ARGONNE LEADERSHIP COMPUTING FACILITY

2015 SCIENCE HIGHLIGHTS

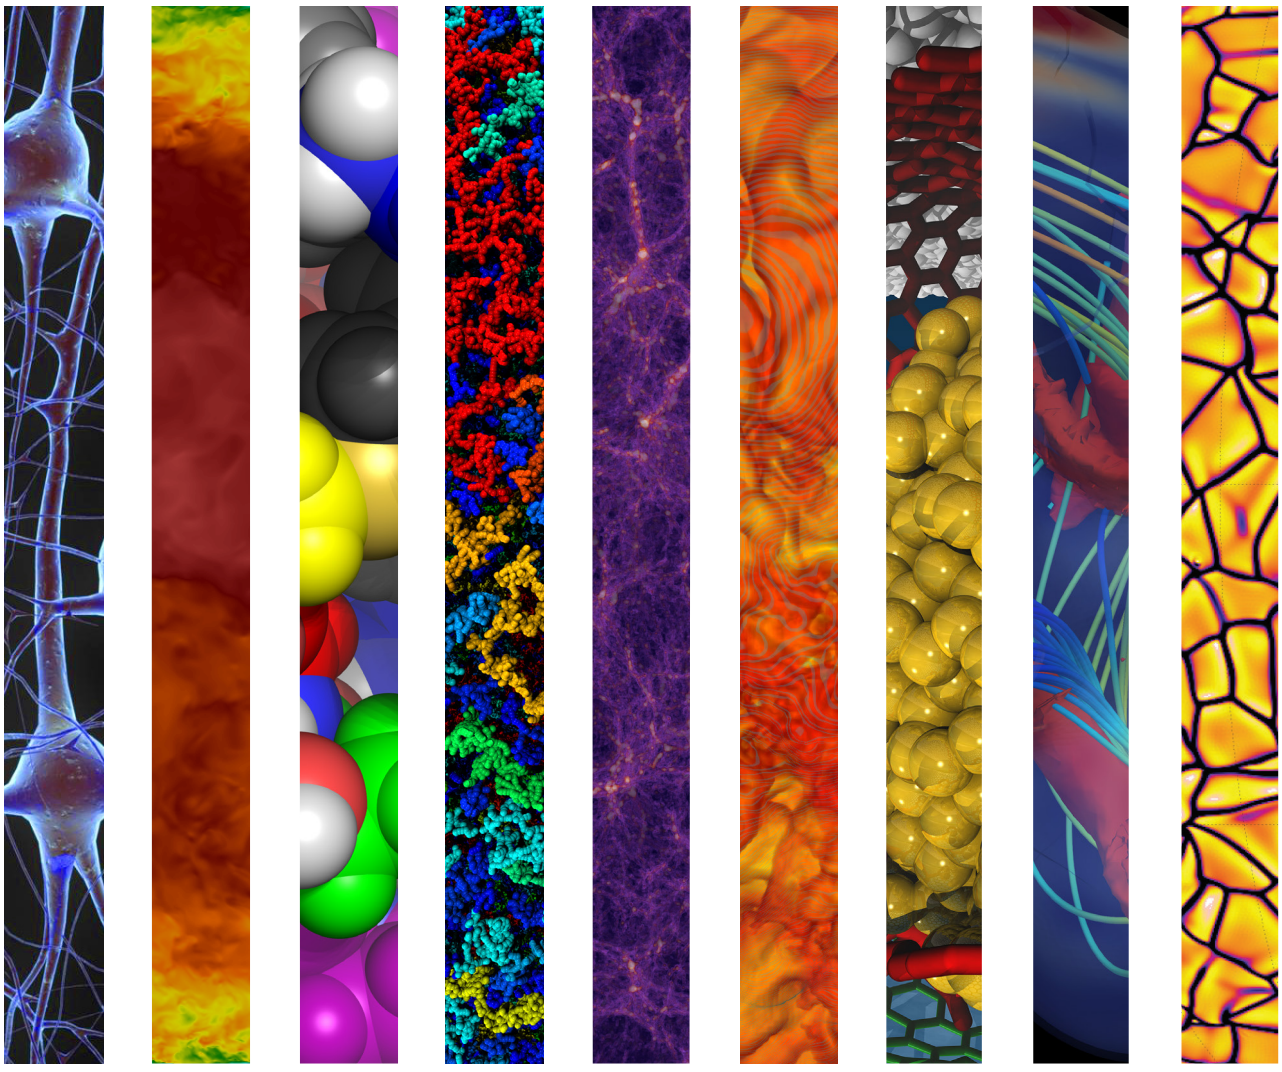

Argonne NATIONAL LABORATORY

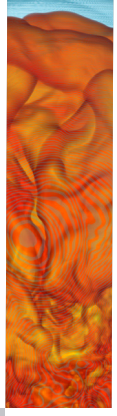
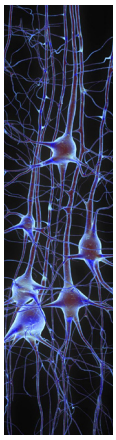
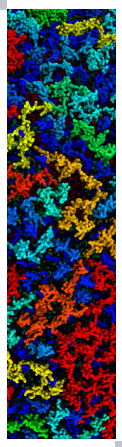
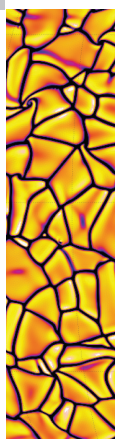
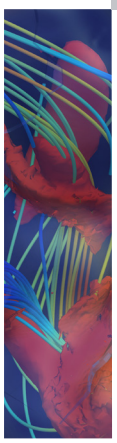
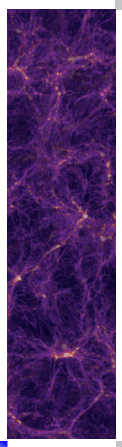
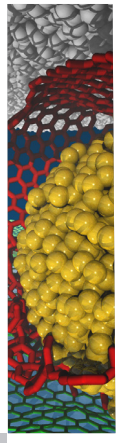
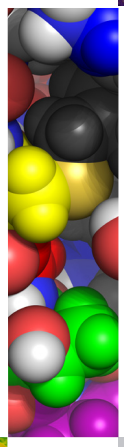
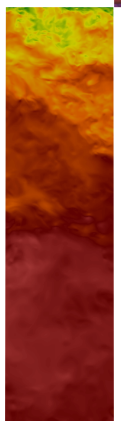

TABLE OF CONTENTS

Materials Science

Physics

ALCF RESOURCES

2015 PROJECTS

About ALCF

The Argonne Leadership Computing Facility (ALCF) provides supercomputing capabilities to the scientific and engineering community to advance fundamental discovery and understanding in a broad range of disciplines.

Supported by the U.S. Department of Energy's Office of Science, Advanced Scientific Computing Research (ASCR) program, the ALCF is one of two sites within the DOE Leadership Computing Facility (the other is located at Oak Ridge National Laboratory).

Available to researchers from universities, industry, and government agencies, the ALCF is a national scientific user facility that helps accelerate the pace of discovery and innovation by providing supercomputing resources that are 10 to 100 times more powerful than systems typically available for open scientific research.

Through substantial awards of supercomputing time and user support services, the ALCF enables large-scale modeling and simulation research aimed at solving some of the world's largest and most complex problems in science and engineering.

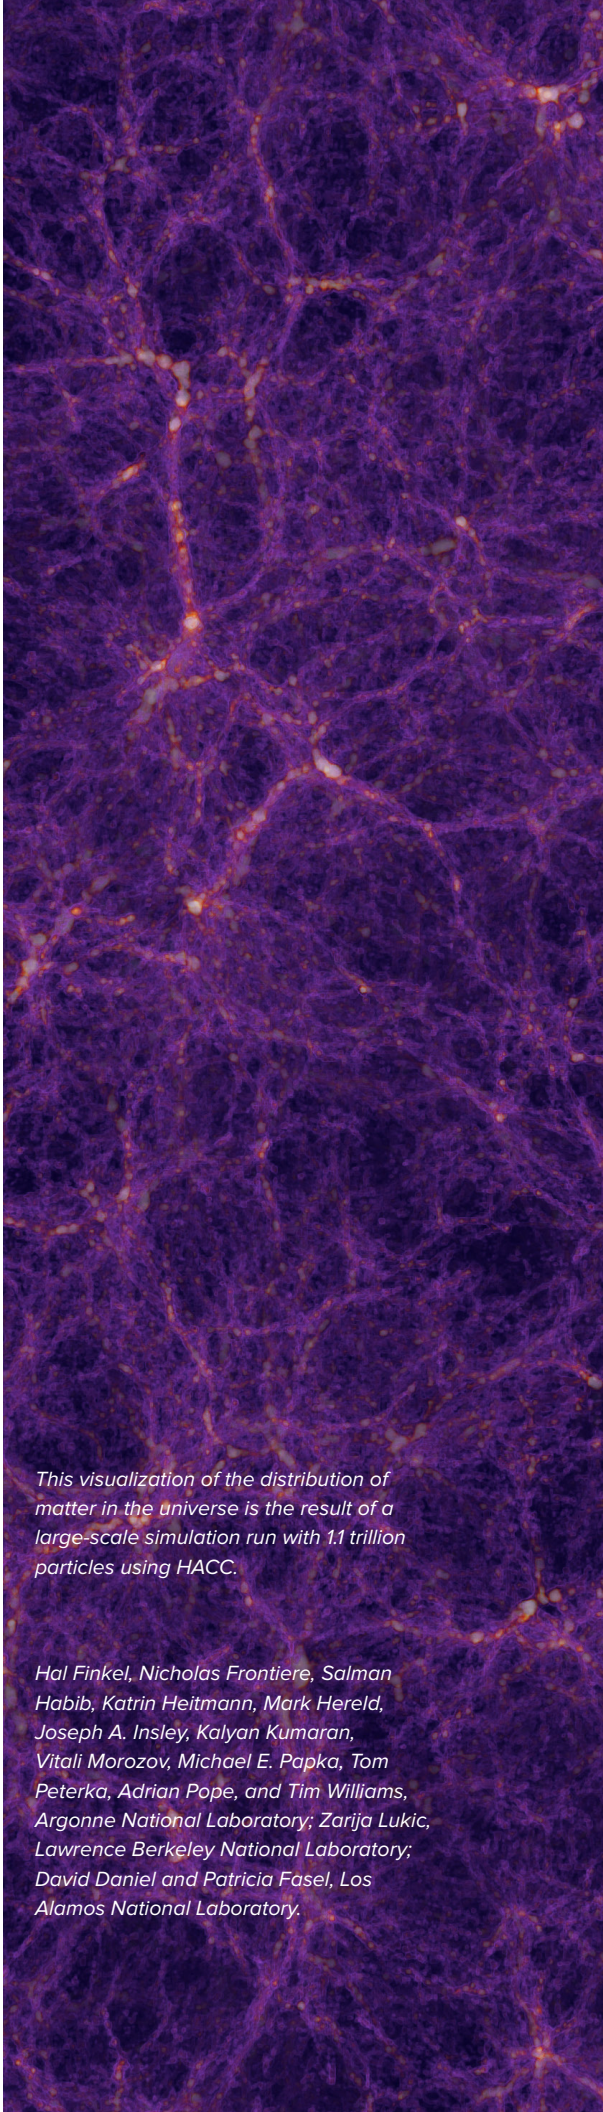

This visualization of the distribution of matter in the universe is the result of a large-scale simulation run with 1.1 trillion particles using HACC.

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

Paul Messina

ALCF Director of Science,
Argonne Distinguished Fellow

Science Director's Message

It has been a great year for science at the ALCF. Our unique combination of resources—leading-edge computing systems and a world-class staff—have enabled dozens of INCITE and ALCC projects to make great strides with their ambitious research goals.

For example, a research team from Argonne used Mira to identify and improve a new mechanism to eliminate friction, which led to the development of a novel material that displayed superlubricity at the macroscale for the first time. Another team from the University of Minnesota used Mira to run a predictive modeling tool that successfully identified new materials with the potential to improve the production of biofuel and petroleum products. As a final example, Argonne researchers collaborated with industry to use ALCF resources to enhance engine modeling capabilities for supercomputers, providing a tool that can help accelerate the development of more energy-efficient engines. You'll read about these projects and several others in the following pages.

We also made a big splash in April with the announcement of our next-generation supercomputers: Aurora and Theta. We're excited about the possibilities that these innovative systems will bring to our users. With more computational power and advanced capabilities, researchers will be able to carry out more realistic simulations of everything from wind turbines and climate systems to brain activity and battery materials.

To prepare for our next supercomputers and the future exascale systems to follow, we've also been active in efforts to ensure that software evolves along with the advances in hardware. In the near term, our Theta Early Science Program (ESP) kicked off this summer to prepare key applications for the new system by providing project teams with training, resources, and pre-production allocations. Led by Tim Williams, ALCF principal project specialist, this program allows researchers to use a leading-edge system for their scientific pursuits, while helping us to identify potential bugs and characterize the behavior of the new hardware and software features. We will also hold an ESP for Aurora with a call for proposals planned for next year.

On a broader scale, we participated in the DOE High-Performance Computing Operational Review on Scientific Software Architecture for Portability and Performance in September. Chaired by Katherine Riley, ALCF principal scientific applications engineer, the three-day workshop brought together application developers, computing facilities, vendors, and library and tool developers to identify approaches and best practices for increasing application portability and performance on diverse architectures in the coming years.

We are also collaborating with our fellow DOE ASCR facilities—OLCF and NERSC—to conduct a series of Exascale Requirements Reviews to determine the requirements necessary for developing an exascale ecosystem that will support forefront scientific research goals for each of the six DOE Office of Science program offices. Forward-looking activities like these are helping to ensure we are on a path to developing future leadership computing systems that have the most appropriate hardware and software for enabling breakthrough discoveries in science and engineering.

Introducing Aurora

In April, DOE announced a $200 million investment to deliver Aurora, the ALCF's next-generation supercomputer.

Designed in collaboration with industry leaders Intel and Cray, Aurora is scheduled for delivery in 2018. The effort is the result of DOE's Collaboration of Oak Ridge, Argonne, and Lawrence Livermore (CORAL) initiative, a joint procurement activity launched in 2014.

Aurora will deliver more than 18 times the computational performance of Mira, its predecessor at the ALCF, using a nearly equal number of compute nodes. Aurora will be a many-core system, but with nearly an order of magnitude more processors.

Aurora's revolutionary architecture features Intel's HPC scalable system framework and second-generation Intel® Omni-Path Fabric, providing a peak performance of 180 petaflops. The system will have a combined total of over 7 petabytes of on-package high-bandwidth memory and persistent memory, connected and communicating via a high-performance system fabric to achieve landmark throughput. The nodes will be linked to a dedicated burst buffer and a high-performance parallel storage solution.

The system will help ensure continued U.S. leadership in high-end computing for scientific research, while also cementing the nation's position as a global leader in the development of next-generation exascale computing systems.

Argonne will also provide an interim Intel system, named Theta, which will be delivered in 2016 to serve as a bridge between Mira and Aurora. Theta will be based on Intel's second-generation Xeon Phi processor and will serve as an early production system for the ALCF, helping users transition their applications to the new technology.

Made for Science

From developing alternative energy sources to treating diseases to advancing our understanding of the physical world, the research done on leadership-class machines yields significant scientific advancements. Key research goals for the next-generation ALCF systems include:

- Materials Science: Designing new classes of materials that will lead to more powerful, efficient, and durable batteries and solar panels.
- Biological Science: Gaining the ability to understand the capabilities and vulnerabilities of organisms that can result in improved biofuels and more effective disease control.
- Transportation Efficiency: Collaborating with industry to improve transportation systems with enhanced aerodynamics features, and to develop better, more efficient, and quieter engines.
- Renewable Energy: Engineering wind turbine design and placement to greatly improve efficiency and reduce noise.

How They Compare	Mira	Theta	Aurora
Peak Performance	10 PF	>8.5 PF	180 PF
Compute Nodes	49,152	>2,500	>50,000
Processor	PowerPC A2 1600 MHz	2nd Generation Intel Xeon Phi	3rd Generation Intel Xeon Phi
System Memory	768 TB	>480 TB	>7 PB
File System Capacity	26 PB	10 PB	>150 PB
File System Throughput	300 GB/s	200 GB/s	>1 TB/s
Intel Architecture (x86-64) Compatibility	No	Yes	Yes
Peak Power Consumption	4.8 MW	1.7 MW	>13 MW
GFLOPS/watt	2.1	>5	>13

"The Leadership Computing Facility has been driving the evolution of supercomputing for over a decade. Its most important function is to align leadership systems with the needs and goals of breakthrough science projects. Aurora will deliver the system capabilities that our users require to expand their investigations in both scale and scope."

- ALCF Director Michael E. Papka

Theta — A Bridge to Aurora

Theta, an early production system based on Intel's second-generation Xeon Phi processor, will arrive in 2016 and serve as a bridge between the ALCF's current supercomputer, Mira, and its next leadership-class supercomputer, Aurora, which is scheduled for delivery in 2018.

Theta is expected to enable much higher performance due to its advanced memory hierarchy, improved single-thread performance, and features that are beneficial to data-intensive simulations.

Theta Early Science Program

This summer, the ALCF selected six projects for its Theta Early Science Program (ESP), a collaborative effort designed to help prepare scientific applications for the architecture and scale of the new supercomputer.

The Theta ESP brings together computational scientists, code developers, and computing hardware experts to optimize key applications for Theta, and to solidify libraries and infrastructure to pave the way for other applications to run on the system.

Modeled after the ALCF's highly successful ESP for Mira, the program also gives researchers substantial allocations of pre-production compute time on Theta to pursue innovative computational science calculations that push the boundaries of what's possible with leadership-class supercomputers.

Like the typical ALCF workload, the six selected ESP projects, known as Tier 1 projects, represent a wide spectrum of scientific areas and numerical methods.

To help develop and optimize their software for Theta, project teams collaborate with ALCF staff, as well as vendor staff through the ALCF Intel-Cray Center for Excellence (AICCE). Four of the six projects were also assigned a dedicated postdoctoral researcher.

In addition, the ALCF will host ESP training sessions, including virtual workshops on system hardware and programming, and a hands-on workshop for testing and debugging of project applications.

Prior to Theta's availability, the ALCF is offering access to small systems based on Theta's CPU, and provide allocations on Mira for development work that does not depend on having the new hardware.

Because of the strong response to the call for proposals, the ALCF expanded the Theta ESP to include six additional Tier 2 projects to help prepare other applications for Theta. These projects did not receive allocations for science runs, but they do have access to ESP training, an ESP discussion forum, early hardware, and to Theta itself for porting, tuning, and debugging.

For one of the Theta ESP projects, University of Chicago scientist Giulia Galli will lead an effort focused on high-performance calculations of nanoparticles and aqueous systems for energy applications.

Nicholas Brawand, University of Chicago

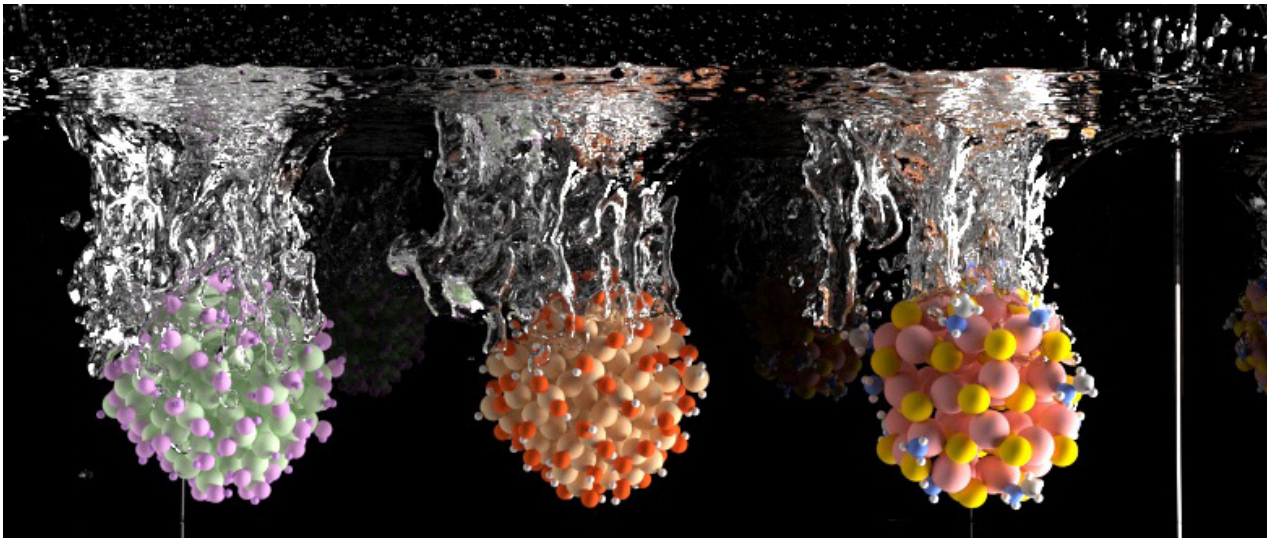

Porting for Performance

ALCF leads collaborative ESP project to develop software tools for Theta.

As high-performance computing architectures continue to get more complex, it is becoming increasingly difficult for scientific applications to fully utilize their capabilities. Programming models, performance tools, debuggers, and software libraries are critical to developing scientific applications for the upcoming generation of leadership-class supercomputers.

To help application developers, the ALCF performance engineering team is leading an omnibus ESP project to port and adapt such tools for Theta so they are available during the system's early deployment phase. This collaborative effort involves many of the major contributors in the HPC tools and system software communities.

Early Science Tools

Theta ESP Projects

Six Tier 1 projects were selected to optimize key scientific codes for Theta while pursuing ambitious science goals on the system. An additional six Tier 2 projects were selected for code development work. The Tier 2 projects did not receive allocations for science runs, but they do have access to ESP training and resources.

TIER 1 SCIENCE PROJECTS

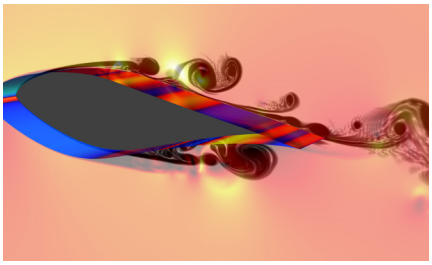

Scale-Resolving Simulations of Wind Turbines with SU2

Juan J. Alonso
Stanford University
Code: SU2

Researchers will develop a simulation capability to design better wind turbines and to lay out large wind farms for maximum energy extraction and improved turbine fatigue life. To do so, the research team will generate a database of large-eddy simulations of various single and multiple wind turbine settings.

Image: Ramesh Balakrishnan, Argonne National Laboratory

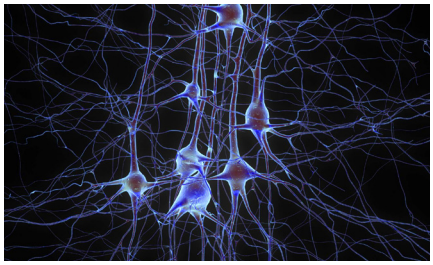

Large-Scale Simulation of Brain Tissue: Blue Brain Project, EPFL

Fabien Delalondre
Ecole Polytechnique Federale de Lausanne
Code: CoreNeuron

This project aims to improve our understanding of the brain using simulations of brain plasticity—experience-dependent changes in synaptic connectivity. Other work will include simulating rodent somatosensory cortex and the electrical activity of the largest possible brain model for several seconds of biological time.

Image: Blue Brain Project, Ecole Polytechnique Federale de Lausanne

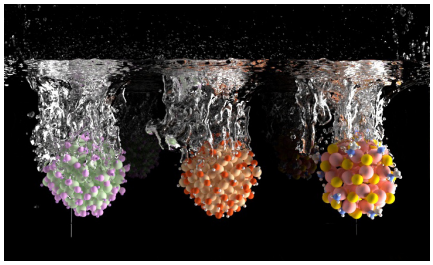

First-Principles Simulations of Functional Materials for Energy Conversion

Giulia Galli
University of Chicago
Codes: Qbox, WEST

Researchers will combine *ab initio* molecular dynamics and post-density functional theory methods to optimize properties of nanostructured materials for use in solar and thermal energy conversion devices at an unprecedented level of accuracy. The ultimate goal is to provide a truly predictive tool for device performance within a Materials Genome Initiative design framework.

Image: Nicholas Brawand, University of Chicago

TIER 2 CODE DEVELOPMENT PROJECTS

Electronic Structure-Based Discovery of Hybrid Photovoltaic Materials on Next-Generation HPC Platforms

Volker Blum
Duke University
Codes: FHI-aims, GAtor

Flow, Mixing, and Combustion of Transient Turbulent Gaseous Jets in Confined Cylindrical Geometries

Christos Frouzakis
Swiss Federal Institute of Technology Zurich
Code: Nek5000

Advanced Electronic Structure Methods for Heterogeneous Catalysis and Separation of Heavy Metals

Mark Gordon
Iowa State University
Code: GAMESS

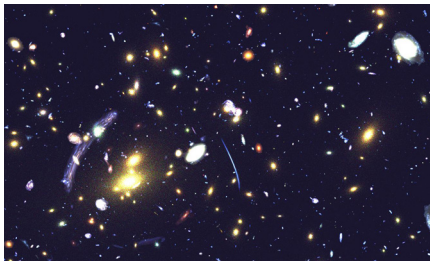

Next-Generation Cosmology Simulations with HACC: Challenges from Baryons

Katrin Heitmann
Argonne National Laboratory
Code: HACC

This project aims to further our understanding of astrophysical processes by performing detailed simulations of the universe for comparison with the latest observational data. The simulations will disentangle astrophysical processes (e.g., galaxy evolution) from fundamental physics (e.g., dark energy), helping mitigate one of the major sources of systematic uncertainties for upcoming cosmological surveys.

Image: Lindsey Bleem, Nan Li, and the HACC team, Argonne National Laboratory; Mike Gladders, University of Chicago

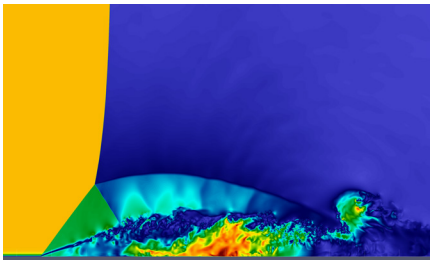

Direct Numerical Simulations of Flame Propagation in Hydrogen-Oxygen Mixtures in Closed Vessels

Alexei Khokhlov
University of Chicago
Code: HSCD

Researchers will perform direct numerical simulations of the flame acceleration and the deflagration-to-detonation transition process in hydrogen-oxygen mixtures in closed spherical vessels—exactly matching experimental apparatus. This research is aimed at improving the industrial and public safety of hydrogen fuels and certain water-cooled nuclear reactors.

Image: Charles Bacon, Marta García, and Joseph A. Insley, Argonne National Laboratory; Alexei Khokhlov, University of Chicago; Joanna Austin and Andrew Knisely, University of Illinois at Urbana-Champaign

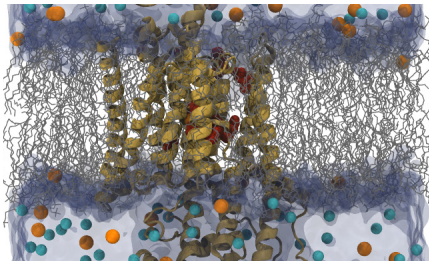

Free Energy Landscapes of Membrane Transport Proteins

Benoît Roux
University of Chicago
Code: NAMD

This project will carry out molecular dynamics simulations to provide detailed visualizations of the large conformational changes of membrane transport proteins and quantitative predictions of the energetics of these processes. This atomistic picture of membrane transport proteins stands to improve our understanding of a broad range of biological functions.

Image: Brian Radak and Huan Rui, University of Chicago

Extreme-Scale Unstructured Adaptive CFD: From Multiphase Flow to Aerodynamic Flow Control

Kenneth Jansen
University of Colorado Boulder
Code: PHASTA

The Hadronic Contribution to the Anomalous Magnetic Moment of the Muon

Paul Mackenzie
Fermilab
Codes: MILC, CPS

Quantum Monte Carlo Calculations in Nuclear Theory

Steven Pieper
Argonne National Laboratory
Code: GFMC

Allocation Programs

Any researcher in the world with a large-scale computing problem can apply for time on ALCF computing resources.

Project proposals are peer-reviewed and assessed for computational readiness. Accepted projects are allocated parcels of computing time that add up to billions of core-hours each year.

A breakdown of how computing time on Mira is allotted among the allocation programs.

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

The DOE's INCITE program provides allocations to computationally intensive, large-scale research projects that aim to address "grand challenges" in science and engineering. The program 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 year-long allocation cycle runs from July 1 to June 30.

Director's Discretionary

The ALCF's Director's Discretionary 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 core-hours.

Early Science Program (ESP)

As part of the process of bringing a new supercomputer into production, the ALCF hosts Early Science Programs to ensure its next-generation systems are ready to hit the ground running. The intent of the ESP is to use the critical pre-production time period to prepare key applications for the architecture and scale of a new supercomputer, and to solidify libraries and infrastructure to pave the way for other production applications to run on the system. In addition to fostering application readiness, the ESP allows researchers to pursue innovative computational science projects not possible on today's leadership-class supercomputers.

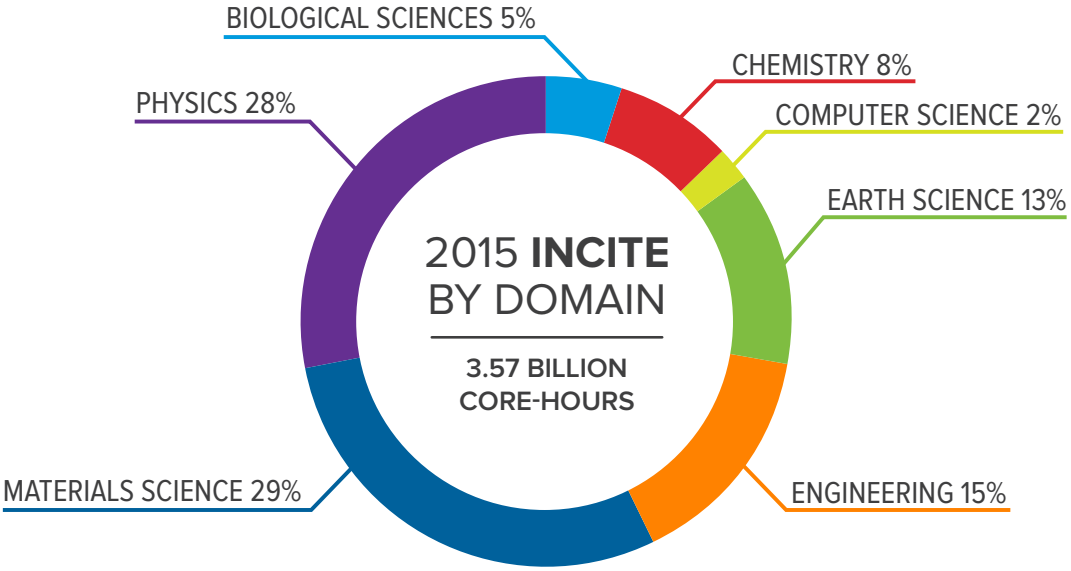

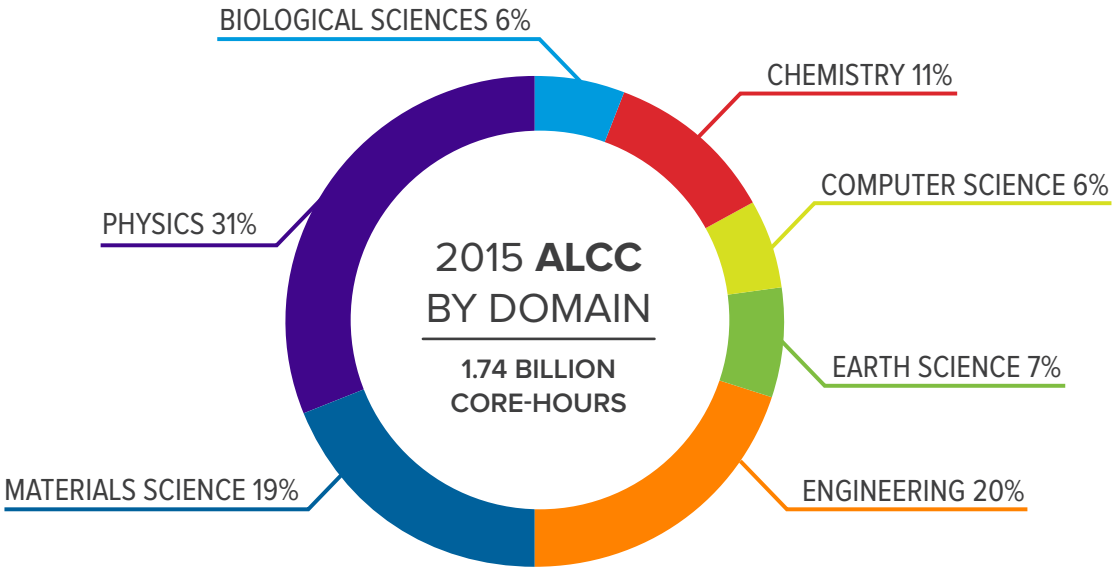

Note: Data is from calendar year 2015.

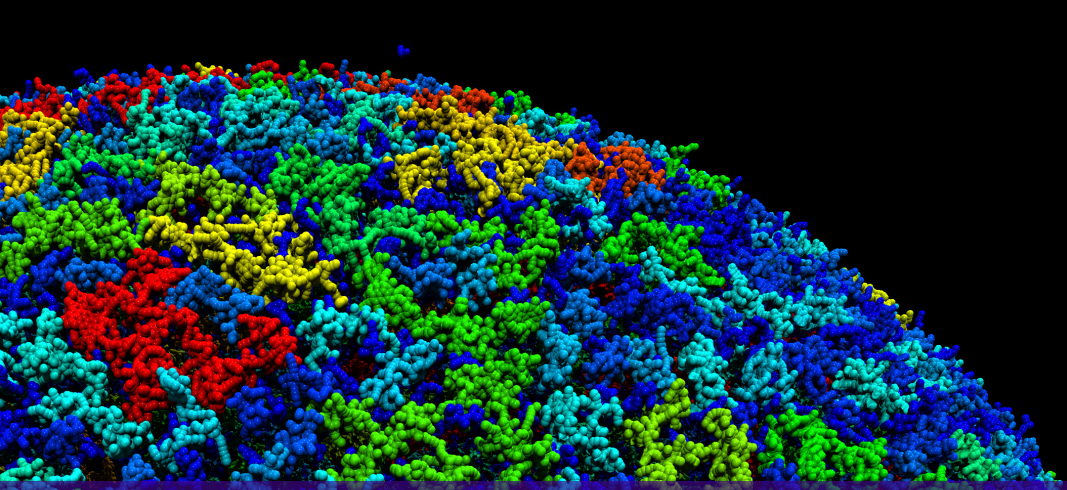

Continuing Discoveries

Past projects provide valuable insights

Supercomputers are powerful research tools that can greatly accelerate scientific discoveries, but analyzing the massive datasets they produce is a process that can take many months. In addition, large-scale computational science projects often require significant time to fully develop the necessary software to achieve research goals.

At the ALCF, users often seek allocation renewals to continue their work from year to year, allowing research projects to evolve along with the facility and our increasingly powerful computing resources. In the following pages, we highlight some past projects that have achieved notable results.

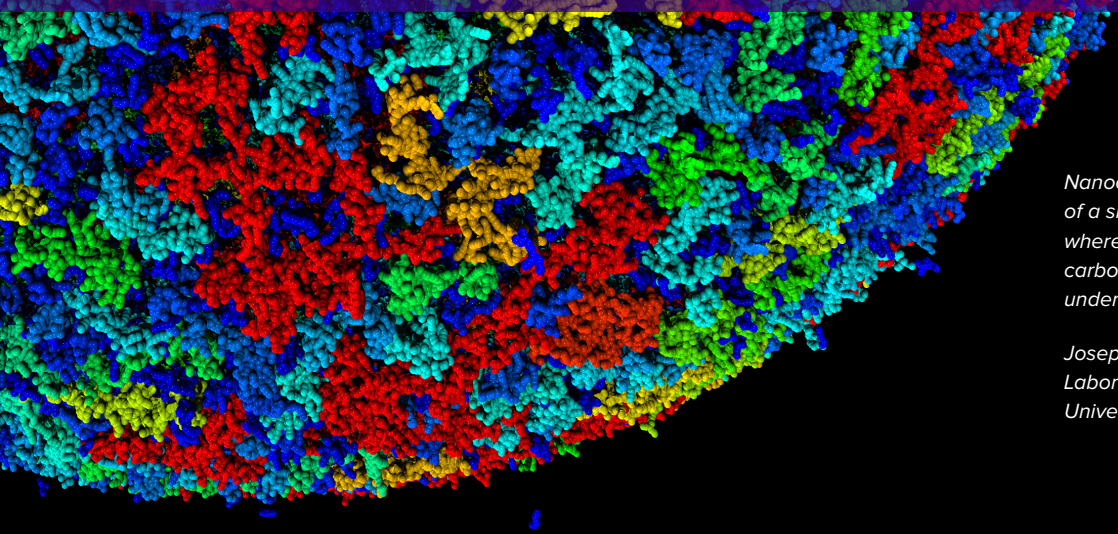

Nanocarbon produced by oxidation of a silicon-carbide nanoparticle, where color represents the size of carbon-clusters just before they undergo a percolation transition.

Joseph A. Insley, Argonne National Laboratory; Ken-ichi Nomura, University of Southern California

HANNING CHEN

George Washington University
chenhanning@gwu.edu

Director's Discretionary
6 Million Core-Hours

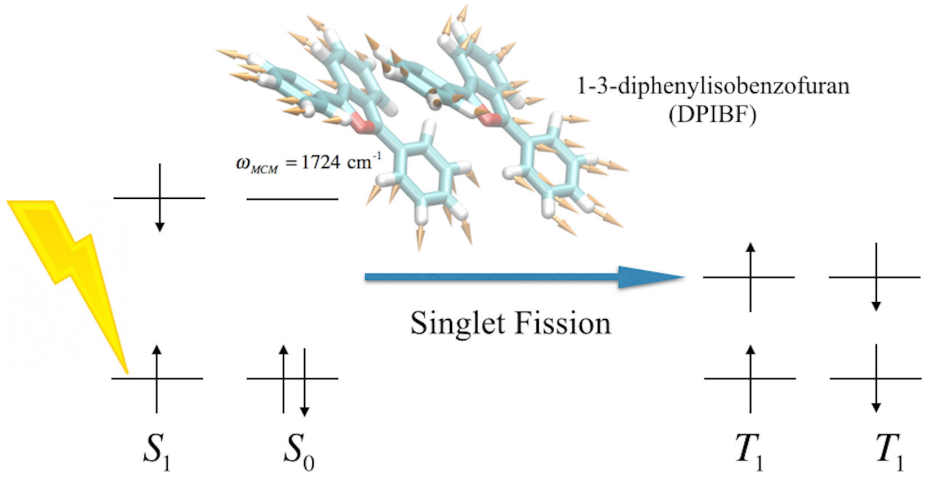

Image: Photo-induced spin polarization within a pair of 1-diphenylisobenzofuran (DPIBF) molecules.

Hanning Chen, George Washington University

CHEMISTRY

MOLECULAR MODELING OF SINGLET FISSION FOR SOLAR ENERGY CONVERSION

Singlet fission is a promising solution to long-standing roadblocks to more efficient solar energy conversion. The process converts two singlet molecules into two triplets upon the absorption of a single incident photon. This project investigated the molecular mechanism of this photo-induced spin separation process in a rich library of crystalline materials by a novel functional mode electron transfer theory.

Research efforts focused on large-scale molecular dynamics calculations to ascertain a full spectrum of vibrational normal modes of the 1-3-diphenylisobenzofuran (DPIBF) thin film before the spin separation process is examined by hybrid quantum mechanics/molecular mechanics simulations. ALCF was integral in helping researchers skillfully run the open-source molecular simulation package CP2K at petascale on Mira, rendering a detailed picture of thermally driven spin transfer in molecular crystals.

Petascale simulations provided free energy profiles along the reaction coordinate of singlet fission in addition to the projection of the associated reorganization energy onto a large number of vibrational normal modes.

It was found that a specific vibrational mode of DPIBF is strongly correlated with the photo-induced spin separation, also suggesting a very strong vibronic coupling during singlet fission in DPIBF, due to the excessively large reorganization energy.

More importantly, the strong vibronic coupling in DPIBF thin film indicated the feasibility of improving singlet fission efficiency by modifying the vibrational density of states by temperature. For example, the pronounced anharmonicity of softened optical phonons at higher temperatures may result in a slower spin relaxation, responsible for the lower triplet quantum yield as observed by experiments.

IMPACT: The results of this research deepen an understanding of singlet fission at the electronic structure level of theory, such as the interplay between atomic displacements and electronic wave functions. Such a detailed comprehension will expedite the journey towards optimal solar energy conversion using environmentally friendly organic materials.

V. RAO KOTAMARTHI
Argonne National Laboratory
vrkotamarthi@anl.gov

Director's Discretionary
32.5 Million Core-Hours

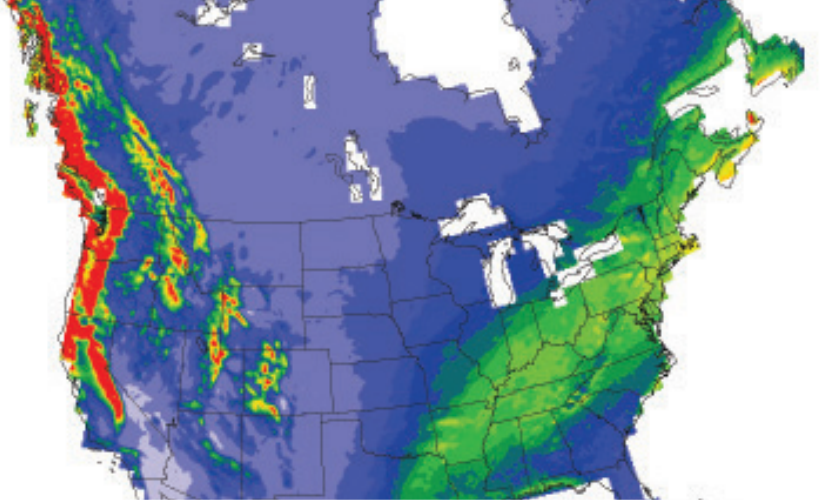

Image: Average winter precipitation rate (mm per day) for a 10-year period (1995 to 2004) as simulated by a regional climate model with 12-km spatial resolution.

Jiali Wang, Argonne National Laboratory

EARTH SCIENCE

DYNAMIC DOWNSCALING OF CLIMATE MODELS

Global climate models are used to simulate large-scale patterns suitable for weather forecasting and large-area climate trends, but they lack the level of detail needed to model changing climate conditions at local and regional scales. More accurate regional climate projections are needed to help policymakers and stakeholders develop adaptation strategies and mitigation measures that address the potential regional effects of climate change, including coastal flood risks, infrastructure vulnerabilities, and agricultural disruptions.

With a method called dynamical downscaling, researchers can use outputs from coarse-resolution global climate models to drive higher-resolution regional climate models. As part of an ongoing project at the ALCF, a research team from Argonne National Laboratory and the University of Chicago is using supercomputing resources to investigate the effectiveness of dynamically downscaled climate models. The researchers leveraged the open-source Weather Research and Forecasting (WRF) model, a widely used numerical weather and climate prediction tool, to develop a regional climate model that covers most of North America with a spatial resolution of 12 km (global models typically operate with a spatial resolution of 100 km to 300 km).

To optimize the regional climate model for Mira, ALCF staff helped the team to overcome an I/O bottleneck by creating an ensemble job script that allowed them to run simulations two- to four-times faster than before. In addition, ALCF researchers are working with WRF developers to explore the use of new parallel I/O algorithms that demonstrated the ability to speed the reading of boundary data by a factor of 10 in benchmark tests on Mira.

In a recent study, the team evaluated their regional climate model's ability to predict precipitation by comparing decades of simulations to observational data. This involved calculating two measures for precipitation patterns over North America: (1) spatial correlation for a range of distances and (2) spatiotemporal correlation over a wide range of distances, directions, and time lags. The researchers found that the correlations in the regional model output show similar patterns to observational data, and exhibit much better agreement with the observations than the global model in capturing small-scale spatial variations of precipitation, especially over mountainous regions and coastal areas. The results demonstrate the ability of downscaled regional climate models to represent the statistical characteristics of precipitation accurately.

IMPACT: This works seeks to validate the use of dynamic downscaling as a means to get climate models to the scales needed for local climate projections. Data from this project will be made publicly available, so local and regional decision makers can use the results to inform climate change adaptation strategies. The results will also be useful to the scientific community for a wide variety of studies on the impacts of changing climate conditions.

ROBERT MOSER

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

INCITE
175 Million Core-Hours

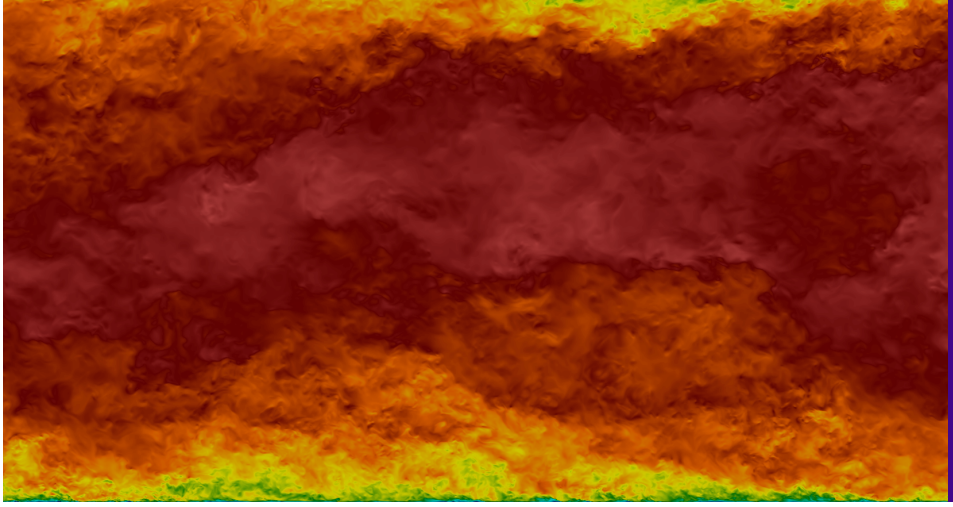

Image: This visualization depicts the instantaneous streamwise velocity component over a section of the simulated channel.

Nicholas Malaya, University of Texas at Austin

ENGINEERING

DIRECT NUMERICAL SIMULATIONS OF HIGH REYNOLDS NUMBER TURBULENT CHANNEL FLOW

A substantial amount of the energy consumed by a moving vehicle or boat is due to the drag and dissipation of energy caused by turbulence as it moves through air or water. The same forces are at work as air or liquid moves through ducts or pipes. A large reference dataset describing the behavior of turbulence is needed to better understand and model this phenomenon.

Researchers from the University of Texas at Austin used Mira to conduct the largest-ever direct numerical simulation (DNS) of this fluid dynamics problem—simulations at a Reynolds number (a dimensionless ratio of inertial forces to viscous forces) of 5,200 on a 15,360 x 1,536 x 11,520 mesh. Of particular interest to the research team was the overlap region, where viscous near-wall turbulence interacts with outer-layer turbulences. Previous simulations did not allow for a high enough Reynolds number to obtain the scale separation needed to understand the complexity of this turbulent interaction.

Due to the substantial power of Mira, the researchers were able to simulate a DNS with a high enough Reynolds number to generate sufficient scale separation. Collaboration with ALCF staff helped improve the management of cache and execution threads, doubling code performance. They also achieved near-perfect OpenMP scalability (99 percent) by minimizing the inter-memory access between OpenMP threads.

These simulations generated a new reference dataset that will remain useful for turbulence research for years to come. The channel flow data is now available for the public to use for calibrating their turbulence models, developing new turbulence models, and validating existing ones. The data is available online at turbulence.ices.utexas.edu.

IMPACT: The models resulting from this project will allow faster, more effective studies of turbulence, such as air-flow over vehicles, and could lead to improved vehicle surfaces and reduced-drag piping and ducts. Data from the project is also publicly available so it can be used by other researchers in their turbulence models and design projects.

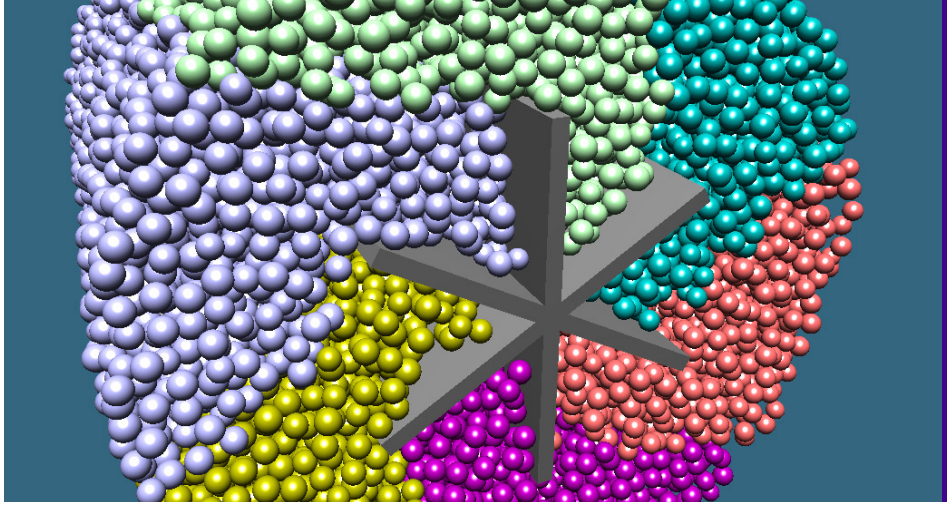

WILLIAM GEORGE

National Institute of Standards and Technology
wgeorge@nist.gov

INCITE
40 Million Core-Hours

Image: This simulation of cement mortar shows suspended particles in a virtual rheometer. The spheres are color coded by their starting location in the rheometer.

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

MATERIALS SCIENCE

COMPUTATIONAL DESIGN OF NOVEL MULTISCALE CONCRETE RHEOMETERS

Concrete is the most widely used building material in the world, but the production of its ingredients, namely cement, is a significant contributor to the world's greenhouse gas emissions. To enable the design of new, more sustainable mixtures of concrete, researchers from the National Institute of Standards and Technology (NIST) used ALCF supercomputers to advance the understanding of concrete's flow properties.

Given the critical importance of concrete to our nation's infrastructure, there is broad interest in making it a more sustainable material by reducing the amount of greenhouse gas created during its production, improving its strength and durability, and finding new ways to recycle it. These and other improvements require that scientists first find a way to accurately measure and control concrete's rheological (i.e., flow) properties to satisfy performance specifications.

Due to the complex nature of concrete, which is a dense suspension comprised of cement, water, and an aggregate, such as sand or gravel, it is a challenge to accurately measure its rheological properties. Through multiple INCITE awards at the ALCF, the NIST research team set out to advance the measurement science of concrete and to gain a fundamental understanding of how it flows. With the ALCF's massively parallel supercomputers, they were able to simulate how a suspension would change if one or more parameters varied, such as the number of suspended particles or their size.

Suspensions have a remarkable property: plotting viscosity vs. shear rate always generates the same shaped curve as plotting them for the suspending fluid alone without added particles. Adding solid particles to the fluid just causes the curve to move to a different location on the X-Y axis without altering its shape.

Through their simulations, the NIST researchers found they could use the microscopic shear rates that existed between neighboring particles to predict the amount that the curves had to be shifted for the suspension to change. Experiments at the University of Strasbourg confirmed the simulated results, which allowed the team to come up with a general theory of suspensions' properties. This discovery enables researchers to predict how fresh concrete will behave by measuring the properties of the fluid that the particles are placed in.

IMPACT: The results should help accelerate the design of a new generation of high-performance and eco-friendly cement-based materials by reducing time and costs associated with R&D. NIST is also using this new knowledge to create Standard Reference Materials for industrial researchers to calibrate concrete rheometers—instruments used to measure the flow of complex fluids—for materials development. Ultimately, this could help researchers zero in on the best new recipes for more sustainable concrete and expand the use of alternative materials.

J. ILJA SIEPMANN
University of Minnesota
siepmann@umn.edu

Director's Discretionary
8 Million Core-Hours

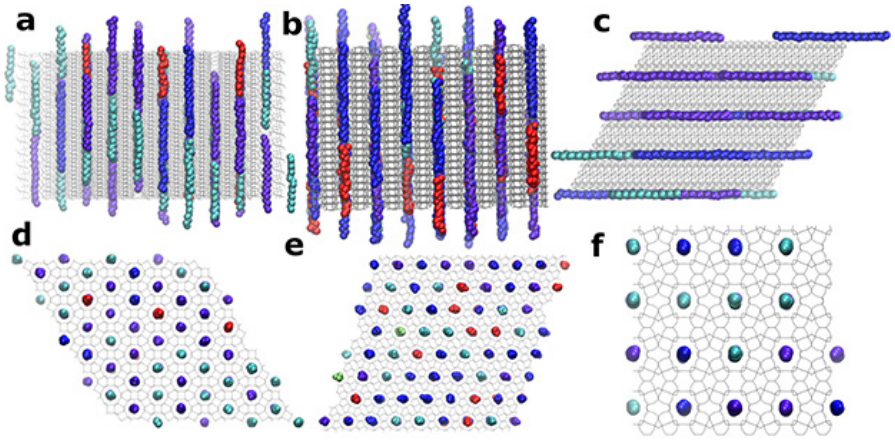

Image: Snapshots of representative hydrocarbon configurations inside zeolite frameworks. Views facing the main channel axis (a–c) and along the main channel axis (d–f) are shown for ATO (a,d), MTW (b,e) and PCOD-8113534 (c,f). Zeolite frameworks are depicted as grey lines, and n-octadecane, n-tetracosane, n-triacontane, (2-methyl and 4-methylheptadecane) and 2,2-dimethylhexadecane molecules as cyan, purple, blue, red and green spheres, respectively.

Peng Bai and J. Ilja Siepmann, University of Minnesota

MATERIALS SCIENCE

COMPUTATIONS FOR THE DEVELOPMENT OF THE NANOPOROUS MATERIALS GENOME

In the petrochemical and biofuel industries, aluminosilicate materials called zeolites are used as catalysts and molecular sieves to aid in the processing of fuels and chemical feedstocks. Researchers from the University of Minnesota and Rice University have used Mira to demonstrate a predictive modeling capability for identifying new zeolites that could improve fuel processing and production efforts.

To date, more than 200 types of zeolites have been synthesized and more than 330,000 zeolite structures have been predicted based on previous computer simulations. With such a large pool of candidates, using traditional laboratory methods to identify the optimal zeolite for a particular application presents a time- and labor-intensive process that could take decades.

To help accelerate this process, the research team used Mira to demonstrate a predictive theory and modeling tool that can quickly screen thousands of materials to pinpoint promising candidates for further research and development. ALCF staff helped ensure optimal performance on Mira by guiding the developers of the MCCCS-MN code in adding OpenMP support to permit hybrid MPI/OpenMP parallelism, and helping design an MPI-based framework to allow high-throughput calculations capable of using all of Mira's 786,432 cores.

Enabled by Mira's massively parallel architecture, the team's simulations successfully identified new zeolites for two important applications. For one study, the researchers discovered a zeolite with the ability to purify ethanol from fermentation broths in a single separation step, demonstrating the potential to replace an energy-intensive, multi-step distillation process currently used by industry. To validate the simulation results, University of Minnesota researchers synthesized and tested the promising zeolite, providing experimental data that was in very good agreement with the predictions.

For the second study, the team investigated potential zeolite-based catalyst systems for a dewaxing process called hydroisomerization, in which linear long-chain alkanes are transformed into slightly branched alkanes to reduce the pour point and increase the viscosity of lubricant oils and other fuel products. Their simulations identified zeolites with up to 100 times better adsorption capability than current technology uses for this process.

IMPACT: The computational screening tool demonstrated in this project has the potential to accelerate the discovery of new zeolites and metal-organic frameworks for gas storage, separations, and catalysis. By identifying optimal zeolites for particular tasks, this predictive modeling capability could help improve the production of biofuel and petroleum products, and the development of gas storage and carbon capture devices, while reducing the time and cost of associated laboratory research and development efforts.

JAMES AMUNDSON
Fermilab
amundson@fnal.gov

INCITE
50 Million Core-Hours

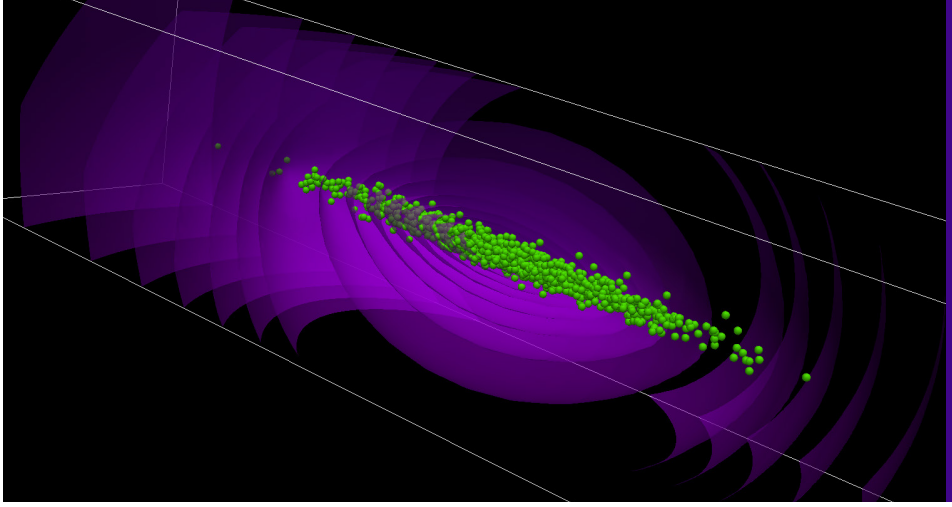

INTENSITY-DEPENDENT DYNAMICS IN FERMILAB AND CERN ACCELERATORS

Image: Synergia simulation of a bunched beam, including particles (green) and self-fields (purple).

James Amundson, Fermilab

PHYSICS

Particle accelerators are an enabling technology for both basic research and the applied sciences. The future of high energy physics requires running today's accelerators at higher intensities than ever before. To accurately understand intensity-dependent effects in their accelerators' complexes, researchers at both Fermilab and CERN are relying on advanced detailed numerical modeling.

For this INCITE project, a research team from Fermilab used ALCF resources to perform complex accelerator simulations aimed at reducing the risks and costs involved in developing the world's highest intensity particle beams. ALCF assisted in porting Synergia, a hybrid C++-Python code to the Blue Gene/Q platform and finding optimizations specific to the Blue Gene/Q architecture.

The Fermilab Recycler and Main Injector form the final high-energy stage of the Fermilab accelerator complex. During each acceleration cycle, the Recycler receives protons in six batches from the booster. Through a radio frequency manipulation procedure known as slip-stacking, the batches are combined to increase the instantaneous intensity. The protons are then transferred to the Main Injector where they are accelerated from 8 GeV to 120 GeV.

Simulations of the booster have been focused on understanding further details of instability generation caused by the wakefield mechanism, in which a charged particle traveling through a conducting pipe generates fields that affect trailing particles. Researchers succeeded in reproducing theoretically-predicted strong space charge modes in Synergia under simplified conditions.

The Fermilab Recycler was used to slip-stack at much higher intensities than previously has been performed. Current slip-stacking simulations of the Recycler aim to understand whether this energy difference combined with the space charge effects of higher intensity beams will result in increased particle losses.

IMPACT: These simulations will advance an understanding of the dynamics of intense particle beams, which will be used to create neutrino sources for the Long Baseline Neutrino Experiment at Fermilab and for the upgrade to the Large Hadron Collider at CERN. Beyond the realm of particle physics, this accelerator technology potentially could impact basic energy applications, such as the transmutation of nuclear waste.

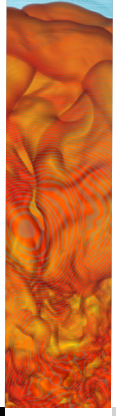
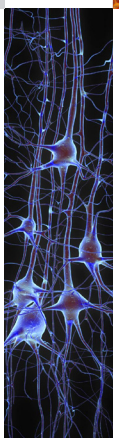

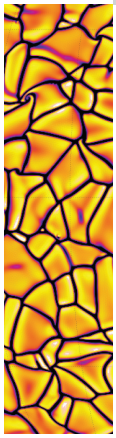

2015 SCIENCE HIGHLIGHTS

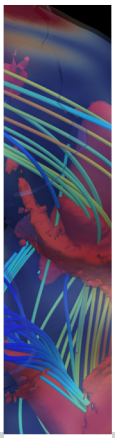

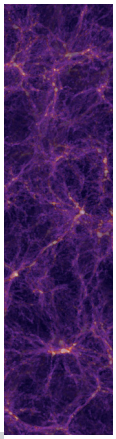
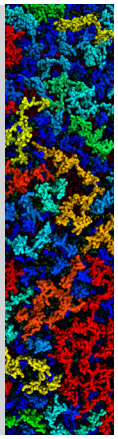
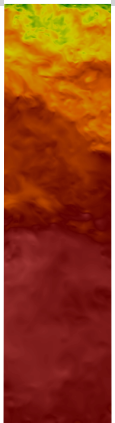

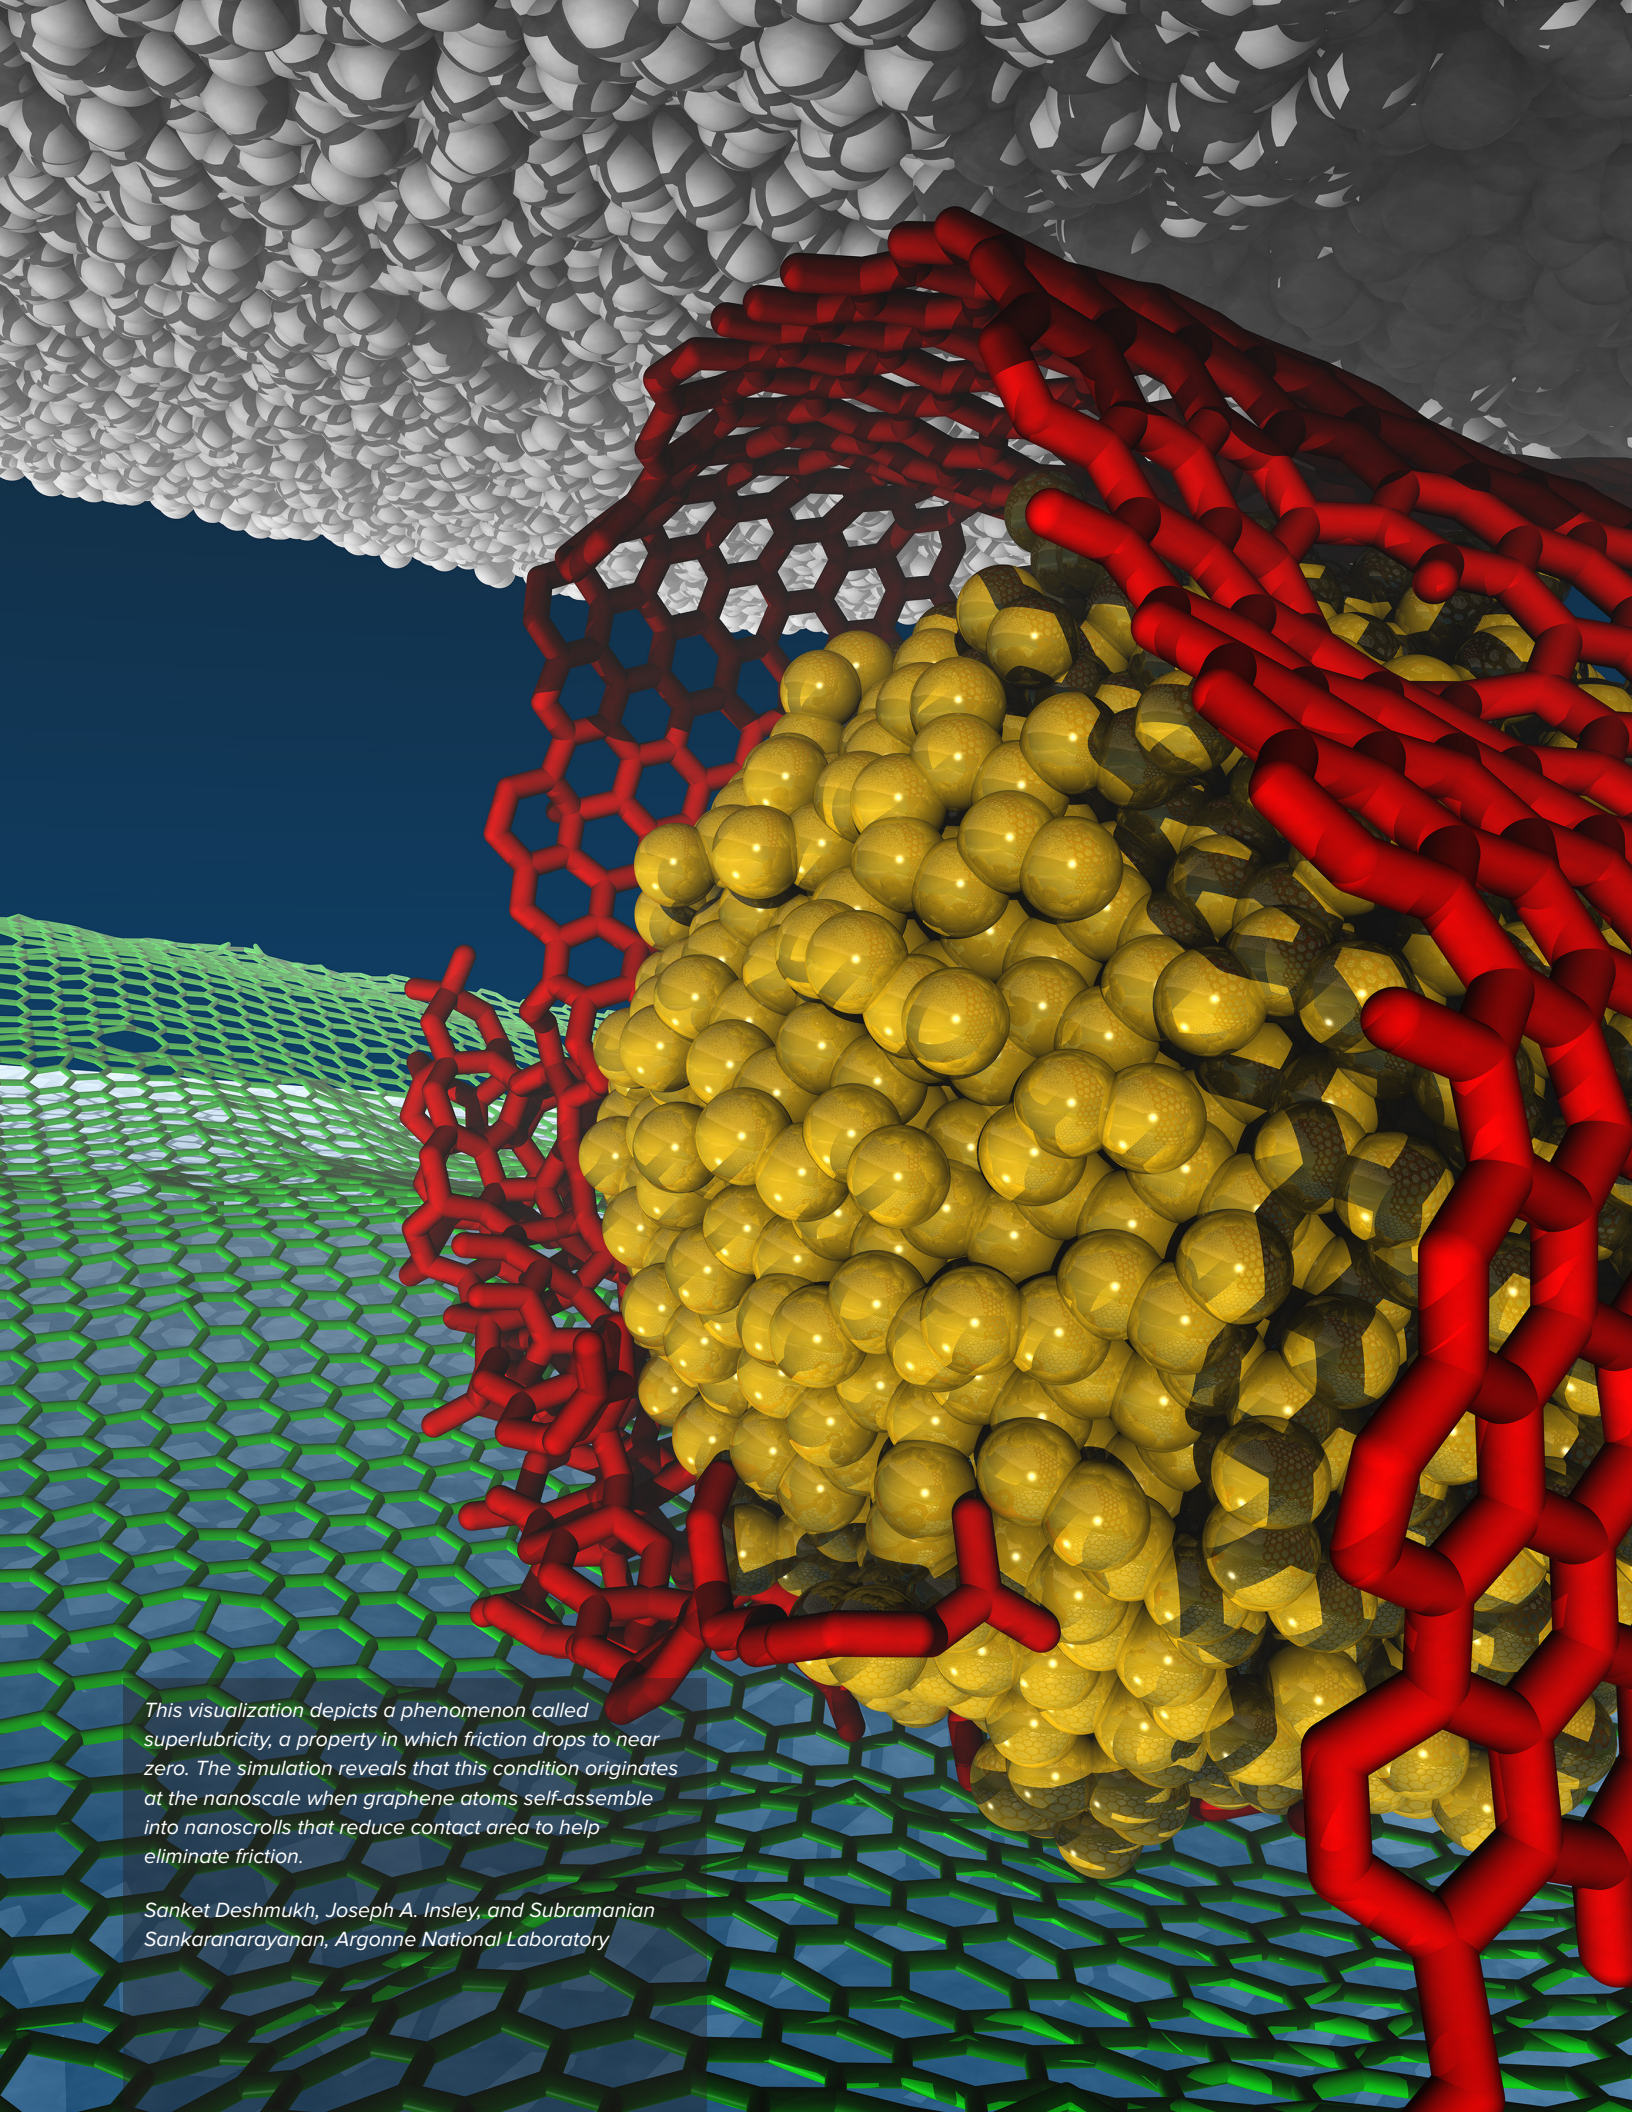

This visualization depicts a phenomenon called superlubricity, a property in which friction drops to near zero. The simulation reveals that this condition originates at the nanoscale when graphene atoms self-assemble into nanoscrolls that reduce contact area to help eliminate friction.

Sanket Deshmukh, Joseph A. Insley, and Subramanian Sankaranarayanan, Argonne National Laboratory

BENOÎT ROUX

University of Chicago
roux@uchicago.edu

INCITE
120 Million Core-Hours

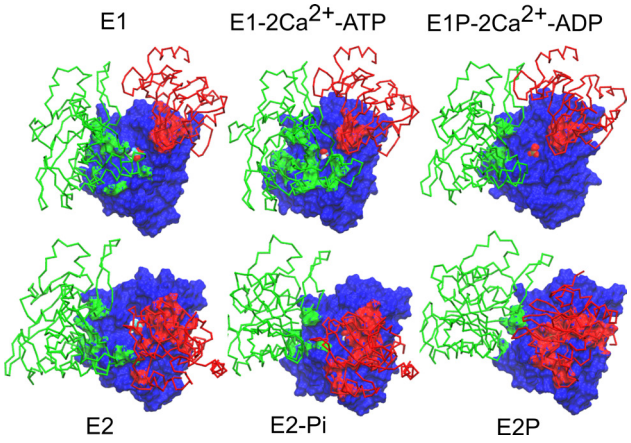

Image: Interaction of cytoplasmic domains in the calcium pump of sarscoplasmic reticulum. These six states represent important intermediates along the reaction cycle. Blue represents the phosphorylation domain (P). Red and green represent actuator (A) and nucleotide binding (N) domains, respectively. The red and green patches in the P domain are interacting with residues in A and N domains, respectively.

Avisek Das, University of Chicago

BIOLOGICAL SCIENCES

STUDIES OF LARGE CONFORMATIONAL CHANGES IN BIOMOLECULAR MACHINES

Complicated macromolecular machines abound in biology. The concerted action of these "molecular machines" underlies all the activities of the living cell. This research aims to gain a mechanistic understanding of the character of the free energy landscape that governs the function of large biomolecular systems, and demonstrate the feasibility of an extremely scalable computational strategy built upon the string method and classical molecular dynamics (MD).

Membrane-associated proteins play an essential role in controlling the bidirectional flow of material and information. Some of these proteins act like electric pumps, consuming adenosine triphosphate (ATP) to carry charged ions against their electrochemical potential gradient. Others, like the Src kinases, are turned on by a chemical modification called phosphorylation.

Understanding the detailed molecular mechanism of ion pumps has proved a long-standing problem in biology. The objective of this project is to simulate large-scale conformational transitions that are responsible for the pumping activity.

Biophysicists from the University of Chicago are working to elucidate the mechanisms of complex membrane proteins through a modified version of the nanoscale molecular dynamics (NAMD) code, ported to Mira. To determine these conformational transition pathways, researchers leverage NAMD with the string method—a computational methodology that achieves extreme scalability on leadership-class supercomputers.

The resulting simulations provide a deep mechanistic perspective of protein function, linking structure to dynamics by characterizing the free energy landscape, or potential of mean force, that governs key functional motions. In the case of the c-Src tyrosine kinase, the effects of its regulatory domains on the activation transition were studied using the string method with swarms-of-trajectory, which requires a set of collective variables to define the transition path. The simulation provides important insight into the mechanism of activation of c-Src tyrosine kinase, which plays a crucial role in cellular signal transduction.

Project researchers also have been studying large domain motions in an adenosine triphosphate (ATP)-driven calcium pump called SERCA, which plays an important role in the relaxation of skeletal muscles. This study has been completed and the final results provide a "molecular movie" of the structural mechanism of this important ion pump.

IMPACT: Results from this project will help develop a roadmap for simulating, visualizing, and elucidating the workings of complex molecular machines using extremely scalable computational strategies on leadership-class resources. It will also assist in the rational design of novel drugs for neurodegenerative diseases, such as Parkinson's.

NOA MAROM

Tulane University
nmarom@tulane.edu

ALCC
120 Million Core-Hours
(ALCF: 100M; NERSC: 20M)

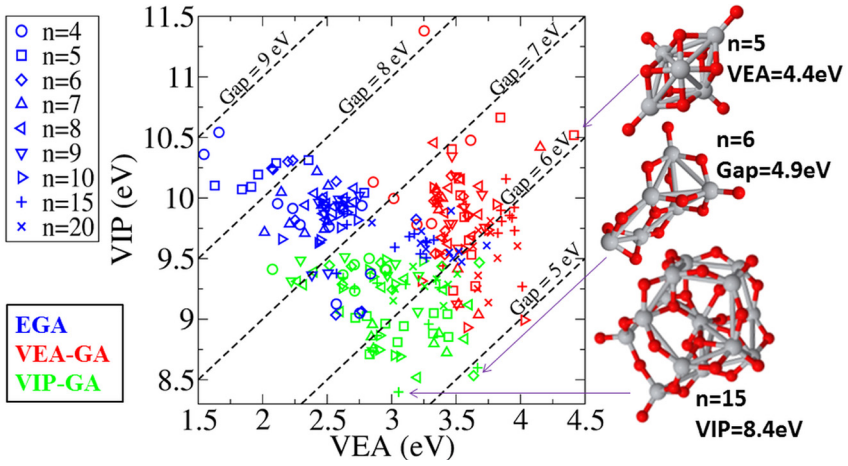

Image: Vertical ionization potential vs. vertical electron affinity for the best 10 structures found by each property-based genetic algorithm (GA) for $(TiO_2)n$ clusters with n=2-20. The loci of constant fundamental gaps are indicated by diagonal lines.

Reproduced from Phys. Rev. B 91, 241115(R) (2015)

CHEMISTRY

COMPUTATIONAL DESIGN OF INTERFACES FOR PHOTOVOLTAICS

The quest for clean sustainable energy is driving the development of emerging technologies, such as organic and hybrid organic-inorganic solar cells. The functionality and efficiency of these devices are often determined by interactions at the interface between two materials. This project is conducting large-scale, massively parallel first-principles quantum mechanical (QM) and molecular dynamics simulations to probe the physical attributes of these critical interfaces.

This comprehensive study will tackle organic-organic, organic-inorganic, and inorganic-inorganic interfaces, which can be found in emerging solar technologies, including organic photovoltaics and dye-sensitized solar cells (DSCs). To meet the rigorous computational demands required to adequately sample the vast configuration space of materials structure and composition, researchers turned to ALCF resources. Ported to Mira, the all-electron, full-potential electronic structure code FHI-aims performed fully QM simulations of the structure and electronic properties of the nano-structured functional interfaces found in organic and hybrid solar cells.

ALCF was also instrumental in helping the team implement a new parallelization scheme for GAtor, a versatile genetic algorithm (GA) package for structure prediction and design of molecular crystals. The new parallelization scheme enables GAtor to scale linearly, with the number of structures sampled on up to 105 cores.

Outcomes include the development of a new method, within the framework of many-body perturbation theory, to accurately calculate the ionization potentials (IPs) and electron affinities (EAs) of donor and acceptor molecules. A comprehensive benchmark study assessed the accuracy of different electronic structure methods for the IPs and EAs of 24 organic semiconductors, chosen from chemical families typically used in organic electronics.

Work on dye-sensitized titanium dioxide (TiO_2) clusters demonstrated that the energy differences at the dye-TiO_2 interface result from an interplay of quantum confinement and electrostatic screening effects at the nanoscale and that they can be manipulated by nano-structuring the oxide. Thus, interface engineering may be a path to more efficient DSCs.

A suite of property-based genetic algorithms was developed to design nano-clusters with tailored electronic properties for applications in nano-catalysis. It was demonstrated that the electronic properties of TiO_2 clusters depend on specific structural motifs more than on size.

IMPACT: This research will advance first-principles and multiscale simulations, and allow researchers to reveal essential details of these interfaces that are difficult to resolve experimentally. The theoretical understanding developed in this project will catalyze the emergence of new design paradigms for next-generation solar cell technologies.

GREGORY A. VOTH

University of Chicago
gavoth@uchicago.edu

ALCC
57.6 Million Core-Hours

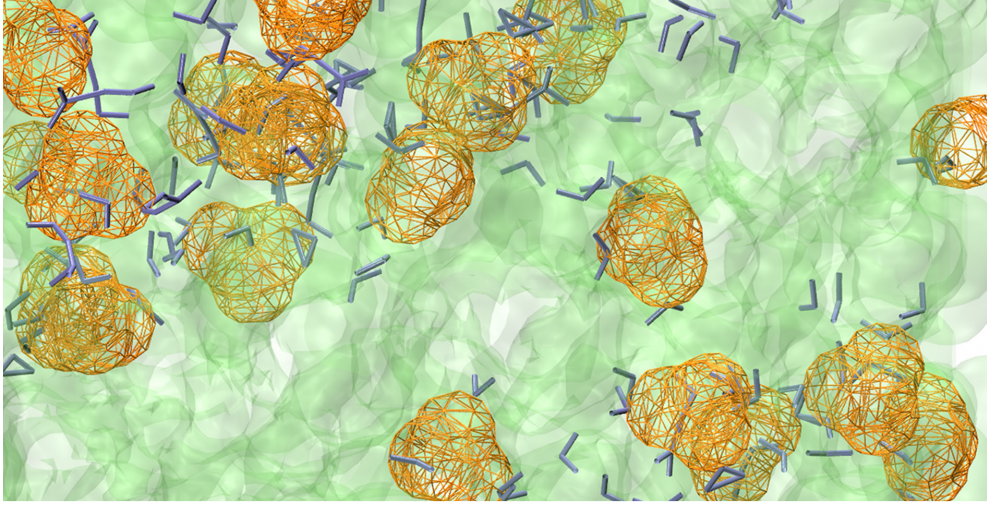

Image: Example of a water pathway that forms in polyelectrolyte membranes connecting two hydrophilic (water) regions. Water molecules are represented as blue stick figures, hydrated excess protons are orange wire-meshes, and Hyflon membrane is translucent green background.

Christopher Knight, Argonne National Laboratory

CHEMISTRY

INFLUENCE OF MORPHOLOGY ON PROTON TRANSPORT IN PROTON EXCHANGE MEMBRANES

A critical challenge in developing electrochemical energy-conversion devices, such as fuel cells, is the ability to produce these technologies at low cost with long lifecycles that can maintain high efficiency. Toward this end, researchers are analyzing detailed computer models to understand the factors that control proton transport in proton exchange membranes in order to improve their performance.

Proton exchange membranes are ion conducting polymers through which ions, such as hydrated excess protons (hydronium cations), flow to complete an electrical circuit. The internal shape, or morphology, of these membranes determines the structure and stability of water and ion pathways within them. Ideal membranes rely on good thermal and chemical stability to achieve highest-efficiency conductivity.

The complex chemistry and physics involved in the dynamics of these structures makes them difficult to study experimentally. Their study requires a molecular simulation tool that incorporates accurate molecular interactions, chemical reactivity, and robust parallelization algorithms to enable sampling of required spatial and temporal scales.

To meet these needs, University of Chicago researchers have teamed Mira with LAMMPS, an open-source molecular simulation code from Sandia National Laboratories, and their own reactive molecular dynamics (MD) simulation code, RAPTOR. The team simulated two perfluorosulfonic acid membrane (PFSA) morphologies: lamellae of polymer rods and bundles of polymer rods. Analysis of the hydronium diffusion constant showed very different behavior between the systems, with the polymer bundle results being in the best agreement with current experimental measurements. However, the large interface of the lamellar system appears to be more beneficial to proton transport than in the bundle and random systems, as proton transport was found to occur more effectively along hydrophobic-hydrophilic interfaces prominent in lamellar structures.

The team also performed additional simulations of a sodium/hydronium mixture to better understand the negative effect sodium has on the diffusion coefficient of hydronium in PFSA membranes. Using the same configurations used for the lamellar and rod simulations, they randomly replaced fractions of the excess protons in each simulation with sodium ions. An unexpected finding was that sodium diffusion was faster than proton diffusion and was sensitive to morphology with different binding strengths of protons to ionic sidechains.

IMPACT: Understanding how membrane properties correlate with the rate of proton transport will provide the knowledge required to design improved energy storage technologies. The detailed atomistic data generated from these reactive MD simulations will also help refine development of mesoscopic models of charge transport at significantly reduced computational cost.

DAVID BAKER

University of Washington
dabaker@u.washington.edu

INCITE
80 Million Core-Hours

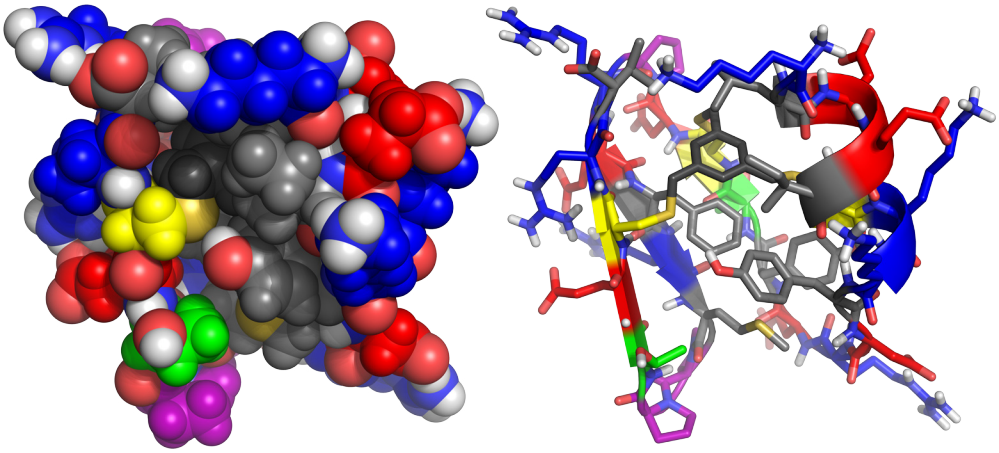

Image: A designed 29-residue peptide that uses the three-way crosslinker 1,3,5-tris(bromomethyl) benzene (dark gray) to covalently constrain the structure and to form hydrophobic packing interactions with other residues. This molecule is currently being synthesized for experimental examination.

Vikram Mulligan, University of Washington

CHEMISTRY

TOWARDS BREAKTHROUGHS IN PROTEIN STRUCTURE CALCULATION AND DESIGN

Proteins are large, complex molecules that drive virtually all cellular functions in living organisms. With the emergence of protein structure modeling tools, researchers have the ability to design proteins with targeted applications, such as treating diseases and catalyzing medically and industrially useful reactions. A research team from the University of Washington is using Mira to develop and apply new computational methods aimed at enhancing protein structure prediction and design capabilities.

The Rosetta software suite, developed at the University of Washington's Baker Laboratory, is designed to tackle to two difficult computational problems: the prediction of protein structure from amino acid sequences, and the design of new amino acid sequences to yield a desired function.

With this INCITE project, researchers are using Mira to enhance Rosetta's ability to sufficiently sample the vast conformational space and address other emerging challenges in protein structure calculation and design. Their work includes improving the Rosetta energy function and further developing homology model refinement methods.

In addition, the researchers are using Mira to design mini-proteins called peptides for a wide range of therapeutic targets, including influenza, Ebola, HIV, and Alzheimer's disease. Like proteins, peptides have larger binding surface areas than conventional small-molecule drugs. This allows peptides to bind to their targets with improved efficacy, reducing the potential for side effects.

To enable the computational design work, the team developed a multistate design approach that allows them to create novel peptides with unique, rigid folds. Their method involves exhaustively enumerating the possible conformations of small peptides with many geometric constraints (e.g., covalent cross-links, terminal peptide bonds). It then uses all of the possible conformational states to search for a sequence that uniquely stabilizes one desired structure and destabilizes all other alternative structures. The search algorithm benefits enormously from Mira's massively parallel architecture, particularly when assigning each conformational state to a separate core.

IMPACT: This project is advancing protein structure modeling capabilities to enable the design of novel proteins, including therapeutic peptides that target diseases such as Ebola, HIV, and Alzheimer's. Artificial peptides represent a new class of drugs that have potential for greater efficacy and fewer side effects than traditional drugs. In addition, the computational tools can be used to design peptide catalysts and enzymes for various environmental, energy, and industrial applications.

RON BRIGHTWELL

Sandia National Laboratories
rbbrigh@sandia.gov

ALCC
40 Million Core-Hours
(ALCF: 5M; NERSC: 5M; OLCF: 30M)

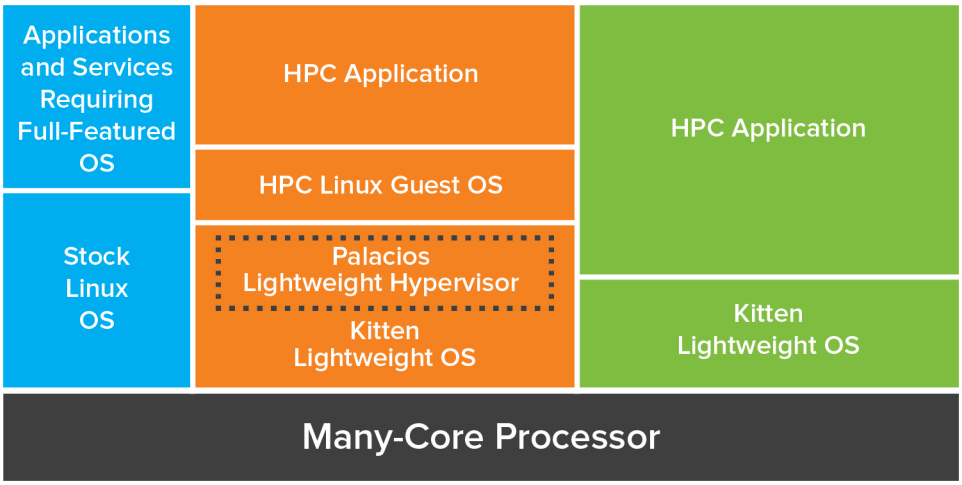

Image: A diagram of the Hobbes operating system multi-stack architecture.

Ron Brightwell, Sandia National Laboratories

COMPUTER SCIENCE

HOBBES: OPERATING SYSTEM AND RUNTIME RESEARCH FOR EXTREME SCALE

Four key challenges facing future large-scale computing systems are: dramatically improve power efficiency; improve resilience in the presence of increasing faults; enable efficient data movement across deepening memory hierarchies and new storage technologies; and manage dramatically increased parallelism. To address these challenges, the operating system and runtime (OS/R) must take on more responsibility for managing more resources, such as power and parallelism, and share more of the burden for insulating applications from the complexities of a system. Much of the focus of extreme-scale system software in the last decade has been on measuring and characterizing the impact that the OS can have on application scalability.

The Hobbes project is a collaboration of four national laboratories and eight universities with the goal of providing a system software environment that enables application composition through lightweight virtualization. Rather than providing a single unified OS/R that supports several parallel programming models, Hobbes is leveraging lightweight virtualization that provides the flexibility to construct and execute custom OS/R environments. While much of the exploration and development of the Hobbes software environment can be carried out at small scales, evaluating the scalability of OS/R interfaces and mechanisms at large scales is crucial.

The team's present focus is the development of a scalable implementation of the basic system architecture. When completed, this component will support the remainder of the work outlined for the project. The team has begun by mapping existing software infrastructure such as Kitten, Palacios, MRNet, Nessie, and ADIOS to the proposed system architecture. It will then supplement these components with additional infrastructure as necessary. Each component on which the basic architecture is based has been demonstrated to scale well to leadership-class systems; nevertheless, maintaining scalability while integrating these components remains a significant challenge.

Use of Mira is important as it features multiple communication topologies—an important component of the team's load-balancing testing—and will be used for both scalability/production runs and for preparation/development runs. In addition, Mira's unique design facilitates the booting of custom kernels.

IMPACT: This project supports large-scale testing and evaluation of operating system and runtime implementation techniques to improve OS/R for future high-performance computing machines. Its aim is to deliver an OS/R environment for extreme-scale scientific computing—one with a versatile structure with a lightweight virtualization layer that provides an ideal research foundation for the entire OS/R community.

ROBERT LATHAM

Argonne National Laboratory
robl@mcs.anl.gov

INCITE
25 Million Core-Hours

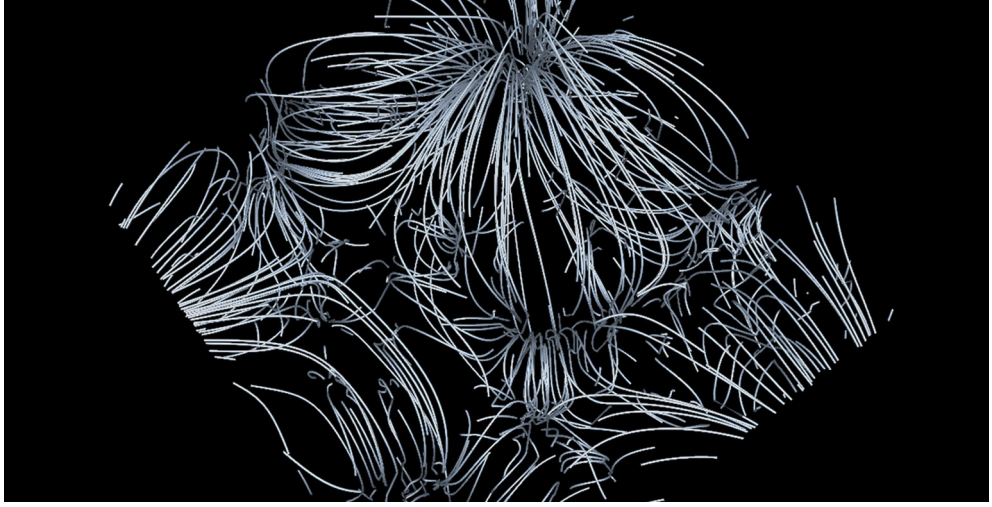

Image: Streamlines from an early time step of the Rayleigh-Taylor instability depend on scalable storage, communication, and data analysis algorithms developed at extreme scale using INCITE resources.

Tom Peterka, Argonne National Laboratory

COMPUTER SCIENCE

SCALABLE SYSTEM SOFTWARE FOR PARALLEL PROGRAMMING

The Scalable Systems Software for Parallel Programming (SSSPP) project designs, develops, and refines the system software available on leadership-class platforms. These efforts cover message passing, I/O, operating systems, and analysis. The tools and techniques developed and deployed through this research provide the foundation for scientific computing on leadership-class machines.

Simulation and modeling is increasingly being used by the research community to analyze HPC system design decisions prior to building the real architectures. It can also be used by researchers to gain insight into system performance, where access to the real architectures is not possible. SSSPP researchers utilized parallel discrete-event simulations to build accurate and efficient models of current and future HPC networks, including torus networks offered by the IBM Blue Gene architectures. The simulations were executed at a large scale, ranging from a few thousand network nodes to the size of a projected exascale system, providing network designers the ability to explore design options with a variety of synthetic and real HPC application network traces.

These I/O trace extrapolations, the ability to easily capture and repeat network communication events, were further investigated in the context of the ScalaTrace MPI/POSIX tracer, inspired by the ScalaExtrap communication extrapolation tool. Results suggest that while it is often difficult or impractical to retrieve application traces at scales of interest for study, having the ability to extrapolate traces based on access pattern analysis allows for the synthesis of high-fidelity, large-scale application workloads, without requiring leadership-scale execution.

While current work has focused on applications for petascale systems, initial evaluations have begun on future exascale designs through the performance of million-node simulations of network topologies. The massively parallel ROSS event simulator scales to significant fractions of Mira, and allows for the evaluation of future designs by using current hardware. Exascale designs will require fault tolerance, resilience, and power-aware approaches in system software. This research will explore that design space with projects like fault tolerant message passing interface MPI (MPIXFT), multi-level check-pointing (FTI), and power-aware monitoring (MonEQ).

IMPACT: Work done through SSSPP will allow computational scientists to simultaneously achieve performance and productivity on extreme-scale systems by effectively allowing applications to utilize the processing, communication, and storage capabilities hidden under the architectural complexity of leadership-class computers.

JONATHAN AURNOU

University of California, Los Angeles
aurnou@ucla.edu

INCITE
83 Million Core-Hours

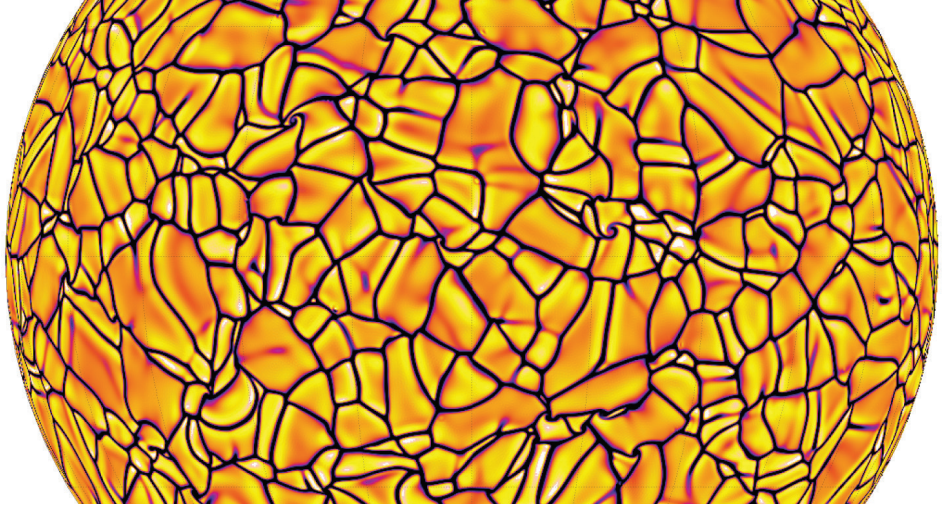

Image: Radial velocity near the surface of a turbulent solar simulation run on Mira. This simulation possesses a density stratification of four density scale heights and a Rayleigh number of about 3x109. Upflows are denoted in yellow-orange tones, and downflows are indicated by dark blue tones.

Nicholas Featherstone, University of Colorado Boulder

EARTH SCIENCE

FRONTIERS IN PLANETARY AND STELLAR MAGNETISM THROUGH HIGH-PERFORMANCE COMPUTING

Magnetic fields are generated deep within the interiors of stars and planets. This process, known as dynamo action, relies on physical mechanisms that remain largely, if not completely, inaccessible to direct observation. The aim of this research is to leverage high-performance computing tools to develop a new generation of dynamo models describing the dynamics at play in the sun, Earth, and Jupiter.

Most planets and stars produce an outward flux of heat throughout some fraction of their radius by the bulk motion of a conducting fluid, whether iron in the Earth's core or dense plasma at the heart of a star. In either case, the motion of the fluid, sustained through this outward flux of heat, gives rise to electrical currents that generate a magnetic field.

Solar and terrestrial magnetism plays an important role in our modern technological society where we are largely shielded from explosive, magnetically driven, solar phenomena by the Earth's magnetic field. Understanding and, ultimately, predicting the potential short- and long-term impacts of these phenomena on the Earth, requires a coherent picture of the dynamo processes at the heart of this two-body system.

A team of geo- and astrophysicists is developing state-of-the-art computational models to describe the interior dynamics of the sun, Earth, and Jupiter, using Rayleigh, a pseudo-spectral code designed to study magnetohydrodynamic convection in spheres. Using the petascale power of Mira, the team has been able to construct high-resolution models and resolve a range of spatial scales previously inaccessible to numerical simulation.

These models are allowing a detailed study of how convection—occurring on spatial scales ranging from the sun's diameter to those roughly 1,000 times smaller—transports energy generated by fusion deep within the sun's core to the solar surface.

Their simulations are now beginning to answer important questions, such as how the velocity amplitude and structure of stellar interiors relates to energy throughput of the system. Answering this question is a necessary first step in understanding the dynamics that give rise to the 22-year magnetic cycle of the sun and, ultimately, the explosive solar weather events that accompany the rise and fall of that cycle.

IMPACT: These models will open new windows into the understanding of the interplay of magnetism, rotation, and turbulent convection occurring within the remote interiors of geophysical and astrophysical bodies. The research will also provide the broader community access to the singular, extreme datasets generated by these massive computational efforts.

THOMAS H. JORDAN

Southern California Earthquake Center
tjordan@usc.edu

INCITE
167 Million Core-Hours
(ALCF: 48M; OLCF: 119M)

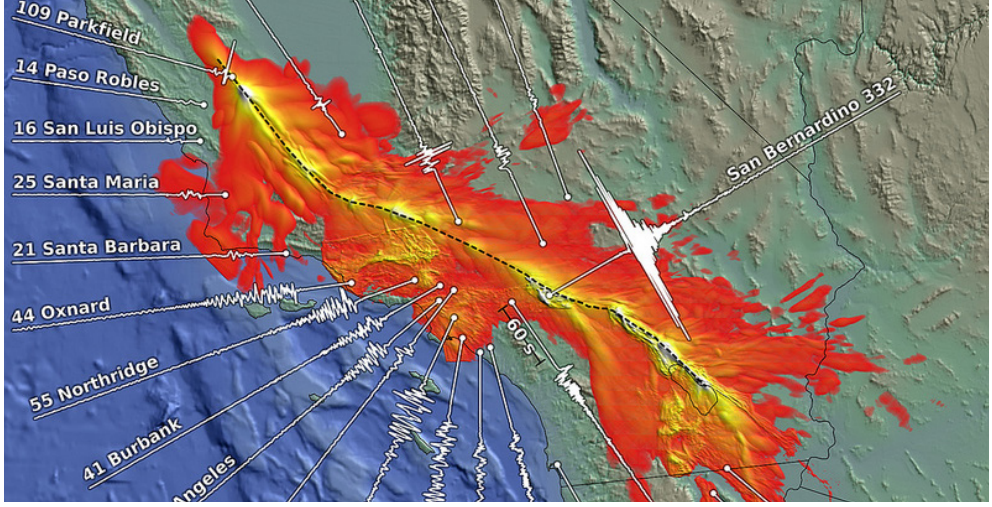

Image: Orange, yellow, and white colors on this map of California reveal regions where strong ground shaking would occur during a magnitude-8 earthquake on the San Andreas Fault. The seismograms shown on the map indicate peak velocity ground motions for selected California cities. SCEC is using Mira to produce detailed 3D seismic velocity models for use in scenario earthquake simulations like this one, which shows where strong ground shaking would produce significant hazards to people and infrastructure. SCEC's CyberShake hazard model calculates long-term seismic hazard estimates by simulating all plausible future earthquakes in a region of interest, which are then combined to produce estimates of peak ground motions expected at sites of interest over the next 50 years.

Geoffrey Ely, Argonne National Laboratory

EARTH SCIENCE

HIGH-FREQUENCY GROUND MOTION SIMULATION FOR SEISMIC HAZARD ANALYSIS

Earthquake ground motions create seismic hazards that put our modern urban infrastructure at risk, and create economic exposure to earthquake devastation. Improved earthquake hazard estimates will help engineers improve the resilience of our cities to future earthquakes. Physics-based modeling and simulation provide a path to more accurate representations of earthquake systems—from the rupture at the fault to the response of the built environment. This approach relies on numerical simulation of rupture dynamics and seismic wave propagation in realistic 3D models of the crust's heterogeneous structure to represent the ground motion during strong earthquakes.

The Southern California Earthquake Center (SCEC)'s research team seeks to produce simulations at a level of resolution valid for engineering applications (i.e., at frequencies higher than previously used). Higher frequency earthquake system modeling needs to be coupled to engineering models of infrastructure systems, such as buildings, bridges, and other critical distributed systems (e.g., lifeline and medical networks), that depend strongly on how complex earthquake wavefields interact with the mechanical heterogeneities of the ground and the built landscape, including both off-fault and near-surface plasticity.

Using Mira, the SCEC team has made rapid progress developing and applying the full 3D tomography (F3DT) computational techniques needed to evaluate and improve 3D seismic velocity models in California and other regions. The SCEC group has produced the fourth F3DT iteration of a new central California crustal seismic velocity model whose purpose is to help generate more accurate ground motion predictions. Many features revealed in the model are consistent with independent geophysical observations in central California, including controlled-source tomography, gravity anomalies, and the locations of active faults. The model covers a geographical volume of central California that is 446 km long by 509 km wide and 50 km deep. It will now be used as the preferred 3D seismic velocity model in future scenario earthquake ground motion simulations and CyberShake probabilistic seismic hazard calculations. These models are registered into the SCEC's open-source Unified Community Velocity Model software framework so they can be queried and used to build velocity meshes for ground motion simulation.

IMPACT: Through careful evaluation and validation of earthquake simulations and close collaboration between seismologists and civil engineers, this effort is helping to establish and expand the use of simulated ground motions for engineering purposes. The project is developing the computational framework and scientific and engineering processes needed to evaluate alternative ground motion simulation methods, leading to the production of more realistic synthetic seismograms and more accurate ground motion forecasts for use in earthquake engineering.

ANNE DORD

GE Global Research
anne.dord@ge.com

INCITE
89 Million Core-Hours

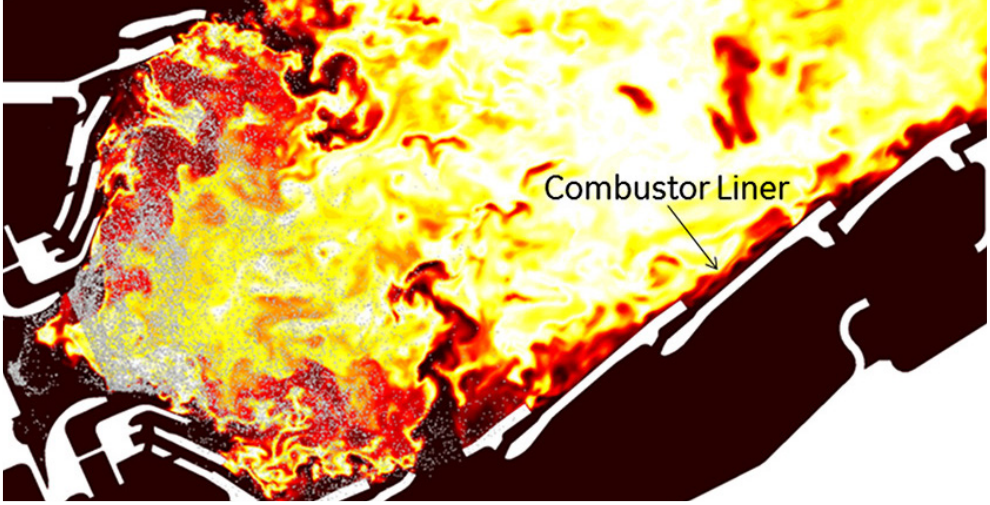

Image: Large-eddy simulations of a GE CFM combustor, temperature distribution, and liquid fuel droplets.

Anne Dord, GE Global Research

ENGINEERING

LARGE-EDDY SIMULATIONS OF COMBUSTOR LINER FLOWS

The commercial aviation industry spent an estimated $207 billion on fuel in 2012, or about 33 percent of operating costs. That amounts to approximately five million barrels of oil per day, or about 6 percent of the world's total oil usage. Aviation gas turbine engines are also significant producers of CO_2 and other harmful greenhouse gas emissions. As the aviation industry continues to grow, the cost and environmental impacts of operating gas turbine engines will continue to increase as well.

In the second year of this multi-year INCITE project, researchers are continuing their studies of the complex near-wall physics of combustor liner flows with a focus on aircraft engine and gas turbine applications. Understanding and predicting the aero-thermal flow field are keys to designing and optimizing combustors for better fuel efficiency, lower emissions, and improved performance. The proposed calculations require significantly greater computational resources than those traditionally used in industry, necessitating the use of DOE's leadership-class supercomputers; massively parallel computing platforms such as Mira are particularly well suited to performing these calculations.

The calculations use large-eddy simulations (LES) to model the behavior of an idealized configuration representative of combustor liners. The same approach will be applied to a more complex configuration, involving a larger domain and more realistic flow characteristics representative of the large-scale flow unsteadiness present in aircraft engines.

The team has studied the aerodynamics characteristics of impingement multi-hole cooling jets on a real GE gas turbine's combustor transition piece, which connects the main combustion zone and the high-pressure turbine system. The profiles developed from the impingement jets can cause impacts on the inlet profiles of the fuel nozzles and therefore affect overall fuel-air mixing performance. The team focused on the effect of using a model of the combustor liner wall, finding that the hybrid wall-modeled LES predicted a more reasonable distribution of wall shear stress on the impingement wall. This result is in marked contrast to the baseline LES case that does not rely on a model of the liner wall and which predicted significantly lower values. Further study will quantify the impact on flow-split characteristics, as well as combustion characteristics.

IMPACT: The new insights gained using this approach will enable modelers to generate high-fidelity datasets that will be used to improve on the low-fidelity models currently available to designers. Designers, in turn, will better understand the complex and unsteady processes governing the aero-thermal field around combustor liners. The results will assist the gas turbine industry in designing and optimizing aircraft engines and gas turbines for better fuel efficiency, lower emissions, and improved performance.

QIQI WANG

Massachusetts Institute of Technology
qiqi@mit.edu

Director's Discretionary
2 Million Core-Hours

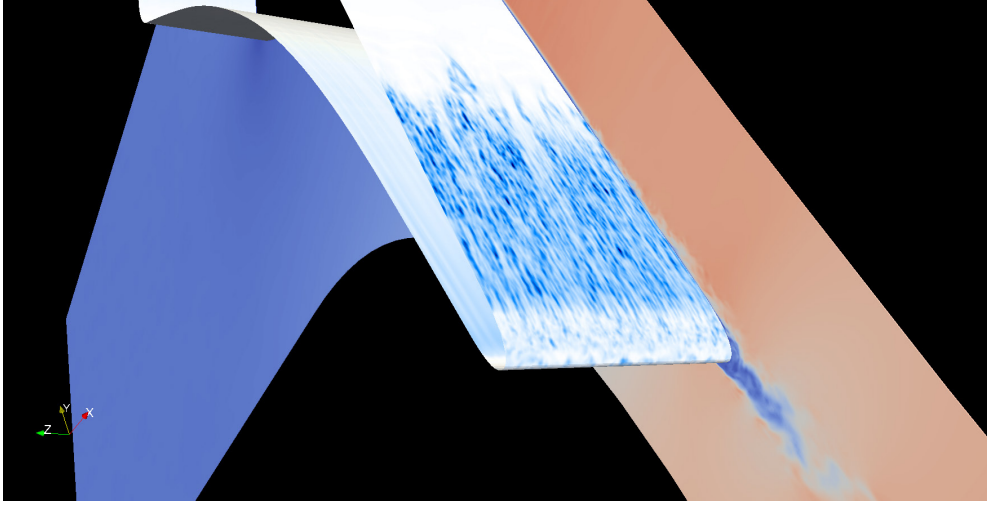

Image: The trailing edge optimized for heat transfer and aerodynamic loss. The image shows skin friction on the blade surface and Mach number contour on the cut plane. The LES-based optimization performed on Mira achieved a 17 percent reduction in heat transfer over the trailing edge region and a 21 percent reduction in pressure loss.

Patrick Blonigan and Chaitanya Talnikar, Massachusetts Institute of Technology

ENGINEERING

PARALLEL OPTIMIZATION OF TURBULENT FLOW SIMULATIONS

Parallel optimization requires additional concurrency in addition to state-of-the-art, grid-level parallelism. An array of simulations, each performed for a different geometry or flow control strategy, run at the same time. During their time integration, the simulations communicate their partially converged objective functions with each other. The inter-simulation communication allows each simulation to choose from the following three options: (1) refinement, which continues converging the current simulation; (2) exploitation, which moves to simulate the best design and is estimated based on the current information from all simulations; or (3) exploration, which moves to simulate an outside-the-box design.

The proposed parallel optimization method will attempt to discover the optimal geometry or flow control strategy by concurrently running an array of simulations, communicating their results, and making selections during runtime. It has many important engineering applications. Further, to run such an optimization quickly enough for these applications requires the use of extreme-scale parallel computers like Mira.

The team has succeeded in testing a parallel optimization prototype for large-eddy simulations (LES), finding the Bayesian optimization framework to be suitable for parallel optimization, both in optimizing the flow control strategy in a turbulent channel and in optimizing the geometry of a turbine blade trailing edge. The LES-based optimization performed on Mira achieved a 17 percent reduction in heat transfer over the trailing edge region and a 21 percent reduction in pressure loss, compared to the baseline semi-elliptic trailing edge. These researchers will continue investigating and demonstrating the optimization in more realistic applications.

IMPACT: Parallel optimization of turbulent flow simulations is an emerging application enabled by extreme-scale computing. Parallel optimization runs an array of flow simulations and adjoint gradient calculations simultaneously, each being massively parallel, to collectively optimize a geometry or flow control strategy.

SIBENDU SOM
Argonne National Laboratory
ssom@anl.gov

MARTA GARCÍA
Argonne National Laboratory
mgarcia@alcf.anl.gov

Director's Discretionary
6 Million Core-Hours

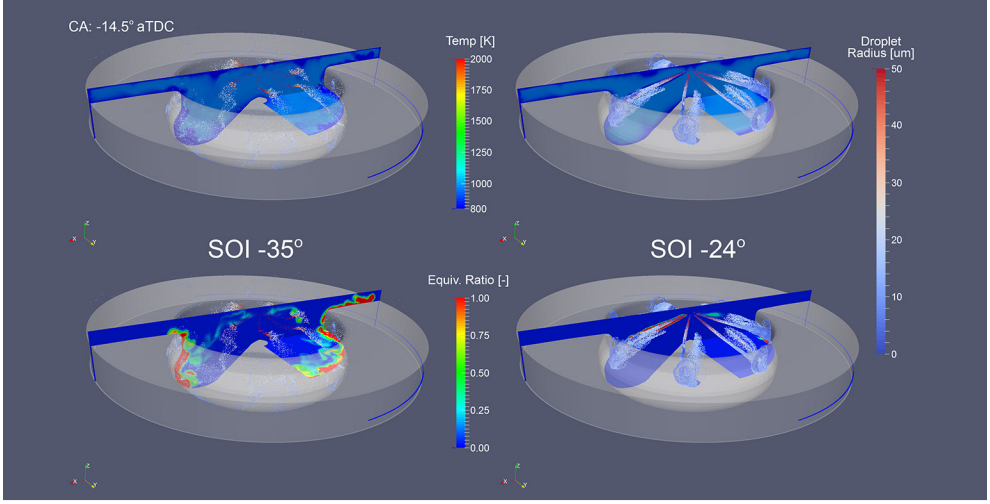

Image: This image compares an earlier Start of Injection (SOI) timing (SOI -35°) to a later SOI timing (SOI -24°). Temperature and equivalence ratio are shown on planes that slice through the domain. With this project, VERIFI researchers demonstrated that in gasoline compression ignition, varying the SOI timing produces significant differences in the reactivity of the fuel mixture, delaying ignition.

Stephen Ciatti, Marta García, Kevin Harms, Joseph A. Insley, Janardhan Kodavasal, Christopher Kolodziej, and Sibendu Som, Argonne National Laboratory

ENGINEERING

PERFORMANCE IMPROVEMENT OF CFD CODE CONVERGE ON BLUE GENE/Q SYSTEMS

Engine modeling and simulation tools have the ability to optimize complex fuel spray and combustion processes for a variety of fuels over a wide range of operating conditions, helping automotive manufacturers improve engine efficiency and performance, while reducing development costs and accelerating time to market. Argonne researchers are using ALCF resources and expertise to develop more powerful modeling and simulation capabilities that can shed new light on the complex processes taking place inside of internal combustion engines.

Researchers from Argonne's Virtual Engine Research Institute and Fuels Initiative (VERIFI) are working with ALCF staff and Convergent Science Inc. to optimize the CONVERGE code, a commercial computational fluid dynamics (CFD) software tool, for high-performance computing systems. Prior to the team's work at the ALCF, CONVERGE was used to run CFD simulations on up to 50 cores. The ongoing collaborative effort has enabled the code to perform its largest engine simulation to date running on 8,192 cores on ALCF systems.

The performance optimizations have been focused on the near-ignition portion of the simulation, which is typically the most computationally intensive due to the small timescales associated with calculating chemical kinetics in every CFD cell. As part of the work, the Argonne team modified CONVERGE to use parallel read/write processes, which allow simultaneous file writing by processors, rather than having to wait for one action to complete before taking on another. This resulted in more than a 100x speedup in writing large data files generated by the software.

The research team also developed a load balancing scheme for carrying out chemistry calculations in CONVERGE. The varying levels of complexity in the chemistry of ignition and combustion meant that some cores were not actively engaged in computation, while other cores were working overtime to handle more complex parts of the simulation. By balancing the computational workload evenly across more cores, this effort resulted in an 8x improvement in load balancing and a 3.4x improvement in time to solution. The enhancements achieved in this project are machine independent, and expected to benefit CONVERGE users running the code at scale on any high-performance computing system. These optimizations can also be applied to other CFD software used for engine simulations.

IMPACT: This project aims to develop novel engine modeling and simulation capabilities for supercomputers, providing automotive manufacturers with a cost-effective tool to accelerate the development of more energy-efficient engines. Ultimately, these tools have the potential to improve the fuel economy of vehicles, thereby reducing U.S. dependence on foreign oil and reducing carbon emissions.

PARVIZ MOIN
Stanford University
moin@stanford.edu

ALCC
120 Million Core-Hours

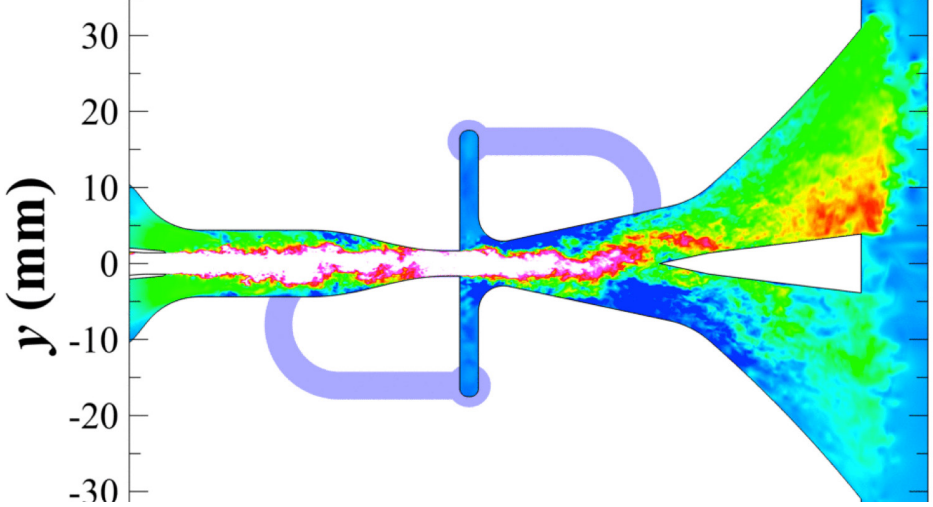

Image: Detailed simulation inside a flow control device. The main jet oscillates between the upper and lower outputs (separated by a triangular block in this cross-section); contours of streamwise velocity are shown. Actual aircraft will have dozens of these ports installed along the tail and wings for control.

Jeonglae Kim, Stanford University

ENGINEERING

PREDICTIVE LARGE-EDDY SIMULATION OF JET FUEL ATOMIZATION, HIGH-LIFT AIRFRAMES, AND REACTING SUPERSONIC TURBULENT FLOWS ON UNSTRUCTURED GRIDS

Accurate predictive modeling is crucial in the design of energy-efficient and environmentally friendly engineering systems, including aircraft propulsion and land-based power generation. High-fidelity, unstructured large-eddy simulation (LES) is emerging as an accurate yet cost-effective computational tool for prediction of several key aircraft components. Combining recent advances in LES modeling with the highly scalable, unstructured code CharLES makes possible the predictive LES of real aircraft geometries at flight Reynolds numbers using today's leadership-class supercomputers like Mira.

In this past year, the team's main focus were simulations to test the flight feasibility of active flow control devices as well as the atomization and mixing processes in high-pressure fuel injectors. There is a great deal of activity, both experimental and computational, in evaluating the ability of flow control devices installed on a plane's tail rudder to help a pilot steer the plane, especially at the low speeds characteristic of takeoff and landing or in a strong cross wind. When installed, these devices act to blow high-speed air across the tail's rudder to increase the tail's steering effect at low speeds. This design might allow aircraft manufacturers to install smaller-size tails and significantly reduce the drag experienced in cruise conditions on long flights.

Several simulations of the flow inside such devices were performed this year on Mira. The results accurately predicted the bi-stable oscillatory blowing, frequency-velocity relationship, and wall-pressure distribution through the flow control actuator in accord with experimental measurements. This detailed information has allowed the team to develop suitable wall boundary conditions to be used in full-scale simulations of a tail rudder where dozens of such actuators may be installed. Using Mira, the team was able to test such models on NASA's energy-efficient transport airfoil at several angles of attack and compare results with ongoing experiments.

IMPACT: This project aims to demonstrate that LES of complex geometries at flight Reynolds numbers is possible with today's large-scale supercomputers, which is a crucial step toward extended use of LES in transportation and power industries.

HASSAN NAGIB
Illinois Institute of Technology
nagib@iit.edu

ALCC
11 Million Core-Hours

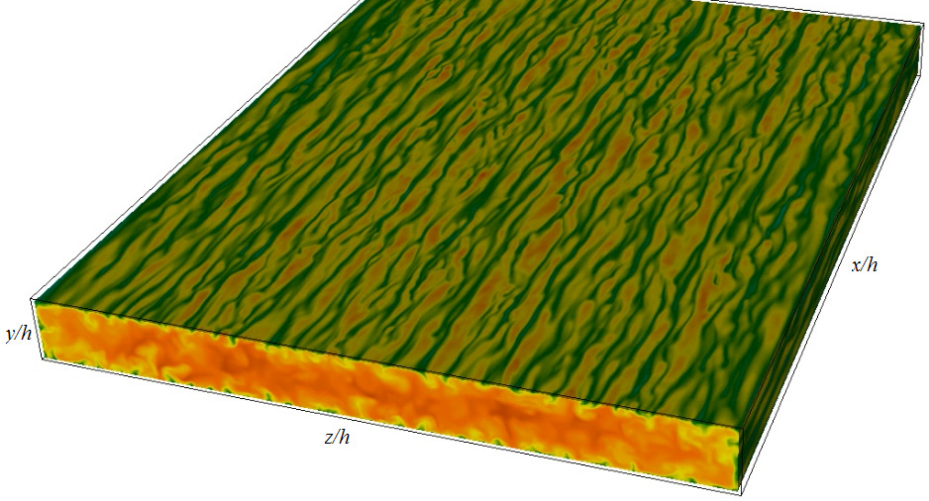

Image: Instantaneous realization of the streamwise velocity field in the AR=10 case at Reτ=180. Flow is from bottom to top, and the four walls have been removed to allow better visualization. Note that the simulation captures the near-wall streaks on the four walls, and the interaction of ejection events from top-bottom and side walls at the corner.

Hassan Nagib, Illinois Institute of Technology; Philipp Schlatter and Ricardo Vinuesa, KTH Royal Institute of Technology

ENGINEERING

UNDERSTANDING SECONDARY MOTIONS AND THEIR IMPACT IN MODELING TURBULENT FLOWS

This numerical study is examining the turbulent flow behavior of a fluid moving inside a duct in order to better understand the physical mechanisms that contribute to friction and energy loss in numerous industrial applications. In particular, secondary motions of Prandtl's second kind have been observed by computing turbulent duct flows at various Reynolds numbers and aspect ratios (ARs), defined as the duct width divided by its height, and then comparing these findings with an ongoing experimental campaign. The data collected during this project will help develop more accurate turbulence models that are needed because current industrial codes still fail to predict 3D flows with secondary motions accurately.

Using the numerical code Nek5000, the team has performed detailed studies of the 3D effects responsible for the skin friction dependence on the AR (i.e., secondary vortices and side-wall boundary layers), as well as how these effects evolve with Reynolds numbers and ARs. Presently, two cases have been finalized: AR=14.4 at Reτ=180, and AR=3 at Reτ=360 with particles. The first case is the widest ever computed with direct numerical simulation (DNS) and directly matches an experiment at the Illinois Institute of Technology, and the second one will be used to further understand the nature of the secondary flow connected with the mechanisms typical of wall-bounded turbulence.

The simulations indicate that the flow changes when it is close to the corner and at the core of the duct. In fact, where one would expect that turbulence would monotonically decrease as a pipe widens, the numerical simulation data show that for intermediate ARs, the turbulence may actually be greater than it is for low ARs. This finding confirms some existing conjectures about duct flow dynamics and sheds light on some new, open issues. Previously, these vortical secondary motions had not been documented in the literature to this degree of detail and are extremely difficult to detect in experiments. Seeing these behaviors emerge in petascale simulations on Mira will be of great benefit to the turbulent flow community.

IMPACT: The results are expected to have a profound impact on our understanding of comparisons between modern and classical canonical experiments, and the largest volume of DNS data in the literature (i.e., channel flows). Insights will enable the development of better simulation models and contribute to a scientific database on turbulent flows.

GIULIA GALLI
University of Chicago
gagalli@uchicago.edu

FRANCOIS GYGI
University of California, Davis
fgygi@ucdavis.edu

INCITE
180 Million Core-Hours

Image: Ab initio *molecular dynamics simulations were employed to study the equilibrium density and compressibility of water and ice; the results highlighted the importance of using hybrid functionals together with dispersion corrections.*

Peter Allen, University of Chicago

MATERIALS SCIENCE

COMPUTATIONAL SPECTROSCOPY OF HETEROGENEOUS INTERFACES

The interfaces between solids, nanoparticles, and liquids play a fundamental role in determining the properties of materials. A better understanding of the microscopic structure of interfaces will enable researchers to predict optimal properties of materials for many applications, including photo-electrochemical water splitting and solar cells. With this project, researchers are tackling the characterization and control of heterogeneous interfaces at the atomic and molecular scale.

While the properties of interfaces are key to materials predictions, it is challenging to include them in *ab initio* models due to the complexity and cost of the associated calculations. The INCITE research team is using the open-source codes Qbox (eslab.ucdavis.edu/software/qbox) and WEST (west-code.org) to study interfaces present in materials of interest for solar energy conversion processes. Their approach involves integrating newly developed algorithms for calculating opto-electronic and vibrational spectra with large-scale *ab initio* molecular dynamics simulations (AIMD).

To validate the theory used in modeling highly complex interfaces, the researchers first carried out AIMD simulations of water and ice with gradient-corrected (PBE) and hybrid (PBE0) functionals, and they are now investigating water interfaced with oxides and functionalized semiconductor surfaces. With Mira, they were able to perform several different simulations of ice and water samples with PBE0, whose efficiency was greatly improved with new algorithms (in particular, the recursive subspace bisection method). They showed that the PBE0 functional, together with an approximate inclusion of dispersion interactions, yielded results in better agreement with experiment for the density and compressibility of water than widely used semi-local functionals. They also demonstrated the feasibility of many body-perturbation theory calculations at the GW level, in conjunction with AIMD for complex aqueous interfaces.

In addition, the team has been improving the performance of their codes on Mira. With assistance from ALCF staff, they achieved a 10 percent speedup of the WEST code by reducing the time needed for I/O. With regard to Qbox, the researchers implemented new techniques for calculating sum frequency generation spectra, which allows them to investigate the spectroscopic signatures of liquid water at oxide and semiconductor interfaces.

IMPACT: This project is providing a better understanding of the properties of heterogeneous interfaces, which will help inform the design of optimal materials for solar energy applications. In addition, the researchers are establishing a strategy to enable the comparison of *ab initio* data with experiments carried out at light sources, such as Argonne's Advanced Photon Source. Ultimately, the results could lead to new analysis tools for spectroscopic data that can be used by theorists and experimentalists alike.

JUSTIN WOZNIAK

Argonne National Laboratory
jmwozniak@anl.gov

Director's Discretionary
6.9 Million Core-Hours

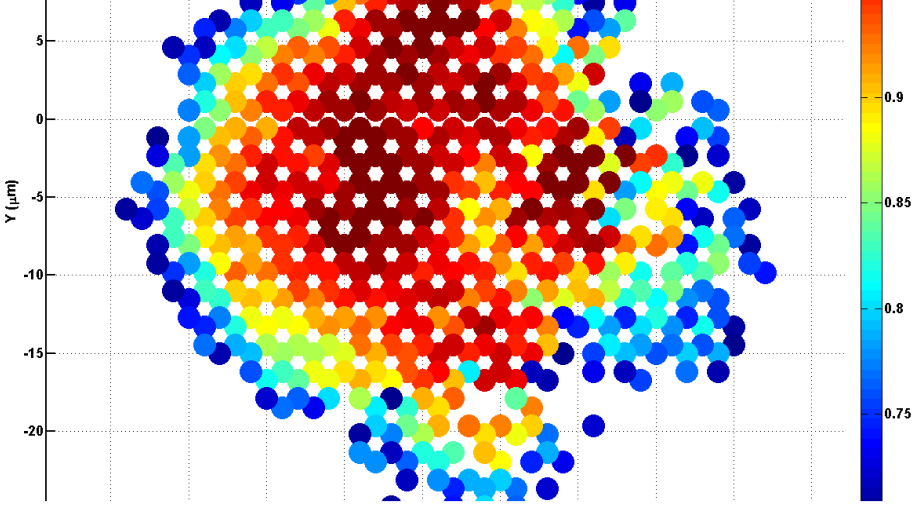

Image: High-energy x-ray diffraction microscopy image of an ~50 micron diameter gold wire (courtesy of B. Suter, Carnegie Mellon University) showing confidence index for each ~1 micron voxel. The data is from the 1-ID beamline of the APS. Reconstructions conducted using APS and ALCF resources were led by Argonne researchers Hemant Sharma and Justin Wozniak.

Jonathan Almer and Hemant Sharma, Argonne National Laboratory

MATERIALS SCIENCE

INTEGRATING SIMULATION AND OBSERVATION: DISCOVERY ENGINES FOR BIG DATA

Argonne's Advanced Photon Source (APS) generates the "brightest" high-energy synchrotron x-rays in the Western Hemisphere, enabling scientists to peer inside atomic-level structures and design next-generation materials. This project aims to accelerate discoveries at the APS by leveraging ALCF resources to address gaps in computation, analysis, simulation, and data management for beamline experiments. The team is employing the cyberinfrastructure that is being developed within the project, both for managing data and developing interactive analysis tools and algorithms capable of handling big data.

The computational developments will be performed in the context of analyzing APS data collected using three different techniques: single crystal diffuse scattering, time-resolved pump-probe techniques, and diffraction microscopy. Data from the beamlines will be streamed to the ALCF, where they will be analyzed using computational techniques ranging from advanced statistical analysis of data correlations to comprehensive simulations using molecular dynamics and *ab initio* modeling.

Four software packages designed for Mira have been developed over the last year: CCTW (crystal coordinate transformation workflow); a NeXus dataset processing workflow; the DISCUS-based genetic optimization algorithm DIFFEV for modelling x-ray and neutron scattering results; and NF-HEDM, a workflow for analyzing near-field, high-energy diffraction microscopy data. NF-HEDM can be used with FF (far field) HEDM to obtain a higher-fidelity study of a given sample. This multimodal approach can determine local strains that lead to cracks, enable microstructure-sensitive modeling, and validate additive manufacturing processes. The overarching programming model is Swift, a parallel scripting language that transforms HEDM analysis software into a scalable application capable of utilizing Mira efficiently.

These codes are now entering production use for live APS data collection runs, in which data are streamed from APS beamline experiments to the ALCF to perform analysis, simulation-based inverse modeling, and data manipulation/management. When in full production, millions of core-hours per week could be used in support of APS experiments. An initial run used Mira to help detect and correct a problem early on, so that a two-day experiment produced publishable results.

IMPACT: This project seeks to demonstrate (1) the practical value of high-performance computing in real-world applications; (2) the ability of materials scientists to extract insights from the data while not needing to possess expertise in high-performance computing; and (3) the value of deploying this framework on the remaining beamlines at the APS.

AIDAN P. THOMPSON
Sandia National Laboratories
athomps@sandia.gov

ALCC
80 Million Core-Hours

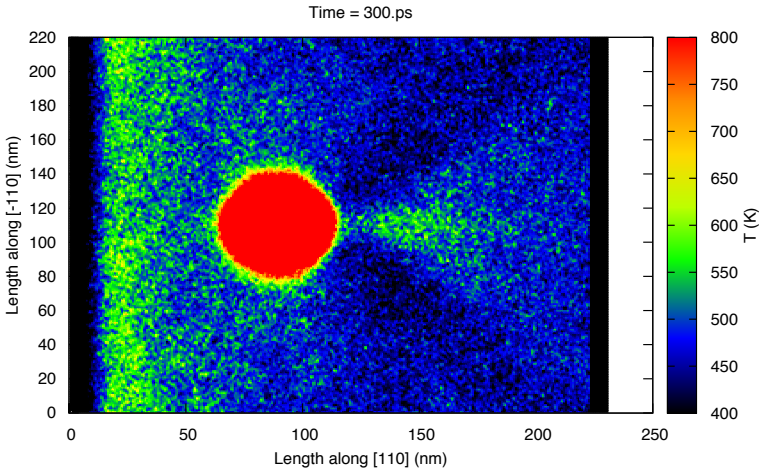

Image: Temperature plot of the PETN crystal at 300 ps after shock. A triangular secondary hot zone is in front of (i.e., to the right of) the circular-shaped primary hot spot.

Tzu-Ray Shan, Sandia National Laboratories

MATERIALS SCIENCE

NANOSTRUCTURE-ENHANCED CHEMICAL REACTIVITY AND DETONATION IN ENERGETIC MATERIALS

Understanding the physics and chemistry of detonation in energetic materials is important across many sectors, including national security applications, civil engineering (mining and excavation), and in specialized applications such as emergency passenger restraint systems and rocket propulsion. Despite recent advances in understanding the initiation of energetic materials, researchers are not yet capable of accurately predicting the shock-to-detonation transition in real explosives. However, they can—using state-of-the-art quantum mechanical dynamics methods—now predict the shock properties of perfect single crystals of energetic materials.

Challenges, nevertheless, remain because material defects, such as entrained solvent, dislocations, porosity, and grain boundaries, are believed to play a key role in the onset of shock-induced chemical reactions and the ignition of hot spots. Furthermore, the spatial scales of these phenomena are too small to resolve in continuum models but too large for molecular models. The research team is using the reactive force field (ReaxFF) method implemented in the LAMMPS parallel molecular dynamics code. ReaxFF bridges the molecular and continuum scales while also retaining the accuracy of quantum methods, enabling the team to run micron-scale atomistic simulations of void collapse and hot spot formation.

Results show spherical and cylindrical voids have similar effects on hot spot formation and initiation in hexanitrostilbene (HNS) crystals under shock conditions, although averaged hot spot temperature is somewhat lower for cylindrical voids. A shocked HNS expands after a relief wave passes through, and temperature and pressure drop significantly. From 0–8 ps, the temperature still increases owing to chemical reactions inside the hot spot; however, following the relief wave, the temperature decreases, and pressure decreases monotonically.

When simulating a shock wave passing through an 8.5-million-atom pentaerythritol tetranitrate (PETN) crystal containing a 20-nm cylindrical void, the team observed the formation of a secondary hot zone located in front of the primary hot spot. Its formation may be attributable to the secondary shock wave generated when upstream void fragments collide with the downstream void surface.

IMPACT: This project supports the investigation of nanostructure-enhanced chemical reactivity and detonation in energetic materials and the study of shock energy localization at materials heterogeneities. The outcome of this work will fill a critical gap in energetic materials modeling, enabling predictive models of performance, reliability, and failure.

RAJIV KALIA

University of Southern California
rkalia@usc.edu

INCITE
180 Million Core-Hours

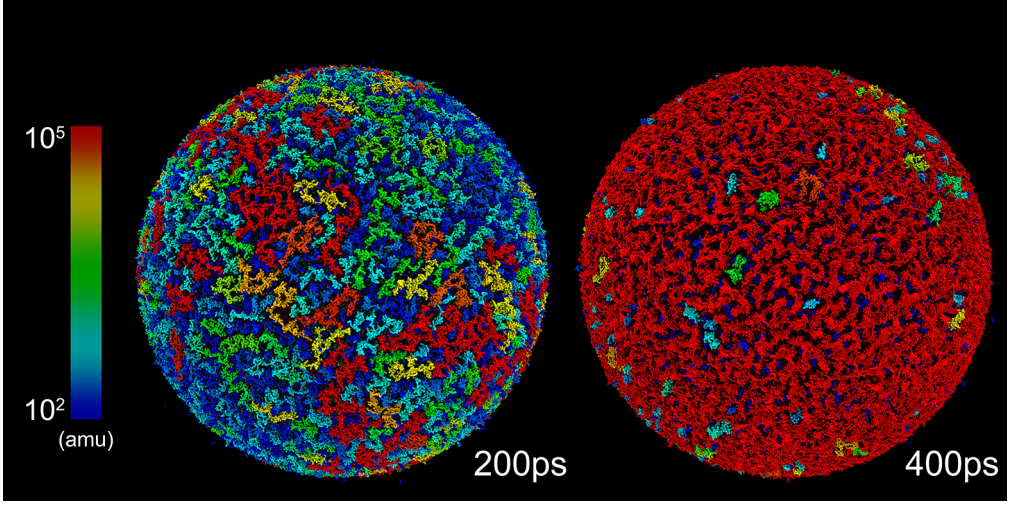

Image: Graphene-like carbon clusters produced on the surface of an oxidizing SiC nanoparticle, exhibiting a percolation transition. The color represents the cluster mass in atomic mass unit (amu).

Joseph A. Insley, Argonne National Laboratory; Ken-ichi Nomura, University of Southern California

MATERIALS SCIENCE

PETASCALE SIMULATIONS OF SELF-HEALING NANOMATERIALS

Self-healing nanomaterials are capable of sensing and repairing damage, such as cracks, in devices operating in harsh conditions. With the ability to enhance reliability and reduce maintenance costs, these materials show great promise for use in energy technologies, such as high-temperature turbines. Researchers from the University of Southern California (USC) are performing massive simulations on Mira to help advance the understanding and viability of self-healing nanomaterials.

The goal of the project is to determine the atomistic mechanisms underlying these remarkable self-healing capabilities. To do so, the USC research team is using ALCF supercomputers to carry out quantum molecular dynamics, reactive molecular dynamics (RMD), and mesoscale simulations of two types of systems: anticorrosion coatings for metals and ceramic nanocomposites.

RMD simulations based on the reactive force field (ReaxFF) approach have emerged as a powerful research tool for describing chemical reactions, but the method's iterative determination of atomic charges at every time step presents a major computational bottleneck. The USC team eliminated this speed-limiting charge iteration by developing a new extended-Lagrangian RMD (XRMD) code that drastically improved energy conservation and time to solution. Working closely with ALCF staff, the researchers optimized XRMD with several code transformations, including the introduction of OpenMP to parallelize the computationally intensive bond-order, potential-energy, and force calculations within each MPI processing element. With the enhancements in place, they achieved a 2.4x speedup over XRMD's original baseline performance. In benchmark simulations of a 67.6 billion-atom system, the code demonstrated a weak-scaling parallel efficiency of 0.977 while running on all of Mira's 786,432 cores.

In their studies of ceramic nanocomposites, the team is performing RMD simulations to examine the self-healing of cracks in an alumina matrix embedded with silicon carbide nanoparticles. At high temperatures, the oxidation of silicon carbide produces silica, which can flow into cracks to heal them. Using the full Mira system, the researchers performed the largest-ever RMD science simulation (106 million atoms) to study this initial oxidation process. They observed the unexpected formation of large graphene-like flakes despite the harsh oxidation conditions. The large system size has, for the first time, enabled statistical analysis of the geometry of the nanocarbon product.

IMPACT: This project will provide a fundamental understanding of self-healing nanomaterials, helping to accelerate the discovery of new materials for use in extreme conditions. Ultimately, this work could help enhance the reliability and reduce the cost of components for many energy applications, including wind and solar energy technologies, and high-temperature turbines.

DAVID CEPERLEY

University of Illinois at Urbana-Champaign
ceperley@illinois.edu

INCITE
185 Million Core-Hours
(ALCF: 100M; OLCF: 85M)

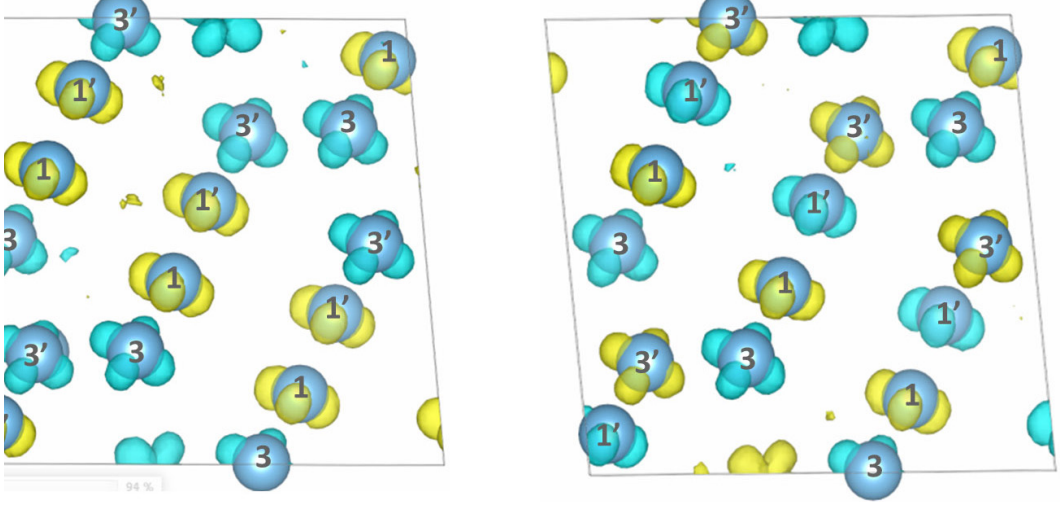

Image: Diffusion Monte Carlo spin densities for antiferromagnetic states AF1 (left) and AF3 (right). Shown only atoms with a non-null spin. The picture shows the localization of the spins in d-orbital. Blue is a spin down, while yellow is a spin up.

Anouar Benali, Argonne National Laboratory

MATERIALS SCIENCE

QMC SIMULATIONS DATABASE FOR PREDICTIVE MODELING AND THEORY

Due to its numerical expense, quantum Monte Carlo (QMC) methods were once limited to model systems of small atoms or molecules. But the arrival of leadership-class computers has led to the routine application of QMC methods in a wide variety of systems, allowing for rigorous calculations of more complicated materials. This INCITE project involves QMC studies of transition metal oxides and lanthanides.

Project researchers are using QMC to model electronic structure systems, calculating properties that emerge from these utilizing the QMCPACK program ported on Mira. Research thrusts include heterogeneous catalysis of transition metal nanoparticles, phase transitions, properties of materials under pressure, and strongly correlated materials.

Earlier investigations into platinum solids, nanoclusters, and surfaces produced a highly accurate model of the face-centered cubic equation of state for platinum. Calculations have also determined the energy needed to form the (111) surface of platinum, a surface important in the catalytic process of many chemical reactions. These results demonstrate the significant potential of QMC simulations to produce highly accurate values for transition metals and surface energetics, an achievement that still eludes most density functional-based approaches. The team has begun calculating the energy of the Pt (100) surface, another important facet for many surface and catalytic applications.

Research into catalysis of transition metal oxides has focused on titanium oxides, which are of interest for a number of applications, ranging from resistive random access memories to paint dyes. What makes them and other transition metal oxides versatile, in part, are the many different oxidation states possible.

When used as a fast switch resistance, the oxygen atoms in titanium dioxide (TiO_2) start to diffuse, creating regions of TiO_2, Ti_2O_3, and Ti_4O_7 in the semiconductor. To perfectly understand the process, researchers are studying each one of these phases starting with Ti_4O_7, for the challenge it offers using mean field theory methods.

The team confirmed that an anti-ferromagnetic state with a long-range spin ordering is the most stable phase in Ti_4O_7. Despite numerous experimental studies, this was the first theoretical study explaining the spin ordering in Ti_4O_7 and its role in the stability of the multiple states in Ti_4O_7.

IMPACT: This project aims to advance the efficiency and global applicability of QMC through the development of tools that make systematic research less time intensive. Its results will provide accurate predictions for energy-related materials and fundamental material properties, as well as establish benchmarks that will provide targets for future developments in related electronic structure approaches.

SUBRAMANIAN SANKARANARAYANAN

Argonne National Laboratory
ssankaranarayanan@anl.gov

INCITE
40 Million Core-Hours

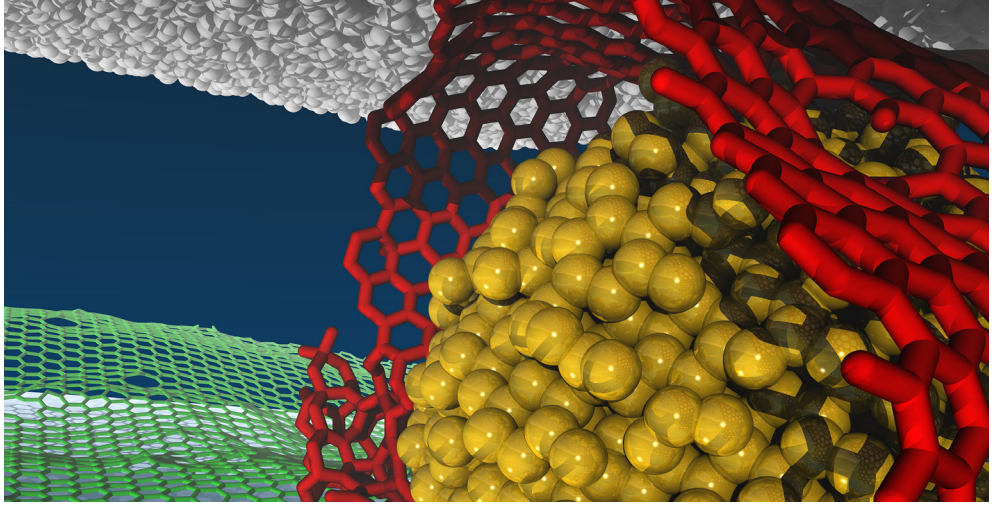

Image: This large-scale simulation depicts a phenomenon called superlubricity, a property in which friction drops to near zero. The simulation reveals that this condition originates at the nanoscale when graphene atoms self-assemble into nanoscrolls that reduce contact area to help eliminate friction.

Sanket Deshmukh, Joseph A. Insley, and Subramanian Sankaranarayanan, Argonne National Laboratory

MATERIALS SCIENCE

REACTIVE MD SIMULATIONS OF ELECTROCHEMICAL OXIDE INTERFACES AT MESOSCALE

Superlubricity—a state in which friction essentially disappears—is a highly desirable property. Considering that nearly one-third of every fuel tank is spent overcoming friction in automobiles, a material that can achieve superlubricity would greatly benefit industry and consumers alike. Argonne scientists used Mira to identify and improve a new mechanism for eliminating friction, which fed into the development of a hybrid material that exhibited superlubricity at the macroscale for the first time.

Prior to the simulation work at the ALCF, a team of Argonne scientists was studying a promising new lubricant material in laboratory tests. The experimental setup, which consisted of small patches of graphene sliding against a ball coated with diamond-like carbon, was resulting in a very low friction coefficient, but the friction levels were fluctuating up and down for no apparent reason.

To shed light on the mysterious behavior, Argonne computational scientists used Mira to replicate the experiments with large-scale molecular dynamics (MD) simulations aimed at understanding the underlying mechanisms of superlubricity at an atomistic level. ALCF staff helped to enable the computationally demanding simulations by optimizing the LAMMPS code and its reactive force field (ReaxFF) module. The optimization work included adding OpenMP threading, replacing MPI point-to-point communication with MPI collectives in key algorithms, and leveraging MPI I/O. Altogether, these enhancements enabled the code to perform twice as fast as before.

The simulations revealed a completely new mechanism for superlubricity. When the graphene and diamond-like carbon slid against each other, the graphene would roll up to form hollow cylindrical nanoscrolls that helped to practically eliminate friction. The researchers then tried incorporating nanodiamond particles into the simulations to help stabilize the nanoscrolls. The modification proved successful as the graphene patches spontaneously rolled around the nanodiamonds, which held the scrolls in place and resulted in sustained superlubricity. The simulation results informed the design of a hybrid material that demonstrated superlubricity in laboratory tests as well. While the material does not work in the presence of water, the discovery of graphene nanoscrolls shows great potential for applications in dry environments.

IMPACT: This work led to the development of a material that exhibited superlubricity at the macroscale for the first time. The material could potentially be used for applications in dry environments, such as computer hard drives, wind turbine gears, and mechanical rotating seals for microelectromechanical and nanoelectromechanical systems. In addition, the knowledge gained from this study is expected to spur future efforts to develop materials capable of superlubricity for a wide range of mechanical applications.

LUCAS K. WAGNER

University of Illinois at Urbana-Champaign
lkwagner@illinois.edu

INCITE
106 Million Core-Hours

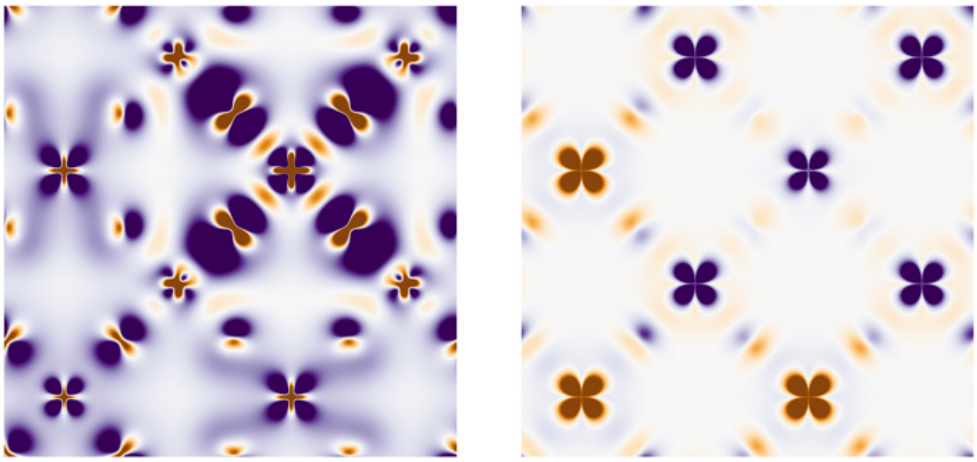

Image: Charge (left) and spin (right) density of a spin polaron in cuprates. This object could result in superconductivity at high temperatures.

Lucas K. Wagner, University of Illinois at Urbana-Champaign

MATERIALS SCIENCE

SIMULATION OF CORRELATED ELECTRONS FOR SUPERCONDUCTING MATERIALS

High-temperature superconductors have the potential to transform many applications from computers to renewable energy technologies, but the mechanism responsible for their unique properties remains a mystery. Researchers from the University of Illinois at Urbana-Champaign (UIUC) are using the computing power of Mira to uncover new microscopic physics in materials that exhibit high-temperature superconductivity.

The study of emergent physics, that is, how unique properties emerge from fundamental physics, is a major frontier in condensed matter physics. In particular, electronic systems are a fundamental part of modern technology and offer great promise for enabling the development of new materials. The understanding of quantum mechanics has advanced to the point that many new electronic systems can be designed for targeted objectives.

However, this understanding is primarily based on a critical assumption—that the interactions between electrons are small. When this assumption is not valid, a challenge and an opportunity arise. The challenge is that the equations used to calculate the electronic structure become very hard to solve. The opportunity is that new and exciting effects, like high-temperature superconductivity, can occur. For this INCITE project, the UIUC research team is using Mira to shed light on high-temperature superconductors by performing computationally intensive calculations to solve the Schrödinger equation with unprecedented accuracy.

Thus far, the researchers used their QWALK code to perform diffusion Monte Carlo calculations of the magnetic state of iron selenide, a known high-temperature superconductor, at varying levels of pressure. Their simulation results added to the existing body of evidence supporting the notion that high-temperature superconductivity is magnetic in origin.

The team is also investigating copper-based superconductors, known as cuprates. In one study, they performed high-fidelity simulations of charge carriers in cuprate materials using quantum Monte Carlo techniques applied to the first-principles Hamiltonian. The simulations led to the discovery of a new feature similar to a spin polaron, in which charge is localized through a strong interaction with the spin. The results provide an accurate first principles-derived paradigm that could have important implications for understanding superconductivity in cuprates.

IMPACT: The goal of this project is to completely characterize, for the first time, the electronic structure of high-temperature superconductors. Findings from this effort will also help improve the fundamental understanding of electronic matter, which will aid scientists in the development of novel materials with new and exotic properties.

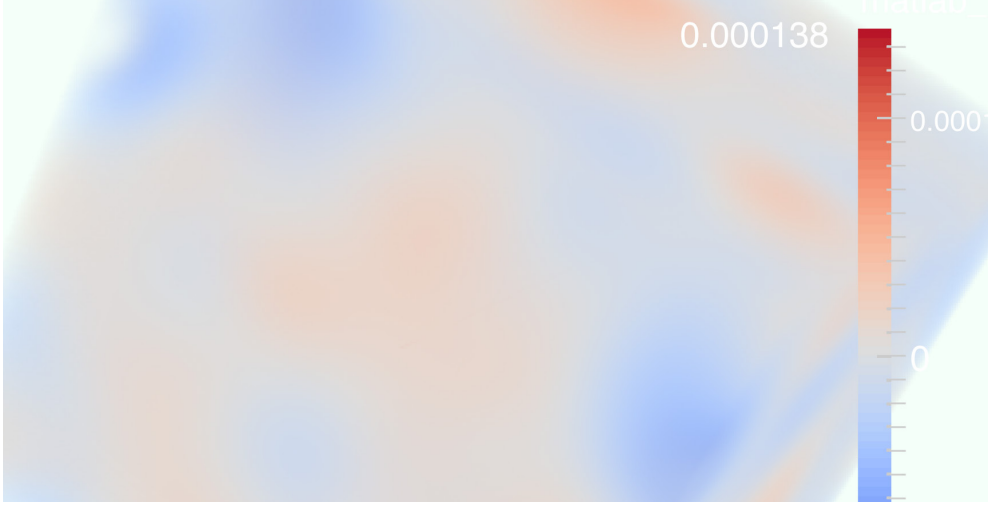

JULIUS KUTI

University of California, San Diego
jkuti@ucsd.edu

ALCC
88.7 Million Core-Hours

Image: The vacuum Higgs condensate from LatHC simulations on Mira.

Julius Kuti, University of California, San Diego; Lattice Higgs Collaboration

PHYSICS

COMPOSITE HIGGS THEORY BEYOND THE STANDARD MODEL AND THE 14 TEV UPGRADE OF THE LARGE HADRON COLLIDER

After its 14-TeV energy upgrade in 2015, the Large Hadron Collider (LHC) at CERN began an era of investigations searching for new physics beyond the Standard Model (BSM). The Lattice Higgs Collaboration (LatHC) is working on large-scale lattice simulations of the BSM paradigm to determine the validity and structure of a composite Higgs mechanism, which could reveal the substructure of the Higgs particle itself.

The international LatHC is testing a candidate theory for a minimal composite Higgs mechanism with a light scalar state (Higgs impostor) emerging at the electroweak scale. This effort potentially could replace the elementary Standard Model Higgs and reveal the nature of new constituents responsible for its origin.

In the simplest BSM theory of the composite Higgs, the vacuum is filled with topological structures leading to the formation of the vacuum Higgs condensate, a universal field created from the condensation of new sub-constituents. This Higgs condensate has very different properties from the ordinary quark condensate that is described by quantum chromodynamics (QCD) of the Standard Model, whose excitations create ordinary matter, like mesons and baryons. The lowest excitations of the composite Higgs condensate create a low-mass scalar with vacuum quantum numbers well below the TeV mass scale. This condensate is formed from new basic constituents of fermion pairs well below the 10^{-13} - 10^{-15} cm scale of quark sub-structures in QCD.

The team developed highly optimized code and performed very large-scale lattice gauge simulations of the new composite Higgs theory to evaluate these new LHC predictions. They identified a low-mass scalar particle well below the TeV energy scale that could act as the Higgs impostor, if the theory becomes successful.

In addition, they predicted a resonance spectrum from the excitations of the Higgs vacuum in the 2 TeV energy range with a rho-like resonance state, whose quantum numbers are consistent with a di-boson anomaly, recently observed at the LHC and consistent with LatHC predictions. This could become the first clue to the composite nature of the Higgs particle.

IMPACT: Simulations of strongly coupled gauge theories, like the one proposed here, will aid experimental searches for compositeness and contribute, in a broader sense, to an understanding of conformal and near-conformal quantum field theories which are among the most important BSM paradigms. If successful, a deeper understanding of the origin of mass in the universe will emerge.

NICKOLAY GNEDIN
Fermilab
gnedin@fnal.gov

INCITE
74 Million Core-Hours

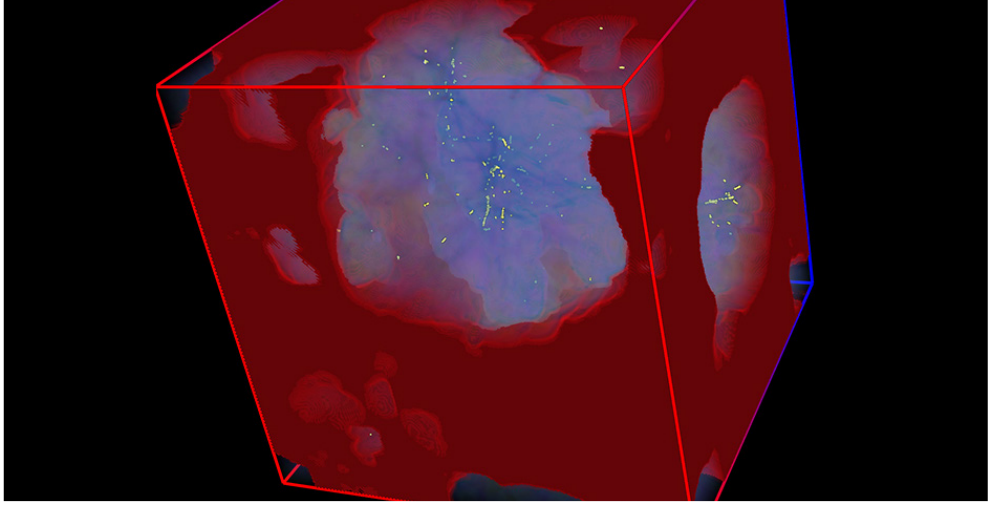

Image: Numerical model of cosmic reionization. Brown non-transparent fog renders neutral gas; glowing blue is dense ionized gas (which becomes completely transparent when it is not dense); and yellow dots are galaxies.

Nickolay Gnedin, Fermilab

PHYSICS

COSMIC REIONIZATION ON COMPUTERS

Cosmic reionization—one of the most promising research areas in astrophysics according to the National Research Council—was one of the two major phase changes in the universe in which neutral cosmic gas was ionized by high-energy radiation from early galaxies. Gaining a better understanding of this process will shed light on many aspects of modern cosmology, including cosmic microwave background observations and the physical state of intergalactic gas in front of high-redshift quasars.

Because the observational constraints on reionization are limited (e.g., distant quasars are very rare), theoretical modeling and numerical simulations play a critical role in reionization studies. Historically, simulations of reionization were limited either to small spatial volumes with detailed physical modeling or to large volumes in which physics was treated highly approximately. Petascale supercomputers like Mira are enabling potential breakthroughs in the field by helping to bridge the gap between small- and large-box simulations.

The Cosmic Reionization On Computers (CROC) project takes advantage of the steady advance in computing power to carry out simulations that model all of the relevant physics, from radiative transfer to gas dynamics and star formation, with sufficient detail. These simulations are being performed in volumes of more than 100 co-moving megaparsecs (Mpc) on a side and with spatial resolution approaching 100 parsecs to cover the full range of spatial, temporal, and mass scales needed to study reionization. Simulations in smaller boxes would not be able to correctly capture the process of overlapping ionized bubbles.

Using the Adaptive Refinement Tree (ART) code, the CROC team has started both of the simulations planned for this year: one at 58 Mpc at high resolution and another at 116 Mpc at medium resolution. The smaller simulation ties together the small-box simulations run previously (i.e., 58 Mpc at medium resolution) with the larger simulation being run in this project. Results obtained so far show a remarkable agreement with observational constraints, and have already overturned a long-standing paradigm about the formation of the smallest galaxies.

IMPACT: Results from this project will be instrumental in achieving a wide range of scientific goals, from understanding the detailed physics of the reionization process, to providing theoretical predictions of the spatial distribution and internal properties of early galaxies for future observations with the James Webb Space Telescope. This work will also help researchers to explore other recent observational campaigns, including the residuals of reionization in the absorption spectra of distant quasars that were discovered by the Sloan Digital Sky Survey.

CHOONG-SEOCK CHANG

Princeton Plasma Physics Laboratory
cschang@pppl.gov

INCITE
270 Million Core-Hours
(ALCF: 100M; OLCF: 170M)

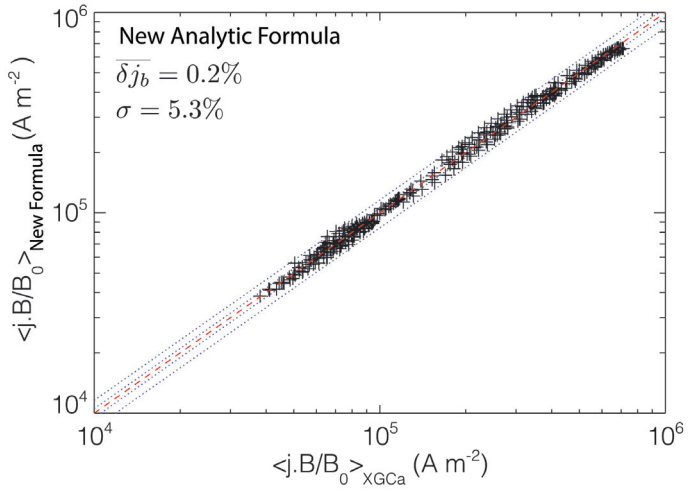

Image: The new analytical current formula (vertical axis) created from numerous large-scale XGCa simulations, using one-third of Mira, agrees excellently with the simulation results (horizontal axis).

Robert Hager, Princeton Plasma Physics Laboratory

PHYSICS

HIGH-FIDELITY SIMULATION OF TOKAMAK EDGE PLASMA TRANSPORT

The success of magnetically confined fusion energy systems, including ITER, is dependent upon how well its plasma edge is confined. A well-confined edge plasma forms a steep pedestal and puts the fusion plasma into a high-confinement mode, which allows an efficient fusion burn in the plasma core. This project will provide first-principles understanding of the underlying plasma physics that govern edge plasma confinement in fusion reactors.

Electric current plays an important role in plasma edge physics. Strong electrical current is critical for better confinement of the edge plasma, hence higher fusion efficiency. Stability of edge plasma is also sensitive to edge electrical current. Thus, accurate predictions for edge electrical current are important to the success of ITER. Due to the overlapping, multiscale, and nonlinear nature of electrical current in steep-edge plasma, extreme-scale computing is needed.

A team, led by researchers from the Princeton Plasma Physics Laboratory, is using the extreme-scale gyrokinetic particle-in-cell family of XGC codes run on Mira. The combination of Mira and XGCa has, for the first time, accurately evaluated a strong self-generated electrical current, called “bootstrap current”, in a steep-edge pedestal plasma.

In a donut-shaped, or torus, tokamak magnetic-field geometry, the magnetic field is stronger at the inner part of the torus than at the outer part. Particles traveling at slower speeds cannot penetrate the higher magnetic field region and are trapped at the exterior of the torus. These are called magnetically “trapped” particles. Particles passing through the high magnetic field region are called “passing” particles.

While previous studies suggest that the bootstrap current exists mostly in the passing electron phase-space, the present discovery shows that, in the edge region where the passing particle fraction is small, most of the bootstrap current exists in the trapped phase-space. Numerous simulations for different geometric and plasma conditions led to a new unified analytical formula that can describe the edge current, as well as the core current.

IMPACT: This new analytical formula, describing both edge and core currents, is expected to be widely used by international fusion researchers to study edge plasma physics. This will enable the fusion research community to better understand and predict edge plasma behavior, and to ensure highly efficient tokamak performance and more economical ITER operation.

ROBERT A. DISTASIO JR.

Cornell University
distasio@cornell.edu

ALCC
175 Million Core-Hours

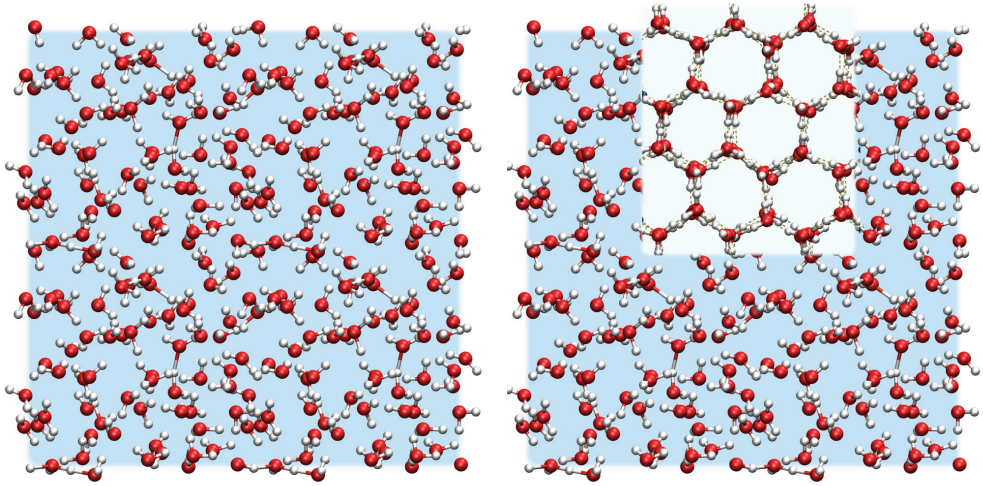

Image: Understanding the microscopic structure and anomalous density properties of water has been made possible through theoretical and algorithmic advances in conjunction with the high-performance computing capabilities at the ALCF.

Robert A. DiStasio Jr., Cornell University; Roberto Car, Hsin-Yu Ko, and Biswajit Santra, Princeton University

PHYSICS

ION SOLVATION, CATALYTIC INTERFACES, AND EXTREME AQUEOUS ENVIRONMENTS: AN AB INITIO STUDY OF LIQUID WATER

Water is arguably the most important molecule on Earth. While the structure of a single water molecule is well-known, liquid water has an intricate and disordered microscopic structure that has been difficult to study. This project addresses key challenges of *ab initio* molecular dynamics (AIMD) simulations of liquid water and ionic solutions, which could significantly impact the development of clean and sustainable energy sources.

An ongoing collaboration between ALCF, Princeton, and Cornell, this project continues to perform highly accurate benchmark atomistic simulations of liquid water and aqueous ionic solutions through a combination of algorithmic advances and the efficient utilization of Mira.

ALCF staff has provided assistance in improving the linear scaling algorithm utilized for the calculations in this project. The code can now be applied to many types of periodic simulation cells, making it invaluable in the study of complex condensed-phase systems.

Researchers have completed several large-scale path integral simulations of liquid water, crystalline ice, pyridine (and pyridine-like molecular crystals), as well as the fundamental hydronium (H_3O^+) and hydroxide (OH^-) aqueous ionic solutions. To do so, they utilized Quantum ESPRESSO, an open-source electronic structure theory code that can perform AIMD simulations utilizing density functional theory.

Results include a highly accurate characterization of the microscopic structures and anomalous density properties of liquid water and crystalline ice, which have posed substantial challenges for theory. In addition, proton transfer rates and diffusivities of the aqueous H_3O^+ and OH^- ions were studied in great detail and found to be in excellent agreement with experiments. The team also determined how van der Waals interactions and nuclear quantum effects affect the structure and equilibrium densities of pyridine and a series of pyridine-like molecular crystals, an advance that will aid in the theoretical description of polymorphism in molecular crystals of biological, pharmaceutical, and technological relevance.

IMPACT: These simulations continue to provide detailed knowledge of the microscopic structure and equilibrium properties of liquid water and aqueous ionic solutions with unprecedented accuracy, addressing important renewable energy research questions underlying the fundamental understanding and rational design of aqueous ion batteries. Results will be stored in a publicly available structural database as an invaluable resource for future developments across many scientific disciplines.

FREDERICO FIUZA

SLAC National Accelerator Laboratory
fiuza@slac.stanford.edu

INCITE
110 Million Core-Hours

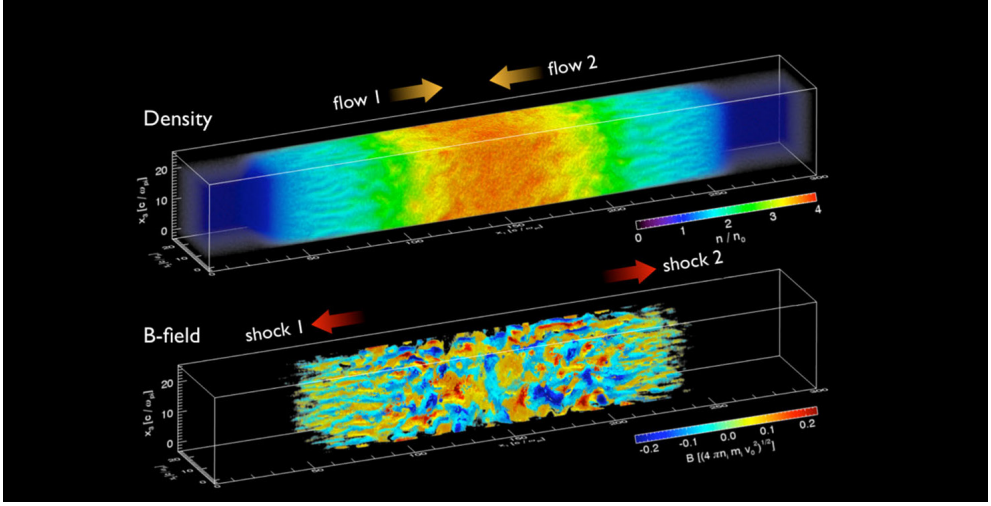

Image: 3D OSIRIS simulation of the formation of a Weibel-mediated collisionless shock for expected NIF conditions. Two laser-driven counter-streaming plasma flows interact in the central region, leading to the development of a Weibel-mediated shock. The strong B-fields thermalize and slow down the initial flows, leading to a density compression consistent with hydrodynamic jump conditions.

Frederico Fiuza, SLAC National Accelerator Laboratory

PHYSICS

PARTICLE ACCELERATION IN SHOCKS: FROM ASTROPHYSICS TO LABORATORY IN SILICO

Particle acceleration in astrophysical shocks is believed to be one of the most important sources of energetic cosmic rays. However, the details of how these shocks are formed and how they efficiently accelerate particles are not fully understood. Project researchers are relying on *ab initio* particle-in-cell (PIC) simulations to study the microphysics of shock formation and particle acceleration.

Leveraging the state-of-the-art relativistic massively parallel PIC code OSIRIS, and its recently incorporated hybrid model, researchers aim to establish a better understanding of the conditions necessary to form collisionless shocks and the dominant particle acceleration mechanisms in shocks.

Using an unprecedented number of particles-per-cell on Mira, researchers conducted large-scale simulations of particle acceleration in non-relativistic shocks generated in initially weakly magnetized plasmas. While acceleration in this regime had not yet been observed, these simulations show a clear and efficient particle injection and acceleration in non-relativistic high-Mach number ($M_A > 100$) shocks due to plasma instabilities.

The interaction of the particles reflected by the shock with the ambient medium gives rise to a Weibel instability, whereby small-scale magnetic-field perturbations are exponentially amplified, reaching energies comparable to that of the plasma flows. These strong magnetic fields can then deflect the energetic particles, allowing them to cross the shock front multiple times and gain energy in a process known as Fermi acceleration. Both electrons and ions are shown to be efficiently accelerated. This result has important consequences for the understanding of particle injection in high-M_A shocks.

The experimental investigation of plasma instabilities, such as the Weibel instability, is relevant for the formation of collisionless shocks in counter-streaming plasmas now underway at several high-energy laser facilities. The team's large-scale 3D simulations of current experimental conditions accurately reproduced experimental measurements, allowing for the first clear demonstration of large magnetic-field generation through the ion Weibel instability in collisionless plasmas. The simulations helped identify the conditions required for the generation of Weibel-mediated shocks in the laboratory and successfully design the first National Ignition Facility (NIF) experiments to study these shocks.

IMPACT: This work is allowing transformative results in the understanding of particle acceleration in shocks and on the ability to study shock waves in the laboratory. The insight gained from these first-principles simulations may prove relevant not only to better understand the physics of cosmic rays, but also to design more efficient terrestrial accelerators for a variety of applications.

SEAN COUCH

Michigan State University
couch@pa.msu.edu

INCITE
50 Million Core-Hours

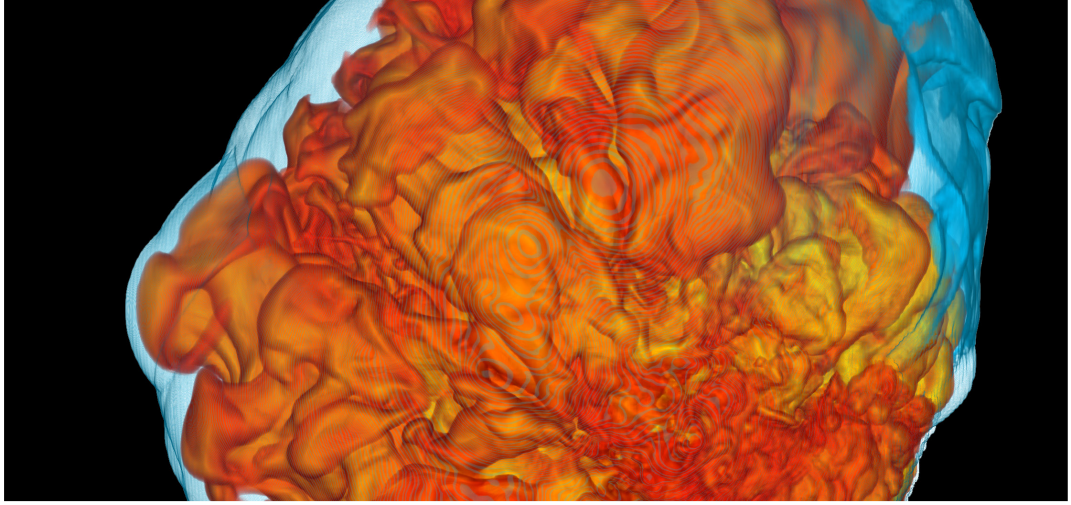

Image: Volume rendering of the specific entropy from a 3D magnetorotational core-collapse supernova simulation, carried out as part of this INCITE project. The stalled supernova shock is visible in blue, while orange and red colors show the turbulent convection below the shock.

S.M. Couch, Michigan State University

PHYSICS

PETASCALE SIMULATION OF MAGNETOROTATIONAL CORE-COLLAPSE SUPERNOVAE

Core-collapse supernovae (CCSNe) are responsible for the production of many elements in the universe. While their importance in galactic chemical evolution cannot be underestimated, the mechanism that reverses stellar core collapse and drives supernova explosions is not fully understood. Researchers will conduct a comprehensive study of the impact of rotation and magnetic fields on CCSNe using 3D magnetohydrodynamics (MHD) simulations of the collapse of rotating, magnetic stellar cores.

CCSNe are the luminous explosions that herald the death of massive stars and serve as the source of origin for cosmic anomalies, such as neutron stars, pulsars, stellar-mass black holes, and, possibly, certain types of gamma-ray bursts. But the mechanism that reverses stellar core collapse and drives supernova explosions is still unclear, making this one of the most important challenges of computational astrophysics.

Project researchers are conducting a comprehensive study of the impact of rotation and magnetic fields on CCSNe using the multi-physics FLASH code on Mira. Code capabilities utilized include adaptive mesh refinement (AMR), an accurate and detailed treatment of nuclear burning, a complex equation of state for matter above nuclear density, and a sophisticated two-moment neutrino transport scheme based on an explicit hyperbolic solver.

Guided by earlier pilot simulations, they ran the first-ever production simulation at extremely high resolution, sufficient to study the MHD turbulence in the CCSNe environment. Using the sophisticated gridding techniques enabled by the code's flexible AMR package, the team achieved a finest resolution element of 125 m, roughly a factor of three better than any 3D CCSNe simulation ever performed. Further analysis will look at growth rates of MHD instabilities and address the impact of rotation and magnetic fields on the favorability for successful supernova explosions.

A large set of 2D CCSNe simulations was conducted and have already yielded significant scientific results. Specifically, it was found that, in purely Newtonian cases, no progenitor successfully explodes, while including an approximate treatment of relativistic gravity results in successful explosions for some progenitors. This may resolve tension amongst recent 2D CCSNe simulation results from different research groups. These simulations also will serve as invaluable guideposts in the design of forthcoming 3D simulations.

IMPACT: Project research will allow for the prediction of the spins, kicks, magnetic field strengths, and alignments of newly formed neutron stars, pulsars, and magnetars, as well as the dependence of these parameters on progenitor conditions. These simulations will be the most physically detailed and accurate CCSNe simulations to include magnetorotational effects ever accomplished, with the potential for uncovering a robust and realistic CCSNe explosion mechanism.

RENE BELLWIED
University of Houston
bellwied@uh.edu

INCITE
150 Million Core-Hours

Image: An illustration of the formation of protons, neutrons, and other hadrons at the QCD crossover transition from Quark-Gluon Plasma to a gas of hadrons, during the evolution of the universe.

Sandor Katz, Eotvos University

PHYSICS

QUARK FLAVORS AND CONSERVED CHARGES AT FINITE DENSITY IN THE QCD PHASE DIAGRAM

The fundamental theory of strong interactions, or quantum chromodynamics (QCD), predicts that ordinary hadronic matter goes through a phase transition to the Quark-Gluon Plasma (QGP), or deconfined phase of matter, in extreme conditions of temperature or density. This project is pursuing a microscopic understanding of the properties of primordial deconfined matter near the transition to ordinary nuclear matter.

An international team of researchers is working toward a precise, continuum-limit determination of high-order fluctuations of light, strange, and charm quark flavors or, equivalently of electric charge, baryon number, and strangeness. Such findings will help determine the microscopic description necessary to understand the underlying physics principles of the QCD transition, currently being realized in heavy ion collisions at the Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory, and at the Large Hadron Collider (LHC) at CERN, near Geneva, Switzerland.

Utilizing a novel lattice QCD approach on Mira, researchers will map out a QCD phase diagram at high temperature and finite density to determine the possibility of a critical point in the phase diagram of matter at extreme conditions that occurred shortly after the Big Bang.

Highly optimized for Mira, the team's Janos code was used to generate zero and finite temperature lattice configurations. The combination is essential for producing large zero-temperature simulations needed for scale setting and comparison for the nonzero temperature runs, which require numerous statistics for accurate results.

As part of a related INCITE project, the researchers previously made fundamental contributions to the field by extracting from first principles the temperature at which hadrons (protons, pions, kaons, etc.) are formed. They have also simulated several observables, which can be compared directly to experimental measurements to determine details of particle formation during the transition. These comparisons can be made today for the first time because of the precision reached in their continuum-extrapolated calculations. This new project extends the previous analysis to larger chemical potentials (i.e., larger densities).

IMPACT: Data from this project may provide guidance on the necessary conditions required to observe a critical phase transition point during the second phase of the RHIC beam energy scan campaign, planned for 2018-2019. Continued research efforts at low densities also will provide benchmarks for LHC, in particular, regarding the heavy quark to hadron transition.

THOMAS LECOMPTE

Argonne National Laboratory
lecompte@anl.gov

ALCC
52 Million Core-Hours
(ALCF: 50M; NERSC: 2M)

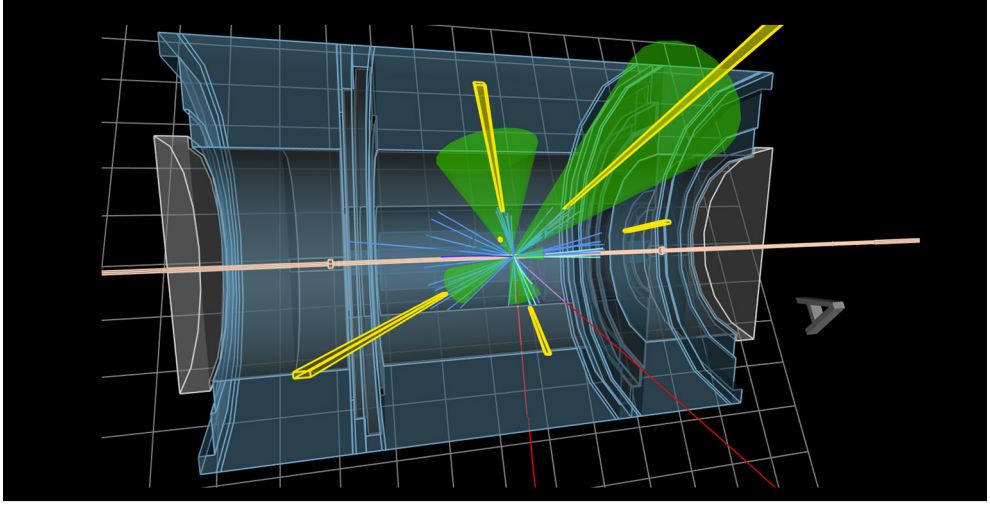

Image: A visualization of a simulated collision event in the ATLAS detector. This simulation, containing a Z boson and five hadronic jets, is an example of an event that is too complex to be simulated in bulk using ordinary PC-based computing grids.

Taylor Childers, Argonne National Laboratory

PHYSICS

SIMULATION OF LARGE HADRON COLLIDER EVENTS USING LEADERSHIP COMPUTING

Simulations are needed to help scientists understand the particle collision experiments taking place at CERN's Large Hadron Collider (LHC), the world's largest and most powerful particle accelerator. The computational requirements for such simulations are extremely large and continue to grow. Project researchers are using Mira to perform large-scale simulations of LHC events with a massively parallel supercomputer for the first time, shedding light on a path forward for interpreting future LHC data.

Scientists from across the world use the LHC to explore the behavior of matter, energy, space, and time at the smallest scales ever probed. Computer simulations are key to understanding the response of the LHC detectors in particle collisions at the facility. Differences between the observations and simulation data can lead to scientific discoveries. However, some LHC events are so complex that it would take weeks to complete the calculations on LHC's computing grid, a system of 100,000 PC-like computers distributed all over the world. In addition, the LHC's computing needs are expected to grow by at least a factor of 10 in the next several years.

With this ALCC project, a team of researchers from Argonne is exploring the use of supercomputers as a tool to enable future discoveries at LHC. In particular, they are focusing on events in the ATLAS experiment that are difficult to simulate with traditional computing resources. To prepare for the simulations, ALCF staff worked closely with Argonne physicists to transform single-threaded LHC simulation codes into massively multi-threaded codes appropriate to run on supercomputers. This effort involved scaling up ALPGEN, a Monte Carlo-based application that generates events in hadronic collisions. By improving I/O and reducing the memory usage, ALPGEN was able to run on the full Mira system and was sped up by a factor of 23.

The code optimization work has enabled the team to routinely simulate millions of LHC collision events in parallel. The improved throughput allowed them to complete two years' worth of ALPGEN simulations in a matter of weeks, freeing the LHC computing grid to run other jobs. Throughout the course of the ALCC project, the team's simulations equated to about 6 percent of the annual computing done by the ATLAS experiment. Ultimately, this effort is helping to accelerate the science that depends on these simulations.

IMPACT: With this project, Argonne researchers are showing that supercomputers like Mira can help drive future discoveries at the LHC. Running particle collision simulations on DOE leadership computing resources provides three key benefits to LHC experiments: increasing the amount of simulated data that can be produced; simulating more complex and realistic events than are currently possible; and helping to evolve LHC's substantial code base for current and future supercomputing platforms.

BRIAN WIRTH

University of Tennessee, Knoxville
bdwirth@utk.edu

ALCC
116 Million Core-Hours
(ALCF: 80; OLCF: 36M)

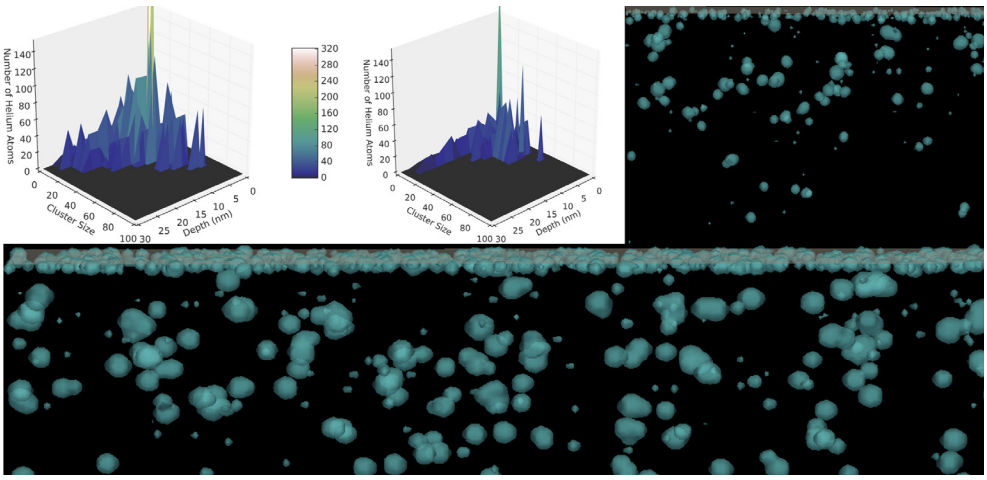

Image: Helium depth and spatial distributions under a W(100)[19] m^{-2} for two different flux values, 3.3 x 10^{25} $m^{-2}s^{-1}$ (left/bottom) and 1.3 x 10^{26} $m^{-2}s^{-1}$ (right/top).

Karl D. Hammond, University of Missouri

PHYSICS

UNDERSTANDING HELIUM-HYDROGEN PLASMA MEDIATED TUNGSTEN SURFACE RESPONSE TO PREDICT FUSION PLASMA-FACING COMPONENT PERFORMANCE IN ITER

Providing energy from nuclear fusion reactions is extremely challenging. One possible route is magnetically confined plasma reactors. This project studies the interaction of helium, the product of fusion, with tungsten surfaces, the proposed plasma-facing component in the ITER divertor. The surface of such plasma-facing materials is damaged by the high-energy ions as they form bubbles and blisters beneath the surface.

Tokamak divertors are responsible for the removal of fusion reaction by-products in magnetic-confinement nuclear fusion reactors, such as ITER. This project has studied the reaction of helium plasma on specific surface structures of tungsten, currently the material of choice for the divertor, due to its high thermal conductivity and related high-temperature properties.

Research efforts on this project involve large-scale molecular dynamics (MD) simulations that serve two purposes: (1) to provide a database of atomistic simulation data to aid in the development of longer-time-scale simulations, and (2) to provide a computer experiment with which to discover new mechanisms and ascertain the relative importance of various phenomena. ALCF facilities have been used to provide the lowest-flux, largest-surface-area simulations relevant to this line of work, providing a very detailed microscopic description of helium-induced structures on and under tungsten surfaces.

Results from additional MD studies performed at ALCF and the National Energy Research Scientific Computing Center (NERSC) show that helium atoms are strongly bound to grain boundaries—the boundary between individual crystallites in a polycrystalline material—and that there is a relatively high migration energy in the plane of the grain boundary. This is consistent with the idea that helium transport occurs via diffusion through the relatively undamaged tungsten bulk until it encounters a grain boundary or a bubble, at which point it effectively stops moving at relevant time scales. This results in effectively stationary helium bubbles that grow at the expense of mobile clusters, collection of helium on grain boundaries, and reduced helium bubble concentrations in the vicinity of grain boundaries.

IMPACT: The simulations at ALCF in particular have demonstrated that flux is an enormously important parameter in plasma-facing materials. Lower fluxes allow helium to reach greater depths before it becomes immobilized. This has implications for experimental comparison, as it indicates that much deeper and larger bubbles will form than those that exist during the high-flux simulations typically published to date.

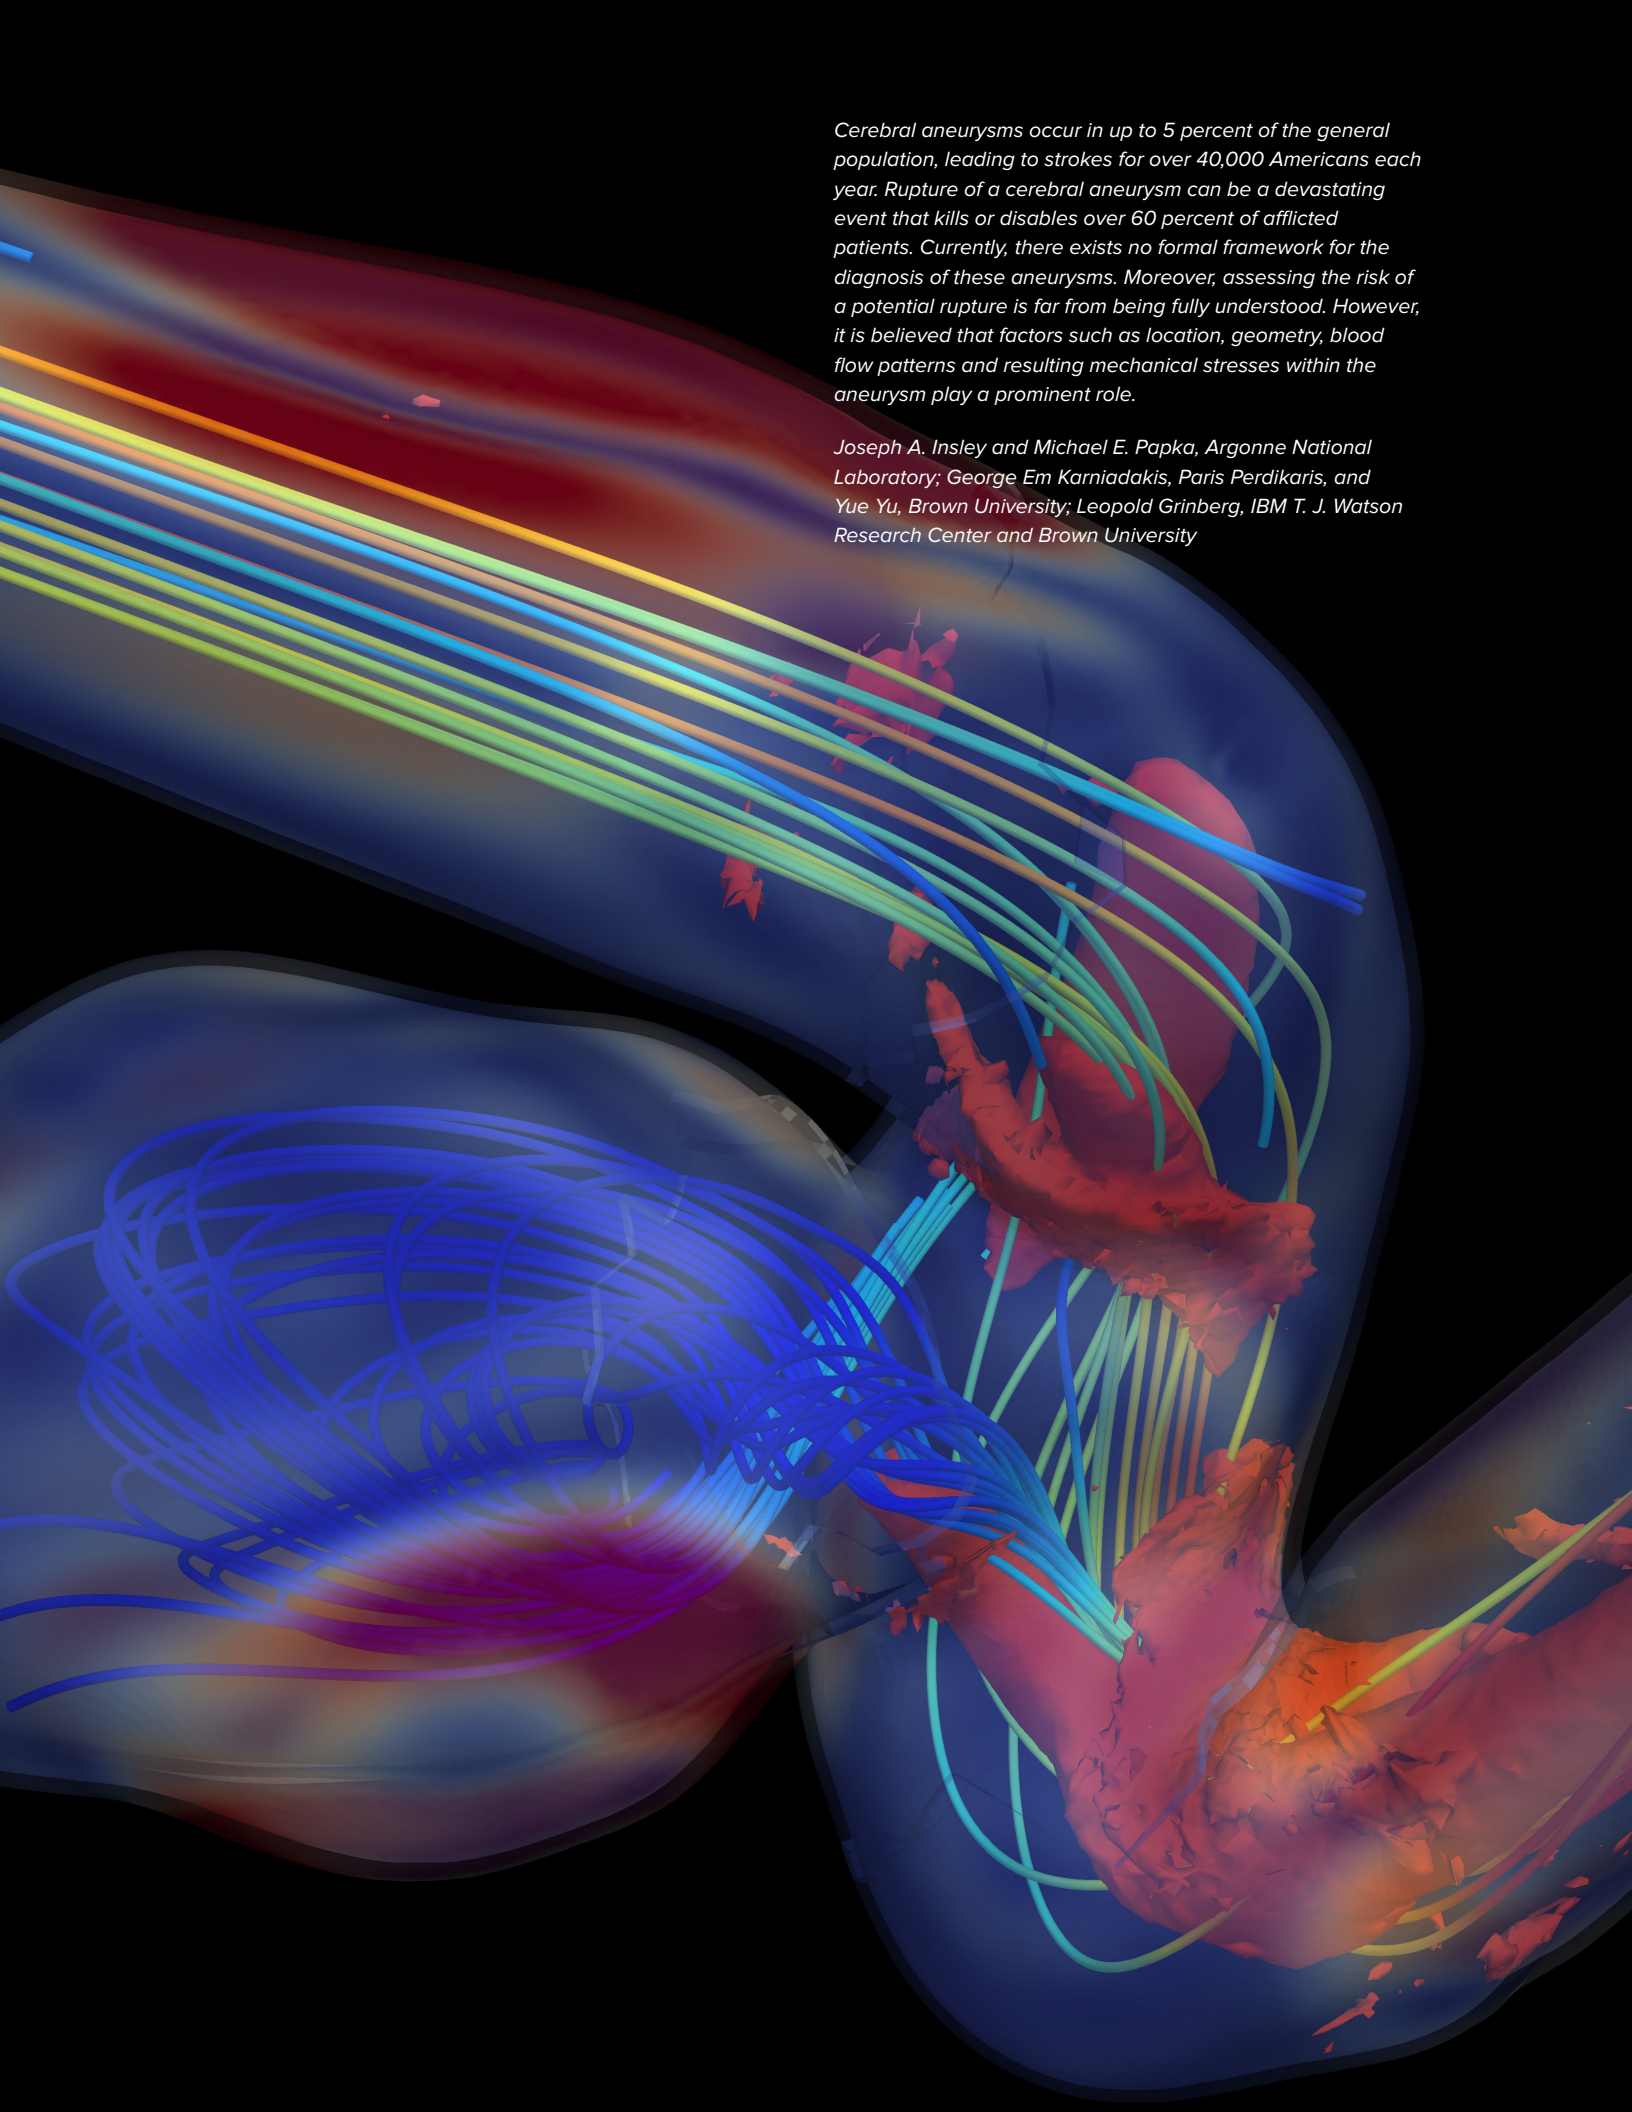

Cerebral aneurysms occur in up to 5 percent of the general population, leading to strokes for over 40,000 Americans each year. Rupture of a cerebral aneurysm can be a devastating event that kills or disables over 60 percent of afflicted patients. Currently, there exists no formal framework for the diagnosis of these aneurysms. Moreover, assessing the risk of a potential rupture is far from being fully understood. However, it is believed that factors such as location, geometry, blood flow patterns and resulting mechanical stresses within the aneurysm play a prominent role.

Joseph A. Insley and Michael E. Papka, Argonne National Laboratory; George Em Karniadakis, Paris Perdikaris, and Yue Yu, Brown University; Leopold Grinberg, IBM T. J. Watson Research Center and Brown University

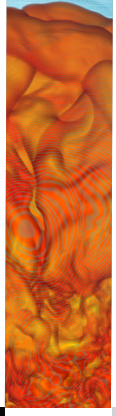
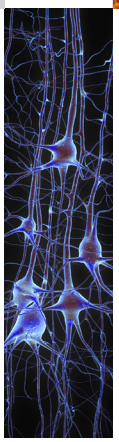

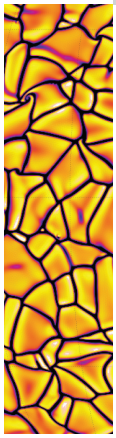

ALCF RESOURCES

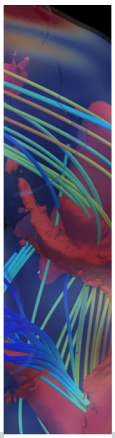

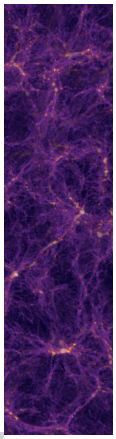
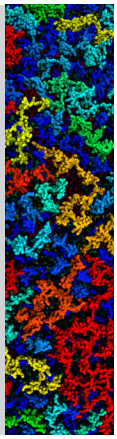
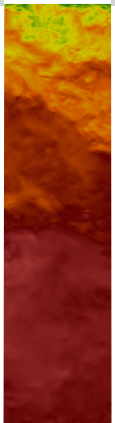

Computing Resources

Mira

Mira, the ALCF's 10-petaflops IBM Blue Gene/Q supercomputer, opens the door for scientists to analyze data more efficiently, design products more quickly, and address some of society's biggest problems in ways that would otherwise be impossible. An engineering marvel, the system is capable of carrying out 10 quadrillion calculations per second. Mira is also among the most energy-efficient supercomputers, saving considerable energy through innovative chip designs and a unique water-cooling system.

- 48 racks
- 16 1600 MHz PowerPC A2 cores per node
- 49,152 nodes
- 786,432 cores
- 768 TiB RAM
- 5D torus interconnect
- 384 I/O nodes
- Peak performance of 10 petaflops

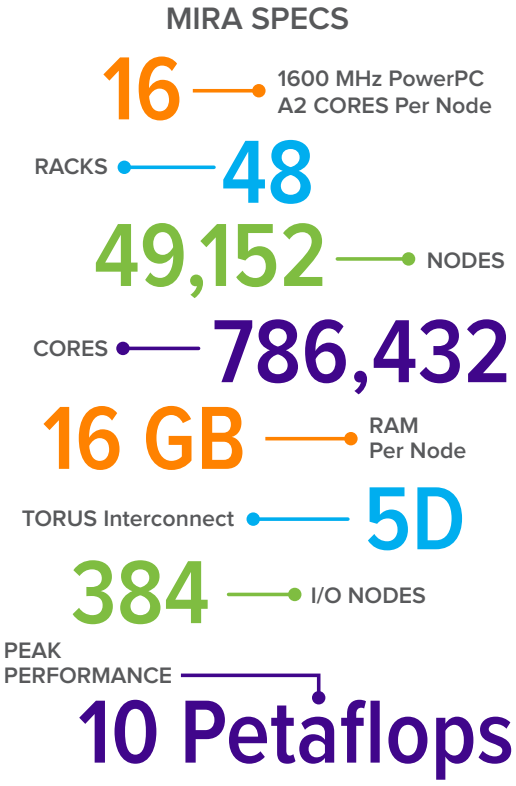

Cetus

The primary role of Cetus is to run small jobs to debug problems that occurred on Mira. Cetus shares the same software environment and file systems as Mira.

- 4 racks
- 4,096 nodes
- 64 TB RAM
- 5D torus interconnect
- 32 I/O nodes
- Peak performance of 838 teraflops

Vesta

As the ALCF's test and development platform, Vesta serves as a launching pad for researchers preparing to use Mira. Vesta has the same architecture as Mira, but on a much smaller scale.

- 2 racks
- 2,048 nodes
- 32 TB RAM
- 5D torus interconnect
- 32 I/O nodes
- Peak performance of 419 teraflops

Cooley

Cooley is the ALCF's analysis and visualization cluster. Equipped with graphics processing units (GPUs), Cooley converts computational data from Mira into high-resolution visual representations. The resulting images and videos help users to better analyze and understand the data generated by Mira. Cooley can also be used for statistical analysis, helping to pinpoint trends in the simulation data. Additionally, the system is capable of preprocessing efforts, such as meshing, to assist users preparing for Mira simulations. Cooley shares file systems with Mira, enabling direct access to Mira-generated results.

Each Cooley node has:

- 2 2.4 GHz Intel Haswell E5-2620 v3 6-core processors
- NVIDIA Tesla K80 GPU accelerator containing two Kepler GK210 GPUs
- 384 GB RAM
- 24 GB GPU RAM

The full Cooley system has:

- 126 nodes
- 1,512 cores
- FDR Infiniband interconnect
- 47 TB RAM
- 3 TB GPU RAM
- Peak performance of 293 teraflops

Data Storage

At the ALCF, disk storage provides intermediate-term storage for active projects, offering a means to access, analyze, and share simulation results. Tape storage is used to archive data from completed projects.

Disk Storage: The Blue Gene/Q data systems consist of 384 I/O nodes that connect to 22 storage arrays that control 13,000 disk drives with a total useable capacity of 27 PB and a maximum aggregate transfer speed of 330 GB/s over two file systems. The ALCF uses the GPFS file system to access the storage.

Tape Storage: The ALCF has two 10,000-slot libraries using LTO 6 tape technology. The LTO tape drives have built-in hardware compression with compression ratios typically between 1.25:1 and 2:1, depending on the data, giving an effective capacity of 26-40 PB.

Networking

The Blue Gene/Q systems have an internal proprietary network for communicating between nodes. InfiniBand enables communication between the I/O nodes and the storage system. Ethernet is used for external user access, and for maintenance and management of the systems.

The ALCF connects to other research institutions using up to 100 Gb/s of network connectivity. Scientists can transfer datasets to and from other institutions over fast research networks, such as the Energy Science Network (ESNet) and Internet2.

Cooley has nearly eight times the memory capacity of Tukey, the ALCF's previous data visualization and analysis system.

ALCF Expertise

The ALCF has assembled a world-class team of experts to help maximize the use of ALCF computing systems.

Catalysts are computational scientists who work directly with project teams to maximize and accelerate their research efforts. With multidisciplinary domain expertise, a deep knowledge of the ALCF computing environment, and experience with programming methods and community codes, the catalyst team helps users surmount any obstacles they may encounter using ALCF computing resources. This includes collaborating with colleagues from other ALCF teams to resolve issues when additional expertise is required.

Data Analytics and Visualization Experts facilitate the use of tools and methods for high-performance post processing of large datasets, interactive data exploration, batch visualization, and production visualization.

Operations ensures that system hardware and software work reliably and optimally; system tools are matched to the unique system architectures and scale of ALCF resources; the entire system software stack works smoothly together; and I/O performance issues, bug fixes, and requests for system software are addressed.

Performance Engineers help users achieve optimal performance on ALCF resources by working with them to port, tune, and parallelize scientific applications and other software on current and emerging systems. The team also helps drive the future of high-performance computing and data-intensive computing through research and development of performance models and scalable algorithms, and the development of standards for key areas, such as benchmarks, compilers, and programming models.

User Experience coordinates communications, support, and outreach for the ALCF. The communications team stewards ALCF user stories and ensures that the facility provides relevant and accurate information to users, DOE, and the public at large. The support team manages the workflows for user accounts and projects, and assists users with support requests related to their ALCF projects. Outreach focuses on organizing training efforts and user-centered events for the ALCF.

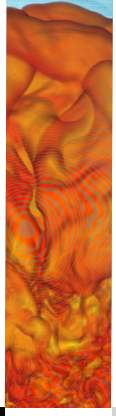
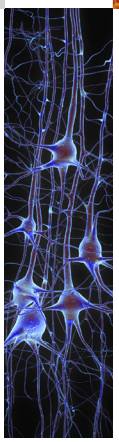

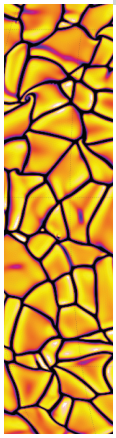

2015 ALCF PROJECTS

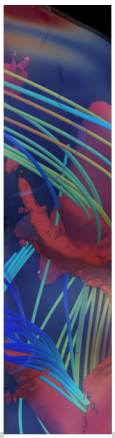

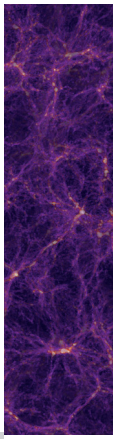
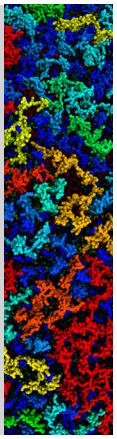
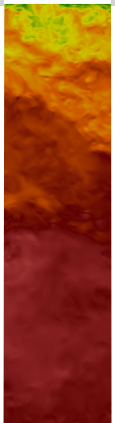

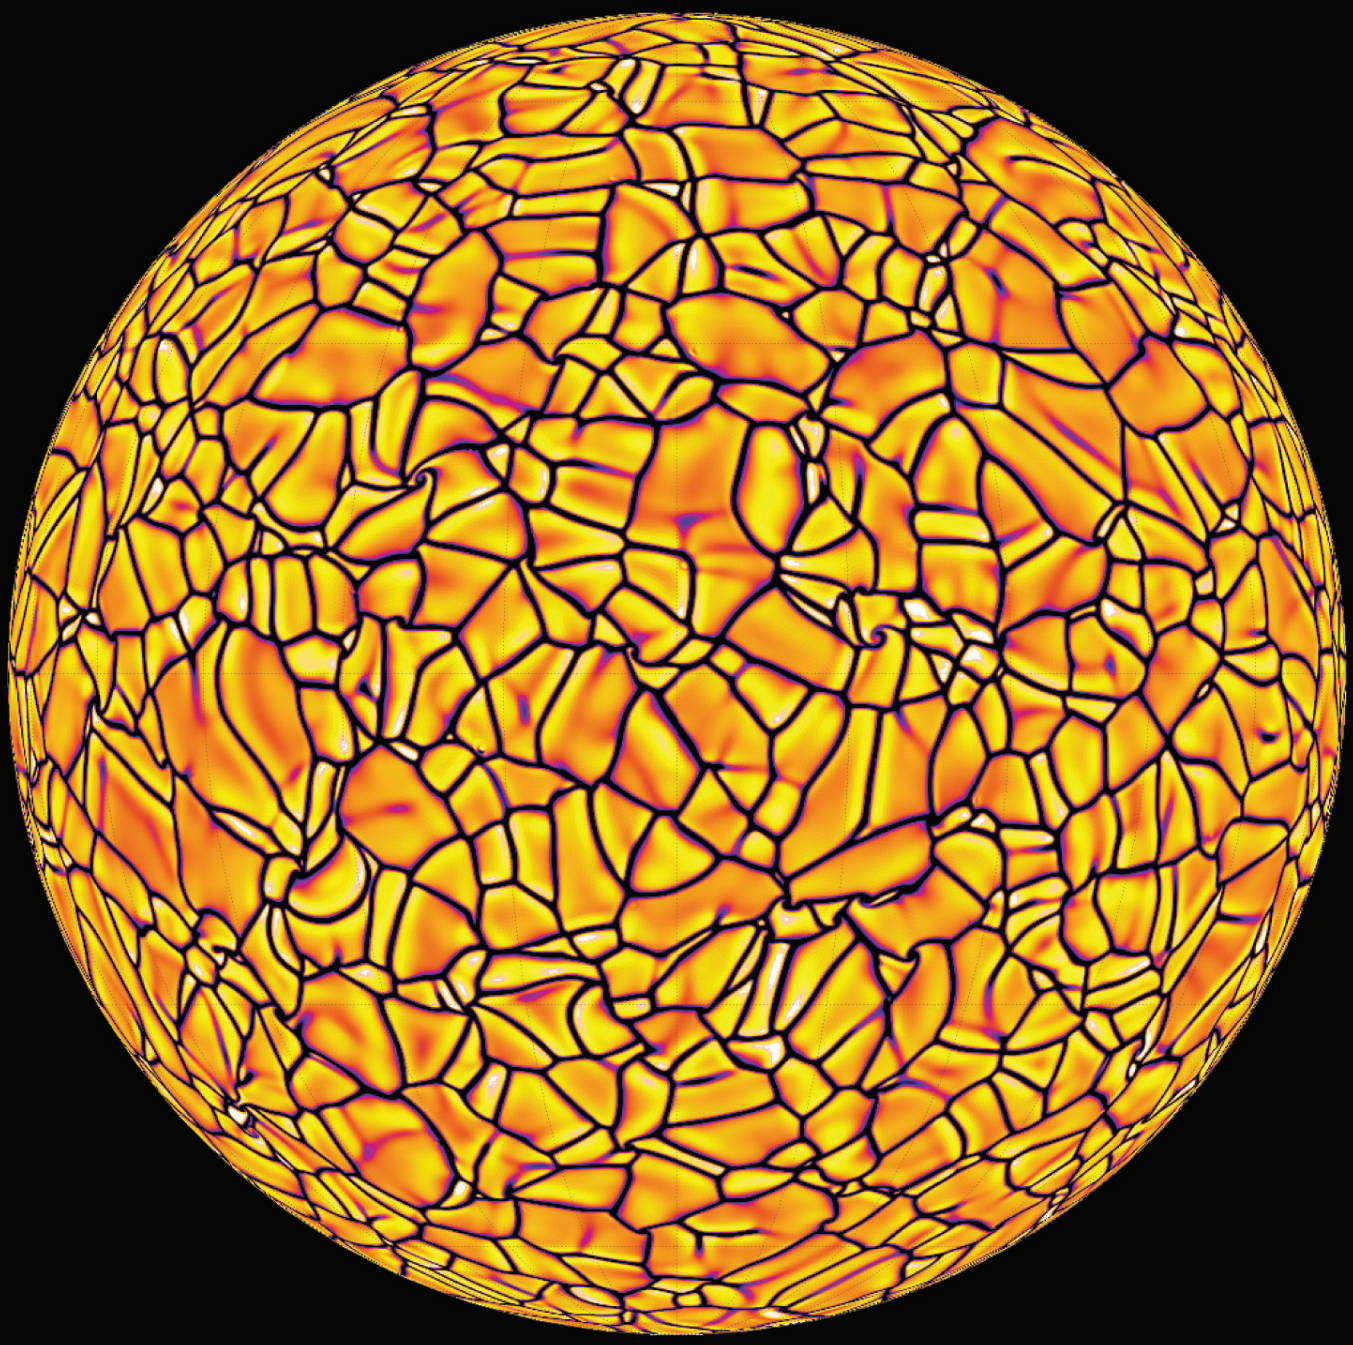

Radial velocity near the surface of a turbulent solar simulation run on Mira. This simulation possesses a density stratification of four density scale heights and a Rayleigh number of about 3x109. Upflows are denoted in yellow-orange tones, and downflows are indicated by dark blue tones.

Nicholas Featherstone, University of Colorado Boulder

2015 INCITE PROJECTS

Biological Sciences

Multiscale Simulations of Human Pathologies
George Karniadakis, Brown University
70 Million Core-Hours (ALCF: 45M; OLCF: 25M)

Studies of Large Conformational Changes in Biomolecular Machines
Benoît Roux, University of Chicago
120 Million Core-Hours

Chemistry

Catalyst Support Interactions
Frank Abild-Pedersen, Stanford University
50 Million Core-Hours

First-Principles Simulations of High-Speed Combustion and Detonation
Alexei Khokhlov, University of Chicago
150 Million Core-Hours

Towards Breakthroughs in Protein Structure Calculation and Design
David Baker, University of Washington
80 Million Core-Hours

Computer Science

Dynamic and Adaptive Parallel Programming for Exascale Research
Robert Harrison, Brookhaven National Laboratory
15 Million Core-Hours

Performance Evaluation and Analysis Consortium End Station
Leonid Oliker, Lawrence Berkeley National Laboratory
90 Million Core-Hours (ALCF: 45M; OLCF: 45M)

Scalable System Software for Parallel Programming
Robert Latham, Argonne National Laboratory
25 Million Core-Hours

Earth Science

Accelerated Climate Modeling for Energy
Mark Taylor, Sandia National Laboratories
190 Million Core-Hours (ALCF: 140M; OLCF: 50M)

CESM Century-Scale Climate Experiments with a High-Resolution Atmosphere
Warren Washington, National Center for Atmospheric Research
200 Million Core-Hours (ALCF: 200M; OLCF: 0)

Frontiers in Planetary and Stellar Magnetism through High-Performance Computing
Jonathan Aurnou, University of California, Los Angeles
83 Million Core-Hours

High Frequency Ground Motion Simulation for Seismic Hazard Analysis
Thomas Jordan, University of Southern California
167 Million Core-Hours (ALCF: 48M; OLCF: 119M)

Engineering

Adaptive Detached Eddy Simulation of a High Lift Wing with Active Flow Control
Kenneth Jansen, University of Colorado Boulder
70 Million Core-Hours

Direct Numerical Simulations and Robust Predictions of Cloud Cavitation Collapse
Petros Koumoutsakos, Swiss Federal Institute of Technology
88 Million Core-Hours

DNS/LES of Complex Turbulent Flows
Krishnan Mahesh, University of Minnesota
100 Million Core-Hours

Large-Eddy Simulations of Combustor Liner Flows
Anne Dord, GE Global Research
INCITE 2015
89 Million Core-Hours

Large-Eddy Simulation of the Bachalo-Johnson Flow, with Shock-Induced Separation
Philippe Spalart, Boeing
135 Million Core-Hours

Parameter Studies of Boussinesq Flows
Susan Kurien, Los Alamos National Laboratory
44 Million Core-Hours

Materials Science

Computational Spectroscopy of Heterogeneous Interfaces
Giulia Galli, University of Chicago
180 Million Core-Hours

Non-Covalent Bonding in Complex Molecular Systems with Quantum Monte Carlo
Dario Alfè, University College London
148 Million Core-Hours (ALCF: 68M; OLCF: 80M)

Petascale Simulations of Self-Healing Nanomaterials
Rajiv Kalia, University of Southern California
180 Million Core-Hours

Predictive Materials Modeling for Li-Air Battery Systems
Larry Curtiss, Argonne National Laboratory
50 Million Core-Hours

QMC Simulations DataBase for Predictive Theory and Modeling
David Ceperley, University of Illinois at Urbana–Champaign
185 Million Core-Hours (ALCF: 100; OLCF: 85)

Reactive MD Simulations of Electrochemical Oxide Interfaces at Mesoscale
Subramanian Sankaranarayanan, Argonne National Laboratory
40 Million Core-Hours

Simulation of Correlated Electrons for Superconducting Materials
Lucas Wagner, University of Illinois at Urbana–Champaign
106 Million Core-Hours

SiO2 Fracture: Chemomechanics with a Machine Learning Hybrid QM/MM Scheme
James Kermode, King's College London
125 Million Core-Hours

State-of-the-Art Simulations of Liquid Phenomena
Mark Gordon, Iowa State University
200 Million Core-Hours

Physics

Accelerator Modeling for Discovery
James Amundson, Fermilab
60 Million Core-Hours

Cosmic Reionization on Computers
Nickolay Gnedin, Fermilab
74 Million Core-Hours

Cosmological Simulations for Large-Scale Sky Surveys
Salman Habib, Argonne National Laboratory
160 Million Core-Hours (ALCF: 80M; OLCF: 80M)

High-Fidelity Simulation of Tokamak Edge Plasma Transport
Choong-Seock Chang, Princeton Plasma Physics Laboratory
270 Million Core-Hours (ALCF: 100M; OLCF: 170M)

Lattice QCD
Paul Mackenzie, Fermilab
280 Million Core-Hours (ALCF: 180M; OLCF: 100M)

Nuclear Structure and Nuclear Reactions
James Vary, Iowa State University
204 Million Core-Hours (ALCF: 100M; OLCF: 104M)

Particle Acceleration in Shocks: From Astrophysics to Laboratory In Silico
Frederico Fiuza, SLAC National Accelerator Laboratory
110 Million Core-Hours

Petascale Simulation of Magnetorotational Core-Collapse Supernovae
Sean Couch, Michigan State University
50 Million Core-Hours

Petascale Simulations of Laser Plasma Interaction Relevant to IFE
Frank Tsung, University of California, Los Angeles
90 Million Core-Hours

Quark Flavors and Conserved Charges at Finite Density in the QCD Phase Diagram
Rene Bellwied, University of Houston
150 Million Core-Hours

2014 ALCC PROJECTS

Biological Sciences

Applying Breakthroughs in Protein Structure Calculation to the Creation of Designer Enzymes
David Baker, University of Washington
200 Million Core-Hours

Chemistry

Influence of Morphology on Proton Transport in Proton Exchange Membrane
Gregory Voth, University of Chicago/Argonne National Laboratory
57.6 Million Core-Hours

Large-Scale Turbulent Clean Coal Combustion
Martin Berzins, University of Utah
40 Million Core-Hours (ALCF: 10M; OLCF: 30M)

Computer Science

Hobbes: Operating System and Runtime Research for Extreme Scale
Ron Brightwell, Sandia National Laboratories
40 Million Core-Hours (ALCF: 5M; NERSC: 5M; OLCF: 30M)

Earth Science

Delivering the Department of Energy's Next-Generation High-Resolution Earth System Model
Peter Thornton, Oak Ridge National Laboratory
137 Million Core-Hours (ALCF: 107M; OLCF: 30M)

Engineering

Amplitude Modulation of Wind Turbine Noise
Sanjiva Lele, Stanford University
36.5 Million Core-Hours

Petascale Simulations in Support of CESAR
Elia Merzari, Argonne National Laboratory
80 Million Core-Hours

Predictive Large-Eddy Simulation of Jet Fuel Atomization, High-Lift Airframes, and Reacting Supersonic Turbulent Flows on Unstructured Grids
Parviz Moin, Stanford University
120 Million Core-Hours

Understanding Secondary Motions and Their Impact in Modeling Turbulent Flows
Hassan Nagib, Illinois Institute of Technology
11 Million Core-Hours

Materials Science

Interfaces in Organic and Hybrid Photovoltaics
Noa Marom, Tulane University
105 Million Core-Hours (ALCF: 75M; NERSC: 30M)

Large-Scale Quantum Simulations of Electrode-Electrolyte Interfaces
Giulia Galli, University of Chicago
35 Million Core-Hours

Nanostructure-Enhanced Chemical Reactivity and Detonation in Energetic Materials
Aidan Thompson, Sandia National Laboratories
80 Million Core-Hours

Prediction and Design of Energy Materials by Petascale Evolutionary Algorithm Simulations
Giancarlo Trimarchi, Northwestern University
30 Million Core-Hours

Revealing the Reversible Electrodeposition Mechanism in Multivalent-Ion Batteries
Gerbrand Ceder, Massachusetts Institute of Technology
98 Million Core-Hours

Physics

Composite Higgs Theory Beyond the Standard Model and the 14 TeV Upgrade of the Large Hadron Collider
Julius Kuti, University of California, San Diego
88.7 Million Core-Hours

Cosmic Frontier Computational End-Station
Salman Habib, Argonne National Laboratory
176 Million Core-Hours (ALCF: 100M; NERSC: 76M)

Ion Solvation, Catalytic Interfaces, and Extreme Aqueous Environments: An *Ab Initio* Study of Liquid Water
Robert A. DiStasio Jr., Cornell University
350 Million Core-Hours

Simulation of Large Hadron Collider Events Using Leadership Computing
Thomas LeCompte, Argonne National Laboratory
52 Million Core-Hours (ALCF: 50M; NERSC: 2M)

Turbulent Multiphase Flows for Nuclear Reactor Safety
Igor A. Bolotnov, North Carolina State University
76.8 Million Core-Hours

Understanding Helium Plasma Mediated Tungsten Surface Response to Better Predict Fusion Plasma Facing Component Performance in ITER
Brian Wirth, University of Tennessee
96 Million Core-Hours (ALCF: 66M; OLCF: 30M)

Validation Studies of Gyrokinetic Simulations to Understand the Coupling of Ion and Electron Scale Turbulence in Tokamak Plasmas
Christopher Holland, University of California, San Diego
140 Million Core-Hours (ALCF: 90M; NERSC: 50M)

2015 ALCC PROJECTS

Chemistry

Anomalous Density Properties and Ion Solvation in Liquid Water: A Path-Integral Ab Initio Study
Robert A. DiStasio Jr., Cornell University
175 Million Core-Hours

Computational Design of Interfaces for Photovoltaics
Noa Marom, Tulane University
120 Million Core-Hours (ALCF: 100M; NERSC: 20M)

PT-Symmetric Quantum Mechanics for Real-Time Electron Transport Simulations
Hanning Chen, George Washington University
16 Million Core-Hours

Computer Science

Demonstration of the Scalability of Programming Environments By Simulating Multi-Scale Applications

Robert Voigt, Leidos Inc.
167 Million Core-Hours (ALCF: 127M; OLCF: 40M)

Performance Analysis, Modeling and Scaling of HPC Applications and Tools
Abhinav Bhatele, Lawrence Livermore National Laboratory
29.4 Million Core-Hours (ALCF: 20.1M; OLCF: 9.3M)

Portable Application Development for Next-Generation Supercomputer Architectures
Tjerk Straatsma, Oak Ridge National Laboratory
160 Million Core-Hours (ALCF: 60M; NERSC: 40M; OLCF: 60M)

Earth Science

Delivering the Department of Energy's Next-Generation High-Resolution Earth System Model
Peter Thornton, Oak Ridge National Laboratory
165 Million Core-Hours (ALCF: 110M; OLCF: 55M)

Validation of RAP/HRRR for the Wind Forecast Improvement Project II
Joe Olson, National Oceanic and Atmospheric Administration
15 Million Core-Hours

Engineering

Advancing Internal Combustion Engine Simulations Using Sensitivity Analysis
Sibendu Som, Argonne National Laboratory
60 Million Core-Hours

Computational Design of Novel Multiscale Concrete Rheometers
William George, National Institute of Standards and Technology
50 Million Core-Hours

Credible Predictive Simulation Capabilities for Advanced Clean Energy Technology Development through Uncertainty Quantification
Aytekin Gel, ALPEMI
111.5 Million Core-Hours

High-Fidelity Computations of Fuel Assemblies Subjected to Seismic Loads
Elias Balaras, George Washington University
34 Million Core-Hours

Large-Eddy Simulation of Turbine Internal Cooling Passages
Gustavo Ledezma, GE Global Research
6 Million Core-Hours

Toward a Longer-Life Core: Thermal-Hydraulic CFD Simulations of Deformed Fuel Assemblies
Elia Merzari, Argonne National Laboratory
72 Million Core-Hours

Materials Science

First-Principles Large-Scale Simulations of Interfaces for Energy Conversion and Storage
Marco Govoni, University of Chicago/Argonne National Laboratory
75 Million Core-Hours

Large-Scale *Ab Initio* Simulation of Crystalline Defects in Mg-alloys
Kaushik Bhattacharya, Caltech
20 Million Core-Hours

Predictive Modeling of Functional Nanoporous Materials
J. Ilja Siepmann, University of Minnesota
120 Million Core-Hours

Revealing the Reversible Electrodeposition Mechanism in Multivalent-ion Batteries
Gerbrand Ceder, Massachusetts Institute of Technology
70 Million Core-Hours

Physics

Cosmic Frontier Computational End-Station
Salman Habib, Argonne National Laboratory
115 Million Core-Hours (ALCF: 65M; NERSC: 15M; OLCF: 35M)

An End-Station for Intensity and Energy Frontier Experiments and Calculations
Thomas LeCompte, Argonne National Laboratory
78 Million Core-Hours (ALCF: 62; NERSC: 16M)

Hadronic Light-By-Light Scattering Contribution to the Muon Anomalous Magnetic Moment from Lattice QCD with Chiral Fermions
Thomas Blum, University of Connecticut
175 Million Core-Hours

Large-Eddy Simulation and Direct Numerical Simulation of Fluid Induced Loads on Reactor Vessel Internals
Milorad Dzodzo, Westinghouse
40 Million Core-Hours

Understanding Helium-Hydrogen Plasma Mediated Tungsten Surface Response to Predict Fusion Plasma Facing Component Performance in ITER
Brian Wirth, University of Tennessee, Knoxville
116 Million Core-Hours (ALCF: 80M; OLCF: 36M)

Validation Simulations of Macroscopic Burning-Plasma Dynamics
Jacob King, Tech-X
40 Million Core-Hours

2015 DIRECTOR'S DISCRETIONARY PROJECTS

The following list provides a sampling of the many Director's Discretionary projects at the ALCF.

Biological Sciences

Binding Affinity Calculations of Estrogen Receptor Against FDA-Approved Drugs
Sichun Yang, Case Western Reserve University
8 Million Core-Hours

Computing Three-Dimensional Structures of Large RNA from Small Angle X-Ray Scattering Data and Secondary Structure
Yun-Xing Wang, National Cancer Institute
15 Million Core-Hours

Correlating Experimentally Measured Molecular Dynamics with Computational Trajectories: Understanding Dynamic Allostery in Ubiqu
R. Andrew Byrd, National Cancer Institute
5 Million Core-Hours

Developing Novel Umbrella Sampling/Solute Tempering Algorithms in NAMD
Sunhwan Jo, Argonne National Laboratory
4.1 Million Core-Hours

Highly Parallel Macromolecular Conformational Searches and Energy Evaluations with the CHARMM Program
Robert J. Petrella, Harvard University
10 Million Core-Hours

Long-Time MD Simulation of Protein Structural Function
Ruth Nussinov, National Cancer Institute
2 Million Core-Hours

Virus Calculations with FMO
Yuri Alexeev, Argonne National Laboratory
15 Million Core-Hours

Chemistry

Large-Scale Combustion Preparatory Access
Gabriel Staffelbach, CERFACS
12 Million Core-Hours

Modeling Nonadiabatic Spin-Forbidden Reaction Mechanisms in Metal-Sulfur Proteins
Sergey Varganov, University of Nevada, Reno
2 Million Core-Hours

Quantum Monte Carlo Applied to Lithium Hyperoxides in Li-Air Batteries
John J. Low, Argonne National Laboratory
10 Million Core-Hours

Shift-and-Invert Parallel Spectral Transformation Eigensolver
Murat Keceli, Argonne National Laboratory
500,000 Core-Hours

Solving Petascale Public Health & Safety Problems Using Uintah
Martin Berzins, University of Utah
1 Million Core-Hours

Computer Science

ExaHDF5: Advancing HDF5 HPC I/O to Enable Scientific Discovery
Venkatram Vishwanath, Argonne National Laboratory
5 Million-Core Hours

Performance Studies of Three-Dimensional Fast Fourier Transforms Using Overlap of Communication with Computation
Dmitry Pekurovsky, University of California, San Diego
2.5 Million Core-Hours

SciDAC Scalable Data Management Analysis and Visualization
Michael E. Papka, Argonne National Laboratory
3.25 Million Core-Hours

Earth Science

Scalability Study for NUMA (Non-Hydrostatic Unified Model of the Atmosphere)
Andreas Mueller, Naval Postgraduate School
5 Million Core-Hours

Energy Technologies

GTRI- and NEAMS-Related Production Tests and Runs
Micheal A. Smith, Argonne National Laboratory
10 Million Core-Hours

Engineering

Aircraft Fuel Burn Reduction Using Minute Roughness Elements
Ali Uzun, Florida State University
1 Million Core-Hours

DNS of Compressible Turbulent Boundary Layers
Jonathan Poggie, Air Force Research Laboratory
5 Million Core-Hours

DNS of Multi-Mode Rayleigh-Taylor Instability
Maxwell Hutchinson, University of Chicago
5 Million Core-Hours

Extreme-Scale Unstructured Mesh CFD Workflow
Cameron Smith, Rensselaer Polytechnic Institute
5 Million Core-Hours

Large-Eddy Simulation of Crackling Supersonic Jets
Joseph Nichols, University of Minnesota
10 Million Core-Hours

Numerical Simulation of Acoustic Radiation from High-Speed Turbulent Boundary Layers
Lian Duan, Missouri University of Science and Technology
2 Million Core-Hours

Performance Improvement of CFD Code CONVERGE on BG/Q Systems
Marta García and Sibendu Som, Argonne National Laboratory
6 Million Core-Hours

Primary Atomization DNS of ECN's Spray A
Marcus Herrmann, Arizona State University
5 Million Core-Hours

Turbulent Rayleigh-Benard Convection at High Rayleigh and Low Prandtl Numbers
Janet Scheel, Occidental College
1 Million Core-Hours

Materials Science

Collective I/O and Bond Analysis Code Development on SiC Nanoparticle Oxidation
Ying Li, Argonne National Laboratory
5 Million Core-Hours

Electronic Response to Particle Radiation in Semiconductor Systems
Andre Schleife, University of Illinois at Urbana-Champaign
16 Million Core-Hours

High-Performance Li-Air Battery
Ying Li, Argonne National Laboratory
5 Million Core-Hours

Integrating Simulation and Observation: Discovery Engines for Big Data
Justin Wozniak, Argonne National Laboratory
6.9 Million Core-Hours

Many-Body Stochastic Analysis of Semiconductor Bulk and Defect Properties
Elif Ertekin, University of Illinois at Urbana-Champaign
2 Million Core-Hours

Mathematics

Parallel Multiscale Simulations of Advanced Steel Materials
Axel Klawonn, University of Cologne;
Oliver Rheinbach, Technische Universität Bergakademie Freiberg
3 Million Core-Hours

Scalable Domain Decomposition Methods for Computational Cardiology and Isogeometric Analysis
Luca F. Pavarino, University of Milan
1.3 Million Core-Hours

Physics

Calculation of Nuclear Matrix Element of Neutrinoless Double-Beta Decay
Jun Terasaki, University of Tsukuba
3.8 Million Core-Hours

DNS Simulations of Turbulent Rayleigh-Taylor Unstable Flames Using Nek5000
Elizabeth P. Hicks, Epsilon Delta Labs
2.5 Million Core-Hours

Effective Interactions in Coulombic Systems with Highly Disparate Particle Sizes
Monica Olvera de la Cruz, Northwestern University
3.1 Million Core-Hours

Extreme-Scale Turbulence Simulations
William M. Tang, Princeton Plasma Physics Laboratory
60 Million Core-Hours

Mira Simulations of High-Intensity Laser Experiments to Study Turbulent Amplification of Magnetic Fields
Don Q. Lamb, University of Chicago
10 Million Core-Hours

Moving Mesh Simulations of Gravito-Turbulence in Global Proto-Planetary Disks
Andrew MacFadyen, New York University
250,000 Core-Hours

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 ALCF is supported by DOE's Office of Science, Advanced Scientific Computing Research (ASCR) program.

About Argonne National Laboratory

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

Availability of this Report (ANL/ALCF-15/2)

Online Access: U.S. Department of Energy (DOE) reports produced after 1991 and a growing number of pre-1991 documents are available free via DOE's SciTech Connect (http://www.osti.gov/scitech/).

Reports not in digital format may be purchased by the public from the National Technical Information Service (NTIS):

U.S. Department of Commerce
National Technical Information Service
5301 Shawnee Rd
Alexandra, VA 22312
www.ntis.gov
phone | 800.553.NTIS (6847) or 703.605.6000
fax | 703.605.6900
orders@ntis.gov

Reports not in digital format are available to DOE and DOE contractors from the Office of Scientific and Technical Information (OSTI):

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

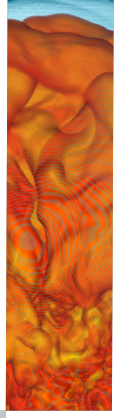
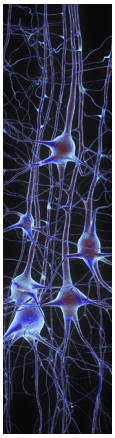
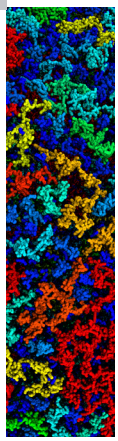
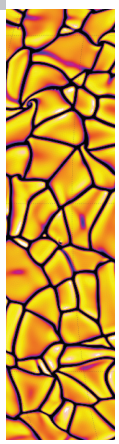
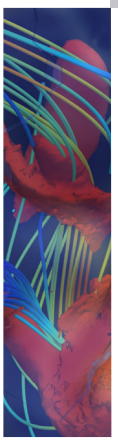
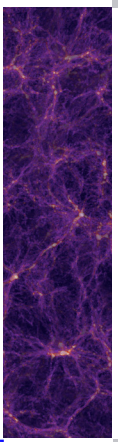
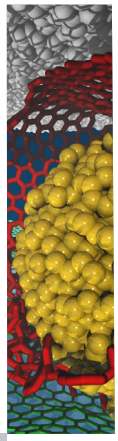
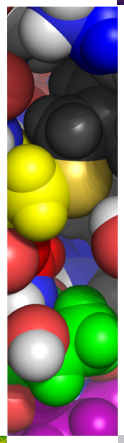
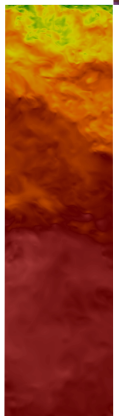

alcf.anl.gov

media@alcf.anl.gov

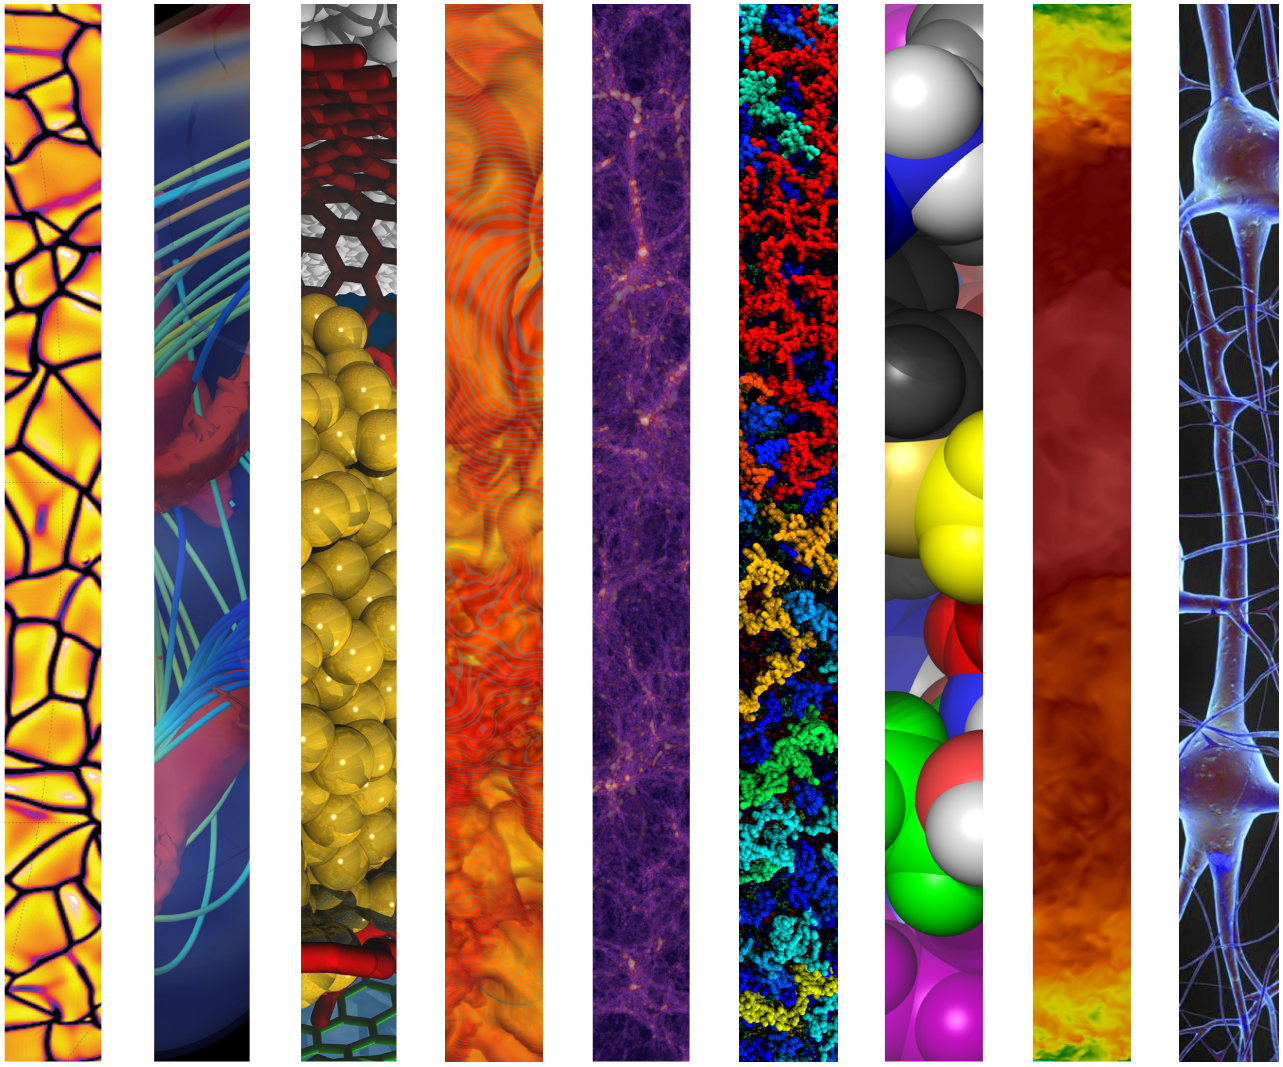

U.S. DEPARTMENT OF ENERGY

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

alcf_science_brochure2015_rc