

**CONTACT**

[media@alcf.anl.gov](mailto:media@alcf.anl.gov)


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

# Science

2023

ARGONNE LEADERSHIP COMPUTING FACILITY

2023 Science Report



Researchers from the University of Dayton Research Institute and the Air Force Research Laboratory are using ALCF supercomputers to carry out quantum mechanically guided simulations of hypersonic flows in thermal and chemical non-equilibrium. This image shows the velocity field from a calculation to simulate a Mach 8.2 experiment. *Image: ALCF Visualization and Data Analytics Team; University of Dayton Research Institute*

## CONTENTS

03 Message from ALCF Leadership

### 04 Argonne Leadership Computing Facility

---

06 About ALCF

07 ALCF Team

08 ALCF Computing Resources

### 10 Advancing HPC and AI for Science

---

12 Standing up Aurora

15 Early Application Performance Highlights

22 ALCF AI Testbed Deploys New Systems

24 Creating an Integrated Research Infrastructure

## 28 Science

---

### 30 Accessing ALCF Resources for Science

---

#### 32 2023 ALCF Science Highlights

#### 33 Biological Sciences

Accelerating Advances in Biomedicine with HPC and AI  
*Ravi Madduri*

Design of Peptides and Proteins on Classical and Quantum Computing Hardware  
*Vikram Mulligan*

Identifying Reaction Coordinates for Complex Protein Dynamics  
*Ao Ma*

Multiphysics Modeling of Blood Flow with Cell Suspensions in Patient-Specific Capillary Network  
*Jifu Tan*

#### 37 Chemistry

Accelerated Catalyst Discovery from First-Principles Simulations and Machine Learning  
*Rejeev Assary*

Microscopic Insight into Transport Properties of Li-Battery Electrolytes  
*Wei Jiang*

#### 39 Computer Science

Linking Scientific Instruments and Computation: Patterns, Technologies, and Experiences  
*Rafael Vescovi*

Standardizing Complex Numbers in SYCL  
*Thomas Applencourt*

#### 41 Earth Science

High-Resolution Regional Climate Model Simulations  
*Rao Kotamarthi*

Multirate Methods for Coupled Compressible Navier-Stokes Systems  
*Shinhoo Kang*

Projected U.S. Drought Extremes through the Twenty-First Century with Vapor Pressure Deficit  
*Rao Kotamarthi*

#### 44 Energy Technologies

Assessment of Turbulent Prandtl Number for Heavy Liquid Metal Flow in a Bare Rod Bundle  
*Yiqi Yu, Emily Shemon*

Automatic Building Energy Modeling and Analysis  
*Joshua New*

High-Fidelity Simulations of Gas Turbine Combustors for Sustainable Aviation Applications  
*Muhsin Ameen, Sicong Wu*

#### 47 Engineering

DNS of Wall-Bounded Magnetohydrodynamic Turbulence at High Reynolds Number  
*Myoungkyu Lee*

First-Principles Simulation of Hypersonic Flight  
*Maninder Grover*

Online Machine Learning for Large-Scale Turbulent Simulations  
*Kenneth Jansen*

Turbulent Rayleigh-Bénard Convection in Suspensions of Bubbles  
*Parisa Mirbod*

#### 51 Materials Science

AI-Guided Exascale Simulations of Quantum Materials Manufacturing and Control  
*Aiichiro Nakano*

End-to-End AI Framework for Small Molecules and Inorganic Crystals  
*Eliu Huerta*

Large-Scale Simulations of Light-Activated Matter  
*Giulia Galli*

Materials Informatics Study of Two-Dimensional Magnetic Materials and Their Heterostructures  
*Trevor Rhone*

Towards DMC Accuracy Across Chemical Space with Scalable  $\Delta$ -QML  
*Paul Kent, Anouar Benali*

#### 56 Physics

Electron Kinetic Plasma Physics of Black Hole Accretion Flows  
*Dmitri Uzdensky*

First-Principles Simulations of Black Hole Accretion Flows and Coronae  
*Luca Comisso*

Physics-Inspired AI Modeling of Complex Multiscale and Multiphysics Systems  
*Eliu Huerta*

## 60 ALCF Projects

---

**MICHAEL E. PAPKA**

ALCF Director

**BILL ALLCOCK**

ALCF Director of Operations

**SUSAN COGLAN**

ALCF-X Project Director

**KALYAN KUMARAN**

ALCF Director of Technology

**JINI RAMPRAKASH**

ALCF Deputy Director

**KATHERINE RILEY**

ALCF Director of Science

## MESSAGE FROM ALCF LEADERSHIP

Last year, the Argonne Leadership Computing Facility (ALCF) introduced Polaris, its largest GPU-accelerated system to date. Sunspot, a testing and development system built from actual Aurora hardware, was also launched. Sunspot's identical architecture is an ideal "on-ramp" for Aurora, helping research teams in DOE's Exascale Computing Project and ALCF's Aurora Early Science Program to optimize allocation performance for the exascale era.

Polaris hosts a range of projects, including machine learning and big data analysis jobs, along with traditional simulation science campaigns from INCITE, ALCC, and other programs. It also serves as a development platform for merging ALCF resources with large-scale experimental facilities.

Polaris and Theta, ALCF's bedrock production system, together push the boundaries of knowledge across disciplines. They help unravel the mysteries of black hole accretion flows and supernovae, identify new materials for the batteries of the future, advance areas of biomedical research, and look at future forecasts of droughts over the remainder of the century. These and many other science highlights are discussed in the pages that follow.

We are pleased to inform you that all 10,624 of Aurora's blades have been installed and ALCF staff continues to work on ensuring applications will be ready for Aurora when it is deployed. We understand that the deployment of Aurora is the facility's most important current activity, and we are proud of our progress so far.

With the monumental deployment of Aurora nearly done, Nexus is a groundbreaking initiative from Argonne National Laboratory. It is a visionary stride towards integrating scientific facilities, supercomputing capabilities, and cutting-edge data technologies to enable next-generation, large-scale science.

Argonne researchers are shaping the future of inter-facility workflows by automating them and identifying ways to make these workflows reusable and adaptable for different experiments. Nexus's ambitious vision also extends to autonomous discovery and self-driving labs to streamline processes, conserve resources, and expedite the pace of discovery. Nexus stands as an emblem of the future of research—a testament to the power of integration and automation in catalyzing scientific discovery.


As we conclude, we want to reiterate our strong dedication to our mission of facilitating scientific breakthroughs through our unparalleled resources, capabilities, and profound research expertise. Our cutting-edge systems are the future of HPC, but our exceptional staff and diverse user community are leading the way in employing HPC for scientific advancements.

Thank you for your ongoing support.

---

## ARGONNE LEADERSHIP COMPUTING FACILITY

The ALCF enables breakthroughs in science and engineering by providing supercomputing resources and expertise to the research community.





The distribution of the debris of an explosion of the core of a 23-solar-mass progenitor model at ~4.3 seconds after core bounce and ~4.0 seconds after explosion. The physical scale is ~40,000 kilometers on a side. The outer translucent veil is the shock wave. The globular structures trace the debris with a preponderance of the isotope Ni-56. The coloring is in mass density, with blue being lower density and red being higher density. Carried out on the ALCF's Theta supercomputer, this is the longest-duration 3D full-physics core-collapse supernova calculation ever performed. *Image: ALCF Visualization and Data Analytics Team; Adam Burrows and Tianshu Wang, Princeton University*

## About ALCF



The ALCF is a national scientific user facility located at Argonne National Laboratory.

The Argonne Leadership Computing Facility (ALCF), a U.S. Department of Energy (DOE) Office of Science user facility at Argonne National Laboratory, enables breakthroughs in science and engineering by providing supercomputing and AI resources to the research community.

ALCF computing resources—available to researchers from academia, industry, and government agencies—support large-scale computing projects aimed at solving some of the world’s most complex and challenging scientific problems. Through awards of supercomputing time and support services, the ALCF enables researchers to accelerate the pace of discovery and innovation across a broad range of disciplines.

As a key player in the nation’s efforts to provide the most advanced computing resources for science, the ALCF is helping to chart new directions in scientific computing through a convergence of simulation, data science, and AI methods and capabilities.

Supported by the DOE’s Advanced Scientific Computing Research (ASCR) program, the ALCF and its partner organization, the Oak Ridge Leadership Computing Facility, operate leadership-class supercomputing resources that are orders of magnitude more powerful than the systems typically used for open scientific research.



# ALCF Team



The ALCF's talented and diverse staff make the facility one of the world's premier centers for scientific computing.

## Science

Computational scientists with multidisciplinary domain expertise work directly with ALCF users to maximize and accelerate their research efforts. In addition, the ALCF team applies broad expertise in data science, machine learning, data visualization and analysis, and mathematics to help application teams leverage ALCF resources to pursue data-driven discoveries.

With a deep knowledge of the ALCF computing environment and experience with a wide range of numerical methods, programming models, and computational approaches, staff scientists and performance engineers help researchers optimize the performance and productivity of simulation, data, and learning applications on ALCF systems.

## Operations

The ALCF's HPC systems administrators manage and support all ALCF computing systems, ensuring users have stable, secure, and highly available resources to pursue their scientific goals. This includes the ALCF's production supercomputers, supporting system environments, storage systems, and network infrastructure. The team's software developers create tools to support the ALCF computing environment, including software for user account and project management, job failure analysis, and job scheduling.

User support specialists provide technical assistance to ALCF users and manage the workflows for user accounts and projects. In the business intelligence space, staff data architects assimilate and verify ALCF data to ensure accurate reporting of facility information.

## Technology

The ALCF team plays a key role in designing and validating the facility's next-generation supercomputers. By collaborating with compute vendors and the performance tools community, staff members ensure the requisite programming models, tools, debuggers, and libraries are available on ALCF platforms. The team also helps manage Argonne's Joint Laboratory for System Evaluation, which houses next-generation testbeds that enable researchers to explore and prepare for emerging computing technologies.

ALCF computer scientists, performance engineers, and software engineers develop and optimize new tools and capabilities to facilitate science on the facility's current and future computing resources. This includes the deployment of scalable machine learning frameworks, in-situ visualization and analysis capabilities, data management services, workflow packages, and container technologies. In addition, the ALCF team is actively involved in programming language standardization efforts and contributes to cross-platform libraries to further enable the portability of HPC applications.

## Outreach

ALCF staff members organize and participate in training events that prepare researchers for efficient use of leadership computing systems. They also participate in a wide variety of educational activities aimed at cultivating a diverse and skilled HPC community for the future. In addition, staff outreach efforts include facilitating partnerships with industry and communicating to external audiences the impactful research enabled by ALCF resources.

# Supercomputing Resources

ALCF supercomputing resources support large-scale, computationally intensive projects aimed at solving some of the world's most complex and challenging scientific problems.

| SYSTEM NAME         | POLARIS                                       | THETA: KNL NODES                           | THETA: GPU NODES            | COOLEY                          |
|---------------------|---|--|-----------------------------|---------------------------------|
| Purpose             | Science Campaigns                             | Science Campaigns                          | Science Campaigns           | Data Analysis and Visualization |
| Architecture        | HPE Apollo 6500 Gen10+                        | Intel-Cray XC40                            | NVIDIA DGX A100             | Intel Haswell                   |
| Peak Performance    | 44 PF<br>(Tensor Core Double Precision)       | 11.7 PF                                    | 3.9 PF                      | 293 TF                          |
| Processors per Node | 3rd Gen AMD EPYC                              | 64-core, 1.3-GHz Intel Xeon Phi 7230       | 2 AMD EPYC 7742             | 2 6-core, 2.4-GHz Intel E5-2620 |
| GPU per Node        | 4 NVIDIA A100 Tensor Core                     | —  | 8 NVIDIA A100 Tensor Core   | NVIDIA Tesla K80                |
| Nodes               | 560   | 4,392                                      | 24                          | 126                             |
| Cores               | 17,920  | 281,088                                    | 3,072                       | 1,512                           |
| Memory              | 280 TB (DDR4); 87.5 TB (HBM)                  | 843 TB (DDR4); 70 TB (HBM)                 | 26 TB (DDR4); 8.32 TB (GPU) | 47 TB (DDR4); 3 TB (GDDR5)      |
| Interconnect        | HPE Slingshot 10 with Dragonfly Configuration | Aries Network with Dragonfly Configuration | NVIDIA QM8700 InfiniBand    | FDR InfiniBand                  |
| Racks               | 40  | 24   | 7                           | 6                               |

## ALCF AI TESTBED

The ALCF AI Testbed provides an infrastructure of next-generation AI-accelerator machines that allows researchers to evaluate the usability and performance of machine learning-based applications running on the systems. AI testbeds include:

| SYSTEM NAME  | CEREBRAS CS-2   | SAMBANOVA CARDINAL SN30                               | GROQRACK  | GRAPHCORE BOW POD-64                                   | HABANA GAUDI-1  |
|--|---|---|---|--|---|
| System Size  | 2 Nodes (Each with a Wafer-Scale Engine) Including MemoryX and SwarmX | 64 Accelerators (8 Nodes and 8 Accelerators per Node) | 72 Accelerators (9 Nodes and 8 Accelerators per Node) | 64 Accelerators (4 Nodes and 16 Accelerators per Node) | 16 Accelerators (2 Nodes and 8 Accelerators per Node) |
| Compute Units per Accelerator                          | 850,000 Cores   | 1,280 Programmable Compute Units                      | 5,120 Vector ALUs                                     | 1,472 Independent Processing Units                     | 8 TPC + GEMM Engine                                   |
| Estimated Performance of a Single Accelerator (TFlops) | > 5,780 (FP16)  | >660 (BF16)   | >188 (FP16)<br>>750 (INT8)                            | >250 (FP16)  | >150 (FP16)   |
| Software Stack Support                                 | Cerebras SDK, TensorFlow, PyTorch                                     | SambaFlow, PyTorch                                    | GroqWare SDK, ONNX                                    | PopART, TensorFlow, PyTorch, ONNX                      | SynapseAI, TensorFlow, PyTorch                        |
| Interconnect   | Ethernet-based  | Ethernet-based  | RealScale™  | IPU Link   | Ethernet-based  |

## DATA STORAGE SYSTEMS

ALCF disk storage systems provide intermediate-term storage for users to access, analyze, and share computational and experimental data. Tape storage is used to archive data from completed projects.

| SYSTEM NAME                  | EAGLE                 | GRAND                 | THETA-FSO         | SWIFT                        | TAPE STORAGE                  |
|------------------------------|-----------------------|-----------------------|-------------------|------------------------------|-------------------------------|
| File System                  | Lustre                | Lustre                | Lustre            | Lustre                       | —                             |
| Storage System               | HPE ClusterStor E1000 | HPE ClusterStor E1000 | HPE Sonexion L300 | All NVMe Flash Storage Array | LTO6 and LTO8 Tape Technology |
| Usable Capacity              | 100 PB                | 100 PB                | 9 PB              | 123 TB                       | 300 PB                        |
| Sustained Data Transfer Rate | 650 GB/s              | 650 GB/s              | 240 GB/s          | 48 GB/s                      | —                             |
| Disk Drives                  | 8,480                 | 8,480                 | 2,300             | 24                           | —                             |

## NETWORKING

Networking is the fabric that ties all of the ALCF’s computing systems together. InfiniBand enables communication between system I/O nodes and the ALCF’s various storage systems. The Production HPC SAN is built upon NVIDIA Mellanox High Data Rate (HDR) InfiniBand hardware. Two 800-port core switches provide the backbone links between 80 edge switches, yielding 1600 total available host ports, each at 200 Gbps, in a non-blocking fat-tree topology. The full bisection bandwidth of this fabric is 320 Tbps. The HPC SAN is maintained by the NVIDIA Mellanox Unified Fabric Manager (UFM), providing Adaptive Routing to avoid congestion, as well as the NVIDIA Mellanox Self-Healing Interconnect Enhancement for IntelLigent Datacenters (SHIELD) resiliency system for link fault detection and recovery.

When external communications are required, Ethernet is the interconnect of choice. Remote user access, systems maintenance and management, and high-performance data transfers are all enabled by the Local Area Network (LAN) and Wide Area Network (WAN) Ethernet infrastructure. This connectivity is built upon a combination of Extreme Networks SLX and MLXe routers and NVIDIA Mellanox Ethernet switches.

ALCF systems connect to other research institutions over multiple 100 Gbps Ethernet circuits that link to many high performance research networks, including local and regional networks like the Metropolitan Research and Education Network (MREN), as well as national and international networks like the Energy Sciences Network (ESnet) and Internet2.

## JOINT LABORATORY FOR SYSTEM EVALUATION

Argonne’s Joint Laboratory for System Evaluation (JLSE) provides access to leading-edge testbeds for exploratory research aimed at evaluating future extreme-scale computing systems, technologies, and capabilities.

---

**Florentia: Test and development system equipped with early versions of the Sapphire Rapids CPUs and Ponte Vecchio GPUs that will power Aurora**

---

**Arcticus, DevEP, Iris: Intel discrete and integrated GPU testbeds to support the development, optimization, and scaling of applications and software for Aurora**

---

**Aurora Software Development Kit: Frequently updated version of the publicly available Intel oneAPI toolkit for Aurora development**

---

**Arm Ecosystem: Apollo 80 Fujitsu A64FX Arm system, NVIDIA Ampere Arm and A100 test kits, and an HPE Comanche with Marvell ARM64 CPU platform provide an ecosystem for porting applications and measuring performance on next-generation systems**

---

**Presque: Intel DAOS nodes for testing the Aurora storage system**

---

**Edge Testbed: NVIDIA Jetson Xavier and Jetson Nano platforms provide a resource for testing and developing edge computing applications**

---

**NVIDIA and AMD GPUs: Clusters of NVIDIA V100, A100, and A40 GPUs, and AMD MI50 and MI100 GPUs for preparing applications for heterogeneous computing architectures**

---

**NVIDIA Bluefield-2 DPU SmartNICs: Platform used for confidential computing, MPICH offloading, and APS data transfer acceleration**

---

**NextSilicon Maverick: First-generation product being tested by Argonne researchers**


---


**Atos Quantum Learning Machine: Platform for testing and developing quantum algorithms and applications**

---

## ADVANCING HPC AND AI FOR SCIENCE

With the launch of new supercomputing and AI resources and capabilities, the ALCF is enabling pioneering research at the intersection of simulation, big data analytics, and machine learning.





A simulated view of Category 5 Hurricane Maria as it approached Puerto Rico on September 20, 2017. The field shows the radar reflectivity (dBZ) from 1 km simulations carried out on the ALCF's ThetaGPU system. This is a snapshot from a three domain setup, which has 50 million grid points in the innermost domain (1000x1000x51) that corresponds to the visualized portion of the hurricane. *Image: ALCF Visualization and Data Analytics Team; Chunyong Jung, Rao Kotamarthi, Gokhan Sever, and Jiali Wang, Argonne National Laboratory*

# Standing up Aurora

The ALCF has made significant progress in deploying its exascale supercomputer, completing the hardware installation and helping application teams get ramped up for early science campaigns.

Occupying the space of two professional basketball courts and weighing 600 tons, the ALCF's massive Aurora exascale supercomputer is now fully installed in the facility's data center.

Aurora will be theoretically capable of delivering more than two exaflops of computing power when it's deployed for science. The new supercomputer will follow the ALCF's previous and current systems—Intrepid, Mira, Theta, and Polaris—to deliver on the facility's mission to provide leading-edge supercomputing resources that enable breakthroughs in science and engineering. Open to researchers from across the world, ALCF supercomputers are used to tackle a wide range of scientific problems including designing more efficient airplanes, investigating the mysteries of the cosmos, modeling the impacts of climate change, and accelerating the discovery of new materials.

As Argonne's largest and most powerful supercomputer to date, the lab had to complete some substantial facility upgrades to get ready for the Intel-Hewlett Packard Enterprise (HPE) system, including adding new data center space, mechanical rooms, and equipment that significantly increase the building's power and cooling capacity. With a footprint of 10,000 square feet, Aurora is made up of eight rows of supercomputer cabinets that stand over eight feet tall. The cabinets are outfitted with more than 300 miles of networking cables, countless red and blue hoses

that pipe water in and out to cool the system, and specialized piping and equipment that bring the water in from beneath the data center floor and the electrical power from the floor above.

Over the past year, the system has taken shape with the delivery and installation of its computer racks and components, including its Intel CPUs and GPUs, the DAOS (Distributed Asynchronous Object Storage) storage system, and the HPE Slingshot interconnect technology. In June, the installation of Aurora's 10,624th and final "blade" marked a major milestone for the highly anticipated exascale supercomputer. As the backbone of the system, Aurora's blades house its processors, memory, networking, and cooling technologies. The machine gets its computational muscle from a combination of state-of-the-art Intel CPUs and GPUs. Each blade is equipped with two Intel Xeon CPU Max Series processors with high bandwidth memory and six Intel Data Center GPU Max Series processors.

The ALCF also deployed the Sunspot test and development system late last year. Outfitted with the exact same technologies as Aurora, the two-rack testbed gives researchers an ideal environment for multi-node testing to help them further optimize applications and workloads for the exascale supercomputer.





The Aurora installation team celebrates a milestone: the installation of the supercomputer's 10,624th and final blade.



A technician uses a specialized machine to install compute blades into the Aurora system.

### Preparing for Science on Day One

To help ensure researchers can use Aurora for science on day one, teams participating in DOE's Exascale Computing Project (ECP) and the ALCF's Aurora Early Science Program (ESP) continue their work to port and optimize dozens of key scientific computing applications. With access to the Aurora software development kit and Sunspot, ECP and ESP researchers have made significant progress in improving the performance and functionality of various codes, frameworks, and libraries using the programming models that will be supported on Aurora.

Training events, such as workshops, hackathons, and webinars, have been an important mechanism for providing guidance on application development and disseminating the latest details on hardware and software. In June, for example, the Argonne-Intel Center for Excellence hosted a multi-day workshop for ECP and ESP teams to provide updates on the system, share approaches and best practices for performance portability, and facilitate hands-on sessions with exascale software tools.

Now that Aurora is fully assembled, ECP and ESP team members are beginning to transition their work to the supercomputer to ready their applications for full system runs. These early users also help to stress test the supercomputer and identify potential bugs that need to be resolved ahead of its deployment.

### Early Performance Results

The ECP and ESP teams' initial runs on the Aurora GPUs have been promising compared to leading alternative GPUs. Some of the early application performance results include:

---

As part of the ECP ExaSMR (Exascale Small Modular Reactor) project, researchers achieved 30-70 percent performance improvements with NekRS, a GPU-oriented thermal-fluids simulation code, across a set of benchmark problems.

---

Another ExaSMR code, OpenMC, which is used for neutron and photon transport simulations, showed a 205 percent performance advantage on the Intel GPUs.

---

Supported by ESP and ECP projects, the Argonne-developed Hardware/Hybrid Accelerated Cosmology Code (HACC) has seen 2.6x speedups in early runs on the hardware.

---

QMCPACK, a quantum Monte Carlo code used for electronic structure calculations, has shown a 50 percent improvement in runs thus far. QMCPACK's exascale development is supported by both ESP and ECP.

---

XGC, a fusion plasma simulation code that is also supported by ESP and ECP, has performed 60 percent faster using an initial test problem.

---

The following pages provide additional details on the teams' efforts to optimize these applications for Aurora.



# Early Application Performance: HACC


PI Katrin Heitmann, Salman Habib,  
Argonne National Laboratory  
PROG Aurora Early Science Program and  
Exascale Computing Project

HACC (Hardware/Hybrid Accelerated Cosmology Code) is a cosmological N-body and hydrodynamics simulation code designed to run at extreme scales on HPC systems operated by DOE national laboratories. HACC computes the complicated emergence of structure in the universe across cosmological history, the core of the code's functionality consisting of gravitational calculations along with the more recent addition of gas dynamics and astrophysical subgrid models.

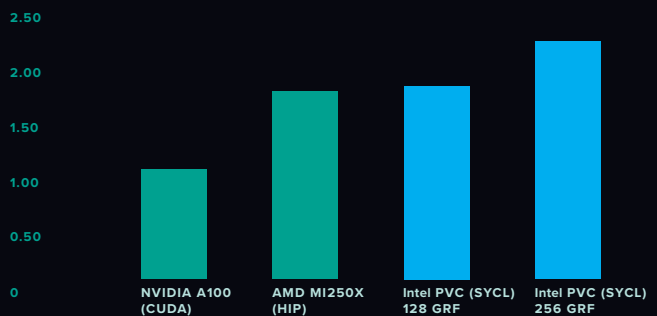
HACC is structured to remain mostly consistent across different architectures such that it requires only limited changes when ported to new hardware; the inter-nodal level of code—the level of code that communicates between nodes—is nearly invariant from machine to machine. Consequently, the approach taken to porting HACC effectively reduces the problem to the node level, thereby permitting concentration of effort on optimizing critical code components with a full awareness of the actual hardware.

In bringing HACC to exascale, the developers have aimed to evaluate Aurora's early hardware and software development kit on a set of more than complex kernels primarily written in CUDA or otherwise under active development, minimize divergence between CUDA and SYCL versions of codebase, identify configurations and implementation optimizations specific to Intel GPUs, and identify more generally applicable implementation optimizations.

Versions of HACC being developed for exascale systems incorporate basic gas physics (hydrodynamics) to enable more detailed studies of structure formation on the scales of galaxy clusters and individual galaxies. These versions also include sub-grid models that integrate phenomena like star formation and supernova feedback, which means the addition of more performance-critical code sections; these sections will be offloaded to run on the GPUs. A GPU implementation of HACC with hydrodynamics was



Single GPU Performance FOM (Higher is Better)



TOP Time evolution of the Farpoint particle data captured on a single Message Passing Interface rank at a number of redshifts. Approximately 75 million particles are depicted. Image: HACC Team

BOTTOM A single-GPU figure of merit problem used 33 million particles per GPU, and Intel PVC results are shown for both small (128) and large (256) General-purpose Register File (GRF) modes.

previously developed for the Summit system using CUDA. All the GPU versions of the code are being rewritten to target Aurora.

IMPACT Modern cosmology provides a unique window to fundamental physics, and has led to remarkable discoveries culminating in a highly successful model for the dynamics of the Universe. Simulations and predictions enabled by the HACC code deployed at exascale will help deepen our understanding of structure of the universe and its underlying physics. generate New generations of cosmological instruments, such as the Vera Rubin Observatory, will depend on exascale systems in order to perform measurements of unprecedented precision.

## Early Application Performance: NWChemEx

PI Theresa Windus, Iowa State University and  
Ames Laboratory  
PROG Aurora Early Science Program and  
Exascale Computing Project

As the original NWChem code is some quarter-century old, in updating the NWChem code, the NWChemEx developers decided to rewrite the application from the ground up, with the ultimate goal of providing the framework for a next-generation molecular modeling package. The new package is capable of enabling chemistry research on a variety of leading-edge computing systems.


The NWChemEX developers aim to restructure core functionality concurrent with the production of sophisticated physics models intended to leverage the computing power promised by the exascale era. As one component of this strategy, the developers have adopted the Aurora-supported DPC++ programming model as one of its development platforms.

The development team gives equal weight and consideration to physics models, architecture, and software structure, in order to fully harness large-scale HPC systems. To this end, NWChemEx incorporates numerous modern software-engineering techniques for C++, while GPU compatibility and support have been planned since the project's initial stages, thereby orienting the code to the demands of exascale as matter of constitution.

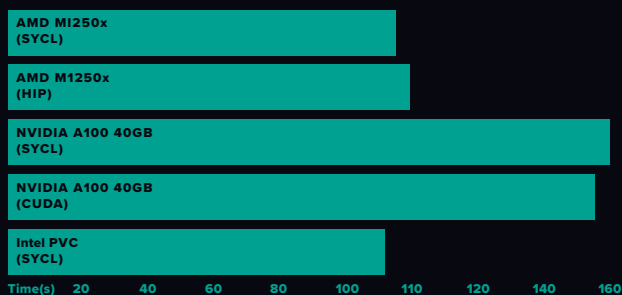
In order to overcome prior communication-related bottlenecks, the developers have localized communication to the greatest possible extent.

To help localize communication and thereby reduce related bottlenecks, NWChemEx is being geared such that CPUs handle communication protocols as well as any other non-intensive components. Anything else—anything “embarrassingly parallel” or computationally expensive—is to be processed by GPU.

For Intel hardware, the developers employ Intel's DPC++ Compatibility Tool to port any existing optimized CUDA code and translate it to DPC++. The Compatibility Tool is sophisticated enough that it reliably determines apposite



Performance on Single GPU (Lower is Better)



**TOP** Ubiquitin—a protein molecule typical of many biomolecular molecules, and with an abundance of experimental data available—was selected as a performance benchmark to evaluate the performance of NWChemEx on biomolecular systems at exascale. *Image: NWChemEx Team*

**BOTTOM** Performance on a single GPU was tested, measuring the time in seconds per iteration for the DLPNO-CCSD physics module, with strong results for the Sunspot accelerators.

syntax in translating abstractions from CUDA to SYCL, greatly reducing the developers' burden. Subsequent to translation, the developers finetune the DPC++ code to remove any redundancies, inelegancies, or performance issues introduced by automation.

**IMPACT** The NWChemEx project, when realized, has the potential to accelerate the development of next-generation batteries, drive the design of new functional materials, and advance the simulation of combusive chemical processes, in addition to addressing a wealth of other pressing challenges at the forefront of molecular modeling, including the development of stress-resistant biomass feedstock and the development of energy-efficient catalytic processes to convert biomass-derived materials into biofuels.

# Early Application Performance: OpenMC


PI Steven Hamilton, Oak Ridge National Laboratory;  
Paul Romano, Argonne National Laboratory  
PROG Exascale Computing Project

OpenMC, originally written for CPU-based high-performance computers and which is capable of using both distributed-memory (MPI) and shared-memory (OpenMP) parallelism, simulates the stochastic motion of neutral particles through a model that, as a representation of a real-world experimental setup, can range in complexity from a simple slab of radiation-shielding material to a full-scale nuclear reactor.


The GPU-oriented version of OpenMC has been completed and is already running on a number of GPU-based supercomputers, including Sunspot—the ALCF’s Aurora testbed and development system—and the ALCF’s NVIDIA-based Polaris. While the team’s goal is focused on honing performance on Aurora, the OpenMP offloading model has resulted in strong performance on every machine on which it was deployed, irrespective of vendor.

Current full-machine projections for OpenMC running on Aurora, based on preliminary simulation runs performed on Sunspot, are in the ballpark of 25 billion particle histories per second—indicating a speedup by some 2500x over what could be achieved at full-machine scale at the time of the ECP’s inception (the goal for which, again, had been a fiftyfold speedup).

**IMPACT** The ECP-supported ExaSMR project aims to use OpenMC to model the entire core of a nuclear reactor, generating virtual reactor simulation datasets with high-fidelity, coupled physics models for reactor phenomena that are truly predictive, filling in crucial gaps in experimental and operational reactor data.



Single GPU Performance FOM (Higher is Better)



TOP 2D slice of the advanced test reactor at Idaho National Laboratory.  
Image: OpenMC Team

BOTTOM Performance as tested on a single GPU using Sunspot is almost twice that achieved using other systems.

# Early Application Performance: QMCPACK

PI Paul Kent, Oak Ridge National Laboratory  
PROG Aurora Early Science Program and  
Exascale Computing Project

Quantum Monte Carlo (QMC) methods are ideal candidates for the next generation of material-design tools, which target not only simple bulk properties but collective effects in strongly correlated materials such as magnetic ordering, phase transitions, and quantum coherence. The open-source QMCPACK simulation code performs electronic structure calculations via accurate approximation of the Schrodinger equation.


With GPU acceleration supported by OpenMP offloading, the port of QMCPACK has proved resilient enough to performantly run on NVIDIA, AMD, and Intel GPUs.

To facilitate a variety of vendor-specific libraries, the developers redesigned the entire infrastructure of QMCPACK, with the intent of ensuring there exists a general framework that enables the flexible implementation of higher-level algorithms while allowing specialized implementations to run at a lower level as needed.

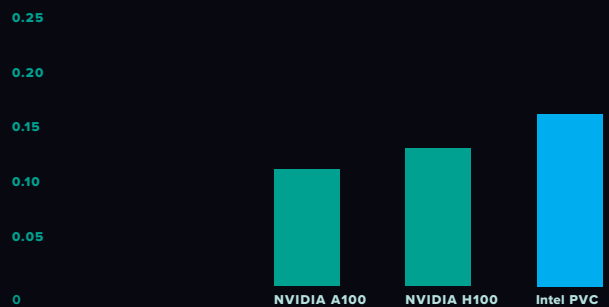
In attempting to attain optimal performance, the developers take a holistic approach that accounts for the entirety of their software stack and hardware. In addition to creating fast compute kernels, the developers focus on minimizing resource idle on host processors.

The team collaborates closely with Intel to modify certain application programming interfaces (APIs) in the Math Kernel Library (MKL). The APIs, as implemented in QMCPACK, are not designed for GPU performance. In addition to generating benchmarks and performing validation for MKL, the developers interact with Intel to design new MKL APIs.

**IMPACT** The greatly expanded computational power and parallelism of exascale will enable predictive capabilities far beyond the capacity of QMCPACK's current implementation, helping to accelerate the discovery of quantum materials that aid the development of sensors, devices, and new forms of low-power electronics.



Single GPU Performance FOM (Higher is Better)



**ТОР** Electron spin density of cobalt-doped silver nickel oxide delafossite where 33 percent of nickel sites (orange) are substituted with cobalt dopants (dark gray). The spin density is fully estimated using QMC. *Image: ALCF Visualization and Data Analytics Team; QMCPACK Team*

**БОТТОМ** Simulating a supercell of nickel oxide with 6144 electrons and 512 NiO atoms total, single-GPU performance on Sunspot was as much as 25 percent better than on the other machines tested.

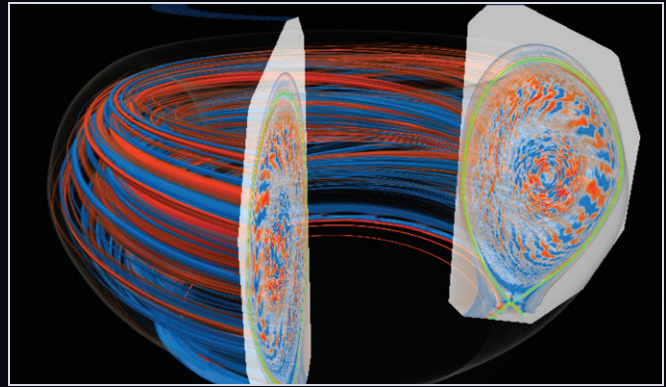
# Early Application Performance: XGC

PI Choongseok Chang,  
Princeton Plasma Physics Laboratory  
PROG Aurora Early Science Program and  
Exascale Computing Project

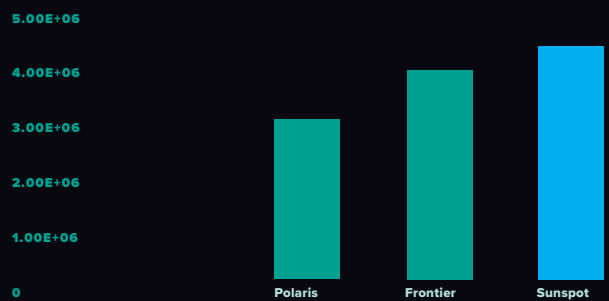
Developed in tandem with the ECP-supported Whole Device Model Application project—which aims to build a high-fidelity model of magnetically confined fusion plasmas to plan experiments with ITER—XGC is a gyrokinetic particle-in-cell code (with an unstructured 2D grid and structured toroidal grid) used to perform large-scale simulations on DOE supercomputers, and optimized for treating edge plasma.

Specializing in edge physics and realistic geometry, XGC is capable of solving boundary multiscale plasma problems across the magnetic separatrix (that is, the boundary between the magnetically confined and unconfined plasmas) and in contact with material wall called divertor, using first-principles-based kinetic equations.

To prepare for the next generation of high-performance computing, the code is being re-implemented for exascale using a performance-portable approach. Running at exascale will yield unique computational capabilities, some of which carry the potential for transformational impacts on fusion science: exascale expansion will make it possible to study, for instance, a larger and more realistic range of dimensionless plasma parameters than has ever been achieved, along with the energy-angle distribution of plasma particles impinging upon the material wall and the full spectrum of kinetic micro-instabilities that control the quality of energy confinement in a toroidal plasma. Further, exascale will enable physics modeling that incorporates multiple-charge tungsten ion species—impurities discharged from the tokamak vessel walls that impact edge-plasma behavior and fusion performance in the core-plasma through migration across the magnetic separatrix. Toward this end, XGC will support a wide array of additional features and modes, including delta-f and full-f, electrostatic and electromagnetic, axisymmetric, neutral particles with atomic cross-sections, atomic number transitions among different impurity states, and coupling Physics in constant state of development.



Single GPU Performance FOM (Higher is Better)



**TOP** Turbulent fluctuations in the electrostatic potential field in an XGC tokamak simulation. The green lines in the cross sections show the separatrix. Inside the separatrix, magnetic fields take the form of nested closed surfaces. *Image: XGC Team*

**BOTTOM** A gyrokinetic particle-in-cell simulation of tokamak plasma using C++ found that XGC demonstrates up to 46 percent greater performance on Sunspot than on other tested machines.

Optimization for exascale has required both GPU offloading and algorithmic flexibility. Performance on the Aurora testbed, Sunspot, has yielded scaling performance comparable to that of GPU systems such as Polaris.

**IMPACT** The resulting exascale application will be unique in its computational capabilities and will have potentially transformational impact in fusion science, for example, by studying a much larger and more realistic range of dimensionless plasma parameters than ever before, by providing the energy-angle distribution of plasmas hitting the material wall, and by assessing the rich spectrum of kinetic micro-instabilities that control the quality of energy confinement in a toroidal plasma.

The Aurora Early Science Program is designed to prepare key applications for the scale and architecture of the ALCF's upcoming exascale supercomputer, and field-test compilers and other software to pave the way for other production applications to run on the system.

### Simulation Projects

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

PI Anouar Benali, Argonne National Laboratory  
SW QMCPACK

#### Extreme-Scale Cosmological Hydrodynamics

PI Katrin Heitmann, Argonne National Laboratory  
SW CRK-HACC, Thrust

#### Extreme-Scale Unstructured Adaptive CFD

PI Kenneth Jansen, University of Colorado Boulder  
SW PHASTA

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

PI C.S. Chang, Princeton Plasma Physics Laboratory  
SW XGC, Kokkos, Cabana, ADIOS2

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

PI Theresa Windus, Iowa State University and Ames Laboratory  
SW NWChemEx

### Data Projects

#### Dark Sky Mining

PI Salman Habib, Argonne National Laboratory  
SW Custom analyses, containers, TensorFlow, hyperparameter optimization, HACC, CosmoTools, JAX, Numba/JIT

#### Data Analytics and Machine Learning for Exascale Computational Fluid Dynamics

PI Ken Jansen, University of Colorado Boulder  
SW SmartSim, RedisAI, PHASTA, VTK, Paraview

#### Exascale Computational Catalysis

PI David Bross, Argonne National Laboratory  
SW NRRAO, RMG, PostgreSQL, Fitpy, KinBot, Sella, Balsam, NWChemEx

#### Extreme-Scale In-Situ Visualization and Analysis of Fluid-Structure-Interaction Simulations

PI Amanda Randles, Duke University  
SW HARVEY, SENSEI, VTK, Kokkos

#### Simulating and Learning in the ATLAS Detector at the Exascale

PI Walter Hopkins, Argonne National Laboratory  
SW AthenaMT, Root, workflows, containers, TensorFlow, Kokkos

### Learning Projects

#### Accelerated Deep Learning Discovery in Fusion Energy Science

PI William Tang, Princeton Plasma Physics Laboratory  
SW FusionDL/FRNN, PyTorch, TensorFlow, DeepHyper

#### Enabling Connectomics at Exascale to Facilitate Discoveries in Neuroscience

PI Nicola Ferrier, Argonne National Laboratory  
SW TensorFlow, Horovod, flood-fill network, AlignTK, Tomosaic

#### Machine Learning for Lattice Quantum Chromodynamics

PI William Detmold, Massachusetts Institute of Technology  
SW MLHMC, USQCD libraries, PyTorch, TensorFlow, HDF5, MongoDB

#### Many-Body Perturbation Theory Meets Machine Learning to Discover Singlet Fission Materials

PI Noa Marom, Carnegie Mellon University  
SW Balsam, SISO, Bayesian optimization, BerkeleyGW, FHI-aims

#### Virtual Drug Response Prediction

PI Rick Stevens, Argonne National Laboratory  
SW CANDLE, Keras, TensorFlow

DOE's Exascale Computing Project (ECP) is a collaborative research, development, and deployment project focused on delivering a capable exascale computing ecosystem that encompasses applications, software, and hardware technologies. Here are some of the key ECP Application Development projects targeting Aurora.

---

### CANDLE: Exascale Deep Learning–Enabled Precision Medicine for Cancer

PI Rick Stevens, Argonne National Laboratory  
SW CANDLE

---

### ExaWind: Exascale Predictive Wind Plant Flow Physics Modeling

PI Michael Sprague, National Renewable Energy Laboratory  
SW Nalu-Wind, AMR-Wind, OpenFAST

---

### E3SM-MMF: Cloud-Resolving Climate Modeling of the Earth's Water Cycle

PI Mark Taylor, Sandia National Laboratories  
SW E3SM

---

### GAMESS: General Atomic and Molecular Electronic Structure System

PI Mark Gordon, Iowa State University  
SW GAMESS, libcchem

---

### EQSIM: High-Performance, Multidisciplinary Simulations for Regional-Scale Earthquake Hazard/ Risk Assessments

PI David McCallen, Lawrence Berkeley National Laboratory  
SW SW4, ESSI

---

### LatticeQCD: Validate Fundamental Laws of Nature

PI Andreas Kronfeld, Fermi National Accelerator Laboratory  
SW Grid, QUDA, Chroma, MILC

---

### EXAALT: Molecular Dynamics at Exascale

PI Danny Perez, Los Alamos National Laboratory  
SW LAMMPS, LATTE, ParSplice

---

### MFIX-Exa: Performance Prediction of Multiphase Energy Conversion Device

PI Jordan Musser, National Energy Technology Laboratory  
SW MFIX-Exa

---

### ExaFEL: Data Analytics at Exascale for Free Electron Lasers

PI Amedeo Perazzo, SLAC National Accelerator Laboratory  
SW CCTBX, spiniFEL

---

### NWChemEx: Tackling Chemical, Materials, and Biomolecular Challenges in Exascale

PI Theresa Windus, Iowa State University and Ames Laboratory  
SW NWChemEx

---

### ExaSky: Computing at the Extreme Scales

PI Salman Habib, Argonne National Laboratory  
SW HACC, CRK-HACC, Nyx

---

### QMCPACK: Quantum Mechanics at Exascale

PI Paul Kent, Oak Ridge National Laboratory  
SW QMCPACK

---

### ExaSMR: Coupled Monte Carlo Neutronics and Fluid Flow Simulation of Small Modular Reactors

PI Steven Hamilton, Oak Ridge National Laboratory  
SW NekRS, OpenMC

---

### WDMApp: High-Fidelity Whole Device Modeling of Magnetically Confined Fusion Plasmas

PI Amitava Bhattacharjee, Princeton Plasma Physics Laboratory  
SW XGC, GENE, GEM

---

### ExaStar: Exascale Models of Stellar Explosions

PI Daniel Kasen, Lawrence Berkeley National Laboratory  
SW Flash-X, Thornado

# ALCF AI Testbed Deploys New Systems

Available to researchers worldwide, the ALCF AI Testbed continues to add new platforms to advance the use of AI for science.

In 2023, the ALCF AI Testbed expanded its offerings to the research community, with the addition of new Graphcore and Groq systems as well as upgraded Cerebras and SambaNova machines.

The testbed is a growing collection of some of the world's most advanced AI accelerators available for open science. Designed to enable researchers to explore next-generation machine learning applications and workloads to advance AI for science, the systems are also helping the facility to gain a better understanding of how novel AI technologies can be integrated with traditional supercomputing systems powered by CPUs and GPUs.

The testbed's newest additions give the ALCF user community access to new leading-edge platforms for data-intensive research projects.

The new Graphcore Bow Pod64 is well-suited for both common and specialized machine learning applications, which will help to facilitate the use of new AI techniques and model types.

The new GroqRack system brings inference-based solutions that will aid in using trained machine learning models to make predictions or discover patterns in complex data.

The upgrade to a Cerebras Wafer-Scale Cluster WSE-2 optimizes the ALCF's existing Cerebras CS-2 system to include two CS-2 engines, enabling near-perfect linear scaling of large language models (LLMs). This capability helps make extreme-scale AI substantially more manageable.


The upgrade to a second-generation SambaNova DataScale SN30 system enables a wider range of AI-for-science applications, making massive AI models and datasets more tractable to users. In this system, each accelerator is allocated a terabyte of memory, which is ideal for applications involving LLMs as well as high-resolution imaging data from experimental facilities.

Together, the ALCF AI Testbed systems provide advanced capabilities that will support Argonne's efforts to develop an integrated research infrastructure that seamlessly connects advanced computing resources with data-intensive experiments, such as light sources and fusion experiments, to accelerate the pace of discovery.

Scientists are leveraging the ALCF AI Testbed systems for a wide range of data-driven research campaigns. The following summaries provide a glimpse of some of the efforts that are benefitting from the AI accelerators' advanced capabilities.







The Graphcore Bow Pod64 is the latest ALCF AI Testbed platform to be deployed to the research community for open science.

#### Experimental Data Analysis

Argonne researchers are leveraging multiple ALCF AI Testbed systems to accelerate and scale deep learning models to aid the analysis of X-ray data obtained at Argonne’s Advanced Photon Source (APS). The team is using the ALCF AI Testbed to train models—too large to run on a single GPU—to generate improved 3D images from x-ray data. They are also exploring the use of the ALCF’s AI platforms for fast-inference applications. Their work has yielded some promising initial results, with various models (PtychoNN, BraggNN, and AutoPhaseNN) showing speedups over traditional supercomputers. ALCF and vendor software teams are collaborating with the APS team to achieve further advances.

#### Neural Networks

Graph neural networks (GNNs) are powerful machine learning tools that can process and learn from data represented as graphs. GNNs are being used for research in several areas, including molecular design, financial data, and social networks. ALCF researchers are working to compare the performance of GNN models across multiple ALCF AI Testbed accelerators. With a focus on inference, the team is examining which GNN-specific operators or kernels, as a result of increasing numbers of parameters or batch sizes, can create computational bottlenecks that affect overall runtime.

#### COVID-19 Research

An Argonne-led team relied on the ALCF AI Testbed when using LLMs to discover SARS-CoV-2 variants. Their workflow leveraged AI accelerators alongside GPU-accelerated systems including the ALCF’s Polaris supercomputer. One of the critical problems the team had to overcome was how to manage extensive genomic sequences, the size of which can overwhelm many computing systems when establishing foundation models. The learning-optimized architecture of the ALCF AI Testbed systems was key for accelerating the training process. The team’s research resulted in the 2022 Gordon Bell Award Special Prize for COVID-19 Research.

#### Battery Materials

Argonne scientists are leveraging the ALCF AI Testbed to aid in the development of an application that combines two types of computations for research into potential battery materials: (1) running physics simulations of molecules under redox and (2) training a machine learning model that predicts that energy quantity. The application uses the machine learning model to predict the outcomes of the redox simulations, helping to identify molecules with the desired capacity for energy storage. The ALCF AI Testbed has enabled shortened latency when cycling between the execution of a new calculation that yields additional training data and when that model is used to select the next calculation.

# Creating an Integrated Research Infrastructure

Through its Nexus effort, Argonne continues to advance research aimed at combining supercomputers and AI with experiments to accelerate new discoveries.

To radically accelerate the pace of discoveries in the face of a growing deluge of scientific data, Argonne is playing a key role in supporting DOE's vision to create an integrated research infrastructure (IRI) that seamlessly melds DOE's cutting-edge experimental user facilities with its world-class supercomputing, AI, and data resources.

An IRI would give researchers new tools and pathways to analyze the massive amount of data being produced by light sources, telescopes, particle accelerators, sensors, and other scientific instruments. It would enable experiment-time data processing, allowing scientists to form hypotheses while experiments are being carried out and providing insights that can be used to design their next experiments. Such an ecosystem would also support the increasingly complex workflows and emerging AI for science capabilities needed for integrative research.


Argonne has long been at the forefront of developing tools and methods to connect its powerful computing resources with large, data-intensive experiments. Through its Nexus effort, the lab is continuing this work to help DOE create a broader research ecosystem that accelerates science via the close integration of AI and HPC resources with experimental facilities.

Over the course of the last several years, the ALCF has been involved in a number of successful collaborations that demonstrate the efficacy of combining its supercomputers

with experiments to analyze large datasets quickly. A major focus of Argonne's efforts has been centered on merging ALCF computing resources with Argonne's Advanced Photon Source (APS), but the work has also involved collaborations with facilities ranging from DIII-D National Fusion Facility in California to CERN's Large Hadron Collider (LHC) in Switzerland.

The projects have led to the creation of multiple tools for managing computational workflows and the development of new functionalities (e.g., preemptable queues for on-demand computing), giving the lab valuable experience to support the broader DOE IRI initiative. Globus and the ALCF Community Data Co-Op (ACDC) are also critical resources in enabling the IRI vision. Globus, a research automation platform created by researchers at Argonne and the University of Chicago, is used to manage high-speed data transfers, computing workflows, data collection, and other tasks for experiments. ACDC provides large-scale data storage capabilities, offering a portal that makes it easy to share data with external collaborators across the globe.

The following examples are just a few of the Argonne activities aimed at paving the way toward an integrated research ecosystem that connects advanced computing systems, data services, and large-scale experiments.





The ALCF's Polaris supercomputer is supporting research to advance the development of an integrated research infrastructure.




Aerial photo of Argonne's Advanced Photon Source.

#### Advanced Photon Source

At Argonne, the co-location of the ALCF and APS provides an ideal environment for testing and fine-tuning approaches to tightly couple the facilities. From developing workflow tools to enabling secure access to on-demand computing, the Argonne team has been building the infrastructure for integrated ALCF-APS research for over a decade.

The collaborative effort made major strides with the launch of the Polaris supercomputer in 2022. To enable rapid autonomous data analysis capabilities that scientists can use to drive their experiments, the integration team deployed preemptable queues on Polaris to facilitate on-demand job submissions. Using data collected during an APS experiment, the team was able to complete their first fully automated end-to-end test of the preemptable queues on Polaris with no humans in the loop. The process leverages Globus to manage the numerous high-speed data transfers, ALCF computations, and data cataloging and distribution steps involved in an experiment.

While work continues to enable this capability at more and more beamlines, the effort points to a future where the integration of the upgraded APS and the ALCF's Aurora exascale supercomputer will transform science at Argonne and beyond.




Graphical representation of the ATLAS detector at CERN's Large Hadron Collider. Image: ALCF Visualization and Data Analytics Team

#### Large Hadron Collider

Working with the LHC's ATLAS experiment, Argonne researchers explored how ALCF supercomputers could be leveraged to help meet the LHC's growing computing needs. As part of this multiyear collaboration, the team used the ALCF's Mira system to simulate particle collision experiments with a massively parallel supercomputer at a large scale for the first time, shedding light on a path forward for interpreting future LHC data.

The team's work involved using the ALCF-developed Balsam tool to enable ATLAS researchers to run jobs automatically through their workflow management system, simplifying the integration of ALCF resources as production compute endpoints for LHC experiments. The collaboration demonstrated how Balsam can help manage large-scale job campaigns, while diminishing the burden on the user and providing opportunities to optimize how jobs are submitted and executed.

As part of an Aurora Early Science Program project, Argonne researchers are now gearing up to use the lab's exascale supercomputer to help accelerate the search for new physics at the ATLAS experiment.



DIII-D, the largest magnetic fusion user facility in the U.S., is a tokamak confinement device that allows researchers to explore the optimization of the advanced tokamak approach to fusion energy production. *Image: General Atomics*

#### DIII-D National Fusion Facility

Argonne researchers also led a successful integration effort with the DIII-D National Fusion Facility. To study magnetic confinement in fusion energy, DIII-D scientists conduct fast-paced plasma physics experiments that involve creating 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.

To help the DIII-D team obtain these results on a between-pulse timescale, the Argonne team automated and shifted the analysis step to ALCF systems, 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. The team used the Balsam workflow tool to facilitate the flow of jobs between DIII-D and the ALCF.

This project was the first instance of an automatically triggered, between-shot fusion science analysis code running on-demand at a remotely located HPC resource.

#### Bringing It All Together

While Argonne and other national labs have been working on projects to demonstrate the promise of an integrated research paradigm for years, DOE's Advanced Scientific Computing Research (ASCR) program made it a formal initiative in 2020 with the launch of the IRI Task Force. The group, which included researchers from the ALCF and several national labs, identified the opportunities, risks, and challenges posed by such an integration.


In 2022, ASCR launched the IRI Blueprint Activity to create a framework for implementing the IRI. The blueprint team released a report that describes a path forward from the lab's individual partnerships and demonstrations to a broader long-term strategy that will work across the DOE ecosystem. Over the past year, the blueprint activities have started to formalize, with testbed resources and environments at each of the DOE computing facilities, where IRI ideas can be explored and evolved in collaboration with teams from DOE experimental facilities.

With the launch of Argonne's Nexus effort, the ALCF team will continue to leverage its knowledge, expertise, and resources to help DOE and the larger scientific community enable and scale this new paradigm across a diverse range of research areas, scientific instruments, and user facilities.

---

## SCIENCE

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





This visualization is the result of a state-of-the-art 3D simulation of supernova explosion and neutron-star birth carried out on ALCF supercomputers. It is a rare instance where the full stellar evolution of such an object, including the physics of the convection and the radiation, has been simulated in three dimensions. The image shows the deep core shrinking after explosion due to neutrino cooling and deleptonization on its way to becoming a cold, compact neutron star. *Image: ALCF Visualization and Data Analytics Team; Adam Burrows and Tianshu Wang, Princeton University*

# Accessing ALCF Resources for Science

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 through competitive, peer-reviewed allocation programs supported by the DOE and Argonne.

## Allocation Programs

### INCITE

The Innovative Novel Computational Impact on Theory and Experiment (INCITE) program aims to accelerate scientific discoveries and technological innovations by awarding ALCF computing time and resources to large-scale, computationally intensive projects that address grand challenges in science and engineering.

### ALCC

The ASCR Leadership Computing Challenge (ALCC) program allocates ALCF computing resources to projects that advance the DOE mission; help to broaden the community of researchers capable of using leadership computing resources; and serve the national interests for scientific discovery, technological innovation, and economic competitiveness.

### Director's Discretionary

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

### ESP

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




## INCITE/ALCC BY DOMAIN

### 2023 INCITE

**19.8M** NODE HOURS


|                              |     |
|------------------------------|-----|
| <b>A</b> Biological Sciences | 8 % |
| <b>B</b> Chemistry           | 10  |
| <b>C</b> Computer Science    | —   |
| <b>D</b> Earth Science       | 8   |
| <b>E</b> Energy Technologies | —   |
| <b>F</b> Engineering         | 14  |
| <b>G</b> Materials Science   | 15  |
| <b>H</b> Physics             | 45  |



### 2023 ALCC

**6.6M** NODE HOURS

|                              |     |
|------------------------------|-----|
| <b>A</b> Biological Sciences | 4 % |
| <b>B</b> Chemistry           | 24  |
| <b>C</b> Computer Science    | 1   |
| <b>D</b> Earth Science       | 7   |
| <b>E</b> Energy Technologies | 20  |
| <b>F</b> Engineering         | 5   |
| <b>G</b> Materials Science   | 13  |
| <b>H</b> Physics             | 26  |



ALCC data are from calendar year 2023.

# 2023 Science Highlights

The ALCF user community continued to push the boundaries of scientific computing, producing groundbreaking studies in areas ranging from climate modeling to materials design and discovery.

ALCF users employ large-scale data analyses, AI-driven methods, and simulations operating in tandem to conduct ambitious investigations and achieve scientific breakthroughs that would not otherwise be possible.

From detailed atomic-level simulations to massive cosmological studies, researchers use ALCF systems to probe extremely complex physical systems and processes that are too small or large, costly, or dangerous to study in a laboratory.

Year after year, ALCF users produce noteworthy results, whether they are developing and demonstrating novel computational methods or publishing papers in high-impact scientific journals.

In the following pages, we present a selection of science highlights from projects supported by ALCF allocation programs. This year's highlights range from simulating the complex aerothermodynamics of hypersonic flight to collaborating with ComED to help them assess the potential impacts of climate change on electricity infrastructure.

ALCF users also developed methods to further integrate supercomputing resources with experimental facilities; employed AI and machine learning techniques to supercharge materials discovery; carried out large-scale simulations to advance our understanding of turbulent flows, black holes, and other complex phenomena; and modeled building energy use to enhance sustainability and decarbonization efforts.

# Accelerating Advances in Biomedicine with HPC and AI

PI Ravi Madduri, Argonne National Laboratory  
 AWARD Director's Discretionary  
 SYSTEM Theta, Polaris, ALCF AI Testbed


**At Argonne National Laboratory, researchers are leveraging HPC and AI to spur innovations in biomedical research targeted towards prognosis, diagnosis, and treatment planning. Their efforts include overcoming data limitation challenges, enabling personalized predictions, and generating risk assessments.**

**CHALLENGE** In the realm of biomedicine, traditional research approaches are often hindered by lack of access to large-scale computing, access to real-world data, and a compendium of complex genetic interactions.

**APPROACH** Argonne scientists are contributing to three collaborative efforts that are using ALCF supercomputing and AI resources to advance biomedical research: (1) training a deep learning model to predict gene expression and the genetic underpinnings of disease; (2) generating synthetic cancer tumor images from genetic data; and (3) generating polygenic risk scores for different diseases.

**RESULTS** Given the pivotal role of transcription factors (TFs) in gene regulation and phenotype expression, predicting their binding accurately is crucial. In this project, the team developed a method called TFPred for individualized TF binding prediction based on genome-wide epigenetic features predicted by Enformer. This innovative approach provides a method for profiling TF binding patterns without the need for extensive ChIP-sequencing experiments.

To address the issue of data scarcity in biomedicine, the team developed the RNA-Cascaded-Diffusion-Model (RNA-CDM) to perform RNA-to-image synthesis in a multi-cancer context. This technique leverages a variational autoencoder to reduce the dimensionality of gene expression profiles and a cascaded diffusion model



Comparison of RNA-CDM generated synthetic whole slide images and real images from the Cancer Genome Atlas (TCGA) project database for Glioblastoma (GBM) and Lung Adenocarcinoma (LUAD). Image: Ravi Madduri and Tarak Nath Nandi, Argonne National Laboratory; Francisco Carrillo-Perez, Olivier Gevaert, Marija Pizurica, Jeanne Shen, and Yuanning Zheng, Stanford University

to synthesize realistic whole-slide image tiles from RNA-Seq data. The generated synthetic samples maintain the distribution of cell types observed in real-world data, and machine learning models pretrained on this synthetic data demonstrate improved performance. This approach shows promise for enhancing machine learning models in data-scarce settings, particularly in the context of cancer diagnosis.

In another effort, the team set out to improve the accuracy of polygenic risk scores (PRS) by exploring efficient strategies for collecting data for genome-wide association studies (GWAS). Using advanced Bayesian PRS methods and large-scale simulations, their research evaluates performance across different populations and trait scenarios. This work sheds light on the potential of more sophisticated PRS methodologies and the implications of multi-population GWAS data on risk prediction accuracy. As part of this project, the team is using advanced tools to simulate genotype population data from multiple ancestries for several million individuals. The simulated genotype data will allow them to test multiple theories about PRS strategies as well as test new developments in GWAS data tools.

**IMPACT** Together, these efforts highlight how supercomputers and advanced AI techniques are helping scientists develop innovative solutions to overcome data scarcity, unravel complex gene regulatory networks, and enhance the accuracy of risk prediction.

## PUBLICATIONS

Carrillo-Perez F., M. Pizurica, Y. Zheng, T. Nandi, R. Madduri, J. Shen, and O. Gevaert. "RNA-to-image multi-cancer synthesis using cascaded diffusion models," *Nature Biomedical Engineering* (preprint; under review with *Nature Biomedical Engineering*).  
<https://doi.org/10.1101/2023.01.13.523899>

Biological Sciences | 🧬🧑🏫 Simulation, Data, Learning

# Design of Peptides and Proteins on Classical and Quantum Computing Hardware

PI Vikram Mulligan, Flatiron Institute


AWARD INCITE

SYSTEM Theta

While existing peptide drugs rank among our best medicines, almost all of them were discovered in nature, as rational design of such drugs has only recently become possible. Led by Flatiron Institute researchers, this project aimed to advance the design of new peptide compounds by leveraging supercomputers for methods including physics-based simulations and machine learning.

**CHALLENGE** This research sought to reduce the computational and energetic costs of producing successful peptide macrocycle drugs or industrial enzymes. The researchers worked to develop low-cost machine learning methods that can approximate the output of computationally expensive validation simulations, ultimately enabling design and validation tasks on much more modest computing systems. Additionally, the researchers explored the use of quantum computing technologies as a means of solving the design problem at much lower energetic cost, which requires large-scale classical computing hardware both for carrying out quantum computing simulations during quantum algorithm development and for performing computational validation of designs produced on current-generation quantum computing hardware.

**APPROACH** Deploying the Rosetta software suite on ALCF supercomputing resources, the researchers used physics-based methods for drug design, including quantum chemistry calculations. Validation strategies involved large-scale conformational sampling of amino-acid sequences in designed peptides. Such approaches, in which countless molecular arrangements are explored, are highly parallelizable, enabling the researchers to carry out the conformation analyses across thousands of ALCF compute nodes. These brute-force calculations laid the foundation for data



An artificially designed symmetric peptide macrocycle. Unlike natural proteins, which are built from amino acids that all have the same handedness, artificial peptides may be synthesized from building blocks that are mirror images of natural amino acids. *Image: Vikram Mulligan, Flatiron Institute*

repositories that the team will use to train machine-learning models.

**RESULTS** In a paper published in *Cell*, the researchers detailed their design of 184 6-12 residue macrocycles and the subsequent experimental determination of 35 structures. The paper shows that membrane permeability can be systematically achieved.

**IMPACT** This work could help substantially reduce the computational and energetic costs associated with rational drug design while also increasing the availability of computational design technologies to the scientific community.

## PUBLICATIONS

Bhardwaj, G., J. O'Connor, S. Rettie, Y.-H. Huang, T. A. Ramelot, V. K. Mulligan, G. G. Alpkilic, J. Palmer, A. K. Bera, M. J. Bick, M. Di Piazza, X. Li, P. Hosseinzadeh, T. W. Craven, R. Tejero, A. Lauko, R. Choi, C. Glynn, L. Dong, R. Griffin, W. C. van Voorhis, J. Rodriguez, L. Stewart, G. T. Montelione, D. Craik, and D. Baker. "Accurate De Novo Design of Membrane-Traversing Macrocycles," *Cell* (August 2022), Cell Press. <https://doi.org/10.1016/j.cell.2022.07.019>


# Identifying Reaction Coordinates for Complex Protein Dynamics

PI Ao Ma, University of Illinois Chicago  
 AWARD Director's Discretionary  
 SYSTEM Theta

Proteins play an essential role in nearly all biological processes. By investigating proteins and their functions, scientists are providing insights to drive drug development, further our understanding of disease mechanisms, and advance many other areas of biomedical research. With help from ALCF supercomputers, a team from University of Illinois Chicago (UIC) has made an important breakthrough in understanding how proteins function.

**CHALLENGE** The primary goal of protein science is to understand how proteins function, which requires understanding the dynamics responsible for transitions between different functional structures of a protein. If the exact reaction coordinates (the small number of essential coordinates that control functional dynamics) were known, researchers could determine the transition rate for any protein configuration and thoroughly understand its mechanism. Despite intensive efforts, identifying the exact reaction coordinates in complex molecules remains a formidable challenge.

**APPROACH** The UIC team employed their generalized work functional (GWF) method to study the flap opening process of HIV-1 protease, a complex protein and major drug target for combatting the HIV virus. GWF is a fundamental mechanical quantity rooted in Newton's law. Using the transition path sampling method, the researchers leveraged the ALCF's Theta supercomputer to generate 2,000 reactive trajectories that start from structures of HIV-1 protease with flaps in the semi-open state and end at structures with flaps in the open state. This data served as the input to the GWF method, which was used to pinpoint the exact reaction coordinates and determine the molecular mechanism of the flap opening process.



Representative open structures from accelerated trajectories along the identified six reaction coordinates. The reference semi-open structure is colored white and atoms are colored green, red, and blue according to its weight in the reaction coordinates. Image: Ao Ma, University of Illinois Chicago

**RESULTS** As detailed in their paper in the *Proceedings of the National Academy of Sciences*, the team was able to identify the exact reaction coordinates for a major conformational change of a large functional protein for the first time. Their results show that the flap opening of HIV-1 protease has six reaction coordinates, providing the precise definition of collectivity and cooperativity in the functional dynamics of a protein. Success in determining the reaction coordinates enabled acceleration of this important process by  $10^3$  to  $10^4$  folds compared to regular molecular dynamics simulations. The team's work demonstrates that the GWF method could potentially be applied to other problems in protein research, such as folding, entropic barriers, and reaction rates.

**IMPACT** By successfully identifying the exact reaction coordinates for a complex protein for the first time, the team has made an important breakthrough toward understanding protein functional dynamics. Their work has far-reaching implications for both biomedical research and protein engineering, providing insights that are crucial for designing drugs, fighting drug resistance, and developing artificial enzymes that can complete desired functions.

**PUBLICATIONS**  
 Wu, S., H. Li, and A. Ma. "Exact Reaction Coordinates for Flap Opening in HIV-1 Protease," *Proceedings of the National Academy of Sciences* (December 2022), PNAS.  
<https://doi.org/10.1073/pnas.2214906119>

# Multiphysics Modeling of Blood Flow with Cell Suspensions in Patient-Specific Capillary Network

PI Jifu Tan, Northern Illinois University

AWARD Director's Discretionary


SYSTEM Cooley

**Cardiovascular disease, including heart attack and stroke, is the leading cause of death in the United States. In this project, simulations of blood flow with deformable red blood cells were performed for the first time in a patient-specific retina vascular network examining the impact of blockages on flow rate and cell transport dynamics.**

**CHALLENGE** Modeling capillary flow accurately is challenging due to the complex structure with various vessel branches and loops, and moving cell suspensions whose size is comparable to vessel diameters. Large three-dimensional (3D) vascular networks, such as this, are typically represented by simplified one-dimensional (1D) models at a much lower computational cost; however, these reduced order models may not accurately describe the flow dynamics.

**APPROACH** Flow dynamics in a patient-specific retina capillary network were simulated through coupling of a lattice Boltzmann method (LBM) based fluid solver with particle-based cell membrane models using the immersed boundary method (IBM). The geometry of the retina network was obtained from the National Institutes Health 3D print database. The red and white blood cells were modeled as thin membranes using a particle-based method implemented in LAMMPS. Collaborating with the ALCF Visualization and Data Analytics team, the team used Cooley to develop scientific visualizations of their blood flow simulations.

**RESULTS** From the 3D simulations, it was found that cells in blood act as moderators of flow. The flow of blood was redistributed from high flow rate regions near the inlet to the distant vessels with lower flow rates. Cell splitting behavior at bifurcations was found to be complex, which depends on many factors such as flow rates, pressure differences, or geometric parameters of the daughter



A snapshot of complex flow with cell suspensions in the retina network showing the transport of cells in different regions of the network. The blue/green coloring indicates the flow velocity within the network and blood cells are colored red. Image: ALCF Visualization and Data Analytics Team; Kacper Ostalowski and Jifu Tan, Northern Illinois University.

branches. From 1D simulations, the steady state flow rate through the network was obtained 1) without any blockages and 2) for blockages in various vessels to assess the severity (i.e., change in flow velocity) and impact in different parts of the network. Several potential improvements to the 3D model were noted as well as the need for efficient post-analysis and visualization tools to enable in-situ visualization and analysis considering the large volume of data generated.

**IMPACT** Inclusion of larger white blood cells was found to significantly increase the transit time of red blood cells through vessels. The simulation of flow under partial vessel blockage (e.g., stenosis) with cells showed that cells could oscillate and be trapped in an adjacent vessel due to the fluctuating flow. The best performing 1D reduced order model still resulted in large errors in both the number of red blood cells and flow rate for short vessels, and such models may be more suitable for networks with larger vessels.

## PUBLICATIONS

Ostalowski, K., and J. Tan. "Direct Simulation of Blood Flow with Heterogeneous Cell Suspensions in a Patient-Specific Capillary Network," *Physics of Fluids* (April 2022), AIP Publishing. <https://doi.org/10.1063/5.0088342>.

Thompson, A. "Capillary Model Tracks Blood Flow in Healthy and Diseased Networks," *Scilight* (April 2022), AIP Publishing. <https://doi.org/10.1063/10.0010412>

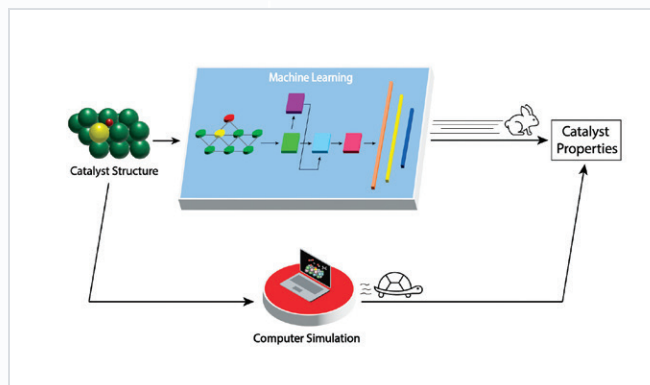
# Accelerated Catalyst Discovery from First Principles Simulations and Machine Learning

PI Rajeev Assary, Argonne National Laboratory  
 AWARD ALCC  
 SYSTEM Theta

**Biomass conversion technologies have the potential to reduce our dependence on fossil fuels and promote a sustainable energy future. New catalysts are needed to improve the conversion of biomass materials to transportation fuels and chemicals with other potential applications, such as biodegradable plastics. With help from ALCF supercomputers, scientists from Argonne National Laboratory developed a machine learning-driven model that can help accelerate the search for efficient and economically viable catalysts.**

**CHALLENGE** Catalyst design guidelines based on first principles simulations and reliable data-driven models are essential to the discovery of promising new catalyst materials. However, determining all properties that control catalytic activity and stability is time consuming and costly due to the complex interactions of the catalysts with chemicals and reaction intermediates. Therefore, further development of computational tools that can quickly and accurately predict the properties and efficacies of catalysts is necessary for accelerating catalyst discovery with minimal cost.

**APPROACH** With this ALCC project, the Argonne team set out to develop an accelerated computational approach that uses high-fidelity first principles simulations coupled with machine learning to provide improved guidelines to the catalyst discovery challenges. Using the ALCF's Theta supercomputer, the team employed density functional theory and machine learning methods to enable exhaustive searches for active catalyst facets and reveal active site motives for deoxygenation and carbon-carbon bond formation, which are essential chemical transformation processes for producing biofuels. High-throughput data generation via first principles simulations followed by deep learning provides a powerful approach to screen



The newly developed machine learning model greatly speeds up assessing the properties of molybdenum carbide catalysts for biomass conversion to useful products (top path) compared with current computer simulation methods (bottom path). Image: Argonne National Laboratory

for optimal candidates from a large search space of catalytic materials.

**RESULTS** In a paper published in *Digital Discovery*, the team detailed how their graph neural network machine learning model can accelerate the evaluation of crucial descriptors for catalyst screening. Their study focused on molybdenum carbides ( $\text{Mo}_2\text{C}$ ), which are considered among the most active and economically viable catalysts for biomass conversion. The team created 20,000 doped  $\text{Mo}_2\text{C}$  catalyst structures and used Theta to compute oxygen binding energies. Their simulations accounted for several dozen dopant elements and over a hundred possible positions for each dopant on catalyst surface. They used this database to train a graph neural network model for fast and accurate predictions of oxygen binding energies as the descriptors for catalytic stability. The team provided their results to the Chemical Catalysis for Bioenergy Consortium, where they are conducting experiments to further evaluate a small set of candidate catalysts.

**IMPACT** The team's model provides an accurate and broadly applicable machine learning approach for accelerating descriptor-based catalyst discovery. Instead of being limited to evaluating a few thousand catalyst structures over months with conventional computational methods, the new deep learning model can be used for accurate and inexpensive calculations for tens of thousands of structures in milliseconds.

## PUBLICATIONS

Doan, H. A., C. Li, L. Ward, M. Zhou, L. A. Curtiss, and R. S. Assary. "Accelerating the Evaluation of Crucial Descriptors for Catalyst Screening via Message Passing Neural Network." *Digital Discovery* (February 2023), Royal Society of Chemistry. <https://doi.org/10.1039/d2dd00088a>

# Microscopic Insight into Transport Properties of Li-Battery Electrolytes


PI Wei Jiang, Argonne National Laboratory  
 AWARD ALCC  
 SYSTEM Theta

The chemistries of non-lithium-ion batteries offer twice or more energy stored in a given volume or weight compared to lithium ion, but researchers need to develop improved electrolytes to pave the way for such high performing and long-lasting batteries. With help from ALCF's Theta supercomputer, Argonne scientists have developed a new fluoride-containing battery electrolyte that offers many advantages for powering electric vehicles.

**CHALLENGE** Fluorides have been identified as a key ingredient in interphases supporting aggressive battery chemistries. While the precursor for these fluorides must be pre-stored in electrolyte components and only delivered at extreme potentials, the chemical source of fluorine so far has been confined to either anions or fluorinated molecules, whose presence in the inner-Helmholtz layer of electrodes, and consequently their contribution to the interphasial chemistry, is restricted. The main problem is that their high energy density declines rapidly with repeated charge and discharge.

**APPROACH** The team's solution involved changing the electrolyte, a liquid through which lithium ions move between cathode and anode to implement charge and discharge. In lithium metal batteries, the electrolyte is a liquid consisting of a lithium-containing salt dissolved in a solvent. Simulations on the ALCF's Theta supercomputer revealed that the fluorine cations stick to and accumulate on the anode and cathode surfaces before any charge-discharge cycling. Then, during the early stages of cycling, a resilient solid electrolyte interphase (SEI) layer forms that is superior to what is possible with previous electrolytes.

**RESULTS** The team was able to tune the proportion of fluoride solvent to lithium salt to create a layer with optimal



Design of lithium metal battery with electrolyte containing a fluorinated cation (atomic structure at center). The "interface" area represents the layer with fluorine that forms on the anode surface, as well as the cathode surface.  
 Image: Argonne National Laboratory

properties, including an SEI thickness that is not too thick or thin. Because of this layer, lithium ions could efficiently flow in and out of the electrodes during charge and discharge for hundreds of cycles. High-resolution electron microscopy at Argonne and Pacific Northwest National Laboratory revealed that the highly protective SEI layer on the anode and cathode led to the stable cycling.

**IMPACT** The team's new electrolyte offers many advantages. It is low cost because it can be made with extremely high purity and yield in one simple step rather than multiple steps. It is environmentally friendly because it uses much less solvent, which is volatile and can release contaminants into the environment. And it is safer because it is not flammable. The chemistries of non-lithium-ion batteries offer twice or more energy stored in a given volume or weight compared to lithium ion. They could power cars for much longer distances and could even power long-haul trucks and aircraft one day. The expectation is that widespread use of such batteries will help address the problem of climate change.

## PUBLICATIONS

Liu, Q., W. Jiang, J. Xu, Y. Xu, Z. Yang, D.-J. Yoo, K. Z. Pupek, C. Wang, C. Liu, K. Xu, and Z. Zhang. "A Fluorinated Cation Introduces New Interphasial Chemistries to Enable High-Voltage Lithium Metal Batteries," *Nature Communications* (June 2023), Springer Science and Business Media LLC. <https://doi.org/10.1038/s41467-023-38229-7>



# Linking Scientific Instruments and Computation: Patterns, Technologies, and Experiences

PI Rafael Vescovi, Argonne National Laboratory  
AWARD Director's Discretionary  
SYSTEM Theta, ThetaGPU

As large-scale experiments generate an ever-increasing amount of scientific data, automating research tasks, such as data transfers and analysis, is critical for helping scientists to sort through the results to make new discoveries. However, new ways of implementing the automation of research processes are necessary to fully utilize the power of leading-edge computational and data technologies. A team of researchers from Argonne National Laboratory and Globus presented new techniques for automating inter-facility workflows by identifying how to express their component processes in abstract terms that make them reusable.

**CHALLENGE** Workflows that support experimental research activities largely feature the same general patterns and consist of the same general steps, referred to as paths or flows. Flows can include data collection, reduction, inversion, storage and publication, machine-learning model training, experiment steering, and coupled simulation. However, the majority of flows are assembled on an ad hoc basis, each built independently of the others to support a particular beamline and specific analysis, and few flows are able to make effective use of modern supercomputers such as those available at ALCF. The researchers sought to make building workflows easier and faster via reuse and automation.

**APPROACH** As part of their efforts to automate inter-facility workflows, the research team captured data at multiple beamlines at Argonne's Advanced Photon Source (APS), implementing flows for several applications. In each case the data were transferred from the APS to the ALCF to make use of computational resources. Analyses were then performed quasi-instantaneously, with visualizations loaded into a data portal to enable near-real-time experiment monitoring.



In a recent study, researchers from Argonne and Globus explored and demonstrated methods for implementing computational flows that link scientific instruments with computing, data repositories, and other resources. *Image: Patterns, vol. 3, no. 10, Oct. 2022, p. 100606.*

**RESULTS** To minimize the duplication of effort necessarily entailed by ad hoc workflow creation, the research team demonstrated, in an article published in *Patterns*, that Globus automation services (such as Globus Flows) can be used to create a language for expressing workflows that makes them reusable. Moreover, many beamline flows are automatable. The researchers demonstrated flow reusability by abstracting the paths and breaking them down into common components that consist of actions like publish, move data, or add metadata to an index. Results were incorporated automatically into a catalog for further exploration (enabling, for example, solution of a sample's protein structure). The diversity of the implementation activities demonstrated the general applicability of the described automation method, which was applied to x-ray photon correlation spectroscopy, serial synchrotron crystallography, ptychography, and high-energy diffraction microscopy.

**IMPACT** Workflow abstraction and reusability enable users at facilities such as the ALCF and APS to perform more effective research and accelerate discovery. The team's methods are helping to advance the development of an integrated research infrastructure that seamlessly connects data-intensive scientific instruments with advanced computing resources, enabling automated data processing and real-time feedback to experiments.

## PUBLICATIONS

Vescovi, R., R. Chard, N. D. Saint, B. Blaiszik, J. Pruyne, T. Bicer, A. Lavens, Z. Liu, M. E. Papka, S. Narayanan, N. Schwarz, K. Chard, and I. Foster. "Linking Scientific Instruments and Computation: Patterns, Technologies, and Experiences," *Patterns* (October 2022), Elsevier.  
<https://doi.org/10.1016/j.patter.2022.100606>


# Standardizing Complex Numbers in SYCL

PI Thomas Applencourt, Argonne National Laboratory  
 AWARD Director's Discretionary  
 SYSTEM Theta

Complex numbers are an integral part of running scientific simulations on HPC systems. The numbers, which represent the sum of real and imaginary numbers, allow researchers to create mathematical models that produce simulations involving quantum mechanics, fluid dynamics, and other phenomena. Researchers from Argonne National Laboratory led an effort to develop a method that streamlines the use of complex numbers in SYCL, a key programming framework for the ALCF's Aurora exascale supercomputer.

**CHALLENGE** While most HPC programming tools are equipped to handle complex numbers, SYCL was originally created to aid in the development of graphics, so it did not account for the use of complex numbers needed in scientific computing. This resulted in fragmented and inconsistent implementations in the SYCL ecosystem. With support from Intel, Codeplay, the Khronos Group, DOE's Exascale Computing Project and many partner organizations, SYCL has evolved considerably in recent years, emerging as an important framework for enabling portable programming across heterogeneous HPC architectures, including DOE's exascale systems.

**APPROACH** To address the issue of fragmentation, the Argonne-led team defined and implemented a "header only" pure SYCL library of the C++ standard complex and math function that can be used in any SYCL implementation. The library is based on the existing LLVM implementation, which is very robust and well tested. Despite this, the LLVM standard C++ library was created with standard math functions on CPU hardware, while the team's new implementation (SycCPLX) utilizes SYCL built-ins and is hardware agnostic. To ensure the SycCPLX built-ins can be executed as both host functions



Box plots of the kernel runtime achieved by three SYCL implementations of Milc-Dslash on a single NVIDIA A100 GPU as a function of work-group sizes. Circles inside the box represent the sample mean. Median values are displayed at the top of the graph. Image: Thomas Applencourt, Argonne National Laboratory

and SYCL device functions, the researchers developed a test-suite that verifies the approach works as intended across multiple devices with various SYCL implementations.

**RESULTS** The team collaborated with the developers of two applications—GENE (a plasma microturbulence code used for fusion energy research) and Milc-Dslash (a benchmark derived from Milc, a simulation code used for lattice quantum chromodynamics research)—to test their approach. Running the applications on supercomputers at ALCF and NERSC, they found that the new SYCL implementation was comparable to existing non-standard implementations with no significant impacts on application performance. The team's work was recognized as Most Outstanding Paper at the 2023 IWOC & SYCLcon conference. The open-source implementation is currently available on GitHub, but the goal is to have the library integrated into the next SYCL specification.

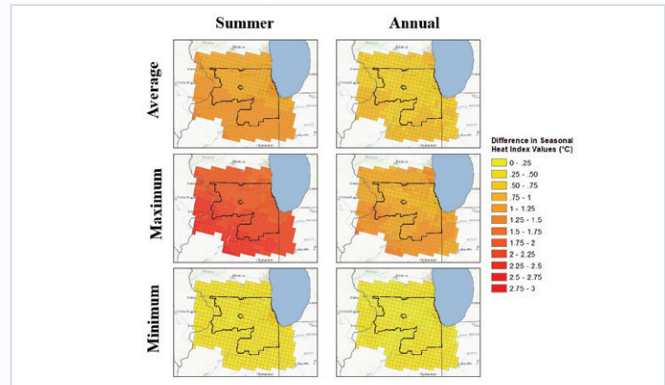
**IMPACT** The team's open-source library of complex numbers and associated math functions for SYCL provides researchers with an accurate and reliable solution that is portable across different HPC systems.

## PUBLICATIONS

Applencourt, T., B. Videau, J. Le Quellec, A. Dufek, K. Harms, N. Liber, B. Allen, and A. Belton-Schure. "Standardizing Complex Numbers in SYCL," *Proceedings of the 2023 International Workshop on OpenCL and SYCL* (April 2023), ACM. <https://doi.org/10.1145/3585341.3585343>

# High-Resolution Regional Climate Model Simulations

PI Rao Kotamarthi, Argonne National Laboratory  
 AWARD Director’s Discretionary  
 SYSTEM Theta



Change in the summer and annual average, maximum, and minimum heat index from the baseline to mid-century periods. Image: Argonne National Laboratory and ComEd

Argonne National Laboratory and Commonwealth Edison Company (ComEd) have joined forces to safeguard the power supply for millions of people in Northern Illinois. Backed by supercomputing, Argonne’s advanced modeling provides a deeper understanding of changing climate conditions, enabling ComEd to better assess its evolving risks, make informed decisions, and strengthen its electricity infrastructure.

**CHALLENGE** ComEd strives to deliver strong reliability to its customers, investing in its power grid to minimize power outages. However, severe weather caused by climate change poses major challenges to the resilience of the electric company’s infrastructure. Conditions such as rising temperatures are already having an impact. As severe weather becomes more prevalent, ComEd’s electrical system must deal with environmental conditions that can directly affect its performance, reliability, and durability.

**APPROACH** To navigate climate risks, ComEd needs a forward-looking approach to grid planning in order to safeguard power supply. The company has traditionally relied on historical weather data for its infrastructure planning—specifically, records of the hottest weather over the past 30 years. With each year bringing warmer temperatures, this data alone is no longer sufficient. In response, the Center for Climate Resilience and Decision Science at Argonne is working with ComEd to understand the potential impact of climate change on its distribution grid and develop adaptation strategies. Fundamentally, this project aims to integrate climate change forecasts into its infrastructure planning.

For the first phase of this collaboration, the researchers analyzed temperatures, heat index, and average wind speeds. Their findings indicate that by mid-century,

Northern Illinois will experience hotter weather, with average daytime temperatures surpassing 93–94 degrees Fahrenheit and rising nighttime temperatures. The research team used high-resolution downscaling approaches and computationally intensive simulations to enhance the accuracy of their predictions, providing information that can be used to provide targeted adaptations to grid infrastructure and operations.

**RESULTS** With this downscaling capability, the team achieved a resolution of 12 km by 12 km for their projections under the first phase of the collaboration. This enabled them to capture climate conditions at a finer scale compared to the typical resolution of about 100 km by 100 km. The added insight has helped ComEd more effectively plan for future climate changes to Northern Illinois.

**IMPACT** Using climate data from the study, ComEd is now able to improve its infrastructure planning to meet projected energy demands and mitigate the impacts of climate change, ensuring a reliable power supply to customers. With insights from the study, ComEd is evolving its infrastructure planning process to better account for the changing climate and potential future conditions.

**PUBLICATIONS**

Burg, R., R. Kartheiser, M. Mondello, M. Blaise, M. Altinakar, J. Branham, C. Burdi, M. Cuadra, B. Gamelin, V. R. Kotamarthi, K. B. Pfeiffer, V. Ramalingam, A. Raffety, A. Ruiz, T. A. Wall, and E. Yan. “ComEd Climate Risk and Adaptation Outlook, Phase 1: Temperature, Heat Index, and Average Wind,” *ANL Technical Report ANL/DIS-22/4* (November 21, 2022), Argonne National Laboratory. <https://doi.org/10.2172/1900595>

Earth Science | 🏔️👤 Simulation

# Multirate Methods for Coupled Compressible Navier-Stokes Systems


PI Shinhoo Kang, Argonne National Laboratory  
 AWARD Director's Discretionary  
 SYSTEM Theta

**Earth system models are complex integrated models of atmosphere, ocean, sea ice, and land surface. Coupling the components can be very challenging because of the difference in physics, temporal, and spatial scales. By exploring advanced algorithms for the fluid-fluid interaction problem and showcasing their scalable parallel performance, Argonne researchers opened the door to more complex climate research.**

**CHALLENGE** Developing multirate methods for large-scale coupled climate models is challenging, not only because of the requirement in conservation and convergence but also because of the demand for scalability in a parallel framework. These coupled models consist of multiple domains and different sets of partial differential equations, and hence the interface treatment poses great numerical and computational difficulties. Cross-domain communication, transition between the coupling components, and load balancing need to be carefully considered.

**APPROACH** Using ALCF supercomputers, the research team applied multirate partitioned Runge–Kutta (MPRK) coupling methods for fluid-fluid interaction problems, and demonstrated its strong and weak scalability on massively parallel computers by using the PETSc library. A buffer region at the interface was used for a coherent transition, which was crucial for achieving conservation and convergence. Domain decomposition was applied to each coupling problem as if the buffer region did not exist. To evaluate computational performance systematically, the researchers developed a theoretical performance model for both serial and parallel cases.

**RESULTS** In a paper published in *Computers and Fluids*, the researchers considered the compressible Navier–Stokes equations with gravity coupled through a rigid-lid interface.



Evolution of temperature field for thermal convection. The cold fluid parcel in the atmosphere drops down to the interface while the warm fluid in the ocean rises. Image: Shinhoo Kang, Argonne National Laboratory

The team's large-scale numerical experiments, carried out on ALCF supercomputers, reveal that multirate partitioned Runge–Kutta coupling schemes can conserve total mass, have second-order accuracy in time, and provide favorable strong- and weak-scaling performance on modern computing architectures.

**IMPACT** The developed multirate coupling scheme supports conservative, stable, and scalable simulations for coupling problems, which can efficiently handle different timescales arising from coupled Earth-system components—thereby enabling more comprehensive climatic research, as well as solutions to partial differential equations in general.

## PUBLICATIONS

Kang, S., A. Dener, A. Hamilton, H. Zhang, E. M. Constantinescu, and R. L. Jacob. "Multirate Partitioned Runge–Kutta Methods for Coupled Navier–Stokes Equations," *Computers and Fluids* (June 2023), Elsevier.  
<https://doi.org/10.1016/j.compfluid.2023.105964>

# Projected U.S. Drought Extremes through the Twenty-First Century with Vapor Pressure Deficit

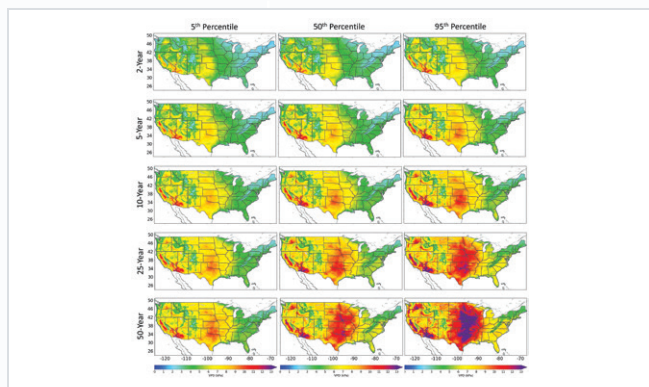
PI Rao Kotamarthi, Argonne National Laboratory  
 AWARD Director's Discretionary, ALCC  
 SYSTEM Theta

Although potential increases in hurricane and severe weather activity have attracted much of the focus of climate researchers looking to model and assess the impact of global warming, drought is another projected climate impact that is expected to have significant economic consequences. Researchers modeling the changing climate over the remainder of this century have developed a new method to assess the likelihood of extreme drought conditions over several different regions of the United States.

**CHALLENGE** Vapor pressure deficit (VPD) is a valuable measure of evaporative demand as moisture moves from the surface into the atmosphere under warming conditions. Considering that saturation vapor pressure is driven by temperature, VPD will likely be a key metric for projecting future extremes and thus the focus of this work. In this project, researchers investigated the utility of VPD in detecting short-term droughts (also called flash droughts) by calculating a drought index with VPD, and by assessing future VPD extremes by applying extreme value theory models to VPD.

**APPROACH** Leveraging ALCC supercomputing resources to run regional climate models at a high spatial resolution, the researchers developed a VPD-based drought index, which is calculated based on a combination of temperature and relative humidity, and consists of the difference between how much water vapor the air can hold when saturated and the total amount of water vapor available.

**RESULTS** Results show the VPD-based index performs well in identifying the timing and magnitude of short-term droughts, and that extreme VPD is increasing across the United States. By the end of the twenty-first century,



The Vapor Pressure Deficit (VPD) 2, 5, 10, 25 and 50-year return periods in the United States for the 5th, median (50th), and 95th percentile of the sampled model ensemble. Image: Brandi L. Gamelin, Argonne National Laboratory

the number of days VPD is above 9 kPa, a threshold over which drought becomes likely, is projected to increase by 10 days along California's ww, 30 to 40 days in the northwest and Midwest, and 100 days in California's Central Valley.

**IMPACT** The projected worsening of droughts because of global warming is likely to have significant consequences in terms of crop loss, wildfires, and demand for water resources. This new predictive technique has the potential to help communities better understand the phenomena induced by high temperatures and high evaporation rates, such as flash droughts.

**PUBLICATIONS**  
 Gamelin, B. L., J. Feinstein, J. Wang, J. Bessac, E. Yan, and V. R. Kotamarthi. "Projected U.S. Drought Extremes through the Twenty-First Century with Vapor Pressure Deficit," *Scientific Reports* (May 2022), Springer Nature. <https://doi.org/10.1038/s41598-022-12516-7>

# Assessment of Turbulent Prandtl Number for Heavy Liquid Metal Flow in a Bare Rod Bundle

PI Yiqi Yu, Argonne National Laboratory  
Emily Shemon, Argonne National Laboratory

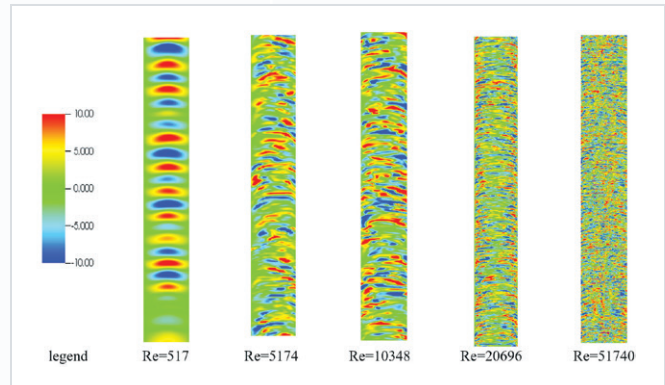
AWARD ALCC  
SYSTEM Theta

Westinghouse Electric Company is working with an international team to develop its next-generation high-capacity nuclear power plant based on lead-cooled fast reactor technology. Using ALCF supercomputers, researchers from Argonne National Laboratory are collaborating with the company to provide insights into the reactor's flow physics and heat transfer mechanisms.

**CHALLENGE** Lead-cooled fast reactors are a type of nuclear reactor design that offer many advantages, including the ability to operate at higher thermal efficiencies than existing commercial light water reactors. Developing these advanced reactors poses challenges due to the unique characteristics of heavy liquid metal (HLM) coolants, such as a low Prandtl number ( $Pr_t$ ) compared to water. Existing turbulence models are inadequate for accurately predicting heat transfer in HLM flows, making the selection of an appropriate turbulent Prandtl number ( $Pr_t$ ) critical. Accurate modeling and simulation of heat transfer and mixing in the HLM coolant is needed to help prepare the technology for licensing.

**APPROACH** For this effort, the team performed large eddy simulations (LES) using the open-source Nek5000 code on ALCF's Theta system to study nuclear fuel rod bundles with HLM flows. LES do not require a  $Pr_t$  to model turbulence-driven heat transfer, and thus can be used as benchmarks for selecting a  $Pr_t$  in a less computationally expensive Reynolds Averaged Navier–Stokes (RANS) model which requires this parameter.

**RESULTS** In a paper published in *Nuclear Engineering and Design*, the researchers showed that the selection of the appropriate  $Pr_t$  significantly impacts the accuracy of simulations for advanced nuclear reactors. By analyzing a prototypical lead-cooled fast reactor assembly with



LES predictions of the turbulent structure (flow oscillations) in a lead-cooled rod bundle with varying Reynolds number (Re). As Re increases, the flow oscillations become smaller scale and more chaotic. Image: Yiqi Yu, Argonne National Laboratory

different  $Pr_t$  values, the team found that inappropriate  $Pr_t$  can introduce error in Nusselt number (a measure of heat transfer) by up to 44 percent. They also compared detailed temperature distributions obtained by computationally expensive LES and less expensive RANS simulations to better understand the deviation introduced by the turbulence model. The analysis shows that the RANS model with  $Pr_t=1.5$  shows the best agreement with LES on the prediction of local temperature distribution and global Nusselt number.

**IMPACT** The team's research is helping to enhance the understanding and modeling of heavy liquid metal flow behavior and heat transfer mechanisms for next-generation nuclear reactors. In addition, their study provides valuable high-fidelity reference data that can be used by the nuclear reactor research community to validate and calibrate less computationally expensive models.

## PUBLICATIONS

Yu, Y., E. Shemon, and E. Merzari. "LES Simulation on Heavy Liquid Metal Flow in a Bare Rod Bundle for Assessment of Turbulent Prandtl Number," *Nuclear Engineering and Design* (April 2023), Elsevier.  
<https://doi.org/10.1016/j.nucengdes.2023.112175>

# Automatic Building Energy Modeling and Analysis


PI Joshua New, Oak Ridge National Laboratory  
AWARD INCITE  
SYSTEM Theta

With U.S. buildings accounting for 73% of the nation's electricity consumption and 39% of greenhouse gas emissions, the building sector presents a significant opportunity to enhance sustainability and decarbonization efforts. To help identify opportunities for reducing energy, peak electric demand, emissions, and cost, researchers from Oak Ridge National Laboratory are using ALCF computing resources to model the energy use of every building in America.

**CHALLENGE** Climate change and shifting weather patterns continue to affect energy consumption and grid resilience. Understanding the relationship between the changing climate and building energy use is crucial to the planning and safeguarding of critical infrastructure. In addition, as urban building energy modeling (UBEM) capabilities continue to develop, it is important to bridge the gap between model predictions and measured data by evaluating the accuracy, bias, and limitations of such models.

**APPROACH** With this project, the team extends the capabilities of their Automatic Building Energy Modeling (AutoBEM) software suite. The researchers have used the ALCF's Theta supercomputer to create and simulate models of 141.5 million buildings in the U.S. and Puerto Rico. The models include unique 3D geometry, building type classification, and internal characteristics based on building codes during the year of construction. The team's software and data are being applied to several building energy modeling efforts, including assessing climate change impacts and sources of bias.

**RESULTS** In a recent study, the team compared buildings' modeled annual energy use against actual utility data for three cities, which revealed an average error rate of 2.3%.



AutoBEM simulations estimate energy, costs, and emissions changes for each building out to the year 2100. Image: Oak Ridge National Laboratory

In a related study, the researchers used higher-resolution utility data from over 50,000 buildings to identify the distribution of error at higher resolutions up to building-specific estimations of energy use every 15 minutes. The team's research indicates a need for additional building data and model improvement for individual building analysis.

The team also conducted a study to understand how future climate scenarios will impact electricity and natural gas use of commercial buildings in the U.S. Using scenarios defined in the IPCC Assessment Report 5, the team analyzed building energy use for 2030, 2045, and 2100. Their research helps identify potential energy use changes due to climate scenarios and regions where energy use patterns could shift, providing insights to help inform infrastructure planning and energy efficiency strategies.

**IMPACT** The AutoBEM project aims to stimulate private sector activity towards more simulation-informed and grid-aware energy efficiency technologies, helping to pave the way to a more sustainable built environment. The team's efforts to provide free, publicly available U.S. building energy models ([bit.ly/ModelAmerica](https://bit.ly/ModelAmerica)) and analyses have resulted in custom data for use by over 50 organizations.

## PUBLICATIONS

Bass, B., J. New, A. Berres, M. Adams, N. Clinton, M. Leung, K. Tuxen-Bettman, and S. Van Groenou "Poster: Environmental Insights Explorer for Buildings." *Geo for Good Summit* (September 2022), Google.

Bass, B., J. New, N. Clinton, M. Adams, B. Copeland, and C. Amoo. "How Close Are Urban Scale Building Simulations to Measured Data? Examining Bias Derived from Building Metadata in Urban Building Energy Modeling." *Applied Energy* (December 2022), Elsevier. <https://doi.org/10.1016/j.apenergy.2022.120049>.

Bass, B., and J. New. "How Will United States Commercial Building Energy Use Be Impacted by IPCC Climate Scenarios?," *Energy* (January 2023), Elsevier. <https://doi.org/10.1016/j.energy.2022.125945>


# High-Fidelity Simulations of Gas Turbine Combustors for Sustainable Aviation Applications

PI Muhsin Ameen, Argonne National Laboratory  
Sicong Wu, Argonne National Laboratory  
AWARD Director's Discretionary  
SYSTEM Theta

**With a push towards decarbonizing the aviation sector, sustainable aviation fuels (SAFs) have gained prominence as a potential replacement for fossil fuels. This project is developing the capabilities to perform fully-resolved simulations of modern gas turbine combustors to enable improved understanding of the multiphysics processes in the context of advancing the development of SAFs.**

**CHALLENGE** To assess the viability of various SAFs, researchers must be able to understand and predict the complex flow, spray, and combustion processes taking place in the gas turbine combustors, and their influence on events such as lean blow out, high altitude relight and cold start, that affect the performance of gas turbines. With recent advances in numerical methods and the availability of HPC resources, computer simulations can provide unprecedented details of the underlying multiphysics processes, but they rely on the complex task of creating a detailed computational model of the gas turbine that is accurate and runs efficiently on modern computers.

**APPROACH** The objective of this research is to develop capabilities to perform fully-resolved simulations of modern gas turbine combustors using Nek5000 to enable improved understanding of the multiphysics processes in the context of advancing the development of sustainable aviation fuels. Nek5000 is a high-order spectral element method (SEM) based code, developed at Argonne, targeted towards exascale systems. Proper orthogonal decomposition of the turbulent flow field were performed to investigate the dynamics of the large- and small-scale turbulence in the combustor. Finally, simulations with fuel injection were used to determine the effect of fuel spray on the turbulent flow structures.



Instantaneous velocity magnitude in ARC-M1 combustor. Image: Sicong Wu, Argonne National Laboratory

**RESULTS** In this project, the team performed the first-ever wall-resolved large eddy simulations of the turbulent flow and spray processes in the Army Research Laboratory's ARC-M1 research combustor. The simulations were validated using particle image velocimetry measurements from a group at the University of Illinois at Urbana Champaign, and showed good agreement. The simulations demonstrated the presence of large and small recirculation regions generated due to mixing between the different flow streams. The accurate prediction of these recirculation regions is key in predicting the flame anchoring and dynamics for reacting simulations.

**IMPACT** These high-fidelity simulations that leverage the DOE supercomputers can help researchers understand the combustion and heat transfer challenges introduced by using low-carbon sustainable aviation fuels. This project will help establish a high-fidelity, scalable, numerical framework that can be used for evaluating the effect of fuel properties on flow and flame dynamics in a practical gas turbine combustor.

## PUBLICATIONS

Wu, S., D. Dasgupta, M. Ameen, and S. Patel. "LES Modeling of Gas Turbine Combustor Using Nek5000," *Bulletin of the American Physical Society* (November 2022), APS.

Wu, S., D. Dasgupta, M. Ameen, and S. Patel. "High-Fidelity Simulations of Gas Turbine Combustor Using Spectral Element Method," *AIAA SCITECH 2023 Forum* (January 2023), AIAA.  
<https://doi.org/10.2514/6.2023-1641>



# DNS of Wall-Bounded Magnetohydrodynamic Turbulence at High Reynolds Number


PI Myoungkyu Lee, University of Houston  
AWARD INCITE  
SYSTEM Theta

Understanding the fundamental physics of wall-bounded magnetohydrodynamic (MHD) turbulence is the key to developing multiple engineering applications, such as devices with liquid metal blankets in nuclear fusion reactors. The study of such applications is challenging because operating conditions are often at a high Reynolds number ( $Re$ ), necessitating either very high-fidelity experimental measurement systems or leadership-scale computing systems. This project performs direct numerical simulations (DNS) of wall-bounded MHD turbulence at high  $Re$  to reveal the spectral behaviors of turbulent kinetic energy (TKE) as functions of fluid speed, the strength and direction of the magnetic field, and the wall-normal distances.

**CHALLENGE** Since DNS requires resolving the entire spectrum of length- and time-scales of turbulent flows, it is impractical to use for complicated scenarios. Instead, the researchers will study incompressible canonical channel flow at different flow speeds and magnetic field strengths and directions in both transient and statistically stationary states.

**APPROACH** Flow with multiple parameters of fluid viscosity and electrical conductivity will be simulated on Theta using PoongBack, an incompressible turbulent flow solver with parallel fast Fourier transforms. The  $Re$  and corresponding grid sizes are selected from previous simulation results with hydrodynamic wall-bounded turbulence.

**RESULTS** The researchers identified the effect of magnetic fields in flows of electrically conductive fluid at high  $Re$ . Generally, the influence of increasing magnetic field strength  $Ha$  originates from the central region of a channel. The mean velocity gradient in the central area becomes notably flattened. In cases where  $Ha$  is substantially higher



Velocity field with (bottom) and without (top) influence of magnetic field at  $Re = 20,000$ . Image: Myoungkyu Lee, University of Houston

than  $Re$ , the mean velocity tends towards uniformity almost everywhere except for a thin near-wall zone where it increases linearly with wall-normal distance. Due to the broadened area of the flattened velocity profile, a rapid velocity change is needed in the near-wall region, leading to a higher mean velocity gradient at the wall in the presence of magnetic fields. Moderate strength magnetic fields suppress turbulence magnitudes originating from outer flows. The influence is particularly significant on wall-normal and spanwise velocity components. A magnetic field enhances anisotropy in the near-wall region and moreover serves to suppress large-scale motion.

**IMPACT** This work will detail the lifecycle of TKE from its production and dissipation at different length scales while displaying the influence of large-scale motion on the near-wall flows in MHD turbulence. Additionally, the resulting data will be useable for developing and improving reduced-order models, such as Reynolds-averaged Navier–Stokes models, subgrid stress models for large eddy simulation, and machine learning surrogates. Improved reduced-order models will shorten the turnaround time to design MHD applications with computational fluid dynamics (CFD) and make CFD solutions more accurate.

Engineering |  Simulation

# First-Principles Simulation of Hypersonic Flight


PI Maninder Grover, University of Dayton Research Institute

AWARD INCITE  
SYSTEM Theta

**Hypersonic flight, the ability to fly at more than five times the speed of sound, has the potential to revolutionize technologies for national security, aviation, and space exploration. However, a fundamental understanding of the aerothermodynamics of hypersonic flight is needed to enable technological advances in this field. A research team from the University of Dayton Research Institute and Air Force Research Laboratory is using ALCF supercomputers to shed light on the complex thermal environment that hypersonic vehicles encounter.**

**CHALLENGE** Strong shockwaves formed during hypersonic flight can cause the excitation of internal energy modes, and chemical reactivity in the shock heated gas. The rate processes for these phenomena competes with the local flow time, causing the flow to be in thermal and chemical nonequilibrium. Proper characterization of this state is important for designing the required thermal protection systems for hypersonic vehicles. A key challenge is to ensure that reduced-order models used in computational fluid dynamics codes can capture the strong coupling between fluid mechanics of the gas flow, gas-phase thermochemistry, and transport properties at high temperatures. Traditionally, these physics have been investigated separately by producing simplified models that tend to reproduce only certain aspects of high-speed, reacting flows.

**APPROACH** With this INCITE project, the team is running a custom version of the SPARTA Direct Simulation Monte Carlo (DSMC) code on ALCF computing resources to carry out direct molecular simulations (DMS) of hypersonic experiments. Their goal is to conduct simulations that rely solely on molecular-level interactions modeled using quantum mechanics, providing a fundamental comparison



Contour of gas velocity along an experimental test article. Image: ALCF Visualization and Data Analytics Team; University of Dayton Research Institute

with experiments, and well-characterized solutions that can be used as benchmarks for reduced-order models.

**RESULTS** In a new study published in the *Journal of Fluid Mechanics*, the team detailed a large-scale, fully resolved DSMC computation of a non-equilibrium, reactive flow of pure oxygen over a double cone (a canonical hypersonic test case). The researchers used their highly accurate DMS method to obtain first-principles data to inform the parameters of the thermochemical and transport collision models. Their computations show good agreement with heat flux and pressure measured on the test article during the experiment. The computation also provided molecular-level insights such as the nonequilibrium distribution of energy in the kinetic and vibrational modes in the shock layer. The team's results show the importance of particle methods in verifying physical assumptions made by reduced-order models.

**IMPACT** The team's research is helping to advance our understanding of the complex aerothermodynamics of hypersonic flight, providing insights that could help inform the design of safer and efficient technologies for space travel and defense.

## PUBLICATIONS

Valentini, P., M. S. Grover, A. M. Verhoff, and N. J. Bisek. "Near-Continuum, Hypersonic Oxygen Flow over a Double Cone Simulated by Direct Simulation Monte Carlo Informed from Quantum Chemistry," *Journal of Fluid Mechanics* (July 2023), Cambridge University Press.  
<https://doi.org/10.1017/jfm.2023.437>

Grover, M. S., A. M. Verhoff, P. Valentini, and N. J. Bisek. "First Principles Simulation of Reacting Hypersonic Flow over a Blunt Wedge," *Physics of Fluids* (August 2023), AIP Publishing.  
<https://doi.org/10.1063/5.0161570>


# Online Machine Learning for Large-Scale Turbulent Simulations

PI Kenneth Jansen, University of Colorado  
AWARD INCITE and Aurora Early Science Program  
SYSTEM Theta, Polaris

Led by University of Colorado Boulder researchers, this project seeks to advance the current state of the art for online data analytics and machine learning (ML) applied to large-scale computational fluid dynamics simulations, as well as to develop more predictive lower-fidelity and less computationally expensive turbulence models for flows relevant to the aerospace, automotive, and renewable energy industries.

**CHALLENGE** Using tools developed under previous allocations to extend beyond canonical turbulent flows neural net sub-grid stress (SGS) models for large eddy simulation (LES), the researchers perform two direct numerical simulations (DNS) coupled with online learning of wall-bounded flows with increasing complexity and scale so as to provide training data for SGS closures and to develop a neural net SGS model capable of accurately predicting flows of increasing complexity and that surpasses current state-of-the-art closures. LES on the same flows computed by DNS validate the accuracy of the newly trained SGS model. To evaluate the trained SGS model on a previously unseen flow, the team performs LES of the turbulent boundary layer over an airfoil with flow separation.

**APPROACH** To perform their computations, the researchers have adopted a mature finite-element based flow solver, PHASTA, coupled with SmartSim to manage the stream from the data producer (the flow solver) to the data consumer (the Python-based ML libraries performing the online training). Making efficient use of SmartSim's Orchestrator, the flow simulation and training are set up to be asynchronous, non-blocking tasks, thus allowing the training to be carried out concurrently while the solver generates new solution states.



Isosurface of instantaneous Q criterion colored by speed over a vertical tail at  $Re = 3.5 \times 10^5$ , with a rudder deflection angle of 30 degrees and a single unsteady jet active (fifth from the top). Image: Kenneth Jansen, University of Colorado Boulder

**RESULTS** As fluid dynamics simulations grow, generating new training datasets for traditional offline learning creates I/O and storage bottlenecks. Additionally, performing inference at runtime requires non-trivial coupling of ML framework libraries with simulation codes. In a new study, the researchers offer a solution to both of these limitations by simplifying this coupling and enabling in situ training and inference workflows on heterogeneous clusters. Leveraging SmartSim, the presented framework deploys a database to store data and ML models in memory, thus circumventing the file system. Using the ALCF's Polaris supercomputer, the team demonstrates perfect scaling efficiency to the full machine size of the data transfer and inference costs, due to a novel collocated deployment of the database, and trains an autoencoder in situ from a turbulent flow simulation.

**IMPACT** Through the integration of a mature flow solver already scaled to more than 3 million processes with distributed and online data analytics and training algorithms, this work will greatly enhance the confidence in lower-fidelity models and enable engineers to obtain more accurate solutions to complex flows outside the reach of today's modeling capabilities.

## PUBLICATIONS

Balin, R., F. Simini, C. Simpson, A. Shao, A. Rigazzi, M. Ellis, S. Becker, A. Doostan, J. A. Evans, and K. E. Jansen. "In Situ Framework for Coupling Simulation and Machine Learning with Application to CFD" <https://doi.org/10.48550/arXiv.2306.12900>

Engineering |  Simulation


# Turbulent Rayleigh-Bénard Convection in Suspensions of Bubbles

PI Parisa Mirbod, University of Illinois Chicago  
 AWARD Director's Discretionary  
 SYSTEM Theta, Polaris

**Turbulent Rayleigh-Bénard convection occurs when a fluid is heated from below and cooled from above. This process drives many phenomena including heat transport in stars, atmospheric flows, and oceanic currents. A research team from the University of Illinois Chicago (UIC), in collaboration with KTH Royal Institute of Technology in Sweden, is using ALCF supercomputers to advance our understanding of turbulent Rayleigh-Bénard convection in flows with bubbles in various natural and industrial settings.**

**CHALLENGE** Turbulent Rayleigh-Bénard convection is a highly complex phenomenon involving multiphase thermal flows across a wide range of scales, from small vortices to large eddies. Accurately modeling such flows is a computationally demanding task that necessitates the use of high-resolution numerical simulations and high-performance computing systems. To fill a knowledge gap in this area of research, the UIC-led team is using ALCF supercomputers to study the behavior of bubbles and their influence on heat transfer modulation in turbulent Rayleigh-Bénard convection flows. Their work seeks to gain a deeper understanding of bubbles dynamics and their impact on heat transfer.

**APPROACH** The team has worked with ALCF staff to optimize FlUTAS, an open-source code for multiphase fluid dynamics simulations, on Theta and Polaris. FlUTAS is designed to run efficiently on many-CPU and many-GPU systems in a single and unified framework. With partial support from the National Science Foundation, the researchers are using the code to perform direct numerical simulations that solve the Navier-Stokes equations, which govern fluid motion, on a grid that accurately captures the smallest flow scales.



Volume rendering of thermal plumes (red and blue) and liquid-liquid interface (grey) of a turbulent Rayleigh-Bénard convection. *Image: Luca Brandt, KTH Royal Institute of Technology and Norwegian University of Science and Technology; Nicolo Scapin, Princeton University; Hal Brynteson, Parisa Mirbod, and Abbas Moradi Bilondi, University of Illinois Chicago*

**RESULTS** As part of their research, the team explored the behavior of emulsions (liquid-liquid bubbles) within turbulent Rayleigh-Bénard, allowing for a detailed analysis of fluid dynamics at the smallest scales. They investigated the influence of key parameters, including the concentration of the dispersed phase and the viscosity ratio between the two fluids, on the dynamics of the emulsions and their impact on overall mass and heat transfer in the system. The team's study provides a comprehensive analysis of the effects of secondary phase deformability on heat transfer modulation, which has not been previously explored. Their findings offer valuable insights into the underlying mechanisms that govern the behavior of emulsions in turbulent Rayleigh-Bénard convection flows.

**IMPACT** A detailed understanding of turbulent Rayleigh-Bénard convection with bubbles is crucial to many areas of research, including weather predictions and climate calculations. It can also help drive advances in optimizing and designing industrial processes that involve multiphase turbulent thermal convection, such as cooling fuel rods in spent fuel pools of nuclear reactors.

## PUBLICATIONS

Moradi Bilondi, A., N. Scapin, L. Brandt, and P. Mirbod. "Multi-Phase Turbulent Rayleigh-Bénard Convection with Bubbles," *Bulletin of the American Physical Society* (November 2022), APS.

Moradi Bilondi, A., N. Scapin, L. Brandt, and P. Mirbod. "Heat Transfer Enhancement in Turbulent Rayleigh-Bénard Convection with Liquid-Liquid Emulsions," *Bulletin of the American Physical Society* (November 2023), APS.

Moradi Bilondi, A., N. Scapin, L. Brandt, and P. Mirbod. "Direct Numerical Simulations of Emulsions in Turbulent Rayleigh-Bénard Convection," (In preparation).

# AI-Guided Exascale Simulations of Quantum Materials Manufacturing and Control


PI Aiichiro Nakano, University of Southern California  
AWARD INCITE  
SYSTEM Polaris

**This project aims to boost scalable manufacturing of quantum materials and ultrafast control of their emergent properties on demand using AI-guided exascale quantum dynamics simulations. Neural-network quantum molecular dynamics (NNQMD) simulations based on machine learning are revolutionizing atomistic simulations of materials by providing quantum-mechanical accuracy at speeds orders of magnitude faster than is possible with traditional methods, but face challenges in scaling properly on massively parallel systems.**

**CHALLENGE** Despite its remarkable computational scalability, massively parallel NNQMD simulations face a major unsolved issue known as fidelity scaling. In such cases, small prediction errors can propagate and lead to unphysical atomic forces that degrade the accuracy of atomic trajectory over time. These force outliers can even cause the simulation to terminate unexpectedly. As simulations become spatially larger and temporally longer, the number of unphysical force predictions is expected to scale proportionally, which could severely limit NNQMD fidelity on new exascale supercomputing platforms, especially for the most exciting far-from-equilibrium applications.

**APPROACH** To solve the fidelity-scaling issue, the researchers implemented the Allegro–Legato model in its NNQMD code, RXMD-NN, which was deployed on the ALCF’s Polaris supercomputer. The model was trained using sharpness-aware minimization to regularize its sharpness along with its training loss and thereby enhance its robustness.

**RESULTS** As shown in an *ISC High Performance 2023* paper, the implemented Allegro–Legato model increases time-to-failure while maintaining the same inference speed and nearly equal accuracy. Specifically, time-to-failure in



Allegro–Legato has significantly increased the time to failure, thereby enabling long-time simulation including nuclear quantum effects, which was necessary to explain recent high-resolution inelastic neutron scattering experiments on ammonia. Image: Anikeya Aditya, Thomas Linker, and Ken-ichi Nomura, University of Southern California

Allegro–Legato is less dependent on problem size, thus allowing larger-scale and longer-duration NNQMD simulations without failure. Additionally, the researchers demonstrate that the fidelity-scalability of the NNQMD model correlates with sharpness of the model more than the number of parameters in the model.

**IMPACT** This work, directly validated by x-ray free electron laser, ultrafast electron diffraction, and neutron experiments at DOE facilities, will enable future production of high-quality custom quantum material architectures for broad and critical applications for continued U.S. leadership in technology development, including that for sustainable ammonia, thereby addressing DOE basic research needs for transformative manufacturing and quantum materials. The Allegro–Legato model exhibits excellent computational scalability and GPU acceleration in carrying out NNQMD simulations, with strong promise for emerging exascale systems.

## PUBLICATIONS

Ibayashi, H., T. M. Razakh, L. Yang, T. Linker, M. Olguin, S. Hattori, Y. Luo, R. K. Kalia, A. Nakano, K. Nomura, and P. Vashishta. “Allegro-Legato: Scalable, Fast, and Robust Neural-Network Quantum Molecular Dynamics via Sharpness-Aware Minimization,” *ISC High Performance 2023* (May 2023), Springer Nature. [https://doi.org/10.1007/978-3-031-32041-5\\_12](https://doi.org/10.1007/978-3-031-32041-5_12)


# End-to-End AI Framework for Small Molecules and Inorganic Crystals

PI Eliu Huerta, Argonne National Laboratory  
 AWARD Director's Discretionary  
 SYSTEM ThetaGPU

AI has emerged as a powerful tool for accelerating the design and discovery of novel functional materials. To help researchers speed up the process of predicting user-specified material properties, a team from Argonne National Laboratory has developed an end-to-end AI framework with transferable applications in the modeling of small molecules, inorganic crystals, and metal-organic frameworks.

**CHALLENGE** Deploying open-source, state-of-the-art AI models and benchmark datasets in modern computing environments is a laborious and time-consuming process that is often complicated when the tools and libraries used to develop AI models are inaccessible, deprecated, or incompatible with the computing resources available to new users. Furthermore, most of the existing packages specialized in predicting quantum mechanical properties of small molecules do not support crystals.

**APPROACH** With help from the ALCF's ThetaGPU system, Argonne researchers developed a computational framework that consolidates libraries, AI models and AI interpretability tools to study molecules, crystals, and metal-organic frameworks. The AI framework: 1) introduces novel data distillation techniques that accelerate model training and provide superior AI inference predictions; 2) enables hyperparameter tuning through the open-source library DeepHyper; 3) enables distributed AI model training; and 4) provides tools for interpretable AI inference of small-molecule quantum mechanics properties from public datasets such as QM9, and crystal properties from datasets such as hMOF.



Example molecules used to demonstrate the transferability of TorchMD-NET. Molecular dynamics trajectories of molecules on the left are used to train the TorchMD-NET models, which are then used to perform molecular dynamics simulations for the molecules on the right. *Image: Mach. Learn.: Sci. Technol. 4 (2023) 025036*

**RESULTS** The team added attention to the MPNN-transformer model, which improved its performance. They also used adjacency list format node and edge embedding schemes to improve the predictive capabilities of their modified crystal graph convolutional neural network model. In addition, they presented a novel application of TorchMD-NET, in which the AI model was finetuned to describe small molecules and then seamlessly used to describe other molecules with different structures, while still capturing physically realistic bond length and angle distributions. Their framework brings together several interpretability AI tools encompassing model performance attribution and scientific visualization; dimension reduction to explore clustering of molecules with similar properties; and metrics such as the roughness index to quantify the predictive performance of their AI models for quantum mechanics properties.

**IMPACT** Using this computational framework, researchers will be able to combine state-of-the-art AI models with datasets that are coupled with advanced supercomputing platforms to accelerate the discovery of new functional materials. This approach will catalyze the sharing of AI knowledge and tools in the context of molecular and crystal property prediction applications.

## PUBLICATIONS

Park, H., R. Zhu, E. A. Huerta, S. Chaudhuri, E. Tajkhorshid, and D. Cooper. "End-to-End AI Framework for Interpretable Prediction of Molecular and Crystal Properties," *Machine Learning: Science and Technology* (June 2023), IOP Publishing.  
<https://doi.org/10.1088/2632-2153/acd434>

# Large-Scale Simulations of Light-Activated Matter

PI Giulia Galli, University of Chicago/Argonne National Laboratory

AWARD INCITE


SYSTEM Theta, Polaris

Spin defects in wide-bandgap semiconductors provide a promising platform for the creation of qubits, or quantum bits, the basic units of quantum information technologies. Their synthesis and optical characterization, however, present considerable challenges, and the mechanisms responsible for their generation or annihilation are poorly understood. Researchers at Argonne National Laboratory and the University of Chicago carried out atomistic and first-principles simulations to elucidate spin defect formation processes in silicon carbide and developed a protocol, based on density functional theory (DFT) calculations, to predict their photoluminescence spectra.

**CHALLENGE** Understanding the formation of defects in semiconductors represents a challenge for both theory and experiments. The researchers focused here on the divacancy complex (VV) in silicon carbide, a key candidate for qubits and investigated its formation process at the atomistic level, using atomistic and quantum simulations coupled with enhanced sampling methods, and DFT electronic structure calculations.

**APPROACH** Leveraging the Theta and Polaris systems to model silicon carbide at the atomistic scale, the researchers employed a combination of codes to study defects in this system. The researchers simulated electronic excited state properties of heterogeneous materials, inclusive of defects, using coupled first-principles molecular dynamics and electronic structure methods beyond DFT, as implemented in the Qbox and WEST open-source codes.

**RESULTS** In a paper submitted to *Nature Communications*, the researchers presented a computational protocol to investigate the synthesis of point-defects at the atomistic level, and applied it to the study of VV in silicon carbide.



Pictorial representation of the free energy of defects in silicon carbide, representing the formation of divacancies (left) from monovacancies (right). Image: Elizabeth Lee, University of Chicago

Combining electronic structure calculations based on DFT and enhanced sampling techniques coupled with first-principles molecular dynamics, the researchers predict the optimal annealing temperatures for the formation of VVs at high temperature and show how to engineer the Fermi level of the material to optimize the defect's yield for several polytypes of silicon carbide. The team's results strongly agree with available experimental data and provide novel atomistic insights into point defect formation and annihilation processes as a function of temperature.

**IMPACT** This work provides predictions of the structural and electronic properties of heterogeneous systems that further develop our understanding of how defective and nanostructured materials interact with light, as well as validated data for systems germane to sustainable and quantum technologies. Its research lays the foundation for an integrated experimental and theoretical strategy for the design and optimization of spin defects for quantum technologies.

## PUBLICATIONS

Zhang, C., F. Gygi, and G. Galli. "Engineering the Formation of Spin-Defects from First Principles" (preprint). <https://doi.org/10.48550/arXiv.2303.16319>

# Materials Informatics

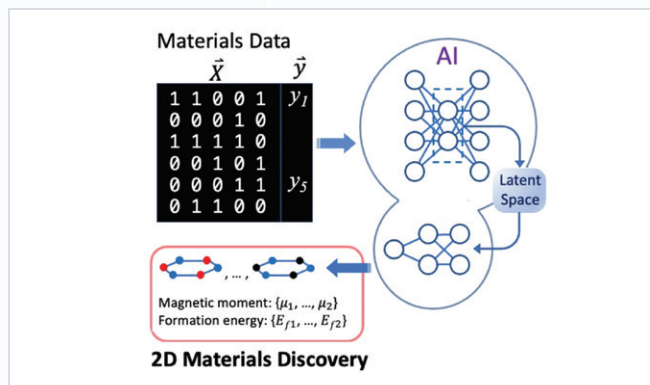
## Study of Two-Dimensional Magnetic Materials and Their Heterostructures

PI Trevor Rhone, Rensselaer Polytechnic Institute  
 AWARD Director's Discretionary  
 SYSTEM Theta

Two-dimensional (2D) materials, also referred to as van der Waals (vdW) materials, have attracted significant attention in recent years due to their unique properties and potential for technological innovation. Leveraging ALCF supercomputers and AI, a team from Rensselaer Polytechnic Institute is developing a materials informatics framework to speed up the search for promising 2D materials that could be used for applications in data storage, spintronics, and quantum computing technologies.

**CHALLENGE** The recent discovery of long-range magnetic ordering in 2D crystals has sparked research efforts to better understand the properties of magnetism in reduced dimensions and to identify additional 2D magnets with desirable properties for various applications. Traditional tools for materials discovery, based on experiments or first-principles calculations, can be costly and time-consuming. Identifying a means of accelerating the discovery of 2D materials with exotic electronic spin and charge degrees of freedom is urgently needed to advance scientific and technological innovation.

**APPROACH** Leveraging earlier work carried out through the ALCF Data Science Program, the team is developing a materials informatics framework that combines high-throughput density functional theory calculations and machine learning to explore a large number of vdW materials. Their approach employs a neural network model to learn the relationship between a crystal structure's chemical composition and its corresponding magnetic and thermodynamic properties. They used semi-supervised learning to train the model using both labeled and unlabeled data. This approach helps mitigate the lack of labeled data by creating a mathematical representation of the materials using unlabeled data. The trained neural



Materials data are combined with AI to implement a semi-supervised learning framework designed to search an enormous number of materials candidates to identify those with desirable properties. *Image: Adv. Theory Simul. 2023, 6, 2300019.*

network facilitates the fast and accurate prediction of materials properties for the entire materials space. This allows the team to quickly identify materials candidates that meet their search criteria.

**RESULTS** In a paper published in *Advanced Theory and Simulations*, the team applied their framework to study a large space of monolayer transition metal halides. The researchers used their AI model to predict the properties of various candidate materials; key properties include the magnetic moment and formation energy. They compared the predictions to validated data to successfully identify previously unexplored vdW magnetic materials with large magnetic moments and high chemical stability, such as  $\text{Cr}_2\text{Fe}_2\text{Br}_6\text{Cl}_6$ ,  $\text{Fe}_2\text{Cr}_2\text{Br}_{12}$ , and  $\text{Mn}_2\text{Cr}_2\text{Cl}_{12}$ .

**IMPACT** The team's materials informatics framework can be used to accelerate the discovery of promising 2D magnetic materials for potential applications in data storage and spintronics. The framework can be easily generalized for the exploration of materials with different crystal structures, extending its utility to other types of materials and applications.

#### PUBLICATIONS

Rhone, T. D., R. Bhattarai, H. Gavras, B. Lusch, M. Salim, M. Mattheakis, D. T. Larson, Y. Krockenberger, and E. Kaxiras. "Artificial Intelligence Guided Studies of van Der Waals Magnets," *Advanced Theory and Simulations* (April 2023), Wiley. <https://doi.org/10.1002/adts.202300019>



# Towards DMC Accuracy Across Chemical Space with Scalable $\Delta$ -QML

PI Paul Kent, Oak Ridge National Laboratory  
Anouar Benali, Argonne National Laboratory


AWARD INCITE  
SYSTEM Theta

Improving our ability to understand and predict the behavior and properties of molecules and materials is crucial to enabling the design and discovery of new materials for batteries, catalysts, semiconductors, and countless other applications. With this INCITE project, a multi-institutional team is advancing the use of quantum Monte Carlo (QMC) methods, coupled with machine learning, to provide accurate and reliable predictions of the fundamental properties of a wide range of molecules and materials.

**CHALLENGE** The predictive accuracy of quantum machine learning (QML) models trained on quantum chemistry data and used for the navigation of the chemical compound space is inherently limited by the predictive accuracy of the approximations used within the underlying quantum theory. To help QML models achieve the coveted threshold of chemical accuracy ( $\sim 1$  kcal/mol average deviation of calculated values from experimental measurements of atomization energies), the INCITE team is leveraging DOE supercomputers to demonstrate the usefulness of recently implemented and numerically efficient QMC methods for generating highly accurate training data.

**APPROACH** The team's primary application is QMCPACK, an open-source code for computing the electronic structure of atoms, molecules, 2D nanomaterials, and solids. As part of a recent study, the researchers used the ALCF's Theta supercomputer to couple QMCPACK with  $\Delta$ -QML-based surrogate methods to predict the energetics of large molecules at chemical accuracy and at a fraction of the computational cost of traditional machine learning methods.

**RESULTS** In a paper published in the *Journal of Chemical Theory and Computation*, the team showed that their



This figure shows the increase of accuracy as the size of the molecular fragment (amon) is increased to reach chemical accuracy. *Image: J. Chem. Theory Comput. 2023, 19, 6, 1711–1721*

$\Delta$ -QML framework can alleviate the computational burden of QMC such that it offers clear potential to support the formation of high-quality descriptions across the chemical space. Their work involved using Theta to conduct diffusion Monte Carlo (DMC) calculations on over 1,000 small amons containing up to five heavy atoms and covering parts of the QM9 database, which is used routinely for machine learning predictions of various chemical properties. This is the largest dataset ever computed with DMC and the first use of such a dataset for machine learning. The team's research suggests that the QMC training datasets of amons can predict total energies with near chemical accuracy throughout chemical space, setting the foundation for the study of larger databases.

**IMPACT** Using the  $\Delta$ -QML approach, the team was able to predict the energetics of large molecules at a reduced computational cost while maintaining chemical accuracy. The high efficiency of the  $\Delta$ -QML framework compared to traditional approaches indicates a path to use the computationally expensive but highly accurate QMC methodology in machine learning. This new method will allow researchers to study larger systems and predict the properties of molecules and materials more accurately, which could lead to significant advances in fields such as materials science, drug discovery, and energy research.

## PUBLICATIONS

Huang, B., O. A. von Lilienfeld, J. T. Krogel, and A. Benali. "Toward DMC Accuracy Across Chemical Space with Scalable  $\Delta$ -QML." *Journal of Chemical Theory and Computation* (March 2023), American Chemical Society. <https://doi.org/10.1021/acs.jctc.2c01058>

Physics |  Simulation


# Electron Kinetic Plasma Physics of Black Hole Accretion Flows

PI Dmitri Uzdensky, University of Colorado  
 AWARD INCITE  
 SYSTEM Theta

Accretion flows around supermassive black holes at the centers of galaxies emit electromagnetic radiation that is critical to understanding these active galactic nuclei, which influence galactic evolution. Interpreting observed radiation, however, requires detailed modeling of the complex multiscale plasma processes in accretion flows. Using petascale 3D particle-in-cell (PIC) simulations, this project investigates electron versus ion energization, nonthermal particle acceleration, and self-consistent synchrotron radiation for plasma processes likely ubiquitous in black-hole accretion, including plasma turbulence driven by the magnetorotational instability (MRI) or other forces, and collisionless magnetic reconnection.

**CHALLENGE** The team has identified three key links in the chain of plasma processes that lead from gravitational attraction of matter around a black hole to accretion and radiation. The development of the MRI leads to outward angular momentum transport that allows accretion; it also generates turbulence and current sheets leading to magnetic reconnection, both of which result in particle energization and hence radiation.

**APPROACH** To perform simulations, the researchers deployed the Zeltron application on ALCF supercomputers. Zeltron models relativistic, radiating, and rotating astrophysical plasmas from first principles using an explicit finite-difference time-domain, radiative electromagnetic PIC code. Zeltron can include the radiation reaction force (due to synchrotron and inverse Compton emission) in the particles' equations of motion, and simulate shearing box boundary conditions appropriate for studying MRI in black hole accretion disks.



Evolution of the pair-plasma MRI in a small 3D box ( $2 \times 2 \times 2 \lambda \text{ MRI}$ , where  $\lambda \text{ MRI}$  is the most unstable MRI wavelength), showing the spatial distribution of the radial ( $B_x$ ) and toroidal ( $B_y$ ) magnetic field at subsequent times during the simulation (in terms of  $P_0$ , the orbital period). Image: Dmitri Uzdensky, University of Colorado

**RESULTS** As detailed in a paper published in *The Astrophysical Journal*, the researchers explored nonlinear development of MRI turbulence in a pair plasma, employing fully kinetic PIC simulations in two and three dimensions carried out on Theta. This included studying the axisymmetric MRI with 2D simulations, explaining how and why the 2D geometry produces results that differ substantially from 3D MHD expectations; and then performing the largest such 3D simulations carried out to date, for which the team employed a novel shearing-box approach, demonstrating that 3D PIC models can reproduce mesoscale MRI dynamics in sufficiently large runs. Using the fully kinetic simulations, the team was able to describe the nonthermal particle acceleration and angular-momentum transport driven by the collisionless MRI.


**IMPACT** The work takes a critical step toward understanding the behavior of black holes in the universe. The simulations of plasma processes and energy conversion mechanisms in black hole accretion flows will be used to inform global magnetohydrodynamics computational and theoretical modeling, thus accounting for kinetic processes to predict radiation output and enable comparison to observations. Moreover, these simulations have the potential to significantly advance computational plasma physics.

## PUBLICATIONS

Bacchini, F., L. Arzamasskiy, V. Zhdankin, G. R. Werner, M. C. Begelman, and D. A. Uzdensky. "Fully Kinetic Shearing-Box Simulations of Magnetorotational Turbulence in 2D and 3D. I. Pair Plasmas," *The Astrophysical Journal* (October 2022), American Astronomical Society.  
<https://doi.org/10.3847/1538-4357/ac8a94>

# First-Principles Simulations of Black Hole Accretion Flows and Coronae

PI Luca Comisso, Columbia University  
AWARD INCITE  
SYSTEM Theta



This visualization shows the particle number density fluctuations in a turbulent region in proximity to the black hole. Lighter regions indicate higher particle density, while darker ones stand for lower particle density. *Image: Luca Comisso, Columbia University*

**Supermassive black holes are believed to exist at the center of every galaxy, where they play a crucial role in the evolution of galaxies. While no particles or electromagnetic radiation can escape a black hole, electromagnetic radiation can be emitted from the accretion material surrounding and feeding the black hole: the so-called accretion flow. The work aims to study how magnetized plasma turbulence and magnetic reconnection—two of the most fundamental and ubiquitous plasma processes, which were historically studied separately, but have recently been shown to be inevitably interconnected—lead to heating and particle acceleration in the accretion flows feeding massive black holes.**

**CHALLENGE** The team’s computational project focuses on two aspects of the physics of plasmas near black holes. The first concerns the exploration of the turbulence and reconnection interplay in low-luminosity accretion flows to understand the efficiency and the mechanisms of electron heating and acceleration, which is critical to producing physically motivated models that can be used to compare with observations such as those generated by the Event Horizon Telescope. The second involves the investigation of the self-consistent interplay of turbulence and radiation in the most magnetized regions around luminous black hole systems. The team is assessing the origin of the observed hard x-ray emission, and testing whether black hole coronae are likely sources of ultrahigh-energy cosmic rays, the most energetic particles in the universe.

**APPROACH** The primary code used for this work is TRISTAN-MP, a Fortran-based particle-in-cell (PIC) code optimized to handle ultra-relativistic flows. Like other PIC codes, TRISTAN-MP can model astrophysical plasmas from first principles as a collection of charged macro-particles that are

moved by leap-frog integration of the Lorentz force. The currents associated with the macro-particles are deposited on a grid, where Maxwell’s equations are discretized and then used to advance electromagnetic fields, with particle currents as the source term. Finally, the updated fields are extrapolated to the particle locations and used to compute the Lorentz force.

**RESULTS** The team has performed PIC simulations of particle acceleration in a turbulent plasma, focusing on the effects of plasma magnetization, temperature, and turbulence fluctuation levels. In tandem with the simulations, the researchers also performed detailed analyses of the resulting turbulence and particle statistics, about which papers are being drafted. Additional simulations were carried out with the newly developed Compton scattering module, so as to explore turbulent comptonization in magnetized coronae of accreting black holes.

**IMPACT** The team’s work will help advance our understanding of the plasma physics of turbulence and reconnection in the extreme relativistic conditions near black holes.

# Physics-Inspired AI Modeling of Complex Multiscale and Multiphysics Systems

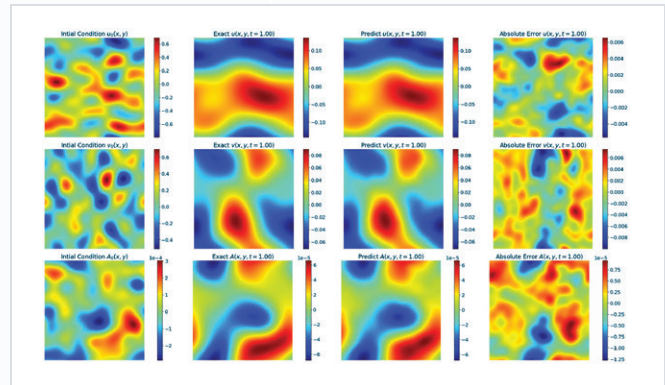
PI Eliu Huerta, Argonne National Laboratory  
 AWARD Director's Discretionary  
 SYSTEM ThetaGPU

AI is transforming the practice of scientific research, seamlessly combining human insight and digital capabilities to enhance our ability to address grand challenges in science and engineering. At Argonne National Laboratory, scientists are creating novel, physics-inspired AI approaches to accurately model complex multiscale and multiphysics systems that describe phenomena ranging from tsunamis to nuclear fusion.

**CHALLENGE** Physical processes evolve at different timescales and require different grid sizes to resolve subtle phenomena. Balancing accuracy, disparate timescales, and grid resolutions naturally leads to a need for compute-intensive simulations. To address this challenge, Argonne scientists have developed physics-informed neural operators (PINOs), incorporating physics and mathematics principles in the design, training, and optimization of AI models. These approaches allow for zero-shot super-resolution learning, which involves using AI surrogates to provide accurate predictions for higher-resolution data despite having only seen low-resolution data.

**APPROACH** Using the ALCF's ThetaGPU system, Argonne researchers have developed PINOs to solve partial differential equations (PDEs) that are ubiquitous in the study and modeling of physics phenomena using carefully curated datasets. They also developed PINOs to solve incompressible magnetohydrodynamics (MHD) equations used in plasma physics and astrophysics simulations. In both cases, the researchers have also released scientific software that unifies data production, construction of boundary conditions, and their use to train, validate, and test the performance of their AI surrogates.

**RESULTS** In the study focused on PDEs, the team demonstrated how PINOs can be used to accurately solve



PINO MHD simulations with Reynolds number of 100. From top to bottom, the rows show the velocity in the x direction,  $u$ ; the velocity in the y direction,  $v$ ; and the magnetic potential,  $A$ . From left to right, the columns show the initial condition; the ground truth; the PINO predictions; and the relative mean squared error. *Image: Mach. Learn.: Sci. Technol.* 4 (2023) 025022

such equations for simple systems (1D wave equation and 1D Burgers equation) and moderately complex systems (2D scalar, 2D inviscid and 2D vector Burgers equation, and 2D wave equation with non-constant coefficients). Their work also applied PINOs, for the first time, to coupled PDEs (2D linear and nonlinear shallow waters equations) that are used for coastal and tsunami modeling.


In another paper, the team detailed the first application of PINOs to model 2D incompressible MHD simulations. These AI models incorporate tensor Fourier neural operators as their backbone, implemented with the TensorLY package. These PINOs accurately capture the physics of MHD simulations for a broad range of laminar and turbulent systems. AI surrogates reported consistent results across three levels of data resolution, namely,  $64^2$ ,  $128^2$ , and  $256^2$ .

**IMPACT** The team's open-source scientific software and AI surrogates offer a promising alternative to the computationally demanding numerical methods traditionally used to solve PDEs for modeling multiscale and multiphysics phenomena. In addition, these software tools are expected to provide data-driven and physics-informed solutions that more accurately describe and identify novel features and patterns in experimental data.

## PUBLICATIONS

Rosofsky, S. G., H. Al Majed, and E. A. Huerta. "Applications of Physics Informed Neural Operators." *Machine Learning: Science and Technology* (May 2023), IOP Publishing.  
<https://doi.org/10.1088/2632-2153/acd168>

Rosofsky, S. G., and E. A. Huerta. "Magnetohydrodynamics with Physics Informed Neural Operators." *Machine Learning: Science and Technology* (July 2023), IOP Publishing.  
<https://doi.org/10.1088/2632-2153/ace30a>



Scientific visualization of Bragg diffraction peaks in a 15x15 pixel patch of an undeformed bicrystal gold sample (height denotes photon counts). Researchers processed the data from Argonne's Advanced Photon Source on ALCF supercomputers using a variety of AI models. *Image: ALCF Visualization and Data Analytics Team; Hemant Sharma, Argonne National Laboratory*

# ALCF Projects

## INCITE 2023

### BIOLOGICAL SCIENCES

#### High-Resolution Ensemble 3D Structures of Genome across Tissues

PI Jie Liang, University of Illinois at Chicago  
HOURS ALCF: 1,625,000 Node-Hours

### CHEMISTRY

#### Heterogeneous Catalysis as a Collective Phenomenon Within Dynamic Ensembles of States

PI Anastassia Alexandrova, University of California  
HOURS ALCF: 2,200,000 Node-Hours

### EARTH SCIENCE

#### Convection-Permitting Climate-Scale Simulations for Extreme Event Modeling

PI Rao Kotamarthi, Argonne National Laboratory  
HOURS ALCF: 1,600,000 Node-Hours

### ENGINEERING

#### DNS of Wall-Bounded Magnetohydrodynamic Turbulence at High Reynolds Number

PI Myoungkyu Lee, University of Alabama  
HOURS ALCF: 700,000 Node-Hours

#### First-Principles Simulation of Hypersonic Flight

PI Maninder Grover, University of Dayton Research Institute  
HOURS ALCF: 1,650,000 Node-Hours

#### Hypersonic Turbulent Boundary Layers Over Parameterized Wall Conditions

PI Lian Duan, The Ohio State University  
HOURS ALCF: 500,000 Node-Hours

### MATERIALS SCIENCE

#### AI-Guided Exascale Simulations of Quantum Materials Manufacturing and Control

PI Aiichiro Nakano, University of Southern California  
HOURS ALCF: 200,000 Node-Hours

#### Exascale Gyrokinetic Study of ITER Challenge on Power-Exhaust and ELM-Free Edge

PI Choongseok Chang, Princeton Plasma Physics Laboratory  
HOURS ALCF: 300,000 Node-Hours  
OLCF: 2,500,000 Node-Hours

#### Exascale Simulations of Quantum Materials

PI Paul Kent, Oak Ridge National Laboratory  
HOURS ALCF: 100,000 Node-Hours  
OLCF: 1,000,000 Node-Hours

### First-Principles Electron Dynamics in Complex Systems

PI Andre Schleife, University of Illinois at Urbana-Champaign  
HOURS ALCF: 1,000,000 Node-Hours

### High-Throughput Calculation of Materials Properties at Finite Temperature

PI Chris Wolverton, Northwestern University  
HOURS ALCF: 1,800,000 Node-Hours

### PHYSICS

#### Ab-initio Nuclear Structure and Nuclear Reactions

PI Gaute Hagen, Oak Ridge National Laboratory  
HOURS ALCF: 2,500,000 Node-Hours  
OLCF: 1,590,000 Node-Hours

#### Combining deep-learning with Physics-Based affinity estimation 2 (COMPBIO2)

PI Peter Coveney, University College London  
HOURS ALCF: 75,000 Node-Hours  
OLCF: 520,000 Node-Hours

#### Electron Kinetic Plasma Physics of Black Hole Accretion Flows

PI Dmitri Uzdensky, University of Colorado  
HOURS ALCF: 2,600,000 Node-Hours

#### Exascale Models of Astrophysical Thermonuclear Explosions

PI Michael Zingale, Stony Brook University  
HOURS ALCF: 1,000,000 Node-Hours  
OLCF: 700,000 Node-Hours

#### Next-Generation 3D Core-Collapse Supernova Simulations

PI Adam Burrows, Princeton University  
HOURS ALCF: 2,800,000 Node-Hours

#### Radiation-Dominated Black Hole Accretion

PI James Stone, Institute for Advanced Study  
HOURS ALCF: 110,000 Node-Hours  
OLCF: 1,000,000 Node-Hours

#### Unraveling How Lasers and Beams with Arbitrary Spatial and Temporal Structure Interact with Plasmas

PI Paulo Alves, University of California Los Angeles  
HOURS ALCF: 200,000 Node-Hours

## ALCC 2023–2024

### BIOLOGICAL SCIENCES

#### Probabilistic Comparative Modeling of Colorectal Cancer Screening Strategies

PI Jonathan Ozik, Argonne National Laboratory  
HOURS ALCF: 160,000 Node-Hours  
NERSC: 100,000

### Scaling Genomic Variant Callers to Leadership-Class Systems: A Collaboration Between VA-MVP and DOE

PI Ravi Madduri, Argonne National Laboratory  
HOURS ALCF: 210,000 Node-Hours

### CHEMISTRY

#### Microscopic Insight Into Transport Properties of Li-Battery Electrolytes

PI Wei Jiang, Argonne National Laboratory  
HOURS ALCF: 500,000 Node-Hours

#### Relativistic Quantum Dynamics in the Non-Equilibrium Regime

PI Eugene DePrince, Florida State University  
HOURS ALCF: 700,000 Node-Hours

### ENERGY TECHNOLOGIES

#### Accelerating Deployment of Next-Generation Nuclear Power Using High-Fidelity CFD

PI Dillon Shaver, Argonne National Laboratory  
HOURS ALCF: 500,000 Node-Hours

#### High Energy Density Physics of Inertial Confinement Fusion Ablator Materials

PI Ivan Oleynik, University of South Florida  
HOURS ALCF: 500,000 Node-Hours  
OLCF: 1,500,000 Node-Hours

#### Large Eddy Simulation on Flow and Heat Transfer Behavior in Involute Plate Research Reactor Supporting the Needs of the Materials Management and Minimization (M3) Reactor Conversion Program

PI Yiqi Yu, Argonne National Laboratory  
HOURS ALCF: 510,000 Node-Hours

### ENGINEERING

#### Two-Phase Flow Interface Capturing Simulations

PI Igor Bolotnov, North Carolina State University  
HOURS ALCF: 200,000 Node-Hours  
NERSC: 300,000 Node-Hours

### MATERIALS SCIENCE

#### Computational Design of Novel Semiconductors for Power and Energy Applications

PI Feliciano Giustino, University of Texas  
HOURS ALCF: 100,000 Node-Hours

#### Large Scale Simulations of Materials for Quantum Information Science

PI Giulia Galli, University of Chicago  
HOURS ALCF: 600,000 Node-Hours  
NERSC: 400,000 Node-Hours

### Quantum Accurate Large-Scale Atomistic Simulations of Advanced Fusion Reactor Materials

PI Aidan Thompson, Sandia National Laboratories  
HOURS ALCF: 850,000 Node-Hours  
OLCF: 500,000 Node-Hours  
NERSC: 250,000 Node-Hours

### PHYSICS

#### Continuum Limit Lattice Calculation of Direct CP-violation in Kaon Decays

PI Christopher Kelly, Brookhaven National Laboratory  
HOURS ALCF: 135,000 Node-Hours

#### Energy Partition and Particle Acceleration in Laboratory Magnetized Shocks

PI Frederico Fiuza, SLAC National Accelerator Laboratory  
HOURS ALCF: 300,000 Node-Hours  
NERSC: 150,000 Node-Hours

#### Hadronic Contributions to the Muon G-2 from Lattice QCD

PI Thomas Blum, University of Connecticut  
HOURS ALCF: 5,000 Node-Hours  
OLCF: 3,283,000 Node-Hours

#### Using GPU to Reconstruct LHC Collisions Recorded with the CMS Detector

PI Dirk Hufnagel, Fermi National Accelerator Laboratory  
HOURS ALCF: 70,000 Node-Hours

#### Short Range Correlations from a Quantum Monte Carlo Perspective

PI Noemi Rocco, Fermi National Accelerator Laboratory  
HOURS ALCF: 730,000 Node-Hours

#### The Spectrum and Structure of Hadrons

PI Robert Edwards, Jefferson Laboratory  
HOURS ALCF: 300,000 Node-Hours

### ALCC 2022–2023

### BIOLOGICAL SCIENCES

#### Probabilistic Comparative Modeling of Colorectal Cancer Screening Strategies

PI Jonathan Ozik, Argonne National Laboratory  
HOURS ALCF: 283,000 Node-Hours

### COMPUTER SCIENCE

#### A Multiscale Surrogate Model for Fracture Evolution Using DeepONet

PI George Karniadakis, Brown University  
HOURS ALCF: 50,000 Node-Hours  
OLCF: 40,000 Node-Hours  
NERSC: 60,000 Node-Hours

### CHEMISTRY

#### Microscopic Insight Into Transport Properties of Li-Battery Electrolytes

PI Wei Jiang, Argonne National Laboratory  
HOURS ALCF: 500,000 Node-Hours

### EARTH SCIENCE

#### A Climate Model Ensemble for Understanding Future Changes to Extreme

PI Paul Ullrich, University of California  
HOURS ALCF: 900,000 Node-Hours  
NERSC: 300,000 Node-Hours

### ENERGY TECHNOLOGIES

#### High-Fidelity CFD Simulations Supporting the Needs of Industry and the DOE

PI Dillon Shaver, Argonne National Laboratory  
HOURS ALCF: 400,000 Node-Hours  
OLCF: 400,000 Node-Hours  
NERSC: 100,000 Node-Hours

#### Investigation of Flow and Heat Transfer Behavior in Involute Plate Research Reactor with Large Eddy Simulation to Support the Conversion of Research Reactors to Low Enriched Uranium Fuel

PI Yiqi Yu, Argonne National Laboratory  
HOURS ALCF: 500,000 Node-Hours

#### Modeling Operating Conditions in the US East Coast Offshore Wind Energy Lease Areas

PI Sara Pryor, Cornell University  
HOURS ALCF: 142,000 Node-Hours

#### Predictive Simulations of Inertial Confinement Fusion Ablator Materials

PI Ivan Oleynik, University of South Florida  
HOURS ALCF: 150,000 Node-Hours

### ENGINEERING

#### HFIR DNS Simulation

PI Emilian Popov, Oak Ridge National Laboratory  
HOURS ALCF: 224,000 Node-Hours

### MATERIALS SCIENCE

#### Computational Design of Novel Semiconductors for Power and Energy

PI Feliciano Giustino, University of Texas  
HOURS ALCF: 883,000 Node-Hours

### PHYSICS

#### Cosmological Hydro Simulations to Explore the High and Low-Redshift Universe

PI Zarija Lukić, Lawrence Berkeley National Laboratory  
HOURS ALCF: 100,000 Node-Hours  
OLCF: 50,000 Node-Hours  
NERSC: 50,000 Node-Hours

**Energy Partition and Particle Acceleration in Laboratory Magnetized Shocks**

PI Frederico Fiuza, SLAC National Accelerator Laboratory  
 HOURS ALCF: 860,000 Node-Hours

**High Precision Hadronic Vacuum Polarization Contribution to the Muon Anomalous Magnetic Moment using Highly Improved Staggered Quarks**

PI Steven Gottlieb, Indiana University  
 HOURS ALCF: 100,000 Node-Hours  
 OLCF: 1,000,000 Node-Hours  
 NERSC: 100,000 Node-Hours

**Particle-in-Cell Simulations of Beam-Driven, Field-Reversed Configuration Plasmas**

PI Jaeyoung Park, TAE Technologies, Inc.  
 HOURS ALCF: 400,000 Node-Hours

**Short Range Correlations from a Quantum Monte Carlo Perspective**

PI Noemi Rocco, Fermi National Accelerator Laboratory  
 HOURS ALCF: 730,000 Node-Hours

**The Spectrum and Structure of Hadrons**

PI Robert Edwards, Jefferson Laboratory  
 HOURS ALCF: 300,000 Node-Hours

**AURORA EARLY SCIENCE PROGRAM****Accelerated Deep Learning Discovery in Fusion Energy Science**

PI William Tang, Princeton Plasma Physics Laboratory

**Dark Sky Mining**

PI Salman Habib, Argonne National Laboratory

**Data Analytics and Machine Learning for Exascale Computational Fluid Dynamics**

PI Kenneth Jansen, University of Colorado Boulder

**Enabling Connectomics at Exascale to Facilitate Discoveries in Neuroscience**

PI Nicola Ferrier, Argonne National Laboratory

**Exascale Computational Catalysis**

PI David Bross, Argonne National Laboratory

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

PI Anouar Benali, Argonne National Laboratory

**Extreme-Scale Cosmological Hydrodynamics**

PI Katrin Heitmann, Argonne National Laboratory

**Extreme-Scale In-Situ Visualization and Analysis of Fluid-Structure-Interaction Simulations**

PI Amanda Randles, Duke University

**Extreme-Scale Unstructured Adaptive CFD**

PI Kenneth Jansen, University of Colorado Boulder

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

PI C.S. Chang, Princeton Plasma Physics Laboratory

**Machine Learning for Lattice Quantum Chromodynamics**

PI William Detmold, Massachusetts Institute of Technology

**Many-Body Perturbation Theory Meets Machine Learning to Discover Singlet Fission Materials**

PI Noa Marom, Carnegie Mellon University

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

PI Theresa Windus, Iowa State University and Ames Laboratory

**Simulating and Learning in the ATLAS Detector at the Exascale**

PI Walter Hopkins, Argonne National Laboratory

**Virtual Drug Response Prediction**

PI Rick Stevens, Argonne National Laboratory

**DIRECTOR'S DISCRETIONARY**

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

**BIOLOGICAL SCIENCES****Bioinformatics Codes for Polygenic Risk Scores**

PI Ravi Madduri, Argonne National Laboratory

**Large Ensemble Model of Single-Cell 3D Genome Structures**

PI Jie Liang, University of Illinois at Chicago

**Predicting Transcription Factor Binding and Other Epigenetic Features to Gain Insight into the Biology of Diseases**

PI Ravi Madduri, Argonne National Laboratory

**RNA-to-Image Synthesis**

PI Ravi Madduri, Argonne National Laboratory

**Targeting Intrinsically Disordered Proteins Using Artificial Intelligence-Driven Molecular Simulations**

PI Arvind Ramanathan, Argonne National Laboratory

**Understanding the Mechanism of Ligand-Induced Conformational Dynamics of HIV-1 Protease and the Effects of Mutations**

PI Ao Ma, University of Illinois Chicago

**CHEMISTRY****Multimodal Imaging with Intense X-ray Pulses**

PI Phay Ho, Argonne National Laboratory

**COMPUTER SCIENCE****APS Beamline Data Processing and Analysis**

PI Rafael Vescovi, Argonne National Laboratory

**Advanced Photon Source Data Processing**

PI Nicholas Schwarz, Argonne National Laboratory

**Optimizing Bayesian Neural Networks for Scientific Machine Learning Applications**

PI Murali Emani, Argonne National Laboratory

**EARTH SCIENCE****Convective-Resolved Regional Climate Simulations**

PI Rao Kotamarthi, Argonne National Laboratory

**Data-Driven Coupling Methods for Atmospheric-Ocean Interactions**

PI Shinhoo Kang, Argonne National Laboratory

**ENERGY TECHNOLOGIES****High-Fidelity Simulations of Gas Turbine Combustors for Sustainable Aviation Applications**

PI Sicong Wu, Argonne National Laboratory

**In Situ Visualization and Analysis of Blood Flow**

PI Jifu Tan, Northern Illinois University

**Investigation of Flow and Heat Transfer Behavior in Involute Plate Research Reactor with Large Eddy Simulation to Support the Conversion of Research Reactors to Low Enriched Uranium Fuel**

PI Yiqi Yu, Argonne National Laboratory

**Multiphase Simulations of Nuclear Reactor Thermal Hydraulics**

PI Igor A Bolotnov, North Carolina State University

**ENGINEERING****Turbulent Rayleigh-Benard Convection in Suspensions of Bubbles**

PI Parisa Mirbod, University of Illinois at Chicago



## **MATERIALS SCIENCE**

### **Adsorptive CO<sub>2</sub> Removal from Dilute Sources**

PI John Low, Argonne National Laboratory

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

PI Jacqueline Cole, University of Cambridge

### **Materials Informatics Study of Two-Dimensional Magnetic Materials and Their Heterostructures**

PI Trevor Rhone, Harvard University

## **PHYSICS**

### **Analytic Continuation of Interacting Fermion Spectra**

PI Adrian Del Maestro, University of Tennessee

### **Neural Network Quantum States for Atomic Nuclei**

PI Alessandro Lovato, Argonne National Laboratory

### **Reproducible and Accelerated Physics-Inspired Neural Networks**

PI Eliu Huerta, Argonne National Laboratory

## About the 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](http://www.anl.gov).

## Availability of this Report (ANL/ALCF-23/2)

---

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.  
Alexandria, 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), Bill Allcock (Director of Operations), Susan Coghlan (ALCF-X Project Director), Kalyan Kumaran (Director of Technology), Jini Ramprakash (Deputy Division Director), and Katherine Riley (Director of Science)

Editorial Team: Beth Cerny, Jim Collins, Nils Heinonen, Logan Ludwig, and Laura Wolf

Design and production: Sandbox Studio, Chicago

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

