2017 ARGONNE LEADERSHIP COMPUTING FACILITY

Argonne NATIONAL LABORATORY

# Science

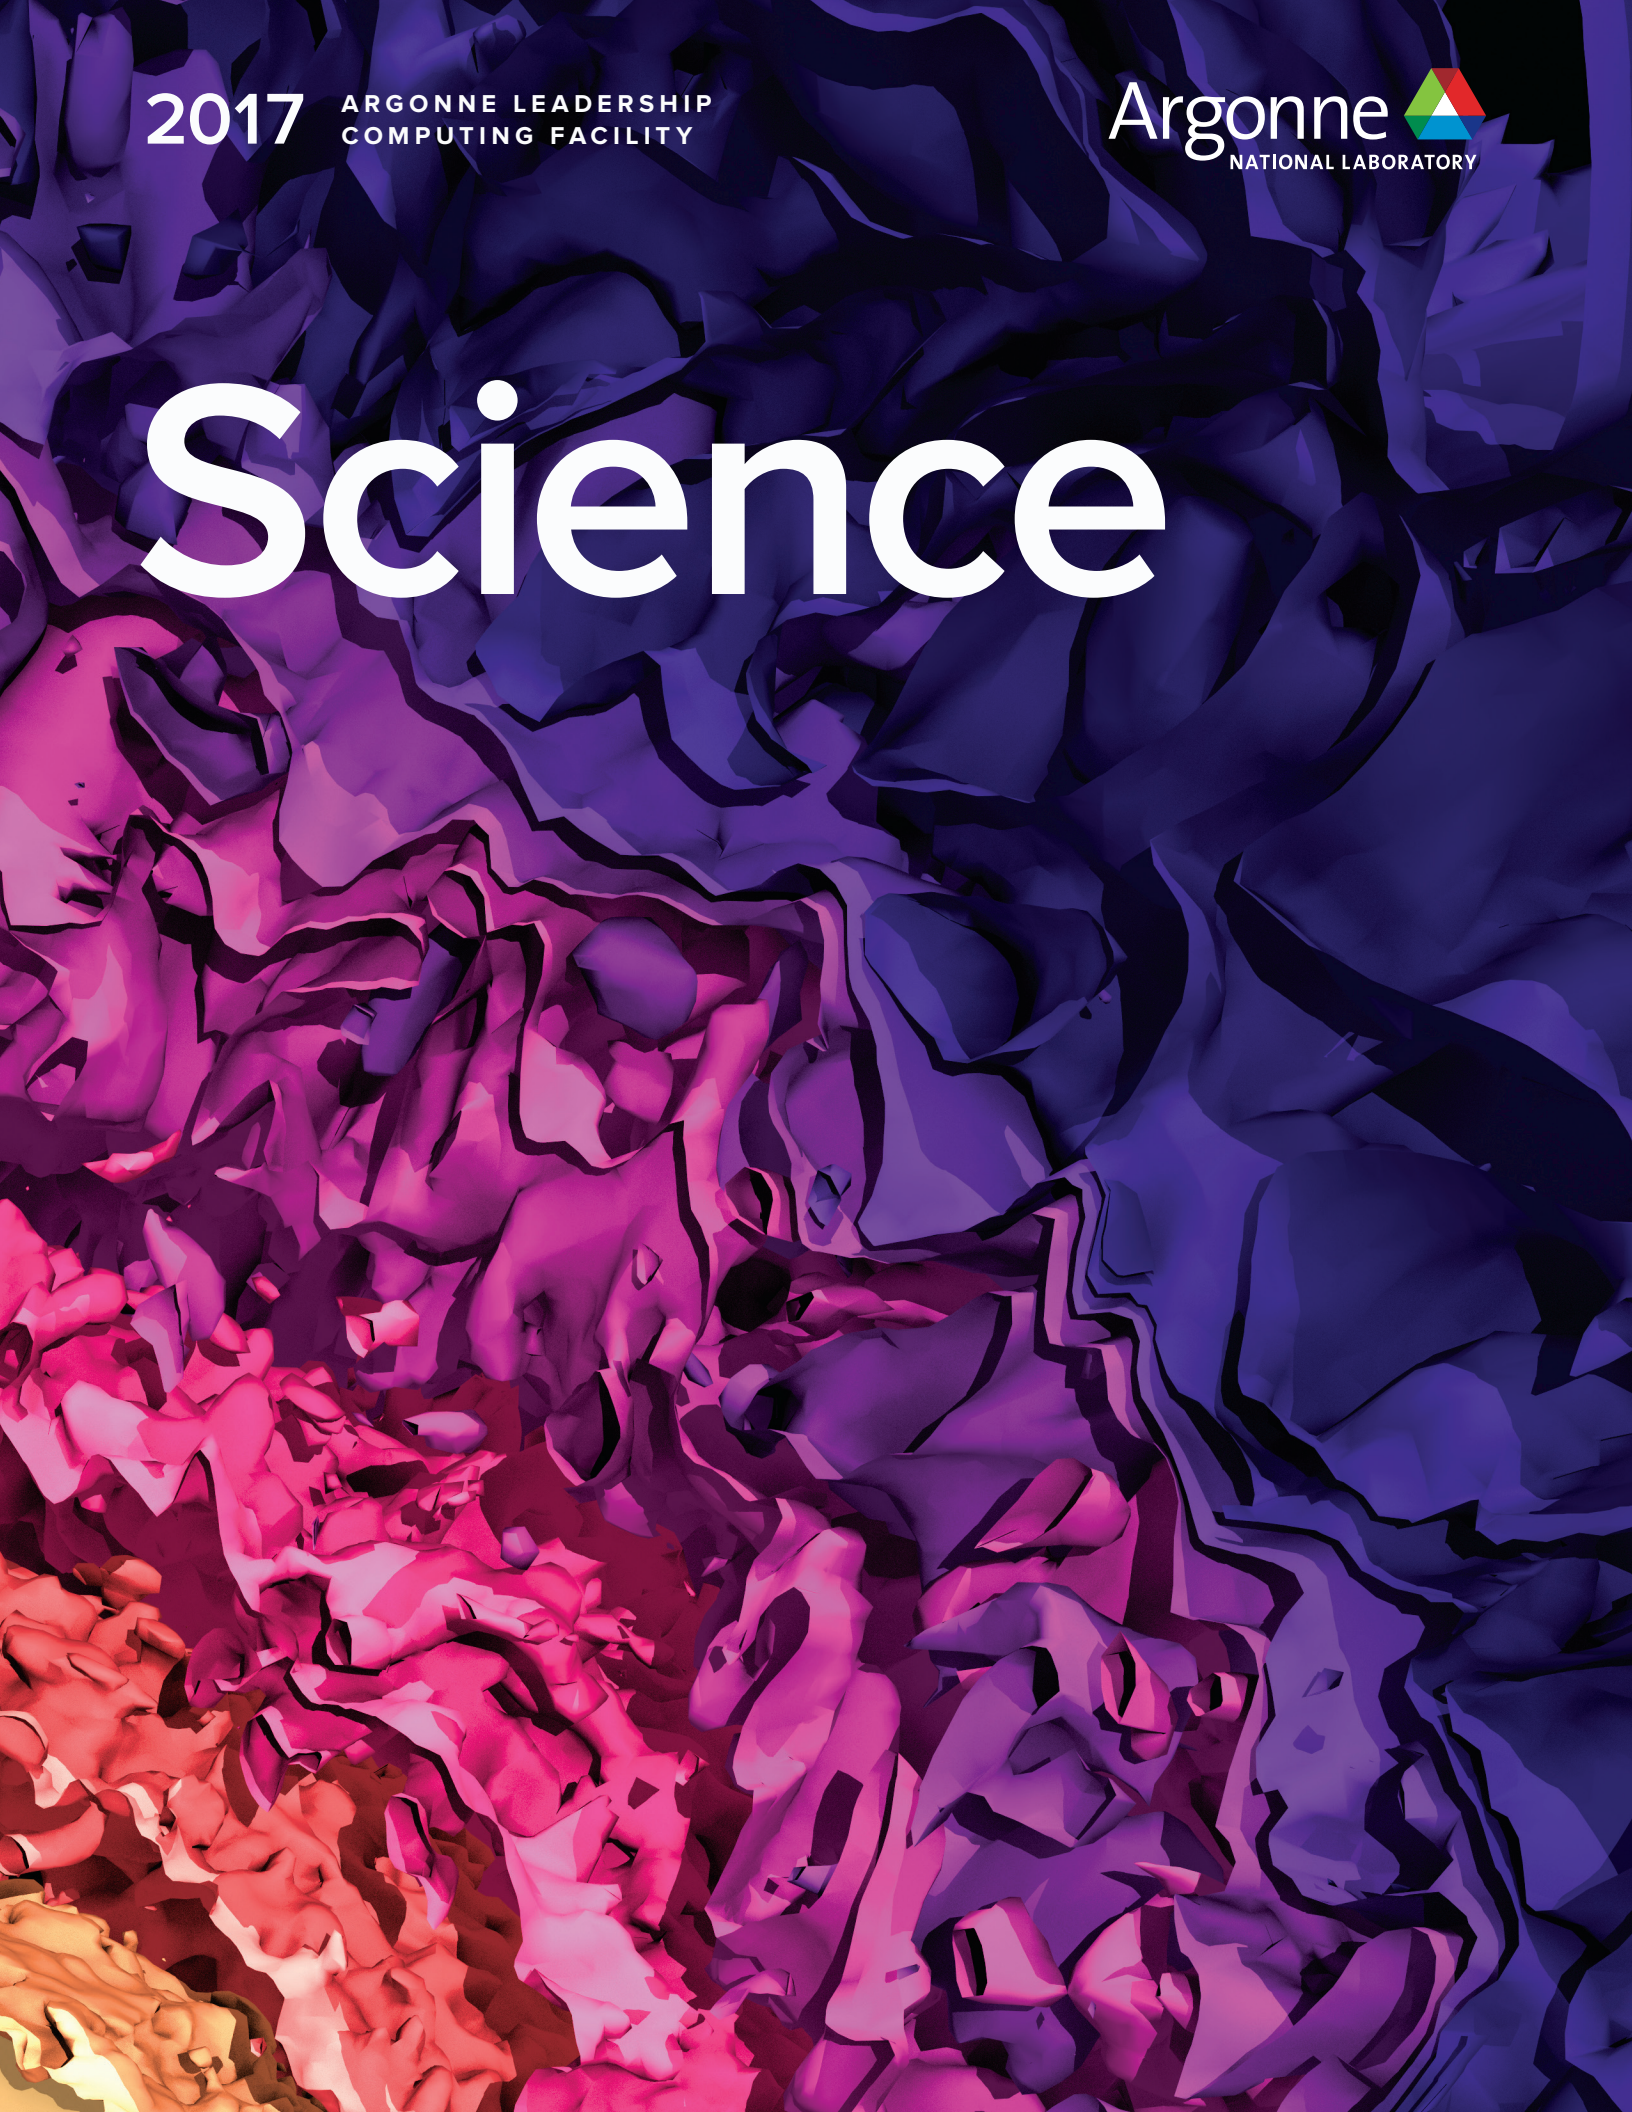

On the cover: Global radiation hydrodynamic simulations of an 80-solar-mass star envelope performed on Mira. Convection around the iron opacity peak region causes the turbulent structures shown here. Density is highest in the yellow region near the core of the star. Low-density wind is launched near the surface. *Image credit: Joseph A. Insley, Argonne National Laboratory; Matteo Cantiello, Flatiron Institute; Lars Bildsten, Omer Blaes, and Yan-Fei Jiang, University of California, Santa Barbara; Eliot Quataert, University of California, Berkeley*

CONTENTS

# 2017 Science

## SCIENCE DIRECTOR'S MESSAGE

**Katherine Riley**
Director of Science

The arrival in July of our newest supercomputer, a Cray XC40 system called Theta, doubled the ALCF's production computing capabilities from one petascale machine to two, added an Intel-based architecture to the mix, and opened up exciting new opportunities for our user community.

All of the science and code development projects that participated in Theta's Early Science Program, expertly administered by Deputy Director of Science Tim Williams, reported significant performance gains. More than half of these projects—selected for their ability to stress-test the machine's infrastructure for a broad range of applications—are already seeing science results.

Meanwhile, Mira has continued to help deliver scientific achievements. A National Cancer Institute team perfected a technique that accurately computes the 3D structure of RNA sequences—giving researchers a window into these complex biological polymers that are fundamentally involved in health and disease, and a leg up in the fight for better treatments and cures.

The ALCF's powerful computing resources garner considerable attention for their vast potential to advance knowledge and innovations. Simulation has brought a new way of doing science, but in many ways still serves the traditional pillars of science that are theory and experiment.

We and others in the HPC community are now rethinking the role that computation plays in scientific discovery. The intersections of simulation, data analytics, and machine learning are becoming more and more ubiquitous. The ALCF Data Science Program is but one effort aimed at improving a wide range of data science techniques for an evolving computing ecosystem. These projects, which you will read about in the following pages, are using both production and experimental resources, new workflows, and machine learning to probe large-scale experimental and observational data.

As a scientist, I am eager to see what capabilities future generations of systems will bring: what boundaries we can push and what problems we can solve as we press on toward achieving exascale capabilities. We are preparing to enter a new era in scientific computing together, and I am thrilled to be part of this organization's future breakthroughs.

Lastly, I am also mindful of the amazing staff we have supporting the science at the ALCF. Each member of our staff performs a role that is critical to our users' ability to achieve their goals. Since assuming the role of director of science, I've had more opportunities to interact with the amazing teams that keep our resources running and our users productive.

Thank you for taking the time to read and learn about our efforts in this year's science report.

# ARGONNE LEADERSHIP COMPUTING FACILITY

The ALCF provides supercomputing resources and expertise to help the research community accelerate the pace of discovery and innovation.

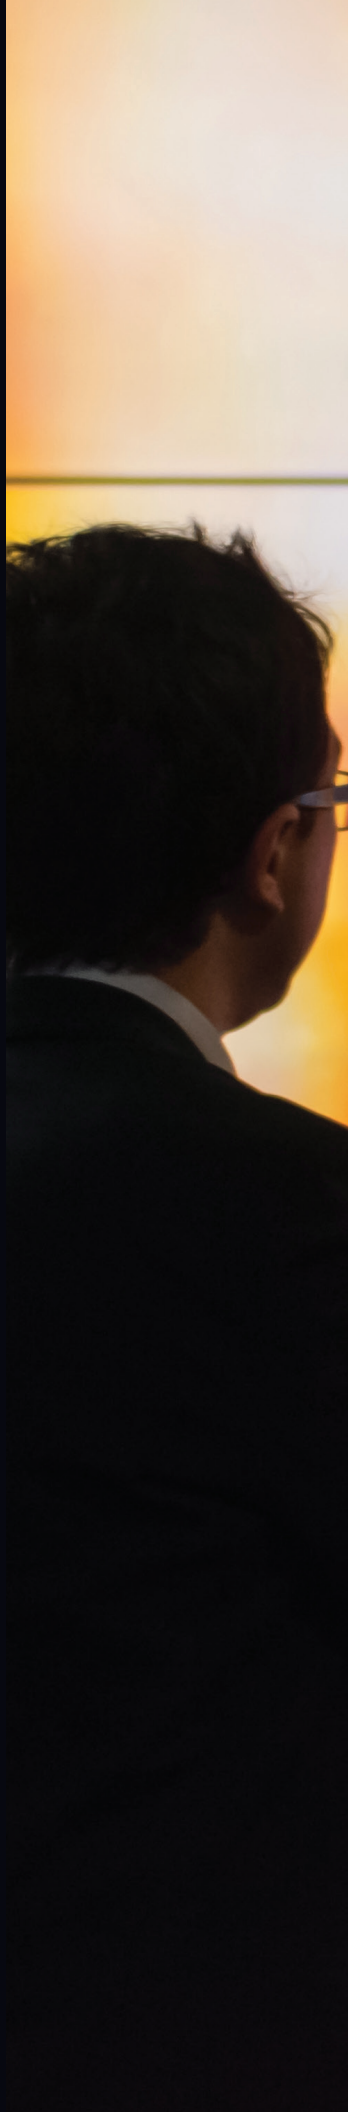

Image caption: ALCF opens its doors to visiting researchers, workshop attendees, and the general public for facility tours and hands-on demonstrations.

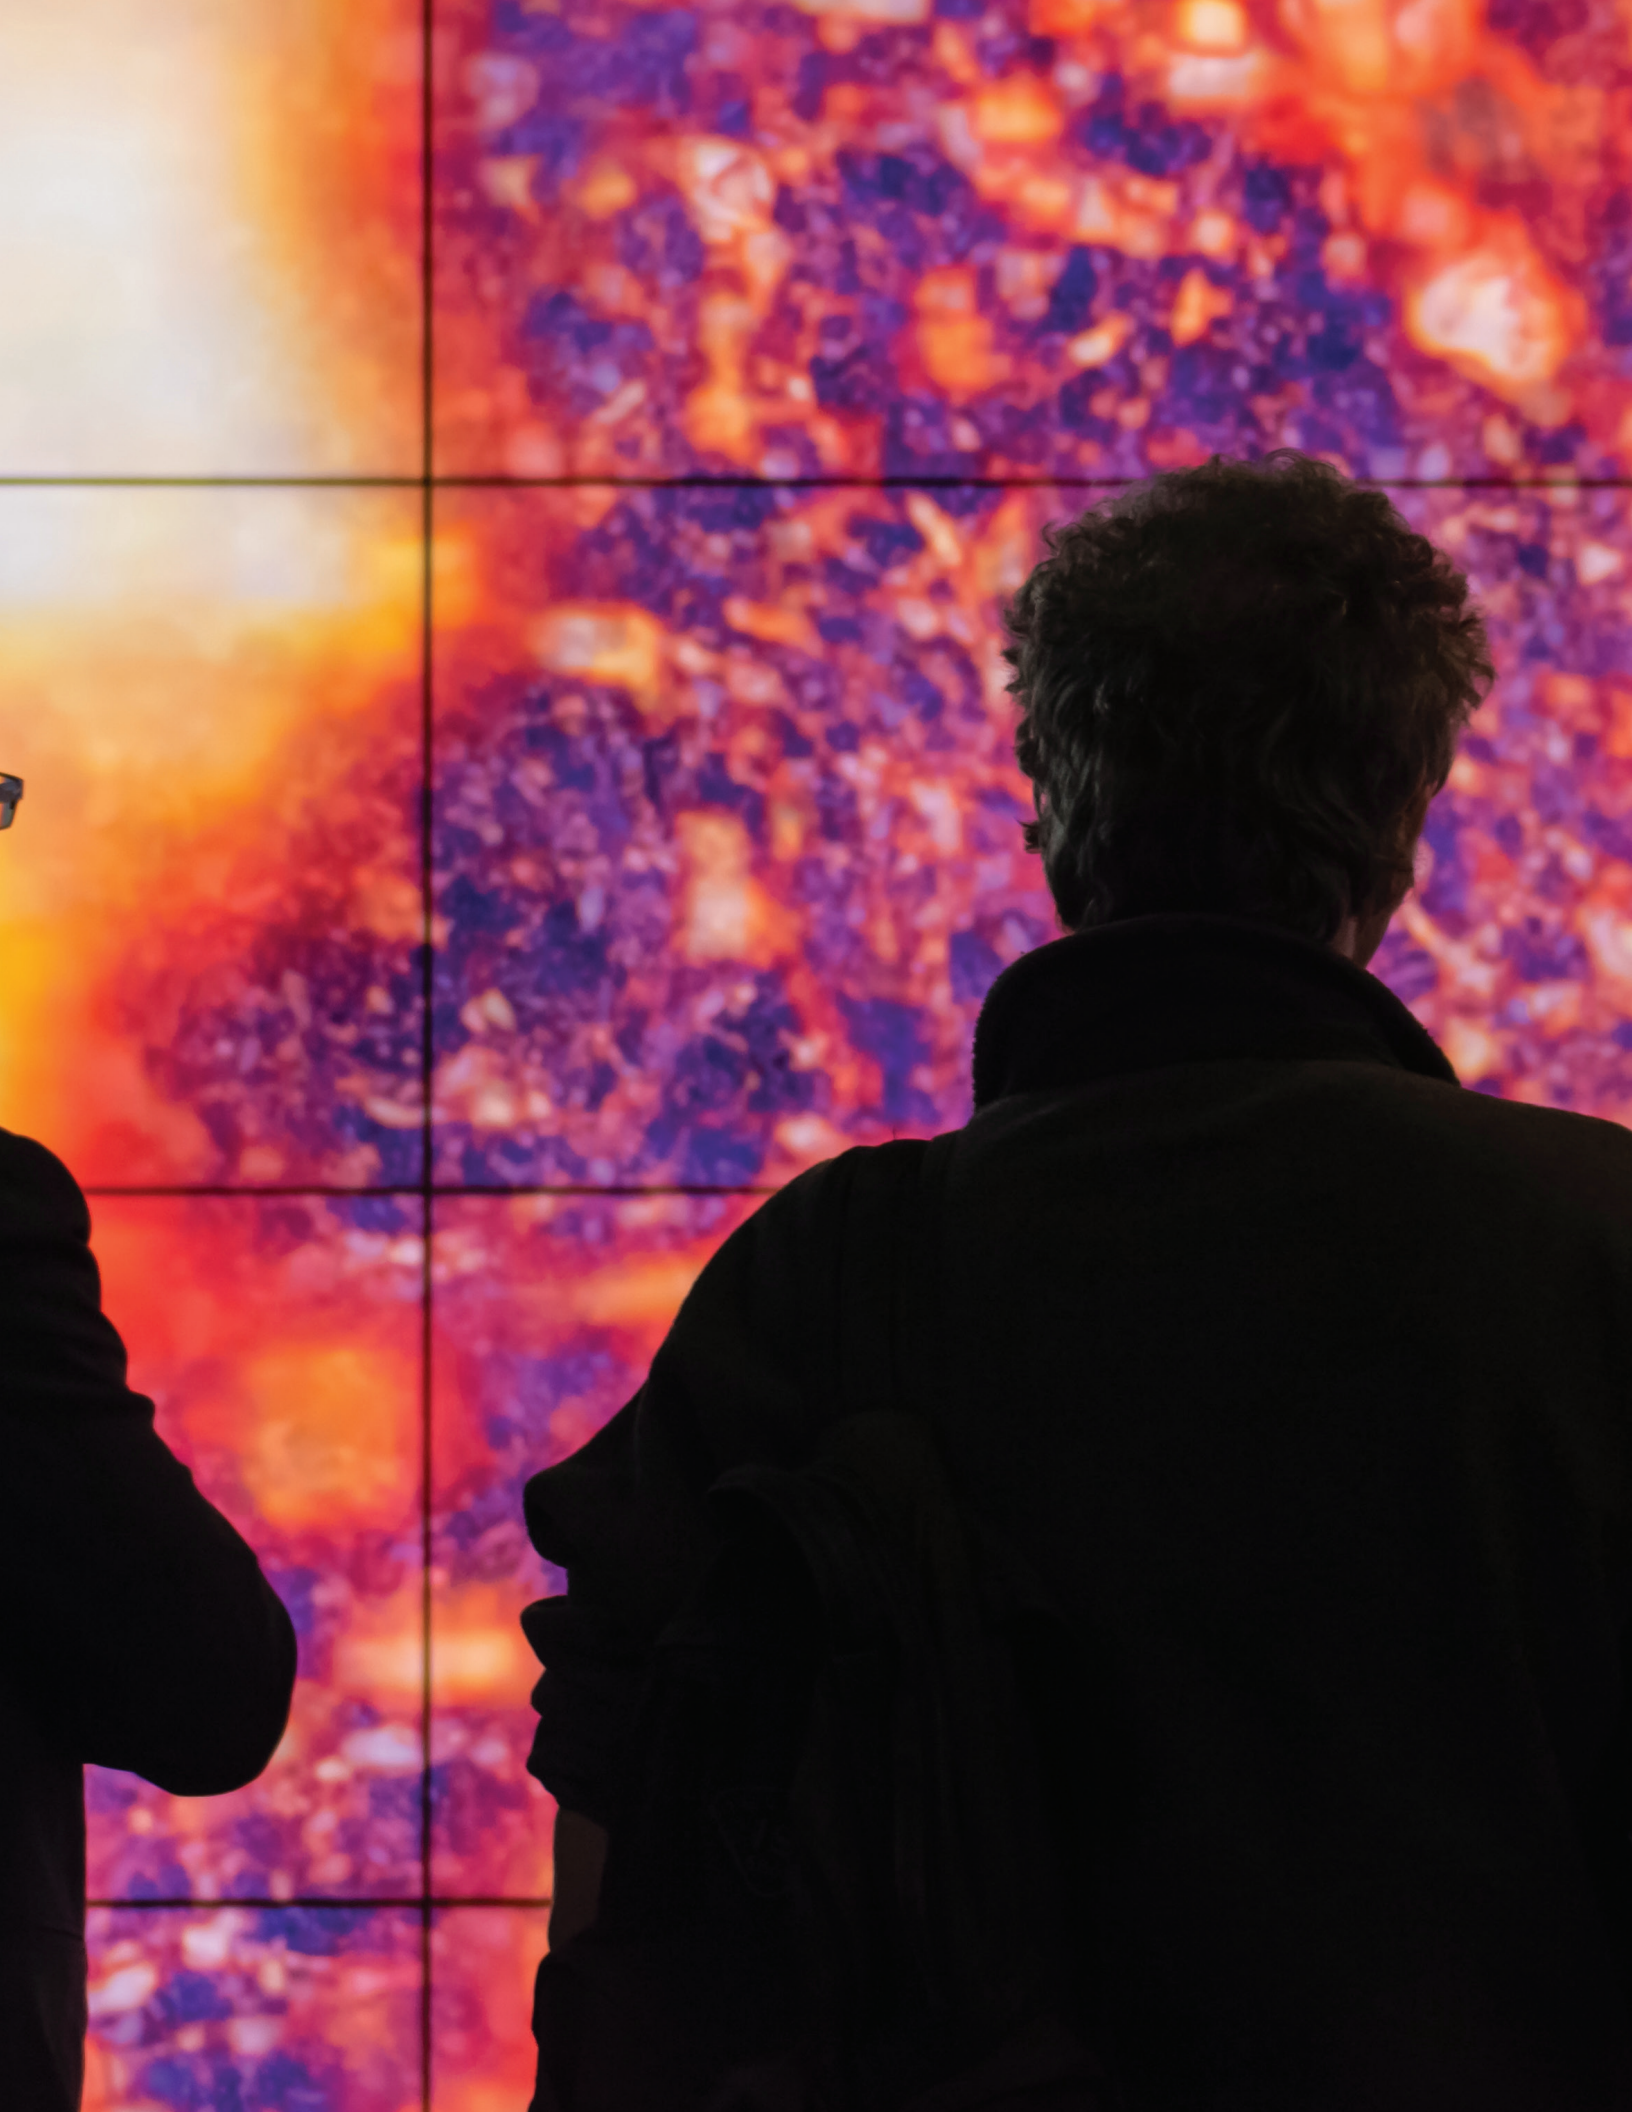

# Our Facility

The Argonne Leadership Computing Facility was established at Argonne National Laboratory in 2006 to provide state-of-the-art computing capabilities to the research community to pursue major discoveries and innovations through open science.

ALCF computing resources are 10 to 100 times more powerful than systems typically used for scientific research and are open to researchers in academia, industry, and government laboratories. The ALCF is supported by the U.S. Department of Energy Office of Science, Advanced Scientific Computing Research Program.

# Our People

**Operations**
The Operations Team manages and supports all ALCF computing systems, network infrastructure, storage, and systems environments, ensuring that users have stable, secure, and highly available resources to pursue their scientific goals. This team also creates accounts for users, provides technical support to research teams, and generates documentation to communicate policies and procedures to the user community.

**Science**
The Science Team works directly with users to maximize and accelerate their research efforts on the facility's computing resources. With multidisciplinary domain expertise, a deep knowledge of the ALCF computing environment, and experience with programming methods and community codes, the Science Team ensures that users are scientifically productive.

**Outreach**
The Outreach Team communicates the work conducted at the ALCF to a wide range of audiences through reports, promotional materials, science highlights, tours, and educational activities. This team also connects the ALCF with industry, coordinates user training events, and conducts user surveys.

# Our Resources


Theta, the ALCF's newest production system, is based on the second-generation Intel Xeon Phi processor and Cray's high-performance computing software stack.

ALCF resources include advanced computing capabilities, visualization hardware, software, and network infrastructure to meet the needs and priorities of its users.

## Production Systems

**Theta**

Theta is a 9.65-petaflops Cray XC40 system based on Intel processors and Cray interconnect technology, a new memory architecture, and a Lustre-based parallel file system, all integrated by Cray's HPC software stack.

**Mira**

Mira is a 10-petaflops IBM Blue Gene/Q machine equipped with 786,432 processors, 768 TB of memory, and IBM's 5D torus interconnect.

**Cooley**

Cooley has a total of 126 compute nodes; each node has 12 CPU cores and one NVIDIA Tesla K80 dual-GPU card. Cooley is coupled to ALCF's production supercomputing environment and is used to analyze and visualize the data produced by these systems.

## Support Resources

**Cetus**

Cetus is an IBM Blue Gene/Q and shares the same software environment and file systems as Mira. Cetus is primarily used to debug problems.

**Vesta**

Vesta is an IBM Blue Gene/Q used for testing and development.

**Iota**

Iota is a Cray XC40 used for testing and development.

**Disk Storage**

Mira has access to ~27 PB of GPFS file system capacity with performance of 240 GB/s on the largest file system (19 PB). Theta has access to ~18 PB of GPFS/Lustre file system capacity; 9 PB is GPFS and 9.2 PB is Lustre. The Lustre file system performance was benchmarked at ~240 GB/s.

**Tape Storage**

The ALCF has three 10,000-slot libraries using LTO 6 tape technology. The LTO tape drives have built-in hardware compression for an effective capacity of 36–60 PB.

**Networking**

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.

## Testbeds

Through Argonne's Joint Laboratory for System Evaluation, the ALCF provides access to next-generation hardware and software to explore low-level experimental computer and computational science, including operating systems, messaging, compilers, benchmarking, power measurements, I/O, and new file systems. These include:

| | |
|---|---|
| **Intel Xeon Phi Knights Landing Cluster** | **AppliedMicro X-C1 Server Development Kit Plus** |
| **Cray Urika-GX Analytics Platform** | **NVIDIA DGX-1** |
| **IBM Power System S822LC** | **IBM Elastic Storage Server GL6** |

# Accessing ALCF Resources

As a national user facility dedicated to open science, any researcher in the world with a large-scale computing problem can apply for time on ALCF computing resources.

Researchers gain access to ALCF systems for computational science and engineering projects, typically with awards of millions of core-hours, through competitive, peer-reviewed allocations programs supported by the DOE and Argonne.

The ALCF also hosts competitive, peer-reviewed application programs designed to prepare key scientific applications and innovative computational methods for the architecture and scale of ALCF supercomputers.

## Application Programs

### ADSP

The ALCF Data Science Program (ADSP) provides allocations on ALCF resources for big data projects that require the scale and performance of leadership computing resources. ADSP projects focus on developing and improving data science techniques and can either achieve a science goal or implement a technology needed to support data science.

### ESP

As part of the process of bringing a new supercomputer into production, the ALCF conducts an Early Science Program (ESP) to prepare applications for the architecture and scale of the new system. The ESP projects represent a typical system workload at the ALCF and cover key scientific areas and numerical methods.

## Allocation Programs

The following chart shows the breakdown of the three allocation programs on Mira.

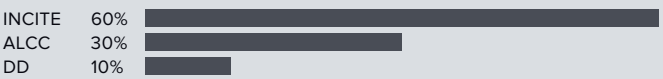


### INCITE

The Innovative Novel Computational Impact on Theory and Experiment (INCITE) program invites proposals for large-scale, computationally intensive research projects to run at the ALCF. The INCITE program awards sizable allocations on some of the world's most powerful supercomputers to address grand challenges in science and engineering.

### ALCC

The ASCR Leadership Computing Challenge (ALCC) program allocates computational resources at the ALCF for projects directly related to the DOE's energy mission, for national emergencies, or for broadening the community of researchers capable of using leadership computing resources.

### DD

Director's Discretionary (DD) projects are dedicated to leadership computing preparation, INCITE and ALCC scaling, and application performance to maximize scientific application efficiency and productivity on leadership computing platforms.

**2017 INCITE by Domain**
3.53 Billion Core-Hours

| | | |
|---|---|---|
| 01 | Biological Sciences | 3% |
| 02 | Chemistry | 7% |
| 03 | Computer Science | 1% |
| 04 | Earth Science | 18% |
| 05 | Engineering | 13% |
| 06 | Materials Science | 25% |
| 07 | Physics | 33% |

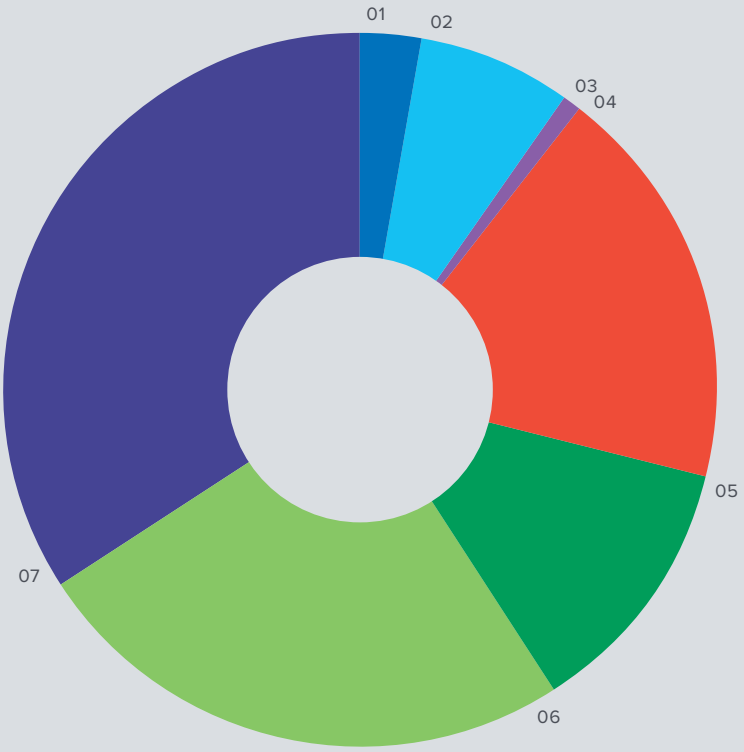


**2017 ALCC by Domain**
1.75 Billion Core-Hours

| | | |
|---|---|---|
| 01 | Biological Sciences | 3% |
| 02 | Chemistry | 5% |
| 03 | Computer Science | 9% |
| 04 | Earth Science | 9% |
| 05 | Engineering | 9% |
| 06 | Materials Science | 19% |
| 07 | Energy Technologies | 11% |
| 08 | Physics | 34% |
| 09 | Mathematics | 1% |

Note: ALCC data are from calendar year 2017.

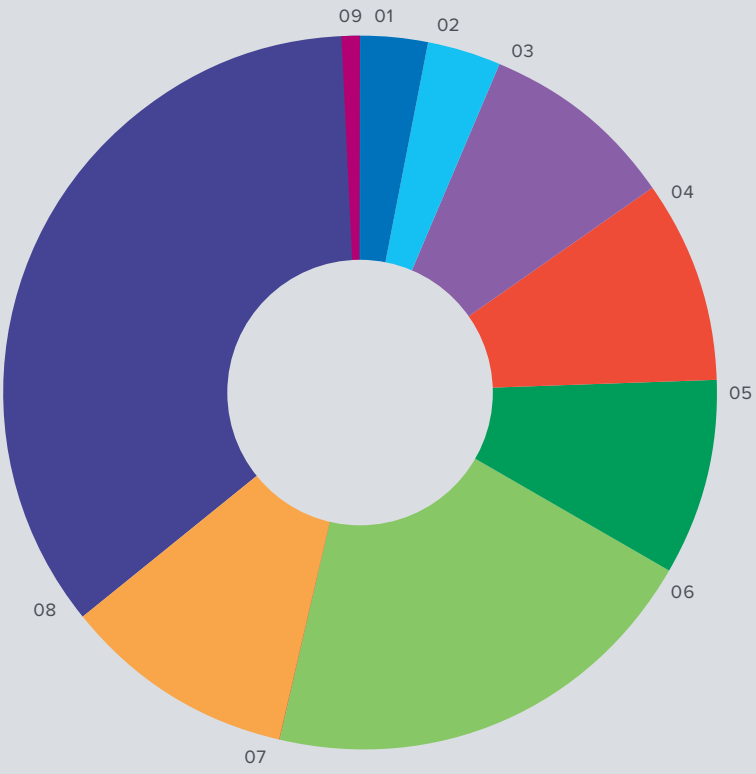

# REDEFINING HPC

The ALCF helps researchers achieve significant breakthroughs in science and engineering, while also advancing the nation's efforts to develop future exascale computing systems.

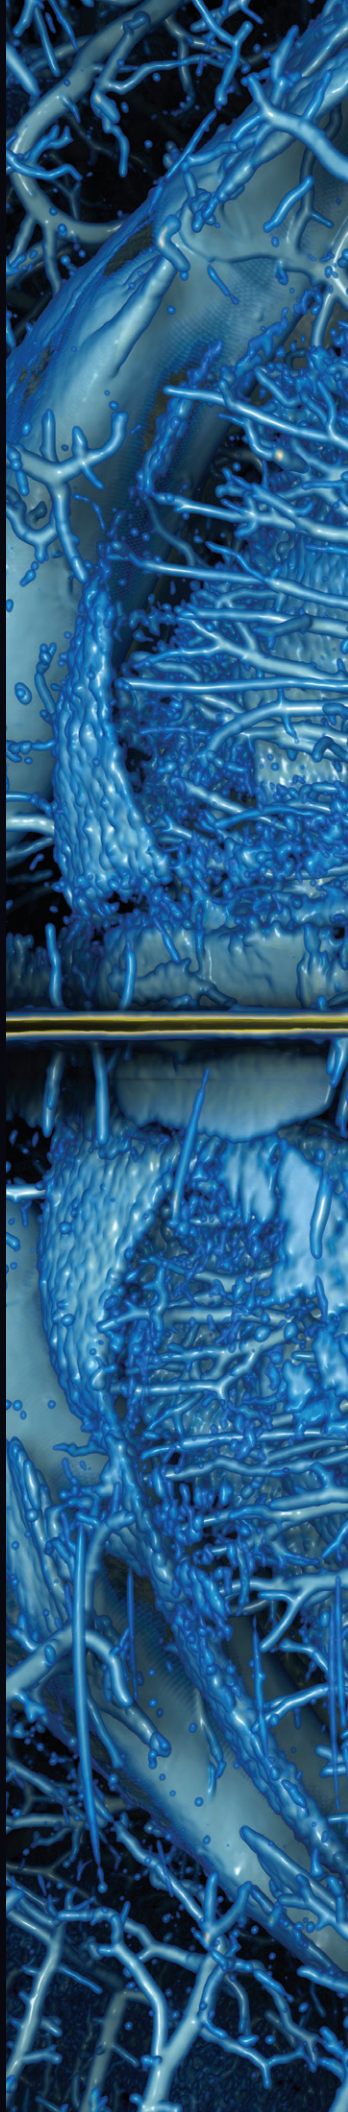

Image caption: Reconstructed cerebrovascular network of the mouse brain using synchrotron microtomography.
*Image credit: Joseph A. Insley, Argonne National Laboratory*

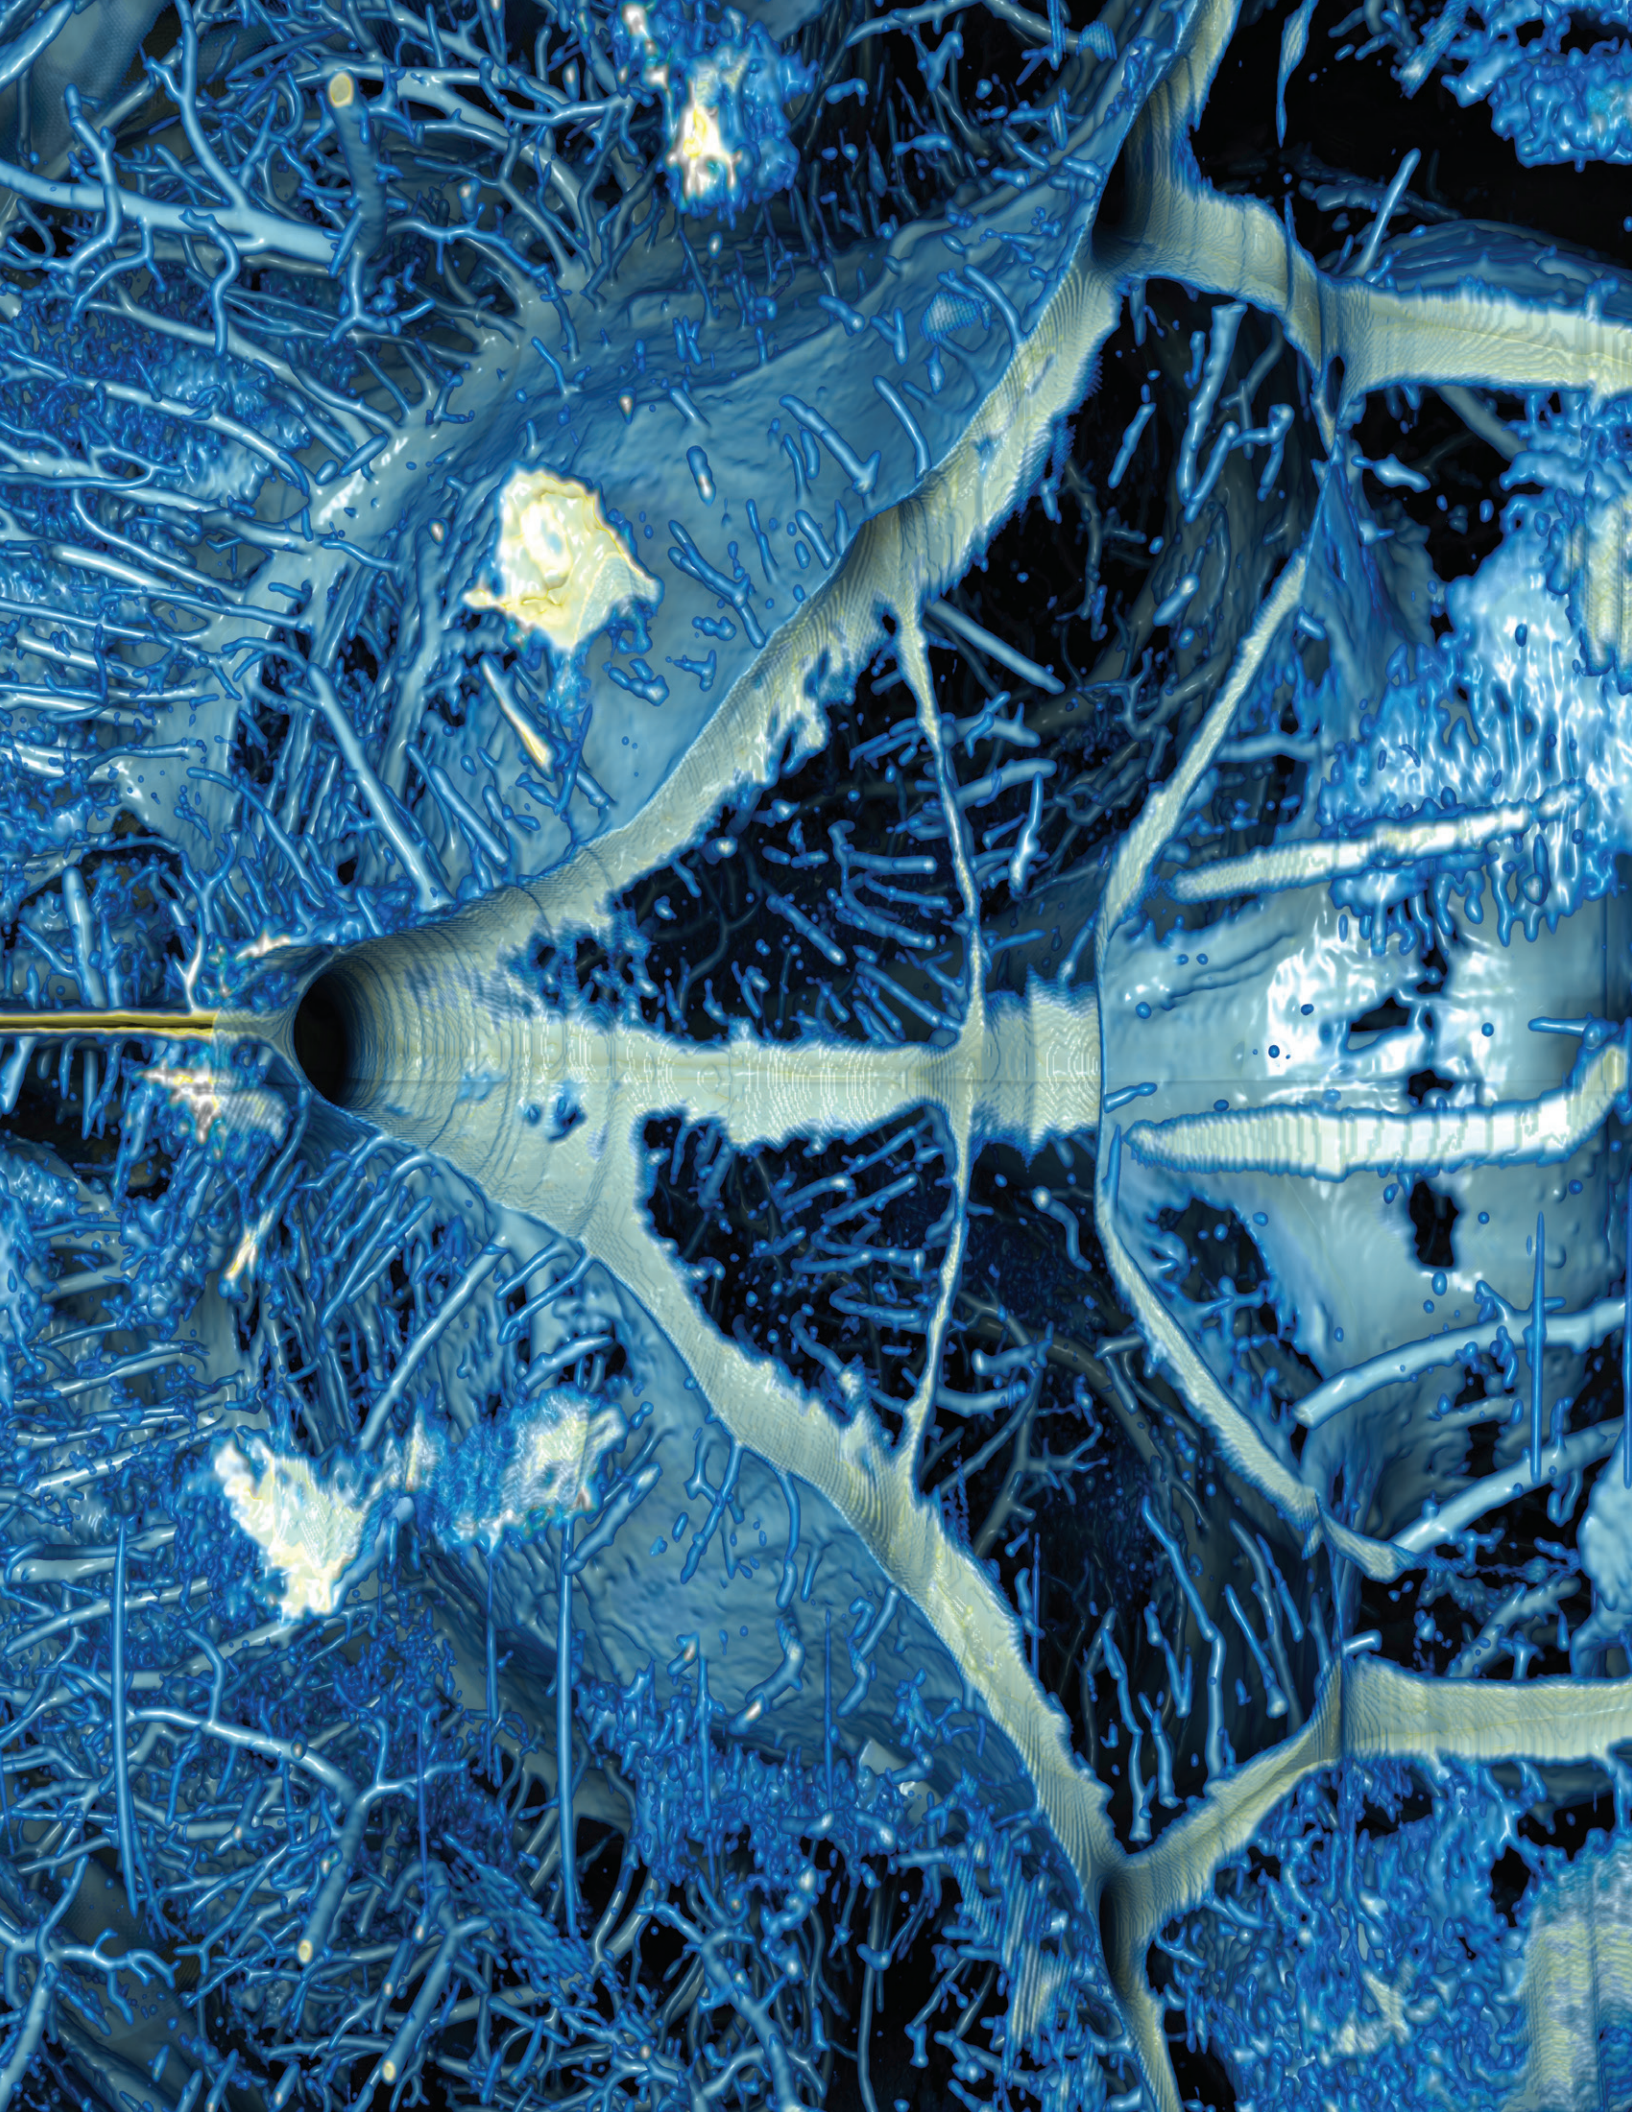

# Preparing for Exascale

## The ALCF is part of a broader community working to achieve a capable exascale computing ecosystem for scientific discoveries.

Every time computing power increases by large factors, new benefits open up to the scientific research community. The benefits of exascale computing—computing capability that can achieve at least a billion *billion* operations per second—is the next waypoint in high-performance computing. The value of exascale computing is primarily in the applications that it will enable, which include fundamental science, industrial design, and health.

Achieving speeds 50 to 100 times more powerful than the nation's fastest supercomputers in use today is a significant challenge. To be useful to a wide spectrum of applications, in addition to peak speed, supercomputers need to have large memories and the ability to store and read vast quantities of data at high speed, and a software environment that facilitates its efficient and productive use.

ALCF researchers lead and contribute to several strategic activities, including DOE's Exascale Computing Project (ECP), that aim to push the boundaries of what's possible in computational science and engineering.

In addition to enabling traditional supercomputing applications at the exascale, the ALCF seeks to contribute to the emerging convergence of supercomputing, big data analytics, and machine learning across a wide variety of science and engineering domains and disciplines to address problems of national significance.

**Exascale Computing Project**
The DOE launched the ECP in 2016 as part of a coordinated effort to achieve the next generation of HPC and to accelerate scientific discovery and economic competitiveness.

This seven-year project aims to advance all aspects of the computing ecosystem: application development, software technology, hardware technology, exascale systems, and workforce development. The long-term benefit of this approach will be a thriving ecosystem of capable, U.S.-based, exascale computing products, and scientists and engineers with expertise in using them.

The ECP is a collaboration among the nation's six premier computing labs with funding from the DOE Office of Science and the National Nuclear Security Administration.

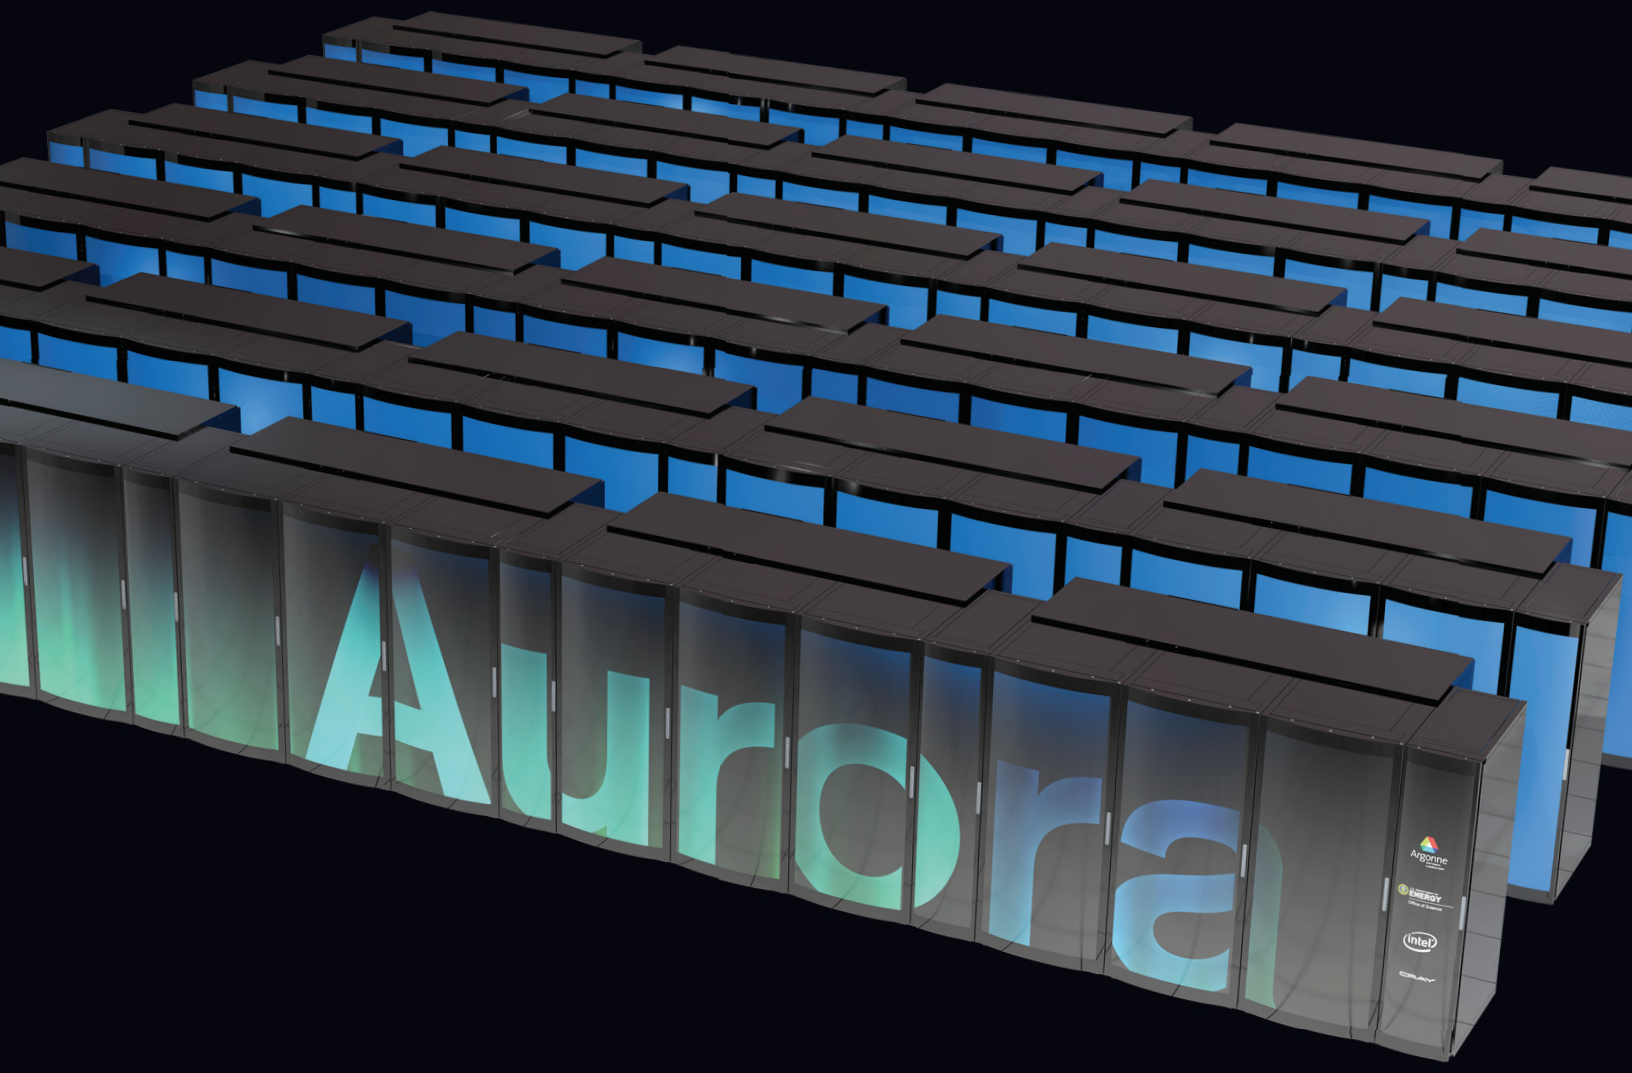


AURORA

**In 2015, the DOE announced major investments in the nation's leadership computing program and put several computing centers on a fast track to exascale computing capabilities. Under the joint Collaboration of Oak Ridge, Argonne, and Lawrence Livermore (CORAL) initiative, the ALCF will build a next-generation supercomputer, known as Aurora.**

# Targeting Big Data

## The ALCF Data Science Program is a pioneering initiative aimed at improving computational methods to advance data-driven discoveries across all scientific disciplines.

In 2016, the ALCF launched an initiative to explore new ways to foster data-driven discoveries, with an eye to growing a new community of HPC users. The ALCF Data Science Program (ADSP), the first of its kind in the leadership computing space, targets users with 'big data' science problems and provides millions of core-hours on ALCF resources, along with staff support, a dedicated postdoc, and training.

The outcomes of these projects—whether to achieve a specific science goal, or to implement a specific technology needed to support data science—focus on proving, and improving, data science techniques that can exploit powerful computing resources to root out significant findings from massive datasets being generated by large-scale simulations, observations, and experiments.

**ADSP projects employ leadership-computing systems and infrastructure to explore, prove, and improve a wide range of data science techniques:**

- uncertainty quantification
- statistics
- machine learning
- deep learning
- databases
- pattern recognition
- image processing
- graph analytics
- data mining
- real-time data analysis
- complex and interactive workflows

# ALCF Data Science Program Projects

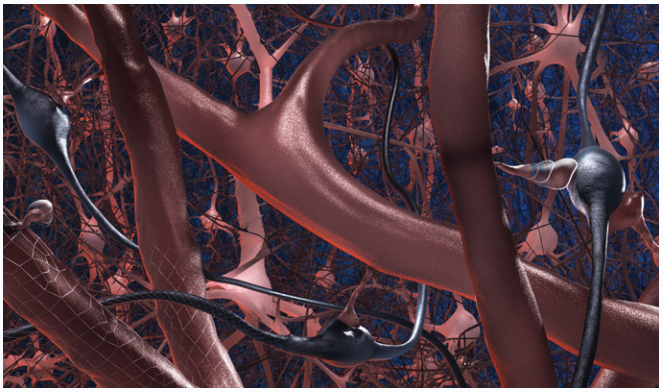

ADSP Project: Leveraging Non-Volatile Memory, Big Data, and Distributed Workflow Technology to Leap Forward Brain Modeling. Fabien Delalondre, Blue Brain, EPFL
*Image credit: Nicolas Antille, Blue Brain, EPFL*

### Advancing the Scalability of LHC Workflows to Enable Discoveries at the Energy Frontier

PI NAME Taylor Childers
INST Argonne National Laboratory

This project is developing an end-to-end workflow to manage the data motion and payload execution of ATLAS detector simulations on ADSP resources. This work will increase the analysis reach of LHC scientists and lead to new discoveries in particle physics.

### Data-Driven Molecular Engineering of Solar-Powered Windows

PI NAME Jacqueline M. Cole
INST University of Cambridge

This project aims to develop a new material-by-design methodology by using natural language processing, machine learning, and data mining in conjunction with large-scale simulation and experiments. This synergistic approach will enable the discovery of new light-absorbing dye molecules for the development of solar-powered windows that have the potential to power buildings in an entirely energy-sustainable fashion.

### Large-Scale Computing and Visualization on the Connectomes of the Brain

PI NAME Doga Gursoy
INST Argonne National Laboratory

This project will enable extreme-scale, data-centric pipelines for brain science. The scalable workflows, focused on analysis and visualization of experimental data, will help researchers gain invaluable knowledge about disease models, such as Alzheimer's, autism spectrum disorder, and many others. Additionally, the insights gleaned will enable transformative advances in neuromorphic computing.

### Leveraging Non-Volatile Memory, Big Data, and Distributed Workflow Technology to Leap Forward Brain Modeling

PI NAME Fabien Delalondre
INST Blue Brain Project, EPFL

This project is developing a computational workflow that will integrate new data storage paradigms, run times, and big data technology. This work will enable the simulation and analysis of brain tissue models at unprecedented scales, paving the way for future brain research and neuroscience breakthroughs.

# Delivering on Day One

## The ALCF's Early Science Program is designed to help prepare scientific applications for the architecture and scale of the facility's new supercomputers.

The Leadership Computing Facility and other DOE-supported computing centers deploy some of the most powerful computers and fastest networks in the world. These machines are largely unique, built in collaboration with leading HPC vendors, to offer greater processing power for less energy per flop.

The Early Science Program prepares key applications for each new supercomputer during the critical period between installation and production. The program, the first of its kind in the nation, provides researchers with early access to the new architecture and allows technical staff to prepare the machine to deliver science on day one.

Early Science teams, whose projects cover key scientific areas and numerical methods, work with ALCF staff and technical vendors to adapt their codes to the architecture and scale of the new machine. Collectively, the projects also represent a typical system workload at the ALCF, allowing vendors and staff to solidify the libraries and infrastructure as the project teams run real science campaigns.

Each Early Science Program instance is slightly different: Mira's consisted of 16 science projects and consumed nearly 2 billion core-hours in the few months between acceptance and production. Theta's program supported science projects with pre-production allocations for science runs, and six code development projects. Aurora's, now underway supports 10 science projects and also encourages application code portability among heterogeneous architectures.

# Aurora Early Science Program Projects

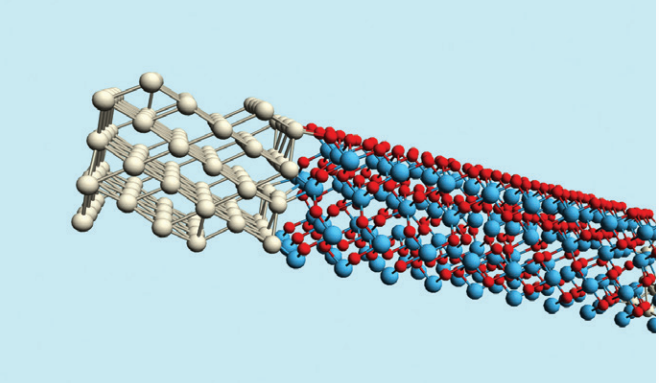

Aurora ESP Project: Extending Moore's Law Computing with Quantum Monte Carlo. Anouar Benali, Argonne National Laboratory.
*Image credit: Olle Heinonen, Argonne National Laboratory*

The Aurora program provides 10 diverse science teams with dedicated time on Theta to do code development work and to conduct real science runs targeted at a new system that will deliver at least 20 times the computational performance of Mira.

The project teams attend training sessions and work closely with staff and vendors Intel and Cray to adapt their high-performance codes to Aurora's future architecture. The teams also have access to training and hardware at the Oak Ridge Leadership Computing Facility and the National Energy Research Supercomputing Center to promote application code portability among heterogeneous architectures.

### Benchmark Simulations of Shock-Variable Density Turbulence and Shock-Boundary Layer Interactions with Applications to Engineering Modeling

PI NAME Sanjiva Lele
INST Stanford University
CODE SU2, PadeOps

What do inertial confinement fusion (ICF) and supersonic aircraft have in common? Both involve the flow of gases in extreme conditions, including shock waves and turbulence. This project advances scientific understanding of variable density turbulence and mixing, including shock interactions and near-wall effects. These apply to the mixing of the fuel capsule surface with the imploding plasma in ICF, and shock interaction with fuel streams in a supersonic jet engine as a way to improve combustion.

### Design and Evaluation of High-Efficiency Boilers for Energy Production Using a Hierarchical V/UQ Approach

PI NAME Martin Berzins
INST The University of Utah
CODE Uintah

This project simulates and evaluates the design of a next-generation, 1,000 MW advanced coal boiler that is ultra supercritical. In a coal-fired power plant, this design promises to reduce the boiler footprint 50 percent, saving costs and improving efficiency (53 percent efficiency, compared with traditional-boiler 35 percent efficiency), and also reducing $CO_2$ emissions by 50 percent relative to a traditional boiler. Simulations on Aurora using the Uintah asynchronous many-task software incorporates validation and uncertainty quantification (V/UQ), predicting thermal efficiency with uncertainty bounds constrained by observed data.

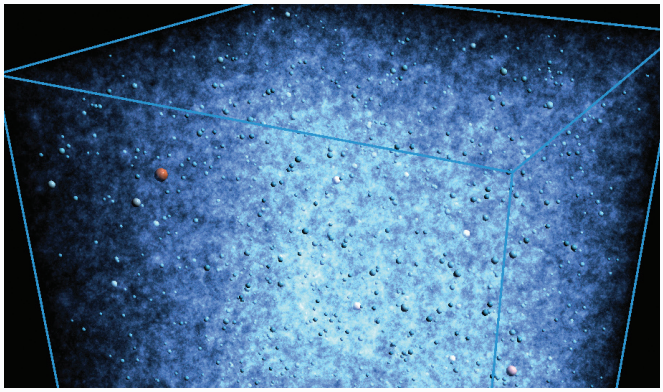

Aurora ESP Project: Extreme-Scale Cosmological Hydrodynamics. Katrin Heitmann, Argonne National Laboratory
*Image credit: Silvio Rizzi and Joseph A. Insley, Argonne National Laboratory*

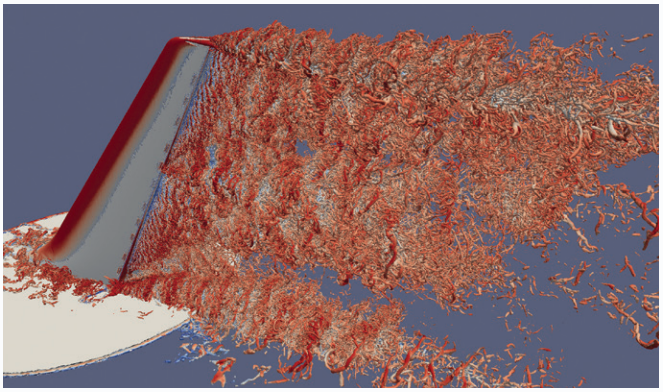

Aurora ESP Project: Extreme-Scale Unstructured Adaptive CFD. Kenneth Jansen, University of Colorado Boulder
*Image credit: Kenneth Jansen, University of Colorado Boulder; Michel Rasquin, Cenaero and University of Colorado Boulder*

### Extending Moore's Law Computing with Quantum Monte Carlo

PI NAME Anouar Benali
INST Argonne National Laboratory
CODE QMCPACK

For decades, massively parallel supercomputers have reaped the benefits—predicted by Moore's Law—of the relentless increase in density of components on chips, which also rapidly improved performance of PCs and smartphones. In this project, supercomputing returns the favor by attacking a fundamental materials problem impacting the latest and future chips: electrical current leakage through the $HfO_2$-silicon interface. $HfO_2$ is used widely as a dielectric in Si-CMOS chips like the Aurora CPUs. Simulating this problem with the highly accurate quantum Monte Carlo method is only now becoming computationally possible with supercomputers like Aurora.

### Extreme-Scale Cosmological Hydrodynamics

PI NAME Katrin Heitmann
INST Argonne National Laboratory
CODE HACC

This project simulates large fractions of the universe, including not only gravity acting on dark matter, but also baryons (which make up visible matter such as stars) and gas dynamics using a new, smoothed particle hydrodynamics method. These simulations are deeply coupled with guiding and interpreting observations from present and near-future cosmological surveys.

### Extreme-Scale Unstructured Adaptive CFD

PI NAME Kenneth Jansen
INST University of Colorado Boulder
CODE PHASTA

This project uses unprecedented high-resolution fluid dynamics simulations to model dynamic flow control over airfoil surfaces at realistic flight conditions and to model bubbly flow of coolant in nuclear reactors. Synthetic jet actuators, tiny cavities with speaker-like diaphragms that alternately expel and intake air, can alter and control airflow across surfaces such as plane tail rudders, allowing much stronger force (turning force, for a rudder). The project uses multiphase flow modeling to simulate realistic reactor geometries with far more accuracy than today's state of the art, yielding valuable information on thermal management to improve the safety of existing light-water reactors and inform the design of next-generation systems.

### Free Energy Landscapes of Membrane Transport Proteins

PI NAME Benoît Roux
INST The University of Chicago and Argonne National Laboratory
CODE NAMD

Membrane transport protein molecules play key roles in cellular biology functions. This includes natural processes as well as drug delivery and drug resistance. How these "molecular devices" perform their function is extremely complex. The proteins move into dramatically different conformations in the process. Modeling the myriad possibilities with atomistic molecular dynamics, even using the best statistical approaches, is at the forefront of what's possible. These calculations on Aurora will advance that front.

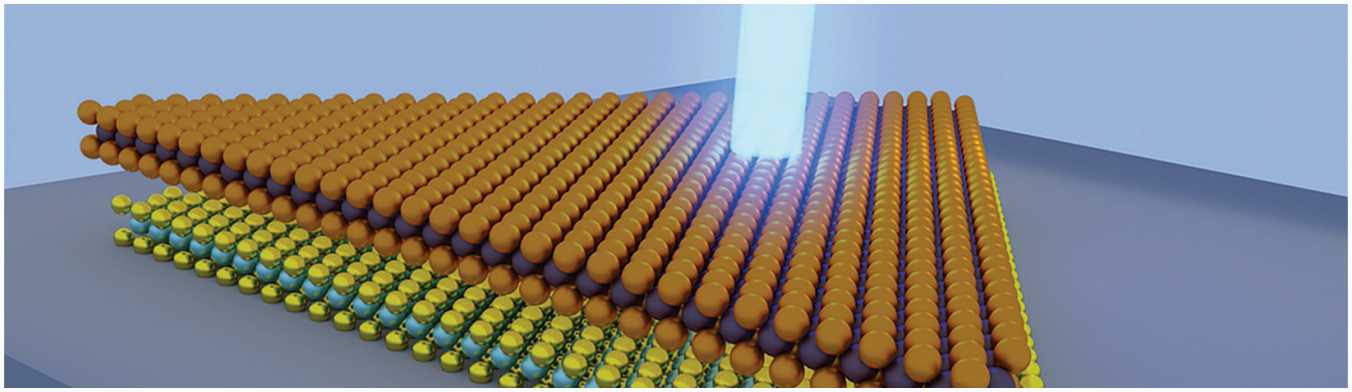

Aurora ESP Project: Metascalable Layered Materials Genome. Aiichiro Nakano, University of Southern California
*Image credit: Aravind Krishnamoorthy, University of Southern California*

### High-Fidelity Simulation of Fusion Reactor Boundary Plasmas

PI NAME C.S. Chang
INST Princeton Plasma Physics Laboratory
CODE XGC

The behavior of plasma at the outer edge of a tokamak fusion reactor is critically important to the success of future fusion reactors such as ITER, now under construction in France. Misbehavior at the edge can lead to disruptions bombarding a small area of the divertor plates—metal structures at the bottom of the tokamak designed to absorb ejected heat—at levels beyond which the divertor material can withstand. This project uses particle simulations of the plasma, including impurities and the important magnetic field geometry at the edge, to predict behavior of ITER plasmas and to help guide future experimental parameters.

### Lattice Quantum Chromodynamics Calculations for Particle and Nuclear Physics

PI NAME Paul Mackenzie
INST Fermilab
CODE USQCD codes

This project delivers calculations urgently needed by experimental programs of high energy and nuclear physics, based on the computational methods of lattice quantum chromodynamics (lattice QCD). QCD embodies our most fundamental understanding of the strong nuclear force and associated particles, a key component of the more general Standard Model of particle physics. In high energy physics, lattice calculations are required to extract the fundamental parameters of the Standard Model (such as quark masses) from experiment. Evidence for physics beyond the Standard Model can be discovered if discrepancies are found between different methods for determining these parameters.

### Metascalable Layered Materials Genome

PI NAME Aiichiro Nakano
INST University of Southern California
CODE NAQMD, RMD

Functional materials, as the name implies, have behaviors useful in science and industry. There is great interest today in engineering materials to have desired behaviors. One approach involves stacking extremely thin layers of different materials to achieve a complex molecular interplay throughout the stack. The resulting behavior of the stack cannot be explained by traditional theories and can only be predicted by directly simulating the layers as collections of molecules interacting with each other. Massive quantum mechanical and reactive molecular dynamics simulations on Aurora will be validated by experiments on the same materials using a free-electron X-ray laser.

### NWChemEx: Tackling Chemical, Materials, and Biochemical Challenges in the Exascale Era

PI NAME Thomas Dunning
INST Pacific Northwest National Laboratory
CODE NWChemEx

The NWChemEx code is providing the framework for next-generation molecular modeling in computational chemistry and for implementing critical computational chemistry methods. This project applies it to two problems in the development of advanced biofuels: design of feedstock for efficient production of biomass; and design of new catalysts for converting biomass-derived chemicals into fuels.

# Theta Early Science Program Projects

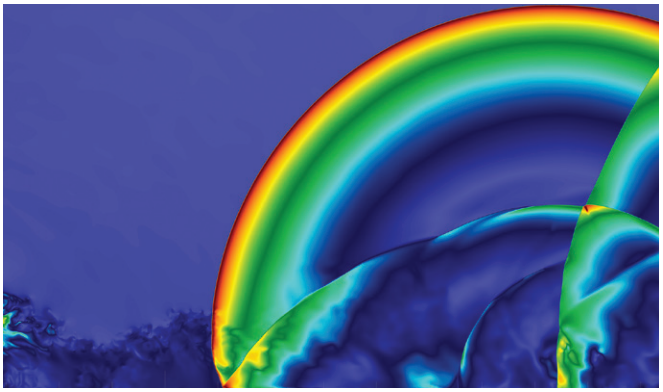

Theta ESP Project: Direct Numerical Simulations of Flame Propagation in Hydrogen-Oxygen Mixtures in Closed Vessels. Alexei Khokhlov, The University of Chicago
*Image credit: Marta Garcia and Joseph A. Insley, Argonne National Laboratory*

The Theta program awarded substantial allocations and support to six science projects and six code development projects. The Early Science teams also had access to Theta-generation hardware in Argonne's Joint Laboratory for System Evaluation for code optimization, testing, and debugging efforts.

Most of the science teams immediately reported significant performance gains on the new Intel-Cray architecture. Several of the projects also produced significant science results.

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

PI NAME Alexei Khokhlov
INST The University of Chicago
CODE HSCD

The goal of the project is to perform direct numerical simulations (DNS) of flame propagation experiments in hydrogen-oxygen mixtures in spherical containers. DNS of the entire experimental apparatuses will be carried out to bypass the traditional steps of extracting the laminar flame speed values from the experimental data and to eliminate the associated uncertainty of the measured laminar flame velocity. Instead, the simulations will directly validate the existing chemical reaction mechanisms, and thermal and species diffusion models. They will also provide a combined sensitivity analysis of flame acceleration in hydrogen-oxygen mixtures on reaction constants and microscopic transport parameters of the mixture at varying ambient pressure and temperature—crucial for carrying out quantitative first-principles predictions of the flame acceleration and the deflagration-to-detonation transition (DDT) in hydrogen mixtures.

IMPACT First-principles understanding and quantitative prediction of flame acceleration and DDT in hydrogen are important for the industrial and public safety of hydrogen fuels and certain types of water-cooled nuclear reactors.

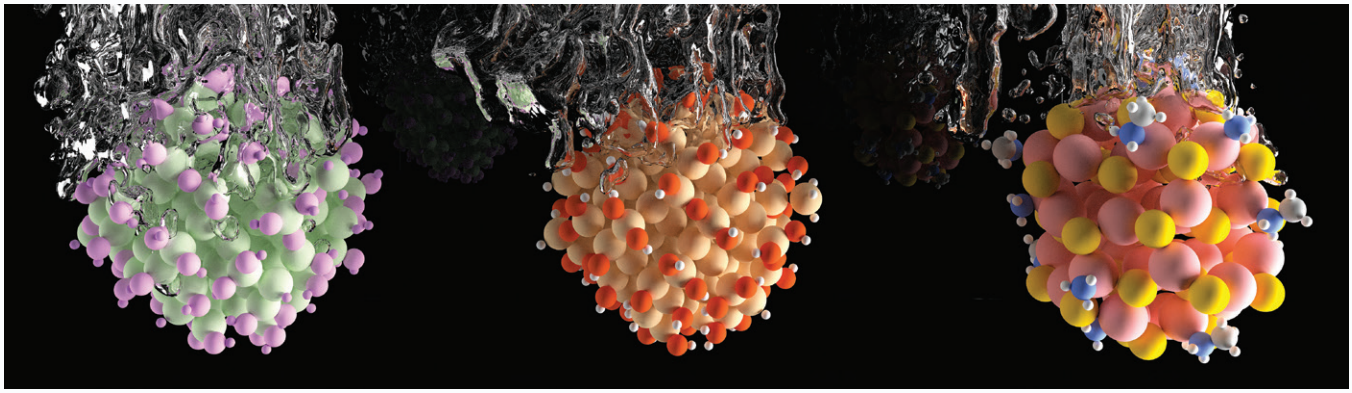

Theta ESP Project: First-Principles Simulations of Functional Materials for Energy Conversion. Giulia Galli, The University of Chicago and Argonne National Laboratory
*Image credit: Nicholas Brawand, The University of Chicago*

### First-Principles Simulations of Functional Materials for Energy Conversion

PI NAME Giulia Galli
INST The University of Chicago and Argonne National Laboratory
CODE Qbox, WEST

This project investigates the properties of nanostructured materials for use in solar and thermal energy conversion devices. For solar cell materials, the team is focusing on group IV and III-V nanoparticles with a variety of ligands that are currently being tested experimentally. For thermoelectric and solar-thermal applications, the team is focusing on silicon-based clathrates, recently shown to exhibit promising solar and thermoelectric properties, and on solar perovskites.

IMPACT This project is focusing on materials exhibiting complex structures on multiple length scales. Researchers hope to predict the electronic and thermal properties of these materials using a theoretical framework that combines *ab initio* molecular dynamics with accurate electronic structure methods.

### Free Energy Landscapes of Membrane Transport Proteins

PI NAME Benoît Roux
INST The University of Chicago and Argonne National Laboratory
CODE NAMD

Membrane transport proteins are a unique class of macromolecular biological systems that play a critical role in numerous aspects of cell function. A distinct feature of those molecular processes is the existence of large conformational changes, often in response to, or in concert with, the surrounding chemical environment. A complete understanding of how these proteins carry out their function will thus rely heavily on characterizing these transitions. This project will use atomistic computational models to determine the free energy landscape of large-scale conformational transitions for two P-type membrane proteins. The necessary simulations will build upon existing structural data from crystallography and provide meaningful connections to measurements on real, dynamical systems.

IMPACT An atomistic picture of membrane transport proteins is a critical component for understanding a broad range of biological functions. This work will utilize computational models to provide both detailed visualizations of large protein motions, as well as quantitative predictions into the energetics of these processes.

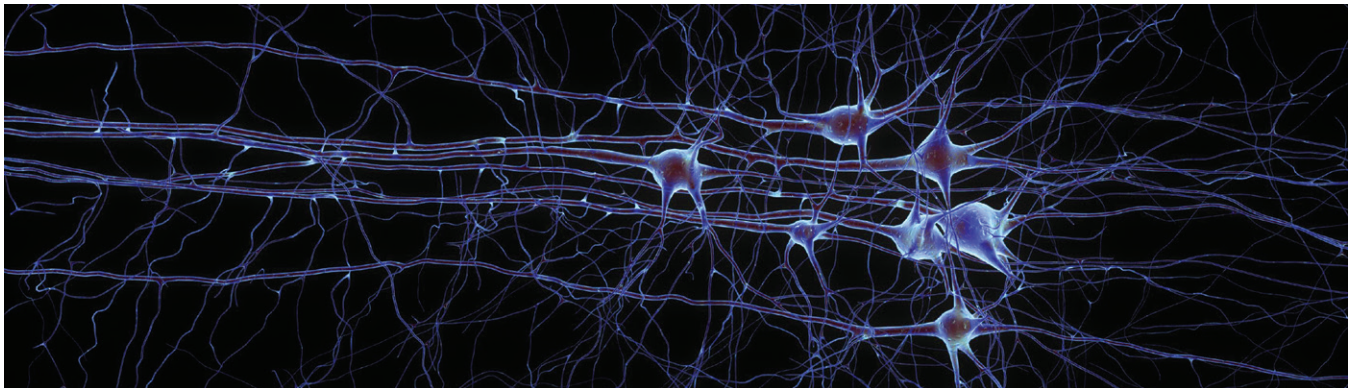

Theta ESP Project: Large-Scale Simulation of Brain Tissue. Fabien Delalondre, Blue Brain Project, EPFL
*Image credit: Blue Brain Project, EPFL*

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

PI NAME Fabien Delalondre
INST Blue Brain Project, EPFL
CODE CoreNeuron

This project will be used to advance four critical scientific applications. The first consists of performing simulations on the order of a few neocortical columns, each including 31,000 neurons and about 37 million synapses. The second and third rely on the simulation of the electrical activity of the rodent somatosensory cortex with morphologically detailed neurons. The fourth consists of simulating the electrical activity of the largest possible brain model for several seconds of biological time.

IMPACT Creating data-driven brain tissue models is a novel method in the scientific repertoire. With this project, researchers are advancing simulation tools, widening the ability for scientific questions on brain plasticity, and triggering valuable discussions with the HPC community.

### Next-Generation Cosmology Simulations with HACC: Challenges from Baryons

PI NAME Katrin Heitmann
INST Argonne National Laboratory
CODE HACC

Today, and in the near future, simulating galaxy formation from first principles remains an intractable task. Therefore, researchers have to rely on so-called subgrid models that allow them to include a range of astrophysical effects in the simulations. The modeling approach relies heavily on empirical findings. With this project, the team has two aims: to further the understanding of astrophysical processes by confronting detailed simulations with the newest observations, and to enable reliable modeling of baryonic physics. The second aim, informed by the first, provides an approach to mitigate possible contamination of cosmological results due to baryonic effects. The researchers are exploring a range of well-motivated subgrid models and extracting observational signatures.

IMPACT By confronting new observations with sophisticated simulations, this research furthers the understanding of astrophysical processes on small scales. At the same time, this work disentangles these processes from fundamental physics and therefore helps mitigate one of the major sources of systematic uncertainties for upcoming cosmological surveys.

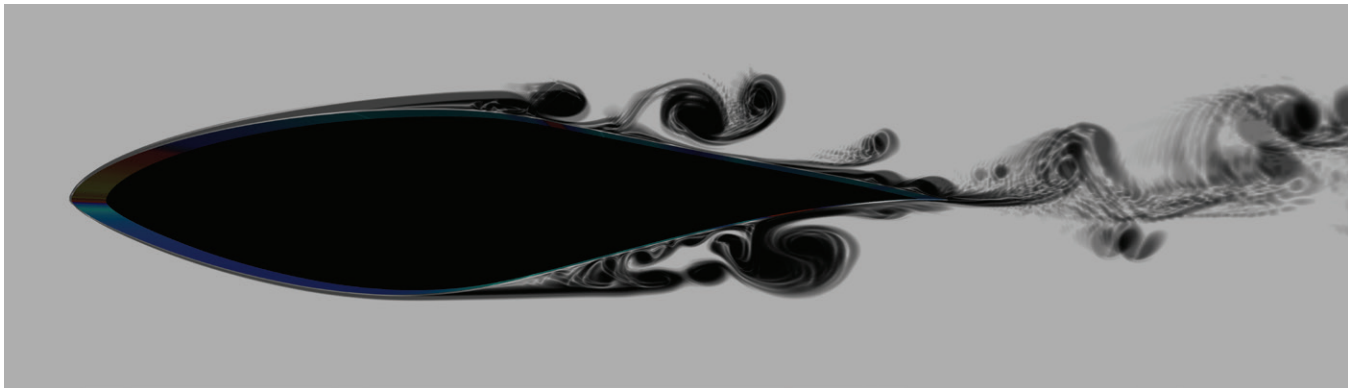

Scale-Resolving Simulations of Wind Turbines with SU2. Juan J. Alonso, Stanford University
*Image credit: Ramesh Balakrishnan, Argonne National Laboratory*

## Scale-Resolving Simulations of Wind Turbines with SU2

PI NAME Juan J. Alonso
INST Stanford University
CODE SU2

The purpose of this project is to create a large-eddy simulation (LES) database of various turbine settings. The long-term goal is to use the kinematic simulation (KS) approach to study fatigue loading of wind turbines as a complex function of many environmental parameters, including wakes of upwind turbines in the wind farm. As input to develop KS models, researchers need a number of simulations of single and multiple turbine settings using LES.

IMPACT Completion of this project will result in a simulation capability that can be used to design better wind turbines and to lay out large wind farms with maximum energy extraction and improved turbine fatigue life.

## Code Development Projects

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

PI NAME Mark Gordon
INST Iowa State University
CODE GAMESS

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

PI NAME Volker Blum
INST Duke University
CODE FHI-aims, GAtor

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

PI NAME Kenneth E. Jansen
INST University of Colorado Boulder
CODE PHASTA

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

PI NAME Christos Frouzakis
INST Swiss Federal Institute of Technology, Zürich
CODE Nek5000

**Quantum Monte Carlo Calculations in Nuclear Theory**

PI NAME Steven Pieper
INST Argonne National Laboratory
CODE GFMC

**The Hadronic Contribution to the Anomalous Magnetic Moment of the Muon**

PI NAME Paul Mackenzie
INST Fermilab
CODE MILC, CPS

# SCIENCE

The ALCF is accelerating scientific discoveries in many disciplines, ranging from chemistry and engineering to physics and materials science.

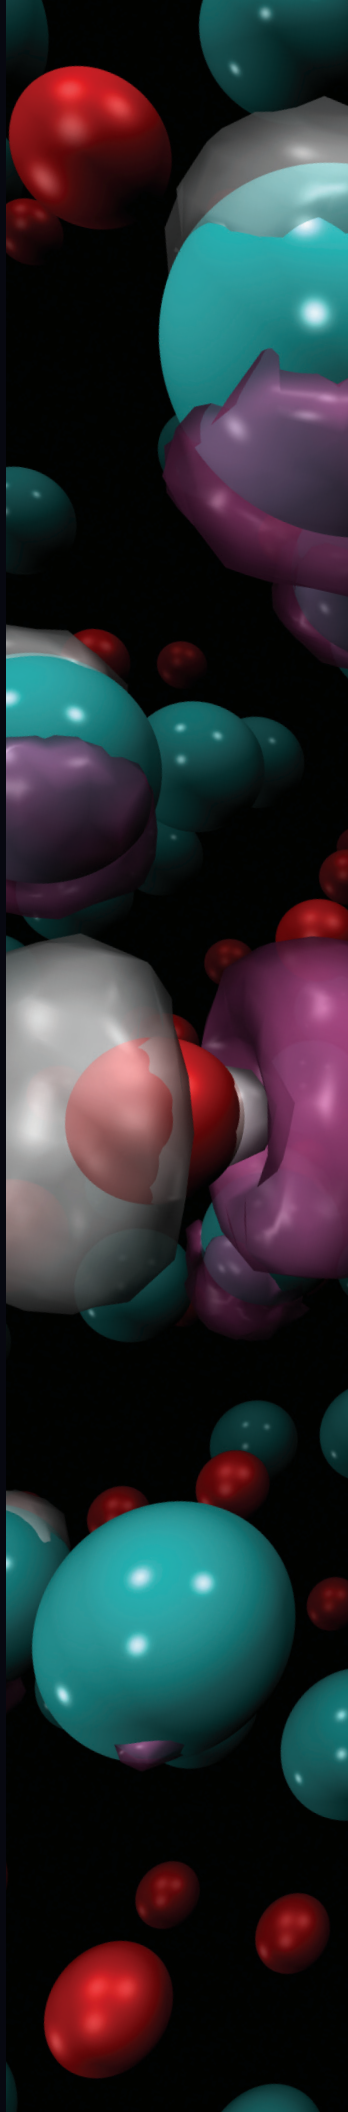

Image caption: Diffusion Monte Carlo spin density difference between bulks of potassium-doped nickel oxide and pure nickel oxide, showing the effects of substituting a potassium atom (center atom) for a nickel atom on the spin density of the bulk. *Image credit: Anouar Benali, Olle Heinonen, Joseph A. Insley, and Hyeondeok Shin, Argonne National Laboratory*

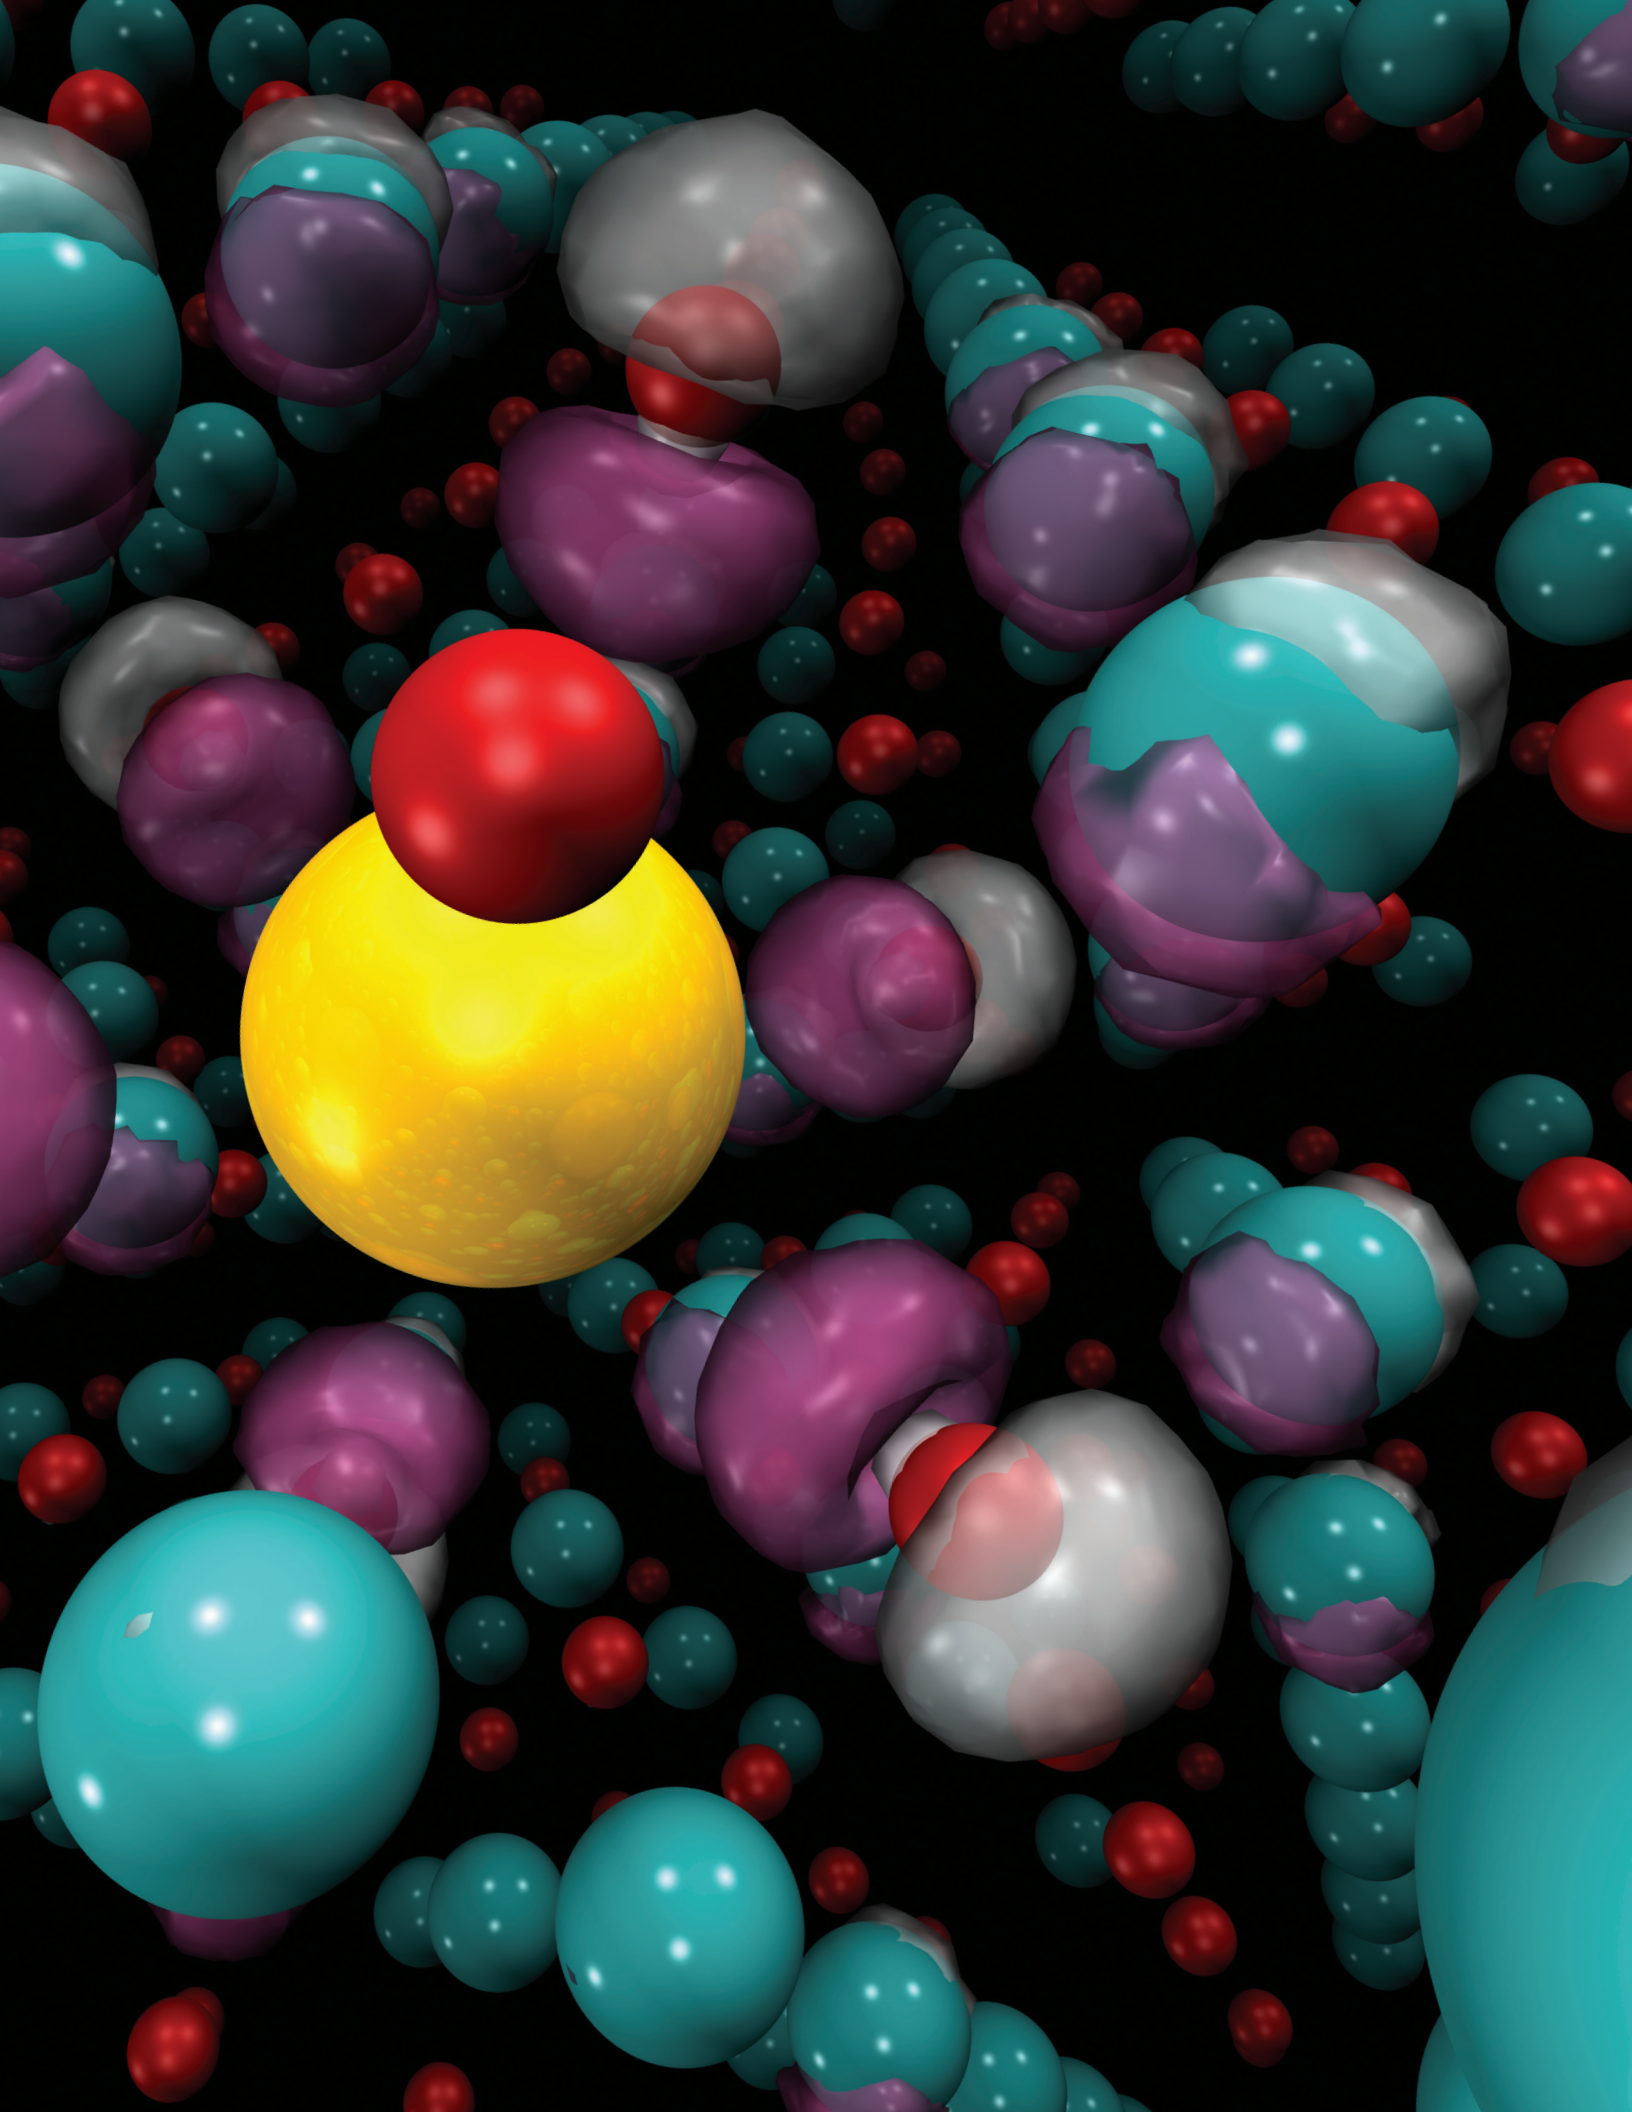

# Large-Scale Computing and Visualization on the Connectomes of the Brain

PI NAME Doga Gursoy and Narayanan Kasthuri
INST Argonne National Laboratory
HOURS ADSP
30 Million Core-Hours

Image caption: Reconstructed cerebrovascular network of the mouse brain using synchrotron microtomography.
*Image credit: Joseph A. Insley, Argonne National Laboratory*

CHALLENGE Through the ALCF Data Science Program, a new initiative targeted at big data problems, this project is developing a large-scale data and computational pipeline that integrates exascale computational approaches for understanding brain structure and pathology. Initial studies will focus on the reconstruction of mice brains utilizing novel imaging and analytical tools to image, for the first time, at the level of individual cells and blood vessels.

APPROACH The team will use X-ray microtomography, a high-resolution 3D imaging technique, to analyze the brain of a petite shrewmouse at submicron resolutions, providing a detailed picture of blood vessels and cell bodies. An electron microscope then will allow for the capture of all the synaptic connections between individual neurons at small targeted regions guided by the X-ray microtomography.

Data provided by these images will require the development of reconstruction, segmentation, and analysis tools, and their seamless integration onto large-scale computing systems. To meet this goal, researchers currently are focused on tweaking their codebase on advanced Intel architectures.

In addition, they are introducing scalable workflows focused on analysis and visualization of experimental data, such as the RhoAna framework, a machine learning-based technique.

These combined techniques will, for the first time, allow researchers to compare potential organizational patterns across brains to determine which are genetic and which are unique.

RESULTS Existing codebases were run effectively on the ALCF's Cooley system using an MPI-based master/worker model. The team developed benchmarking tools to evaluate the performance of these applications.

Benchmarking efforts have focused on making it easy to rapidly explore combinatoric build and run options, while automating analysis. This led to the packaging of code for use in the Spack package manager framework.

IMPACT The images produced by this work will provide a clearer understanding of how even the smallest changes to the brain play a role in the onset and evolution of neurological diseases, such as Alzheimer's and autism, and perhaps lead to improved treatments or even a cure.

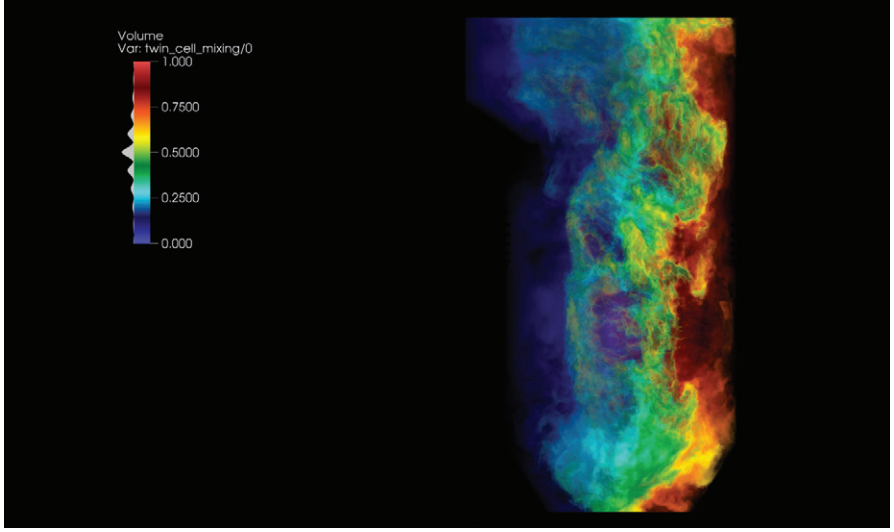


# Evaluation of a 1,000 MW Commercial Ultra-Supercritical Coal Boiler

PI NAME Martin Berzins
INST The University of Utah
HOURS INCITE
280 Million Core-Hours

Image caption: Mixing of the gases generated during the combustion process in a two-cell pulverized coal boiler. All of the inlet jets in the left cell are blue (0); all in the right are red (1). The green region represents the interaction of the gases between each cell.
*Image credit: Ben Isaac, The University of Utah*

CHALLENGE Pulverized coal power plants currently account for 39 percent of the power on the grid in the United States. The desire to develop more fuel-efficient and cleaner coal boilers is driving the power generation industry to use computational simulations to investigate the accelerated adoption of new boiler designs and technologies. This INCITE project simulates and evaluates cleaner, next-generation ultra-supercritical coal boiler designs from General Electric (GE), through the DOE National Nuclear Security Administration-funded Carbon Capture Multidisciplinary Simulation Center at the University of Utah.

APPROACH The model for the study is GE's 60-meter tall "twin-fireball" boiler, which is stoked when coal and air are injected from a main windbox, and over-fire air (OFA) is injected from above. Large-eddy simulations (LES) of different OFA configurations will validate and further improve the boiler design.

The studies rely on LES-based codes in the scalable Uintah open-source framework, which is comprised of computational components that simulate turbulent combustion and address fluid-structure interaction problems at the largest computational scales possible on Mira. The team also integrated the PIDX I/O data format, which resulted in a significant improvement in I/O times, allowing them to output data at more frequent intervals.

RESULTS Four large simulations of this GE boiler provided detailed results on the performance of its design. Analysis indicated that a significant fraction of energy is removed in the convection pass of the boiler, suggesting that the energy distribution at the furnace outlet planes is important. The results also have provided a better understanding of the complex multiphysics occurring in these large boilers. Comparisons with experiments indicated good overall agreement between simulations and experiments for relevant quantities of interest.

Additional INCITE hours allowed the team to simulate a new boiler using an oxygen-rich chemistry. The 500 MW high-efficiency ultra-supercritical oxy-coal power boiler operates at an efficiency of up to 53 percent, compared to the overall U.S. average of about 30 percent.

IMPACT Results demonstrate the role that LES simulations can have on analysis and design of massive operational commercial boilers. They can also serve as design tools for future systems that can operate at much higher efficiencies with the potential to dramatically impact $CO_2$ levels for a highly available carbon-based fuel.

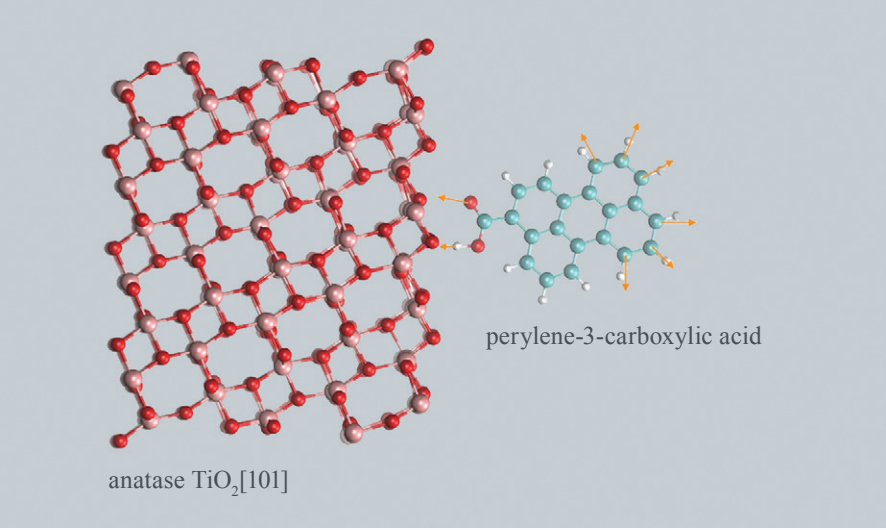


# Molecular Modeling of Hot Electron Transfer for Solar Energy Conversion

PI NAME Hanning Chen
INST George Washington University
HOURS ALCC
16 Million Core-Hours

Image caption: Driving vibrational mode (orange arrows) that promotes interfacial electron injection from a photo-excited perylene-3-carboxylic acid to an anatase $TiO_2$ [101] surface.
*Image credit: Hanning Chen, George Washington University*

CHALLENGE Charge carriers that are driven out of thermal equilibrium by excessive kinetic energy are vibrationally excited electrons, or hot carriers. While they are beneficial in photovoltaic devices, carrier-collecting electrodes can rapidly cool the hot carriers, resulting in a substantial loss of solar energy. But thermally insulated hot carriers can reach a maximum photon conversion efficiency of 66 percent, doubling that of conventional solar cells. Researchers aim to develop a vibrationally resolved electron transfer theory to facilitate ultra-efficient photovoltaic conversion in emerging hot-carrier solar cells.

APPROACH To promote the development of new methods for solar energy conversion through computer-aided optimization of photon-electron coupling, several first-principles methodologies have been developed to model photo-excited electrons. This includes the functional mode electron transfer theory, which identifies the driving vibrational mode for an electron transfer process between any pair of reactant and product states.

The team will conduct thermodynamic integration simulations to determine the energy diagram of the interfacial hot electron injection. Specifically, they will construct 11 thermodynamic integration windows along an ascertained reaction coordinate to determine both the electron injection driving force and its associated reorganization energy. The ultimate goal is to reveal the dependence of the electron injection rate on the incident wavelength under non-thermalized conditions, paving the way for systematic design of hot-carrier solar cells.

RESULTS The team has successfully calculated the diabatic energy gap of 5,000 snapshots of titanium oxide ($TiO_2$)-bound perylene-3-carboxylic acid that were randomly extracted from a 1.0 ns *ab initio* molecular dynamics trajectory followed by functional mode analysis (FMA). Results of an FMA calculation helped identify the reaction coordinate of the photo-induced electron injection from the perylene dye to a $TiO_2$ semiconductor. A ring-tethering vibrational mode with an effective angular frequency of 984 cm-1 was found to promote the interfacial photo-induced electron injection. Its thermal relaxation rate of 0.8 ps-1 was also determined by a phonon scattering simulation, suggesting that it could be outpaced by the ultrafast electron injection under non-thermalized conditions.

IMPACT The new methodology will enable computer simulations of hot electron transfer at various sensitizer/semiconductor interfaces, and deliver a more complete picture of solar energy harvesting when adjoined with other condensed phase theories. Ultimately, the research will pave the way for the systematic design of hot-carrier solar cells.

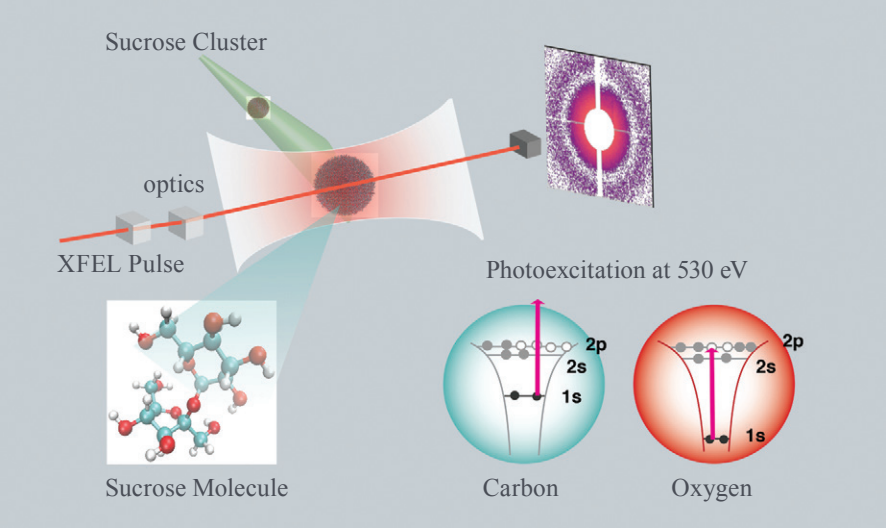


# Modeling of Intense X-ray Laser Dynamics in Nanoclusters

PI NAME Phay Ho
INST Argonne National Laboratory
HOURS ALCC
10 Million Core-Hours

Image caption: High-intensity X-ray free-electron laser pulses exciting 1s electrons in sucrose sugar clusters in coherent X-ray diffractive imaging experiment.
*Image credit: Phay Ho, Argonne National Laboratory*

CHALLENGE High-brightness X-ray free-electron laser (XFEL) pulses provide researchers with unprecedented tools to follow the dynamics of atoms and electrons with atomic spatial resolution. However, the intensities of the XFEL pulses also exhibit a strong interaction with the targeted sample, inducing complex dynamics. This project aims to develop a quantitative and predictive understanding of X-ray matter interactions in nano-sized heterogeneous systems at high intensities, then characterize the complex phenomena that emerge.

APPROACH The challenges associated with tracking the motion of particles and evolution of electronic configurations are addressed with a novel Monte Carlo/molecular-dynamics (MC/MD) simulation algorithm implemented in the simulation code LAMMPS. The quantum nature of the initiating ionization process is accounted for by an MC method to calculate probabilities of electronic transitions and track the transient electronic configurations explicitly.

Researchers modeled the scattering response of large sugar balls (sucrose molecular clusters with sizes reaching 50 nm in diameter) in the soft X-ray regime that includes and exceeds the water window. The sucrose clusters were irradiated with intense XFEL 180-fs pulses. MC/MD calculations were employed to investigate the X-ray processes as a function of pulse parameters and cluster size.

RESULTS Analysis revealed the important role of resonant excitation in the molecular scattering response in the water window. In particular, a high-intensity XFEL pulse can efficiently enable quantum excitation processes that promote electrons from the 1s quantum orbital to the 2p excited orbital in the oxygen atom/ion. These processes provide additional ionization pathways which, combined with the long pulse duration, lead to substantial reduction in the scattering power of sugar clusters.

After extending the 180-fs calculation to photon energies below 530 eV, and adding calculations for 5-fs pulse durations, calculations showed that resonance excitation channels are suppressed at photon energies below 530 eV. They also predict that the reduction of scattering signals will only be exhibited over a narrow energy range in the water window.

IMPACT This novel simulation method will facilitate the efficient investigation of the truly multiscale nature of complicated processes in heterogeneous systems. Results will further establish the applied methodology suitable for next-generation leadership computing resources, and provide predictions and new concepts to support and guide current and future XFEL experiments.

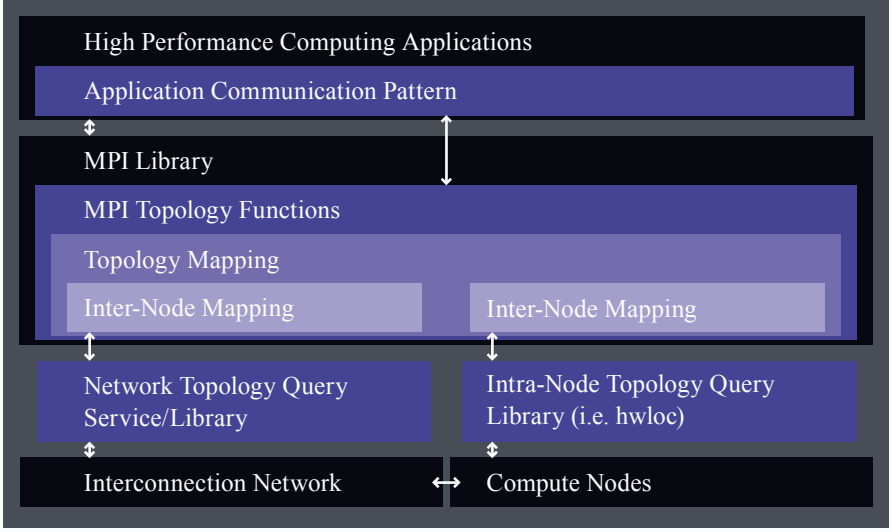


# Topology Mapping of Irregular Applications

PI NAME Zhiling Lan
INST Illinois Institute of Technology
HOURS Director's Discretionary 250,000 Core-Hours

Image caption: Graphical representation of hierarchical task mapping for parallel applications on supercomputers.
*Image credit: Zhiling Lan, Illinois Institute of Technology; Xuanxing Xiong, Synopsys; Jingjin Wu, University of Electronic Science and Technology of China*

CHALLENGE Today's systems, and those anticipated in the future, are largely bound by their communication infrastructure and the cost of data movement across a large number of nodes. Topology aware mapping, i.e., the optimized mapping of parallel application tasks onto physical processors so as to minimize data transmission and network sharing, is a promising approach to address the problem.

Topology aware mapping is an NP-hard problem. While the mapping of regular Cartesian structures to grid-based topologies is well understood, the general problem of mapping application processes onto physical processors with arbitrary topologies, is not. This project aims to develop a cooperative framework for topology awareness to improve communication efficiency and power efficiency.

APPROACH This research consists of four tightly coupled tasks: (1) to develop a lightweight MPI profiling and tracing library to capture application communication patterns; (2) to develop hierarchical task mapping methods to efficiently assign application processors onto allocated computer nodes through inter-node mapping and intra-node mapping; (3) to develop an analytical topology mapping algorithm for torus-connected systems by exploring methods from very-large-scale integration (VLSI) physical design; and, (4) to investigate inter-application and intra-application interference on the emerging Dragonfly system. The team uses Mira for the development of tasks 1-3, and will use Theta for task 4.

RESULTS The team has already studied various topology mapping algorithms for a suite of parallel applications with irregular communication patterns on Intrepid and Mira, including the NAS Parallel Benchmarks, the University of Florida Sparse Matrix Computation Benchmarks, and the cosmology application ART. Preliminary results show that the hierarchical task mapping methods can find high-quality mappings with low runtime overhead, which improves application communication performance by up to 77 percent. The analytical mapping algorithm can reduce the network traffic by up to 83 percent, achieving communication time reduction by up to 72 percent. TOPOMap, the hierarchical task mapping library and the analytical topology mapping tool, are publicly available to the research community.

IMPACT The goal of this work is to develop efficient task mapping algorithms to manage parallel workflows on high-performance computing systems. This work not only promotes productivity of such systems, but also impacts a broad range of application domains that use parallel systems for simulations.

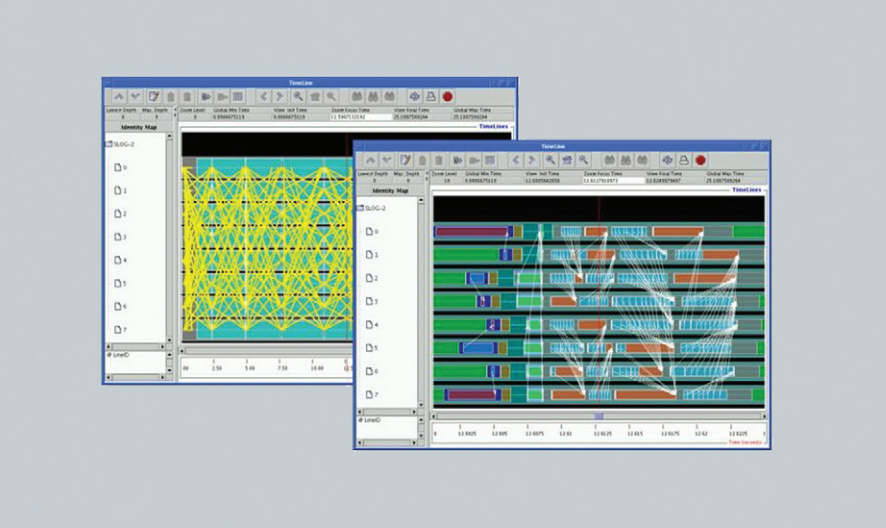

# MPICH—A High-Performance and Widely Portable MPI Implementation

PI NAME Ken Raffenetti
INST Argonne National Laboratory
HOURS Director's Discretionary
26 Million Core-Hours

Image caption: MPI Communication Visualization
*Image credit: MPICH Team*

CHALLENGE MPICH is a high-performance implementation of MPI that originated from Argonne's Mathematics and Computer Science Division over three decades ago. Today, MPICH and its derivatives run on many of the top supercomputers in the world. A team from Argonne is now adapting the code to extreme-scale machines, with a focus on the following challenges: adapting MPICH for increased on-node parallelism; utilizing advanced fabric features; and minimizing the memory used by the MPICH library.

APPROACH To address increased on-node parallelism in future hardware, the team will improve support for threads, including advanced locking and lock-free data-structures. Mira's thread density and parallel network capabilities make it an ideal platform for analysis and development. Since several applications running on the Blue Gene/Q are moving towards driving communication through multiple threads, this effort will benefit them as well.

Next-generation network hardware will provide richer communication interfaces and hardware offload for MPI semantics normally handled in software. To take advantage of these interfaces, the team will implement the thinnest software layer possible between MPI applications and the network. Support for hardware-tagged messaging and remote direct memory access is important for applications to get the most performance out of MPI. The team also plans to leverage modern compiler/linker features like static inlining and interprocedural optimization to build optimized libraries for MPI users. Addressing memory usage will require a refactoring of several MPICH data structures. Things like the MPICH "virtual connection" data structure will be made optional by removing them from internal MPICH application programming interfaces. This enables connectionless networks to store the minimum amount of data needed to communicate with peers by simply translating from MPI rank and communicator to a network address.

RESULTS Work to address these challenges has begun with the implementation of CH4, a new abstract communication device interface. The effort is a collaboration between Argonne, Intel, Cray, Mellanox, Sandia National Laboratories, and RIKEN. CH4 will be available in MPICH 3.3, due for release in mid-2017. Preliminary measurements of the raw instruction counts going from MPI call to network operation show significant improvement, which are reflected in message-rate benchmarks at small scale.

IMPACT Using the Open Fabrics Interface provider for Blue Gene/Q, the team can use Mira to demonstrate the viability of their approaches at large scale and beyond. Experiments will include typical MPI benchmarks, as well as application analysis using DOE applications targeted for DOE's next-generation machines.

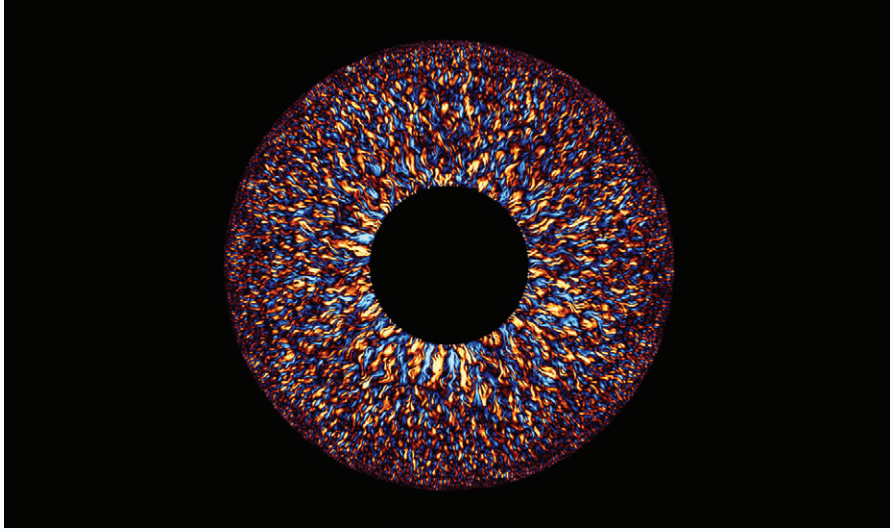

# Frontiers in Planetary and Stellar Magnetism Through High-Performance Computing

PI NAME Jonathan Aurnou
INST University of California, Los Angeles
HOURS INCITE
260 Million Core-Hours

Image caption: Radial velocity field (red = positive; blue = negative) on the equatorial plane of a numerical simulation of Earth's core dynamo. These small-scale convective flows generate a strong planetary-scale magnetic field.
*Image credit: Rakesh Yadav, Harvard University*

CHALLENGE Magnetic fields are generated deep in the cores of planets and stars by a process known as dynamo action. This phenomenon occurs when the rotating, convective motion of electrically conducting fluids converts kinetic energy into magnetic energy. To better understand the dynamics that lead to magnetic field generation, a research team is using ALCF computing resources to simulate the turbulent interiors of Earth, Jupiter, and the Sun at an unprecedented level of detail.

APPROACH With this three year INCITE project, researchers are developing high-resolution 3D models of planetary and stellar dynamo action in turbulent fluids with realistic material properties. Working with ALCF staff, the team has optimized Rayleigh, an open-source code designed to study magnetohydrodynamic (MHD) convection in spherical geometries, for Mira, allowing them to resolve a range of spatial scales previously inaccessible to numerical simulation.

RESULTS The team successfully implemented rotation into their solar dynamo model, which led to some of the highest-resolution and most turbulent simulations of solar convection ever performed. In a paper published in *Astrophysical Journal Letters*, they used the simulations to place upper bounds on the typical flow speed in the solar convection zone—a key parameter to understanding how the Sun generates its magnetic field and transports heat from its deep interior.

In addition, the researchers have made significant progress with their Jupiter convection zone model, enabling the highest-resolution giant-planet simulations yet achieved. The Jupiter simulations will be used to make detailed predictions of surface vortices, zonal jet flows, and thermal emissions that will be compared to data from NASA's Juno mission.

Finally, ongoing simulations of Earth's geodynamo are showing that flows and coupled magnetic structures develop on both small and large scales, revealing new MHD processes that do not appear to develop in lower resolution computations.

IMPACT This project is developing advanced dynamo models to shed light on the interplay of magnetism, rotation, and turbulent convection occurring within the remote interiors of planets and stars. The resulting datasets, which will be made publicly available to the broader research community, will help advance the understanding of dynamo processes and provide new insights into the birth and evolution of the solar system.

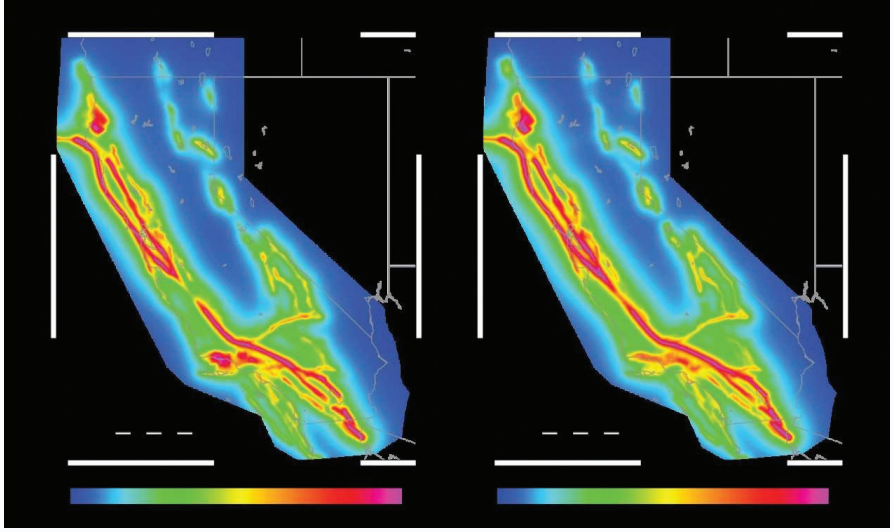

# Quantification of Uncertainty in Seismic Hazard Using Physics-Based Simulations

PI NAME Thomas H. Jordan
INST University of Southern California

HOURS INCITE
141 Million Core-Hours
ALCF: 45M; OLCF: 96M

Image caption: California seismic hazard maps generated with the physics-based RSQSim and empirical UCERF3 models. Maps show the peak ground acceleration, plotted in units of surface gravity, expected with 2 percent probability of exceedance in 50 years. The hazard-map agreement between the empirical and physics-based models provides substantial support for both types of forecasting models.
*Image credit: Kevin Milner, University of Southern California*

CHALLENGE Human and economic risks in seismically active regions continue to increase as urban areas and their dependence on interconnected infrastructure networks continue to grow. Understanding seismic hazards across a wide spectrum of forecasting and response times, including a proper assessment of modeling uncertainties, is the foundation on which most seismic risk-reduction strategies are developed. The Southern California Earthquake Center (SCEC) is using DOE resources to improve the accuracy of earthquake simulations to better understand seismic hazard and assess seismic risk.

APPROACH Mira was used to run the earthquake simulator code RSQSim, which produces long-term synthetic earthquake catalogs. A baseline simulator model was developed with several global uniform model parameters tuned to match earthquake scaling observations. The team also turned to Mira to run the new dynamic rupture simulation code, Waveqlab3D. The code requires relatively few points per wavelength, allowing for both physics-based spontaneous dynamic rupture as well as wave-propagation for ground motion prediction. This eliminates the need for a two-step approach, simplifying the workflow in addition to reducing data storage.

RESULTS Using RSQSim, the team further developed and refined earthquake simulator models of the California fault system, comparing results from the simulator against state-of-the-art UCERF3 (Uniform California Earthquake Rupture Forecast 3). As an initial comparison between RSQSim results and the UCERF3 model, a series of hazard-relevant measures were made, combining both earthquake rupture forecasts with ground motion prediction equations. While the initial intent was to evaluate the magnitude of differences between the models, the SCEC team found, instead, very good agreement. Averaged over a representative set of sites, the RSQSim-UCERF3 hazard-curve differences are comparable to the small differences between UCERF3 and its predecessor, UCERF2. This agreement provides a fundamental verification at the system level for both approaches and is anticipated to be extremely productive in further advancing seismic hazard estimates and their uncertainties.

IMPACT This research will produce more accurate physics-based simulations of earthquake ruptures and wave propagation, and improve physics-based analyses by adding new physical elements and enhancing existing features in the models. Results are expected to improve broad impact seismic hazard information products, including seismic hazard maps and civil engineering design codes.

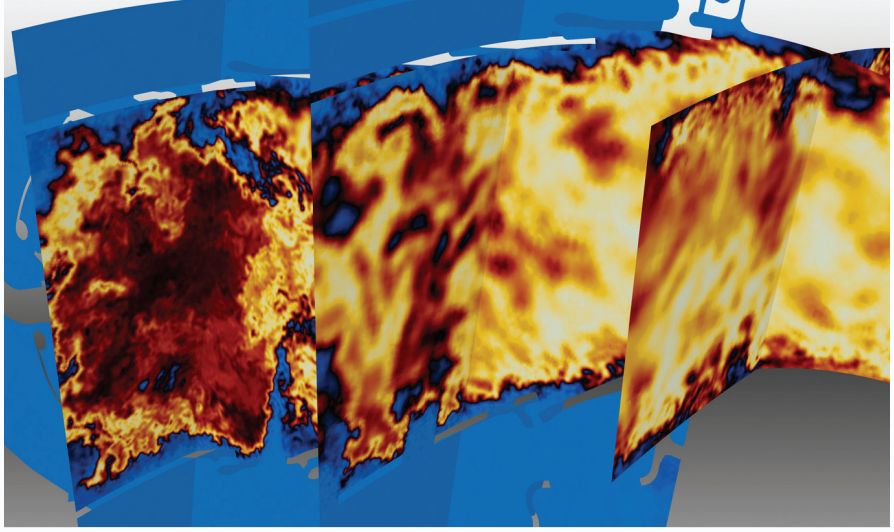

# Large-Eddy Simulations of Combustor Liner Flows

PI NAME Anne Dord
INST GE Aviation
HOURS INCITE
240 Million Core-Hours

Image caption: Temperature distribution through a series of planar cuts of a highly resolved GE combustor.
*Image credit: Lee Shunn, Cascade Technologies; Matthieu Masquelet, GE Global Research*

CHALLENGE This project looks at the complex physics of combustor liner flows to help designers develop cleaner, more efficient engines. Recent simulations used wall-modeled large-eddy simulations (LES) to analyze flow in single- and multi-cup combustors. These set the stage for ongoing and future calculations in the area of turbulent reacting flow modeling.

APPROACH An in-depth study of the detailed geometries of combustor systems requires computational fluid dynamics (CFD) resolutions achievable only on massively parallel computing platforms, such as Mira. This collaborative effort utilizes Cascade Technologies' CFD code, CharLES low-Mach Helmholtz solver. The code performs reacting LES and reproduces the important interactions between chemistry, turbulence, and acoustics that occur in real-world gas turbine engines.

In 2016, the team moved to a more efficient tool suite developed by Cascade Technologies, which includes new features and improved performance. Cascade also developed a massively parallel mesh generator to construct LES-suitable grids in complex geometries based on Voronoi diagrams, providing multiple advantages for both stationary and moving geometry.

RESULTS In the latter part of this three-year project, the team focused on building an appropriate grid for selected effusion-cooled geometries—wherein many small holes connect the combustor to the surrounding cold air passages to insulate the metal from hot combustion temperatures. The team performed two separate simulations; one a passive scalar run to analyze mixing, and the other a fully reacting run. By analyzing the field of the passive scalar, whether on the inner liner or on the exit plane, they found it was possible to evaluate the influence of mesh resolution and near-wall model on heat transfer and exit profiles.

In performing uncertainty quantification on a realistic reacting combustor, areas of high-temperature variability were highlighted, especially downstream of the dilution holes. It was determined that uncertainty in fuel flow rate had a greater impact on outlet quantities, whereas uncertainty in air inlet temperature had a greater impact on key combustor liner quantities.

IMPACT The proposal aims to accelerate the industrial impact of coupling petascale computing with high-fidelity numerical approaches to guide the design of future aircraft engines. In addition, understanding and optimizing the design of turbine engine combustor liners can lead to significant breakthroughs in part life and emissions.

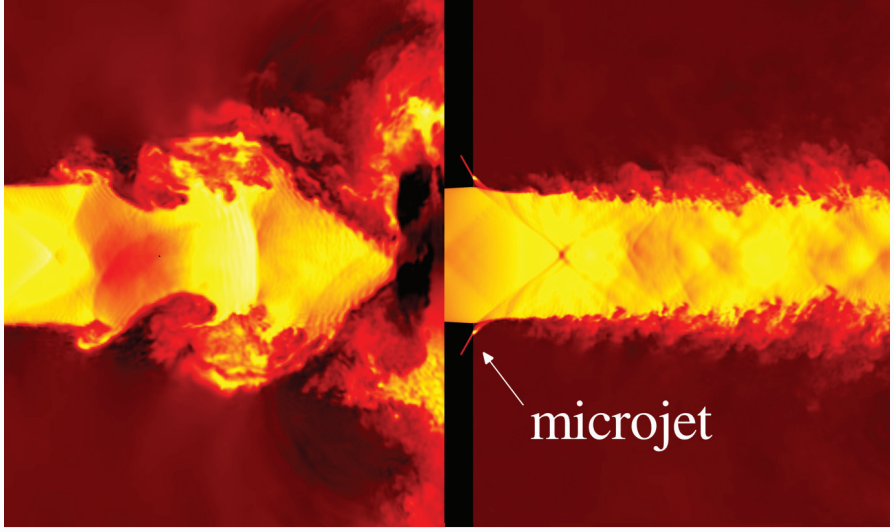


# Novel Reduced-Order Models of Turbulent Jet Noise from High-Fidelity Simulation

PI NAME Joseph Nichols
INST University of Minnesota
HOURS Director's Discretionary
10 Million Core-Hours

Image caption: Researchers performed high-fidelity simulations to reproduce laboratory experiments of a supersonic jet striking a flat wall (left). Analysis of reduced-order models extracted from simulations led to predictions of modified nozzle designs (right). The simulations confirmed drastically altered dynamics and low noise levels.
*Image credit: Joseph Nichols, University of Minnesota*

CHALLENGE The reduction of noise is an important design parameter in the research and development of jet engines. There remains, however, a lack of basic understanding of the mechanisms by which high-speed turbulent flows, like engine exhausts, generate sound. Researchers are harnessing large-scale simulations at ALCF to construct novel, computationally inexpensive techniques to successfully predict and understand the aeroacoustics of high-speed jets.

APPROACH Early models of jet noise involve acoustic analogy, useful in identifying compact acoustic sources. In this framework, non-compact effects important to jet flows in and from ducts can only be recovered by complicated two-point two-time correlations.

Methods inspired by instability theory, such as the parabolized stability equations (PSE), can model non-compact acoustic sources, but do not yet fully recover the far-field acoustics of high-speed jets, especially in the sideline direction, perpendicular to the flow. Researchers applied a technique known as input-output analysis to understand the shortcomings of the PSE and to predict sound in this important direction.

Taking another step, the team constructed a model of a supersonic jet that strikes a flat surface, using global modes based on simulation data. Optimal flow control was extracted from the model, and tested with high-fidelity simulation.

RESULTS Researchers found that the new input-output modes contribute significantly to sound generation and explain the missing sound in the PSE analysis. This is significant because it is the first direct evidence that sideline noise can be explained in terms of coherent, non-compact acoustic sources rather than incoherent motions, as previously thought.

Simulations precisely reproduced the dynamics of a laboratory supersonic jet striking a flat wall. Simulation data informed a reduced-order model that identified the key instability mechanism and provided a systematic way to control it. Simulations confirmed the effectiveness of the control, and suggested new laboratory experiments targeted at previously unknown physics.

IMPACT Guided by high-fidelity simulation, this approach could lead to a more universal acoustic source model than those commonly used in industry, and could become an important design tool. The systematic simulation and modeling methodology could revolutionize laboratory experiments by targeting them at otherwise undetected physics.

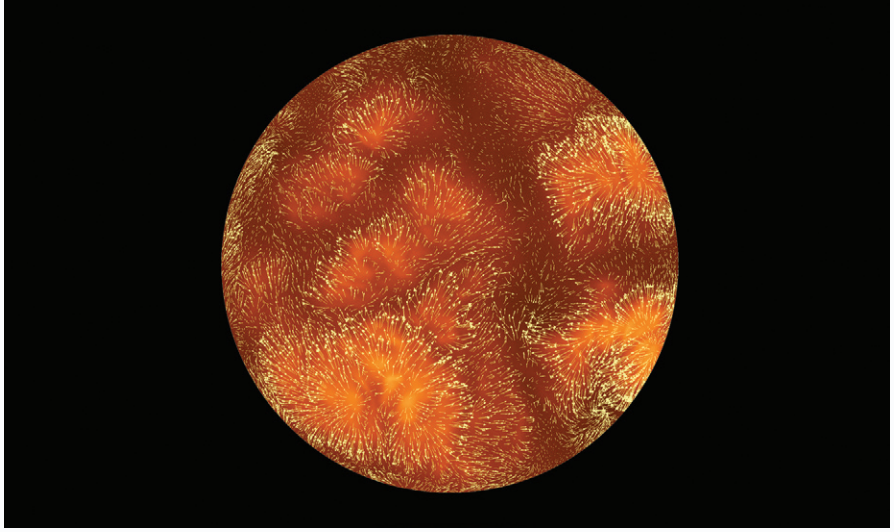

# Convective Turbulence in Liquid Gallium and Sodium

PI NAME Janet Scheel
INST Occidental College
HOURS INCITE
80 Million Core-Hours

Image caption: Temperature (color density plot) and velocity (vector plot) for a horizontal cut through the bottom boundary layer for Pr = 0.005 and Ra = 5 x 107. For the temperature plot, yellow indicates lower temperature and orange, higher temperatures. The size of the arrows indicates the magnitude of the velocity.
*Image credit: Janet Scheel, Occidental College*

CHALLENGE Turbulent Rayleigh-Bénard convection (RBC) is a type of natural convection in which a constant temperature difference is maintained for a horizontal layer of fluid, heated from below. Because the fluid is confined, thin thermal and velocity boundary layers form near the top and bottom plates. One of the central unsolved problems in turbulent RBC is finding the transition to turbulence in these thin boundary layers. Simulations of turbulent RBC at high Rayleigh and low Prandtl numbers will help determine heat transport properties and predict transition parameters.

APPROACH Numerical simulations of turbulent RBC are based on the Nek5000 spectral element software package developed for solving the fluid flow equations on massively parallel supercomputers. The high-resolution simulations will help develop a better understanding of the global structure of the convective flow and the boundary layer dynamics for this new parameter regime.

RESULTS The team performed very finely resolved direct numerical simulations of RBC for Prandtl number (Pr) = 0.021 and Rayleigh number (Ra) = 4 × 108 and Pr = 0.005 and Ra = 5 × 107. The scaling of heat transport and momentum with Ra were determined for Pr = 0.005 and 0.021 using these new parameters.

The velocity and thermal boundary layers were then extracted for a comparison between the turbulent boundary layers in shear flow and the boundary layers for RBC convection. They found that the boundary layers for low Prandtl number convection systems are not fully turbulent in terms of a canonical logarithmic law, but saw a systematic approach to this law as the Rayleigh number increased.

Finally, using three different statistical analysis methods, the parameters for the transition to a turbulent boundary layer were determined for Pr = 0.005, 0.021, and 0.7. All three methods indicate that the range of the Rayleigh number for the transition is shifted to smaller magnitudes as the Prandtl number becomes smaller.

IMPACT Simulations of turbulent heated convection at high Rayleigh and low Prandtl numbers will help researchers determine heat transport properties, in general, and the transition parameters for a turbulent boundary layer. Low-Prandtl-number convection has direct application to convection in liquid metals, such as gallium or sodium, and can shed light on convection in Earth's liquid metal core and in the Sun.

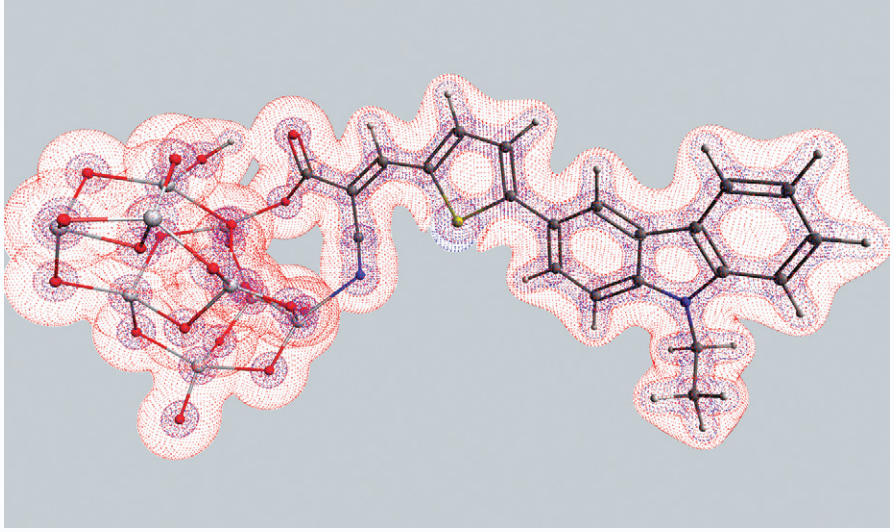

# Data-Driven Molecular Engineering of Solar-Powered Windows

PI NAME Jacqueline M. Cole
INST University of Cambridge
HOURS ADSP
134 Million Core-Hours

Image caption: Isosurface of the Laplacian of the electron density of MK44, an organic dye, attached to the surface of a nanocluster of titania.
*Image credit: Álvaro Vázquez Mayagoitia, Argonne National Laboratory; Jacqueline M. Cole, University of Cambridge*

CHALLENGE Dye-sensitized solar cells (DSC) are a strong contender for next-generation solar cell technology. Their transparent and low-cost nature makes them niche prospects for electricity-generating windows that will equip buildings for energy-sustainable future cities. Despite their vast industrial potential, DSC innovations are being held up by a lack of suitable light-harvesting dye chemicals.

APPROACH This work aims to discover new DSC dyes via a large-scale data mining approach. Prior to mining, suitable data must be sourced. Such data are being generated using a toolkit called ChemDataExtractor. This enables the automated extraction of chemical entities and their associated properties, measurements, and relationships from scientific documents that can be used to populate structured chemical databases. Natural language processing and machine learning lie at the core of this tool. The chemistry literature is a good target for this automated extraction as it is typically comprised of formulaic, data-rich language that is well-suited for machine analysis with the potential for high recall and precision.

The data-mining process probes the databanks that have been constructed, with encoded forms of structure-property relationships which explain the underlying physics and chemistry that govern good DSC dyes. These relationships are established using case studies on known DSC dyes which are built up from a range of materials characterization efforts to support this dye discovery.

RESULTS Within this scope, case studies on dye molecules are being performed using calculations on Theta to complement the synchrotron-based experiments on known high-performance DSC dyes. This work has revealed new molecular design rules for DSC dyes that will help future materials prediction efforts.

The team has also developed a new charge-transfer algorithm which could provide a better way to encode certain descriptors for electronic structure calculations. They are currently verifying the algorithm via tests on case studies of experimentally determined molecular structures. Pending the final assessment of this evaluation process, this new algorithm development will represent a scientific contribution in its own right.

IMPACT Materials discovery of better performing light-absorbing dye molecules will be enabled via a synergistic computational and experimental science approach, wherein machine learning and data mining are used in conjunction with large-scale simulations and experiments to facilitate a materials-by-design workflow.

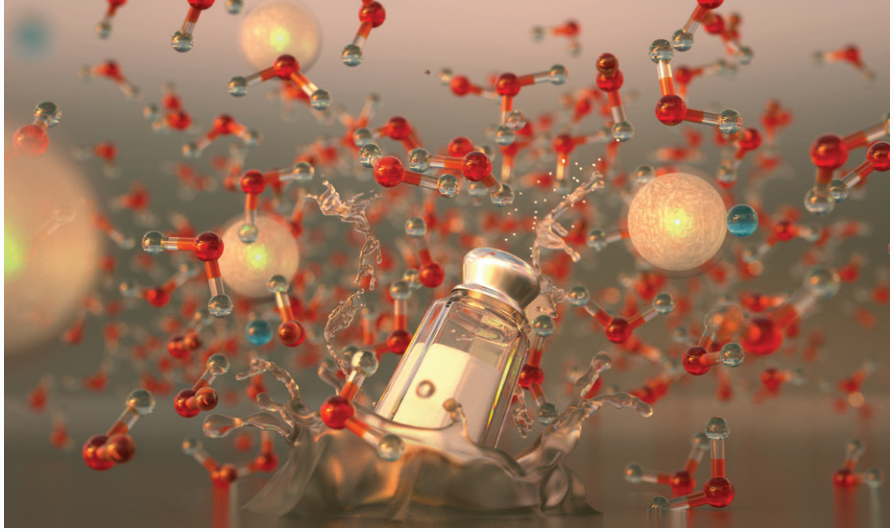

# Computational Spectroscopy of Heterogeneous Interfaces

PI NAME Giulia Galli and François Gygi
INST The University of Chicago, Argonne National Laboratory; and University of California, Davis
HOURS INCITE
200 Million Core-Hours

Image caption: Using Mira, researchers performed extensive first-principles simulations to determine the effects that dissolved salts have on the structure of water.
*Image credit: Peter Allen, The University of Chicago*

CHALLENGE Heterogeneous interfaces between solids, nanoparticles, and liquids play a fundamental role in determining materials properties. With an understanding of the microscopic structure of solid-water and solid-electrolyte interfaces, researchers can better predict the properties of optimal materials for solar energy applications, clean fuel production, and energy storage. However, the properties of interfaces are seldom explicitly included in computational models due to the complexity and cost of the associated calculations.

APPROACH With this multi-year INCITE project, researchers are developing advanced computational methods to study the physical properties of aqueous interfaces with solid oxides and semiconductors, and of inorganic interfaces at the nanoscale. The team is using the open-source codes Qbox and WEST, supported by the Midwest Integrated Center for Computational Materials (MICCoM), to perform calculations of optoelectronic and vibrational spectra integrated with large-scale *ab initio* molecular dynamics simulations. This methodology, combined with the petascale power of Mira, provides a practical time-to-solution for identifying vibrational signatures of specific reaction pathways that occur at interfaces.

RESULTS The team recently published a paper in the *Journal of Physical Chemistry Letters*, detailing the effects that dissolved salts have on the structure of water. For this study, the team performed extensive first-principles simulations of simple salt solutions. They found that, while the cation does not significantly change the structure of water beyond the first solvation shell, the anion has a farther reaching effect, modifying the hydrogen-bond network even outside its second solvation shell. The researchers determined that a distinctive fingerprint of hydrogen bonding modification is responsible for the change in polarizability of water molecules. Molecular dipoles, extensively used in the past to analyze salt solutions, were instead found to be rather insensitive probes. These findings will help scientists to better understand the solvation of biomolecules, as well as the role played by salts at interfaces and in biochemical processes.

IMPACT This project is developing open-source computational protocols that enable accurate calculations of the electronic properties of electrolyte solutions and interfaces important to energy-related applications. Theorists and experimentalists alike can use analysis tools derived from first-principles calculations to interpret experiments and to optimize materials for use in clean fuel production and solar energy applications.

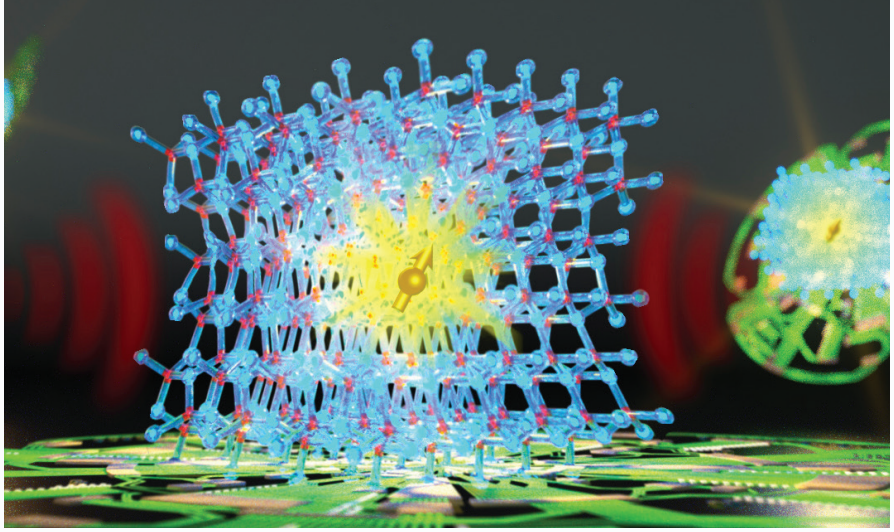

# Computational Engineering of Defects in Materials for Energy and Quantum Information Applications

PI NAME Marco Govoni
INST The University of Chicago and Argonne National Laboratory
HOURS ALCC
53.7 Million Core-Hours

Image caption: This figure illustrates the spin of a nitrogen vacancy in aluminum nitride designed for quantum bit applications.
*Image credit: Hosung Seo, The University of Chicago; Giulia Galli and Marco Govoni, The University of Chicago and Argonne National Laboratory*

CHALLENGE Defects in materials can have unexpected and peculiar properties that may degrade a material's performance. However, the unusual properties of defects can be used to design new functionalities that are not present in their host materials. With this ALCC project, researchers are developing a first-principles theoretical framework to study the properties of defects in soft and hard matter. This will enable the investigation of new materials for applications, including solar-powered fuel production from water and solid-state quantum computing and sensing.

APPROACH The team's computational framework is based on *ab initio* molecular dynamics and accurate electronic structure simulations using the open-source Qbox, Quantum Espresso, and WEST codes. Optimized to run on Mira, these scalable, integrated first-principles codes are capable of tackling systems of unprecedented size (several thousands of electrons) that are crucial to simulating realistic liquid electrolytes and isolated atom-like defects in bulk crystals. The researchers are using the framework to obtain atomic trajectories; compute thermodynamic and electronic properties of electrolytes; and perform high-throughput calculations to identify new potential spin quantum bits (qubits) in semiconductors, such as silicon carbide and aluminum nitride.

RESULTS In a paper published in *Science Advances*, the team demonstrated its efficient and accurate approach for predicting the electronic properties of aqueous solutions. The researchers studied the photoelectron spectra of a broad range of solvated ions, showing that first-principles molecular dynamics simulations and electronic structure calculations using dielectric hybrid functionals provide a quantitative description of the electronic properties of the solvent and solutes, including excitation energies. Their theoretical results were validated through comparison to data from state-of-the-art photoelectron liquid-jet experiments.

In addition, the team published a paper in *Physical Review X*, detailing their development of a generalization of dielectric hybrid functionals that yielded accurate electronic and optical properties for organic and inorganic molecules and semiconducting nanocrystals.

IMPACT By enabling computational characterization of defective states at a microscopic level, this project will provide knowledge and tools to interpret ongoing experiments on fuel production from water. This work will also help establish design rules for predicting robust qubits in technologically mature and functional wide-gap semiconductors.

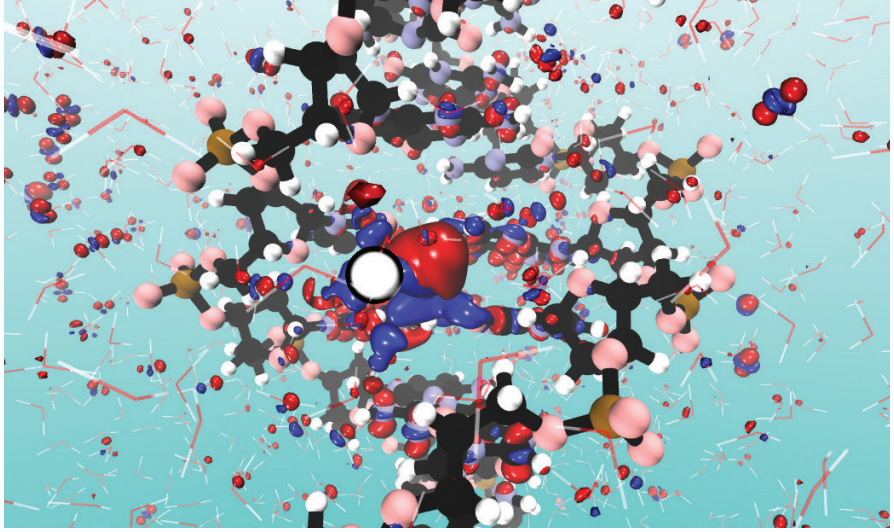

# Electronic Stopping in Condensed Matter Under Ion Irradiation

PI NAME Yosuke Kanai
INST University of North Carolina at Chapel Hill
HOURS INCITE
130 Million Core-Hours

Image caption: A snapshot from non-equilibrium electron dynamics simulation of solvated DNA under proton irradiation. The blue and red isosurface represents areas with positive and negative changes in electron density with respect to the equilibrium density.
*Image credit: Dillon Yost, University of North Carolina at Chapel Hill*

CHALLENGE Electronic stopping describes the transfer of kinetic energy from a highly energetic ion to electrons. Understanding this phenomenon in condensed matter is central to advancing various technological and medical applications, including focused-ion beam microscopes and proton-beam cancer therapy. However, existing mathematical models lack predictive capability and do not provide insights at the atomistic level. To overcome these limitations, researchers are using ALCF computing resources to develop a predictive computational framework to enable accurate studies of electronic stopping and the underlying energy transfer rate, called stopping power.

APPROACH This research team continues previous INCITE work to enhance a highly scalable implementation of real-time, time-dependent density functional theory based on the Qbox/Qb@ll code. The researchers have implemented new numerical integrators and further optimized the code for Mira. Their research is currently focused on simulating the electronic stopping process in the context of two complex phenomena: (1) secondary electron emission from aluminum and graphite due to ion radiation, and (2) electronic excitations in water and DNA under proton irradiation.

RESULTS In their studies of secondary electron emission, the team has carried out massively parallel simulations on Mira to track the electron dynamics during proton irradiation in magnesium oxide, and to explore how the existence of neutral oxygen vacancy affects the non-adiabatic electron-ion dynamics. Their simulations are helping to illustrate the role of the kinetic energy of ions from radiation sources and of projectile-defect interactions in damage formation in semiconductor materials during ion irradiation.

For the team's other research thrust, their quantum-mechanical simulations have revealed how electronic excitations induce electron ionization of individual water molecules in liquid water under proton irradiation. Further investigation with DNA is currently underway.

IMPACT This project is establishing a computational framework that enables researchers to predict electronic stopping power accurately and study the underlying energy transfer mechanism at an unprecedented level of detail. Understanding these processes at the atomistic level will help advance the development of a number of applications, including focused-ion beam technology, proton-beam therapy, and materials for advanced nuclear reactors.

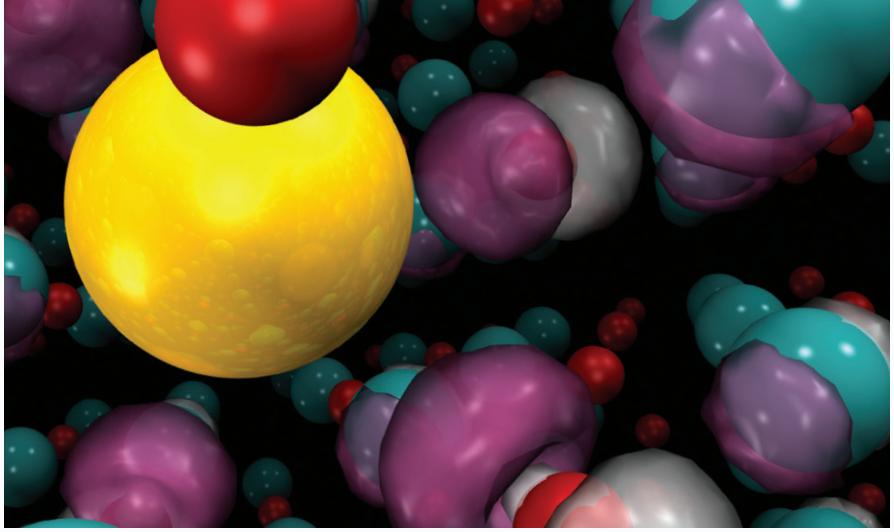

# Predictive Simulations of Functional Materials

PI NAME Paul Kent
INST Oak Ridge National Laboratory
HOURS INCITE
138 Million Core-Hours
ALCF: 98M; OLCF: 40M

Image caption: Diffusion Monte Carlo spin density difference between bulks of potassium-doped nickel oxide and pure nickel oxide, showing the effects of substituting a potassium atom (center atom) for a nickel atom on the spin density of the bulk. *Image credit: Anouar Benali, Olle Heinonen, Joseph A. Insley, and Hyeondeok Shin, Argonne National Laboratory*

CHALLENGE Today, the computational design of functional materials is hindered by the limited predictive power of established modeling and quantum mechanics-based approaches. However, the quantum Monte Carlo (QMC) method has emerged as a powerful tool for providing realistic predictions of materials properties, due to its ability to accurately simulate the complex interactions between many electrons. This computationally demanding method was once limited to modeling systems of small atoms or molecules, but with supercomputers like Mira, researchers can now use QMC methods for rigorous calculations on complex materials.

APPROACH This work aims to improve predictive modeling capabilities and simulate promising new technological materials that cannot be simulated with current techniques. Using the open-source QMCPACK code, the team is demonstrating and validating new QMC methods and algorithms. Their simulations are focused on functional materials such as transition metal oxides that display a wide spectrum of properties, including magnetism, charge ordering, metal-insulator transitions, and large-band gap variation. These materials are of great interest to new electronics, energy conversion, and energy transmission technologies. The team's predictions will be validated by new experimental and characterization efforts that use the full range of DOE user facilities, including neutron scattering, soft and hard X-rays, and high-resolution electron microscopy for imaging and spectroscopy.

RESULTS As part of their research on transition metal oxides, the team used QMCPACK to predict and describe the electronic properties of potassium-doped nickel oxide. The researchers validated their findings through comparisons to experimental data. Their simulation results provide a better understanding of the evolution of electronic and magnetic structure of a correlated oxide under doping, and shed light on how the electronic and magnetic properties can be controlled with oxygen vacancies and strain.

IMPACT The ability to quantitatively predict, analyze, and, therefore, design functional materials with tailored properties is an essential component of materials innovation. This project aims to establish the QMC method as an accurate predictive computational approach for accelerating the design and discovery of new functional materials for energy-related technologies.

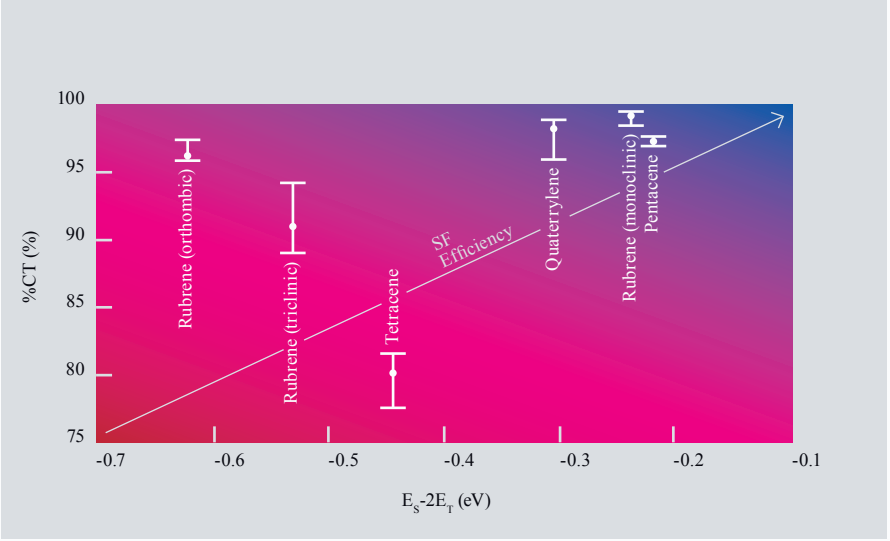


# Materials and Interfaces for Organic and Hybrid Photovoltaics

PI NAME Noa Marom
INST Carnegie Mellon University
HOURS INCITE
160 Million Core-Hours

Image caption: Known and predicted singlet fission (SF) materials ranked with respect to a two-dimensional descriptor based on maximizing the energy conservation criterion and the degree of charge transfer (CT) character of the singlet exciton. Quaterrylene is predicted to exhibit SF with high efficiency, close to that of pentacene, the quintessential SF material. Quaterrylene has the additional advantages of high stability and a narrow band gap that would enhance solar conversion efficiency, compared to pentacene.
*Image credit: Xingyu Liu, Noa Marom, and Xiaopeng Wang, Carnegie Mellon University; Cameron Cook and Bohdan Schatschneider, California State Polytechnic University, Pomona*

CHALLENGE While today's solar energy market is largely dominated by silicon-based photovoltaics, solar cells based on organic and hybrid organic-inorganic materials offer the potential to generate electricity at a lower cost. In addition, their electronic properties can be engineered at the molecular level to improve performance and reliability. With this INCITE project, researchers are using ALCF computing resources to advance the efficiency of organic and hybrid solar cells through computer-aided discovery and design of new materials, crystal forms, and functional nanostructured interfaces.

APPROACH The INCITE team has developed first-principles approaches, based on density functional theory and many-body perturbation theory, to describe materials and interfaces on the most fundamental level at which solar energy conversion takes place and on which structure-function relationships are established. Their broadly applicable framework is achieved through the seamless integration of an electronic structure layer based on the FHI-aims all-electron code with a structure prediction layer comprised of the GAtor code genetic algorithm, the SAMPLE Gaussian process regression algorithm, and the BOSS Bayesian optimization code. ALCF staff has been instrumental in helping parallelize these codes for optimal performance on the facility's systems.

RESULTS The team's large-scale computational campaign is focused on three areas that are critical to organic and hybrid photovoltaics: molecular crystals, interfaces, and hybrid perovskites. Thus far, they have incorporated machine learning techniques into the GAtor and SAMPLE codes to enable structure prediction for organic compounds at interfaces with various metals. The researchers also derived a new concept of "pair modes" to describe the distribution of relative nearest-neighbor cation orientations, enabling a statistical description of the properties of disordered hybrid perovskite materials at finite temperatures. In addition, the team identified crystalline quaterrylene as a possible new singlet fission (SF) material. The realization of SF-based solar cells may significantly increase conversion efficiency by harvesting two charge carriers from one photon.

IMPACT This research is enabling transformative advances in the understanding of organic and hybrid solar photovoltaic materials, and the technology needed to harness their power for affordable, large-scale solar cells. The project's methodological advances will catalyze the emergence of new computational design paradigms that help advance materials science for energy applications.

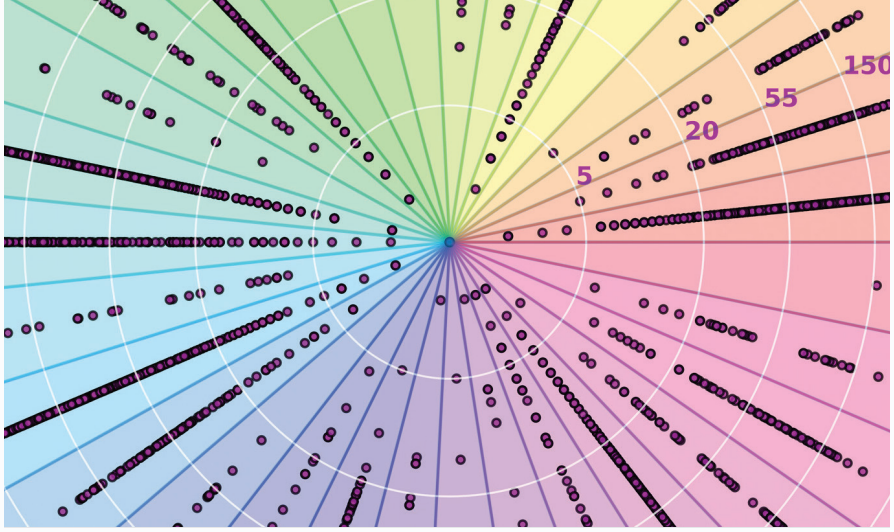


# The Materials Project—Completing the Space of Elastic and Piezoelectric Tensors

PI NAME Kristin Persson
INST Lawrence Berkeley National Laboratory
HOURS ALCC
36 Million Core-Hours

Image caption: The bulk modulus for various symmetries grouped by crystal system represented in the Materials Project.
*Image credit: Shyam Dwaraknath, Joseph Montoya, and Kristin Persson, Lawrence Berkeley National Laboratory*

CHALLENGE Piezoelectric materials, which produce an electrical charge in response to mechanical stress, are central to many important applications, including ultrasonic surgical tools and microelectronic power harvesters. However, because most current piezoelectric technologies are based on lead-based materials, researchers are searching for new materials that are more efficient and environmentally friendly.

To help accelerate the design of novel materials, a research team from the DOE's Materials Project is using ALCF computing resources to perform high-throughput calculations for elastic and piezoelectric tensors (an array of numerical values used to describe a material's properties). The resulting data will be disseminated in the Materials Project's extensive, open-access database containing the thermodynamic and electronic properties of virtually all known and predicted materials.

APPROACH The Materials Project's computational infrastructure is designed to automatically initiate, perform, and post-process simulations with minimal user intervention. This ALCC project is developing a framework for leadership-class machines using the following set of software packages: Fireworks for workflow management; VASP for electronic structure calculations; Pymatgen for materials analysis; and MongoDB for data management. Working with ALCF staff, the researchers are scaling and optimizing these software tools to run efficiently on the facility's petascale computing resources.

RESULTS The team has created an infrastructure to parallelize separate VASP runs, enabling them to perform tensor calculations on Mira in a high-throughput manner. In addition, they have developed methodologies to increase their workflow yield significantly by reducing the number of VASP calculations required to determine elastic and piezoelectric tensors. Following the completion of calculations and subsequent filtering, the results will be uploaded into the Materials Project's online database.

IMPACT This project aims to accelerate the design and discovery of advanced materials for piezoelectric applications. With ALCF computing resources, the team is charting the space of elastic and piezoelectric tensors more thoroughly than ever before, delivering new insights into how their properties emerge from crystal and electronic structures, and, ultimately, constructing a framework for the design of new candidate materials for emerging piezoelectric technologies.

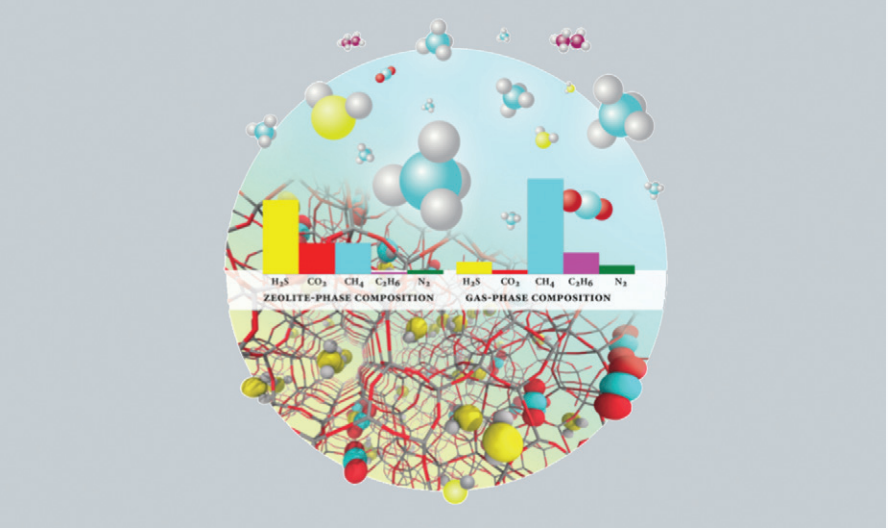


# Predictive Modeling of Functional Nanoporous Materials, Nanoparticle Assembly, and Reactive Systems

PI NAME J. Ilja Siepmann
INST University of Minnesota
HOURS ALCC
117 Million Core-Hours

Image caption: Adsorption-based process for the removal of hydrogen sulfide and carbon dioxide from sour natural gas.
*Image credit: Greg Chung, Pusan University; Evgenii O. Fetisov, Mansi Shap, and J. Ilja Siepmann, University of Minnesota*

CHALLENGE Nanoporous materials, such as metal-organic frameworks and zeolites, are of great interest to the biofuel and petrochemical industries because of their ability to act as sponges for gas storage, as molecular sieves for separations, and as catalysts that aid in the processing of fuels and chemical feedstocks. However, finding an optimal material and process conditions for a given application is a time- and labor-intensive process that could take decades using traditional laboratory methods.

APPROACH Scientists participating in the DOE-funded Nanoporous Materials Genome Center are using Mira to demonstrate and develop predictive theory and modeling tools that can rapidly screen thousands of materials to pinpoint promising candidates for further research. This project uses hierarchical screening workflows that involve machine learning, evolutionary algorithms, molecular simulations, and high-level electronic structure calculations. Leveraging the MCCCS-MN, RASPA, and CP2K codes, the team's research is focused on performing high-throughput screening to discover nanoporous materials with specific functions; first-principles simulations of reactive phase and adsorption equilibria, and of electrolyte solutions in slit pores; and electronic structure calculations to explore nanoporous materials and photovoltaic materials.

RESULTS The team used Mira to interpret experimental findings and predict the performance of zeolite-based separation membranes created using a newly developed synthesis method (patent pending). The groundbreaking fabrication process offers a more efficient and cost-effective method for producing zeolite nanosheets, which can be used to make ultra-selective, high-flux membranes. Calculations on Mira helped elucidate the nanosheet structure and transport barriers for xylene isomers, providing data that corroborated experimental measurements of ultra-high selectivities under operational conditions. This work was published in the journal *Nature*. In other studies, the team's simulations have led to the discovery of optimal zeolites for sweetening of highly sour natural gas streams and novel microporous materials for upgrading the research octane number of fuel mixtures.

IMPACT This project is developing improved predictive modeling capabilities to accelerate the discovery and design of nanoporous materials for complex chemical separation and transformation applications. The ability to identify optimal zeolites and metal-organic frameworks for energy applications has the potential to improve the production of biofuel and petroleum products, and to advance the development of gas storage and carbon capture devices.

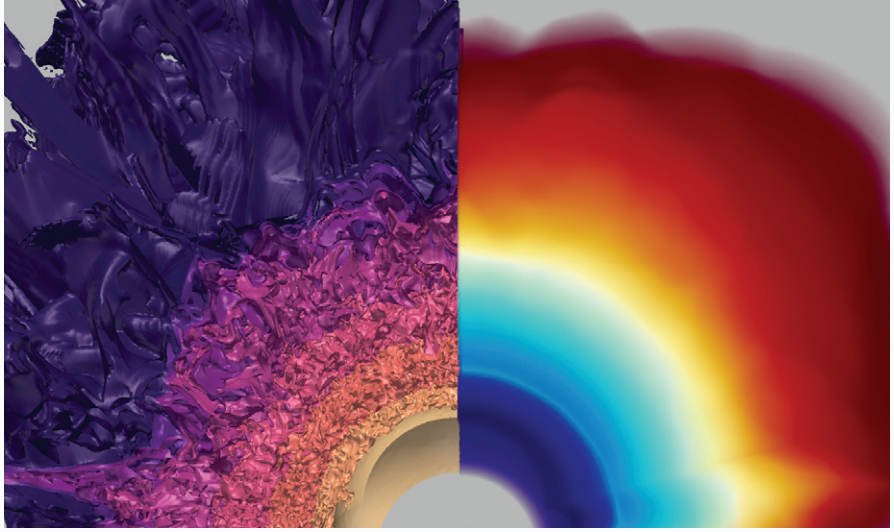

# Global Radiation MHD Simulations of Massive Star Envelopes

PI NAME Lars Bildsten
INST University of California, Santa Barbara
HOURS INCITE
60 Million Core-Hours

Image caption: Global radiation hydrodynamic simulation of an 80-solar-mass star envelope with an effective temperature of 1.67x10^4 K and a solar luminosity of 6.93x10^5. The simulation domain covers from 14 to 300 solar radii with the photosphere located at 100 solar radii. Convection around the iron opacity peak region causes the turbulent structures in the envelope. Episodic mass loss is also formed from the surface of the star. The left panel shows density in one quarter of the star, while the corresponding radiation energy density is shown in the right panel.
*Image credit: Joseph A. Insley, Argonne National Laboratory; Yan-Fei Jiang, University of California, Santa Barbara*

CHALLENGE Massive stars play an important role in many astrophysical environments, but poor understanding of mass loss is one of the biggest uncertainties in massive star evolution. Winds from these stars depend on their own surface layer structure, including the effects of hydrodynamic and/or magnetohydrodynamic (MHD) instabilities that can only be understood via 3D global radiation MHD simulations. These simulations study the global structure of the gaseous outer layers, or envelopes, of massive stars.

APPROACH Early calculations focused on the structure of massive star atmospheres, when the local flux becomes super-Eddington—outward radiation pressure force by photons exceeds the inward gravitational force on matter—which could lead to mass outflows. This is a result of increases in opacity due to atomic transitions of iron, or the iron opacity peak. A key goal is to determine whether the super-Eddington luminosity at the iron opacity peak drives significant mass loss for stars with different masses and at different evolutionary stages.

The team is now using the mesh refinement capabilities of its grid-based code, Athena++, to study stellar structures over a large radial range. Higher resolution near the iron opacity peak resolves 3D turbulent structures due to convection, and relatively lower resolution at large radii captures the extended envelope and radiation driven outflow. Particularly, they hope to capture the photosphere within the simulation to predict observational properties directly.

RESULTS The team completed the first-ever 3D simulations of radiation pressure-dominated massive star envelopes for two 80 solar mass stars; a hot blue star and a yellow supergiant. The simulations capture the turbulence driven by convection at the iron opacity peak. The yellow supergiant's extended, optically thick envelope displays very clumpy structures with significant inflows and outflows, which can explain many observational properties of these stars. The hot blue star has a compact photosphere, but shows significant outflows with mass flux comparable to the observationally inferred value.

IMPACT The project will dramatically improve the understanding of the surface layers of massive stars, including mass loss via radiation pressure-driven winds. The results will be used to create more realistic massive star models and supernovae progenitors, which will significantly improve understanding of the structure and evolution of massive stars.

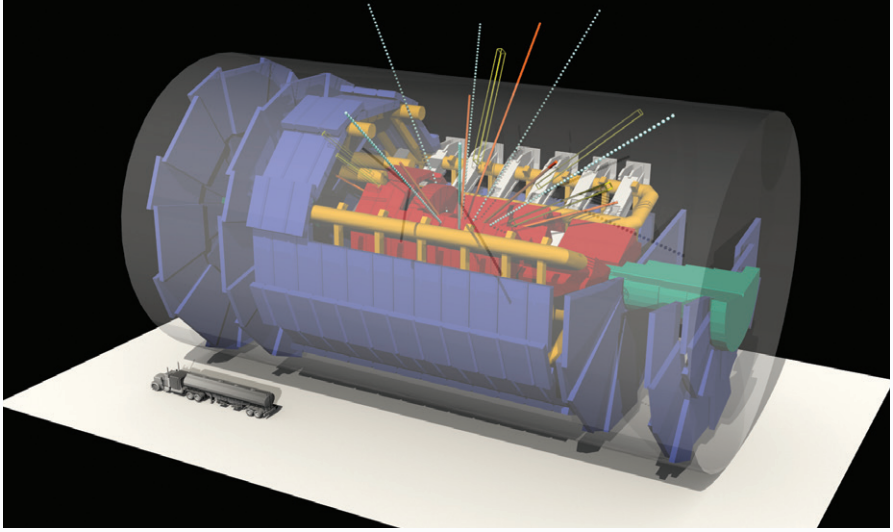

# Advancing the Scalability of LHC Workflows to Enable Discoveries at the Energy Frontier

PI NAME Taylor Childers
INST Argonne National Laboratory
HOURS ADSP
20 Million Core-Hours

Image caption: Artist's representation of the ATLAS detector at CERN's Large Hadron Collider, showing particles produced in the aftermath of the collision between two high-energy protons (the truck shown in lower left is depicted for scale). ATLAS researchers have simulated billions of LHC collision events on ALCF resources.
*Image credit: Taylor Childers, Joseph A. Insley, and Thomas LeCompte, Argonne National Laboratory*

CHALLENGE Particle collision experiments at CERN's Large Hadron Collider (LHC) generate around 50 petabytes of data each year that must be processed and analyzed to aid in the facility's search for new physics discoveries. When CERN upgrades to the High Luminosity LHC in 2025, experimental data produced by the facility is expected to increase by a factor of 20. To help meet the LHC's growing computing needs, a research team is exploring the use of ALCF supercomputers to perform increasingly precise simulations, as well as calculations that are too intensive for traditional computing resources.

APPROACH This research builds on the progress of an ongoing ALCC project that is using Mira to perform event generation simulations for the LHC's ATLAS experiment. With the ADSP allocation, the research team is expanding their use of ALCF resources through the deployment of Athena, the ATLAS simulation and analysis framework, on Theta. Their ultimate goal is to develop an end-to-end workflow on ALCF computing resources that is capable of handling the ATLAS experiment's intensive computing tasks—event generation, detector simulations, reconstruction, and analysis.

RESULTS As a first step, the ALCF is working to deploy HTCondor-CE, a "gateway" software tool developed by the Open Science Grid to authorize remote users and provide a resource provisioning service. This effort requires writing code to allow HTCondor-CE to interact with the ALCF's job scheduler, Cobalt, as well as making modifications to the facility's authentication policy. To test this setup, the Argonne team has successfully carried out end-to-end production jobs on Cooley, paving the way for larger runs on Theta in the near future. With this tool in place, ATLAS researchers will be able to run jobs automatically through their workflow management system, simplifying the integration of ALCF resources as production compute endpoints for LHC experiments.

IMPACT The development of a workflow that integrates event generation and simulation on ALCF computing resources will increase the scientific reach of the ATLAS experiment. By enabling a portion of the LHC computing workload to run on ALCF supercomputers, researchers can speed the production of simulation results and accelerate efforts to search for evidence of new particles.

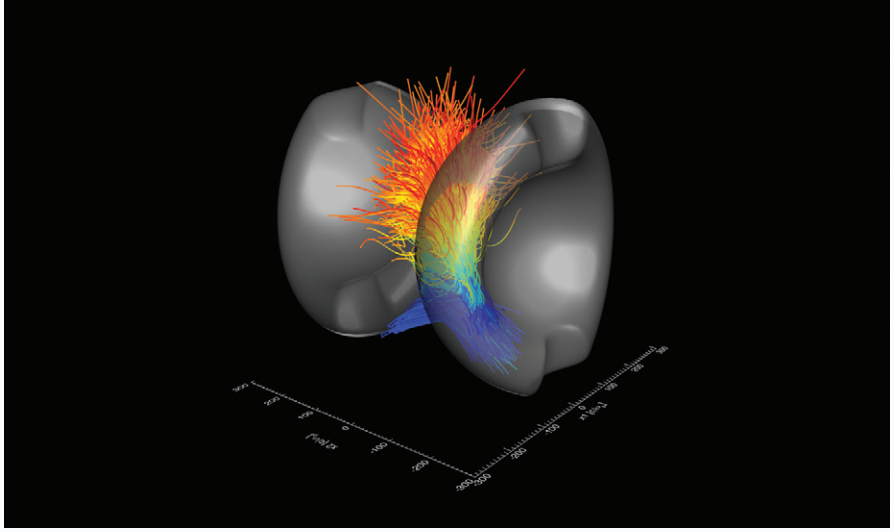

# *Ab Initio* Modeling of the Dynamical Stability of High-Energy Density Plasmas

PI NAME Frederico Fiuza
INST SLAC National Accelerator Laboratory
HOURS ALCC
60 Million Core-Hours

Image caption: A visualization from a 3D OSIRIS simulation of particle acceleration in laser-driven magnetic reconnection. The trajectories of the most energetic electrons (colored by energy) are shown as the two magnetized plasmas (grey isosurfaces) interact. Electrons are accelerated by the reconnection electric field at the interaction region and escape in a fan-like profile.
*Image credit: Frederico Fiuza, SLAC National Accelerator Laboratory*

CHALLENGE Magnetic reconnection is a fundamental process in astrophysical and fusion plasmas in which magnetic energy is converted into plasma energy. It is responsible for a violent acceleration of particles that produces spectacular displays in space, from the dancing lights of aurorae to solar flares. The mechanisms by which these occur are not well understood, but are important in explaining particle acceleration in explosive cosmic phenomena.

APPROACH Recent experiments based on laser-plasma interactions have shown the possibility of studying magnetic reconnection by using powerful lasers to drive two plasma plumes that carry magnetic fields with opposite polarities. A laser blasts two foil targets, which are instantly turned into plasmas that expand and interact with each other.

This ALCC project, an extension of similar work involving the development of particle acceleration in astrophysical shocks, is simulating these laboratory plasma experiments to help shed light on the astrophysical processes that accelerate cosmic ray particles.

Leveraging the relativistic massively parallel particle-in-cell code OSIRIS, researchers conducted detailed simulations to study the possibility of observing particle acceleration in such an experimental setup—arguably the most important signature of reconnection. These were the first 3D simulations of laser-driven reconnection and the first to show the possibility of particle acceleration in this laboratory setting.

RESULTS The simulation results show that laser-driven reconnection leads to strong non-thermal particle acceleration. As the two expanding plasma plumes interact with each other, they form a thin current sheet, or reconnection layer, which becomes unstable, breaking into smaller sheets. During this process, the magnetic field is annihilated and a strong electric field is excited in the reconnection region, strongly accelerating electrons as they enter the region. The accelerated electrons then escape in a fan-like profile which leads to the development of a non-thermal particle distribution associated with some of the most explosive phenomena in the universe, such as solar flares and gamma-ray bursts.

IMPACT Results confirm that magnetic reconnection can lead to the efficient acceleration of astrophysical particles. The insight gained from these 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.

# Fundamental Properties of QCD Matter Produced at RHIC and the LHC

PI NAME Claudia Ratti
INST University of Houston
HOURS INCITE
194 Million Core-Hours

Image caption: The image illustrates how protons, neutrons, and other hadrons formed from quarks and gluons during the QCD transition as the universe expanded. Since the transition is a crossover, there is no sharp temperature, only a broad range where the transition happened.
*Image credit: Sandor Katz, University of Budapest, Hungary*

CHALLENGE A few seconds after the Big Bang, the building blocks of matter emerged from a hot, energetic state known as the quark-gluon plasma (QGP). These building blocks of matter, called hadrons, form when gluons bind quarks together. Physicists are recreating the primordial conditions of the QGP in the Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory, and the Large Hadron Collider (LHC) at CERN. Results suggest that the QGP is a strongly interacting system, which can be studied computationally through lattice quantum chromodynamics (QCD).

APPROACH Using experimental results and the JANOS QCD simulation code, researchers working on this continuing INCITE project are addressing the question of how the transition is modified if the pressure of the deconfined system is raised and the temperature is held constant. This is of particular interest since the collisions generated at RHIC at lower energies produce such a system, and it has been postulated that QCD matter will exhibit critical behavior near a critical point at finite density.

RESULTS Through expansion techniques, the team mapped out the phase diagram to densities characterized by chemical potentials that are twice the pseudo-critical temperature. No evidence for a critical point has been found, but the precision of the calculations enabled the development of effective theories based on equivalent gravitational equations used in black hole calculations. These theories now predict a critical point at chemical potentials as high as eight times the critical temperature.

Another key result is the calculation of observables, which verified the existence of a large number of additional hadronic states carrying strange quarks. These particles are needed to explain the experimental results of the transition behavior between the quark and hadronic phases. These short-lived states have to exist and they have to be populated during the transition. The calculations specify their mass, charge, quark composition, and decay properties, which will allow experimentalists to discover particles never seen before in nature.

IMPACT The resulting data should help scientists determine the strength of coupling between quarks and gluons in the QGP, how charges propagate through it, and whether the QGP is the most ideal fluid ever observed.

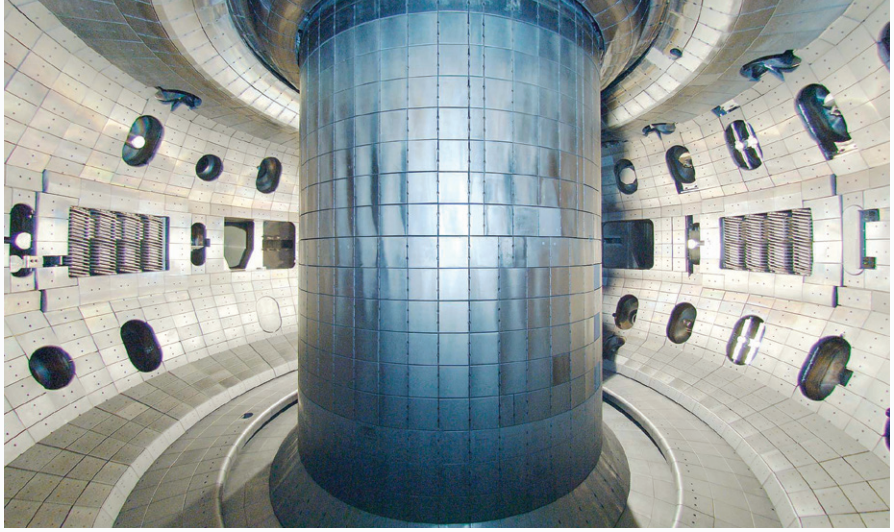

# DIII-D Computing

PI NAME David Schissel
INST General Atomics
HOURS Director's Discretionary 8,000 Core-Hours

Image caption: Inside view of the DIII-D Tokamak
*Image credit: General Atomics*

CHALLENGE Scientists at the DIII-D National Fusion Facility are exploring magnetic confinement of fusion energy by conducting fast-paced plasma physics experiments involving the creation of six-second pulses of confined plasma every 15 to 20 minutes.

Planning for each new pulse is informed by data analysis of the previous pulse, and for some experiments, a fusion science analysis code called SURFMN can assist in that planning. However, as the fine-grid analysis by SURFMN itself takes 20 minutes to complete on dedicated local resources, it hasn't been possible to obtain these results on a between-pulse timescale.

APPROACH An Argonne team of researchers devised a new service for DIII-D by leveraging code from a previous ALCF project that used Mira to simulate particle collision events generated by Large Hadron Collider experiments. For that project, the Balsam service was developed to seamlessly interface with LHC systems and execute the workflow on ALCF resources. For DIII-D, Balsam was adapted to interact with the DIII-D systems and to run fusion analyses instead.

RESULTS The ALCF solved the DIII-D team's timing mismatch problem by automating and shifting the analysis step to its data analysis cluster, Cooley, which computed the analysis of every single pulse and returned the results to the research team in a fraction of the time required by the computing resources locally available at DIII-D.

Now, a fine-grid SURFMN run takes under three minutes to complete and interacts directly with the DIII-D databases for automated input/output. The additional computing power also allowed SURFMN runs to be conducted with higher complexity, using a finer grid for the Fourier analysis, yielding more precise results than researchers typically obtain on local systems.

IMPACT Leadership computing facilities are expanding their services to include big data analytics and machine learning capabilities that can help large experimental and observational efforts make better use of their resources. The ALCF's integration with the DIII-D facility is one example of a partnership that is accelerating frontier science results.

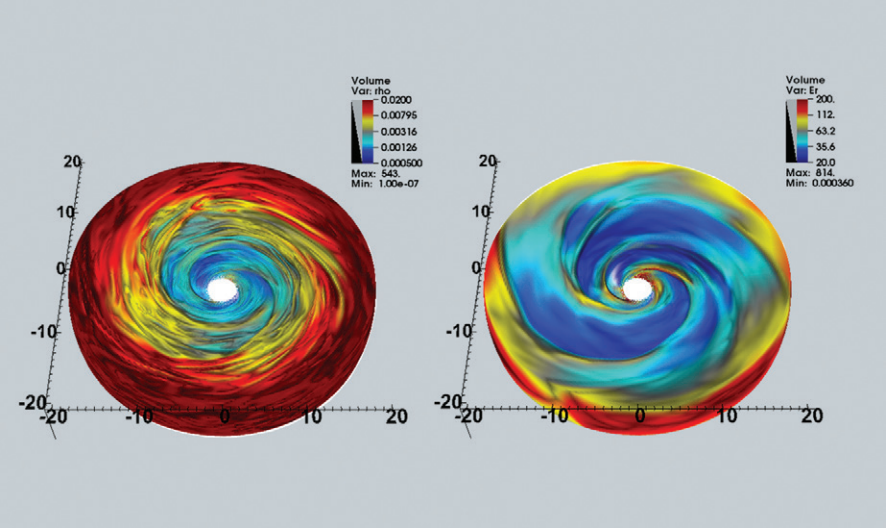


# Magnetohydrodynamic Models of Accretion Including Radiation Transport

PI NAME James Stone
INST Princeton University
HOURS INCITE
54 Million Core-Hours

Image caption: A snapshot of the inner region of an accretion disk onto a 5x10^8 solar mass black hole. The left panel shows the gas density (in unit of 10^-8 g/cm^3), while the right panel shows the radiation energy density (in unit of 1.21x10^7 erg/cm^3). The spiral structures in the images are the density waves excited in the disk, which are responsible for the accretion. The length unit in the image is the Schwarzschild radius.
*Image credit: Yan-Fei Jiang, University of California, Santa Barbara*

CHALLENGE The most luminous objects in the universe, such as quasars and gamma-ray bursts, are powered by accretion, the process of accumulating matter under the influence of gravity. A solid theoretical understanding of the structure, dynamics, and evolution of accretion flows is crucial for interpreting observations of a range of astrophysical sources. However, the important role that radiation transport plays in determining the physics of accretion remains largely unexplored.

With this INCITE award, a research team is using ALCF computing resources to investigate radiation-dominated accretion and make predictions about how accreting sources evolve and affect their environment.

APPROACH The team has developed new and accurate numerical algorithms for time-dependent radiation transport that can be integrated into existing magnetohydrodynamic (MHD) codes to study the physics of radiation-dominated flows. These methods have been implemented in a new code called Athena++, a compressible MHD code, with features that include mesh refinement and new physics, such as algorithms for general relativistic MHD in stationary spacetimes. The team has optimized the code for ALCF supercomputers, achieving excellent performance and scaling.

RESULTS Using Mira, the researchers have completed a survey of accretion around supermassive black holes as the mass accretion rate is varied from highly super-Eddington (the radiation-dominated, slim disk regime) to sub-Eddington (the standard, thin disk regime). The team's simulations, which included realistic radiation transport for the first time, led to the discovery of a new mode of accretion in which spiral density waves excited by baroclinic effects drive the accretion process. In addition, the researchers are investigating whether thermal and/or viscous instabilities, predicted over 40 years ago, actually occur in global models of MHD turbulent disks.

IMPACT Through the development of accurate algorithms for radiation MHD that perform well on leadership computing systems, this project has opened up entirely new regimes of astrophysical accretion for exploration. Results from this effort will be important for testing the basic theory of accretion in luminous sources, understanding galaxy formation and the growth of supermassive black holes in the early universe, and improving the interpretation of observational data.

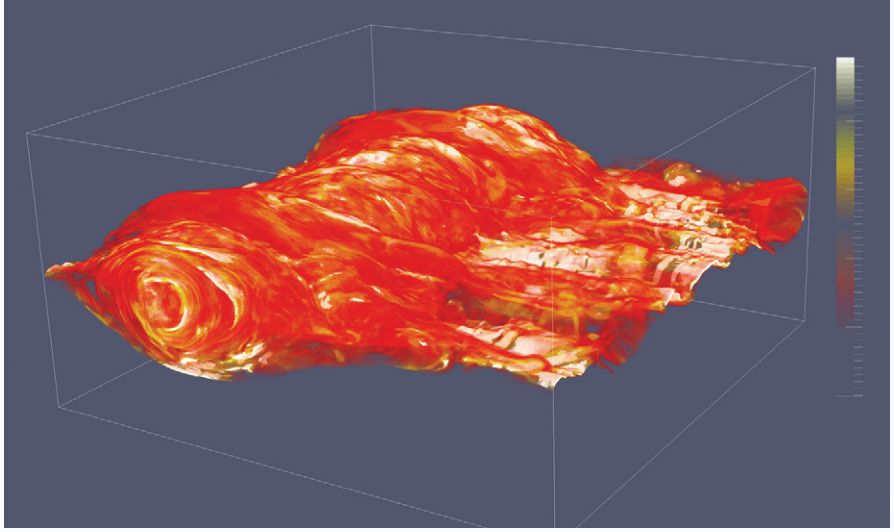

# Kinetic Simulations of Relativistic Radiative Magnetic Reconnection

PI NAME Dmitri Uzdensky
INST University of Colorado Boulder
HOURS INCITE
90 Million Core-Hours

Image caption: The electric current during relativistic magnetic reconnection in electron-positron plasma. *Image credit: Greg Werner, University of Colorado Boulder*

CHALLENGE Magnetic reconnection is a fundamental plasma physics process important in a great variety of both natural and laboratory plasma environments, from magnetic fusion devices to solar flares to intense high-energy X-ray and gamma-ray emissions from exotic astrophysical sources. Using very large kinetic numerical simulations, this project investigates high-energy particle acceleration and radiation caused by magnetic reconnection in the astrophysically relevant regime of relativistic, magnetically dominated plasma.

APPROACH To conduct its simulations, the team developed the radiative particle-in-cell code, Zeltron, which self-consistently incorporates the radiation reaction force on relativistic particles. The code was recently optimized to improve load balancing, greatly improving its performance on Mira.

RESULTS A comprehensive investigation of 3D magnetic reconnection in relativistic electron-positron pair plasmas demonstrated robust acceleration of particles into a nonthermal power law, despite potentially disruptive 3D instabilities. The nonthermal particle spectrum suggests that reconnection may explain observations of nonthermal spectra from distant astrophysical sources, such as the Crab Nebula's gamma-ray flares. A separate investigation of kinetic turbulence in relativistic electron-positron plasma has provided, for the first time using first-principles kinetic simulation, clear evidence that turbulent fluctuations can efficiently drive nonthermal particle acceleration to high energies, much like magnetic reconnection.

Both investigations show a similar dependence of the spectral power law index on the effective magnetization parameter, which characterizes the magnetic energy density relative to the relativistic inertia of the plasma. This similarity provides an important insight and benchmark for testing theoretical explanations of the underlying acceleration mechanisms. A comprehensive study of 2D electron-ion reconnection focused, for the first time, on the semi-relativistic regime of electron-ion reconnection. Results include the observation of electron acceleration to a nonthermal power law, the slope of which depends on magnetization parameters. Additionally, they found that reconnection decreases as magnetization parameters are lowered.

IMPACT Project results will lead to advances in the understanding of fundamental plasma physics processes and have important implications for modern high-energy astrophysics. It will enable first-principles modeling of plasma energization and nonthermal particle acceleration in systems, such as pulsar wind nebulae and accretion flows, and jets powered by black holes.

# ALCF Projects

## 2017 INCITE Projects

### Biological Sciences

**Biophysical Principles of Functional Synaptic Plasticity in the Neocortex**

PI NAME Eilif Muller
INST Blue Brain Project, EPFL
HOURS 100 Million Core-Hours

### Chemistry

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

PI NAME David Baker
INST University of Washington
HOURS 150 Million Core-Hours

**Understanding the Molecular Origin of Climate Change**

PI NAME Subramanian Sankaranarayanan
INST Argonne National Laboratory
HOURS 100 Million Core-Hours

### Computer Science

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

PI NAME Leonid Oliker
INST Lawrence Berkeley National Laboratory
HOURS 60 Million Core-Hours
ALCF: 25M; OLCF: 35M

### Earth Science

**Frontiers in Planetary and Stellar Magnetism Through High-Performance Computing**

PI NAME Jonathan Aurnou
INST University of California, Los Angeles
HOURS 260 Million Core-Hours

**Quantification of Uncertainty in Seismic Hazard Using Physics-Based Simulations**

PI NAME Thomas Jordan
INST University of Southern California
HOURS 141 Million Core-Hours
ALCF: 45M; OLCF: 96M

**High-Resolution Climate Change Simulations with the CESM**

PI NAME Gerald Meehl
INST National Center for Atmospheric Research
HOURS 215 Million Core-Hours

**Accelerated Climate Modeling for Energy**

PI NAME Mark Taylor
INST Sandia National Laboratories
HOURS 278 Million Core-Hours
ALCF: 128M; OLCF: 150M

### Engineering

**Multiscale Physics of the Ablative Rayleigh-Taylor Instability**

PI NAME Hussein Aluie
INST University of Rochester
HOURS 47 Million Core-Hours

**Adaptive DES of a Vertical Tail/Rudder Assembly with Active Flow Control**

PI NAME Kenneth E. Jansen
INST University of Colorado Boulder
HOURS 90 Million Core-Hours

**LES to Characterize Shock Boundary Layer Interaction in a 3D Transonic Turbofan**

PI NAME Umesh Paliath
INST GE Global Research
HOURS 90 Million Core-Hours

**Direct Numerical Simulation of Compressible, Turbulent Flow**

PI NAME Jonathan Poggie
INST Purdue University
HOURS 200 Million Core-Hours

### Materials Science

**Combining High-Accuracy Electronic Structure Methods to Study Surface Reactions**

PI NAME Maria Chan
INST Argonne National Laboratory
HOURS 100 Million Core-Hours

**Computational Spectroscopy of Heterogeneous Interfaces**

PI NAME Giulia Galli
INST The University of Chicago and Argonne National Laboratory
HOURS 200 Million Core-Hours

**Electronic Stopping in Condensed Matter Under Ion Irradiation**

PI NAME Yosuke Kanai
INST University of North Carolina
HOURS 130 Million Core-Hours

**Predictive Simulations of Functional Materials**

PI NAME Paul Kent
INST Oak Ridge National Laboratory
HOURS 138 Million Core-Hours
ALCF: 98M; OLCF: 40M

**Materials and Interfaces for Organic and Hybrid Photovoltaics**

PI NAME Noa Marom
INST Carnegie Mellon University
HOURS 160 Million Core-Hours

**Petascale Simulations for Layered Materials Genome**

PI NAME Aiichiro Nakano
INST University of Southern California
HOURS 140 Million Core-Hours

**Reactive Mesoscale Simulations of Tribological Interfaces**

PI NAME Subramanian Sankaranarayanan
INST Argonne National Laboratory
HOURS 50 Million Core-Hours

### Physics

**The Rate of Spontaneous Plasma Reconnection**

PI NAME Andrey Beresnyak
INST U.S. Naval Research Laboratory
HOURS 52 Million Core-Hours

**Global Radiation MHD Simulations of Massive Star Envelopes**

PI NAME Lars Bildsten
INST University of California, Santa Barbara
HOURS 60 Million Core-Hours

**Collider Physics at the Precision Frontier**

PI NAME Radja Boughezal
INST Argonne National Laboratory
HOURS 100 Million Core-Hours

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

PI NAME Sean Couch
INST Michigan State University
HOURS 100 Million Core-Hours

**Cosmic Reionization on Computers**

PI NAME Nickolay Gnedin
INST Fermilab
HOURS 75 Million Core-Hours

**Lattice QCD**

PI NAME Paul Mackenzie
INST Fermilab
HOURS 348 Million Core-Hours
ALCF: 240M; OLCF: 108M

**Fundamental Properties of QCD Matter Produced at RHIC and the LHC**

PI NAME Claudia Ratti
INST University of Houston
HOURS 194 Million Core-Hours

**Magnetohydrodynamic Models of Accretion Including Radiation Transport**

PI NAME James Stone
INST Princeton University
HOURS 54 Million Core-Hours

**Petascale Simulations of Laser Plasma Interaction Relevant to IFE**

PI NAME Frank Tsung
INST University of California, Los Angeles
HOURS 147 Million Core-Hours

**Nuclear Structure and Nuclear Reactions**

PI NAME James Vary
INST Iowa State University
HOURS 170 Million Core-Hours
ALCF: 80M; OLCF: 90M

**PICSSAR—Particle-In-Cell Spectral Scalable Accurate Relativistic**

PI NAME Jean-Luc Vay
INST Lawrence Berkeley National Laboratory
HOURS 100 Million Core-Hours

## 2016–2017 ALCC Projects

### Biological Sciences

**Molecular Dynamics Studies of Biomass Degradation in Biofuel Production**

PI NAME Klaus Schulten
INST University of Illinois at Urbana-Champaign
HOURS 50 Million Core-Hours

### Chemistry

**Molecular Modeling of Hot Electron Transfer for Solar Energy Conversion**

PI NAME Hanning Chen
INST George Washington University
HOURS 16 Million Core-Hours

### Computer Science

**Demonstration of the Scalability of Programming Environments by Simulating Multiscale Applications**

PI NAME Robert Voigt
INST Leidos
HOURS 191 Million Core-Hours
ALCF: 151M; OLCF: 40M

### Earth Science

**Delivering the Department of Energy's Next-Generation High-Resolution Earth System Model**

PI NAME Peter Thornton
INST Oak Ridge National Laboratory
HOURS 211 Million Core-Hours
ALCF: 158M; OLCF: 53M

### Engineering

**Multiphase Simulations of Nuclear Reactor Flows**

PI NAME Igor Bolotnov
INST North Carolina State University
HOURS 72.1 Million Core-Hours

**Computational Study of Cycle-to-Cycle Variation in Dual-Fuel Engines**

PI NAME Ravichandra Jupud
INST GE Global Research
HOURS 25 Million Core-Hours

**Unraveling Silent Owl Flight to Develop Ultra-Quiet Energy Conversion Machines**

PI NAME Anupam Sharma
INST Iowa State University
HOURS 25 Million Core-Hours

**Adjoint Based Optimization Via Large Eddy Simulation of a Fundamental Turbine Stator-Rotor**

PI NAME Qiqi Wang
INST Massachusetts Institute of Technology
HOURS 15 Million Core-Hours

### Materials Science

**First-Principles Design and Analysis of Energy-Efficient NanoElectronic Switches**

PI NAME Sefa Dag
INST GlobalFoundries
HOURS 10 Million Core-Hours

**Computational Engineering of Defects in Soft and Hard Materials for Energy and Quantum Information Applications**

PI NAME Marco Govoni
INST The University of Chicago and Argonne National Laboratory
HOURS 53.7 Million Core-Hours

**Modeling of Intense X-ray Laser Dynamics in Nanoclusters**

PI NAME Phay Ho
INST Argonne National Laboratory
HOURS 10 Million Core-Hours

**The Materials Project—Completing the Space of Elastic and Piezoelectric Tensors**

PI NAME Kristin Persson
INST Lawrence Berkeley National Laboratory
HOURS 36 Million Core-Hours

**Predictive Modeling of Functional Nanoporous Materials, Nanoparticle Assembly, and Reactive Systems**

PI NAME J. Ilja Siepmann
INST University of Minnesota
HOURS 117 Million Core-Hours

**First-Principles Design of Magnetic Materials, Models, and Mechanisms**

PI NAME Lucas Wagner
INST University of Illinois at Urbana-Champaign
HOURS 30 Million Core-Hours

**Modeling Helium-Hydrogen Plasma-Mediated Tungsten Surface Response to Predict Fusion Plasma-Facing Component Performance in ITER**

PI NAME Brian Wirth
INST University of Tennessee
HOURS 95 Million Core-Hours
ALCF: 70M; OLCF: 25M

## Physics

**High-Intensity Multibunch Physics in the Fermilab Accelerator Complex**

PI NAME James Amundson
INST Fermilab
HOURS 50 Million Core-Hours

**Hadronic Light-by-Light Scattering Contribution to the Muon Anomalous Magnetic Moment from Lattice QCD with Chiral Fermions**

PI NAME Thomas Blum
INST University of Connecticut
HOURS 180 Million Core-Hours

**Extreme-Scale Gyrokinetic Particle Simulations to Complete the 2016 OFES National Theory/ Simulation Performance Target and to Study the Fundamental Edge Physics**

PI NAME Choong-Seock Chang
INST Princeton Plasma Physics Laboratory
HOURS 175 Million Core-Hours
ALCF: 100M; OLCF: 75M

**An End-Station for Intensity and Energy Frontier Experiments and Calculations**

PI NAME Taylor Childers
INST Argonne National Laboratory
HOURS 106.5 Million Core-Hours
ALCF: 93.5M; NERSC: 13M

***Ab Initio* Modeling of the Dynamical Stability of HED Plasmas: From Fusion to Astrophysics**

PI NAME Frederico Fiuza
INST SLAC National Accelerator Laboratory
HOURS 60 Million Core-Hours

**Exploring Higgs Compositeness Mechanism in the Era of the 14 TeV LHC**

PI NAME George Fleming
INST Yale University
HOURS 55 Million Core-Hours

**Nuclear Structure for Tests of Fundamental Symmetries and Astroparticle Physics**

PI NAME Calvin Johnson
INST San Diego State University
HOURS 30 Million Core-Hours
ALCF: 6M; NERSC: 24M

**Muon g-2 Hadronic Vacuum Polarization from Lattice QCD**

PI NAME John Laiho
INST Syracuse University
HOURS 66 Million Core-Hours

**61-Pin Wire-Wrap Turbulent Conjugate-Heat Transfer: V&V for Industry and SESAME**

PI NAME Elia Merzari
INST Argonne National Laboratory
HOURS 120 Million Core-Hours

**Numerical Simulation of Turbulent Flows in Advanced Steam Generators**

PI NAME Aleksandr Obabko
INST Argonne National Laboratory
HOURS 80 Million Core-Hours

**Simulations of Laser Experiments to Study the Origin of Cosmic Magnetic Fields**

PI NAME Petros Tzeferacos
INST The University of Chicago
HOURS 60 Million Core-Hours

## 2017–2018 ALCC Projects

## Biological Sciences

**Multiscale Simulations of Hematological Disorders**

PI NAME George Karniadakis
INST Brown University
HOURS 46 Million Core-Hours
ALCF: 20M; OLCF: 26M

**Protein-Protein Recognition and HPC Infrastructure**

PI NAME Benoît Roux
INST The University of Chicago and Argonne National Laboratory
HOURS 80 Million Core-Hours

## Chemistry

**Quantum Monte Carlo Computations of Chemical Systems**

PI NAME Olle Heinonen
INST Argonne National Laboratory
HOURS 5 Million Core-Hours

**Spin-Forbidden Catalysis on Metal-Sulfur Proteins**

PI NAME Sergey Varganov
INST University of Nevada, Reno
HOURS 42 Million Core-Hours

## Computer Science

**ECP Consortium for Exascale Computing**

PI NAME Paul Messina
INST Argonne National Laboratory
HOURS 969 Million Core-Hours
ALCF: 530M; OLCF: 300M; NERSC: 139M

**Portable Application Development for Next-Generation Supercomputer Architectures**

PI NAME Tjerk Straatsma
INST Oak Ridge National Laboratory
HOURS 60 Million Core-Hours
ALCF: 20M; OLCF: 20M; NERSC: 20M

**Demonstration of the Scalability of Programming Environments by Simulating Multiscale Applications**

PI NAME Robert Voigt
INST Leidos
HOURS 156.8 Million Core-Hours
ALCF: 110M; OLCF: 46.8M

## Earth Science

**Large-Eddy Simulation Component of the Mesoscale Convective System Climate Model Development and Validation (CMDV-MCS) Project**

PI NAME William Gustafson
INST Pacific Northwest National Laboratory
HOURS 74 Million Core-Hours

**Understanding the Role of Ice Shelf-Ocean Interactions in a Changing Global Climate**

PI NAME Mark Petersen
INST Los Alamos National Laboratory
HOURS 87 Million Core-Hours
ALCF: 25M; OLCF: 2M; NERSC: 60M

## Engineering

**Non-Boussines Effects on Buoyancy-Driven Variable Density Turbulence**

PI NAME Daniel Livescu
INST Los Alamos National Laboratory
HOURS 60 Million Core-Hours

**Numerical Simulation of Turbulent Flows in Advanced Steam Generators—Year 3**

PI NAME Aleksandr Obabko
INST Argonne National Laboratory
HOURS 50 Million Core-Hours

## Materials Science

**Computational Engineering of Electron-Vibration Coupling Mechanisms**

PI NAME Marco Govoni
INST Argonne National Laboratory
HOURS 75 Million Core-Hours
ALCF: 60M; OLCF: 15M

**Imaging Transient Structures in Heterogeneous Nanoclusters in Intense X-ray Pulses**

PI NAME Phay Ho
INST Argonne National Laboratory
HOURS 68 Million Core-Hours

**Predictive Modeling of Functional Nanoporous Materials, Nanoparticle Assembly, and Reactive Systems**

PI NAME J. Ilja Siepmann
INST University of Minnesota
HOURS 146 Million Core-Hours
ALCF: 130M; NERSC: 16M

**Modeling Helium-Hydrogen Plasma-Mediated Tungsten Surface Response to Predict Fusion Plasma-Facing Component Performance**

PI NAME Brian Wirth
INST Oak Ridge National Laboratory
HOURS 173 Million Core-Hours
ALCF: 98M; OLCF: 75M

## Physics

**Hadronic Light-by-Light Scattering and Vacuum Polarization Contributions to the Muon Anomalous Magnetic Moment from Lattice QCD with Chiral Fermions**

PI NAME Thomas Blum
INST University of Connecticut
HOURS 220 Million Core-Hours

**High-Fidelity Gyrokinetic Study of Divertor Heat-Flux Width and Pedestal Structure**

PI NAME Choong-Seock Chang
INST Princeton Plasma Physics Laboratory
HOURS 269.9 Million Core-Hours
ALCF: 80M; OLCF: 100M; NERSC: 89.9M

**Simulating Particle Interactions and the Resulting Detector Response at the LHC and Fermilab**

PI NAME John T. Childers
INST Argonne National Laboratory
HOURS 188 Million Core-Hours
ALCF: 58M; OLCF: 80M; NERSC: 50M

**Studying Astrophysical Particle Acceleration in HED Plasmas**

PI NAME Frederico Fiuza
INST SLAC National Accelerator Laboratory
HOURS 50 Million Core-Hours

**Extreme-Scale Simulations for Multi-Wavelength Cosmology Investigations**

PI NAME Katrin Heitmann
INST Argonne National Laboratory
HOURS 125 Million Core-Hours
ALCF: 40M; OLCF: 10M; NERSC: 75M

**Nuclear Spectra with Chiral Forces**

PI NAME Alessandro Lovato
INST Argonne National Laboratory
HOURS 35 Million Core-Hours

**High-Fidelity Numerical Simulation of Wire-Wrapped Fuel Assemblies**

PI NAME Elia Merzari
INST Argonne National Laboratory
HOURS 85 Million Core-Hours

**Elimination of Modeling Uncertainties Through High-Fidelity Multiphysics Simulation to Improve Nuclear Reactor Safety and Economics**

PI NAME Emily Shemon
INST Argonne National Laboratory
HOURS 44 Million Core-Hours

**Nucleon Structure and Electric Dipole Moments with Physical Chirally Symmetric Quarks**

PI NAME Sergey Syritsyn
INST RIKEN BNL Research Center
HOURS 135 Million Core-Hours

# 2017 Director's Discretionary Projects

## Biological Sciences

**Exascale Deep Learning and Simulation-Enabled Precision Medicine for Cancer**

PI NAME Rick Stevens
INST Argonne National Laboratory
HOURS 1 Million Core-Hours

**Angora Scaling Study**

PI NAME Allen Taflove
INST Northwestern University
HOURS 9.9 Million Core-Hours

## Chemistry

**Ensemble Representation for the Realistic Modeling of Cluster Catalysts at Heterogeneous Interfaces**

PI NAME Anastassia N. Alexandrov
INST University of California, Los Angeles
HOURS 2.5 Million Core-Hours

**Spectrum Slicing Eigensolver for *Ab Initio* Simulations**

PI NAME Murat Keceli
INST Argonne National Laboratory
HOURS 2 Million Core-Hours

**Accurate Calculations of the Binding Energies of Dipole-Bound Anions**

PI NAME Brenda M. Rubenstein
INST Brown University
HOURS 7.5 Million Core-Hours

**Comparison of Quantum Methods for Investigation of Protein-Ligand Binding Interactions**

PI NAME Vipin Sachdeva
INST IBM Research
HOURS 1 Million Core-Hours

**Fragment Molecular Orbital (FMO) and Effective Fragment Molecular Orbital (EFMO) Methods for Exascale Computing**

PI NAME Federico Zaharie
INST Iowa State University, Ames
HOURS 6 Million Core-Hours

## Computer Science

**SciDAC Scalable Data Management Analysis and Visualization**

PI NAME Michael E. Papka and Joseph A. Insley
INST Argonne National Laboratory
HOURS 1 Million Core-Hours

**MPICH—A High-Performance and Widely Portable MPI Implementation**

PI NAME Ken Raffenetti
INST Argonne National Laboratory
HOURS 26 Million Core-Hours

**Interfacial Behavior of Alcohol at Water/Organic Biphasic System**

PI NAME Baofu Qiao
INST Argonne National Laboratory
HOURS 5 Million Core-Hours

**Portable Application Development for Next-Generation Supercomputer Architectures**

PI NAME Tjerk Straatsma, Katerina B. Antypas, Timothy J. Williams
INST Oak Ridge National Laboratory, Lawrence Berkeley National Laboratory, Argonne National Laboratory
HOURS 5 Million Core-Hours

**ExaHDF5: Advancing HDF5 HPC I/O to Enable Scientific Discovery**

PI NAME Venkat Vishwanath
INST Argonne National Laboratory
HOURS 2 Million Core-Hours

**Extreme Many-Task Computing with Swift**

PI NAME Justin M. Wozniak
INST Argonne National Laboratory
HOURS 1.2 Million Core-Hours

## Earth Science

**Simulating Global Terrestrial Carbon Sequestration and Carbon Transport to Aquatic Ecosystems—Pilot Study**

PI NAME Jinxun Liu
INST U.S. Geological Survey
HOURS 2 Million Core-Hours

## Energy Technologies

**Tri Alpha Energy FRC Whole Device Modeling —Phase 1**

PI NAME David Leinweber
INST Lawrence Berkeley National Laboratory
HOURS 5 Million Core-Hours

**Modeling and Prediction of Nanocatalyst for Fuel Cells**

PI NAME Binay Prasai
INST Central Michigan University
HOURS 2 Million Core-Hours

**NEAMS Neutronics Verification and Validation Simulations**

PI NAME Emily Shemon
INST Argonne National Laboratory
HOURS 7 Million Core-Hours

**Domain Decomposed 3D Method of Characteristics Reactor Simulation Study**

PI NAME John Tramm
INST Massachusetts Institute of Technology
HOURS 1 Million Core-Hours

**Advancing Understanding of Fission Gas Behavior in Nuclear Fuel Through Leadership Class Computing**

PI NAME Brian Wirth
INST The University of Tennessee, Knoxville
HOURS 3.5 Million Core-Hours

## Engineering

**Multiphase Simulations of Nuclear Reactor Thermal Hydraulics**

PI NAME Igor A. Bolotnov
INST North Carolina State University
HOURS 5 Million Core-Hours

**Variable-Density Fluid Dynamics**

PI NAME Paul E. Dimotakis
INST California Institute of Technology
HOURS 10 Million Core-Hours

**Numerical Simulation of Acoustic Radiation from High-Speed Turbulent Boundary Layers**

PI NAME Lian Duan
INST Missouri University of Science and Technology
HOURS 3 Million Core-Hours

**Variable-Density Turbulence Under Variable Acceleration**

PI NAME Daniel Livescu
INST Los Alamos National Laboratory
HOURS 5 Million Core-Hours

**Simulation of Supersonic Combustion**

PI NAME Farzad Mashayek
INST University of Illinois at Chicago
HOURS 4 Million Core-Hours

**Data Analysis of Turbulent Channel Flow at High Reynolds Number**

PI NAME Robert D. Moser
INST The University of Texas at Austin
HOURS 7 Million Core-Hours

**Influence on Duct Corner Geometry on Secondary Flow: Convergence from Duct to Pipe Flow**

PI NAME Hassan M. Nagib
INST Illinois Institute of Technology
HOURS 3 Million Core-Hours

**Large-Eddy Simulation for the Prediction and Control of Impinging Jet Noise**

PI NAME Joseph W. Nichols
INST University of Minnesota
HOURS 5 Million Core-Hours

**Investigation of a Low-Octane Gasoline Fuel for a Heavy-Duty Diesel Engine in a Low-Temperature Combustion Regime**

PI NAME Sibendu Som
INST Argonne National Laboratory
HOURS 5.9 Million Core-Hours

## Materials Science

**Large-Scale *Ab Initio* Simulations of Crystal Defects**

PI NAME Kaushik Bhattacharya
INST California Institute of Technology
HOURS 2 Million Core-Hours

**Quantum Monte Carlo Study of Spin-Crossover Transition in Fe(II)-Based Complexes**

PI NAME Hanning Chen
INST George Washington University
HOURS 4 Million Core-Hours

**Unveiling the Behavior of UO2 under Extreme Physical Conditions**

PI NAME Peter Littlewood and Laura Ratcliff
INST Argonne National Laboratory
HOURS 5 Million Core-Hours

**Structure and Properties of Grain Boundaries in Materials for Energy Applications**

PI NAME Wissam Saidi
INST University of Pittsburgh
HOURS 2 Million Core-Hours

## Mathematics

**CEED**

PI NAME Misun Min
INST Argonne National Laboratory
HOURS 2 Million Core-Hours

## Physics

**Galaxies on FIRE: Shedding Light on Dark Matter**

PI NAME Philip F. Hopkins
INST California Institute of Technology
HOURS 2 Million Core-Hours

**Plasma Turbulence Near the Sun**

PI NAME Jean C. Perez
INST Florida Institute of Technology
HOURS 2 Million Core-Hours

**DIII-D Computing**

PI NAME David Schissel
INST General Atomics
HOURS 8,000 Core-Hours

**Extreme-Scale Turbulence Simulations**

PI NAME William M. Tang
INST Princeton Plasma Physics Laboratory
HOURS 5 Million Core-Hours

## 2016–2017 ADSP Projects

**Advancing the Scalability of LHC Workflows to Enable Discoveries at the Energy Frontier**

PI NAME Taylor Childers
INST Argonne National Laboratory-

**Data-Driven Molecular Engineering of Solar-Powered Windows**

PI NAME Jacqueline M. Cole
INST University of Cambridge

**Leveraging Non-Volatile Memory, Big Data, and Distributed Workflow Technology to Leap Forward Brain Modeling**

PI NAME Fabien Delalondre
INST Blue Brain Project, EPFL

**Large-Scale Computing and Visualization on the Connectomes of the Brain**

PI NAME Doga Gursoy
INST Argonne National Laboratory

## Theta ESP Projects

**Scale-Resolving Simulations of Wind Turbines with SU2**

PI NAME Juan J. Alonso
INST Stanford University

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

PI NAME Volker Blum
INST Duke University

**Large-Scale Simulation of Brain Tissue**

PI NAME Fabien Delalondre
INST Blue Brain Project, EPFL

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

PI NAME Christos Frouzakis
INST ETH Zürich

**First-Principles Simulations of Functional Materials for Energy Conversion**

PI NAME Giulia Galli
INST The University of Chicago and Argonne National Laboratory

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

PI NAME Mark Gordon
INST Iowa State University

**Next-Generation Cosmology Simulations with HACC: Challenges from Baryons**

PI NAME Katrin Heitmann
INST Argonne National Laboratory

**Adaptive DES of a Vertical Tail/Rudder Assembly with Active Flow Control**

PI NAME Kenneth E. Jansen
INST University of Colorado Boulder

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

PI NAME Alexei Khokhlov
INST The University of Chicago

**Quantum Monte Carlo Calculations in Nuclear Theory**

PI NAME Steven C. Pieper
INST Argonne National Laboratory

**Free Energy Landscapes of Membrane Transport Proteins**

PI NAME Benoît Roux
INST The University of Chicago and Argonne National Laboratory

## Aurora ESP Projects

**Extending Moore's Law Computing with Quantum Monte Carlo**

PI NAME Anouar Benali
INST Argonne National Laboratory

**High-Fidelity Simulation of Fusion Reactor Boundary Plasmas**

PI NAME C.S. Chang
INST Princeton Plasma Physics Laboratory

**NWChemEx: Tackling Chemical, Materials and Biochemical Challenges in the Exascale Era**

PI NAME Thomas Dunning
INST University of Washington

**Extreme-Scale Simulations for Multi-Wavelength Cosmology Investigations**

PI NAME Katrin Heitmann
INST Argonne National Laboratory

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

PI NAME Kenneth Jansen
INST University of Colorado Boulder

**Lattice QCD**

PI NAME Paul Mackenzie
INST Fermilab

**Metascalable Layered Materials Genome**

PI NAME Aiichiro Nakano
INST University of Southern California

**Free Energy Landscapes of Membrane Transport Proteins**

PI NAME Benoît Roux
INST The University of Chicago and Argonne National Laboratory

## About Argonne Leadership Computing Facility

Argonne's Leadership Computing Facility Division operates the Argonne Leadership Computing Facility (ALCF) as part of the U.S. Department of Energy's effort to provide leadership-class computing resources to the scientific community. The ALCF is supported by the DOE 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 of Chicago, at 9700 South Cass Avenue, Argonne, Illinois 60439. For information about Argonne and its pioneering science and technology programs, visit www.anl.gov.

## Availability of this Report (ANL/ALCF-17/6)

Online Access: U.S. Department of Energy (DOE) reports produced after 1991 and a growing number of pre-1991 documents are available for 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**

**phone | 800.553.NTIS (6847) or 703.605.6000**
**fax | 703.605.6900**
**orders@ntis.gov**
**www.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**
**www.osti.gov**


ALCF Leadership
Michael E. Papka
*Division Director; Deputy Associate Laboratory Director, Computing, Environment, and Life Sciences*

Jini Ramprakash
*Deputy Division Director*

Katherine Riley
*Director of Science*

Mark Fahey
*Director of Operations*

Editorial Team
Beth Cerny (project lead), Laura Wolf (lead writer), Jim Collins, and John Spizzirri

Design and Production
Sandbox Studio, Chicago

Argonne NATIONAL LABORATORY

**CONTACT**
ALCF Communications
media@alcf.anl.gov
alcf.anl.gov

U.S. DEPARTMENT OF ENERGY

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