCITE project builds on earlier successes and increases the scope of research from previous INCITE awards.*

The primary bottleneck to both structure calculation and protein design is the search of conformational space: Both the vast number of possible conformations and rugged energy landscape make searching for the lowest free-energy conformation a formidable task. While much progress has been made recently in protein structure modeling, the ability to sufficiently sample conformation space is still a limiting factor for many protein-modeling applications, notably protein structure prediction and design. Making significant advancements in this field requires a methodological breakthrough in sampling algorithms using high performance computing resources.

A team led by David Baker of the University of Washington is using ALCF resources to address a variety of applications for which the resolution of the conformational sampling bottleneck could lead to important scientific breakthroughs. The vast number of highly connected processors make the leadership-class resources available through INCITE crucial to continued progress in these areas.

Building on the success of work achieved through three previous INCITE awards, the team recently described a breakthrough in conformational sampling which leveraged highly parallel computations on the ALCF's Blue Gene supercomputers. The results of this work will enable advances in many challenging problems in computational structural biology, including the *ab initio* prediction of proteins larger than 15 kDa, the calculation of structures of proteins larger than

20 kDa using sparse nuclear magnetic resonance data, the determination of membrane protein structures (of specific relevance to pharmaceutical research), and finally the design of a novel enzyme system to fix carbon dioxide to produce biofuels.



Successful CASP10 contact-assisted structure prediction of a 384-residue homooligomeric tetramer, Tc680o (PDB code 4fm3). (A) The submitted model (red) has a root-mean-square deviation from the crystal structure (blue) of 2.9 angstroms. (B) The interface between chains has near atomic-level accuracy; core-side chains are highlighted.

**Image Credit**  
David Baker, Frank DiMaio, David E. Kim, Yifan Song, Ray Yu-Ruei Wang, University of Washington

**IMPACT** » Proteins serve crucial functions in nearly all biological processes. A protein's function and its structure go hand in hand. By improving our ability to predict the 3D structure of proteins, researchers at the ALCF are paving the way for the development of new antibiotics, cancer treatments, and biofuels.

ALEXEI KHOKHLOV

The University of Chicago | roux@uchicago.edu

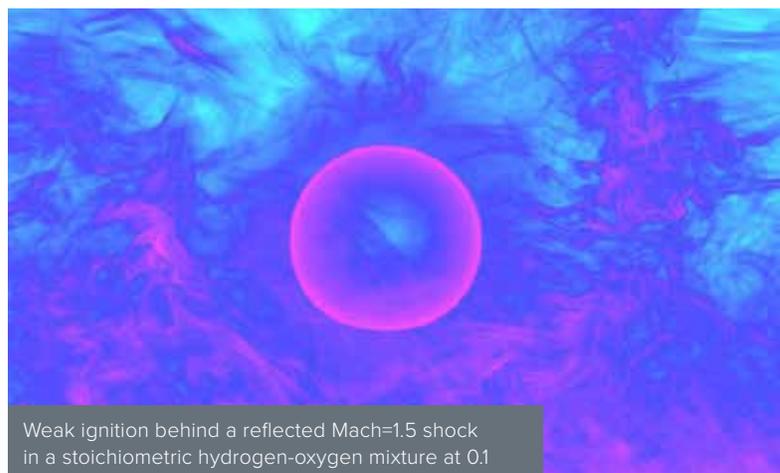
## Simulations of Deflagration-to-Detonation Transition in Reactive Gases

*Deflagration-to-detonation transition (DDT) and the resulting detonation waves in hydrogen can have catastrophic consequences in a variety of industrial and energy-producing settings, including the production, transportation, and use of hydrogen fuel, and safety on nuclear reactors where hydrogen can be accumulated in cooling pipe systems due to radiolysis of water. Researchers at the ALCF are simulating the high-speed combustion and detonation of hydrogen-oxygen mixtures to enable safer and more widespread use of hydrogen as an alternative fuel.*

Numerically, researchers have also long sought to study the DDT process in terms of first-principles calculations based on the Navier-Stokes equations, which contain most of the physics of fluid motion. These equations, however, are prohibitively difficult to solve when modeling the turbulent, high-speed burning of hydrogen-oxygen mixtures without the aid of high-performance computing. First-principles simulations based on the Navier-Stokes equations have been conducted but in compromised fashion because of a lack of computer power. This has resulted in models with unrealistically undersized versions of DDT experiments, and turbulence being simulated in an artificially constrained manner.

To address this issue, a research team, led by Alexei Khokhlov at the University of Chicago, is using ALCF supercomputers to conduct direct numerical simulations aimed at understanding and predicting high-speed combustion and DDT. Their goal is to make first-principles simulations of DDT for hydrogen-oxygen mixtures that burn as they would in typical experiments, permitting a high degree of comparison between the simulations and actual experiments. These are direct numerical simulations, in that turbulence scales are fully resolved rather than modeled.

The group's latest 3D simulations were able to resolve the reflected shock bifurcation phenomena where the boundary layer is lifted off the wall and a near-sonic turbulent recirculation jet is created with order-of-one pressure fluctuations, strong vortices, and multiple shock-lets. The simulations were first validated with non-reactive experiments in CO<sub>2</sub>, showing excellent agreement with the experiment, and then carried out for reactive hydrogen-oxygen mixtures. With groundwork complete, the team embarked on high-resolution simulations of DDT in a hydrogen-oxygen mixture in long pipes. Researchers hope this will allow them to view and study phenomena not observable in experiments.



Weak ignition behind a reflected Mach=1.5 shock in a stoichiometric hydrogen-oxygen mixture at 0.1 atm initial pressure. This picture shows a numerically generated pseudo-schlieren image of the onset of a detonation in a turbulent boundary layer.

### Image Credit

Charles Bacon, Argonne National Laboratory; Alexei Khokhlov, The University of Chicago; Joanna Austin and Andrew Knisely, University of Illinois at Urbana-Champaign

**IMPACT** » Hydrogen is an abundant, environmentally friendly fuel with the potential to reduce our dependence on foreign oil, improve the environment, and boost our economy. Results from this ALCF project will ultimately feed into the design of safer hydrogen fuel systems, addressing a major impediment to widespread adoption of hydrogen as a nonpolluting energy source to help replace fossil fuels.

MAYA GOKHALE

Lawrence Livermore National Laboratory | gokhale2@llnl.gov

## Fault-Oblivious Exascale (FOX) Computing Environment

*Next-generation exascale systems—machines with 100 million cores—will represent a one-thousand-fold increase in computing power over today's state-of-the-industry petascale systems first introduced in 2008. Researchers are using the resources at the ALCF now to develop fundamentally new operating systems and applications required by these future extreme-scale machines.*

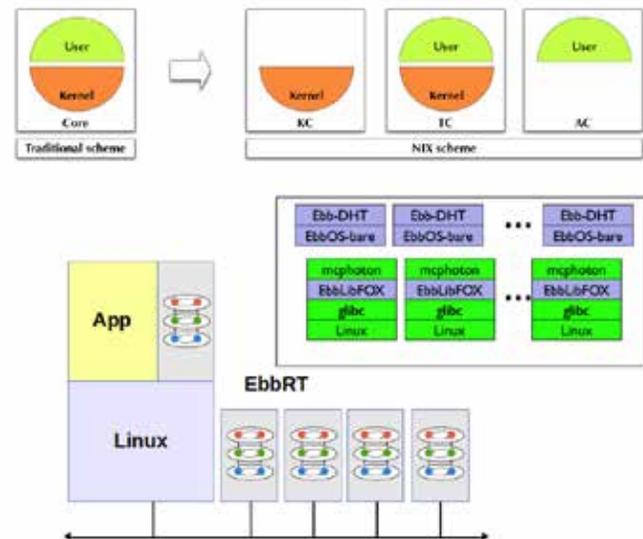
Experts predict that exascale computers will be available to the scientific computing community within the decade. However, developing and harnessing the unprecedented capacity of exascale brings unique challenges, including a fundamental rethinking of the application design to make them resilient in the high-fault environment of exascale.

A research team led by Maya Gokhale of the Lawrence Livermore National Laboratory is using ALCF resources today to address the unique issues inherent in exascale computing. Their work involves developing a prototype exascale environment for use in the development of applications, operating systems, and hardware at the extreme scale.

Most recently, the team developed a “bare metal” distributed hash data table (DHT) as an OS/runtime resource to store and search process and task metadata. The DHT is a prototype primitive for exploring the construction of fault-tolerant applications that benefit from system-level optimization. In particular, the team has ported a Monte Carlo-based simulation to the DHT that can survive faults of the nodes.

The group is also working on an auto-parallelization framework that exploits a task model. Specifically, the team is developing a uniprocessor virtual machine that automatically generates speculative execution tasks. These tasks populate a data store of current/future system-state pairs. The virtual machine periodically consults the data store, and if

a current-state match is found, the true system state is advanced to the corresponding future state. In this manner, the execution model of the virtual machine automatically converts execution into gangs of speculative tasks that parallelize the application. In the future, the DHT will be used as the data store for the auto-parallelization framework.



Hybrid exascale OS combining “bare metal” execution with traditional OS services by assigning roles to cores in the many-core exascale CPU. An application can choose to run its customized libraries under a full-featured OS or directly on the hardware cores. The graphic (above) shows kernel, time-sharing, and application cores; (below) the Monte Carlo photon simulation application’s custom library, in this case a Distributed Hash Table (DHT) Elastic building block (Ebb), can run under the OS or, for high performance, directly on the hardware.

### Image Credit

NIX scheme (top): Jim McKee, Bell Laboratories; Francisco Ballesteros, Rey Juan Carlos University; Ron Minnich, Sandia National Laboratories. Ebb approach (bottom): Jonathan Appavoo, Boston University

**IMPACT** » Access to state-of-the-art supercomputers is key to U.S. competitiveness. To keep pace, the world-class computing resources at the ALCF are being used today to develop the next-generation exascale supercomputers that will one day replace them. Researchers are working to create resilient operating systems for tomorrow’s machines, and to address the inherent challenges of manipulating, managing, and transferring big data.

MICHAEL WEHNER

Lawrence Berkeley National Laboratory | mfwehner@lbl.gov

## Attributing Changes in the Risk of Extreme Weather and Climate

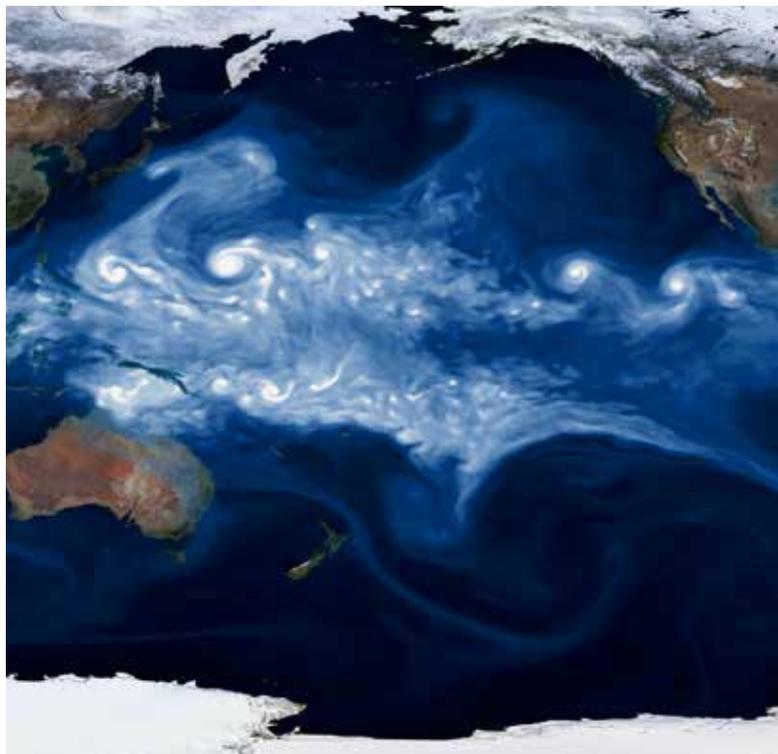
*Researchers from Lawrence Berkeley National Laboratory are using Mira to generate and analyze high-resolution global climate models to resolve many of the processes responsible for extreme precipitation and storms. The results of these studies will be made available to the larger climate community to better understand causes of changes in the risk of localized extreme events.*

The climate system is a naturally chaotic system that is being driven by human caused and natural factors. Variations in extreme weather events due to climate change might have already impacted society, and their effects are likely to increase in the future.

Modeling the statistics of extreme weather is computationally intensive. Because of their rarity, ensembles of simulations are required to build a dataset of extreme events large enough for robust statistical analysis. For this study, a global model is preferable to a regional model, for its ability to quantify past and future extreme weather statistics at many locations around the planet.

A team from Lawrence Berkeley National Laboratory is using Mira to generate and analyze two large ensembles of a ~25-km resolution global climate model. At this resolution, they are able to resolve many of the processes responsible for extreme precipitation and storms. Using the Community Atmospheric Model (CAM5.1) options for the fully coupled Community Earth System Model (CESM), the first ensemble emulates the “world that actually was,” while the second ensemble emulates “the world that might have been” had human activities not interfered with the climate system.

The research team is working with scientists at the National Center for Atmospheric Research and Argonne National Laboratory to improve inter- and intra-component load balancing in CESM by using mathematical modeling, as an alternative to the current manual approach.



CAM5 high-resolution simulations: Total column integrated water vapor.

### Image Credit

Chris Algieri, Bill Collins, Fuyu Li, Prabhat, and Michael Wehner, Lawrence Berkeley National Laboratory; Julio Bacmeister, Andrew Gettelman, and Richard Neale; National Center for Atmospheric Research; Kevin Reed, University of Michigan

**IMPACT** » The simulations enabled by this INCITE project will advance our ability to understand recent changes in extreme weather events. The data and analysis from this study will be made available to the larger climate science community and used to better inform decision makers about society’s ability to adapt to existing and future climate change.

THOMAS JORDAN

University of Southern California | tjordan@usc.edu

## CyberShake 3.0: Physics-Based Probabilistic Seismic Hazard Analysis

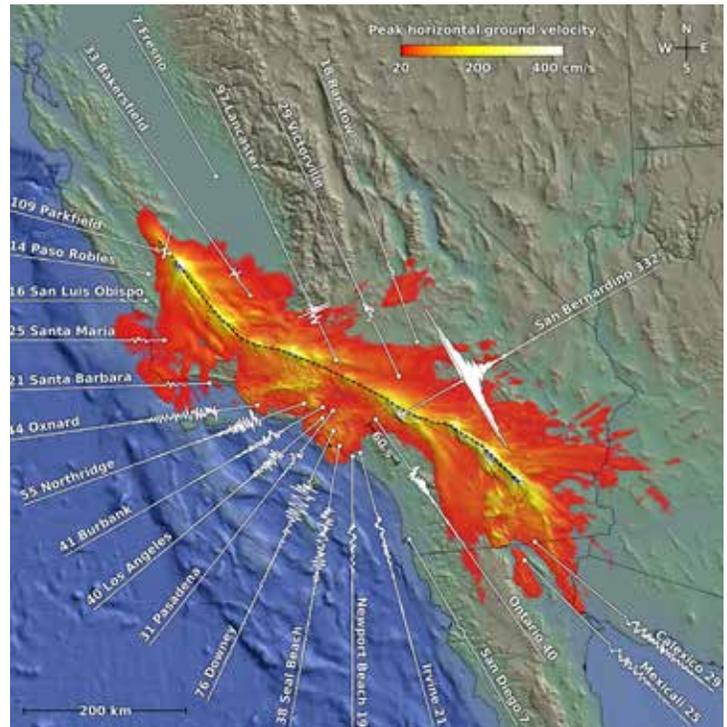
*Recent destructive earthquakes including Haiti (2010), Chile (2010), New Zealand (2011), and Japan (2011) highlight the global need for improved seismic hazard information. Supercomputer simulations from the ALCF are helping scientists model what goes on inside an earthquake and better understand seismic hazard. Their findings are expected to inform the work of building designers and emergency planners worldwide.*

A team led by Thomas Jordan of the Southern California Earthquake Center at the University of Southern California is using the ALCF's high-performance computing resources to calculate a physics-based probabilistic seismic hazard map for California. Dubbed "CyberShake 3.0," this first-ever 3D wave propagation computational approach shows the seismic hazard implications of the most advanced earthquake forecasts of the U.S. Geological Survey (USGS), the government body that produces the nation's official seismic hazard forecasts. These forecasts are used nationwide in the development of building codes to minimize seismic risk.

Produced with the cooperation of the USGS and based on its official hazard maps and site-specific hazard curves, CyberShake 3.0 improves on existing models by better capturing the impact of rupture directivity and sedimentary basin structures on earthquake peak ground motions. In addition, CyberShake introduces never-before-available seismic hazard data products including an extended earthquake rupture forecast and shaking duration forecasts. CyberShake 3.0 represents a pioneering effort to use high-performance computing to improve the broadest-impact seismic hazard data products.

While CyberShake research focuses on the well-studied, seismically vulnerable region of California, the technique is applicable to any seismically active region in the world. A physics-based hazard map for

California represents a transition point for seismic hazard and risk calculations, adding value to the world's most advanced, newest earthquake rupture forecast.



Ground motion simulations reveal regions at risk of strong shaking during a possible magnitude-8 earthquake on the San Andreas fault. For the CyberShake project, reciprocal simulations of all known possible quakes are combined to estimate the total probabilistic hazard for California.

Image Credit

Geoffrey Ely, Argonne National Laboratory

**IMPACT** » This work is enabling scientists to develop more accurate estimates of ground motion in Southern California and apply their findings 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.

ROBERTO PAOLI  
CERFACS | paoli@cerfacs.fr

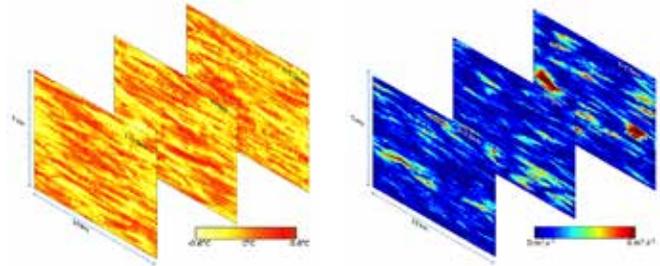
## Evaluation of Mesoscale Atmospheric Model for Contrail Cirrus Simulations

*Contrails, line-shaped ice clouds that appear behind aircraft, are formed from water exhaust. Contrails can spread to form cirrus clouds that affect how the Earth's atmosphere heats and cools. As the volume of air travel continues to grow, researchers are working to fine-tune global climate models to include the effects of contrail cirrus to better inform scientists and policymakers.*

Contrails, the visible white lines in the sky left behind an airplane, are ice clouds made by water exhaust from the aircraft's engine. These ice clouds can spread to form contrail cirrus clouds that may persist for hours and extend over several square miles, and are nearly indistinguishable from naturally occurring cirrus clouds. Because these clouds artificially increase the Earth's cloudiness, they can affect global climate by altering the balance of incoming and outgoing heat in the atmosphere from solar energy. As air traffic volume continues to grow, the impact of more contrail cirrus is of increasing concern for scientists and policymakers.

Roberto Paoli from CERFACS is leading a team of scientists using resources at the ALCF to fine-tune numerical models that will allow for a more accurate understanding of the impact of contrail cirrus on global climate. With a previous INCITE award, these researchers used high-resolution large eddy simulations to characterize the mechanisms that control the transition from the contrail stage to the young cirrus stage as a function of the wake age. Researchers observed that atmospheric turbulence is the main driver of contrail evolution after vortex break-up, whereas at later stages, radiative transfer and sedimentation affect the transition by controlling the plume horizontal and the vertical spreadings.

Next, Paoli's team will validate their model against data collected from a research aircraft flying behind commercial airliners. The research plane will measure velocity fields and particle characteristics inside and outside the contrail at increasing distances from the airliner. Researchers will then design the first numerical simulations using atmospheric conditions that are representative of flights operated by commercial airliners. The computational domain is a box of 10 km x 10 km in the horizontal directions and 5 km in the vertical direction where the background atmospheric turbulence is generated and sustained so as to match the local ambient conditions.



Snapshots of atmospheric turbulence for contrail simulations: potential temperature fluctuations (left) and kinetic energy fluctuations (right) at selected vertical planes.

Image Credit  
Roberto Paoli and Odile Thouron, CERFACS

**IMPACT** » By working to provide a better understanding of the physics of the formation of contrail cirrus, researchers are equipping the scientific community with a more complete evaluation of aviation impact in next-generation climate models. Information obtained from the study of these models will be used to better inform scientists and policymakers.

ROBERT MOSER

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

## Direct Numerical Simulations of High Reynolds Number Turbulent Channel Flow

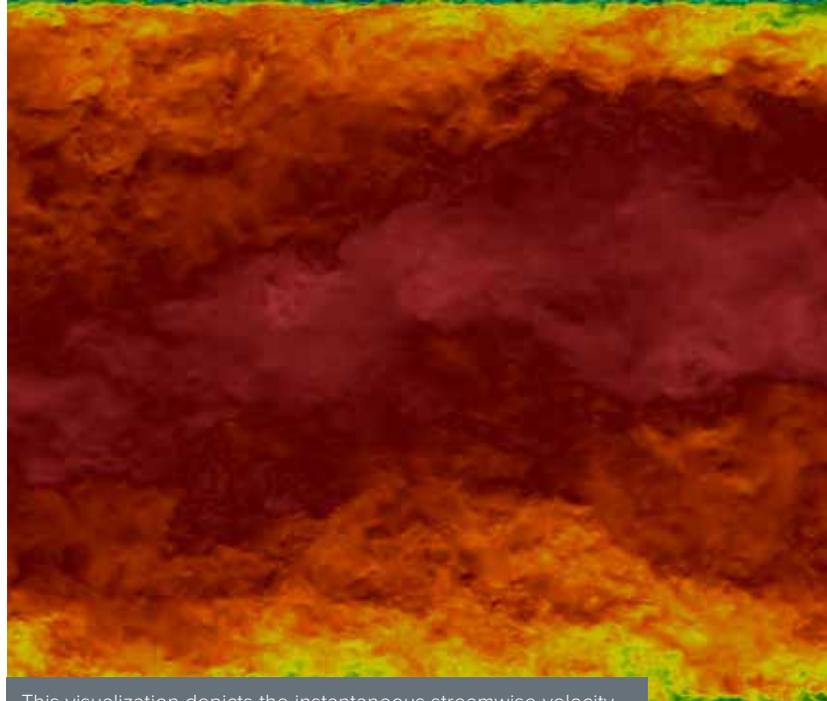
*A substantial fraction of the energy consumed by a moving vehicle is due to drag and is dissipated by turbulence as it moves through air or water. The same forces are at work as air or liquid moves through ducts or pipes. This project is tapping Mira to advance the understanding of turbulent flows and contribute to the improved efficiency of transportation systems.*

Wall-bounded turbulence is at the heart of the interaction between solid surfaces (such as vehicles and pipes) and the fluid flowing past them, leading to drag and the dissipation of energy. Engineering developments to reduce drag and energy consumption are greatly impeded by the lack of accurate models of the turbulence phenomena involved.

The goal of this project, led by Robert Moser at the University of Texas at Austin, is to perform direct numerical simulations (DNS) of high Reynolds number fluid flow to examine the complex physics of wall-bounded turbulence. The subsequent analysis of the simulation data can provide the insights needed to develop improved turbulence models, as well as new concepts for manipulating wall-bounded turbulence.

Of particular interest to Moser's research team is the overlap region, where the viscous near-wall turbulence interacts with the outer-layer turbulences. This region is currently not well understood because simulations to date have not allowed for a sufficiently high Reynolds number to obtain the scale separation needed to shed light on the complexity of this multiscale turbulent structure. However, with the computational power of Mira, this project is currently running DNS with a high enough Reynolds number to generate sufficient scale separation. Given the mesh size (15360 x 1536 x 11520), the simulations are believed to be the largest production DNS ever performed. Prior to this research, the standard

reference data set for turbulence research covered the friction-Reynolds-number range  $Re_\tau = 180 - 2000$ . Supplementing these data with  $Re_\tau = 5200$  from the simulations on Mira will establish a reference data set that will remain useful for the turbulence research community for many years to come.



This visualization depicts the instantaneous streamwise velocity component over a section of streamwise length of the simulated channel. The image is from a simulation of incompressible turbulent flow ( $Re_\tau = 5200$ ) between two parallel smooth planes. Simulations of wall-bounded turbulence such as this are analogous to an experimental laboratory, permitting scientists to probe the interaction between a vehicle (or a pipe) and the fluid flowing past it, transferring momentum and determining the drag.

Image Credit

Nicholas Malaya, University of Texas at Austin

**IMPACT** » Results from this project will provide insights necessary to develop more accurate turbulence models. Ultimately, this work could lead to more energy-efficient transportation through the design of improved vehicle surfaces and reduced-drag piping and ducts.

MIHAI ANITESCU

Argonne National Laboratory | anitescu@mcs.anl.gov

## Optimization of Complex Energy System Under Uncertainty

*The national imperative to use increased levels of renewable sources of energy, such as wind and solar, has resulted in a significant challenge for the operation and design of the U.S. power grid. Mathematical expertise and supercomputing are required to explore the benefits of using optimization under uncertainty as the paradigm for managing the uncertainty in renewable energy supply. This project combines both of these Argonne strengths.*

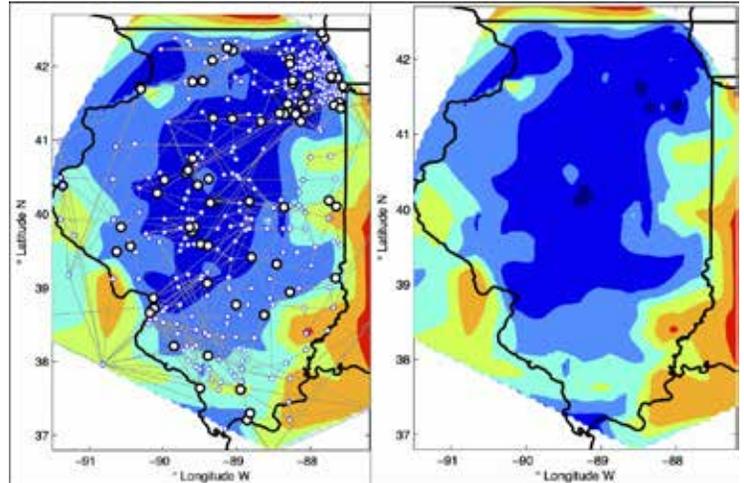
As opposed to fossil-fuel generation systems, the amount of renewable energy available to supplement the power grid at any given time is uncertain and unmatched with demand changes.

This leads to the risk of not satisfying demand and thus having to “shed load,” that is, to cut consumers from the power grid. Uncertainty in weather and other risks, such as generator and transmission line failures, are currently mitigated by using conservative and expensive reserve units — coal or natural gas generators — that can increase electricity generation on short notice. Such reserves are both economically and environmentally costly.

Argonne computational mathematician Mihai Anitescu is exploring the benefits of using optimization under uncertainty as the paradigm for managing the uncertainty in renewable energy supply in order to reduce reserve requirements and stabilize electricity markets in the next-generation power grid.

As a part of this project, stochastic programming — a leading paradigm for optimization under uncertainty — formulations of the decision process that schedules supply and matches demand are solved by a new scalable solver, PIPS. The magnitude of such formulations for the State of Illinois alone can reach billions of variables and constraints once the uncertainty in the supply is taken into account. Incorporation of new algorithmic research results into PIPS made it possible to efficiently simulate such large systems on ALCF’s Intrepid in an amount of

time considered “real-time” by current power industry practices. Strong scaling was obtained by PIPS in solving a 12-hour horizon stochastic problem with 2 billion constraints and 2 billion inequalities.



On the left: This shows the Illinois power grid system overlaid on fields portraying electricity prices under a deterministic economic dispatch scenario. Dark blue areas have the lowest prices while red and yellow have the highest. Argonne researchers use a model of the Illinois grid to test algorithms for making power dispatch decisions under uncertainty. On the right: This shows electricity prices in Illinois under a stochastic economic dispatch scenario. The dark blue region, representing areas with the lowest prices, is larger compared with that in the deterministic economic dispatch scenario. These simulations show that stochastic dispatch leads to more homogeneous prices and a more even spread of social welfare than the deterministic approach.

Image Credit

Victor Zavala, Argonne National Laboratory

**IMPACT** » ALCF supercomputing resources and novel optimization-under-uncertainty formulations have allowed the team to demonstrate that, at least on some configurations, even 20% wind penetration — the U.S. stated policy for the year 2030 — can be accommodated without the need for significant reserve increase if using stochastic optimization, something which would not be achievable with the traditional formulations.

MISUN MIN

Argonne National Laboratory | mmin@mcs.anl.gov

## Predictive Modeling for Complex Phenomena in Electromagnetics and Fluid Systems

*A team led by Misun Min of Argonne National Laboratory is exploring computational methodologies for efficiently and accurately predicting the optical, electrical, mechanical, thermal, and transport properties of electromagnetics and fluids systems.*

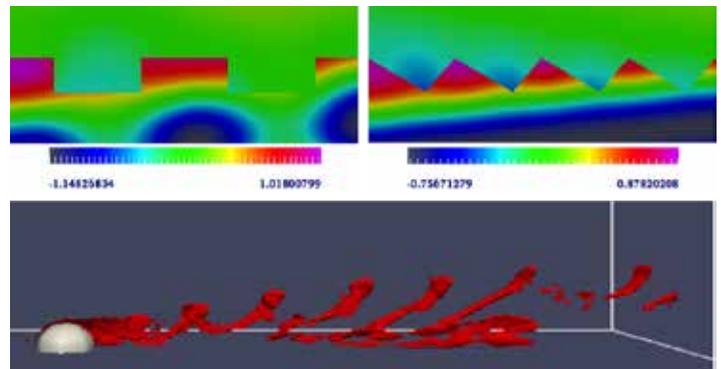
This team is developing scalable, high-order algorithms based on spectral element, spectral element Discontinuous Galerkin (DG), and spectral element DG lattice Boltzmann approaches. Algorithmic efforts include convergence and stability analysis, efficient parallelization and performance, and realistic simulations with validation.

Methodologies are integrated into the software packages, NekCEM and NekLBM, working toward extreme-scale computing capability. Using ALCF resources, they carried out large-scale performance runs on Intrepid and Mira, including threading approaches for computation and I/O algorithms. NekCEM achieved 90% parallel efficiency on Intrepid with 131,072 cores for the problem size of 2.2 billion grid points. NekLBM achieved 78% efficiency for 57 million points on Intrepid with 131,072 cores. Production performance was improved with a threaded reduced-blocking I/O approach by overlapping the I/O latency with computation, achieving the application-perceived write bandwidth of 70 GB/s on Intrepid with 32,768 cores for the problem file size of 52 GB.

The team has recently added new simulation capabilities to these software packages. Key features include high-order solvers for Schrödinger and Helmholtz equations. The latest studies demonstrate a spectral element frequency-domain solver involving transparent boundary treatment for quasi-periodic scattering solutions on multilayered structures with rough interfaces. Some of these algorithms will be extended to study nonlinear optical response for a

hybrid system of a semiconductor quantum dot in the near field of a plasmonic metal nanostructure.

High-order lattice Boltzmann models are also developed for turbulence and heat transfer simulations with validated computational results. The team will extend the algorithms to solve multiphase systems. A binary lattice Boltzmann model for fluid motion and a phase field model for the evolution of interfaces will be constructed in a coupled form for two-phase flow simulations.



Quasi-periodic scattering solutions at oblique incidence on the gratings with square- and sawtooth-shaped interfaces, computed with NekCEM (top); Vortices profile for flows past a hemisphere, computed with NekLBM (bottom).

### Image Credit

Paul Fischer, Misun Min, Argonne National Laboratory; Taehun Lee, City College of New York; Ying He, Purdue University; David Nicholls, University of Illinois at Chicago.

**IMPACT** » The methodologies produced by the team and the high-fidelity, high-efficiency, open source codes will benefit a wide range of relevant research communities and industries involved in the production of plasmonic devices, photovoltaic cells, semiconductors, batteries, and reactors.

MICHAEL CROWLEY

National Renewable Energy Laboratory | michael.crowley@nrel.gov

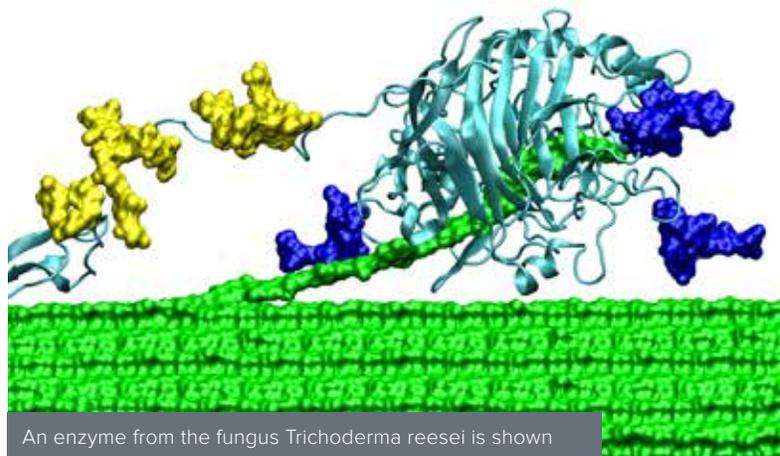
## Thermodynamics of Binding Biomass to Cellulases for Renewable Fuel

*Biofuel feedstocks, such as wood chips and switchgrass, naturally resist being broken down into fermentable sugars, making it difficult to produce biofuels at a cost and pace that can compete with fossil fuels. Scientists working at the ALCF are using Mira to study nature's catalysts, enzymes, for inspiration in their quest to find a more effective means of converting biomass into renewable fuel.*

Scientists from the National Renewable Energy Laboratory (NREL) are conducting large-scale simulations of the physical behavior of cellulase enzymes. Naturally produced by some fungi and bacteria, these particular enzymes are being modeled because they effectively trigger the chemical changes necessary to degrade hardy plant materials into life-sustaining sugars.

The research team is carrying out the simulations to gain a fundamental understanding of the complex cellulose-to-sugar conversion process, known as enzymatic hydrolysis. With this information, the research team hopes to uncover how these enzymes can be manipulated to develop superior biological catalysts for improved biofuel production. The results can then be fed into experiments aimed at developing and validating improved catalysts.

With high-performance computers like Mira, scientists are able to glean data that is difficult to obtain through conventional experimental approaches (e.g., binding free energy), helping to accelerate the process of screening and testing new enzymes. The research team has begun development of a molecular-level theory of enzyme "processivity" (the ability to catalyze consecutive reactions) that relates directly to the structural features of enzymes. This work will provide accurate binding free energy values on various enzymes of importance to the biofuels industry, shedding light on a key parameter used to compare the function of enzymes.



An enzyme from the fungus *Trichoderma reesei* is shown digesting cellulose (green). The enzyme consists of three sub-domains: a small carbohydrate-binding module; a long, flexible linker decorated with O-linked glycosylation (yellow); and a large catalytic domain (CD) with N-linked glycosylation (blue) and a 50 Å tunnel for the threading of cellodextrin for catalytic cleavage.

Image Credit

Gregg Beckham, National Renewable Energy Laboratory

**IMPACT** » The project addresses a national goal to make biofuels a more prevalent and reliable option as an alternative transportation fuel. In fact, the DOE has stipulated that 30 percent of our nation's gasoline demand be displaced by biofuels by 2030.

LEE SHUNN

Cascade Technologies, Inc. | shunn@cascadetechnologies.com

## Combustion Stability in Complex Engineering Flows

*Many industrial combustion systems are subject to instabilities that complicate operations and limit safety and efficiency. This project intends to numerically reproduce common combustion instabilities to better understand the physical mechanisms behind their generation and to develop strategies to eliminate or control them.*

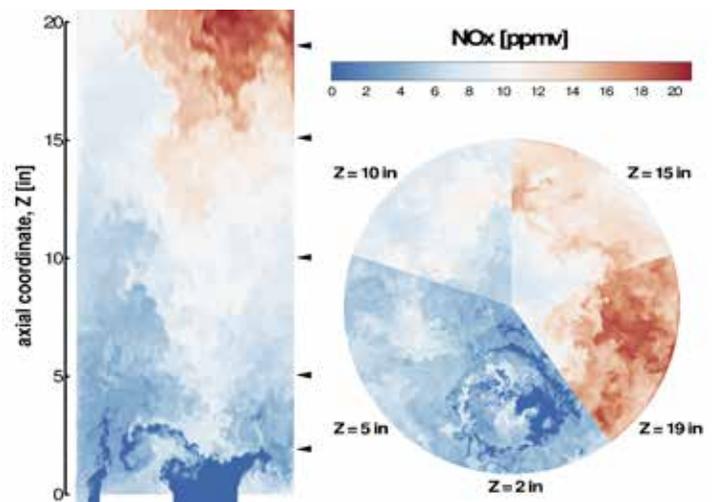
Many analytical and empirical methods have been developed to explain the growth and propagation of combustion instabilities. The predictive capability of these theories, however, remains extremely limited due to the inherently complex nature of this multiscale, multiphysics phenomenon.

Research led by Lee Shunn of Cascade Technologies, Inc., and Shoreh Hajiloo of GE Global Research, will develop and validate a high-fidelity large eddy simulation (LES) combustion model for accurately predicting unsteady combustion problems based on the flamelet/progress-variable (FPV) approach. The LES results will contribute to a more thorough understanding of combustion physics in systems that are prone to instability-induced failures.

Two industrially relevant configurations have been selected to highlight recurrent types of combustion instabilities with broad impact on a range of applications: modern Dry Low-NOX (DLN) combustors used in land-based gas turbines and a high-speed combustor inspired by the Hypersonic International Flight Research Experimentation program.

The simulations are being performed on the ALCF's Blue Gene/P and Blue Gene/Q high-performance computing platforms using CharLES, the parallel unstructured LES solver developed by Cascade Technologies, Inc. CharLES's unique low-dissipation numerics, dynamic turbulence models, and massive scalability enable these problems to be studied under realistic operating conditions in complex engineering geometries.

The project is focusing on predicting emissions, operability limits, and engine-level stability in DLN gas turbines and high-speed propulsion systems. Simulation results will be used to create a high-quality numerical database for assessment of industrial combustor performance and evaluation of lower-fidelity prediction tools.



NO<sub>x</sub> prediction in a swirl-type DLN combustor: (left) zoom near one-of-six swirling premixers, (right) several axial locations. NO<sub>x</sub> accumulates in the low-velocity region at the center of the swirling vortex. Virtual testing enables engineers to design next-generation, low-emission combustion systems.

Image Credit

Lee Shunn, Cascade Technologies; Shoreh Hajiloo, GE Global Research

**IMPACT** » This research represents the frontier of complex, multiphysics combustion simulations involving transient reaction phenomena, non-equilibrium chemistry, compressible turbulence, and multi-mode heat transfer. Successful high-fidelity simulations of these scenarios will help build the scientific insight and understanding necessary to engineer practical solutions to the ubiquitous challenges of combustion instabilities.

CHRISTOS E. FROUZAKIS

Swiss Federal Institute of Technology Zurich | frouzakis@lav.mavt.ethz.ch

## Direct Numerical Simulation of Autoignition in a Jet in a Cross-Flow

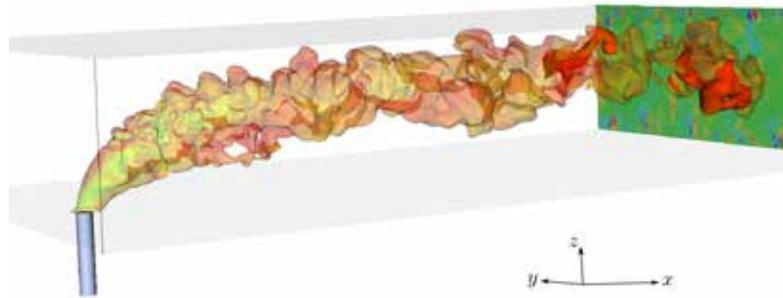
*Advanced gas turbines provide excellent opportunities for environmentally friendly propulsion and electricity generation, but the technology is severely limited by the danger of unwanted autoignition of the fuel-air mixture. For this ESP project, researchers are studying the propensity of fuel-air mixtures to autoignite in a jet in cross-flow (JICF) configuration, and creating a direct numerical simulation (DNS) database that provides significant insight into the complex fluid and flame dynamics of such flows.*

Enabling the next generation of low-emission, fuel-flexible gas turbine technology relies heavily on the design of rapid premixing sections, which are often influenced by JICF. For highly reactive hydrogen or hydrogen-rich fuels, rapid mixing can result in autoignition in the premixing section, leading to combustor damage. The autoignition process, which is determined by complex coupling between turbulent mixing and chemistry, is currently not fully understood.

To address this issue, researchers from the Swiss Federal Institute of Technology Zurich and Argonne National Laboratory are using Mira to investigate the complex dynamics of a diluted hydrogen jet issuing in a turbulent hot stream of air in a channel at different operating conditions. Highly nonlinear chemistry-turbulence coupling and a significant scale separation complicate the dynamics of this process, necessitating the use of an efficient high-order, method-based DNS code, such as Nek5000, and the computational power of Mira.

In the first series of studies, the research team used the low Mach number reactive flow solver based on the Nek5000 spectral element code to run large-scale simulations for two cross-flow temperatures (930 K and 950 K) and a channel friction Reynolds number of 180. In both cases, autoignition kernels initially formed far downstream from the fuel nozzle, creating propagating fronts that are convected out of the domain by the flow. Radical buildup in the

recirculation zone near the nozzle exit results in a flame stabilized around the fuel jet. The DNS database is being examined to unravel the complex interaction of the local flow and mixing conditions with hydrogen chemistry leading to autoignition, the propagation characteristics of the flames around the kernels, and the flame stabilization mechanism. The chemical explosive mode analysis technique is being employed to obtain insight into the local reactivity of the fuel-air mixture, and to identify the regions of burned and unburned mixtures, local combustion mode, and flame stabilization mechanisms.



Autoignition spots (indicated by red isosurfaces of high OH mass fraction) which formed on the most reactive mixture fraction isosurface (colored by the z-velocity) following the mixing of the fuel stream (hydrogen diluted in nitrogen at 850 K injected from the circular pipe) with the hot turbulent air stream flowing into the channel. On the outflow plane, vortical structures are shown by the contour plot of the y-component of vorticity.

### Image Credit

Ammar Abdilghanie, Argonne National Laboratory; Christos E. Frouzakis, Swiss Federal Institute of Technology Zurich

**IMPACT** » The DOE is interested in developing advanced gas turbines to produce clean and less costly energy. This project is helping to advance the fundamental understanding of autoignition in a jet in cross-flow setup, which could ultimately feed into the design of improved and safer stationary gas turbines. Additionally, the project's DNS databases will provide a valuable resource that can be further explored by researchers to validate and tune advanced turbulent combustion models.

UMESH PALIATH

GE Global Research | paliath@ge.com

## Enabling Green Energy and Propulsion Systems via Direct Noise Computation

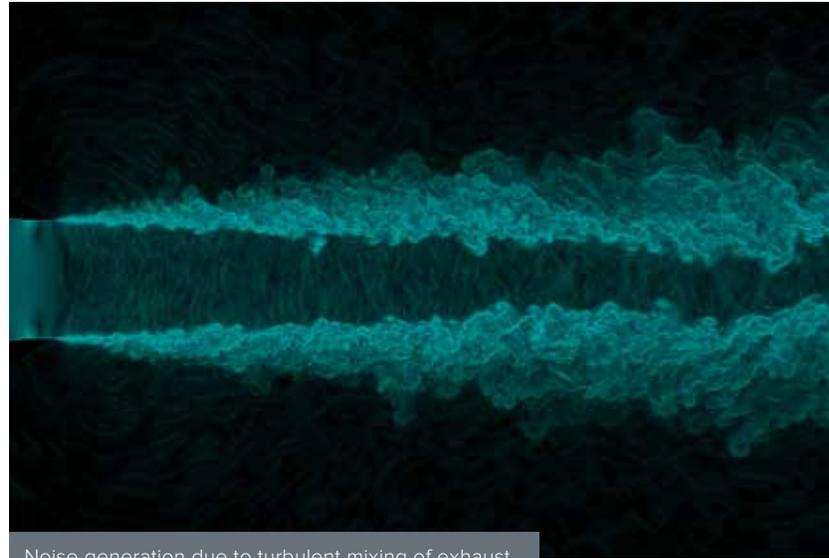
*High-performance computing at the ALCF is enabling a team from GE Global Research to visualize the complex phenomena of turbulent mixing in jet exhaust flow and flow over wind turbine airfoils in ways not possible in experiments.*

A GE Global Research team is conducting simulations at the ALCF to better understand and predict the complex turbulent flow in jet exhaust nozzles and wind turbine airfoils. The research under way for jet engines 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. For wind turbines, simulations are demonstrating readiness in guiding low-noise design, validating the complex scaling of turbulent self-noise, and using effective wall-models to enable large-span blade computations.

Driven by recent jet engine technology's push into ultra-high bypass ratios for increased propulsive efficiency, fundamental work in fan broadband noise has also been initiated. In addition, simulations are being conducted to demonstrate how the validated large eddy simulation solver, coupled with high-performance computational capability, can be used as a numerical rig to gather high-resolution data sets. These data sets are essential in guiding the improvement/development of reduced-order models for these complex flow mechanisms. This two-part strategy (direct computation and numerical rig) to generate data for improved reduced order models can significantly enhance the impact of increasing computational resources.

Researchers continue to conduct large eddy simulations of turbulent wake and jet mixing phenomena with a focus on realistic jet engine/wind turbine geometry and operating conditions. The

GE team is leveraging Mira to demonstrate and accelerate the industrial impact of high-fidelity numerical methods. This approach enables component level to system level optimization thereby accelerating advanced product design.



Noise generation due to turbulent mixing of exhaust flow. Large eddy simulation of exhaust flow from a 2-inch conic nozzle with exit Mach number of 0.97. Visualizing density gradients in the flow.

### Image Credit

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

**IMPACT** » This research 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. The direct transfer of advanced simulation technology will also feed into the design of other next-generation green energy products.

ELIA MERZARI

Argonne National Laboratory | emerzari@anl.gov

## Petascale Thermal Hydraulic Simulations in Support of CESAR

*Numerical simulation is an intrinsic part of nuclear engineering research and the design of nuclear power plants. However, there are many limitations for nuclear modeling and simulation tools. The Center for Exascale Simulation for Advanced Reactors (CESAR) aims to develop a coupled, next-generation nuclear reactor core simulation tool capable of efficient execution on exascale computing platforms.*

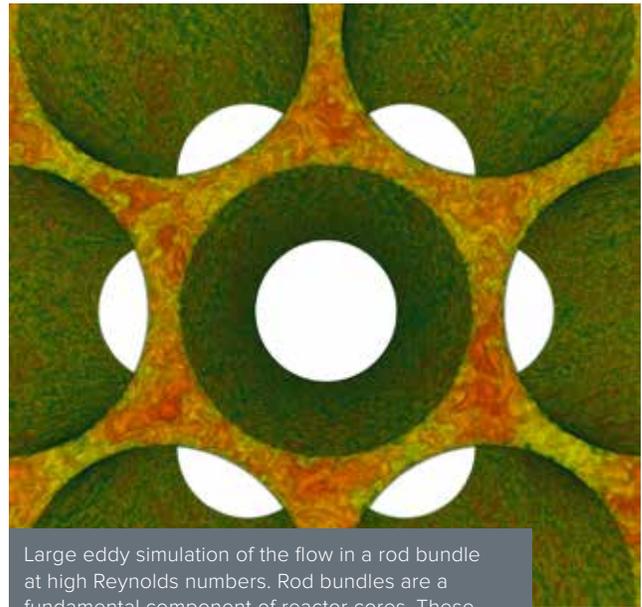
The development, deployment, verification, and validation of higher-fidelity computational capabilities for analyzing, modeling, simulating, and predicting complex thermo-fluid phenomena will help advance nuclear power capabilities by resolving technical, cost, and safety issues.

A team led by Elia Merzari of Argonne National Laboratory is using ALCF resources to create extreme-fidelity physics models that will enable new reactor design analysis in existing and next-generation reactors on exascale computing platforms.

The current state of computational fluid dynamics (CFD) is such that empirical/low-resolution formulations can be run on most computers. These formulations are only known to be valid in certain regimes and, as such, cannot be truly predictive in arbitrary problems. In order to probe beyond recognized regimes, high-fidelity tools must be used, though they are very computationally expensive.

The team's simulations address the limitations of current methods as well as potential scaling to larger machines. Data from one simulation is providing insight into the challenges of exascale simulations. Data from another simulation is being used to examine rod-bundle flows at a level never achieved before, helping to determine computational cost in the exascale limit.

This project is using the spectral element code Nek5000 to focus on the thermal-hydraulic issues in CESAR by examining computational requirements for high-fidelity reactor simulations. The project is running large calculations to test the limits of ALCF's petascale computer Mira and use that information to inform code developers about how the current codes are performing and extrapolate how they might perform on future architectures.



Large eddy simulation of the flow in a rod bundle at high Reynolds numbers. Rod bundles are a fundamental component of reactor cores. These very large and unprecedented calculations (over 8 billion collocation points) will shed light on the performance of current algorithms, help estimate scaling to larger machines, and gain unprecedented insight into the physics of turbulence.

Image Credit

Elia Merzari and Justin Walker, Argonne National Laboratory

**IMPACT** » This project supports the mission of CESAR with large petascale-level simulations needed to address the limitations of current methods and to address the potential scaling to larger machines. The resulting improved state-of-the-art simulation code will aid advancements in nuclear engineering and nuclear energy.

PARVIZ MOIN

Stanford University | moin@stanford.edu

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

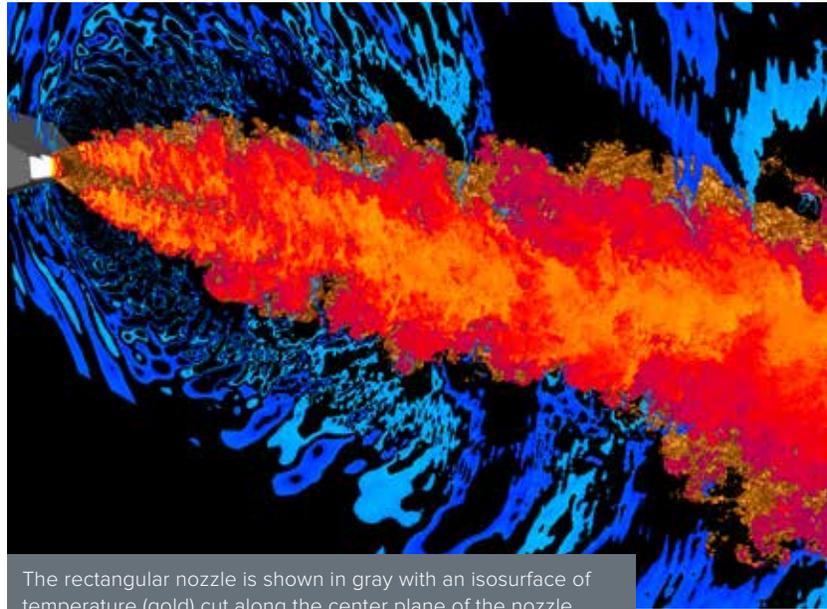
*Stanford University's Parviz Moin is using his ALCC allocation to make a new and potentially industry-changing discovery about the source of crackle in hot supersonic jet engines – a major source of engine noise that causes hearing damage among aircraft carrier deck personnel and significantly impacts communities surrounding military bases.*

Crackle is supersonic jet noise associated with shock-like “N-shaped” acoustic waveforms consisting of sudden strong compressions followed by more gradual expansions, and account for up to 30% of the overall sound pressure level along the direction of peak jet noise. However, the mechanism by which N-shaped waves are generated — directly by the supersonic jet, or as a result of nonlinear acoustic propagation — is not yet fully understood. One of the objectives of this research is to apply large eddy simulation (LES) to determine the source of crackle in hot supersonic jets issuing from military-style faceted nozzles. Eliminating crackle has the potential to help meet the U.S. Navy’s near-term jet noise reduction goal of 3 dB in the peak noise.

Based on simulations performed at the ALCF, an understanding of the source of crackle noise is now emerging. These high-fidelity simulations have allowed Moin’s team to directly observe “crackle waves” being emitted intermittently from the jet turbulence.

Moreover, in comparison to laboratory experiments, the simulations provide a unique ability to revert “backwards in time,” reconstructing high-fidelity flow fields at precise instants preceding the emission of a crackle wave. Using this technique, the team has been able to unravel the physical processes leading up to crackle noise and identify patterns in how crackle might form. Several such crackle events for the base

jet noise case were detected and stored for future analysis. Once the source of crackle is determined, new nozzle designs can be simulated to minimize the source of the noise.



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, Stanford University

**IMPACT** » This work aims to develop turbulence simulation codes containing state-of-the-art numerics, turbulent transport, mixing, aeroacoustic, and combustion capabilities; and to broaden the community of researchers able to fully use leadership computing resources to predict and quantify the uncertainty of multiscale, multiphysics turbulent flow phenomena.

ALEKSANDR OBABKO

Argonne National Laboratory | obabko@mcs.anl.gov

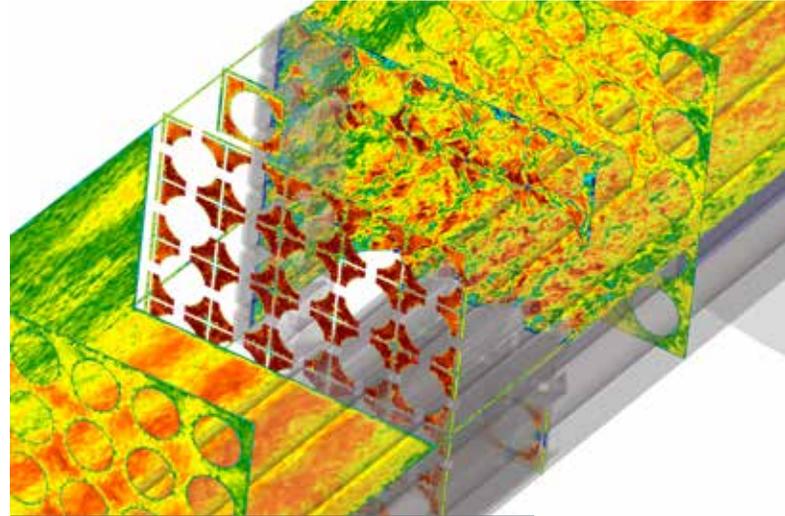
## U.S. - Russia Collaboration on Cross-Verification and Validation in Thermal Hydraulics

*In collaboration with the Moscow Institute of Nuclear Energy Safety (IBRAE), Aleksandr Obabko and the SHARP team at Argonne National Laboratory is using ALCF resources to improve high-fidelity computational capabilities for research into complex thermo-fluid phenomena to advance the case for safe nuclear energy as a reliable, carbon-free energy source.*

One of the critical safety parameters in a nuclear power plant involves maintaining the peak material temperature well below the melting point of the core structural components. Prediction of the peak temperature, as well as the associated uncertainty, involves an accurate computation of the thermal mixing governed by conduction and convection in a complex geometry. While modeling and simulation are intrinsic to engineering research, design, and licensing of existing and proposed nuclear power plants, today's tools are mostly low dimensional, empirically based, valid for conditions close to the original experiments, and in many cases, only incremental improvements on decades-old legacy codes.

Researchers are using the resources of the ALCF to develop improved, higher-fidelity thermal hydraulics codes that will allow for the simulation of nuclear systems with well-defined and validated prediction capabilities. ALCF staff has assisted these efforts by improving the performance of Nek5000 on the Blue Gene/P and Blue Gene/Q platforms to demonstrate its capability to solve thermo-hydraulic problems at extremely large scale.

As part of an international engagement mandated by DOE in support of the safe, secure, and peaceful use of nuclear energy, the team initiated a collaborative activity with IBRAE that focuses on the validation and cross-verification of three thermal hydraulics large eddy simulation codes: Nek5000, Cfoam-CABARET, and Conv3D, for a set of international validation experiments and benchmarks.



Axial velocity in Nek5000 LES simulation of the 2012 OECD/NEA Matis benchmark experiment—one of the experiments chosen for cross-verification and validation in this ALCC project.

Image Credit

Paul Fischer and Aleksandr Obabko, Argonne National Laboratory

**IMPACT** » Nuclear power is one of the safest and least environmentally damaging ways to generate electricity. The safe production of nuclear energy requires systems designed with a solid understanding of the thermal dynamics involved. U.S. and Russian scientists are using ALCF resources in a collaboration to improve the thermal dynamics codes used in models that drive next-generation nuclear systems design.

SUBRAMANIAN SANKARANARAYANAN

Argonne National Laboratory | skrssank@anl.gov

## Dynamics of Conformational Transition in Thermo-Sensitive Polymers and Hydrogels

*Researchers at Argonne National Laboratory are using ALCF resources to study the behavior of polymers to inform the scientific community at work on the creation of new materials that may have diverse practical applications—including drug delivery and oil recovery systems.*

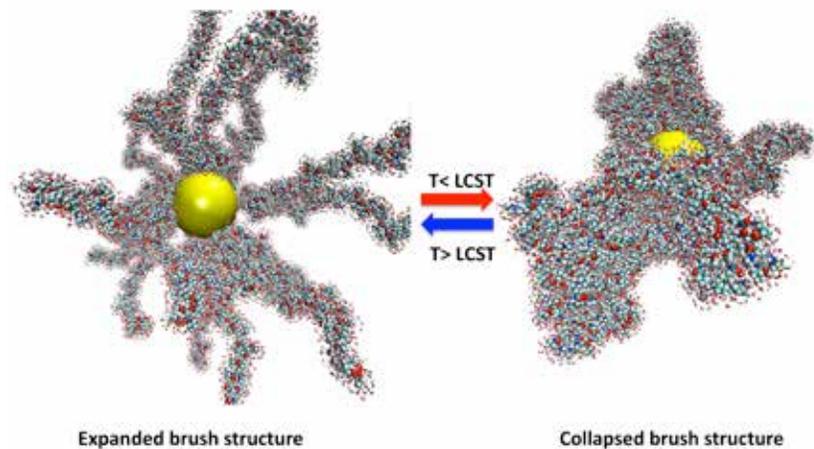
Polymer science—the study of macromolecules—is responsible for synthetic materials like plastics, fibers, and membranes. A better understanding of how macromolecules change with the introduction of variables such as temperature, light, pH, and magnetic fields will allow for the development of new materials useful in a number of practical applications including medical diagnostics, bio-engineering, drug delivery, and enhanced oil recovery.

Subramanian Sankaranarayanan with Argonne's Center for Nanoscale Materials is making use of an ALCC award to study conformational transformations in thermo-sensitive oligomers and their macroscopic architectures such as polymer brushes and polymer gels.

To this end, the group has leveraged the expertise of ALCF staff to pave the way for dramatic enhancements to their research capabilities. First, ALCF staff assisted the team in maximizing the performance of the LAMMPS classical molecular dynamics (MD) software for their work, resulting in a 6x reduction in the time to solution in the presence of intensive I/O operations. Further, the ALCF staff recommended the team move to NAMD, a more scalable classical MD code with high percentage of peak. ALCF staff worked to tune the performance of NAMD by fixing input bugs, optimizing parameters for load balance and parallel I/O, and addressing a significant memory problem plaguing large-scale runs with many fixed atoms.

With the assistance provided by the ALCF performance engineers, the team is currently

simulating multi-million polymer systems at 10x their original calculation rate. These improvements will not only make their scientific campaign feasible but may also enable researchers to expand the scope of their project in ways that were previously not possible.



Expanded brush structure

Collapsed brush structure

Conformational dynamics across lower critical solution temperature in a polymer brush structure grafted on metal nanoparticle.

### Image Credit

Sanket Deshmukh, Ganesh Kamath, Derrick Mancini, and Subramanian Sankaranarayanan, Argonne National Laboratory

**IMPACT** » Researchers at Argonne National Laboratory are using ALCF resources to improve our understanding of how macromolecules change with the introduction of variables. Recent code improvements will enhance their ability to inform the scientific community at work on the creation of a variety of new materials, including those used in drug delivery, medical diagnostics, and oil recovery systems.

WILLIAM GEORGE

National Institute of Standards and Technology | wgeorge@nist.gov

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

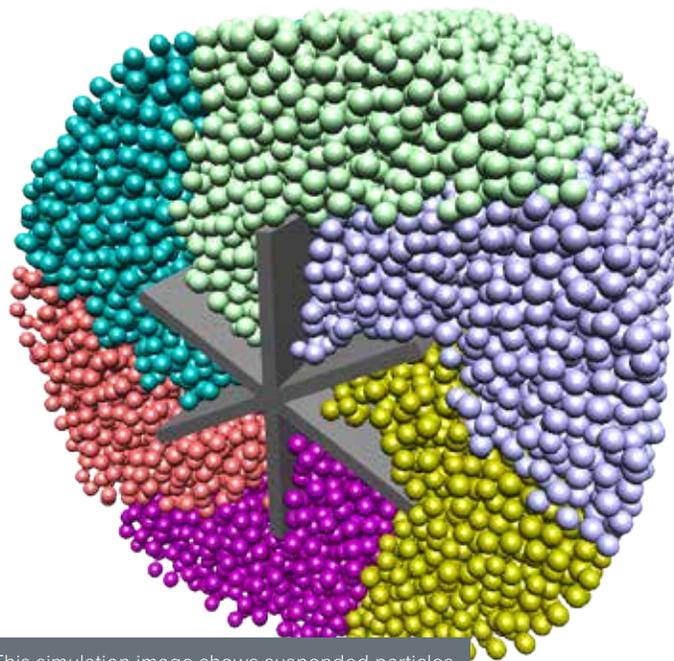
*Highly versatile and stone-like in strength, concrete is the most widely used fabricated building material, representing a multibillion dollar industry. Researchers from the National Institute of Standards and Technology (NIST) are using ALCF supercomputers to conduct large-scale simulations to advance the materials and measurement science of concrete.*

Given the critical importance of concrete to our nation's infrastructure, there is broad interest in making it a more sustainable material by improving its strength and durability, reducing the amount of greenhouse gas created during its production, and by finding new ways to recycle it. Architects also continue to push for enhanced workability that will allow concrete to flow more 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.

Due to the complex nature of concrete, which is a dense suspension comprised of water, cement, and an aggregate, such as sand or gravel, it is a challenge to accurately measure its rheological properties. Led by researchers from NIST, this project aims to advance the measurement science of concrete and to gain a fundamental understanding of how it flows. By simulating the flow of concrete in virtual rheometers, researchers are addressing the problem of relating measured quantities (e.g., torque and angular velocity measurements) to fundamental rheological properties (e.g., viscosity versus strain rate).

Researchers are combining data from these large-scale simulations with theoretical work and physical experiments to create Standard Reference Materials (SRMs) for concrete to allow for more accurate viscosity measurements. SRMs are certified and issued by NIST for use in quality control, regulatory compliance, product development, and scientific research for a wide variety

of materials, including steel, rubber, and plastics. The NIST research team is currently running simulations that will feed into the development of SRMs to calibrate rheometers for mortar (cement and sand) and concrete (mortar and gravel). The mortar SRM is expected to be available soon.



This simulation image shows suspended particles in a rheometer for NIST's proposed mortar SRM. The spheres, which are color coded by their starting location in the rheometer, are suspended in a cement paste with properties derived from NIST's cement paste SRM.

### Image Credit

Nicos Martys and Steven G. Satterfield, National Institute of Standards and Technology

**IMPACT** » The development of Standard Reference Materials (SRMs) will enable accurate predictions of the flow of concrete, which is essential to exploring the use of new, more environmentally friendly ingredients for concrete mixtures. Additionally, the SRMs will help to improve the workability of concrete by creating standardized measurements that will allow builders to request a specific concrete formulation with reliable, repeatable results.

TIMOTHY GERMANN

Los Alamos National Laboratory | tcg@lanl.gov

## Multibillion-Atom MD Studies of the Mechanical Response of Nanocrystalline Ta

*Los Alamos National Laboratory's Timothy Germann is using ALCF resources to conduct first-of-its-kind multibillion-atom simulations of Tantalum (Ta) nanocrystalline samples to determine how well the universal Hall-Petch relationship of steel strength and grain size holds up at or approaching micron-sized grain levels.*

Some 50 years ago, British metallurgists E. O. Hall and N. J. Petch independently discovered a striking relationship between the strength of steel samples and their grain size that has since been found to be universal; yield, flow, and fracture stresses for a wide variety of materials obey a similar relationship down to the smallest grain sizes.

However, this dependence eventually breaks down as the grain sizes approach the atomic (sub-nm) scale, since then materials become amorphous. Where this breakdown occurs remains difficult to probe experimentally due to challenges in synthesizing samples with well-controlled nanoscale grain sizes.

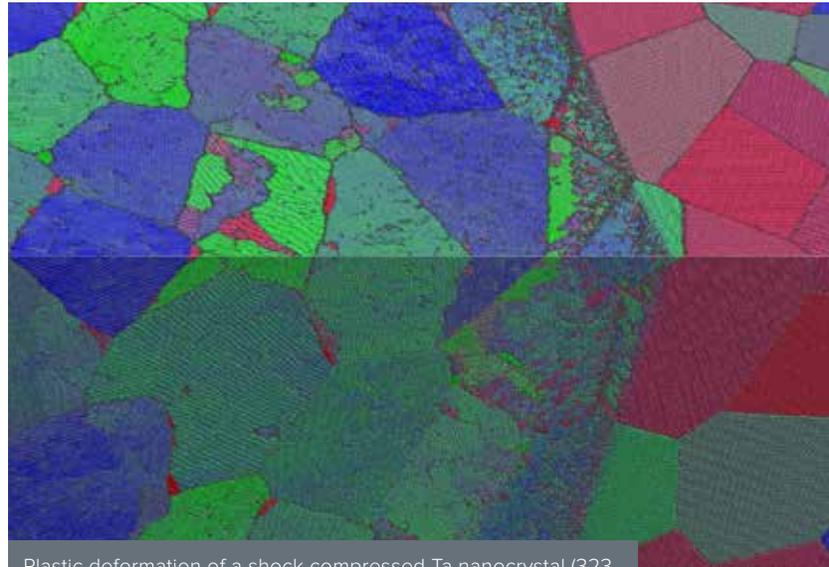
This project will conduct molecular dynamics (MD) simulations of compression under uniaxial strain of Tantalum (Ta) nanocrystalline samples, modeling grain sizes between 70 nanometers and 0.4 micrometers. Grain sizes of this order have never been modeled before with 3D MD simulations, and are the first step towards micron-size simulations.

Findings will be compared with experiments conducted at Argonne's Advanced Photon Source and SLAC's Linac Coherent Light Source to help fill the length and time scale gap between simulations and experiments.

These multibillion-atom simulations will allow researchers to model processes that can only be observed at larger grain sizes. By choosing a body-centered cubic (BCC) such as Tantalum, they

are studying a class of metals that have not been modeled as extensively as face-centered cubic metals, such as copper, aluminum, or nickel.

In demonstrating the Hall-Petch and inverse Hall-Petch relations for Ta, researchers expect the simulations to characterize the grain size at which changes occur both in tension and compression, and reveal the atomic mechanisms associated with such changes in BCC metals.



Plastic deformation of a shock-compressed Ta nanocrystal (323 million atoms, 50 nm grains). Atoms are colored according to their local crystallographic orientation; a shock wave is traveling from left to right, and has traveled about three-quarters of the sample length. A high density of defects form and are annealed behind the shock front, including both slip (appearing as individual points where the dislocation lines intersect the periodic boundaries on the top and front surfaces) and twinning (regions within a grain with a secondary, twin orientation color).

### Image Credit

Timothy Germann, Los Alamos National Laboratory;  
Ramon Ravelo, University of Texas, El Paso

**IMPACT** » The results of this research will have a significant impact in the materials science community, both because of the scale of the simulations proposed, the grain sizes investigated, and the physics insight that can be gained from such simulations.

GREGORY VOTH

The University of Chicago | gavoth@uchicago.edu

## Multiscale Modeling of Energy Storage Materials

*The development of advanced simulation methods will provide new tools for the design of next-generation batteries and fuel cells, which can help reduce U.S. reliance on fossil fuels. By combining a powerful multiscale modeling approach with the leading-edge resources at the ALCF, researchers are conducting transformative simulations to answer questions regarding the fundamental processes leading to improved energy storage technologies.*

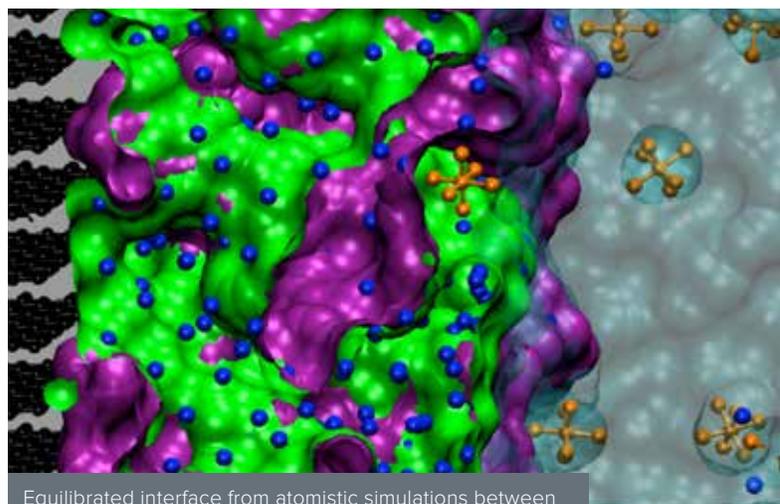
Led by Gregory Voth, a team of scientists from the University of Chicago and Argonne National Laboratory are carrying out simulations to advance the understanding of ion transport in the large-scale ion-exchange membranes of fuel cells and the solid electrolyte interphases of batteries. The primary goals of this work are to fill the existing gaps between experimental measurements and to provide detailed molecular insight into the processes governing charge transport in energy storage devices. The development of novel multiscale methods that are systematic across several length scales and predictive for materials design forms the first step in a potential feedback loop with ongoing experimental efforts.

As part of the project, researchers have conducted a systematic comparison of enhanced sampling techniques with Hamiltonian replica exchange, temperature replica exchange, simulated annealing, and temperature-accelerated molecular dynamics. Comparing potential energies and structural distributions, they identified distinct differences between the methods, implying that enhanced sampling schemes may prove a viable route to providing new insight into the performance of ion-exchange membranes.

The researchers also incorporated results from these equilibrated membrane morphologies into the design of both reactive models and large-scale coarse-grained models to understand how membrane structure relates to transport properties important

for fuel cell performance. For battery systems, the research team is using recently developed replica-exchange enhanced sampling algorithms to calculate ion transport pathways of lithium-ion insertion from the bulk electrolyte to the solid electrolyte interphase in the presence of polarizable electrodes with applied bias potentials.

With recent algorithmic advances, the work accomplished here provides a strong foundation for future leadership computing projects to further understand charge transport phenomena in batteries, fuel cells, and at metal-aqueous interfaces.



Equilibrated interface from atomistic simulations between a solid electrolyte interphase consisting of lithiated components (green and purple) and electrolyte (aqua), where lithium (blue) and hexafluorophosphate (orange) ions are found at the interface.

Image Credit

Ryan Jorn, Argonne National Laboratory/The University of Chicago

**IMPACT** » The development of next-generation batteries and fuel cells can provide viable clean-energy alternatives for replacing internal combustion engines in automobiles and powering personal electronics. Additionally, the project's multiscale and reactive simulation methods carry the potential for significant impacts in the chemistry, biology, and materials science communities.

DARIO ALFÈ

University College London | d.alfè@ucl.ac.uk

## Non-Covalent Bonding in Complex Molecular Systems with Quantum Monte Carlo

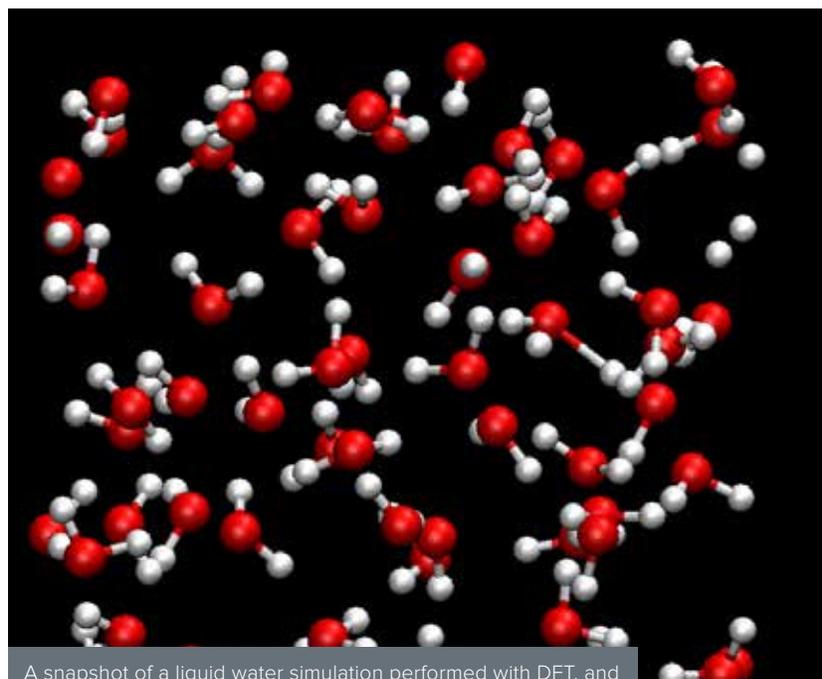
*Ubiquitous and essential to life on earth, water is still not fully understood by modern science. Using the highly accurate electronic structure theory method, quantum Monte Carlo (QMC) and ALCF resources, researchers are modeling liquid water with unprecedented accuracy, generating results applicable to aqueous chemistry across scientific disciplines.*

Density functional theory (DFT) is a modeling method used to describe the electronic structure of systems at a quantum mechanics level with a relatively low computational cost. However, it fails to accurately describe strongly correlated systems, weak interactions, and non-covalent interactions on which water molecules depend.

Over the course of three previous INCITE awards, researchers led by Dario Alfè, University College London, demonstrated that the use of QMC methods to solve electronic structure problems is far better suited for studies of strongly correlated systems: This many-body theory carries all the ingredients to calculate the exact exchange and correlation energies at the chemical accuracy ( $\sim 1$  kcal/mol), making it a key method to investigate the role of hydrogen bonding and other non-covalent interactions in the cohesion of molecular materials.

Because QMC is more computationally expensive, the thrust of Alfè's latest INCITE project aims to address practically important scientific problems concerning non-covalent bonding for a range of paradigm molecular materials, as well as QMC benchmarking for water clusters in thermal equilibrium. The group will use QMC calculations to create an extensive database of energy benchmarks for molecular systems bound by non-covalent forces, particularly hydrogen bonding and van der Waals dispersion. The benchmarks will be used to develop improved versions of DFT and to assess and improve parameterized force fields for molecular systems.

The accurate description of cooperative many-body effects in hydrogen bonding, and the inclusion of screening and many-body effects in dispersion corrected versions of DFT are important aspects of this work. Building on past INCITE success, this groundbreaking research on non-covalent bonding is designed to exploit the greatly increased power of leadership-class systems.



A snapshot of a liquid water simulation performed with DFT, and a box including 64 molecules. By performing benchmark QMC calculations on snapshots of this type, researchers are able to ascertain DFT errors.

Image Credit

Dario Alfè, University College London

**IMPACT** » The benchmarks derived from this INCITE project will improve research capabilities for a diverse array of scientific problems—from the nucleation and growth of ice in the upper atmosphere or on aircraft wings, to the dynamics of glaciers and the adsorption of water on surfaces.

PRIYA VASHISHTA

University of Southern California | priyav@usc.edu

## Petascale Simulations of Stress Corrosion Cracking

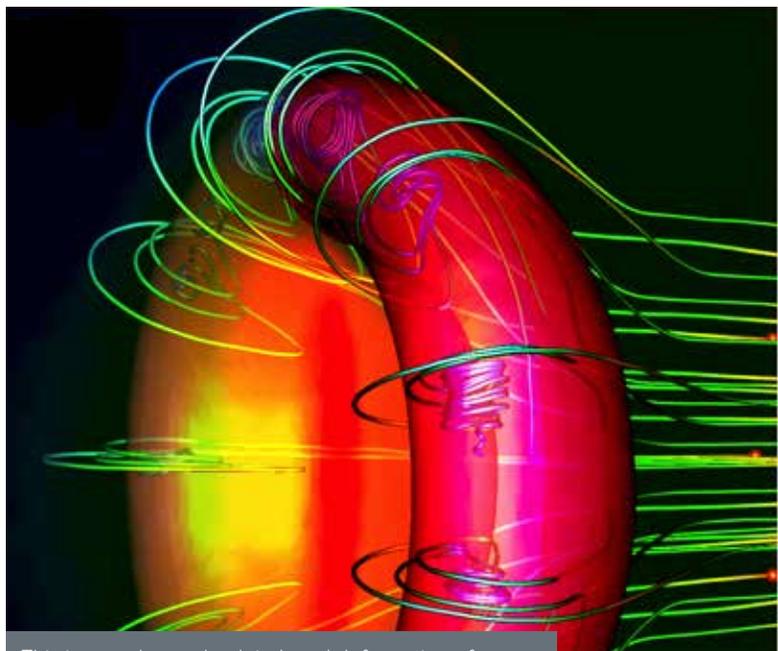
*The performance and lifetime of materials widely used in energy and nuclear technologies are often limited by corrosion under stress loads. Simulations performed at the ALCF are revealing the atomistic mechanisms that control stress-induced corrosion within nuclear reactors—which is key to understanding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

Molecular dynamics (MD) simulation studies of the complex electrochemical processes that lead to the degradation of nickel-based alloys and silicate glass are among the largest and longest running projects at the ALCF. A team from the University of Southern California led by Priya Vashishta has been building an understanding of how a small amount of impurities within these materials can cause fractures and potentially catastrophic failures.

This year, using ALCF's new petascale system, Mira, the team will conduct quantum mechanical (QM) MD simulations to reveal the atomistic mechanisms underlying SCC to prevent such cracking and to predict the lifetime beyond which SCC may cause failure.

This research requires huge computational resources to explore the atoms and molecules interacting over a time span that matches the kinetics of the natural process. The team aims to exploit Mira's capabilities to attempt simulations that encompass several billion atoms for nanosecond timescales, and several million atoms for microseconds.

One of the simulation challenges lies at the QM/MD interface where the charge flow and bond-breaking events at the QM electronic scale are difficult to translate to the rigid bonding/fixed point charges regime usually used in the effective force-field MD approach. To facilitate this translation process, the team is adding another computational method – a first-principles-based reactive force field – that has the ability to handle charge flow by employing a chemical environment-dependent model.



This image shows shock-induced deformation of a gas-filled nanobubble and the pressure distribution in a cross-section of the system. Snapshot of vortex rings at the periphery of the nanobubble. Streamlines represent molecular velocities.

### Image Credit

Rajiv K. Kalia, Aiichiro Nakano, Adarsh Shekhar, and Priya Vashishta, Collaboratory for Advanced Computing and Simulations, University of Southern California

**IMPACT** » To prevent stress corrosion cracking and to predict the lifetime beyond which it may cause failure requires an understanding of the conditions influencing its initiation, dynamics, and growth rates. A holistic understanding of these mechanisms will aid in the design of safe, next-generation nuclear reactors and nuclear waste containers, and may ultimately be applied to other corrosion problems that costs the nation an estimated 3% of its gross domestic product annually.

GIULIA GALLI

University of California, Davis | gagalli@ucdavis.edu

## Vibrational Spectroscopy of Liquid Mixtures and Solid/Liquid Interfaces

*Understanding pure water is an essential prerequisite for grasping the behavior of aqueous solutions interacting with surfaces. This ongoing INCITE project led by Giulia Galli of the University of California, Davis, makes use of ALCF resources to better understand these interactions, including the spectroscopic and dielectric properties of solutions.*

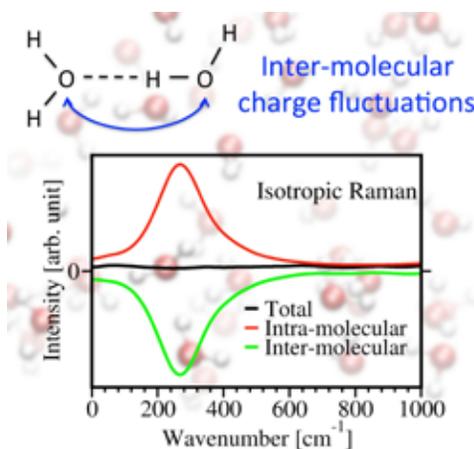
With earlier INCITE allocations, several accurate flavors of density functional theory (e.g., hybrid functional PBE0 and van der Waals functionals) were used to investigate water and ions in water with atomistic resolution. Researchers carried out *ab initio* molecular dynamics (MD) calculations using the Qbox code on ALCF's Intrepid, an IBM Blue Gene/P, to determine the structural and vibrational properties of liquid water and simple solvated anions at several temperatures. Findings suggest that high-performance computing can significantly improve the predictive power of the structural and electronic properties of aqueous environments.

Mira's Blue Gene/Q architecture permits a large increase in computational power and allows researchers to fully exploit the parallel efficiency of the Qbox code. A new version of the code added density functional perturbation theory capabilities, which were used to compute, for the first time, the Raman spectra of liquid water from first principles.

A comprehensive analysis of the water Raman spectra, based on the calculation of maximally localized Wannier functions and effective molecular polarizabilities, found an occurrence of charge fluctuations accompanying hydrogen bonding stretching modes in liquid water, despite the absence of any Raman activity. Such an analysis highlighted some fundamental properties of hydrogen bonding in water and is of general applicability to molecular solids and molecular liquids.

The team is continuing its research to understand the electronic properties and spectroscopic signatures of anions in water, particularly chloride and sulfate, and to couple the Raman capability of Qbox with algorithms recently implemented in the code to carry out efficient and scalable *ab initio* MD simulations with hybrid functionals.

Future work includes expanding computational spectroscopy capabilities to study water at solid oxide interfaces and water/low Z liquid mixtures on solid surfaces.



Computed low-frequency portion of the isotropic Raman spectrum (black line) of liquid water, showing the intra- and inter-molecular contributions indicative of inter-molecular charge fluctuations between hydrogen bonded molecules. Calculations were carried out using *ab initio* molecular dynamics, coupled with density functional perturbation theory, using the Qbox code.

Image Credit

Quan Wan, University of California, Davis

**IMPACT** » This research helped determine the fundamental effects responsible for several observed phenomena (e.g., spectroscopic signatures of water and solvated ions) that cannot be probed directly in experiments. The results represent a stepping stone towards understanding the structure and dissociation of water in contact with solid oxide surfaces, which are of interest to many energy- and environmental-related problems, encompassing solar to fuel production, and reaction of water on rocks in the deep Earth.

ROBERT A. DISTASIO, JR.  
Princeton University | [distasio@princeton.edu](mailto:distasio@princeton.edu)

## Ab Initio Quantum Liquid Water and Aqueous Ionic Solutions

*A highly accurate and detailed understanding of the microscopic structure of liquid water is of great importance to a number of scientific fields. This project will address key challenges of ab initio molecular dynamics (AIMD) liquid water simulations and then apply those results to a series of benchmark simulations on liquid water and ionic solutions, which could impact renewable energy research.*

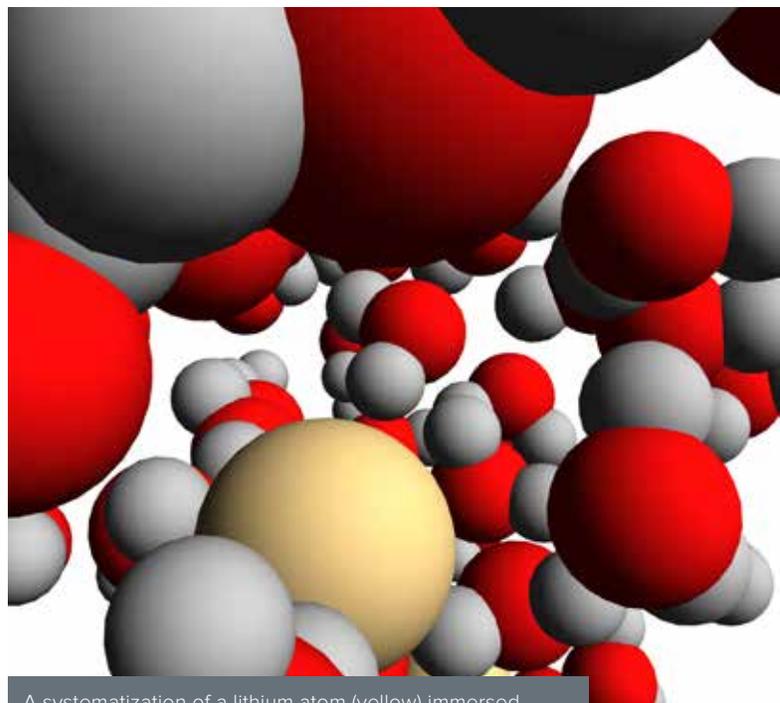
Water is arguably the most important molecule on Earth. While a single water molecule has a simple and well-known structure, the liquid phase of water has an intricate and disordered microscopic structure that has been difficult to study both experimentally and theoretically.

At present, there is no experimental methodology available to directly obtain the real-space microscopic structure of liquid water. Computer-based simulations can furnish such structural information in a relatively straightforward manner, however, the state-of-the-art computational approach for microscopic molecular modeling, or AIMD, has severe limitations when applied to liquid water.

The result of a collaboration between ALCF and Princeton University, this project aims to perform highly accurate benchmark atomistic simulations of liquid water that extend the potential of current state-of-the-art technologies by directly addressing earlier approximation limitations. In addition, researchers will investigate a series of fundamental aqueous ionic solutions, starting with the hydronium and hydroxide ions (i.e., models of acidic and basic conditions, respectively), and continue with ions most relevant to the design of novel clean energy materials.

These simulations will require a substantial amount of computational resources that can only be accomplished by combining new algorithmic advances with the efficient utilization of the high-performance massively parallel Blue Gene/Q platform at ALCF.

Given the potential for obtaining highly refined structural details, long standing questions regarding the coexistence of low- and high-density liquid water phases can be resolved by investigating both the real-time dynamics and the corresponding inherent potential energy surface.



A systematization of a lithium atom (yellow) immersed into an aqueous environment. The water molecules are simulated with periodic continuous properties and light nuclei with quantum effects.

Image Credit  
Alvaro Vazquez-Mayagoitia, Argonne National Laboratory

**IMPACT** » These simulations will provide detailed knowledge of the coordination shell structure and atomic pair correlation functions of liquid water and aqueous ionic solutions with unprecedented accuracy, and address important renewable energy research issues, particularly the understanding and rational design of aqueous ion batteries. The results will be stored in a publicly available structural database to serve as an invaluable resource for further simulation developments.

JAMES AMUNDSON  
Fermilab | amundson@fnal.gov

## Accelerator Simulations for the Intensity Frontier of Particle Physics

*Particle accelerators are an enabling technology for both basic and applied sciences. Researchers from Fermilab are using ALCF resources to perform complex accelerator simulations aimed at reducing the risks and costs involved in developing the world's highest intensity particle beams.*

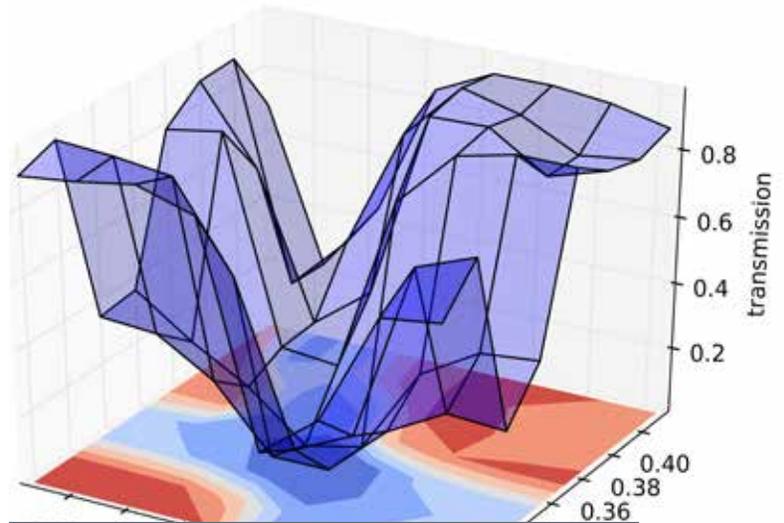
Led by James Amundson, this project will advance particle physics by assisting Fermilab in developing world leadership in the Intensity Frontier (i.e., the production of high-intensity proton beams) through the laboratory's Project X, a proposed proton accelerator complex. These proton beams will, in turn, be used to create the world's best neutrino, kaon, and muon beams, allowing particle physicists to dramatically expand our understanding of rare and hard-to-detect phenomena.

The research will also advance accelerator science by conducting simulations at an unprecedented level of detail, enabling scientists to test simplified theoretical models of complex phenomena. With access to Mira, the research team is able to combine multiple complex physical effects, including magnet nonlinearities, space charge, and wakefields in single high-precision simulations.

A broad course of simulations is necessary to aid in the plans for upgrades to the existing Fermilab Booster and Main Injector, which form crucial parts of Project X in various stages. The massive simulations conducted on Mira will be the most detailed accelerator beam dynamics simulations ever performed.

Already, the research team has performed full 84-bunch simulations of the Fermilab Booster with realistic wake fields and space charge that are in excellent agreement with the experimental measurements of the coherent tune shift, kick

decoherence, and multibunch transverse instabilities. Massive parameter scans of the Main Injector have produced a map leading to the optimal working point for the Project X era. The new working point will be used in future Main Injector simulations.



Results of a massively parallel parameter scan of transmission in the Fermilab Main Injector as a function of ( $Q_x$ ,  $Q_y$ ) working point under Project X conditions.

Image Credit  
Eric Stern, Fermilab

**IMPACT** » Results from these simulations will be used to minimize the risk of problems and maximize operating efficiency of high-intensity proton beams in the Project X era, enabling explorations of the Intensity Frontier of particle physics. The work will also help to advance accelerator physics by furthering our understanding of intensity-dependent effects in particle accelerators.

SALMAN HABIB

Argonne National Laboratory | habib@anl.gov

## Computing the Dark Universe

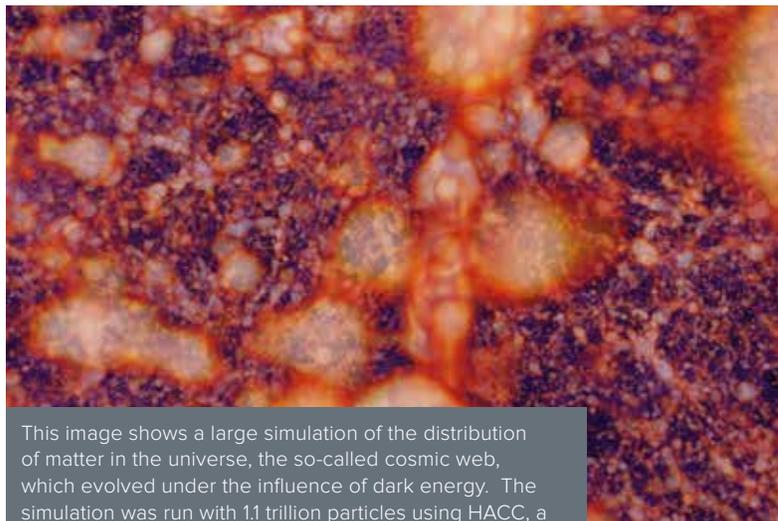
*In the standard model of cosmology, dark energy and dark matter together account for 95% of the mass energy of the universe; however, their ultimate origin remains mysterious. The ALCF is allocating significant supercomputing resources toward unraveling one of the key puzzles—the nature of the dark 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 conducting some of the most detailed, largest-scale simulations of structure formation in the universe ever performed. The simulations target an improved understanding of the nature of dark energy and dark matter.

Led by physicist Salman Habib of Argonne National Laboratory, the project seeks to establish a computation-based discovery capability for critical cosmological probes by exploiting next-generation supercomputing architectures. Using Mira, the research team is performing simulations to study the clustering of matter in the universe and the evolution of the clustering as a multi-dimensional cosmological probe.

Ultimately, the researchers hope to resolve galaxy-scale mass concentrations over observational volumes representative of state-of-the-art sky surveys. 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. To achieve these goals, the project team is focusing their effort on the HACC (Hardware/Hybrid Accelerated Cosmology Code) framework for N-body simulations.

For the second consecutive year, a paper on the team's research earned a finalist designation for the ACM Gordon Bell Prize and will be included in the Gordon Bell Prize sessions at SC13. The paper, "HACC: Extreme Scaling and Performance Across Diverse Architectures," discusses the application of HACC, which has attained unprecedented levels of scalable performance on supercomputers at several national laboratories, including Mira at the ALCF.



This image shows a large simulation of the distribution of matter in the universe, the so-called cosmic web, which evolved under the influence of dark energy. The simulation was run with 1.1 trillion particles using HACC, a new simulation framework developed with the challenges of future supercomputing architectures in mind.

### Image Credit

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

**IMPACT** » This work will allow researchers to compare the best telescope observations of the sky to the structures found in computer simulations of the universe, thereby testing various cosmological models and possible explanations for dark energy and candidates for dark matter. 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.

WILLIAM TANG

Princeton Plasma Physics Laboratory | tang@pppl.gov

## Kinetic Simulations of Fusion Energy Dynamics at the Extreme Scale

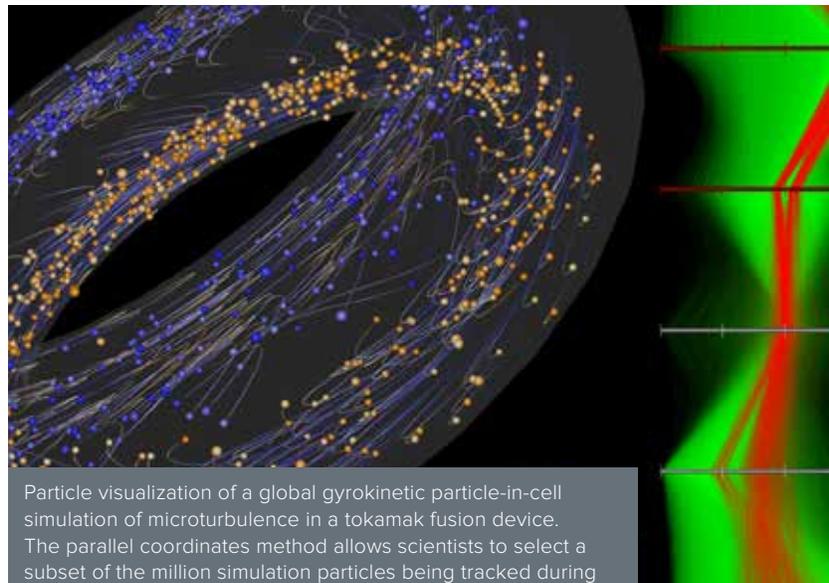
*The quest for alternatives to fossil fuels to meet our energy needs is generating increasing interest in nuclear fusion. William Tang and researchers from the Princeton Plasma Physics Laboratory are using ALCF resources to simulate complex dynamical systems to contribute new insights to the design and operation of fusion power reactors.*

Building the scientific foundations needed to develop fusion power demands high-physics-fidelity predictive simulation capability for magnetically confined fusion energy (MFE) plasmas. To do so in a timely way requires using the power of modern supercomputers to simulate the complex dynamics governing MFE systems, including ITER, a multibillion dollar international burning plasma experiment.

Unavoidable spatial variations in such systems produce microturbulence, which can significantly increase the transport rate of heat, particles, and momentum across the confining magnetic field in tokamak devices. Since the balance between these energy losses and the self-heating rates of the actual fusion reactions will ultimately determine the size and cost of an actual fusion reactor, understanding and possibly controlling the underlying physical processes is key to achieving the efficiency needed to help ensure the practicality of future fusion reactors.

The goal of this project is to gain new physics insights on MFE confinement scaling by making effective use of the ALCF's Mira supercomputer. The knowledge gained will also address the key question of how turbulent transport and associated confinement characteristics scale from present generation devices to the much larger ITER-scale plasmas. This involves the development of modern software capable of using low memory-per-core leadership-class supercomputers to carry out reliable first principles-based simulations of multiscale tokamak plasmas.

The fusion physics challenge is that the key decade-long MFE estimates of confinement scaling with device size (the so-called "Bohm to Gyro-Bohm" rollover trend caused by the ion temperature gradient instability) demands much higher resolution to be realistic and reliable. However, results from this project have revealed that this rollover is much more gradual than established earlier (in far lower resolution, shorter duration studies) with magnitude of transport now reduced by a factor of two.



Particle visualization of a global gyrokinetic particle-in-cell simulation of microturbulence in a tokamak fusion device. The parallel coordinates method allows scientists to select a subset of the million simulation particles being tracked during the calculation. Each horizontal line on the right represents a different particle parameter, such as the position, velocity, statistical weight, etc. The red lines indicate the particles being selected. In this example, the particles are being chosen for their magnetically trapped state and for their parallel velocities close to zero. This occurs in tokamak devices when magnetically trapped ions are at the tip of their orbits and are in the process of changing direction.

### Image Credit

Chad Jones and Kwan-Liu Ma, University of California, Davis;  
Stephane Ethier, Princeton Plasma Physics Laboratory

**IMPACT** » Results from this project will provide transformational scientific insights to help accelerate progress in worldwide efforts to harness the power of nuclear fusion as an alternative to fossil fuels.

PAUL MACKENZIE

Fermilab | mackenzie@fnal.gov

## Lattice QCD

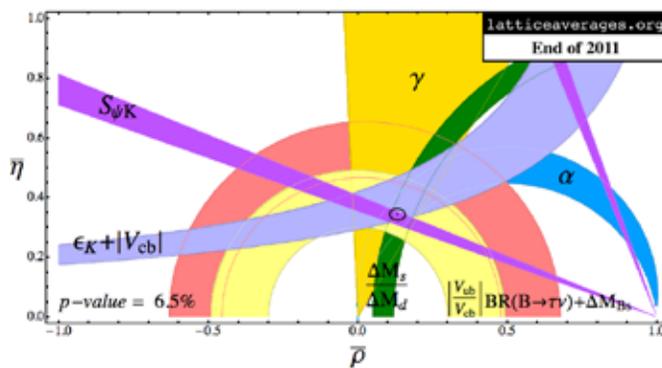
*The United States Quantum Chromodynamics (USQCD) Collaboration carries out extensive lattice gauge theory calculations at the ALCF. This theory describes how quarks and gluons, the fundamental entities of protons and neutrons, help form atomic nuclei. These efforts are the heart of DOE's large experimental 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. For this project, researchers are advancing exploration in lattice QCD and other strongly coupled field theories of importance to the study of high energy and nuclear physics. The team, led by Paul Mackenzie of Fermilab, is using ALCF supercomputers to generate gauge configurations with up, down, strange, and, in one case, charmed quarks on lattices that are sufficiently fine grained and have sufficiently small up and down quark masses to enable full control of systematic errors for a number of key quantities.

A major goal of the project is to determine a number of underlying parameters of the standard model of subatomic physics, including the masses of the quarks, the strong coupling constant and elements of the Cabibbo-Kobayashi-Maskawa (CKM) matrix. Other goals are to determine the mass spectrum of strongly interacting particles; to calculate the electromagnetic properties of low-lying baryons and meson radiative transitions; to obtain a quantitative understanding of the internal structure of the nucleon; to predict the interactions among strongly interacting particles; and to investigate the properties of strongly interacting matter at high temperatures and densities.

In 2013, the MILC Collaboration (a subgroup of USQCD) published the first work of analysis from the Highly Improved Staggered Quark (HISQ) configurations recently generated at the ALCF in Physical Review Letters. A calculation of the ratio of leptonic decay constants  $f_K/f_\pi$  made possible a precise determination of the ratio of CKM matrix

elements  $|V_{us}|/|V_{ud}|$  in the standard model, and places a stringent constraint on the scale of new physics that would lead to deviations from unitarity in the first row of the CKM matrix. This result is the most precise lattice-QCD determination of  $f_K/f_\pi$ , with an error comparable to the current world average.



The current summary of Lattice QCD averages.

### Image Credit

Ruth S. Van de Water, Brookhaven National Laboratory; E. Lunghi, Indiana University; Jack Laiho, Washington University in St. Louis/University of Glasgow (Phys.Rev.D81:034503,2010)

**IMPACT** » The project's goals are essential for meeting the milestones set out by the DOE's Office of Science. In addition to directly advancing the understanding of physics, this work is providing crucial high-precision lattice QCD calculations needed for new or in-progress experiments and for analyzing results from completed experiments.

DON LAMB

The University of Chicago | lamb@oddjob.uchicago.edu

## Next-Generation Petascale Simulations of Type Ia Supernovae

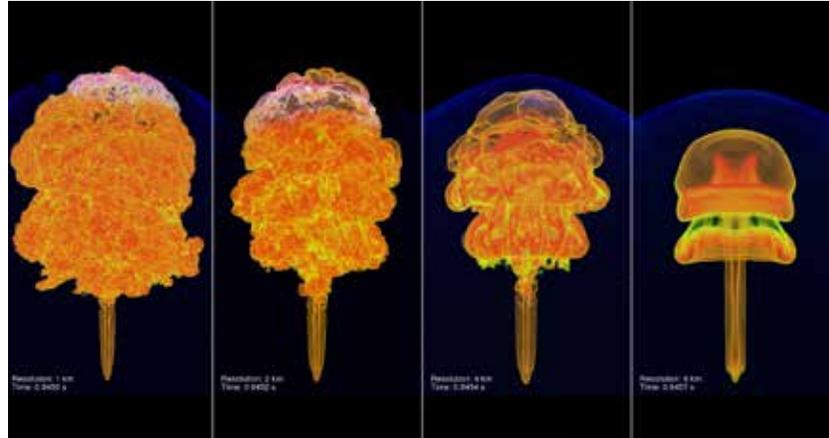
*Among the brightest and most energetic events in the universe, Type Ia supernovae (SNe Ia) are important tools for cosmology research. Observations of SNe Ia have revealed that the expansion of the universe is accelerating and led to the discovery of dark energy. This project is leveraging ALCF supercomputers to better understand these explosions and improve their use as standard candles (objects with known luminosities).*

The scientific goal of this INCITE project is to better understand SNe Ia explosions, and by doing so, allow astronomers to enhance the accuracy with which these events can be calibrated as standard candles. Improving the use of SNe Ia as standard candles will enable a better determination of the properties of dark energy.

The research team, led by Don Lamb at the University of Chicago, is accomplishing this goal by using the FLASH code and the petascale computing made possible by Mira to perform higher fidelity simulations of SNe Ia. Their next-generation SNe Ia models have many features that will contribute to better simulation results, including much higher spatial resolution during the early phase of buoyancy-driven turbulent nuclear combustion, when it is required. The inclusion of the effect of convective turbulence on the thermonuclear flame provides a more accurate thermonuclear burning rate and enables researchers to achieve convergence of this rate with resolution.

The models also have much larger spatial domains that allow the explosion phase to be followed to the point where expansion of the ejecta is fully homologous, increasing the accuracy of the predicted light curves and spectra. Additionally, having 10+ M Lagrangian tracer particles placed proportional to volume and mass is resulting in a more accurate determination of the elements produced in the explosion, and therefore more accurate light curves and spectra.

The higher fidelity simulations will enable researchers to explore the effect of different initial conditions on the explosion. These initial conditions include the mass of the progenitor white dwarf star, the location of the ignition point, and uniform and differential rotation of the progenitor.



This image is from the simulation "Flame Bubble in an Open Computational Domain for Constant Flame Speed, Gravitational Acceleration, and Changing Density."

### Image Credit

Ray Bair, Randy Hudson, John Norris, Michael E. Papka, and Katherine Riley, Argonne National Laboratory; Alan Calder, Stony Brook University; Dean Townsley, University of Alabama; Anshu Dubey, Jonathan Gallagher, Carlo Graziani, George Jordan, Don Lamb, Dongwook Lee, Min Long, Daan van Rossum, and Klaus Weide, The University of Chicago; Robert Fisher, University of Massachusetts at Dartmouth

**IMPACT** » The project's high-fidelity simulations are resulting in a better understanding of Type Ia supernovae. This will enhance their use as standard candles, giving astronomers an improved tool to investigate dark energy, one of the most compelling problems in the physical sciences.

SEAN COUCH

University of Chicago | smc@flash.uchicago.edu

## Petascale Simulation of Magnetorotational Core-Collapse Supernovae

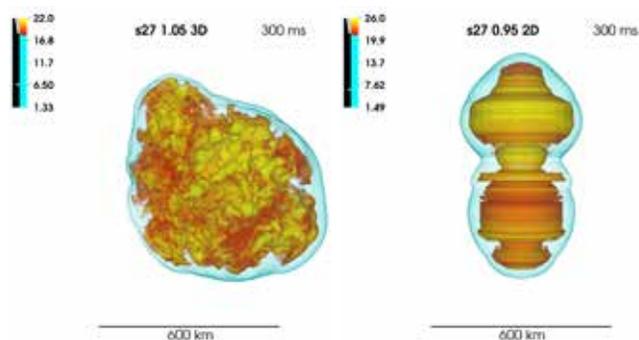
*A team of researchers led by Sean Couch at the University of Chicago, is using the ALCF's cutting-edge resources to conduct first-ever simulations for exploring the influence of magnetic fields and rotation on the mechanisms that reverse stellar core collapse and drive a supernova. These groundbreaking 3D 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.

Couch has used ALCF resources to conduct a series of high-resolution 3D neutrino-radiation-hydrodynamic simulations of the collapse of massive stellar cores using the FLASH multiphysics, adaptive mesh refinement simulation framework. These simulations have shown, for the first time ever, that realistic asphericity in the progenitor star can trigger shock revival in 3D simulations of CCSNe. This work points the way to a plausible, potentially robust explosion mechanism.

In 2012, Couch completed a sequence of 1D, 2D, and 3D simulations on the ALCF's Intrepid that showed the substantial differences in the behavior of the 2D and 3D buoyant plumes—the absence of forced-symmetry in 3D makes a tremendous impact on the character of the shock motion and the development of neutrino-driven convection and turbulence. These findings, published in *Astrophysical Journal* in July 2013, add important evidence regarding the significant differences between 2D and 3D CCSN simulations and further emphasize the need for efforts within the community to be fully 3D.

Threading and single core improvements made to FLASH through the ALCF's Early Science Program have reduced the time to solution by a factor of four. Additionally, new physics packages for treating neutrino transport and a completely re-written gravity solver have dramatically increased the fidelity of Couch's simulations.



Volume renderings of the specific entropy from 3D and 2D simulations of the explosion of a 27-solar mass star. The surface of the outgoing shock wave is visible as the faint blue surface. The dramatic differences between 2D and 3D underscore the necessity of studying supernovae in 3D.

Image Credit

Sean Couch, The University of Chicago

**IMPACT** » Much of the matter that makes up our planet was produced billions of years ago by core-collapse supernovae—the explosions that mark the deaths of massive stars. Researchers are using the resources of the ALCF to accurately simulate these explosions to provide scientists with a greater understanding of the chemical evolution of the universe.

STEVEN LANGER

Lawrence Livermore National Laboratory | langer1@llnl.gov

## Simulation of Laser-Plasma Interaction in National Ignition Facility Experiments

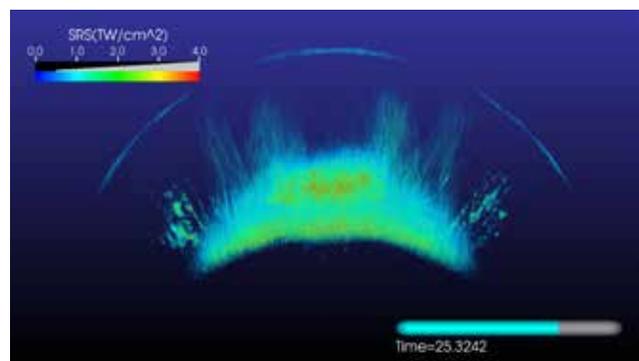
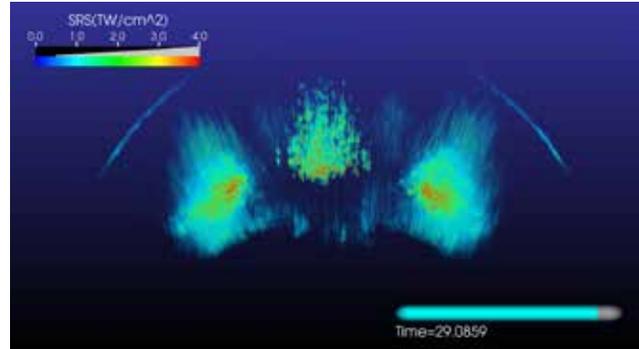
*Scientists have been working to achieve self-sustaining nuclear fusion and energy gain in the laboratory for more than half a century. Researchers at Lawrence Livermore National Laboratory have been tasked with achieving ignition at the National Ignition Facility (NIF) using the inertial confinement fusion process to bring the goal of fusion energy closer to realization.*

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, and environmentally sustainable energy. NIF is currently carrying out experiments to compress a mixture of deuterium and tritium to temperatures and densities high enough that fusion ignition occurs.

The laser intensity at NIF is high enough that some of the laser energy backscatters off the target. Previous research quantified how overlapping beam quads (groups of 2x2 beams propagating in the same direction) impacted backscatter at NIF. This demonstrated that overlapping quads could share a reflected light wave, thereby enhancing reflectivity. The simulations produced results similar to National Ignition Campaign experimental data and have allowed identification of several areas where the simulations need to be improved.

This project is performing a simulation of backscattered light generated from three overlapping quads over the full volume in which laser-plasma interactions (LPI) are expected to occur. The studies are producing synthetic data to be compared with experimental data from NIF. The simulations are also providing detailed information about where LPI occurs and the interaction between stimulated Raman scattering and stimulated Brillouin scattering; the two main plasma instabilities acting in these experiments. This research will provide a better understanding of the conditions where overlapping quads generate

more backscattered light than from the quads acting independently.

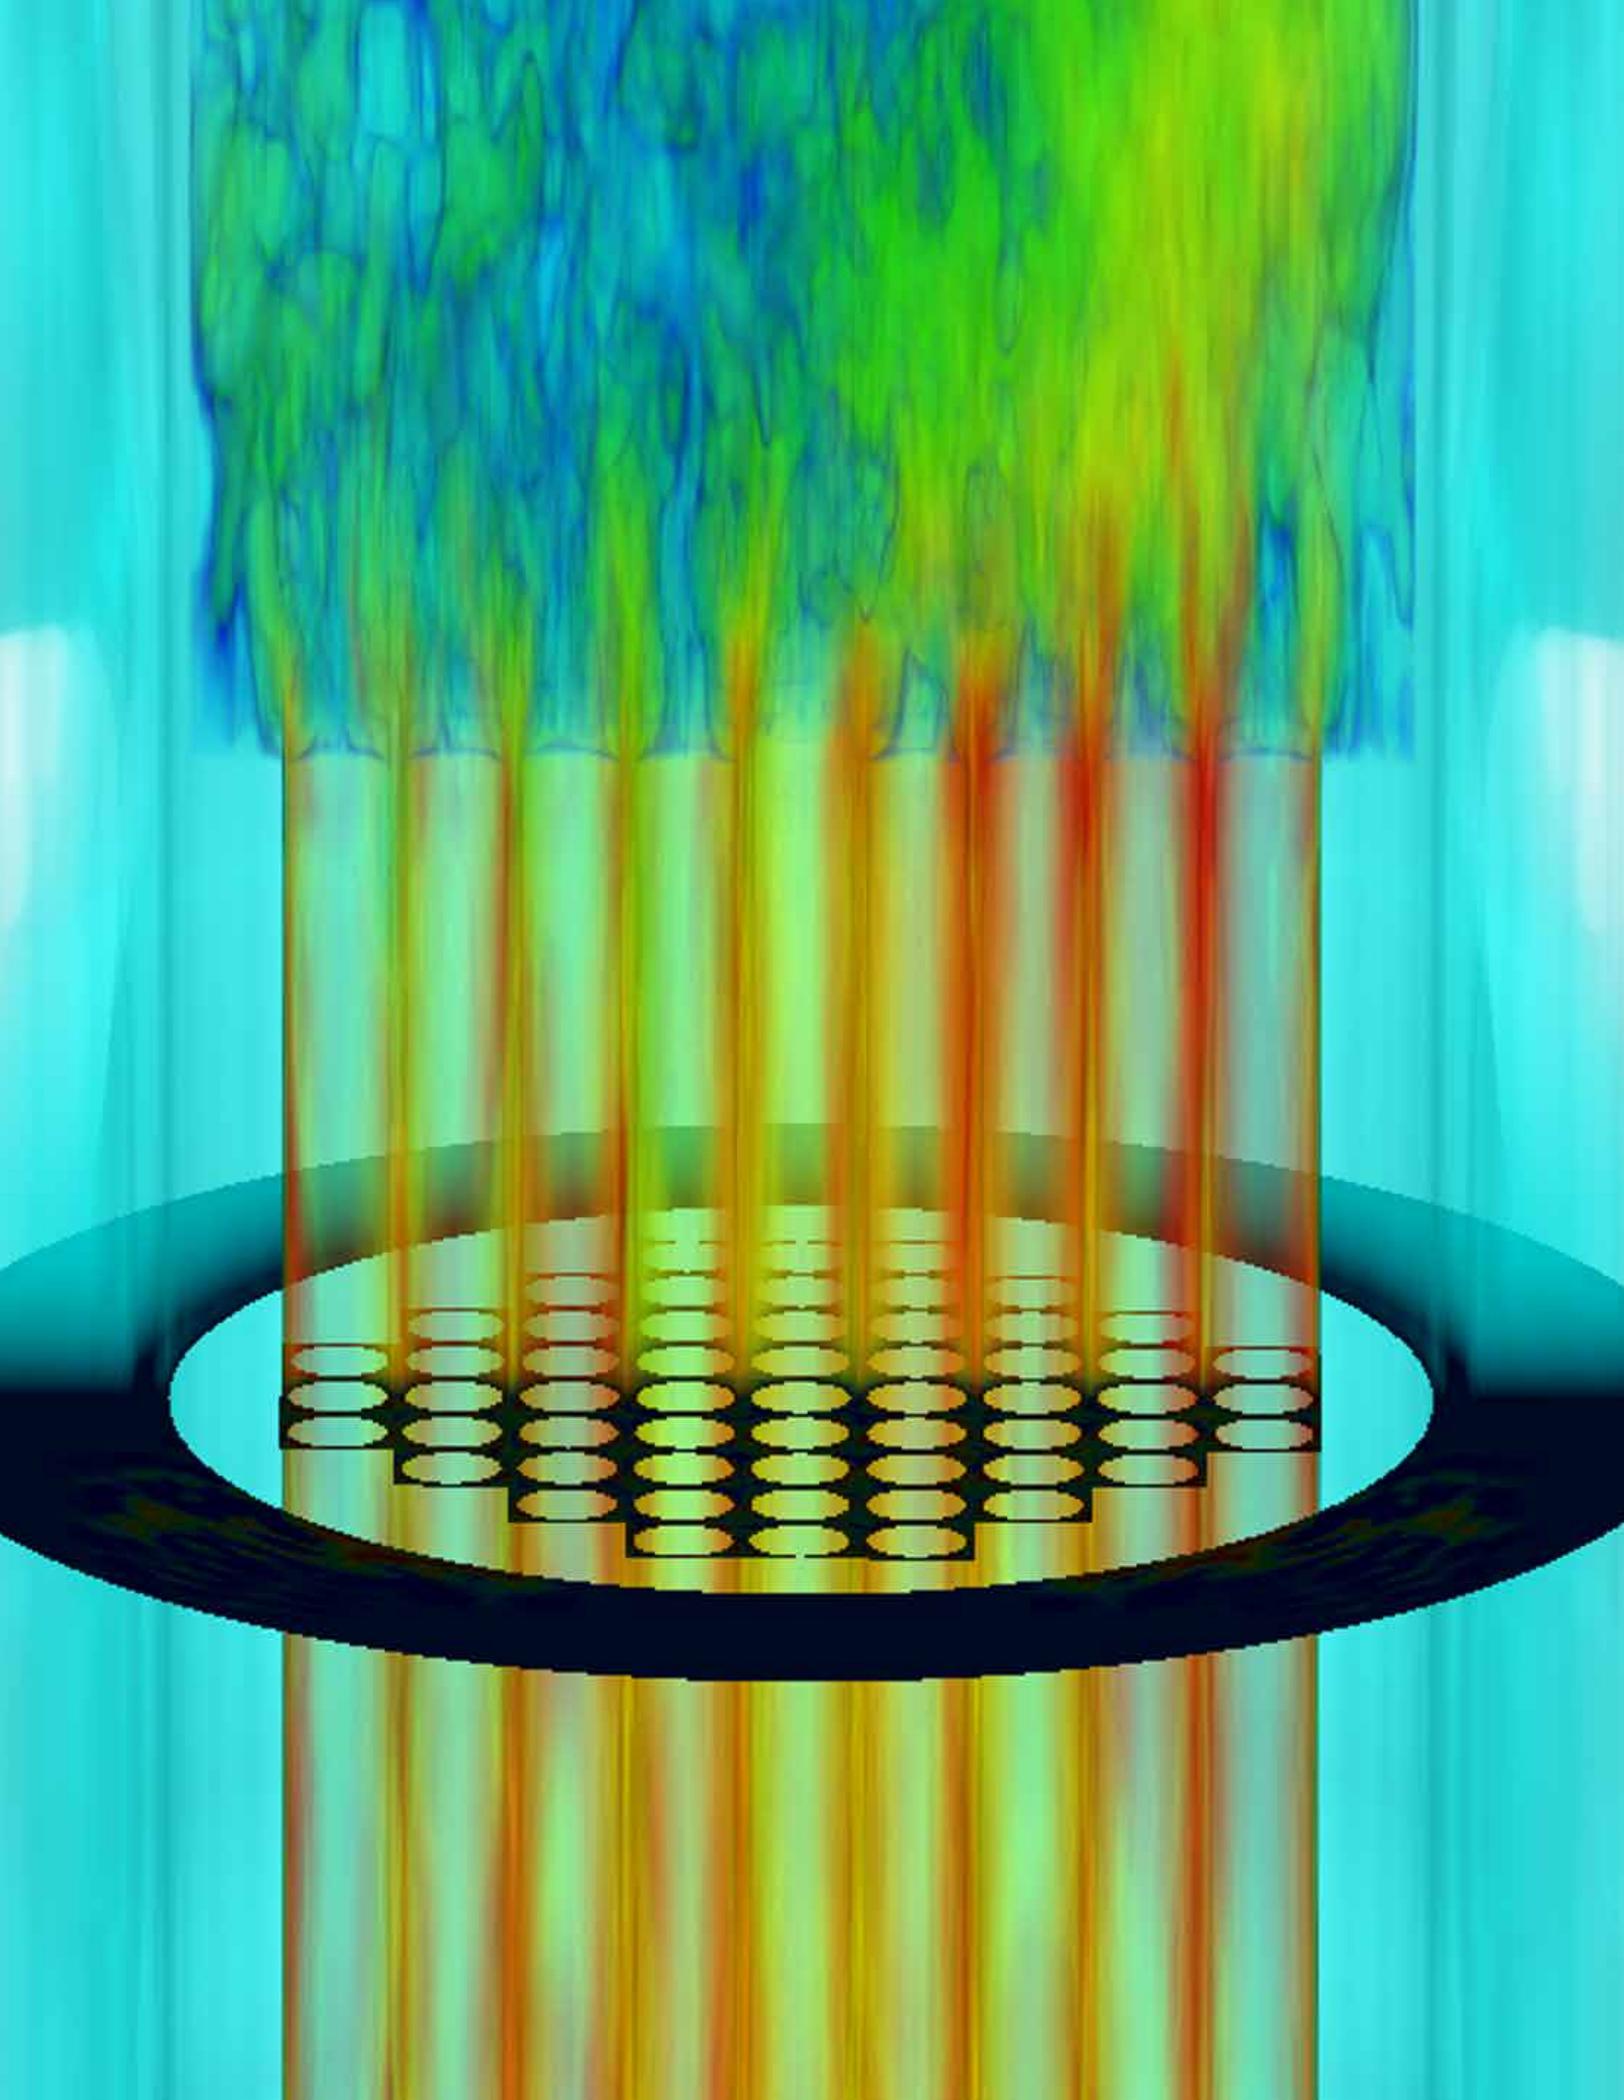


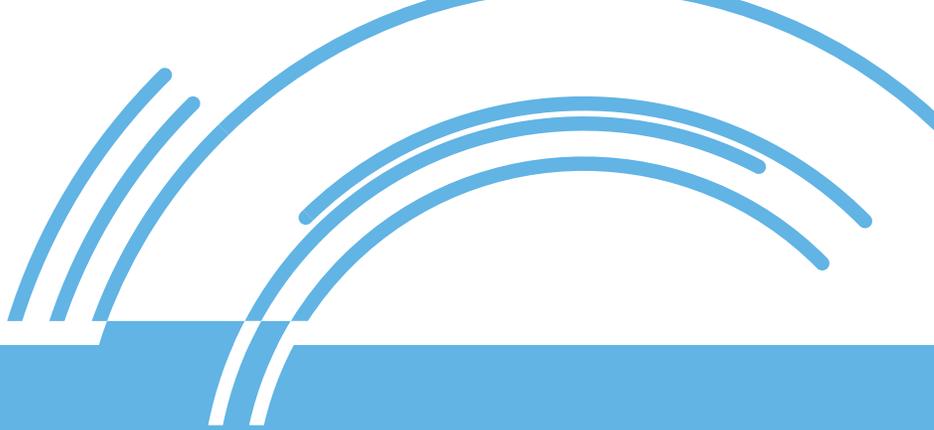
The top image is a volume visualization of the laser light that has been backscattered by Stimulated Raman Scattering in a simulation of a National Ignition Facility inertial fusion experiment. At this point in time, light is scattering independently from the three laser quads (a quad is a group of 2x2 beams propagating in the same direction). The bottom image is a volume visualization from the same simulation, but at a different time. The band of scattered light near the bottom of the simulation extends across all three quads and is an example of "cooperative backscattering."

### Image Credit

Eric Brugger and Steven Langer, Lawrence Livermore National Laboratory

**IMPACT** » Simulations from this project are providing a better understanding of laser backscattering. The results will feed into improved experimental setups aimed at advancing inertial fusion energy. Achieving ignition at the National Ignition Facility will resolve the question of whether fusion energy can help meet the global demand for safe, secure, and environmentally sustainable energy.





# PROJECT AWARDS

## Innovative & Novel Computational Impact on Theory and Experiment (INCITE)

### Biological Sciences

#### Computational Studies of Nucleosome Stability

*George Schatz, Northwestern University*  
Allocation: 20 Million Core-Hours

#### Multiscale Blood Flow Simulations

*George Karniadakis, Brown University*  
Allocation: 71 Million Core-Hours  
ALCF: 20 Million; OLCF: 51 Million

#### Studies of Large Conformational Changes in Biomolecular Machines

*Benoit Roux, The University of Chicago*  
Allocation: 55 Million Core-Hours

#### Towards Breakthroughs in Protein Structure Calculation and Design

*David Baker, University of Washington*  
Allocation: 140 Million Core-Hours

### Chemistry

#### From LES to DNS of Explosions in Semi-Confined Domains

*Thierry Poinsot, CERFACS*  
Allocation: 20 Million Core-Hours

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

*Alexei Khokhlov, The University of Chicago*  
Allocation: 130 Million Core-Hours

### Computer Science

#### Developing and Testing Future Applications and Operating Systems for Exascale

*Maya Gokhale, Lawrence Livermore National Laboratory*  
Allocation: 50 Million Core-Hours

#### Performance Evaluation and Analysis Consortium (PEAC) End Station

*Leonid Oliker, Lawrence Livermore National Laboratory*  
Allocation: 85 Million Core-Hours  
ALCF: 40 Million; OLCF: 45 Million

#### Scalable System Software for Parallel Programming

*Robert Latham, Argonne National Laboratory*  
Allocation: 20 Million Core-Hours

### Earth Science

#### Attributing Changes in the Risk of Extreme Weather and Climate

*Michael Wehner, Lawrence Berkeley National Laboratory*  
Allocation: 150 Million Core-Hours

#### Climate-Science Computational Development Team: The Climate End Station II

*Warren Washington, University Corporation for Atmospheric Research*  
Allocation: 215 Million Core-Hours  
ALCF: 155 Million; OLCF: 60 Million

#### CyberShake 3.0: Physics-Based Probabilistic Seismic Hazard Analysis

*Thomas Jordan, University of Southern California*  
Allocation: 68 Million Core-Hours  
ALCF: 23 Million; OLCF: 45 Million

#### Evaluation of Mesoscale Atmospheric Model for Contrail Cirrus Simulations

*Roberto Paoli, CERFACS*  
Allocation: 25 Million Core-Hours

### Energy Technologies

#### Optimization of Complex Energy System Under Uncertainty

*Mihai Anitescu, Argonne National Laboratory*  
Allocation: 14 Million Core-Hours

#### Thermal Hydraulic Modeling: Cross-Verification, Validation and Co-Design

*Paul F. Fischer, Argonne National Laboratory*  
Allocation: 50 Million-Core Hours

#### Thermodynamics of Binding Biomass to Cellulases for Renewable Fuel

*Michael Crowley, National Renewable Energy Laboratory*  
Allocation: 70 Million Core-Hours

### Engineering

#### Combustion Stability in Complex Engineering Flows

*Lee Shunn, Cascade Technologies, Inc.*  
Allocation: 35 Million Core-Hours

#### Direct Numerical Simulations of High Reynolds Number Turbulent Channel Flow

*Robert Moser, The University of Texas at Austin*  
Allocation: 175 Million Core-Hours

### Enabling Green Energy and Propulsion Systems via Direct Noise Computation

*Umesh Paliath, GE Global Research  
Allocation: 105 Million Core-Hours*

### Materials Science

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

*William George, National Institute of Standards and Technology  
Allocation: 20 Million Core-Hours*

#### Multibillion-Atom MD Studies of the Mechanical Response of Nanocrystalline Ta

*Timothy Germann, Los Alamos National Laboratory  
Allocation: 20 Million Core-Hours*

#### Multiscale Modeling of Energy Storage Materials

*Gregory A. Voth, The University of Chicago  
Allocation: 25 Million Core-Hours*

#### Non-Covalent Bonding in Complex Molecular Systems with Quantum Monte Carlo

*Dario Alfé, University College London  
Allocation: 95 Million Core-Hours  
ALCF: 40 Million; OLCF: 55 Million*

#### Petascale Simulations of Stress Corrosion Cracking

*Priya Vashishta, University of Southern California  
Allocation: 240 Million Core-Hours*

#### Predictive Materials Modeling for Li-Air Battery Systems

*Larry Curtiss, Argonne National Laboratory  
Allocation: 100 Million Core-Hours*

#### Vibrational Spectroscopy of Liquid Mixtures and Solid/Liquid Interfaces

*Giulia Galli, University of California, Davis  
Allocation: 100 Million Core-Hours*

### Physics

#### Accelerator Simulations for the Intensity Frontier of Particle Physics

*James Amundson, Fermilab  
Allocation: 80 Million Core-Hours*

#### Computing the Dark Universe

*Salman Habib, Argonne National Laboratory  
Allocation: 40 Million Core-Hours*

#### Kinetic Simulations of Fusion Energy Dynamics at the Extreme Scale

*William Tang, Princeton Plasma Physics Laboratory  
Allocation: 40 Million Core-Hours*

#### Lattice QCD

*Paul Mackenzie, Fermilab  
Allocation: 430 Million Core-Hours  
ALCF: 290 Million; OLCF: 140 Million*

#### Next-Generation Petascale Simulations of Type Ia Supernovae

*Don Lamb, The University of Chicago  
Allocation: 105 Million Core-Hours*

#### Nuclear Structure and Nuclear Reactions

*James Vary, Iowa State University  
Allocation: 155 Million Core-Hours  
ALCF: 81 Million; OLCF: 74 Million*

#### Petascale Simulations of Inhomogeneous Alfvén Turbulence in the Solar Wind

*Jean C. Perez, University of New Hampshire  
Allocation: 53 Million Core-Hours*

#### Simulation of Laser-Plasma Interaction in National Ignition Facility Experiments

*Steven Langer, Lawrence Livermore National Laboratory  
Allocation: 200 Million Core-Hours*

#### Spectral Slope of MHD Turbulence

*Andrey Beresnyak, Los Alamos National Laboratory  
Allocation: 35 Million Core-Hours*

#### Three-Dimensional Simulations for Core Collapse Supernovae

*Anthony Mezzacappa, Oak Ridge National Laboratory  
Allocation: 65 Million Core-Hours  
ALCF: 30 Million; OLCF: 35 Million*

#### Transformative Simulation of Shock-Generated Magnetic Fields

*Milad Fatenejad, The University of Chicago  
Allocation: 40 Million Core-Hours*

## ASCR Leadership Computing Challenge (ALCC)

### Computer Science

HPC Colony: Adaptive System Software for Improved Resiliency and Performance

*Terry Jones, Oak Ridge National Laboratory*

*Allocation: 3 Million Core-Hours*

### Energy Technologies

Validation Work for Heterogeneous Nuclear Reactor Calculations

*Micheal Smith, Argonne National Laboratory*

*Allocation: 30 Million Core-Hours*

### Engineering

Petascale Thermal-Hydraulic Simulations in Support of CESAR

*Elia Merzari, Argonne National Laboratory*

*Allocation: 30 Million Core-Hours*

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

*Parviz Moin, Stanford University*

*Allocation: 60 Million Core-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 Core-Hours*

### Materials Science

Dynamics of Conformational Transition in Thermo-Sensitive Polymers and Hydrogels

*Subramanian Sankaranarayanan, Argonne National Laboratory*

*Allocation: 50 Million Core-Hours*

Liquid-Solid Interfaces in Electrocatalysis from First Principles

*Jeffrey Greeley, Argonne National Laboratory*

*Allocation: 20 Million Core-Hours*

### Physics

Ab Initio Hyper-Nuclear Physics

*Kostas Orginos, College of William & Mary*

*Allocation: 20 Million Core-Hours*

ALCC: Exploring the Nature of the Lightest Massive Particles in the Universe

*Katrin Heitmann, Argonne National Laboratory*

*Allocation: 4 Million Core-Hours*



## Early Science Program (ESP)

### Biological Science

Multiscale Molecular Simulations at the Petascale

Gregory Voth, The University of Chicago

*Allocation: 150 Million Core-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 Core-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 Core-Hours*

High-Accuracy Predictions of the Bulk Properties of Water

Mark Gordon, Iowa State University

*Allocation: 150 Million Core-Hours*

High-Speed Combustion and Detonation (HSCD)

Alexei Khokhlov, The University of Chicago

*Allocation: 150 Million Core-Hours*

### Earth Science

Climate-Weather Modeling Studies Using a Prototype Global Cloud-System Resolving Model

Venkatramani Balaji, Geophysical Fluid Dynamics Laboratory

*Allocation: 150 Million Core-Hours*

Using Multiscale Dynamic Rupture Models to Improve Ground Motion Estimates

Thomas Jordan, University of Southern California

*Allocation: 150 Million Core-Hours*

### Energy Technologies

Materials Design and Discovery: Catalysis and Energy Storage

Larry Curtiss, Argonne National Laboratory

*Allocation: 50 Million Core-Hours*

Petascale Direct Numerical Simulations of Turbulent Channel Flow

Robert Moser, The University of Texas at Austin

*Allocation: 60 Million Core-Hours*

### Engineering

Direct Numerical Simulation of Autoignition in a Jet in a Cross-Flow

Christos Frouzakis, Swiss Fed. Inst. Tech.

*Allocation: 150 Million Core-Hours*

Petascale, Adaptive CFD

Kenneth Jansen, University of Colorado-Boulder

*Allocation: 150 Million Core-Hours*

### Physics

Ab initio Reaction Calculations for Carbon-12

Steven C Pieper, Argonne National Laboratory

*Allocation: 110 Million Core-Hours*

Cosmic Structure Probes of the Dark Universe

Salman Habib, Los Alamos National Laboratory

*Allocation: 150 Million Core-Hours*

Global Simulation of Plasma Microturbulence at the Petascale and Beyond

William Tang, Princeton Plasma Physics Laboratory

*Allocation: 50 Million Core-Hours*

LatticeQCD - Early Science

Paul Mackenzie, Fermilab

*Allocation: 150 Million Core-Hours*

Petascale Simulations of Turbulent Nuclear Combustion

Don Lamb, The University of Chicago

*Allocation: 150 Million Core-Hours*

## Director's Discretionary (DD)

### Biological Sciences

Coarse-Grained Simulations of Sec-Facilitated Protein Translocation and Membrane Integration

*Thomas F. Miller, III, California Institute of Technology*  
Allocation: 6 Million Core-Hours

Molecular Dynamics Simulations for Exploring Androgen Receptor Antagonism, Drug-Resistant Mutations, and Antagonist Design

*Yang Shen, Toyota Technological Institute at Chicago*  
Allocation: 3 Million Core-Hours

Nanoscale Biomolecular Simulation for Understanding of Peptide Recognition by Epigenetic Protein.

*Nadeem Vellore, The University of Utah*  
Allocation: 300,000 Core-Hours

Simulation of Cell Coupling in Arterial Bifurcation

*Timoty David, University of Canterbury*  
Allocation: 1.1 Million Core-Hours

Simulation of Large-Scale Biomolecular Systems

*Rommie E. Amaro, University of California-San Diego*  
Allocation: 20 Million Core-Hours

Towards a Model of the Replisome

*Aleksii Aksimentiev, University of Illinois at Urbana-Champaign*  
Allocation: 4.2 Million Core-Hours

### Chemistry

Investigation of Catalytic Properties of Nanoclusters

*Leonardo Spanu, Shell International E&P, Inc.*  
Allocation: 2 Million Core-Hours

Minimum Energy Path with Quantum Monte Carlo

*Leonardo Guidoni, University of L'Aquila*  
Allocation: 1 Million Core-Hours

Probing the Free Energy Surface of Spin Separation in Singlet Fission

*Hanning Chen, The George Washington University*  
Allocation: 6 Million Core-Hours

Quantum Monte Carlo Simulations of Biochemical and Catalysis-Related Systems

*William Parker, Argonne National Laboratory*  
Allocation: 10 Million Core-Hours

Scaling and Baseline Study of Coupled Simulations using A Very Big Program (AVBP) on O(B) Finite Volumes on the Blue Gene/Q

*Gabriel Staffelbach, Argonne National Laboratory/CERFACS*  
Allocation: 5 Million Core-Hours

### Computer Science

Charm++ and its Applications

*Laxmikant V. Kale, University of Illinois at Urbana-Champaign*  
Allocation: 2.5 Million Core-Hours

Distributed File Systems for Exascale Computing

*Ioan Raicu, Illinois Institute of Technology*  
Allocation: 500,000 Core-Hours

### Earth Science

Climate Sensitivity Experiments Using a Multiscale Modeling Framework with a Higher-Order Turbulence Closure in its CRM

*Anning Cheng, Science Systems and Applications, Inc./ National Aeronautics and Space Administration*  
Allocation: 3 Million Core-Hours

Dynamic Downscaling of Climate Models

*V. Rao Kotamarthi, Argonne National Laboratory*  
Allocation: 37 Million Core-Hours

Sensitivity and Uncertainty of Precipitation of the GFDL High Resolution Model

*Laura Zamboni, Argonne National Laboratory*  
Allocation: 3.5 Million Core-Hours

### Engineering

Aspect Ratio Effects in Turbulent Duct Flows Studied Through Direct Numerical Simulation

*Hassan Nagib, Illinois Institute of Technology/Argonne National Laboratory/KTH Royal Institute of Technology*  
Allocation: 1.5 Million Core-Hours

Atomistic Simulation of Laser Processing of Metal Multilayers

*Leonid V. Zhigilei, University of Virginia*  
Allocation: 1 Million Core-Hours

DNS Study of a Spatially Developing Compressible Mixing Layer with Non-Unity Density and Temperature Ratios

*Francesco Grasso, Conservatoire National des Arts et Metiers*  
Allocation: 2.5 Million Core-Hours

**DNS of Wall-Bounded Turbulence**

*Gary N Coleman, National Aeronautics and Space Administration*  
Allocation: 1 Million Core-Hours

**High-Fidelity, High-Energy-Density Hydrodynamics Simulations of Shocks Interacting with Material Discontinuities**

*Eric Johnsen, University of Michigan*  
Allocation: 250,000 Core-Hours

**Highly Resolved LES of a GE 3-Cup Combustor System**

*Anne L. Dord, General Electric Company*  
Allocation: 10 Million Core-Hours

**LES of Turbulent Jet Noise**

*Marlene Sanjose, Universite de Sherbrooke*  
Allocation: 1.5 Million Core-Hours

**Predictive Modeling for Complex Phenomena in Electromagnetics and Fluid Systems**

*Misun Min, Argonne National Laboratory*  
Allocation: 2 Million Core-Hours

**Porting Uintah to the Blue Gene/Q Architecture**

*Martin Berzins and John Schmidt, The University of Utah*  
Allocation: 10 Million Core-Hours

**Quantum Lattice Algorithm for Quantum Turbulence**

*George Vahala, William & Mary*  
Allocation: 2.9 Million Core-Hours

**Stochastic (w\*) Convergence for Turbulent Combustion**

*James Glimm, Stony Brook University*  
Allocation: 15.75 Million Core-Hours

**Mathematics****Large and Accurate Numerical Solutions of Partial Differential Equations**

*Benson Muite, King Abdullah University of Science and Technology*  
Allocation: 3 Million Core-Hours

**Materials Science****Chemo-Mechanical Properties of Complex Oxide Interfaces**

*Alessandro De Vita and James Kermode, King's College London/Argonne National Laboratory*  
Allocation: 2 Million Core-Hours

**LAMMPS Performance Optimization**

*Paul Coffman, IBM*  
Allocation: 2 Million Core-Hours

**Liquid Crystal Based Functional Materials**

*Juan J. de Pablo, The University of Chicago, Argonne National Laboratory*  
Allocation: 4 Million Core-Hours

**Mesoscale Modeling of Self-Assembly of Polymer Grafted Nanoparticles**

*Derrick Mancini, Argonne National Laboratory*  
Allocation: 2 Million Core-Hours

**Nanostructure-Enhanced Detonation in Energetic Materials**

*Tzu-Ray Shan, Sandia National Laboratories*  
Allocation: 600,000 Core-Hours

**Phonon Parallelization in Quantum ESPRESSO**

*William Parker, Argonne National Laboratory*  
Allocation: 500,000 Core-Hours

**Radiation Tolerance and Thermo-Mechanical Properties of Amorphous SiCO Glasses and SiCO/Fe Composites**

*Kan-Ju Lin, Massachusetts Institute of Technology*  
Allocation: 1.25 Million Core-Hours

**Toward Crystal Engineering From First Principle**

*James R. Chelikowsky, University of Texas at Austin*  
Allocation: 1 Million Core-Hours

**Physics****Petascale Simulation of Magnetorotational Core-Collapse Supernovae**

*Sean M. Couch, The University of Chicago/Fermilab*  
Allocation: 30 Million Core-Hours

**Computational Studies of the Topological Properties of Micellar Solutions**

*Subas Dhakal, Syracuse University*  
Allocation: 3 Million Core-Hours

**Diffraction-Limited Light Source for APS Upgrade**

*Michael Borland, Argonne National Laboratory*  
Allocation: 10 Million Core-Hours

**DNS Simulations of Turbulent Convection**

*Janet Scheel, Occidental College*  
Allocation: 1 Million Core-Hours

**First Principles Quantum Monte Carlo for Superconducting Materials**

*Lucas K. Wagner, University of Illinois at Urbana-Champaign  
Allocation: 500,000 Core-Hours*

**Global Simulations of Accretion Discs**

*Fausto Cattaneo, The University of Chicago  
Allocation: 7 Million Core-Hours*

**Grid-Enabling High-Performance Computing for ATLAS**

*Thomas J. LeCompte, Argonne National Laboratory  
Allocation: 5 Million Core-Hours*

**High-Fidelity Simulation of Small Modular Reactors with Monte Carlo**

*Benoit Forget, Massachusetts Institute of Technology  
Allocation: 10 Million Core-Hours*

**Neutronics Simulation of FFTF and EBR-II**

*Micheal A Smith, Argonne National Laboratory  
Allocation: 5 Million Core-Hours*

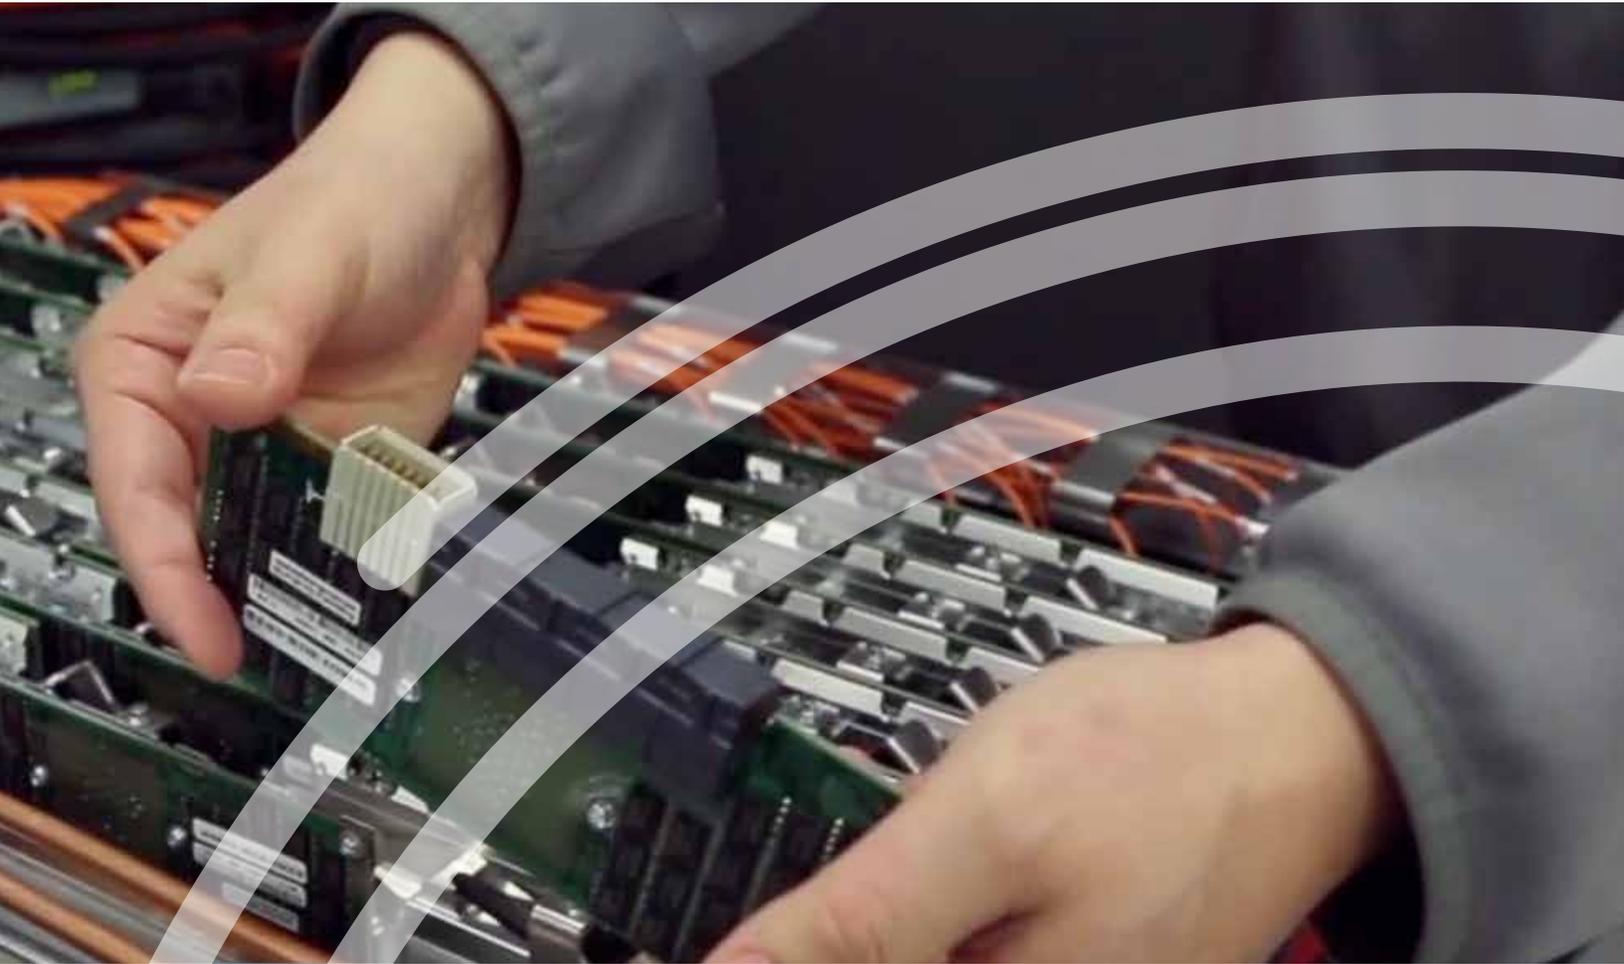
**Nuclear Structure and Nuclear Reactions**

*James Vary, Iowa State University  
Allocation: 1.5 Million Core-Hours*

**Shearing Box Simulations of Accretion Discs**

*Eric Blackman, University of Rochester  
Allocation: 300,000 Core-Hours*





A look inside Mira, the ALCF's IBM Blue Gene/Q supercomputer.

## 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](http://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](mailto:reports@adonis.osti.gov)

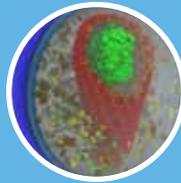
### DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor UChicago Argonne, LLC, nor any of their employees or officers, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of document authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof, Argonne National Laboratory, or UChicago Argonne, LLC.

This simulation image shows suspended particles in a rheometer for NIST's proposed mortar SRM. The spheres, which are color coded by their starting location in the rheometer, are suspended in a cement paste with properties derived from NIST's cement paste SRM.

**Image Credit**

Nicos Martys and Steven G. Satterfield, National Institute of Standards and Technology



This snapshot of a numerical simulation of the MASLWR experiment shows a cross-section with view of the velocity magnitude.

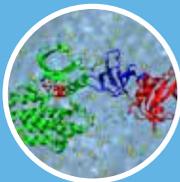
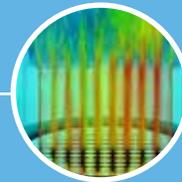
**Image Credit**

Elia Merzari, Argonne National Laboratory.

Coarse-grained representation of encapsulated HIV virion (red) with RNA (green) used in cellular-scale molecular simulations.

**Image Credit**

John Grime, Argonne National Laboratory/  
The University of Chicago



The Src-family of kinases are important enzymes for cellular signaling that can transfer a phosphoryl group from ATP to a target protein. The catalytic domain (green) is tightly regulated by the SH2 (blue) and SH3 (red) domains. The active state is shown here with a large water box and a concentration of ion-mimicking biological conditions.

**Image Credit**

Avishek Das, Mikolaj Fajer, and Benoît Roux, The University of Chicago

Particle visualization of a global gyrokinetic particle-in-cell simulation of microturbulence in a tokamak fusion device.

**Image Credit**

Chad Jones and Kwan-Liu Ma, University of California, Davis; Stephane Ethier, Princeton Plasma Physics Laboratory



Large eddy simulation of the flow in a rod bundle at high Reynolds numbers. Rod bundles are a fundamental component of reactor cores. These very large and unprecedented calculations (over 8 billion collocation points) will shed light on the performance of current algorithms, help estimate scaling to larger machines, and gain unprecedented insight into the physics of turbulence.

**Image Credit**

Chad Jones and Kwan-Liu Ma, University of California, Davis; Stephane Ethier, Princeton Plasma Physics Laboratory

[www.alcf.anl.gov](http://www.alcf.anl.gov)



## CONTACT

[media@alcf.anl.gov](mailto: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-SciBrochure-11/13