Argonne Leadership Computing Facility

2020 Annual Report

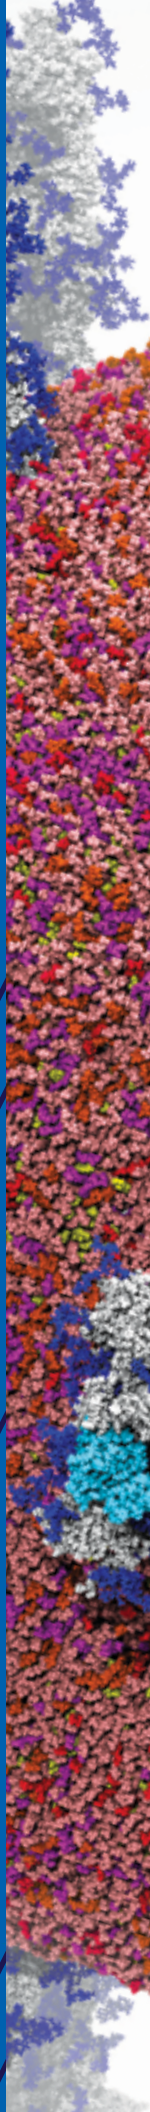

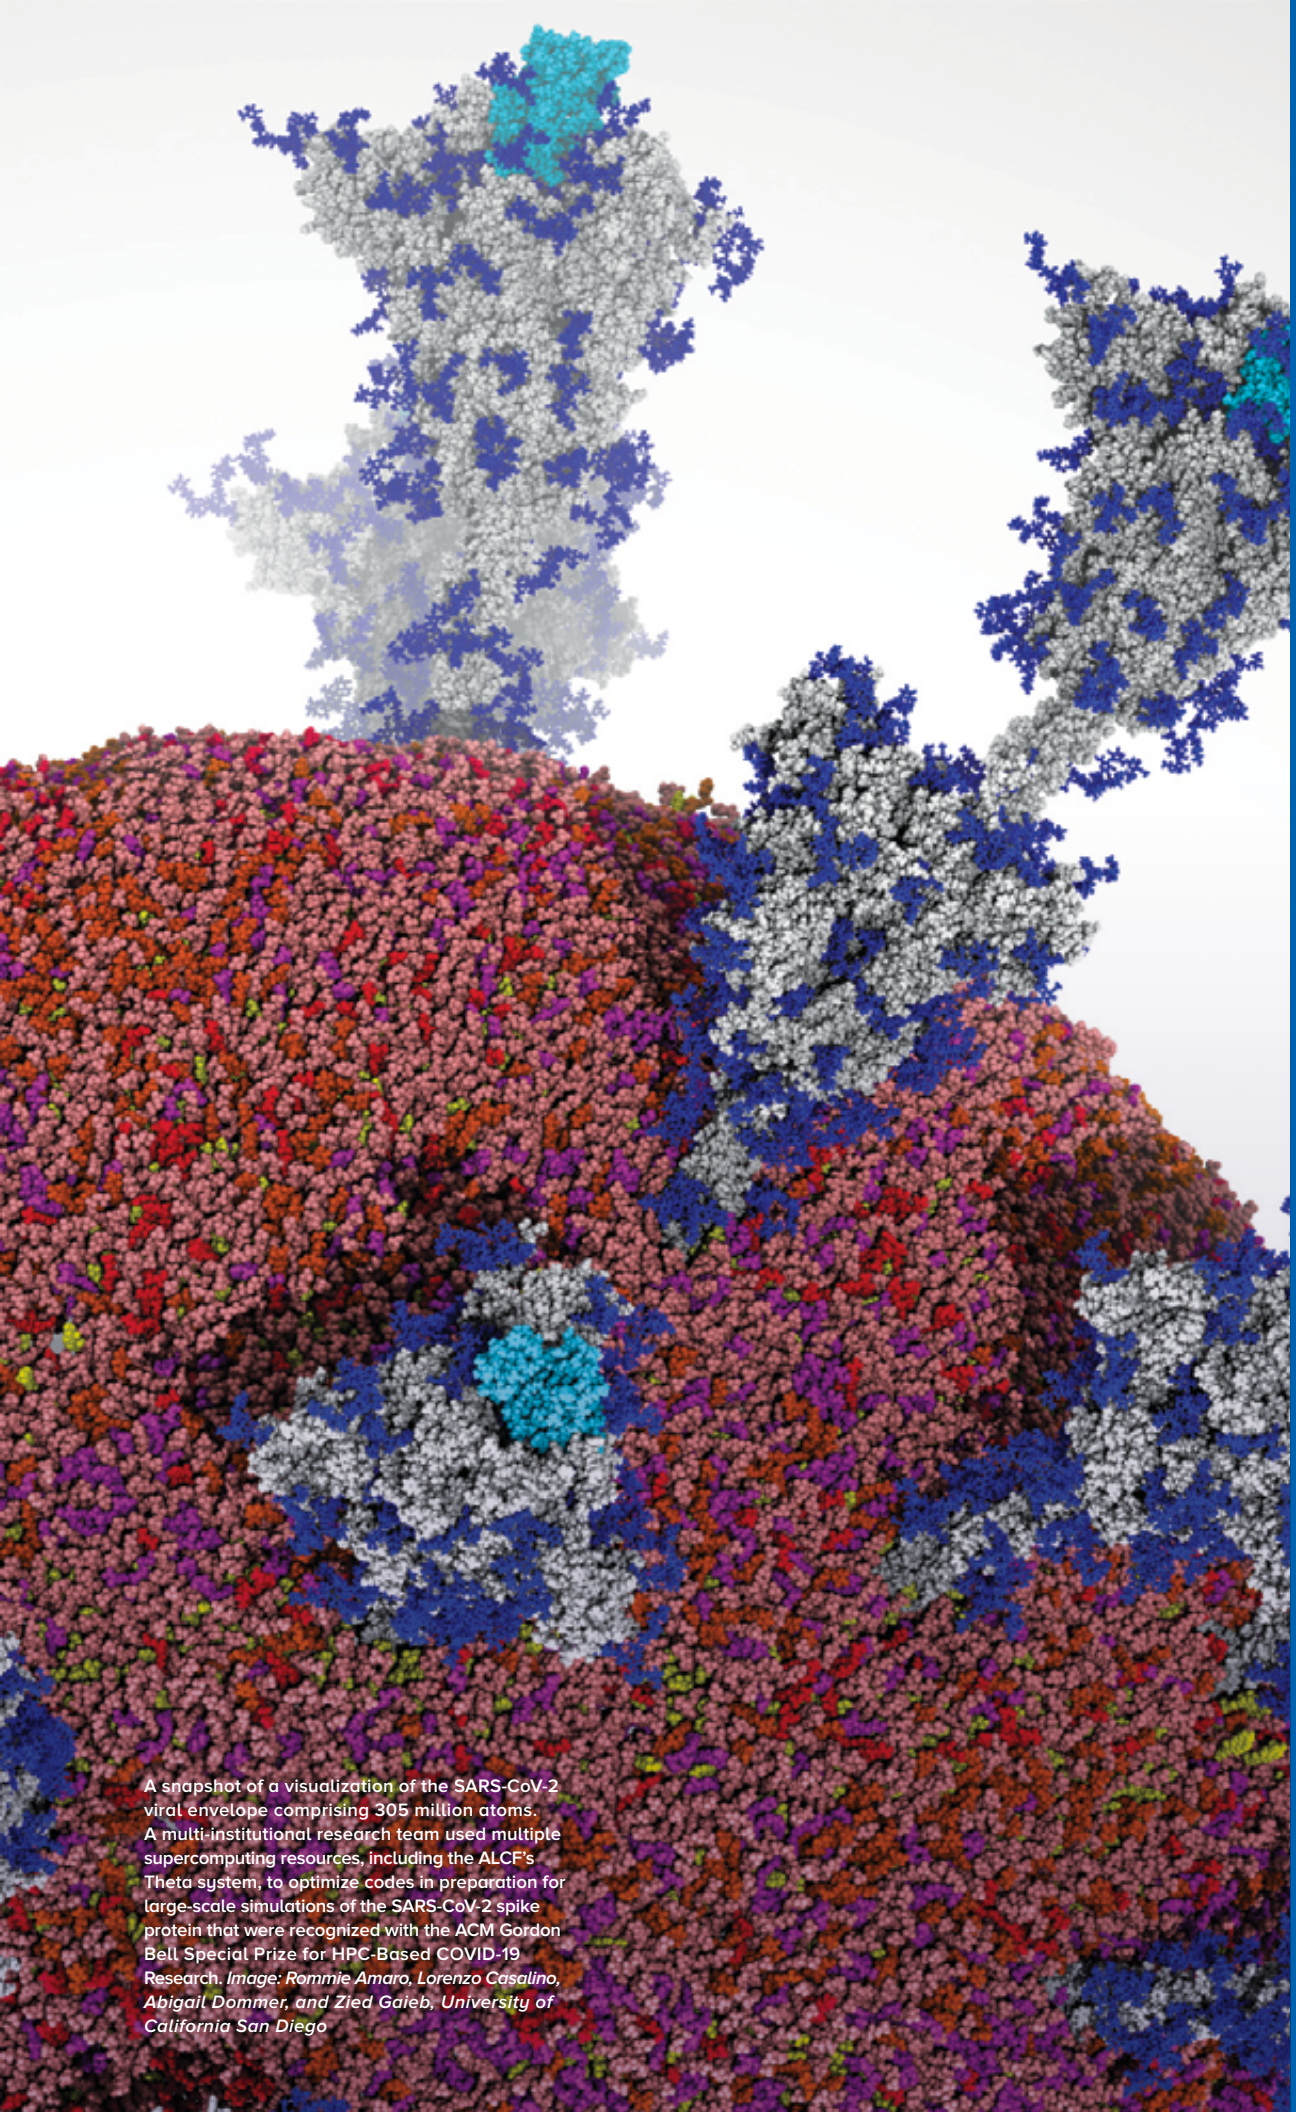

A snapshot of a visualization of the SARS-CoV-2 viral envelope comprising 305 million atoms. A multi-institutional research team used multiple supercomputing resources, including the ALCF's Theta system, to optimize codes in preparation for large-scale simulations of the SARS-CoV-2 spike protein that were recognized with the ACM Gordon Bell Special Prize for HPC-Based COVID-19 Research. *Image: Rommie Amaro, Lorenzo Casalino, Abigail Dommer, and Zied Gaieb, University of California San Diego*

CONTENTS

2020 Annual Report

YEAR IN REVIEW

The Argonne Leadership Computing Facility enables breakthroughs in science and engineering by providing supercomputing resources and expertise to the research community.

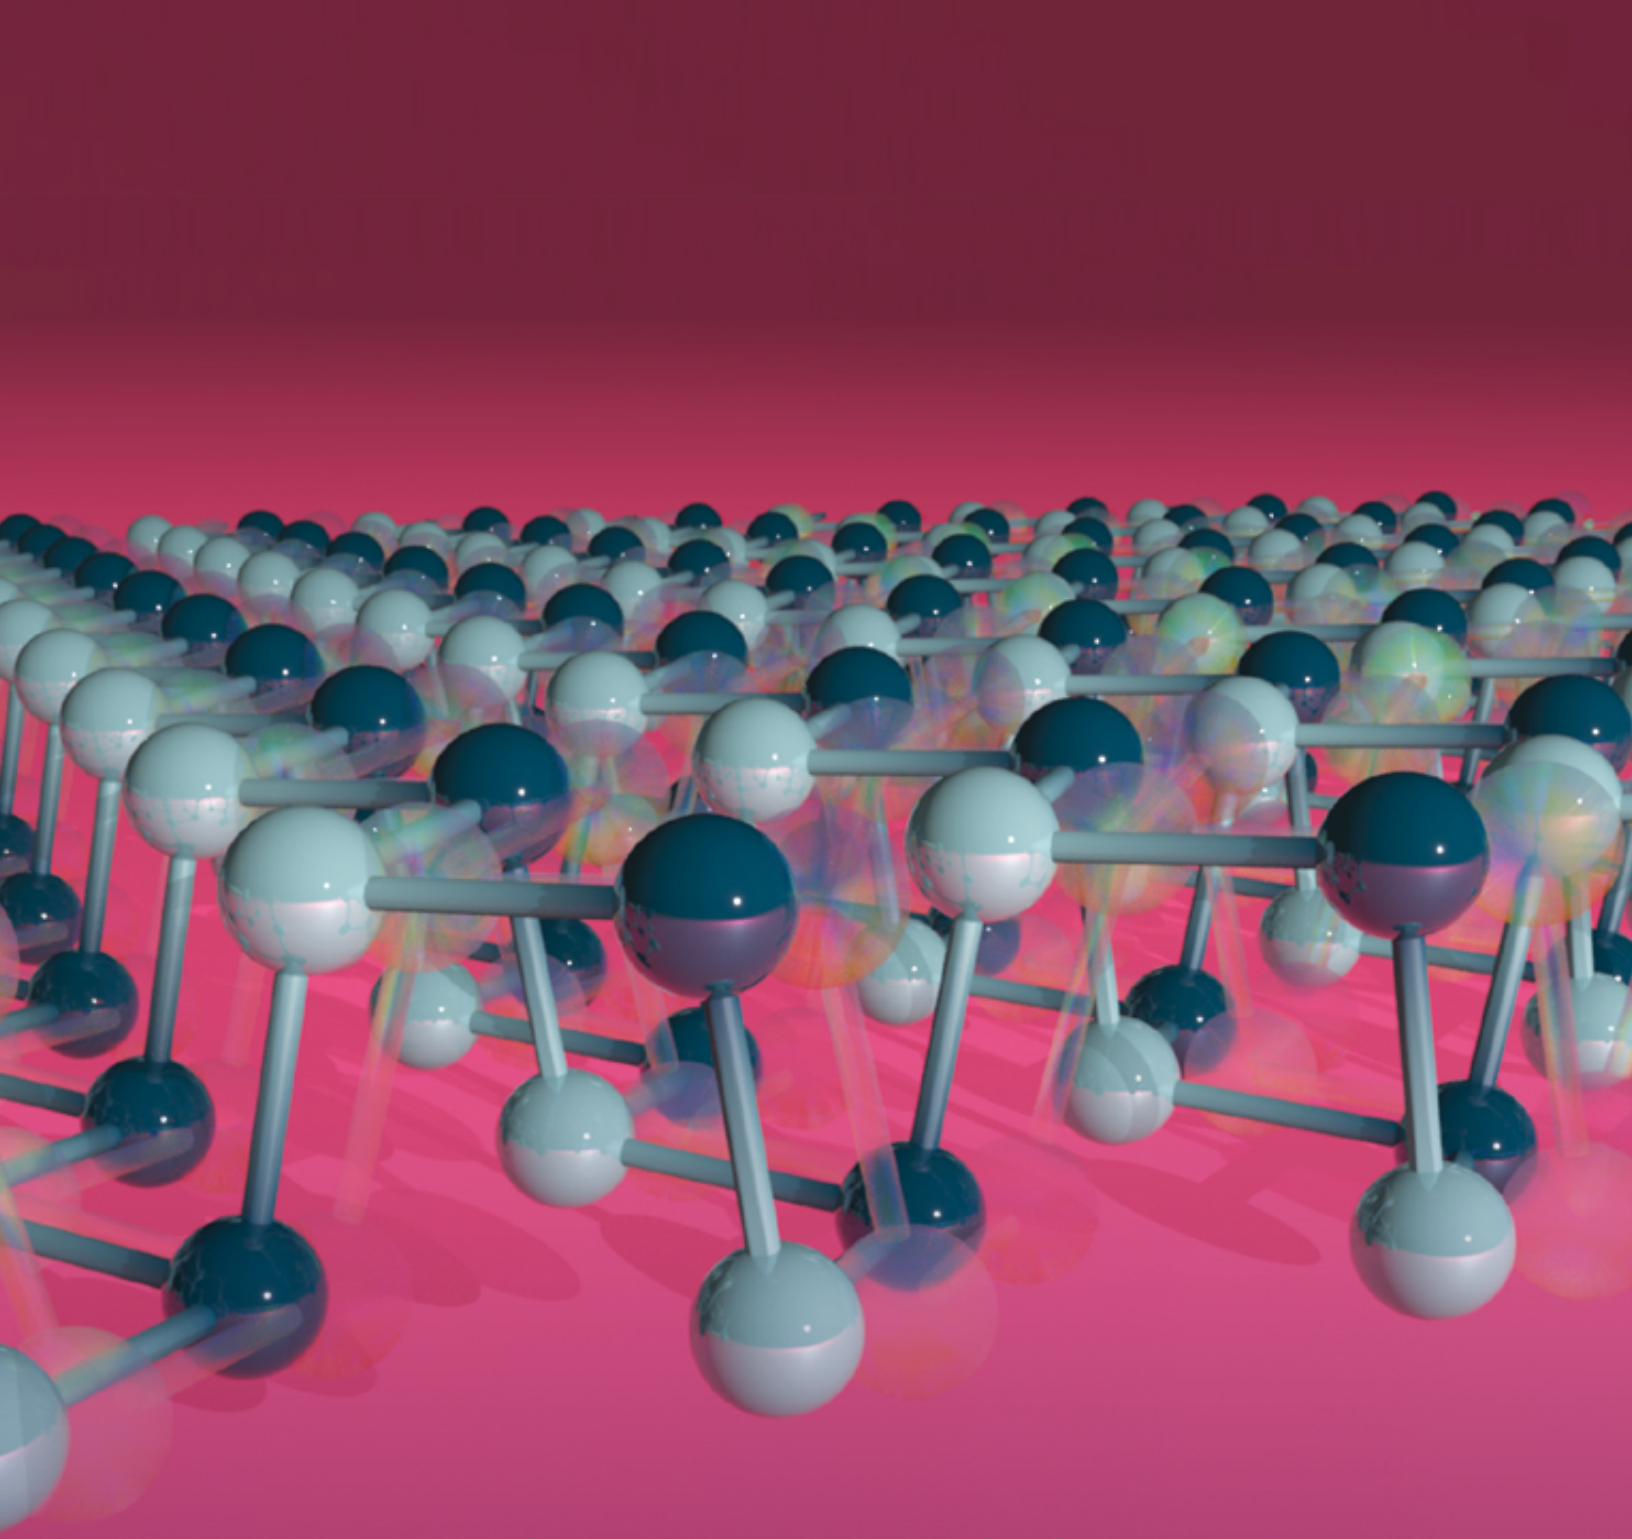

Using a newly developed structural optimization algorithm within Quantum Monte Carlo, researchers from Argonne and Oak Ridge National Laboratories were able to determine the optimized geometry of a 2D nanomaterial called germanium selenide (GeSe) for the first time. This image shows the fully optimized GeSe structure (colored) from the initial structure (clear). *Image: ALCF Visualization and Data Analysis Team*

ALCF Leadership

MICHAEL E. PAPKA
ALCF Director

The past year was marked by unexpected challenges, but the ALCF made meaningful progress without pause, both as a workforce and as an advanced scientific computing center.

After quickly pivoting to an alternative work mode last spring, we allocated computing cycles for COVID-19 research, expanded Theta with GPUs to bring more computational power to addressing the pandemic, and joined a national effort as part of the COVID-19 High Performance Computing (HPC) Consortium to coordinate a unified computational response.

All the while, and despite the pandemic's impact on our day-to-day lives, most ALCF users, already accustomed to working on a remote basis, continued to produce world-class science. And while some very noteworthy research carried out at the ALCF this year were projects that aimed to shed light on the coronavirus, there were also many, many other achievements unrelated to COVID-19. Theta delivered over 20 million node-hours alone to 14 INCITE projects, advancing studies ranging from modeling supernova explosions to developing a leading-edge climate and Earth system model to exploring potential novel cancer therapies.

Mira, the ALCF's previous production system, delivered an additional 67 million node-hours in support of two machine-scale projects before terminating all operation in early March 2020. One research project in lattice quantum chromodynamics (QCD) moved the needle further in efforts to address fundamental questions in high energy and nuclear physics that are essential to meeting a number of DOE Office of Science milestones. The other project, called the Last Journey, was a six-month, massive simulation of the cosmos—one of the five largest cosmological simulations ever—and provided data that will form the basis for sky maps used by numerous surveys. It also prepared the science team to run even more extensive cosmological simulations on future machines, like our Aurora exascale system.

As preparations for Aurora continued to ramp up, our collaborative work through the Early Science Program and DOE's Exascale Computing Project, along with the emerging testbeds in Argonne's Joint Laboratory for System Evaluation (JLSE), helped science teams from industry, academia, and government labs to advance their code development efforts for research. JLSE currently operates four testbeds used to prepare for Aurora, alongside many others used to support more than 80 projects ranging from application portability to software development to tools and compiler development.

The ALCF's newly established AI pathfinding efforts are allowing researchers to evaluate machine-learning-based HPC science applications running on system accelerators from a variety of vendors. This experimental hardware and software environment is now available to the Argonne community (with plans to expand to a wider audience in the near future).

On the data sharing front, progress toward a new global filesystem in 2020 included the configuration of two systems, Eagle and Grand, each with a capacity of 100 petabytes, and peak Input/Output speeds of 650 gigabytes per second. In addition, Eagle will soon start supporting data sharing via Globus, beginning in 2021.

This Annual Report showcases the progress we've made readying our facility for exascale science, and the latest high-impact science produced by our user community. In all, I am proud of how we finished 2020, and I remain profoundly thankful to our staff members for keeping the science flowing.

JINI RAMPRAKASH
ALCF Deputy Director

MARK FAHEY
ALCF Director of Operations

While most of the ALCF team spent the bulk of 2020 working remotely due to COVID-19, we made several adaptations to successfully carry out our scientific mission and protect the health and safety of the ALCF community. Using technologies like Zoom and Slack, we were able to stay connected through regular team meetings, online chats and collaborations, and virtual happy hours. We transitioned our annual summer student program to an all-virtual format, supporting more than 30 students from across the country. And we continued to host interactive facility tours for various outreach events via live broadcasts from the ALCF data center and visualization lab.

In October, we participated in a DOE roundtable, "Lessons from the COVID Era and Visions for the Future," to create a long-term vision for building an even stronger and more resilient scientific ecosystem for all of DOE's user facilities. We also contributed to a number of activities to broaden the ALCF's presence in the broader research community. For example, we helped moderate an AI panel at the DOE-sponsored Community of Interest Workshop on the Future of Scientific Methodologies, and took part in a virtual User Facilities Awareness Workshop designed to forge new partnerships with researchers from historically underrepresented groups. In addition, ALCF staff attended a series of Argonne townhall meetings to develop roadmap for integrating experimental work at the Advanced Photon Source with the lab's supercomputing capabilities to accelerate data-driven science.

We also joined the U.S. Government ORCID consortium, launched by DOE's Office of Scientific and Technical Information (OSTI) in 2020, to stay better connected with our user community's achievements and progress. As a member organization, we have access to the ORCID registry, which includes trusted data on researchers' publications, affiliations, and more.

In March and April of 2020, we physically dismantled Mira, our retired IBM Blue Gene/Q supercomputer. The decommissioning effort involved many steps including draining and disconnecting all the water connections; removing the power whips, 40 miles of optical cables from under the floor, and the compute node boards from each rack to lighten the rack weight from 4,400 pounds by roughly half; and, finally, working with a recycling company to remove and recycle all of the parts.

We also deployed multiple new resources in 2020, including an HDR Infiniband Storage Area network providing 200 Gb/s and two 100 PB filesystems, named Eagle and Grand, based on Cray's ClusterStor E1000 technology providing 650 GB/s bandwidth each; a global storage space for ALCF projects; and a novel community data sharing capability. In addition, we deployed a HPE Cray EX testbed system that has proved quite useful in preparing for our Aurora exascale machine. And finally, we upgraded Theta by augmenting it with NVIDIA DGX A100 nodes to meet DOE's high-priority request to provide an essential resource to support COVID-19 research. Our team acted quickly to make the enhanced system available to researchers despite a delay in receiving all of the hardware due to the pandemic.

In July, we completed a massive software upgrade to Theta to ensure the software versions were compatible with the other machines we installed this past year. Also in 2020, we implemented several important improvements to UserBase3, the ALCF's user accounting system. To name a few, we migrated the "user agreement and user acknowledgement" database to simplify the account request process and we worked with Argonne's Business Information Systems team to make significant improvements to the lab's workflow for foreign nationals requesting access to ALCF resources.

KATHERINE RILEY
ALCF Director of Science

KALYAN KUMARAN
ALCF Director of Technology

While the COVID-19 pandemic made for a challenging year, we are fortunate it has not impacted the ALCF's ability to enable groundbreaking science. Due to the nature of our facility, ALCF users have always accessed our supercomputers remotely, so their scientific progress kept moving forward despite the pandemic. In fact, the Theta expansion, along with our staff expertise, supported 10 projects that helped to advance our understanding of the virus. This included an agent-based modeling campaign to simulate and predict the spread of COVID-19 through populations and an effort to use AI and molecular dynamics simulations to search for promising drug candidates.

Our user community also continued to pursue scientific discovery and innovation across all disciplines. Leveraging Theta, researchers carried out several impactful studies, including the creation of automated databases for materials discovery and an exploration of electronic stopping dynamics in DNA that could help inform improved cancer treatments. And in the final runs on our now-retired Mira supercomputer, an Argonne-led team carried out some of the largest cosmological simulations ever performed, generating a massive data set that will serve the scientific community for years to come. All in all, our users published approximately 250 peer-reviewed publications in 2020.

As we look ahead to Aurora's arrival in 2022, we continued to ramp up our activities to prepare users for the exascale system. Our collaborative work through the Early Science Program and DOE's Exascale Computing Project has helped numerous teams advance code development efforts for Aurora. We also hosted a number of workshops, webinars, and hackathons to provide details and guidance on using various aspects of the Aurora hardware and software environment.

Despite 2020 being a difficult year for obvious reasons, I'm proud to say that we continued to efficiently support science on our production resources and advance our preparations for Aurora through collaborations with Intel on the non-recurring engineering (NRE) contract and with several teams participating in DOE's Exascale Computing Project (ECP) and the Aurora Early Science Program (ESP).

As part of the Aurora efforts, we deployed two Intel discrete GPU testbeds, Yarrow and Arcticus, to allow ECP and ESP teams to develop, optimize, and scale applications and software for Aurora. Our team collaborated closely with Intel on developing oneAPI HPC components, including compilers, math libraries, performance tools, and debugger support for Intel GPUs; articulating the vision for the data and learning software stack and the Python ecosystem for oneAPI data and AI components; porting Kokkos and Raja programming models; and enabling in-situ data rendering capabilities on the Aurora testbeds.

To support the COVID-19 research on ThetaGPU, we installed and optimized various data and learning libraries and developed user documentation for the enhanced system. Our team also recommended changing the adaptive routing settings on Theta from its default value, leading to improved application runtimes for users. Additionally, we continued to develop a variety of tools and services for the ALCF user community. Our staff created open-source tools, such as THAPI (Tracing Heterogenous APIs), Iprof (Intel GPU profiling), and OpenMP Validator for Offload (OVO), and deployed them on the Aurora testbeds to help users port applications to heterogenous node architectures. We also enhanced the DeepHyper AutoML tool to advance AI-driven science, and supported several projects using our Balsam workflow service, including an effort to enable near-real-time analysis of experimental data from DOE light sources.

ALCF at a Glance

The Argonne Leadership Computing Facility (ALCF) is a U.S. Department of Energy (DOE) Office of Science User Facility that enables breakthroughs in science and engineering by providing supercomputing resources and expertise 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 its users to accelerate the pace of discovery and innovation across disciplines.

As a key player in the nation's efforts to deliver future exascale computing capabilities, the ALCF is helping to advance scientific computing through a convergence of simulation, data science, and machine learning methods.

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 supercomputers that are orders of magnitude more powerful than the systems typically used for open scientific research.

Node-hours of compute time

101M

Active projects*

369

Facility users*

1,174

**Fiscal year 2020*

Publications

246

2020 ALCF Users by Affilliation

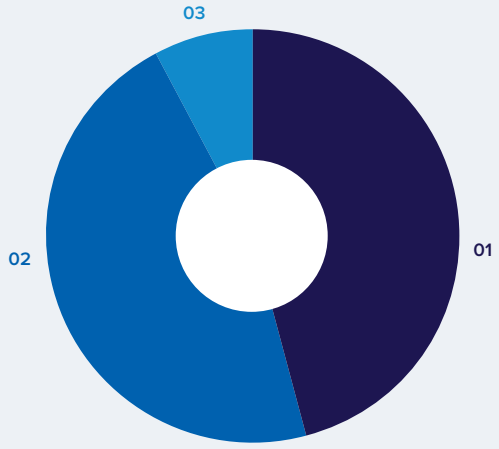

2020 U.S. ALCF Users by State

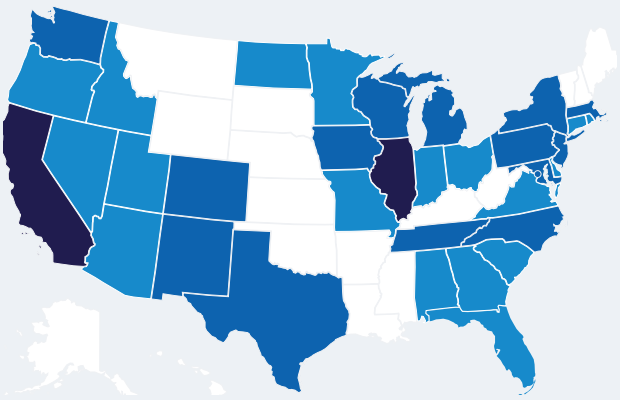

01 Academia

545

02 Government

539

03 Industry

90

100+ Users

California
Illinois

11–100 Users

Colorado
Washington D.C.
Iowa
Maryland
Massachusetts
Michigan
North Carolina
New Jersey
New Mexico
New York
Pennsylvania
Tennessee
Texas
Washington
Wisconsin

01–10 Users

Alabama
Arizona
Connecticut
Delaware
Florida
Georgia
Idaho
Indiana
Minnesota
Missouri
Nevada
North Dakota
Ohio
Oregon
Rhode Island
South Carolina
Utah
Virginia

ADVANCING SCIENCE WITH HPC

The ALCF's high-performance computing (HPC) resources and expertise help scientists tackle some of the world's most complex and challenging scientific problems.

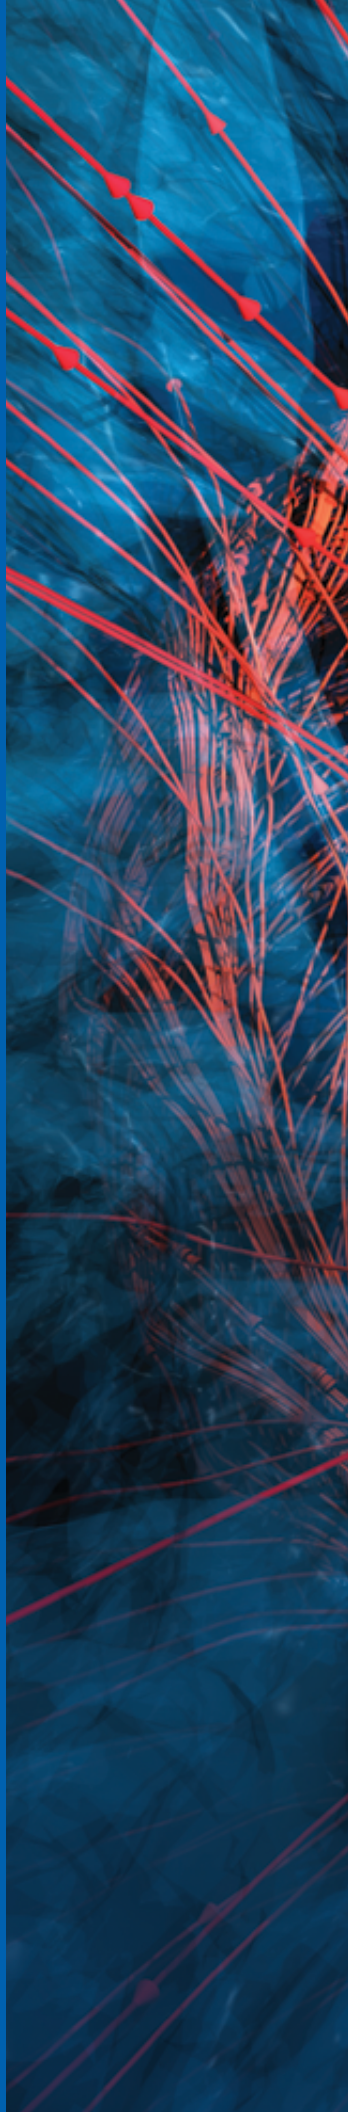

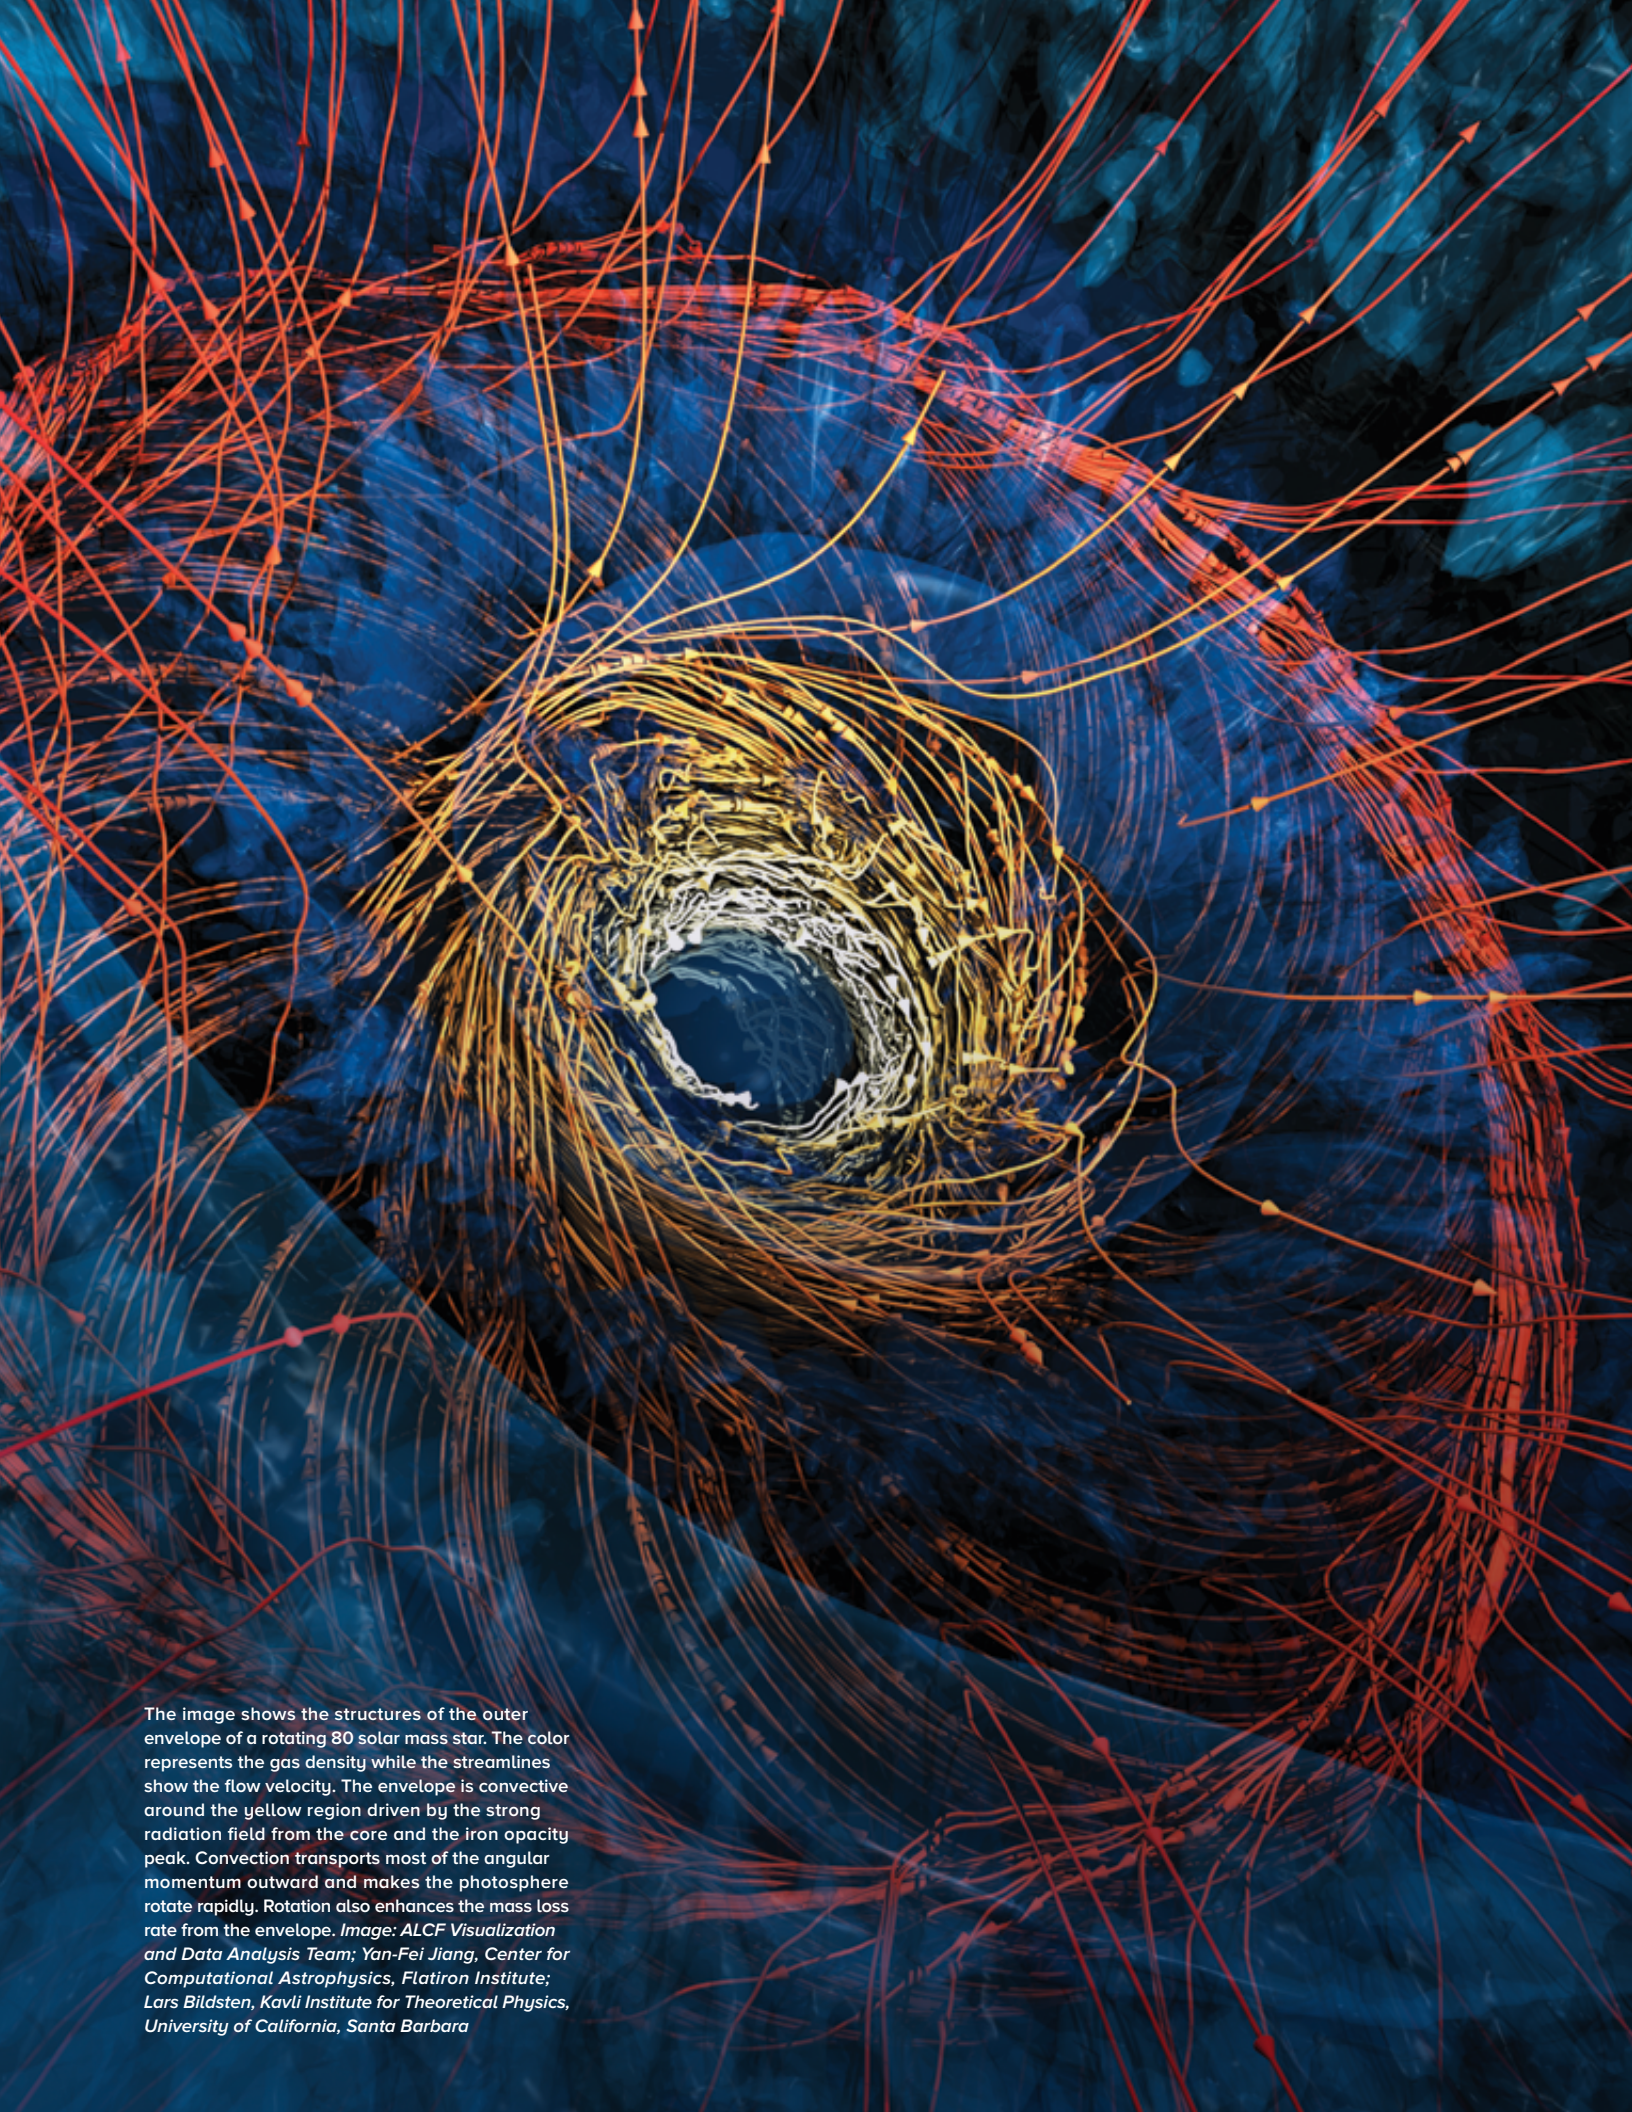

The image shows the structures of the outer envelope of a rotating 80 solar mass star. The color represents the gas density while the streamlines show the flow velocity. The envelope is convective around the yellow region driven by the strong radiation field from the core and the iron opacity peak. Convection transports most of the angular momentum outward and makes the photosphere rotate rapidly. Rotation also enhances the mass loss rate from the envelope. *Image: ALCF Visualization and Data Analysis Team; Yan-Fei Jiang, Center for Computational Astrophysics, Flatiron Institute; Lars Bildsten, Kavli Institute for Theoretical Physics, University of California, Santa Barbara*

Combating COVID-19 at the ALCF

ALCF systems and expertise are accelerating research aimed at developing treatments and strategies to combat the COVID-19 pandemic.

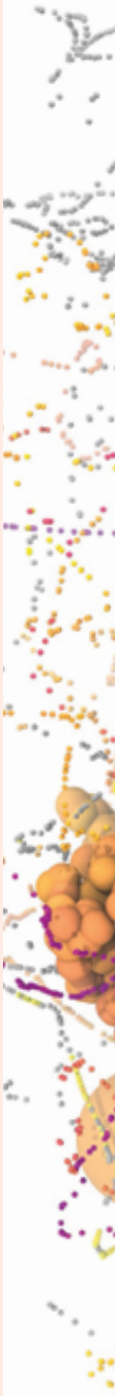

When the COVID-19 pandemic hit early in 2020, Argonne National Laboratory immediately joined the national effort to mitigate this global health problem, leveraging its world-class user facilities and multidisciplinary expertise to advance our understanding of SARS-CoV-2, the virus that causes COVID-19.

As part of the lab's response, the ALCF's powerful computing resources enabled a number of studies, including molecular modeling and epidemiology research, aimed at accelerating the development of treatments and strategies to combat the pandemic. Research campaigns have included performing AI-driven molecular dynamics simulations to identify potential COVID-19 targets and therapeutics; running highly detailed epidemiological models that simulate the spread of the virus throughout the population; and developing a computational pipeline to Argonne's Advanced Photon Source to enable on-demand analysis of experimental data on the crystal structure of COVID-19 proteins.

Much of this work was made possible by initiatives that helped ensure sufficient computational resources were available to support COVID-19 research. With funding from the Coronavirus Aid, Relief and Economic Security (CARES) Act, the ALCF was able to augment its Theta supercomputer with graphics processing units (GPUs) to provide additional computing power dedicated to COVID-19 studies. Deployed rapidly in response to the pandemic, the Theta upgrade offered new capabilities for training AI datasets, while also enabling GPU-specific and GPU-enhanced HPC applications for modeling and simulation.

Early in the pandemic, Argonne joined the COVID-19 High Performance Computing Consortium, a unique private-public effort, established by the White House Office of Science and Technology Policy, DOE, and IBM, that brought together government, industry, and academic leaders to provide access to world-class supercomputing resources in support of COVID-19 research. With more than 40 member institutions, the consortium aggregated an impressive pool of computing capabilities that have enabled COVID-19 researchers to carry out approximately 100 complex computational research projects.

In 2020, the ALCF supported a total of 10 COVID-19 research projects through both the consortium and Director's Discretionary allocations. ALCF staff put significant effort into collaborating with many of the project teams to ensure a quick start and good performance. The following project summaries highlight some of the notable COVID-19 studies carried out on ALCF computing resources.

Identifying COVID-19 targets and therapeutics

An Argonne-based team leveraged multiple supercomputing resources, including the ALCF's Theta system, to employ AI approaches that integrate information from experimental observations, and rigorous, physics-based virtual screening and molecular simulations, to identify viable drugs that can

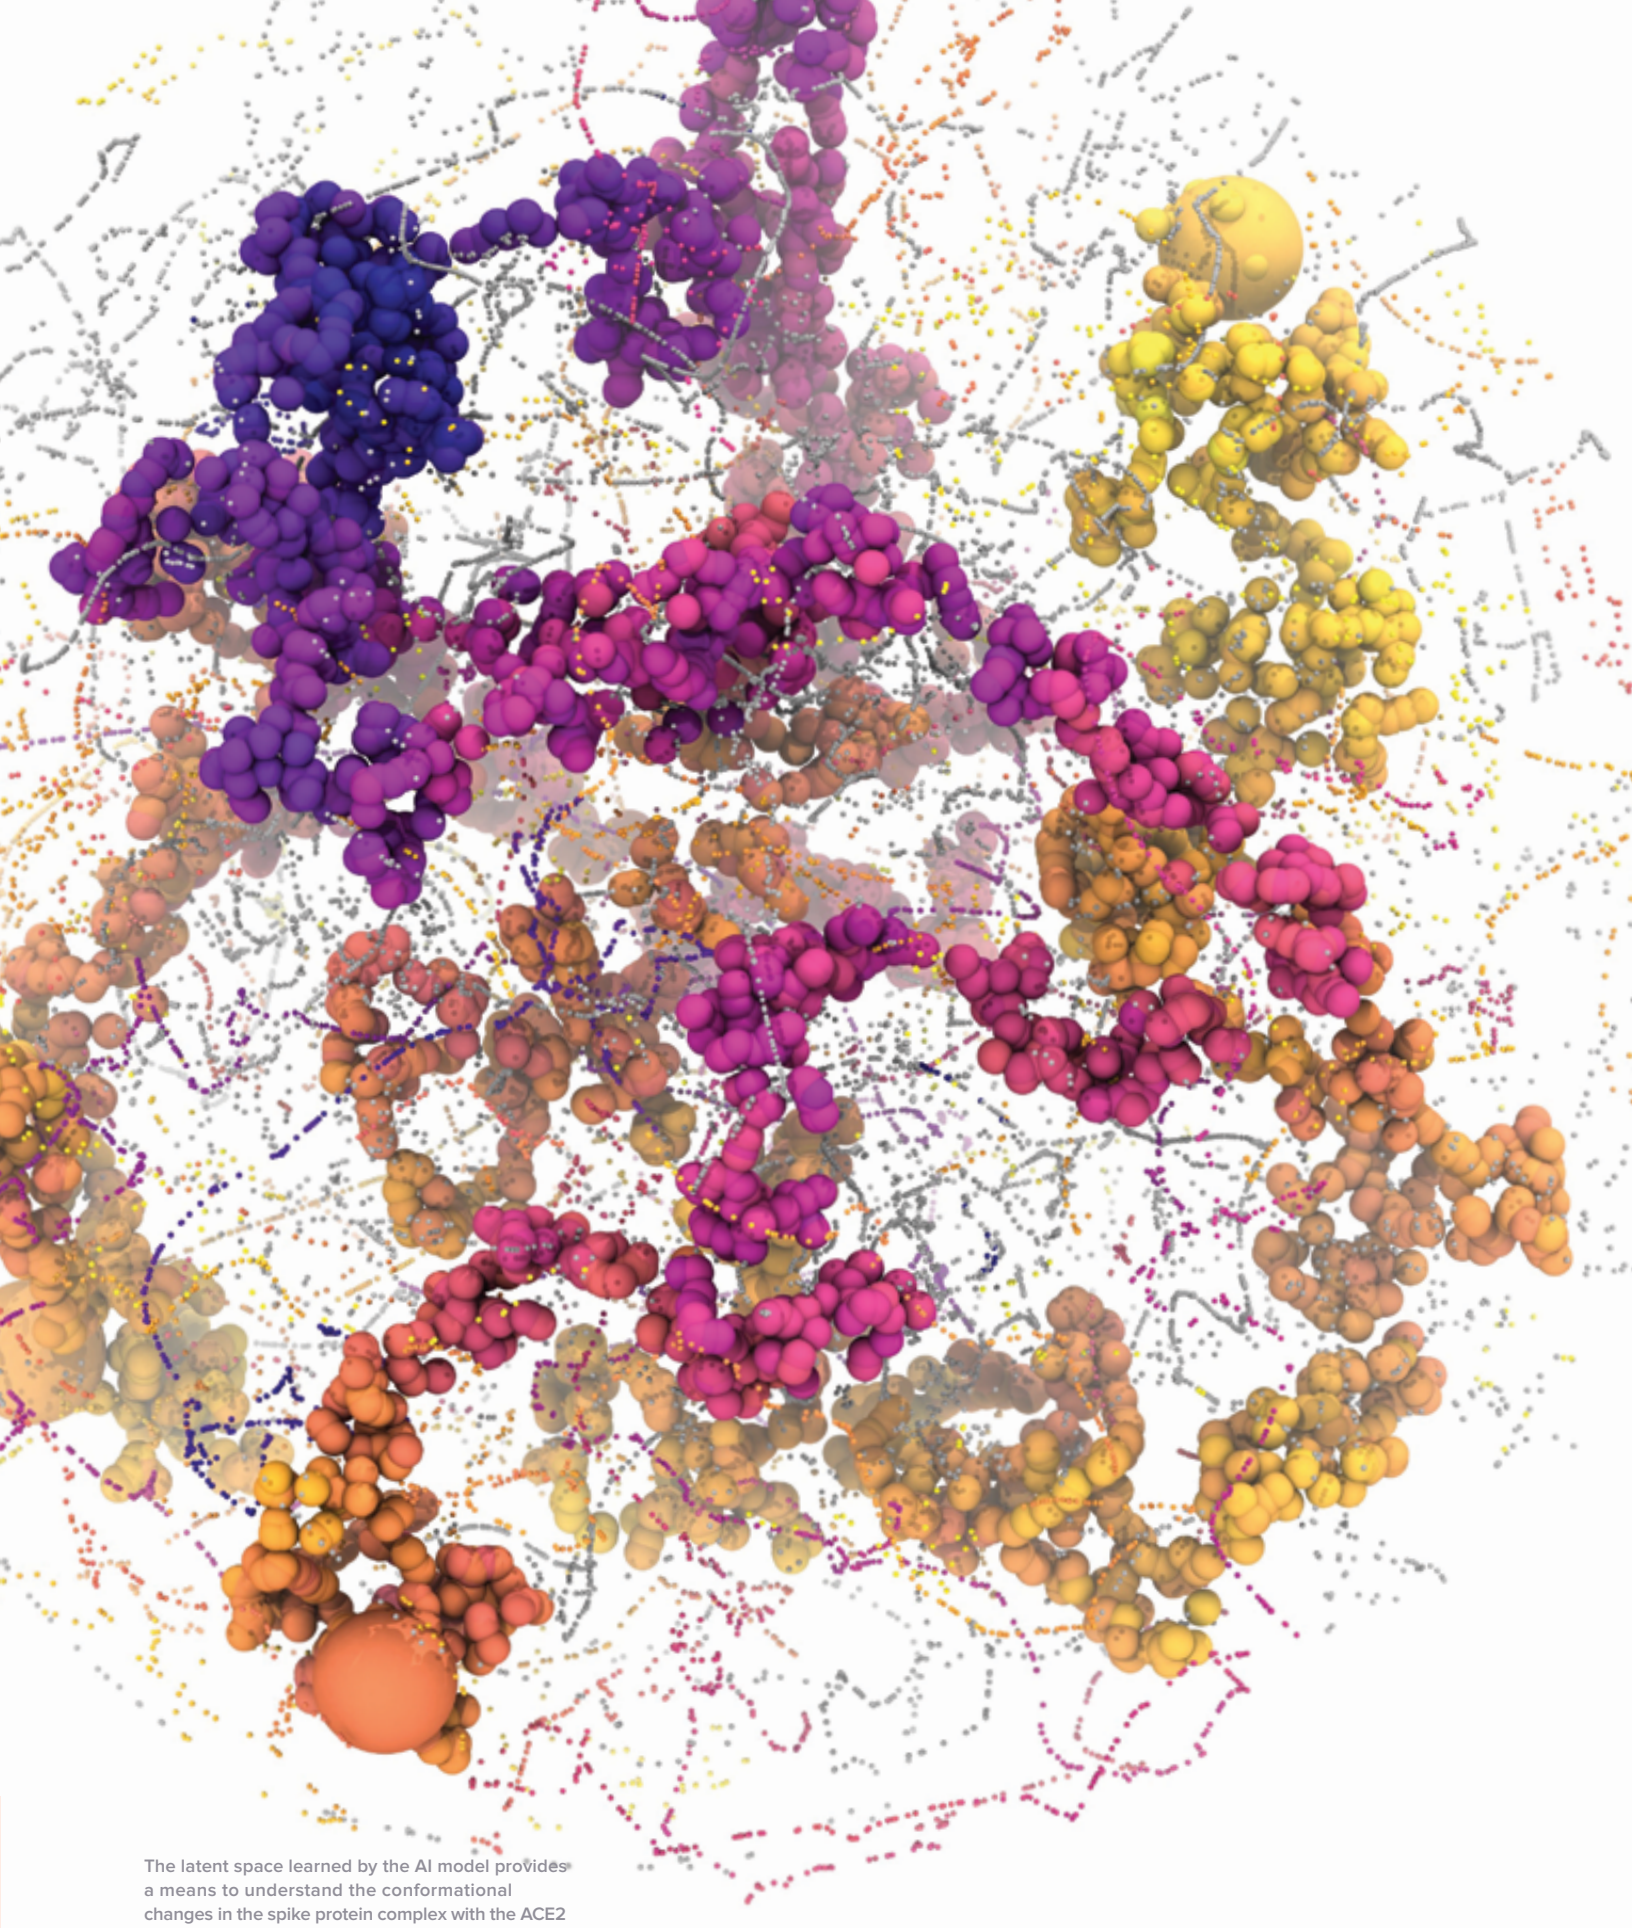

The latent space learned by the AI model provides a means to understand the conformational changes in the spike protein complex with the ACE2 receptor. *Image: Alex Brace and Arvind Ramanathan, Argonne National Laboratory; Anda Trifan, Argonne/University of Illinois at Urbana Champaign (UIUC); Lorenzo Casalino and Rommie Amaro, University of California San Diego; John Stone, UIUC*

Ryan Milner, ALCF HPC systems and network administration lead, installs NVIDIA DGX A100 nodes as part of the Theta upgrade.

inhibit viral proteins. These AI approaches, based on advances in deep learning and reinforcement learning, can predict how strongly a small molecule will bind to a protein, and explore the structural space of compounds that are predicted to bind to find more suitable variants. One of the team's studies, "AI-Driven Multiscale Simulations Illuminate Mechanisms of SARS-CoV-2 Spike Dynamics," was recognized with the first ACM Gordon Bell Special Prize for HPC-Based COVID-19 Research for their massive simulations carried out on OLCF's Summit system. This particular study also utilized many of top U.S. supercomputers, including systems at the ALCF, Texas Advanced Computing Center, San Diego Supercomputing Center, and Lawrence Livermore National Laboratory, to uncover alternate ways to handle the deluge of data.

Simulating the spread of COVID-19

To better understand the spread of COVID-19, Argonne scientists created CityCOVID, a detailed agent-based model that represents the behaviors, contacts, and disease progression of millions of residents in a large urban area. Using ALCF computing resources, the CityCOVID team simulated the movement of millions of people in the Chicago area as they went about their daily lives, predicting how COVID-19 would spread and the impact of preventative measures. Over the course of the pandemic, the team has been using CityCOVID to support city, county, and state stakeholders by providing them with model forecasts. The team's work was nominated as a finalist for the ACM Gordon Bell Special Prize for HPC-Based COVID-19 Research.

Analyzing COVID-19 experimental data in real time

Unraveling the structure of the SARS-CoV-2 virus is key to understanding the disease and potential therapeutics. Using high-intensity x-rays at Argonne's Advanced Photon Source (APS), scientists employed a technique called serial synchrotron crystallography to elucidate the crystal structure of COVID-19 proteins. To support the experimental work's rapid, on-the-fly data processing requirements, a team of Argonne researchers deployed an automated data acquisition, analysis, curation, and visualization pipeline, leveraging Theta for high-speed, on-demand analysis. The automated pipeline was able to move results to a repository and extract metadata for publication in a data portal, which scientists could monitor during experiments.

Identifying existing drugs that are effective against COVID-19

A team led by researchers at Johns Hopkins School of Medicine used a joint computational and experimental approach to identify FDA-approved drugs that possess antiviral activity against SARS-CoV-2. They used the ALCF's Theta system to screen a library of compounds that included existing FDA-approved drugs against ten intracellular catalytic SARS-CoV-2 protein targets for binding, with the goal of identifying drugs that disrupt viral protein function and diminish viral viability. Reducing the number of compounds and protein targets to an experimentally manageable number helps the team establish priorities for the subsequent workflow, in which target proteins will be expressed and purified, and activity assays will be developed. Finally, prioritized compounds in the Johns Hopkins Drug Library will be screened against the purified target proteins to identify those that exhibit antiviral activity.

Evaluating airline boarding changes in response to COVID-19

Using ALCF computing resources, a research team led by scientists from University of West Florida performed pedestrian dynamics simulations to analyze the spread of COVID-19 during air travel. Their work aimed to better understand the impact of new boarding processes on exposure to the virus. To do that, the team developed a computationally efficient constrained linear movement model and code and carried out a large parameter space sweep to simulate realistic boarding scenarios. The simulation

results show that back-to-front boarding roughly doubled the infection exposure compared with random boarding. It also increased exposure by around 50 percent compared to a typical boarding process prior to the outbreak of COVID-19.

Exploring social distancing measures worldwide via computer vision

Researchers from Purdue University built a website that pools together live footage and images from approximately 30,000 network cameras in more than 100 countries, providing a resource to study social distancing behaviors and evaluate the effectiveness of lockdowns and restrictions. Access to ALCF systems provided the team with sufficient computational power and storage to capture the vast amount of visual data and run object detection methods on a significantly large subset of data recorded from the cameras.

COVID-19 RESEARCH BY THE NUMBERS

Node-hours of compute time

1.7M

Research projects

10

Finalists for the Gordon Bell Special Prize for HPC-Based COVID-19 Research

1

Additional petaflops provided by Theta upgrade*

3.6

**Funded by CARES Act*

Preparing for Aurora

As the future home to the Aurora exascale system, the ALCF is driving efforts to prepare for science in the exascale era.

From mapping the human brain to accelerating the discovery of new materials, the power of exascale computing promises to advance the frontiers of some of the world's most ambitious scientific endeavors.

But applying the immense processing power of DOE's upcoming exascale systems to such problems is no trivial task. Researchers from across the high-performance computing (HPC) community are working to develop software tools, codes, and methods that can fully exploit the innovative accelerator-based supercomputers that will be deployed at DOE national laboratories.

With Aurora, an Intel-Hewlett Packard Enterprise (HPE) system, scheduled to arrive in 2022, the ALCF has been ramping up efforts to ready the system's future users for science in the exascale era. Argonne researchers are engaged in a broad range of preparatory activities, including exascale code development, hardware technology evaluations, user training, and collaborations with vendors, fellow national laboratories, and DOE's Exascale Computing Project (ECP).

Building a capable exascale ecosystem

The ECP and the ALCF's Aurora Early Science Program (ESP) have helped drive multi-institutional collaborations that are critical to realizing the full potential of exascale computing. Launched in 2016, the ECP is a multi-lab initiative aimed at accelerating the delivery of a capable exascale computing ecosystem that encompasses applications, system software, hardware technologies, architectures, and workforce development. Argonne—one of the six ECP core labs—has a strong presence on the ECP leadership team and has several researchers engaged in ECP projects and working groups.

The Aurora ESP is designed to prepare key applications for the scale and architecture of the exascale machine. Through open calls for proposals, the ESP has awarded pre-production computing time and resources to five simulation projects, five data projects, and five learning projects. The diverse set of projects is helping to create an environment that supports emerging data science and machine learning approaches alongside traditional modeling and simulation-based research. The ESP teams also field-test compilers and other software, helping to pave the way for other production applications to run on the system.

In the application development space, researchers participating in the ECP and the ALCF's Aurora Early Science Program (ESP) are working to port and optimize dozens of scientific computing applications for Aurora in advance of the system's arrival. With access to the early Aurora software development kit (a frequently updated version of the publicly available oneAPI toolkit) and Intel Gen9, DG1, and X^{e}HP GPUs through Argonne's Joint Laboratory for System Evaluation, ESP and ECP teams are able to test the performance and functionality of various codes, mini-apps,

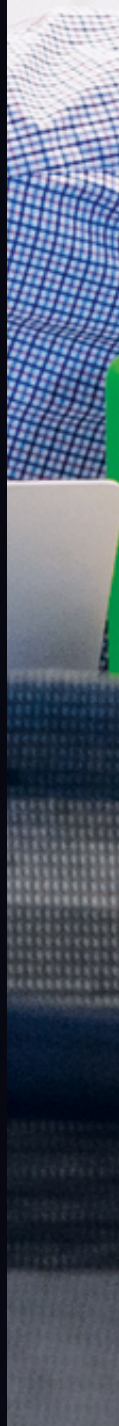

The Argonne-Intel Center of Excellence (COE) has held multiple workshops to provide details and instruction on various aspects of the Aurora hardware and software environment.

frameworks, and libraries using the programming models that will be supported on Aurora.

Working in concert with the ECP, Argonne researchers are also contributing to the advancement of programming models (OpenMP, SYCL, Kokkos, Raja), language standards (C++), and compilers (Clang/LLVM) that are critical to developing efficient and portable exascale applications.

Another area of collaboration involves deploying and integrating ECP software, vendor software, and facility-based software environments to ensure exascale software stacks meet application requirements while allowing for optimal operations at each facility. Activities include developing a common Continuous Integration strategy to drive automation of recurring building and testing across DOE's exascale site environments; using the Spack package manager as a tool for build automation and final deployment of software; managing the ECP allocation program that awards time at DOE computing facilities for the testing and development of new features and functionality at scale; exploring and potentially enhancing the installation, upgrade, and monitoring capabilities of the HPE Module Software Stack; and enabling container support and tools on Aurora and other exascale systems.

The ALCF team also continued its collaborative efforts with Aurora vendors, Intel and HPE, to test and develop various components to ensure they meet the requirements of the facility's user community. By analyzing the performance of key benchmarks and applications on early hardware, ALCF researchers are developing a broad understanding of the system's architecture and capabilities. This effort helps to identify best practices for optimizing codes, and, ultimately, a roadmap for future users to adapt and tune their software for the new system. In addition, the team's preparatory work

Aurora's innovative design will be based on Intel's Xeon Scalable processors, high-performance Intel X^{e} GPU compute accelerators, and Optane DC persistent memory. The system will rely on HPE Cray EX architecture and HPE Slingshot technology, which can provide concurrent support for advanced simulation and modeling, AI, and analytics workflows. Aurora will leverage historical advances in software investments along with increased application portability via Intel's oneAPI. The supercomputer will also be equipped with a new I/O platform called Distributed Asynchronous Object Storage (DAOS) to meet the needs of exascale workloads.

Sustained Performance

≥1 Exaflop DP

Platform

HPE Cray EX

Delivery

CY 2022

Aggregate System Memory

>10 PB

Compute Node
2 Intel Xeon scalable "Sapphire Rapids" processors; 6 X^{e} arch-based GPUs; Unified Memory Architecture; 8 fabric endpoints; RAMBO

GPU Architecture
X^{e} arch-based "Ponte Vecchio" GPU; Tile-based chiplets, HBM stack, Foveros 3D integration, 7nm

CPU-GPU Interconnect
CPU-GPU: PCIe
GPU-GPU: X^{e} Link

System Interconnect
HPE Slingshot 11; Dragonfly topology with adaptive routing

Network Switch
25.6 Tb/s per switch, from 64–200 Gbs ports (25 GB/s per direction)

High-Performance Storage
≥230 PB, ≥25 TB/s (DAOS)

Programming Models
Intel oneAPI, MPI, OpenMP, C/C++, Fortran, SYCL/DPC++

Node Performance
>130 TF

System Size
>9,000 nodes

includes providing feedback to Intel about the alpha/beta compilers and software stack for Aurora.

Training future exascale users

Training researchers and developers to use exascale tools and technologies has been another important thrust at the ALCF. In 2020, the facility hosted a number of workshops, webinars, and hackathons designed to help attendees prepare for Aurora.

The new Aurora Early Adopter webinar series was launched to introduce researchers to programming models, exascale technologies, and other tools available for testing and development work. Open to the public, the quarterly webinars have covered topics like Aurora's DAOS I/O platform, oneAPI's OpenMP offload capabilities, and the new oneAPI Math Kernel Library (oneMKL).

In 2020, the Argonne-Intel Center of Excellence (COE) hosted two Aurora workshops targeted at ESP and ECP research teams. Around 100 attendees visited the ALCF in February for a three-day workshop focused on applications and software development for Aurora. The event included substantial hands-on time for attendees to work with ALCF and Intel experts on developing, testing, and profiling their codes, as well as presentations on the Aurora software development kit, the system's memory model, and available machine learning tools and frameworks. In October, the COE held a virtual Aurora workshop for ESP and ECP teams, providing another opportunity for researchers to receive updates on the latest Aurora hardware and software developments, connect with ALCF and Intel staff for assistance in developing and testing their codes, and share their experiences working with various exascale tools and programming models.

The COE also hosted five virtual hackathons with individual ESP teams in 2020. These multi-day collaborative events paired teams with Argonne and Intel researchers to help advance their application development efforts using the Aurora software development kit. In December, the COE held its first dungeon programming event with four ESP simulation projects. The intensive three-day event focused on developing, porting, and profiling codes on pre-Aurora GPUs.

Together, these activities and collaborations are preparing the scientific community to harness exascale computing power to drive a new era of scientific discoveries and technological innovations.

Aurora Hackathons

06/22–06/23
High-Fidelity Simulation of Fusion Reactor Boundary Plasmas

08/17–08/18
Exascale Computational Catalysis

09/17–09/18
Metascalable Layered Materials Genome

09/25 AND 09/28
Many-Body Perturbation Theory Meets Machine Learning to Discover Singlet Fission Materials

12/07
Simulating and Learning in the ATLAS Detector at the Exascale

12/08–12/10
Aurora COE Dungeon Session 1: Intensive Development on Pre-Aurora GPUs (Four Simulation Applications)

Aurora Webinars

03/25
DAOS: Next-Generation Data Management for Exascale

06/02
Preparing Applications for Aurora Using the Intel DPC++ Compatibility Tool

06/24
OpenMP Offload Capabilities in the oneAPI HPC Toolkit

09/30
Overview of the New Intel oneAPI Math Kernel Library

Aurora Workshops

02/25–02/27
Aurora COE Workshop 2: Aurora Software Development Kit and Hardware Overview

10/21–10/22
Aurora COE Workshop 3: Intel Hardware Update and Pre-Aurora GPU Development

Additional Events

07/15–07/17
ALCF-ECP CMake Workshop

09/01–09/02
2020 Performance, Portability, and Productivity at HPC Forum

10/13–10/16
Intel eXtreme Performance Users Group (IXPUG) Annual Meeting

Aurora Early Science Program Projects

Simulation Projects

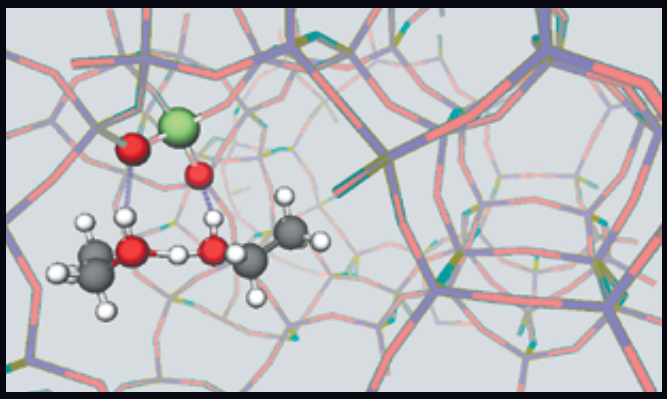

NWChemEx will provide the understanding needed to control molecular processes underlying the production of biomass. *Image: Thom H. Dunning Jr., University of Washington and Pacific Northwest National Laboratory*

Extending Moore's Law Computing with Quantum Monte Carlo

PI Anouar Benali
INST Argonne National Laboratory

High-Fidelity Simulation of Fusion Reactor Boundary Plasmas

PI C.S. Chang
INST Princeton Plasma Physics Laboratory

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

PI Theresa Windus
INST Iowa State University and Ames Laboratory

Extreme-Scale Cosmological Hydrodynamics

PI Katrin Heitmann
INST Argonne National Laboratory

Extreme-Scale Unstructured Adaptive CFD

PI Ken Jansen
INST University of Colorado Boulder

Data Projects

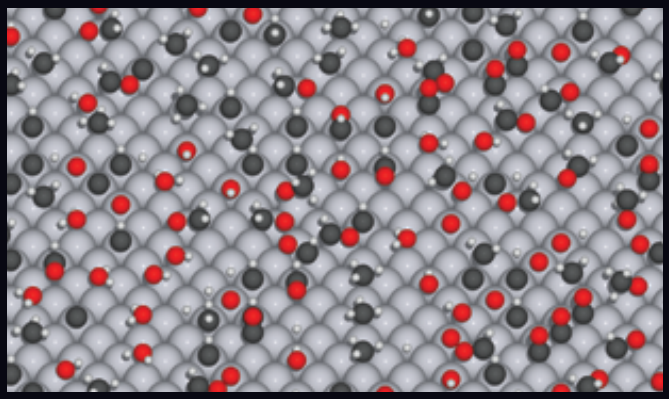

The catalysis project will combine data science techniques and quantum chemistry simulations to explore the otherwise intractable phase space resulting from gas phase molecules on catalyst surfaces to find relevant configurations and the lowest transition states between them. *Image: Eric Hermes, Sandia National Laboratories*

Exascale Computational Catalysis

PI David Bross
INST Argonne National Laboratory

Dark Sky Mining

PI Salman Habib
INST Argonne National Laboratory

Data Analytics and Machine Learning for Exascale Computational Fluid Dynamics

PI Ken Jansen
INST University of Colorado Boulder

Simulating and Learning in the ATLAS Detector at the Exascale

PI Walter Hopkins
INST Argonne National Laboratory

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

PI Amanda Randles
INST Duke University and Oak Ridge National Laboratory

Learning Projects

The PPPL team's Fusion Recurrent Neural Network uses convolutional and recurrent neural network components to integrate both spatial and temporal information for predicting disruptions in tokamak plasmas. *Image: Julian Kates-Harbeck, Harvard University; Eliot Feibush, Princeton Plasma Physics Laboratory*

Machine Learning for Lattice Quantum Chromodynamics

PI William Detmold
INST Massachusetts Institute of Technology

Enabling Connectomics at Exascale to Facilitate Discoveries in Neuroscience

PI Nicola Ferrier
INST Argonne National Laboratory

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

PI Noa Marom
INST Carnegie Mellon University

Virtual Drug Response Prediction

PI Rick Stevens
INST Argonne National Laboratory

Accelerated Deep Learning Discovery in Fusion Energy Science

PI William Tang
INST Princeton Plasma Physics Laboratory

Deploying Leading-Edge Hardware and Software

The evolution of the hardware and software deployed by the ALCF operates on a continuum as the progressively more advanced technology made available moves into the exascale era and beyond.

Recent milestones for the ALCF include innovations resulting from AI pathfinding experiments; the GPU expansion of Theta's architecture; the deployment of the Eagle and Grand filesystems; the deployment of the Cray EX testbed; and, cutting across all platforms, infrastructure for continuous integration (CI) and deployment (CD). Soon to join and further interconnect this array at its nexus, following their impending deployment, are the mutually adjacent Polaris and Aurora systems.

The deployment of leading-edge hardware is aided by extensive CI/CD infrastructure. GitLab instances allow ALCF users to run CI testing on HPC hardware prior to deployment. In this environment, builds either occur in accordance with a schedule or by trigger, and are often subsequent to CI jobs sent from JLSE testbeds. This infrastructure, through constant debugging and regular compiling, enables vital improvements to code as it is ported from machine to machine, ensuring performance and scalability on exascale computers as soon as they are available. Not only do application benefit, but users can test, debug, and deploy new features and functionality altogether.

Adding 24 NVIDIA DGX-A100 GPU systems (192 A100 GPUs) to what was previously an entirely CPU-based compute, the Theta upgrade resulting from the U.S. CARES Act greatly expanded the machine's capabilities and helped drive urgent COVID-19 research throughout the summer before being opened to general production use.

Eagle and Grand, meanwhile, beyond transforming standard user experience in the immediate term, helped prepare the facility for the arrival of exascale machinery by establishing 200 petabytes of storage—100 petabytes of which are available for sharing—as part of the baseline for ALCF science.

As the next GPU-powered system set to arrive at the ALCF, four of the racks approved for installation in Polaris are aimed at the continued development of our understanding of how to integrate with other experimental facilities. In particular, they will help evolve real-time processing of data from the Center for Nanoscale Materials, Advanced Photon Source, and Argonne's Atlas while advancing and deepening the ways by which the facilities can function in tandem with and benefit each other.

The transition to exascale and Aurora's arrival is further supported by the Cray EX testbed, now known as Crux. The Crux machine bears infrastructural similarities to Aurora's anticipated design, including analogous plumbing and electrical systems and, moreover, is capable of running the same software stack that Aurora will run. These correspondences are invaluable for preparing the Operations team in the management and maintenance of the Aurora machine prior to its delivery.

Future architectures—not only CPU- and GPU-based architectures, but those based on more exotic accelerators as well—are considered and investigated via the facility's AI

pathfinding activities. Using the systems acquired as part of the pathfinding effort, ALCF will construct a testbed to conduct experiments to investigate the applicability to science and scalability for even larger deployments. While from a broader standpoint they pave the pathway to a future of extreme compute heterogeneity and multifaceted architectures at the convergence of CPUs, GPUs, AI, and other accelerators, the testbed's capabilities will provide the ALCF with the immediate opportunity to incorporate methods and techniques new to the facility, such as text-based processing for state-of-the-art journal-mining, thereby expanding the range of science possible at the facility.

Furthermore, these technologies point toward what the future will bring, as once-distant concepts and processes like self-driving laboratories and autonomous discovery begin to enter the conversation and become realities. by masking some of the more overt complexity of the backend, such hardware helps grow the scientific computing community, in turn reinforcing the ALCF's crucial role in the research landscape.

What comes next—after GPUs—for high-performance scientific computing?
It was close to 15 years ago that researchers began to first seriously investigate the potential uses of GPUs for high-performance computing; in so doing, experiments with the ALCF AI testbed laid the groundwork for architectures being deployed in the Polaris and Aurora systems.

The testbed provides users with a platform with which to evaluate the usability and performance of machine-learning-based HPC applications, allowing benchmarks for applications, programming models, and machine learning frameworks; support for science application teams; and coordination with vendors during product development.

The testbed currently is being used to evaluate a deep learning accelerator, reconfigurable dataflow units, and intelligent processing unit- (IPU) based systems. This ongoing work is guiding the facility toward a future marked by extreme heterogeneity in the compute: CPUs, GPUs, AI, and other accelerators.

Introducing Polaris

Developed in collaboration with HPE, the ALCF's new GPU-driven supercomputer will serve as a bridge to the exascale era.

In 2021, the ALCF will deploy a new supercomputer named Polaris, the facility's largest GPU-based system to date.

The HPE machine, equipped with AMD processors and NVIDIA GPUs, gives the ALCF and its user community a testbed to help with the transition toward GPU-accelerated systems, including its upcoming exascale supercomputer, Aurora, while providing another powerful platform for enabling scientific breakthroughs.

Beyond the ALCF, Polaris will also be a valuable resource for the Exascale Computing Project (ECP) and other DOE computing facilities as they prepare for the next generation of leadership-class supercomputers.

Polaris will be powered by AMD Rome high-performance server microprocessors and NVIDIA A100 GPUs, providing a peak performance of 40 petaflops. Like Aurora, the system will be based on HPE's Cray EX architecture and Slingshot interconnect, which is designed to support the simulation, data, and learning workloads that will drive science in the exascale era.

The Polaris software environment provides an excellent platform for preparing codes to run on DOE's upcoming exascale systems. It will come equipped with the HPE Cray programming environment, HPE Performance Cluster Manager (HPCM) system software, and the ability to test all of the programming models that will be available on Aurora and the next-generation supercomputers at OLCF and NERSC. NVIDIA's HPC software development kit offers advanced capabilities for diverse applications including artificial intelligence and data analysis campaigns.

After the system is installed and accepted, it will initially be dedicated to research teams participating in the Aurora Early Science Program, DOE's Exascale Computing Project, and the ALCF Data Science Program. In 2022, Polaris will be made available to the broader research community for a wide range of science and engineering projects.

With many similarities at the system and user level, Polaris will be a key resource for researchers preparing to use the ALCF's Aurora exascale supercomputer.

COMPONENT	POLARIS	AURORA
System Software	HPCM	HPCM
Programming Models	MPI, OpenMP, DPC++, Kokkos, RAJA, CUDA, OpenACC, HIP	MPI, OpenMP, DPC++, Kokkos, RAJA, HIP
Tools	CrayPat, gdb, Cray ATP, NVIDIA Nsight, cuda-gdb	CrayPat, gdb, Cray ATP, Intel VTune
MPI	CrayMPI, MPICH	CrayMPI, MPICH, Intel MPI
Multi-GPU	1 CPU : 4 GPU	2 CPU : 6 GPU
Data and Learning	DL frameworks, Cray AI stack, Python/Numba, Spark, Containers, Rapids	DL frameworks, Cray AI stack, Python/Numba, Spark, Containers, OneDAL
Math Libraries	cu* from CUDA	oneAPI

Shared components are shown in white.

GROWING THE HPC COMMUNITY

As a leader in the HPC community, the ALCF is actively involved in efforts to broaden the impact of supercomputers and grow the body of researchers who can use them to advance science.

The ALCF's summer student program gives students an opportunity to tackle research projects that address issues at the forefront of scientific computing.

Partnering with Industry

The ALCF's industry partnerships help to strengthen the nation's innovation infrastructure and expand the use of supercomputers for technological and engineering advances.

Focused on growing the facility's community of industry users, the ALCF's Industry Partnerships Program engages with prospective companies of all sizes, from start-ups to Fortune 500 corporations, that could benefit from leadership computing resources and collaborative opportunities with the ALCF and across Argonne.

The ALCF's leadership computing resources—equipped with advanced simulation, data, and learning capabilities—enable companies to tackle problems that are too computationally demanding for traditional computing clusters.

Access to ALCF systems and expertise allows industry researchers to make predictions with greater accuracy, rapidly analyze massive datasets, and create higher-fidelity models of everything from manufacturing processes to fusion energy devices. The results permit companies to accelerate critical breakthroughs, verify uncertainties, and drastically reduce or eliminate the need to build multiple prototypes.

The ALCF has enhanced its industry outreach program by partnering with other Argonne user facilities and divisions, including the Technology Commercialization and Partnerships Division. By providing more complete picture of the laboratory's resources, this collaborative approach has resulted in broader engagements across Argonne with a number of companies.

The ALCF is also actively involved in directing the DOE Exascale Computing Project's Industry Council, an advisory group of senior executives from prominent U.S. companies interested in working with Argonne and other DOE laboratories to deploy exascale computing to improve their products and services.

DRIVING INNOVATION FOR INDUSTRY

The following project summaries illustrate how ALCF resources are helping companies to advance their R&D efforts.

Advancing Fusion Energy Research

Researchers from TAE Technologies are performing simulations on Theta to accelerate their experimental research program aimed at developing a clean, commercially viable, fusion-based electricity generator. The team will use the simulation results to optimize a device for studying the confinement of energy with high plasma temperatures, and to inform the design of a future prototype reactor.

Designing a More Efficient Aircraft Engine

Raytheon Technologies Research Center is working with Argonne to design a more efficient aircraft engine. They are leveraging HPC and machine learning to create accurate models that predict air flow and heat transfer inside a gas turbine engine, identifying fine-scale surface effects that could not have been discovered with physical experiments alone. These simulations will enable Raytheon to modify their manufacturing process to minimize heat loss and maximize durability of their engine components.

Optimizing Manufacturing Processes

Using machine learning and computational fluid dynamics, Argonne helped 3M optimize a fiber spinning manufacturing process used in the production of filters, fabrics and insulation. The collaborative effort, conducted through the HPC4Mfg program, helped minimize the amount of energy used in producing these materials, reducing the overall cost of production.

Reducing Emissions in Diesel Engines

Argonne worked with Caterpillar Inc. to improve efficiency and reduce emissions in their heavy-duty diesel engines. The team ran hundreds of high-fidelity combustion simulations to identify promising piston bowl designs—the combustion chambers in diesel engines—that could improve fuel efficiency while reducing harmful emissions.

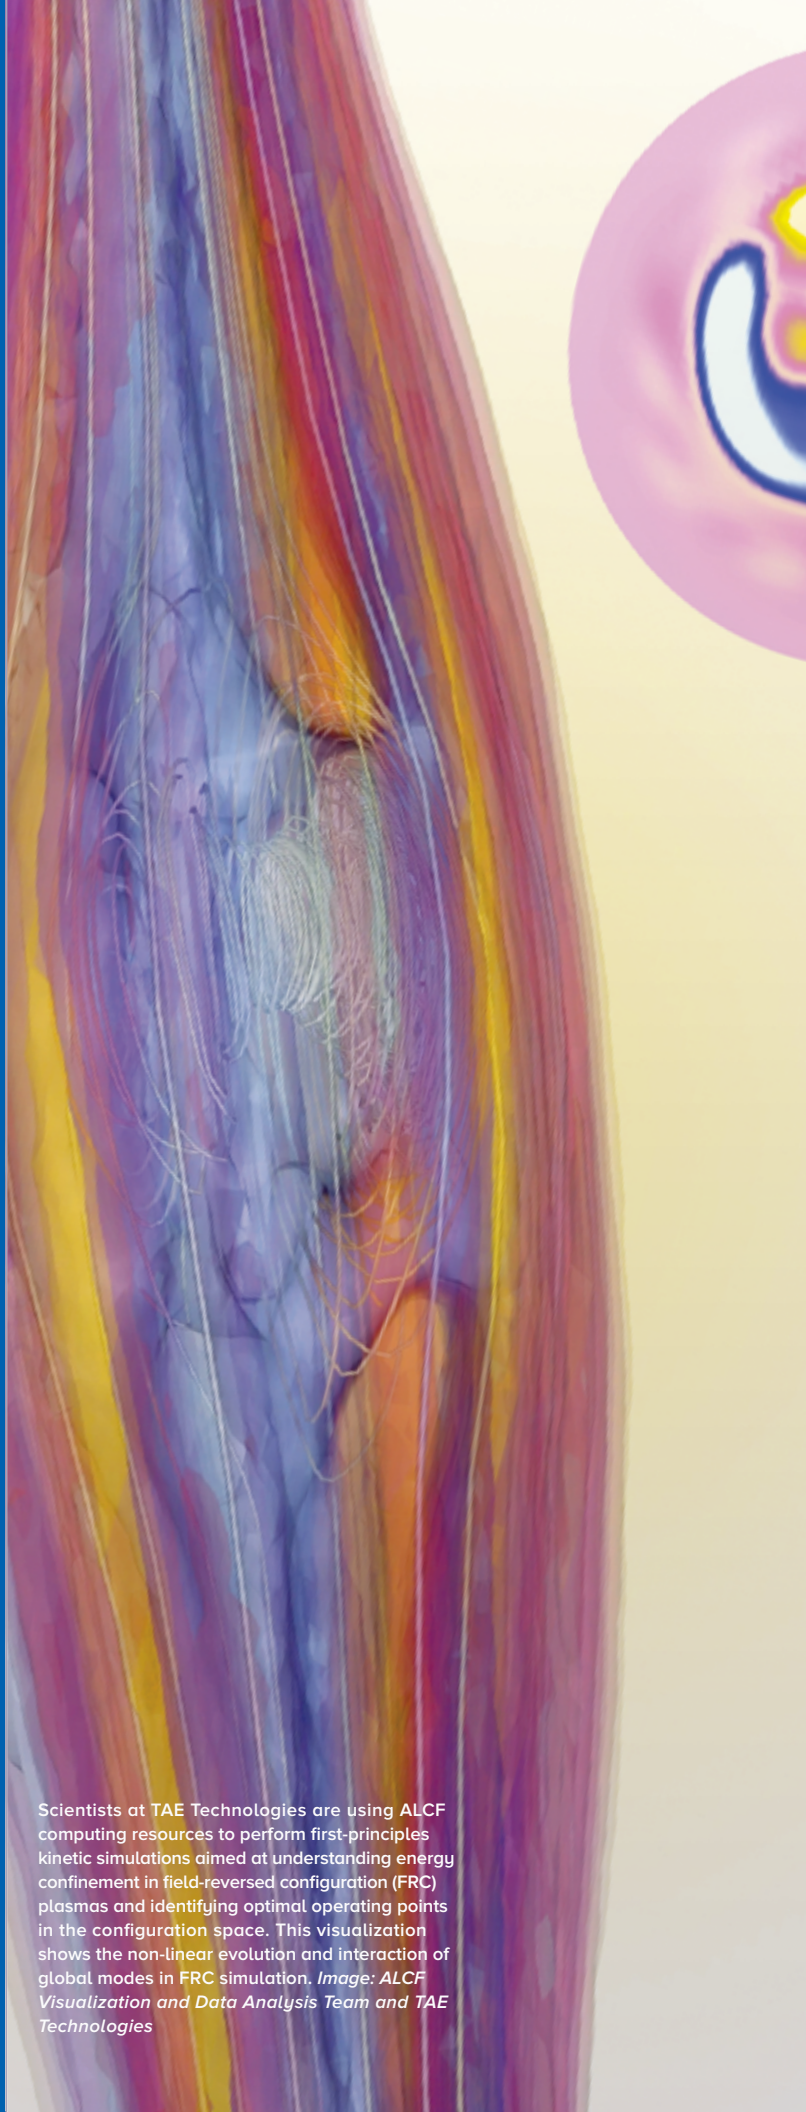

Scientists at TAE Technologies are using ALCF computing resources to perform first-principles kinetic simulations aimed at understanding energy confinement in field-reversed configuration (FRC) plasmas and identifying optimal operating points in the configuration space. This visualization shows the non-linear evolution and interaction of global modes in FRC simulation. *Image: ALCF Visualization and Data Analysis Team and TAE Technologies*

Shaping the Future of Supercomputing

ALCF researchers participate in numerous activities that aim to advance the forefront of scientific computing.

As a leadership computing facility, the ALCF is continually breaking new ground with the development and deployment of leading-edge HPC systems and capabilities.

ALCF staff members, collaborating with the researchers who use leadership-class systems to pursue scientific breakthroughs, are involved in the development and testing of new HPC hardware and software. This unique position the ALCF occupies affords the facility an important perspective on the trends, methods, and technologies that will define the future of supercomputing.

Leveraging this knowledge and expertise, ALCF researchers contribute to many forward-looking activities aimed at advancing the use of supercomputers for discovery and innovation.

These efforts include organizing workshops and meetings on topics like artificial intelligence and quantum computing; engagement in leading user groups and conferences; and contributions to the development of standards, benchmarks, and technologies that help propel continued improvements in supercomputing performance.

ALCF researchers had a strong presence at a series of “AI for Science” town hall meetings hosted by Argonne, Oak Ridge, and Berkeley National Laboratories in 2019.

Community Activities

Diversity and Inclusion

As part of the ALCF's commitment to building a diverse and inclusive work environment, the facility continues to invest in efforts to attract and recruit candidates from historically underrepresented groups in the scientific computing domain. Staff members are actively engaged in promoting career opportunities to women through contributions to Argonne's Women in Science and Technology program, AnitaB.org's Top Companies for Women Technologists program, and the Grace Hopper Celebration of Women in Computing. The ALCF also takes part in the annual Richard Tapia Celebration of Diversity in Computing Conference to recruit attendees from a diverse set of backgrounds and ethnicities.

Engaging with DOE Light Sources

The ALCF continues its efforts to accommodate the increasing computing needs of DOE light sources, such as Argonne's Advanced Photon Source (APS). This includes working with the DOE Light Sources and Computing Facilities Working Group to define future experimental computing requirements and how DOE leadership computing facilities can meet them. In addition, the ALCF participated in a series of Argonne town hall meetings with the APS and the laboratory's computing divisions to develop a common vision of the challenges and opportunities associated with the upcoming APS upgrade, and a roadmap detailing work that must be undertaken over the next decade and the near-term steps required to get started.

Exascale Computing Project

DOE's Exascale Computing Project (ECP) is a multi-lab initiative aimed at accelerating the delivery of a capable exascale computing ecosystem. Launched in 2016, the ECP's mission is to pave the way for the deployment of the nation's first exascale systems by building an ecosystem that encompasses applications, system software, hardware technologies, architectures, and workforce development. Researchers from the ALCF and across Argonne—one of the six ECP core labs—are helping the project deliver on its ambitious goals. The laboratory has a strong presence on the ECP leadership team and has several researchers engaged in ECP projects and working groups focused on application development, software development, and hardware technology. In the workforce development space, the ECP continues to fund the annual Argonne Training Program on Extreme-Scale Computing (ATPESC), which is organized and managed by ALCF staff.

ALCF staff members regularly attend the Grace Hopper Celebration of Women in Computing.

HPC Standards and Community Groups

ALCF staff members remain actively involved in a number of HPC standards and community groups that help drive continued improvements in supercomputing performance. Staff activities include contributions to the C++ Standards Committee, Cray User Group, HPC User Forum, Intel eXtreme Performance Users Group, Khronos OpenCL and SYCL Working Groups, MLPerf (HPC and Science Working Groups), MPI Forum, oneAPI Technical Advisory Board (DPC++ and oneMKL), OpenMP Architecture Review Board, OpenMP Language Committee, Open Scalable File Systems (OpenSFS) Board, and SPEC High-Performance Group.

InCREASE Workshop

In November, Argonne partnered with the Interdisciplinary Consortium for Research and Educational Access in Science and Engineering (InCREASE), a group dedicated to promoting research and education in minority-serving institutions, to host a virtual User Facilities Awareness Workshop. ALCF Deputy Director Jini Ramprakash was among the speakers at the event aimed at building new research collaborations with scientists from historically underrepresented groups.

Performance Portability

In 2020, ALCF staff members organized the annual Performance, Portability, and Productivity in HPC (P3HPC) Forum. Held virtually, the meeting brought together researchers and developers to discuss ideas and progress toward the goal of performance portability across current and future HPC platforms. The ALCF also continued its collaboration with NERSC and OLCF to operate and maintain a website dedicated to enabling performance portability across the DOE Office of Science HPC facilities (performanceportability.org). The website serves a documentation hub and guide for applications teams targeting systems at multiple computing facilities.

SC20

Continuing a long history of participation in the Supercomputing (SC) conference series, Argonne researchers had a strong presence at the virtual SC20 event in November. More than 100 Argonne staff members contributed to various conference activities, including technical paper presentations, invited talks, workshops, birds of a feather sessions, panel discussions, and tutorials, on topics ranging from exascale computing and big data analysis to AI and quantum computing.

Engaging Current and Future HPC Users

The ALCF provides HPC training opportunities for scientists and developers, while also participating in outreach activities to inspire a new generation of computing researchers.

The convergence of artificial intelligence (AI) and data science with traditional modeling and simulation is changing the way researchers use supercomputers for scientific discovery.

To help scientists find their footing in this ever-evolving landscape, the ALCF has placed a premium on training researchers to efficiently use the latest simulation, data analysis, AI, and machine learning tools and techniques on its supercomputing resources.

In 2020, the ALCF hosted a number of training events, including virtual workshops aimed at improving application performance, interactive webinars that connect attendees with the developers of HPC systems and software, and hackathons designed to help researchers prepare their codes for exascale.

The ALCF also supports a wide variety of outreach activities directed at students, with staff members volunteering to engage participants. Multi-day camps centered around programming and big data visualization, as well as external events like Hour of Code, spark students' interest in different aspects of scientific computing and introduce them to exciting career possibilities.

Additionally, the ALCF's annual summer student program gives college students the opportunity to work side-by-side with staff members on real-world research projects and utilize some of the world's most powerful supercomputers, collaborating in areas like computational science, system administration, and data science.

In 2020, the ALCF hosted multiple training events, including a hands-on workshop in February, to help researchers advance code development efforts for its upcoming exascale system, Aurora.

Training Users

ALCF Computational Performance Workshop

Held virtually in May, the annual ALCF Computational Performance Workshop is designed to help attendees boost application performance on ALCF systems. With dedicated access to ALCF computing resources, the three-day workshop allowed workshop participants to work directly with ALCF and invited experts to test, debug, and optimize their applications. One of the workshop's primary goals is to help researchers demonstrate code scalability for INCITE, ALCC, and ADSP project proposals, which are required to convey both scientific merit and computational readiness.

ALCF Developer Sessions

The facility continued to host its monthly webinar series, ALCF Developer Sessions, aimed at training researchers and increasing the dialogue between HPC users and the developers of leadership-class systems and software. Speakers in the series included developers from Intel, ARM, Globus, and Argonne, covering topics such as JupyterHub, Spack, automated research data workflows, and data visualization and analysis.

ALCF-ECP CMake Workshop

In July, the ALCF co-hosted a virtual workshop with Kitware and the Exascale Computing Project (ECP) to help exascale code developers advance their use of CMake on ALCF computing resources, including Aurora. CMake is a cross-platform, open-source build-system generator that is maintained and supported by Kitware. Attendees learned how to write a build-system generator capable of seamlessly configuring for multiple unique architectures with a variety of compilers.

ATPESC 2020

The annual Argonne Training Program on Extreme-Scale Computing (ATPESC) was held virtually for the first time in 2020. The two-week training event provides intensive, hands-on instruction on the key skills, approaches, and tools needed to carry out research on current supercomputers and future exascale systems. Organized by ALCF staff and funded by the ECP, ATPESC has a core curriculum that covers computer architectures; programming methodologies; data-intensive computing and I/O; numerical algorithms and mathematical software; performance and debugging tools; software productivity; data analysis and visualization; and machine learning and data science. More than 70 graduate students, postdocs, and career professionals in computational science and engineering attended this year's program. ATPESC has now hosted more than 500 participants since it began in 2013. To further extend the program's reach, ATPESC lecture videos are made publicly available on YouTube each year.

Aurora Hackathons

The Argonne-Intel Center of Excellence (COE) continued to host a series of hackathons to help Aurora Early Science Program (ESP) teams port and optimize their applications for the forthcoming exascale system. The intensive, hands-on sessions pair individual ESP teams with experts from the ALCF and Intel to advance code development efforts using the Aurora software development kit, early hardware, and other exascale programming tools. In December, the COE hosted its first Dungeon Session with four ESP simulation projects. Similar to hackathons, the programming event spanned three days and focused on developing, porting, and profiling codes on pre-Aurora GPUs.

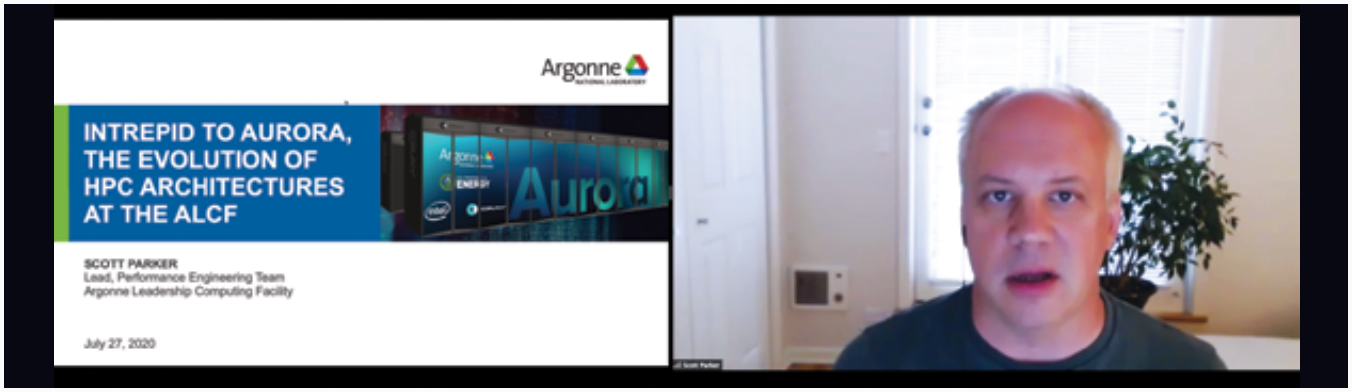

Scott Parker, ALCF performance engineering lead, provides an overview of the facility's supercomputing resources at ATPESC 2020.

Aurora Early Adopter Webinars

The ALCF expanded its webinar program with the launch of the Aurora Early Adopter series. Open to the public, the quarterly webinars introduce researchers to programming models, exascale technologies, and other tools available for testing and development work. Topics included Aurora's DAOS I/O platform, oneAPI's OpenMP offload capabilities, and the new oneAPI Math Kernel Library (oneMKL).

Aurora Workshops

The Argonne-Intel Center of Excellence (COE) hosted two Aurora workshops this year aimed at ESP and ECP research teams. The February workshop, held at the ALCF, was a three-day event focused on the Aurora software development kit and system hardware. In October and November, the COE hosted a two-part virtual workshop to provide updates on hardware and guidance on pre-Aurora GPU development. The invitation-only events gave attendees an opportunity to learn about the latest Aurora hardware and software developments, collaborate with ALCF and Intel staff on application development, and share experiences working with various exascale tools and programming models.

Best Practices for HPC Software Developers

In 2020, the ALCF, OLCF, NERSC, and ECP continued their collaboration with the Interoperable Design of Extreme-Scale Application Software (IDEAS) project to deliver a series of webinars—Best Practices for HPC Software Developers—to help users of HPC systems carry out their software development more productively. Webinar topics included reproducible containers, testing and code review practices, using proxy applications as benchmarks, scalable precision tuning, and software design with performance portability.

Simulation, Data, and Learning Workshop

In October, the ALCF hosted its annual Simulation, Data, and Learning Workshop to help users improve the performance and productivity of simulation, data science, and machine learning applications on ALCF systems. The virtual workshop was structured entirely around hands-on tutorials with opportunities to interact throughout. Day one was designed to help participants get distributed deep learning code and data pipelines running on ALCF systems, day two covered using DeepHyper for hyperparameter optimization and how to profile and improve application performance, and day three was dedicated to getting the attendees' deep learning networks deployed at scale in a simulation.

Inspiring Students

Big Data Camp

Argonne hosted a group of local high school juniors and seniors for its annual Big Data Camp in July. The week-long virtual event, organized by Argonne's Educational Programs Office and led by Argonne computer scientists, taught the students techniques for probing and analyzing massive scientific datasets, including data visualization methods. The attendees worked with data from the Array of Things (AoT) project and its follow-on Software-Defined Sensor Network (SAGE) project, an Argonne–University of Chicago urban sensor project, to gain hands-on experience with data-driven research.

CodeGirls@Argonne Camp

The annual CodeGirls@Argonne Camp hosts sixth- and seventh-grade girls each summer for a three-day event dedicated to teaching them the fundamentals of coding. Taught by Argonne computing researchers and staff from the lab's Learning Center, the virtual camp gave students an opportunity to try out creative and computational thinking through activities that include programming robots. The camp also allowed participants to meet women scientists, who use code to solve problems, and take part in a virtual tour the ALCF's machine room and visualization lab.

Coding for Science Camp

In July, Argonne hosted its Coding for Science Camp for 30 high school freshmen and sophomores who were new to coding. The week-long camp, a joint initiative of Argonne's Educational Programs Office and the ALCF, promotes problem solving and teamwork skills through hands-on coding activities, such as coding with Python and programming a robot, and interactions with Argonne staff members working in HPC and visualization.

Jini Ramprakash, ALCF deputy director, leads a virtual facility tour for students participating in Argonne's Science Careers in Search of Women event.

Hour of Code

As part of the national Computer Science Education Week (CSEdWeek) in December, several ALCF staff members provided virtual talks and demos to Chicago area schools to spark interest in computer science. Working with students in classes from elementary to high school, the volunteers led a variety of activities designed to teach the basics of coding. CSEdWeek was established by Congress in 2009 to raise awareness about the need to elevate computer science education at all levels.

Introducing Girls to STEM Careers

Through participation Argonne's annual Introduce a Girl to Engineering Day (IGED) and Science Careers in Search of Women (SCSW) events, ALCF staff members have the opportunity to connect with young women and introduce them to potential career paths in science, technology, engineering, and mathematics (STEM). IGED pairs approximately 100 local eighth graders with Argonne engineers and scientists for a day of presentations and hands-on activities focused on STEM careers. SCSW hosts high school students for a day of STEM activities, including Q&A panel discussions with Argonne women scientists, facility tours, and a career fair.

Summer Student Program

Every summer, the ALCF opens its doors to a new class of student researchers who work alongside staff mentors to tackle research projects that address issues at the forefront of scientific computing. This year, the internship program went virtual. The facility hosted 33 students ranging from high school seniors to Ph.D. candidates. From enabling quantum computer simulations to applying machine learning to weather simulations, the interns had the opportunity to gain hands-on experience with some of the most advanced computing technologies in the world. The summer program culminated with a series of seminars in which the students presented their project results to the ALCF community.

EXPERTISE AND RESOURCES

The ALCF's unique combination of supercomputing resources and expertise enables researchers to accelerate the pace of scientific discovery and innovation.

ALCF staff members augment Theta with the installation of NVIDIA DGX A100 nodes to provide a dedicated resource for COVID-19 research in 2020.

ALCF Staff and Systems

The ALCF's exceptional staff and powerful supercomputing resources make the facility one of the world's premier centers for scientific computing.

The ALCF's Theta supercomputer, an Intel-Cray XC40 system, was the engine that drove scientific discoveries for facility users in 2020.

With funding from the Coronavirus Aid, Relief and Economic Security (CARES) Act, Theta was augmented with NVIDIA GPUs and an additional four petaflops of computing power to provide a dedicated resource for COVID-19 research. The new integrated Theta supercomputer, which provides enhanced capabilities for data analytics and AI training and learning, was later made available to the broader user community.

Cooley, the facility's visualization and analysis cluster, enables users to transform data into high-resolution images, videos, and animations, helping users to better analyze and understand the results from simulations and experiments.

The ALCF provides high-performance data storage and networking capabilities to help users manage, store, and transfer large-scale datasets generated on the facility's computing resources.

Additionally, Argonne's Joint Laboratory for System Evaluation (JLSE) maintains a range of leading-edge hardware and software environments to enable researchers to evaluate and assess next-generation platforms and technologies.

Expertise and Support

The ALCF has assembled a world-class team of HPC system and network administrators, computational scientists, computer scientists, data scientists, performance engineers, visualization experts, software developers, and support staff, to ensure facility users are able to get the most out of its supercomputers.

Theta is the ALCF's Intel-Cray XC40 supercomputer.

ALCF Systems

Theta

KNL nodes

- Intel-Cray XC40 architecture
- 11.69 petaflops
- 64-core, 1.3-GHz Intel Xeon Phi 7230 processor per node
- 4,392 nodes
- 281,088 cores
- 843 TB of memory
- 70 TB of high-bandwidth memory
- Aries interconnect with Dragonfly configuration
- 24 racks

GPU nodes

- NVIDIA DGX A100 architecture
- 3.9 petaflops
- AMD EPYC 7742
- 24 nodes
- 24 TB of DDR4 memory
- 7, 680 GB of GPU memory
- 7 racks

Iota

Test and development platform

- Intel-Cray XC40 architecture
- 117 teraflops
- 64-core, 1.3-GHz Intel Xeon Phi 7230 processor per node
- 44 nodes
- 2,816 cores
- 12.3 TB of memory
- 1 TB of high-bandwidth memory
- Aries interconnect with Dragonfly configuration
- 1 rack

Cooley

Data analysis and visualization cluster

- Intel Haswell architecture
- 293 teraflops
- Two 6-core, 2.4-GHz Intel E5–2620 processors per node
- 1 NVIDIA Tesla K80 GPU per node
- 126 nodes
- 1,512 cores
- 47 TB of memory
- 3 TB of GPU memory
- FDR InfiniBand interconnect
- 6 racks

DATA STORAGE SYSTEMS

Eagle

Storage system for data sharing

- HPE ClusterStor E1000
- 100 petabytes of usable capacity
- 8,480 disk drives
- Lustre filesystem
  - 160 Object Storage Targets
  - 40 Metadata Targets
- HDR Infiniband network
- Sustained 650 GB/s rate on data transfers

Grand

Storage system for project science campaigns

- HPE ClusterStor E1000
- 100 petabytes of usable capacity
- 8,480 disk drives
- Lustre filesystem
  - 160 Object Storage Targets
  - 40 Metadata Targets
- HDR Infiniband network
- Sustained 650 GB/s rate on data transfers

Data Storage

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

DISK STORAGE

The Theta storage system consists of 30 I/O nodes that connect to a storage array that controls 2,300 disk drives with a total useable capacity of 9 PB and a maximum aggregate transfer speed of 240 GB/s. Theta uses Lustre to access this storage. The ALCF also utilizes a 10 PB filesystem based on an IBM Elastic Storage Server (ESS) to host data for science running on the Theta and Cooley systems. The ESS system is a software defined storage system based on IBM's GPFS filesystem and consists of 60 I/O nodes controlling 7,260 disk drives.

The ALCF supports two 100 PB globally accessible Lustre filesystems named Grand and Eagle. Each storage array controls 8,480 disk drives with a sustained transfer speed of 650 GB/s in the current environment. Eagle storage allocations will be granted to projects focused primarily on data sharing and will be accessible by non-ALCF users via Globus. Grand storage will be provided with standard compute allocations and will only be accessible to ALCF users who are project members.

A new home filesystem, homefs, will reside on an all-flash storage array with a capacity of 220 TB and a transfer rate up to 24 GB/s. A new filesystem named Swift will reside on an all flash storage array with a capacity of 123 TB and transfer rate up to 48 GB/s. It is intended to be targeted by GPU-based workflows.

TAPE STORAGE

The ALCF has three 10,000-slot libraries. The tape technology is currently undergoing an upgrade to replace LTO-6 tape drives with LTO-8 tape drives. The upgrade should ultimately provide up to 300 PB of effective storage (approximately five times the amount provided by the LTO-6 tapes).

Networking

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

The ALCF connects to other research institutions using up to 100 Gb/s of network connectivity. Scientists can transfer datasets to and from other institutions over fast research networks, such as ESnet and Internet2.

Testbeds

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

Arcticus, Yarrow: Intel discrete GPU testbeds for ECP and ESP projects to develop, optimize, and scale applications and software for Aurora

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

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

HPE Comanche Prototype: ARM64 platform for exploring the Marvell Arm architecture

Intel Cooper Lake: Intel Xeon cluster for testing data types for AI and learning applications Cluster

Iris: Intel integrated Gen9 GPUs for ECP and ESP projects to develop, optimize, and scale applications and software for Aurora

NVIDIA GPUs: Clusters of P100, V100, A100 GPUs for preparing applications for heterogenous computing architectures

Presque: Intel DAOS nodes for testing the Aurora storage system

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 assitance 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 the impactful research enabled by ALCF resources to external audiences.

Staff News

ALCF staff members Yasaman Ghadar (left) and Christopher Knight participated in an Argonne promotional campaign for "Thank Your Mentor Day." Knight has served as a mentor to Ghadar since she joined the lab.

ALCF team recognized with DOE Secretary of Energy's Achievement Award
ALCF Director Michael Papka, ALCF Director of Science Katherine Riley, and the lab's High-Performance Computing Resource Team were recognized with a DOE Secretary of Energy's Achievement Award for their work in mobilizing supercomputing resources to accelerate research into treatments and strategies to combat the COVID-19 pandemic.

Emani receives DOE funding award to advance AI for science
ALCF computer scientist Murali Emani received a DOE award for a three-year project aimed at making AI models and data more accessible and reusable to accelerate AI research and development. Emani will co-lead "HPC-FAIR: A Framework Managing Data and AI Models for Analyzing and Optimizing Scientific Applications," one of five new projects that will apply Findable, Accessible, Interoperable, and Reusable (FAIR) Data Principles so that science data can drive innovations in AI.

Papka named NIU Presidential Research, Scholarship, and Artistry Professor
ALCF Director Michael Papka was named a 2020 Presidential Research, Scholarship, and Artistry Professor by Northern Illinois University. In addition to his leadership roles at Argonne, Papka serves as a professor of computer science at NIU and co-director of its Data, Devices and Interaction Laboratory (ddiLab). The award is the university's top recognition for outstanding research or artistry. Award winners receive special financial support of their research for four years, after which they carry the title of Distinguished Research Professor.

Rajak and Luo contribute to "Best Paper" at HPC Asia conference
ALCF researchers Pankaj Rajak and Ye Luo were part of a multi-institutional team that was recognized with the Best Paper Award at the HPC Asia 2020 conference. Co-written by researchers from the University of Southern California, Kumamoto University (Japan), and NSTDA Supercomputer Center (Thailand), the team's paper, "Quantum Dynamics at Scale: Ultrafast Control of Emergent Functional Materials," describes their efforts to prepare a quantum molecular dynamics engine to run on the ALCF's upcoming exascale system, Aurora.

Staff Spotlights

KEVIN BROWN
Postdoctoral Appointee

The term "leadership" does not just apply to the ALCF's leadership computing resources. It also applies to the people working to support facility users and maintain ALCF systems.

With a shared passion for research and innovation, ALCF staff members are helping to shape the future of supercomputing. The following pages highlight six staff members and some of their notable contributions in 2019.

Since joining the ALCF in October 2019, Kevin has been working to identify suitable network and workload configurations for Aurora and DOE's other upcoming exascale supercomputers. Kevin's work will help guide the deployment of these systems, aiming to reduce network interference and to improve user productivity. Specifically, Kevin has been modeling novel network architectures with the Co-Design of Exascale Storage Architectures (CODES) simulation toolkit to evaluate features of interest to the ALCF, such as network quality-of-service (QoS) classes, and to improve the accuracy and fidelity of network simulations. In progressing this work, Kevin has actively collaborated with engineers from HPE and researchers from the Rensselaer Polytechnic Institute, the University of Maryland, and the Illinois Institute of Technology.

In addition, Kevin served as co-chair of the 2020 Argonne Postdoctoral Research and Career Symposium and on the organizing committee of the 2020 ALCF Computational Performance Workshop, transitioning both from in-person activities to virtual events. Kevin also served on several international conference committees, including the Tutorials Committee at SC20, and mentored students through various outreach activities within and outside of Argonne.

BRIAN HOMERDING
Computer Scientist

JANET JASECKAS
Software Development Specialist

Brian first joined the ALCF as a research aide in the summer of 2017. During this time, he worked on improving the accessibility of DOE proxy applications through the Spack package manager and their use through the LLVM test suite. Following this, Brian joined ALCF full time to continue working with the Exascale Proxy Application Project within DOE's Exascale Computing Project (ECP). In this role, he conducted performance analyses on many proxy applications to provide a clear understanding of their computational characteristics. Additionally, Brian contributed to the ECP team's effort to investigate the representativeness of various Monte Carlo neutron transport proxy application with their full science applications. Throughout this time, he has also contributed to the LLVM community by giving several talks and helping to mentor several Google Summer of Code students.

In 2019, Brian joined the ALCF's Performance Engineering group and began to support the ECP's EQSim team, whose goal is to advance regional-scale ground motion simulation capabilities. Recently, Brian was a guest on the ECP's "Let's Talk Exascale" podcast to discuss the EQSim code development efforts for Aurora. This was due to Brian leading the development of a new SYCL backend for the RAJA programming model to provide support for RAJA applications on Aurora, including EQSim. The development of the RAJA SYCL backend has led to investigations of performance across programming models. Brian and his team published some of their findings at the 8th International Workshop on OpenCL (IWOCL) and SYCLcon 2020. In 2021, Brian will continue his efforts to develop and prepare both EQSim and RAJA for Aurora.

Janet joined the ALCF in 2014 as a member of the User Experience team. While in that role, Janet helped manage user accounts and system access. She worked to improve the shortcomings of the existing account and project management software by enhancing security and adding needed features for the User Experience team. In 2017, Janet transitioned to the role of principal software specialist in the ALCF's Advanced Integration Group (AIG). With the knowledge of processes she acquired during her time on the User Experience team, and her previous experience in building software applications in private industry, she was instrumental in rewriting ALCF's new account and project management software (Userbase3) from the ground up.

Since the launch of Userbase3 in April 2019, Janet has been busy integrating new features into the software for ALCF and accommodating other Argonne divisions that are interested in using the platform. She is responsible for supporting all aspects of the software development lifecycle, including gathering requirements, facilitating the conversations between the AIG and User Experience teams, and writing and testing software.

As a member of the Awards Committee for Argonne's Computing, Environment, and Life Sciences (CELS) directorate, Janet has helped with identify deserving members of the Argonne team for nominations to industry-wide awards. Janet also enjoys working with students as a volunteer for Argonne's annual Introduce a Girl to Engineering Day and Science Careers in Search of Women programs.

NEVIN LIBER
Computer Scientist

MISHA SALIM
Assistant Computational Scientist

Nevin's first and foremost responsibility is working on the SYCL/DPC++ backend for Kokkos for initial use on the Aurora exascale system. As Argonne's representative on the C++ Standards Committee, Nevin serves as the vice chair of the Library Evolution Working Group Incubator (LEWGI). This study is responsible for looking at proposals for libraries to be added to C++, in order to both determine if the group really wants them and then to get them in shape for the Library Evolution Working Group (LEWG) to evaluate. Nevin also writes and champions proposals for making C++ a better language for HPC.

In addition, Nevin is Argonne's representative on the SYCL Committee, where they just released SYCL 2020. Prior to the pandemic, he regularly gave C++ talks at various conferences, both locally and internationally. During the pandemic, Nevin continued to deliver talks virtually with presentations on SYCL for the ALCF's Aurora Early Adopter webinar series and at the ECP Annual Meeting. As you can tell, Nevin's passion and expertise is C++. He firmly believes we should give back to the communities we are a part of, and this job encourages him to give back to C++ while working on libraries and infrastructure that benefit HPC developers. Part of the missions for both SYCL and Kokkos is to take the team's abstractions and put them into C++ proper.

Misha joined the ALCF as a postdoc with the data science team in 2017. He designed the current generation of the Balsam workflow manager, which is used by ALCF projects spanning domains including materials science, high energy physics, structural biology, cosmology, and hyperparameter optimization for deep learning. As a workflow developer with the data science team, Misha routinely partners with researchers to design new workflows and run them at scale on ALCF systems. A few recent applications include high-throughput docking of small molecules with SARS-CoV-2 protein targets, real-time coupling of AI-driven molecular dynamics simulations on ThetaGPU with model training on the Cerebras-CS1 system, and scaling containerized simulations of neutrino events in liquid argon time-projection chambers (LArTPC) on Theta.

Misha also developed the parallel model evaluation interface of the DeepHyper framework and worked with Prasanna Balaprakash on a scalable implementation of asynchronous model-based search for hyperparameter optimization. Currently, Misha is focused on developing a Balsam web service that enables experimental scientists to tap into HPC resources and scale out real-time data analysis workflows with seamless cross-facility job submission. This work builds on a toolkit of portable components for data transfer, scheduling, and high-throughput job execution designed to work on current and future GPU-based systems.

EMILY SHEMON
ALCF Catalyst/Nuclear Engineer

Emily has held a joint appointment in Argonne's Nuclear Science and Engineering division as a nuclear engineer and at the ALCF as a catalyst since 2011. Her scientific research lies at the intersection of high-performance computing and computational methods development for nuclear reactor applications. She has been the catalyst, or point-of-contact, for several nuclear energy-related INCITE and ALCC projects throughout the years. She also led her own ALCF project for nuclear energy applications. Her most recent research involves the calculation of "hot channel factors" for fast spectrum nuclear reactors using high-fidelity neutronics and thermal hydraulics multiphysics simulation. Emily has led the Additional Protocol Reporting effort for ALCF whereby all nuclear energy-related work is reported to the International Atomic Energy Agency. Recently, she led the science and allocations sections of the 2020 Operational Assessment Review (OAR) report. She has hosted numerous tours of the ALCF for nuclear energy visitors and collaborators. Emily recently developed resources for staff in Argonne's Energy and Global Security (EGS) directorate to better engage with the ALCF, including improved communication about the facility's resources and preparing successful project proposals.

In addition, Emily has served as the multiphysics technical area deputy for DOE's Nuclear Energy Advanced Modeling and Simulation Program since 2019, and has led several computing projects within that program centered around the Multiphysics Object Oriented Simulation Environment (MOOSE) framework. Emily has also volunteered in Argonne outreach programs including Introduce a Girl to Engineering Day and Science Careers in Search of Women, and has been a postdoctoral mentor for several years.

ALCF STAFF NUMBERS

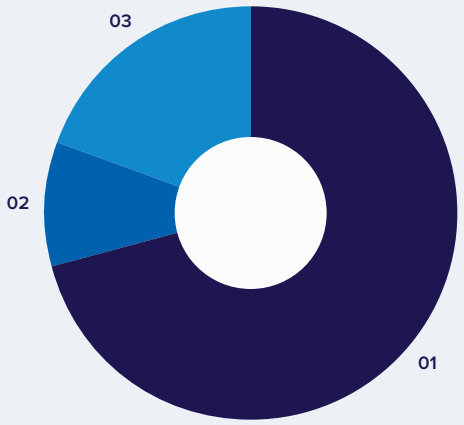

01 Staff Members

113

02 Postdoctoral Researchers

16

03 Summer Students

33

SCIENCE

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

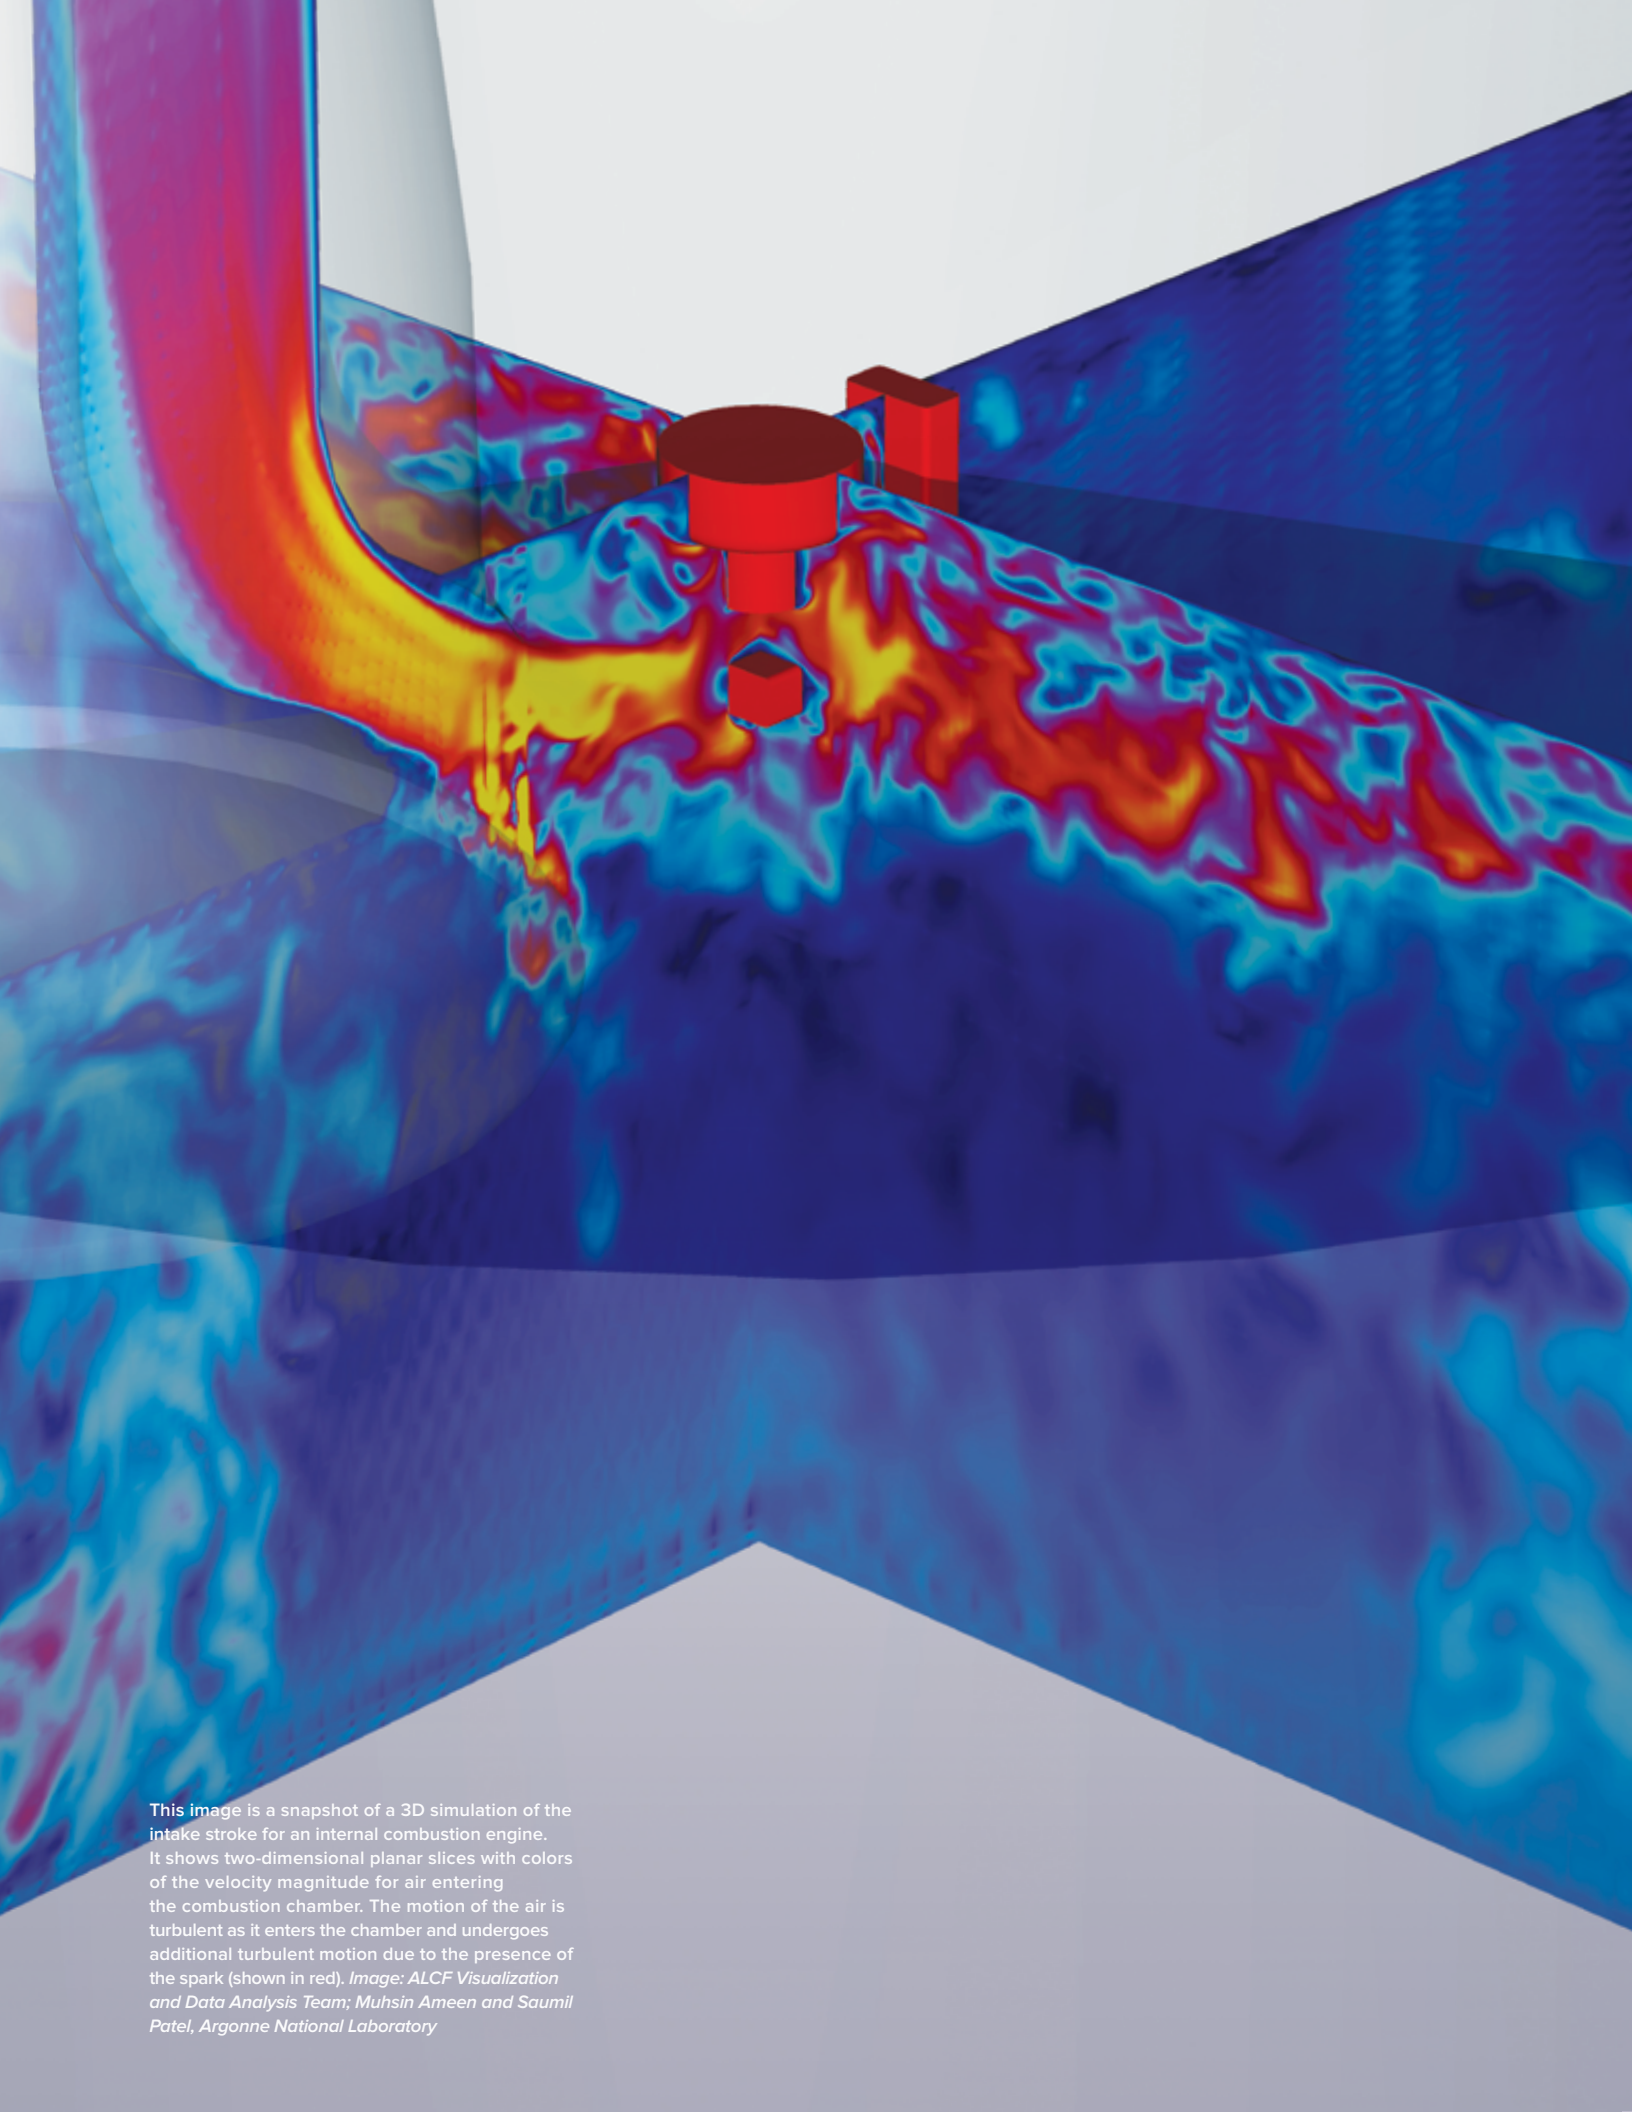

This image is a snapshot of a 3D simulation of the intake stroke for an internal combustion engine. It shows two-dimensional planar slices with colors of the velocity magnitude for air entering the combustion chamber. The motion of the air is turbulent as it enters the chamber and undergoes additional turbulent motion due to the presence of the spark (shown in red). *Image: ALCF Visualization and Data Analysis Team; Muhsin Ameen and Saumil Patel, Argonne National Laboratory*

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—typically with awards of millions of core-hours—through competitive, peer-reviewed allocations programs supported by the DOE and Argonne.

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

Application Programs

ADSP

The ALCF Data Science Program (ADSP) supports big data projects that require the scale and performance of leadership computing resources. ADSP projects focus on developing and improving data science techniques that will enable researchers to gain insights into very large datasets produced by experimental, simulation, or observational methods.

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.

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 application performance to maximize scientific application efficiency and productivity on leadership computing platforms.

INCITE/ALCC BY DOMAIN

2020 INCITE

17.8M NODE HOURS

A	Biological Sciences	– %
B	Chemistry	5
C	Computer Science	–
D	Earth Science	12
E	Energy Technologies	–
F	Engineering	10
G	Materials Science	28
H	Physics	45

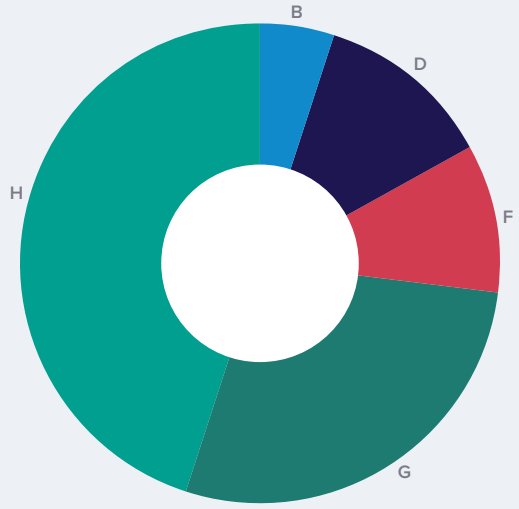

2020 ALCC

5.87M NODE HOURS

A	Biological Sciences	– %
B	Chemistry	13
C	Computer Science	1
D	Earth Science	3
E	Energy Technologies	15
F	Engineering	–
G	Materials Science	19
H	Physics	49

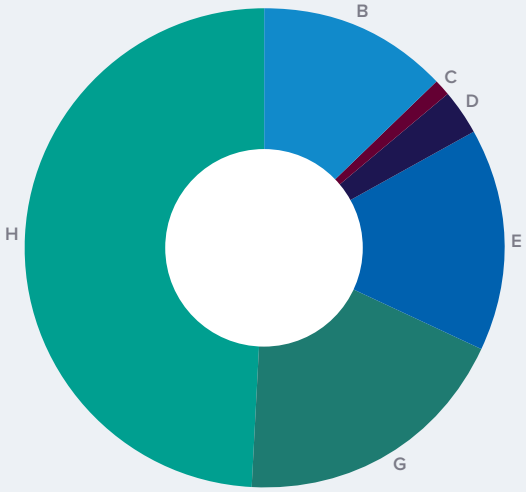

ALCC data are from calendar year 2020.

Biological Sciences | Simulation, Learning

AI-Driven Drug Discovery for SARS-CoV-2 Proteome

PI Arvind Ramanathan, Argonne National Laboratory
AWARD Director's Discretionary
HOURS Theta: 81,000 Node-Hours

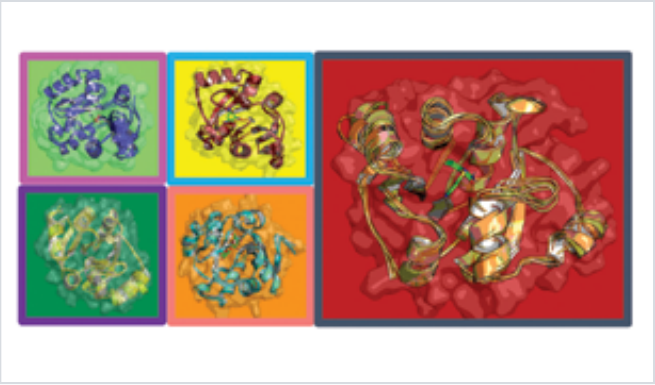

AI-driven MD simulations provide insights into how different ligands modulate the binding region of the viral ADP-ribose-1"-phosphatase protein. Ligands are shown in stick like representation and the protein is shown as a cartoon ensemble. Note that each ligand has an effect on distinct regions of the protein. *Image: Argonne National Laboratory*

This project seeks to address the fundamental biological mechanisms of the SARS-CoV-2 virus and associated COVID-19 disease, while simultaneously targeting the entire viral proteome to identify potential therapeutics.

CHALLENGE

The researchers leverage leadership-class machines to design novel therapeutics against SARS-CoV-2 using AI approaches that integrate information from experimental observations, and rigorous, physics-based virtual screening and molecular simulations, to identify viable drugs that can inhibit viral proteins.

APPROACH

The researchers, working closely with colleagues at Argonne's Advanced Photon Source, use what can be described as snapshots of the virus to determine its crystal structure. From this they try and identify sites of interest—potential targets for other molecules to bind to and/or attack. Simulations are used to provide details, including the effects of other molecules.

The AI approaches employed in this work integrate information from experimental observations, and rigorous, physics-based virtual screening and molecular simulations, to identify viable drugs that can inhibit viral proteins. These AI approaches, based on advances in deep learning and reinforcement learning, are capable of predicting how strongly a small molecule will bind to a protein as well as exploring the structural space of compounds that are predicted to bind to find more suitable variants.

RESULTS

Computational screening of small molecules has resulted in identifying small molecules that can potentially inhibit viral function in wet-lab experiments. These experiments involve live human lung cell cultures being exposed to small molecules followed by subsequent measurements that monitor viral replication. The molecules are being further refined to optimize them for binding to specific viral target proteins. Using AI techniques, the team has screened millions of small molecules and is validating them at Argonne for activity against the virus.

IMPACT

This work potentially could lead to the design of new generative models based on reinforcement learning for both small molecules and antibodies; and to the development of large-scale, AI-driven simulations of the entire viral particle and drugs bound to the various viral targets, as a better pathway to an antiviral drug.

PUBLICATIONS

Casalino, L., A. Dommer, Z. Gaieb, E. P. Barros, T. Sztain, S.-H. Ahn, A. Trifan, A. Brace, A. Bogetti, H. Ma, H. Lee, M. Turilli, S. Khalid, L. Chong, C. Simmerling, D. J. Hardy, J. D. C. Maia, J. C Phillips, T. Kurth, A. Stern, L. Huang, J. McCalpin, M. Tatineni, T. Gibbs, J. E. Stone, S. Jha, A. Ramanathan, and R. E. Amaro. "AI-Driven Multiscale Simulations Illuminate Mechanisms of SARS-CoV-2 Spike Dynamics," *International Journal of High Performance Computing Applications* (preprint), SAGE Publishing.

Biological Sciences | Simulation

COVID-19 Spread and Effectiveness of Interventions

PI Jonathan Ozik and Charles Macal, Argonne National Laboratory

AWARD Director's Discretionary
HOURS Theta: 357,000 Node-Hours

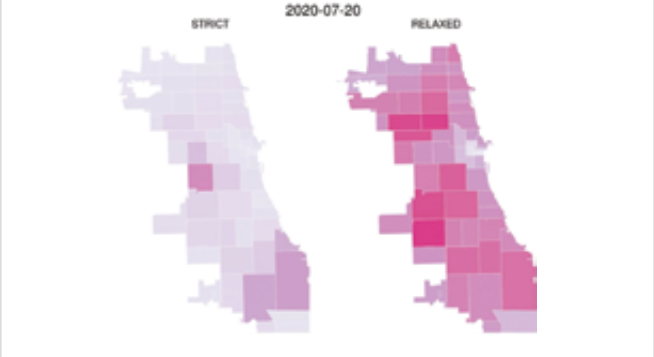

Weekly newly infected counts by ZIP code under differing individual behavior scenarios. *Image: Jonathan Ozik, Argonne National Laboratory*

This project, funded by the Joint DOE Laboratory Plan for Pandemic Modleing and Analysis Capability, oversees the development of epidemiological models to simulate the spread of COVID-19 throughout the population.

CHALLENGE

Argonne researchers have developed CityCOVID, an agent-based model capable of tracking detailed COVID-19 transmission. Agent-based modeling is an approach for capturing the dynamics of heterogeneous, interacting, adaptive agents at an individual, granular level of detail. When applied to a city like Chicago, CityCOVID includes a synthetic population representing the 2.7 million residents of Chicago and the 1.2 million geo-located locations—including households, schools, workplaces, hospitals, nursing homes, dormitories, and jails—where they can co-locate. Throughout a simulated day, each agent moves from place-to- place, hour-by-hour, engaging in social activities and interactions with co-located agents, where COVID-19 exposure events can occur. The COVID-19 disease progression is modeled within each agent, including differing symptom severities, hospitalizations, and age-dependent probabilities of transitions between disease stages.

APPROACH

The models pursue lines of inquiry familiar to anyone following the virus in news media—for example, the difference in outcomes that result from implementing various nonpharmaceutical interventions (NPIs) and how to safely ease off of the NPIs.

The project's significant computational demands result from the models' stochastic components, which encapsulate the underlying uncertainties and parameters of the simulation, and from the complexity of the population-level outcomes effected by the interactions of millions of individual software agents. The Argonne-developed technologies Repast HPC, ChiSIM (https://github.com/Repast/chiSIM), EMEWS, and Swift/T are used for model development and to run the large-scale parameter estimation and NPI scenario workflows on Theta.

RESULTS

CityCOVID is being used to calibrate unobserved model parameters, such as the time-varying degree of individual self-protective behaviors across the population, and to simulate a variety of interventions and future scenarios.

IMPACT

Model results are being provided to the City of Chicago and Cook County Public Health Departments as well as to the Illinois Governor's COVID-19 Modeling Task Force. While the Argonne team has been using Chicago as a testbed for developing these capabilities, CityCOVID is being extended to other regions as well.

PUBLICATIONS

Ozik, J., J. M. Wozniak, N. Collier, C. M. Macal, and M. Binois. "A Population Data-Driven Workflow for COVID-19 Modeling and Learning," *SC20* (November 2020), IEEE.

Biological Sciences | Data, Learning

X-ray Microscopy of Extended 3D Objects: Scaling Towards the Future

PI Chris Jacobsen, Argonne National Laboratory and Northwestern University

AWARD ADSP

HOURS Theta: 187,500 Node-Hours

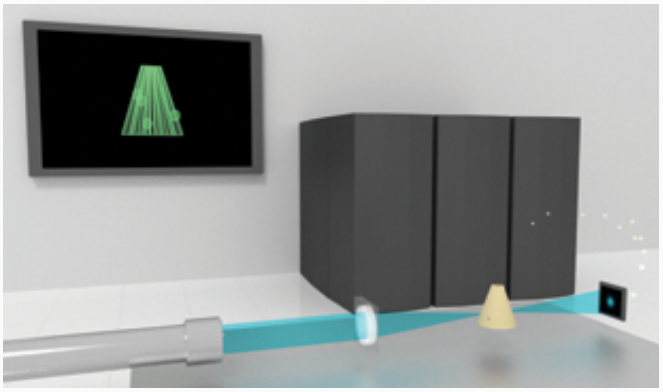

This illustration shows a coherent x-ray beam focused on a large-scale specimen while recording far-field diffraction patterns as the specimen is scanned and rotated. In the background is a computing system using automatic differentiation approaches to reconstruct a 3D image. *Image: Ming Du, Argonne National Laboratory*

The upcoming upgrade of Argonne's Advanced Photon Source will increase the brightness of its x-ray beams by up to 500 times. This will allow for research projects that are impossible at the current intensity, such as tracing the neural connections inside a mouse's brain to learn more about neurological disorders. But the enhanced experimental capabilities will also increase the need for more advanced reconstruction tools. A research team from Argonne National Laboratory and Northwestern University is leveraging ALCF supercomputing resources to develop a novel approach that can help complete a 3D reconstruction of x-ray images with more flexibility and less human effort than traditional computational methods.

CHALLENGE

The APS Upgrade project will enable researchers to image centimeter-thick samples with sub-100-nanometer resolution. However, the multiple scattering of x-ray inside a thick specimen's volume becomes increasingly problematic, as this is not accounted for by most conventional x-ray 3D reconstruction algorithms. This problem must be addressed before one can fully utilize the power of the upgraded light source.

APPROACH

To address this issue, the team is leveraging ALCF computing resources to develop a new approach that employs the "multislice tomography" method to simulate x-ray multiple scattering in a continuous 3D volume, and uses automatic differentiation toolboxes to implement the mathematics for solving the reconstruction problem. In preparation for the reconstruction of large-sized specimens, the team has also explored scalable wavefield simulation algorithms to enable the reconstruction software to take full advantage of hundreds of nodes on Theta.

RESULTS

In a paper published in *Science Advances*, the team detailed how the new algorithm can provide high-quality 3D reconstructions of heavily scattering samples measured under complicated imaging scenarios. The approach was shown to provide good reconstruction results that are free of diffraction-induced artifacts as seen in traditional reconstruction methods. With very minimal adaption, the code can work for both holography and ptychography. Moreover, in another paper recently accepted by *Optics Express*, the team demonstrated two scalable methods for x-ray wave propagation, namely the "tiling-based Fresnel multislice" method and the "finite-difference" method, both working efficiently while running over 16,000 processes on Theta.

IMPACT

By providing a new computational approach for reconstructing x-ray images of thick, real-life materials, this research aims to advance the full range of future nanoscale imaging activities, including cell and brain research, at Argonne's Advanced Photon Source and other DOE light sources. Furthermore, the power of machine learning-enabled automatic differentiation makes their software highly versatile and adaptable to many variants and setups of x-ray microscopes.

PUBLICATIONS

Du, M., Y. S. G. Nashed, S. Kandel, D. Gürsoy, and C. Jacobsen. "Three Dimensions, Two Microscopes, One Code: Automatic Differentiation for X-Ray Nanotomography Beyond the Depth of Focus Limit," *Science Advances* (March 2020), AAAS.

Ali, S., M. Du, M. Adams, B. Smith, and C. Jacobsen. "A comparison of distributed memory algorithms for x-ray wave propagation in inhomogeneous media," *Optics Express* (September 2020), OSA Publishing.

Earth Science | Simulation

Energy Exascale Earth System Model

PI Mark Taylor, Sandia National Laboratories
AWARD INCITE
HOURS Theta: 1,800,000 Node-Hours

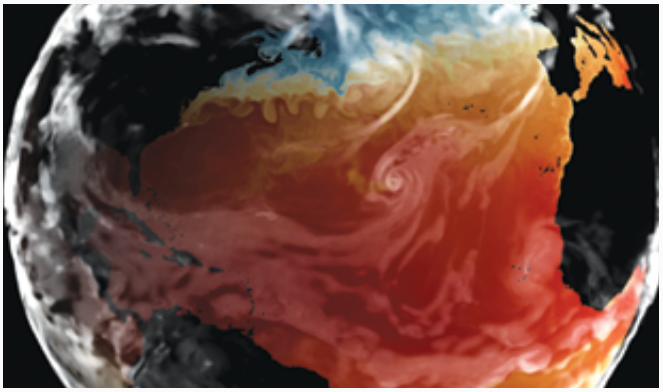

A snapshot of the atmosphere and ocean conditions in the high-resolution E3SM v1 simulation being run on ALCF Theta. Cloud fraction (grey-scale) shown on top of the sea surface temperature (color contours) with land masses in black. *Image: Mark Bolstad, Sandia National Laboratories*

With the coming paradigm shift in computer architectures and programming models as capability moves to the exascale era, the Energy Exascale Earthy System Model (E3SM) project aims to develop a cutting-edge climate and earth system that can tackle the most demanding climate research imperatives. Harnessing ALCF supercomputing resources, the E3SM project is addressing questions concerning the water cycle and cryosphere systems.

CHALLENGE

The research team is focusing on two grand challenge questions: (1) In the water cycle, how will more realistic portrayals of its important features (e.g., resolution, clouds, aerosols) affect river flow and associated freshwater supplies at the watershed scale? (2) In cryosphere systems, could a dynamic instability in the Antarctic Ice Sheet be triggered within the next 40 years? For (1), the team's objective is to simulate changes in the hydrological cycle with a specific focus on precipitation and surface water in orographically complex regions, such as the western United States and Amazon headwaters. For (2), in the first fully coupled global simulations to include ocean circulation and dynamic sub-ice shelf melting, the team aims to examine the near-term risk of initiating the West Antarctic Ice Sheet instability or collapse due to rapid melting by adjacent warming waters. Moreover, the simulations will focus on assessing the potential for significant increase in sub-ice shelf melt rates in response to past, ongoing, and future climatic changes.

APPROACH

The team makes extensive use of Theta to run the ES3M code, which comprises component models for atmosphere, ocean, sea ice, and land. Typically sixty-four tasks are assigned to every node, each with two OpenMP threads so as to establish intranodal parallelism. The majority of the cores are allocated to the atmosphere model, a subset of which also run the land and sea models; the remaining cores are allocated to the ocean model, which runs concurrently with the atmosphere model.

E3SM v1's high-resolution coupled configuration runs atmosphere and ocean components on a grid with 25-kilometer average spacing at the Equator with 72 vertical layers. The ocean and ice run on a variable resolution grid, with horizontal resolution ranging from 18 to 6 kilometers and 100 vertical layers.

RESULTS

Leveraging Theta, the researchers have completed 125 years of a 150 year pre-industrial control simulation. A transient or historical simulation representing the years 1970-2010 is also in progress.

IMPACT

In addition to further advancing the predictive power of climate models and providing insight into the climatic effects of rapid changes in the earth's ice content, the ES3M simulations have the potential to answer how water resources and the hydrological cycle interact with the climate system on both local and global scales. Hurricane hindcast simulations performed for this project demonstrated the high fidelity with which extreme weather can be modeled, while exposing parametric weaknesses that need improvement.

PUBLICATIONS

Bertagna, L., O. Guba, M. A. Taylor, J. G. Foucar, J. Larkin, A. M. Bradley, S. Rajamanickam, and A. G. Salinger. "A Performance-Portable Nonhydrostatic Atmosphere Dycore for the Energy Exascale Earth System Model Running at Cloud-Resolving Resolutions," *Multiphysics Applications, Supercomputing 2020* (November 2020), IEEE.

Caldwell, P., A. Mametjanov, Q. Tang, L. P. Van Roekel, J.-C. Golaz, W. Lin, D. C. Bader et al. "The DOE E3SM Coupled Model Version 1: Description and Results at High Resolution," *Journal of Advances in Modeling Earth Systems* (November 2019), John Wiley and Sons.

Engineering | Simulation

Shock-Induced Multi-Material Mixing

PI Sanjiva Lele, Stanford University
AWARD INCITE
HOURS Mira: 4,500,000 Million Node-Hours

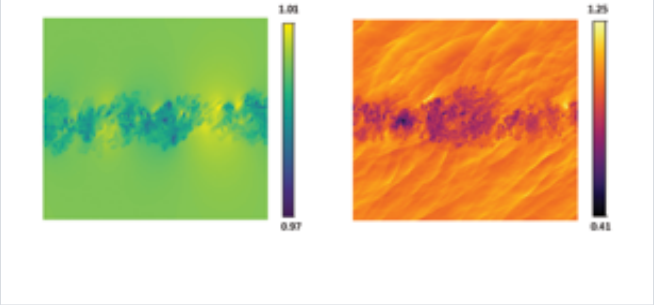

These images show the density fields at a low convective Mach number (left) and a higher convective Mach number (right). At high convective Mach number significant acoustic emission, including eddy shocklets, result from the turbulent mixing. *Image: Sanjiva Lele, Stanford University*

A detailed understanding of shock-induced turbulent mixing is critical to many engineering and scientific applications, including high-speed propulsion in hypersonic air-breathing engines for aerospace vehicles and energy generation through inertial confinement fusion. However, the turbulent flow physics involved these applications remains poorly understood. With this INCITE project, researchers from Stanford University have used ALCF computing resources to explore important aspects of the flow physics that are inaccessible to experiments.

CHALLENGE

Shock-induced turbulent mixing arises from a fluid flow phenomenon known as the Richtmyer-Meshkov instability, which occurs at the interface of two fluids when a shock passes through a fluid. The fast time scales and extreme conditions that characterize the flow physics in the vicinity of the shock-interface interaction region limit the scope of experimental measurements. Numerical simulations therefore play an important role in advancing understanding of compressible turbulent mixing.

APPROACH

Using the ALCF's Mira supercomputer, the team carried out high-resolution simulations of shock-induced turbulence resulting from the Richtmyer-Meshkov instability, with a focus on compressible mixing with shear and variable density effects. The simulations were performed on a domain with a spanwise extent that was four times larger than previous simulations. The larger span simulations have made it possible to obtain reliable results that are free from domain bounding effects until late times. The high-resolution data also makes it possible to quantify higher order statistics better and to make stronger quantitative claims about the novel findings that were suggested by previously obtained lower-resolution data from experiments and simulations.

RESULTS

The team's simulations have established a new regime of mode-coupling in shock-induced mixing and helped to quantify the effect of compressibility and the fundamental compressible energy transfer mechanisms. The researchers are using the simulation results to develop a benchmark-quality database for the model problem of a temporal mixing layer that is validated against experiments and for which the effects of the numerical modeling errors are thoroughly evaluated and demonstrated to be small. Having such a database on turbulence in the compressible regime will help advance the development of models for lower-fidelity methods.

IMPACT

The team's findings will advance the state of the art in simulations of shock-induced turbulent mixing and help improve engineering models of variable density turbulence. The enhanced models that result from this work could help enable innovations in new combustion devices and propulsion systems.

PUBLICATIONS

Jeun, J., G. J. Wu, and S. K. Lele. "Large Eddy Simulations of Screeching Twin Rectangular Jets," *AIAA Scitech 2020 Forum* (January 2020), AAIA.

Matsuno, K., and S. K. Lele. "Compressibility Effects in High Speed Turbulent Shear Layers – Revisited," *AIAA Scitech 2020 Forum* (January 2020), AAIA.

Data-Driven Materials Discovery for Optoelectronic Applications

PI Jacqueline Cole, University of Cambridge
AWARD ADSP
HOURS Theta: 468,750 Node-Hours

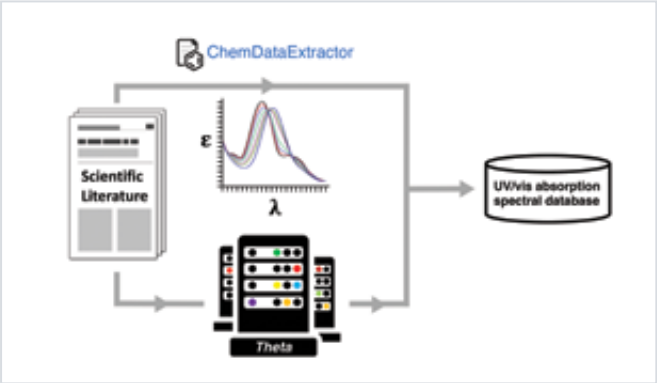

Auto-generating a UV/vis absorption spectral database via a dual experimental and computational chemical data pathway using the Theta supercomputer. *Image: Jacqueline Cole and Ulrich Mayer, University of Cambridge*

New materials are needed to drive advances in solar cells, data storage technologies, and other optoelectronic applications. With this ADSP project, researchers are developing data-driven, materials-by-design capabilities to accelerate the discovery of novel materials for targeted applications.

CHALLENGE

A majority of functional materials are discovered by trial-and-error methods, but this unpredictable approach can present a bottleneck to technological innovation. To overcome this issue, the ADSP team is using systematic molecular design and engineering strategies to develop algorithms that mine massive chemical datasets to enable property prediction and materials discovery.

APPROACH

The workflow for this project involves four key steps: data extraction, data enrichment, materials prediction, and experimental validation. Using text-mining tools, the researchers extract data from scientific literature to build a comprehensive repository of materials properties. To enrich the dataset, the team uses ALCF supercomputing resources to perform high-throughput electronic structure calculations, which produce paired quantities of experimental and computational data. The researchers then employ analytical methods to determine patterns in data that can be used to predict structure-function relationships for new materials. Finally, the team carries out experimental work to validate the candidate materials.

RESULTS

As detailed in a paper in *Scientific Data* (2019), the researchers developed a pipeline to auto-generate a comprehensive database of ultraviolet-visible (UV-vis) absorption spectra. They used their text-mining tool, ChemDataExtractor, on more than 400,000 scientific documents to identify 18,309 experimentally determined UV-vis absorption maxima. The team found large numbers of paired quantities of experimental and computational physical properties, laying the path for reliable in silico calculations of additional optical properties.

The team recently extended the application of their optoelectronics workflow to additional materials, including magnetic and superconducting compounds and battery chemicals. In a paper published in *npj Computational Materials*, the researchers successfully reconstructed the phase diagrams of well-known magnetic and superconducting materials, and demonstrated the ability to predict the phase-transition temperatures of novel compounds. For the battery materials study, they published a paper in Scientific Data (2020), detailing the creation of a database that contains over 250,000 records for a vast range of chemicals.

IMPACT

The team's use of data mining, in conjunction with large-scale simulations and experiments, offers a novel approach to advance the design and discovery of new functional materials. In addition, the project's development of open-source databases and data-extraction software tools will help accelerate materials discoveries by removing the hurdle of manual database creation.

PUBLICATIONS

E. J. Beard, G. Sivaraman, Á. Vázquez-Mayagoitia, V. Vishwanath, and J. M. Cole. "Comparative Dataset of Experimental and Computational Attributes of UV/vis Absorption Spectra," *Scientific Data* (December 2019), Springer Nature.

C. J. Court, and J. M. Cole. "Magnetic and Superconducting Phase Diagrams and Transition Temperatures Predicted Using Text Mining and Machine Learning," *npj Computational Materials* (March 2020), Springer Nature.

Huang, S., and J. M. Cole. "A database of battery materials auto-generated using ChemDataExtractor," *Scientific Data* (August 2020), Springer Nature.

Materials Science | Simulation

Electronic Stopping Simulation of Complex Systems

PI Yosuke Kanai, University of North Carolina at Chapel Hill

AWARD INCITE

HOURS Theta: 2,000,000 Node-Hours

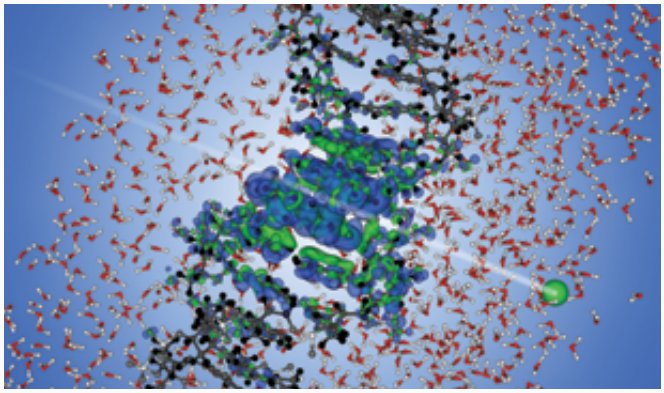

Rendering of solvated DNA undergoing proton irradiation. *Image: Christopher C. Shepard, University of North Carolina at Chapel Hill*

Electronic stopping describes the dynamical transfer of kinetic energy from energetic charged particles (e.g. protons) to electrons in a target matter, consequently inducing massive electronic excitations therein. Elucidation of this phenomenon as it occurs in various systems under ion irradiation contributes to impactful breakthroughs in a number of modern technologies. A team of researchers is developing and using predictive simulation methods to model and understand electronic stopping dynamics in DNA at the molecular level due to their importance in various applications such as proton-beam cancer therapy.

CHALLENGE

Understanding the electronic excitation dynamics of DNA under ion irradiation such as those of protons, alpha-particles, and carbon ions in the context of emerging ion beam cancer therapy and comparing to the electronic excitations under typical photon irradiation require a detailed atomistic model of water-solvated DNA as in the physiological condition. While quantum-mechanical simulations can provide such detailed description, complexity of the chemical system requires new innovations in massively-parallel time-dependent electron dynamics simulation.

APPROACH

This project continues to develop new approaches for a highly scalable implementation of real-time, time-dependent density functional theory using the massively parallelized Qb@ll electronic-structure code. Many thousands of processors in the Theta system are used to simulate the quantum-mechanical electronic response of complex systems to the ion irradiation.

RESULTS

In a recent paper published in *Journal of Physical Chemistry Letters*, the researchers describe recent progress they have made in simulating and understanding electronic stopping phenomena throughout a variety of systems. They describe how their quantum-mechanical method is used to investigate electronic stopping dynamics in relatively simple solids like silicon and aluminum, and how the work has been extended further to study complex systems like solvated DNA under ion irradiation in recent years.

IMPACT

Electronic stopping of protons and other light ions in solvated DNA, particularly in the context of ion-beam cancer therapies, represents a problem of great societal significance. The simulations performed in this work will provide much-needed insights into how DNA-damaging electronic excitation might differ under proton irradiation and traditional x- and gamma-ray irradiation in radiation oncology.

PUBLICATIONS

Yost, D. C., Y. Yao, and Y. Kanai. "First-Principles Modeling of Electronic Stopping in Complex Matter under Ion Radiation," *Journal of Physical Chemistry Letters* (December 2019), ACS Publications.

Modeling the Response of Fusion Plasma Facing Components

PI Brian Wirth, Oak Ridge National Laboratory
AWARD ALCC
HOURS Mira: 5,000,000 Node-Hours

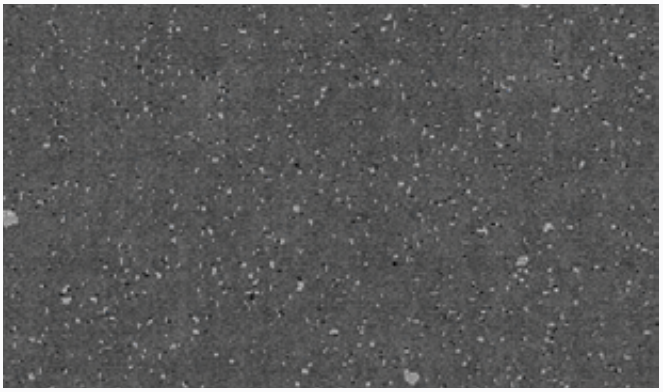

Snapshot from a simulation of helium atoms beneath a tungsten surface exposed to 100 eV helium plasma for 840 ns (fluence of 8.4×1018 m^{-2}. *Image: K. D. Hammond, University of Missouri, et al.,* Nuclear Fusion *(2019) vol. 59, no. 6.*.

Fusion as a practical energy source requires greater knowledge of plasma-surface interactions and materials engineering design of components to withstand the extreme heat and particle flux exposure conditions of a fusion power plant. This project, led by researchers from Oak Ridge National Laboratory, leveraged large-scale molecular dynamics (MD) simulations to model difficult-to-study phenomena that have limited our understanding of fusion.

CHALLENGE

A key complication to understanding fusion is the formation of "fuzz" on the surface of tungsten (that is, the primary plasma-facing divertor material in ITER) following exposure to material-destabilizing energetic helium ions. The mechanism driving this formation is challenging to probe, both experimentally and computationally. Large-scale MD simulations—relatively free of finite-size effects—were thus used to the mechanisms of helium transport and tungsten surface deformation.

APPROACH

MD simulations, rendered with the LAMMPS code, were used to model helium flux effects for two surface orientations. Lateral dimensions of up to 100 nanometers were considered to reduce finite-size effects, such as the formation of bubbles spanning the full width of the supercell. Four distinct helium fluxes were considered, ranging from 1025 to 1028 m−2s−1, with simulation times ranging from 500 nanoseconds up to 2.5 microseconds, depending on the system size. A temperature of 933 Kelvin was chosen so as to be at the low end of the fuzz-forming regime. The helium insertion depth was sampled from a distribution consistent with 100 electronvolt incident energy, which results in 90 percent of helium atoms being implanted within 5 nanometers of the surface.

RESULTS

Pronounced changes across three orders of magnitude of helium flux were observed for helium retention and helium bubble depth distribution, with deep bubbles tending to grow larger and create more prominent surface features over time. These results suggest that nearly all prior reported molecular simulations suffer from finite-size effects and that helium flux is a very important parameter in determining the behavior of helium in plasma-facing components.

IMPACT

The simulations provide much-needed benchmarks for larger-scale models, which are necessary to better understand the mechanisms of helium transport, retention, and agglomeration. This, in turn, will help to identify material design strategies and potentially mitigate issues related to the extreme conditions of a fusion energy environment.

PUBLICATIONS

Hammond, K. D., D. Maroudas, and B. D. Wirth. "Theoretical Model of Helium Bubble Growth and Density in Plasma-Facing Metals," *Scientific Reports* (February 2020), Springer Nature.

Hammond, K. D., I. V. Naeger, W. Widanagamaachchi, L.-T. Lo, D. Maroudas, and B. D. Wirth. "Helium Flux Effects on Bubble Growth and Surface Morphology in Plasma-Facing Tungsten from Large-Scale Molecular Dynamics Simulations," *Nuclear Fusion* (May 2019), IOP Publishing.

Materials Science | Simulation, Learning

Predictive Modeling and Machine Learning for Functional Nanoporous Materials

PI J. Ilja Siepmann, University of Minnesota
AWARD ALCC
HOURS Theta: 620,000 Node-Hours

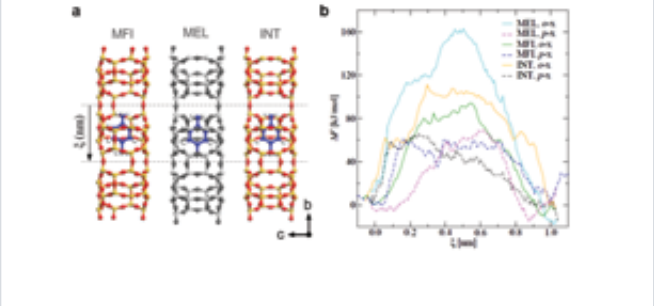

Diffusivity of *p*-xylene and *o*-xylene molecules through zeolite pores. (a) Projection of MFI, MEL and MEL/MFI interfacial pores with *p*-xylene at the center. (b) Free energy barriers for *p*-xylene and *o*-xylene transport. *Image: Evgenii Fetisov and J. Ilja Siepmann, University of Minnesota*

Nanoporous materials, such as zeolites, tailored for specific separation and catalytic processes have the potential to significantly improve energy efficiency in the chemical, biorenewable, and petrochemical industries. An interdisciplinary team led by researchers from the University of Minnesota is using predictive modeling and machine learning techniques on ALCF computing resources to accelerate the discovery and design of nanoporous materials for a variety of energy-related applications.

CHALLENGE

While commercially available zeolite membrane technologies are already enabling improved energy efficiency in gas separations and other industrial uses, efforts are underway to explore their application to hydrocarbon separations. In particular, mixtures of hydrocarbon isomers, like the three xylene isomers, are difficult to separate by distillation and their production could become more efficient by employing membranes that can operate with sufficiently high flux and selectivity at the temperatures and pressures required for membrane–reactor configurations. The zeolite MFI, a framework for widely used industrial catalysts, has shown promise as a thin-film membrane, but the crystal structure of 2D-MFI nanosheets and their relationship to separation performance remain elusive.

APPROACH

In a recent study, the multi-institutional team carried out transmission electron microscopy experiments that revealed intergrowths of the zeolite MEL exist within 2D-MFI. To assess the role of MEL intergrowths, the researchers performed first-principles molecular dynamics simulations on ALCF supercomputers to model the transport of xylene isomers through MFI-MEL nanosheets.

RESULTS

The team's atomistic simulations revealed that commensurate knitting of 1D-MEL within 2D-MFI creates more selective pores compared to pristine MFI nanosheets. Due to slight variations in pore cross section, transport of o-xylene through MEL and MEL-MFI interface pores can only be accomplished via significant framework distortions, whereas transport of *p*-xylene is less hindered which, in turn, leads to dramatic differences in permeability. In permeation experiments using an industrially relevant feed, the researchers achieved a separation factor of 60, which far exceeds previous experimental results. The observation of these intergrowths opens the door to the development of ultra-selective zeolite membranes. The team's findings were published in *Nature Materials*.

IMPACT

The team's research into novel zeolite materials is helping to advance the development of membrane technologies that can provide more energy-efficient separations processes for industry. More broadly, their work to improve the understanding and selection of nanoporous materials for energy applications, such as the production of biofuel and petroleum products, could lead to significant economic and environmental benefits.

PUBLICATIONS
Kumar, P., D. W. Kim, N. Rangnekar, H. Xu, E. O. Fetisov, S. Ghosh, H. Zhang, et al. "One-Dimensional Intergrowths in Two-Dimensional Zeolite Nanosheets and Their Effect on Ultra-Selective Transport." *Nature Materials* (February 2020), Springer Nature.

Physics | Simulation, Data

The Last Journey

PI Katrin Heitmann, Argonne National Laboratory
AWARD ALCC
HOURS Mira: 50,000,000 Node-Hours

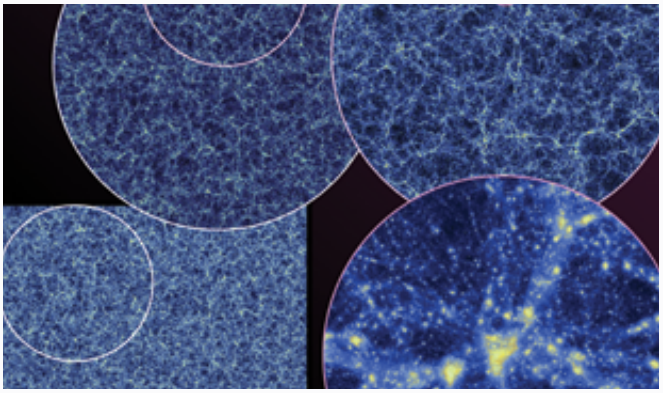

Dark-matter-dominated halos from a small region of the simulation. The radius of the spheres is proportional to halo mass; the dominant halo is the simulation's largest, at ~6×10^{15} solar masses. *Image: ALCF Visualization and Data Analysis Team and the HACC Team*

This project, led by Argonne National Laboratory researchers, is one of the world's five largest cosmological simulations of sufficient resolution and volume to permit the generation of detailed sky maps across multiple wavebands that are targeted to upcoming cosmological surveys. Running on the entirety of the Mira system, the simulation modeled the lifespan of the universe to help answer some of science's deepest questions, including the origins of dark matter and dark energy and what it means for the universe to expand at an accelerating pace.

CHALLENGE

The research team structured the simulation to begin 50 million years after the Big Bang, with conditions that agree with the most up-to-date cosmological theories. Billions of years of evolution between then and now were subsequently modeled in order to create a high-resolution model of what a large portion of the universe should look at present day.

APPROACH

The project was implemented using the Hardware/Hybrid Accelerated Cosmology Code (HACC) simulation and analysis framework on the full Mira system. It evolved more than 1.24 trillion particles to resolve the cosmological structures which host faint galaxies that will be observed by the Legacy Survey of Space and Time (LSST) project when it is carried out at the Vera Rubin Observatory. Cosmological parameters chosen to be consistent with the results from the Planck satellite. Analysis outputs were generated such that synthetic galaxy catalogs could be constructed using a semi-analytic modeling approach in post-processing.

RESULTS

The detailed history of the evolution of cosmological structures is now being processed to create synthetic sky maps for optical and cosmic microwave background surveys. Results are being documented in a sequence of papers to be published in *The Astrophysical Journal Supplement Series*.

IMPACT

This simulation was designed to address numerous fundamental questions in cosmology; the data produced are essential for enabling the refinement of existing predictive tools and aid the development of new models, impacting both ongoing and upcoming cosmological surveys, including the Dark Energy Spectroscopic Instrument (DESI), the LSST, SPHEREx, and the "Stage-4" ground-based cosmic microwave background experiment (CMB-S4).

PUBLICATIONS

Heitmann, K., N. Frontiere, E. Rangel, P. Larsen, A. Pope, I. Sultan, T. Uram, S. Habib, H. Finkel, D. Korytov, E. Kovacs, S. Rizzi, and J. Insley. "The Last Journey. I. An Extreme-Scale Simulation on the Mira Supercomputer," *The Astrophysical Journal Supplement Series* (January 2021), IOP Publishing.

Physics | Simulation

Towards a Definitive Model of Core-Collapse Supernova Explosions

PI Adam Burrows, Princeton University
AWARD INCITE
HOURS Theta: 2,000,000 Node-Hours

Fractured surface of the exploding core of a 16 solar-mass star computed on Theta. Bluish outer veil is the shock wave traveling ~10,000 kilometers per second. The surface is colored by electron fraction: purple is proton rich matter; blue is neutron rich. *Image: ALCF Visualization and Data Analysis Team*

Core-collapse supernovae dramatically announce the death of massive stars and the birth of neutron stars. These violent explosions, which produce the highest densities of matter and energy in the universe, are responsible for creating most of the elements in nature. A fundamental theoretical understanding of such explosions is needed to advance research in nuclear and particle physics, and to inform the interpretation of data from large-scale experiments. To shed light on this mysterious cosmological phenomenon, a research team led by Princeton University is using ALCF supercomputers to address whether and how 3D supernova explosion models differ from their 2D counterparts.

CHALLENGE

2D simulations of supernovae have supported the theory that capturing a small fraction of the neutrinos emitted during collapse powers supernova explosions, but detailed 3D calculations proving this paradigm are lacking. With the power of leadership-class supercomputers and continued advances in software, researchers now have the capabilities to tackle this longstanding challenge in nuclear astrophysics.

APPROACH

To carry out this research, the team is using FORNAX, their highly scalable, 3D radiation-hydrodynamics code. By addressing the transport operator with an explicit method, the code significantly reduces the computational complexity and communication overhead of traditional multidimensional radiative transfer solutions by bypassing the need for global iterative solvers. The team is running FORNAX on Theta to perform large-scale simulations aimed at determining if the neutrino mechanism is a robust explanation for supernova explosions and the stellar progenitor dependence of their observables.

RESULTS

After completing and analyzing the largest suite of 3D simulations ever performed, spanning a broad range of progenitor masses and structures, the team is now focusing on a few long-term simulations to witness the achievement of the asymptotic explosion state. Simulating a full-physics 3D model for more than two seconds of physical time has never before been done, but the more massive star models seem to require such durations to reach the final explosion energy. In addition, the team has completed a comprehensive study of proto-neutron star (PNS) convection in the deep interior of the newly born neutron star. Such core convection may be implicated in the generation of pulsar magnetic fields and affects the neutrino-cooling rate of the residue. Hence, an exploration of the variety of convective behaviors experienced in the core can speak volumes about the hand-off between the explosive and the PNS phases. Both studies have been published in the *Monthly Notices of the Royal Astronomical Society*.

IMPACT

The team's efforts to advance the fundamental theoretical understanding of supernova explosions will benefit ongoing research efforts to determine the origin of the elements in the universe, measure gravitational waves, and interpret laboratory nuclear reaction rate measurements in light of stellar nucleosynthesis.

PUBLICATIONS

Burrows, A., D. Radice, D. Vartanyan, H. Nagakura, M. A. Skinner, and J. C. Dolence. "The Overarching Framework of Core-Collapse Supernova Explosions as Revealed by 3D Fornax Simulations," *Monthly Notices of the Royal Astronomical Society* (January 2020), Oxford University Press.

Nagakura, H., Burrows, A., Radice, D., and Vartanyan, D. "A Systematic Study of Proto-Neutron Star Convection in Three-Dimensional Core-Collapse Supernova Simulations," *Monthly Notices of the Royal Astronomical Society* (January 2020), Oxford University Press.

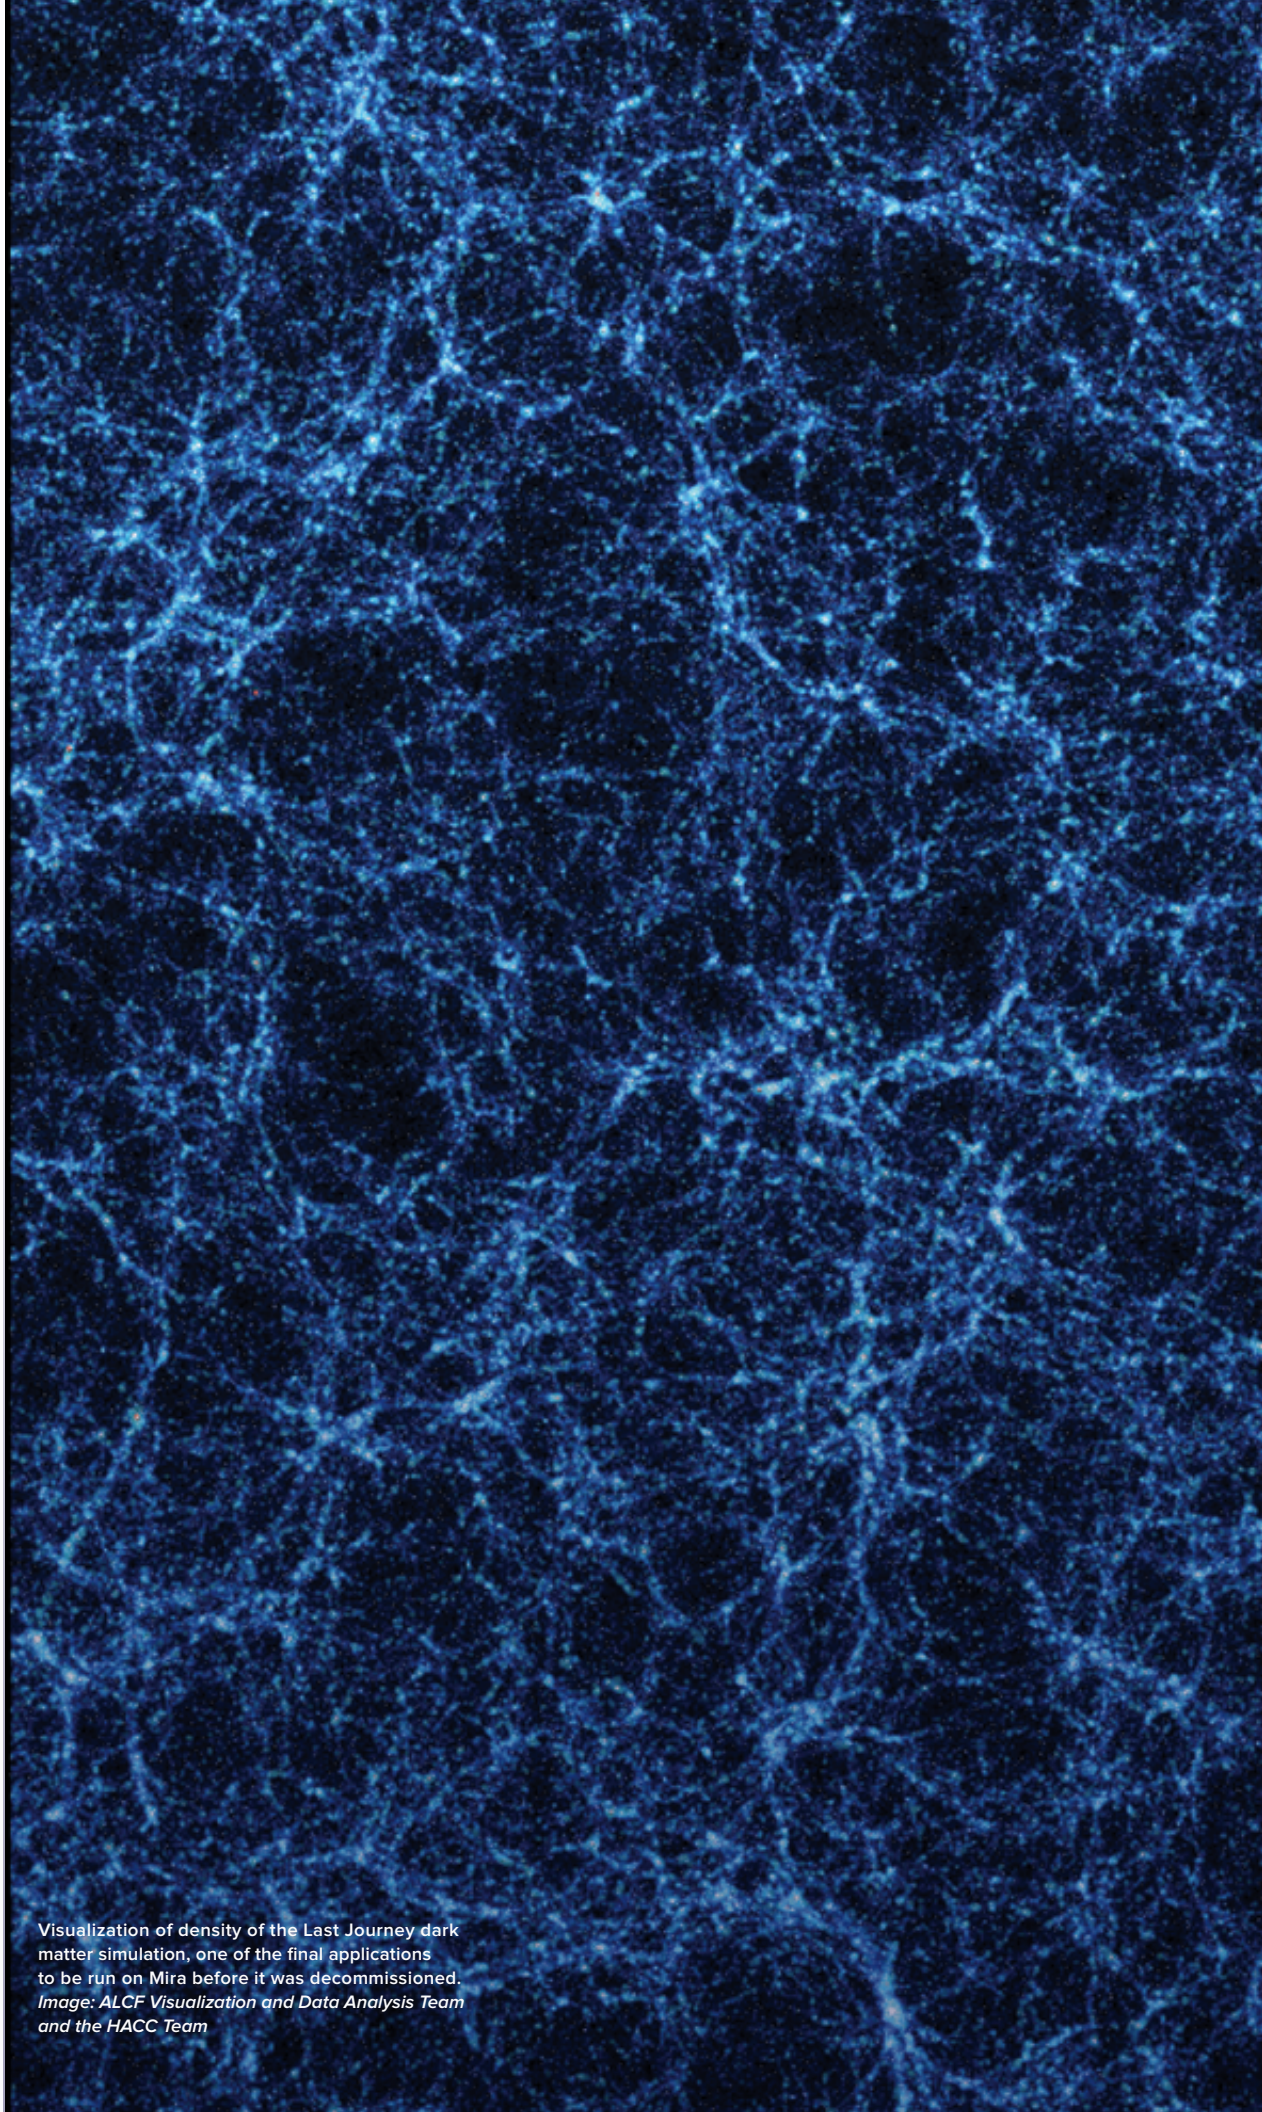

Visualization of density of the Last Journey dark matter simulation, one of the final applications to be run on Mira before it was decommissioned. *Image: ALCF Visualization and Data Analysis Team and the HACC Team*

ALCF Projects

2020 INCITE

BIOLOGICAL SCIENCES

Computational Physical Genomics: Exploring Potential Novel Cancer Therapies

PI Allen Taflove, Northwestern University
HOURS ALCF: 1,000,000 Node-Hours

EARTH SCIENCE

Energy Exascale Earth System Model

PI Mark Taylor, Sandia National Laboratories
HOURS ALCF: 1,800,000 Node-Hours
OLCF: 1,000,000 Node-Hours

Extreme-Scale Simulations for Advanced Seismic Ground Motion and Hazard Modeling

PI Christine Goulet, University of Southern California
HOURS ALCF: 500,000 Node-Hours
OLCF: 300,000 Node-Hours

ENGINEERING

High-Speed Turbulence with Shocks over Non-Adiabatic and Flexible Walls

PI Johan Larsson, University of Maryland
HOURS ALCF: 2,000,000 Node-Hours

MATERIALS SCIENCE

Electronic Stopping Simulation of Complex Systems

PI Yosuke Kanai, University of North Carolina at Chapel Hill
HOURS ALCF: 2,000,000 Node-Hours

Towards Predictive Simulations of Functional and Quantum Materials

PI Paul Kent, Oak Ridge National Laboratory
HOURS ALCF: 1,500,000 Node-Hours
OLCF: 400,000 Node-Hours

Ultrafast Control of Functional Materials

PI Priya Vashishta, University of Southern California
HOURS ALCF: 2,000,000 Node-Hours

PHYSICS

Ab-initio Nuclear Structure and Nuclear Reactions

PI Gaute Hagen, Oak Ridge National Laboratory
HOURS ALCF: 1,000,000 Node-Hours
OLCF: 550,000 Node-Hours

Extreme-Scale Simulation of Supernovae and Magnetars from Realistic Progenitors

PI Sean Couch, Michigan State University
HOURS ALCF: 1,000,000 Node-Hours

High-Fidelity Gyrokinetic Simulation of Tokamak and ITER Edge Physics

PI Choongseock Chang, Princeton Plasma Physics Laboratory
HOURS ALCF: 1,500,000 Node-Hours
OLCF: 970,000 Node-Hours

Petascale Simulations of Kinetic Effects in IFE Plasmas

PI Frank Tsung, University of California, Los Angeles
HOURS ALCF: 1,000,000 Node-Hours

PLASM-IN-SILICO: HPC Modeling of High-Intensity Laser-Solid Interaction

PI Jean-Luc Vay, Lawrence Berkeley National Laboratory
HOURS ALCF: 600,000 Node-Hours
OLCF: 110,000 Node-Hours

Radiation Hydrodynamic Simulations of Massive Stars with Rotation

PI Lars Bildsten, Kavli Institute for Theoretical Physics
HOURS ALCF: 1,800,000 Node-Hours

Towards a Definitive Model of Core-Collapse Supernova Explosions

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

ALCC 2019–2020

CHEMISTRY

Accelerated Catalyst Discovery from First Principles Simulations and Machine Learning

PI Rajeev Surendran Assary, Argonne National Laboratory
HOURS ALCF: 240,000 Node-Hours

Towards Exascale Internal Combustion Engine Simulation with In-Situ Analysis

PI Muhsin Ameen, Argonne National Laboratory
HOURS ALCF: 630,000 Node-Hours

ENERGY TECHNOLOGIES

High-Fidelity Physics Simulations for DOE and Industry Fast Spectrum Nuclear Reactor Systems

PI Emily Shemon, Argonne National Laboratory
HOURS ALCF: 880,000 Node-Hours

Nuclear Energy Industry Validation of Nek5000: ALAIN and HYMERES

PI Aleksander Obabko, Argonne National Laboratory
HOURS ALCF: 340,000 Node-Hours

MATERIALS SCIENCE

Predictive Modeling and Machine Learning for Functionally Nanoporous Materials

PI J. Ilja Siepmann, University of Minnesota
HOURS ALCF: 620,000 Node-Hours
NERSC: 200,000 Node-Hours

Supercomputing for Automotive High-Temperature Alloy Design

PI Dongwon Shin, Oak Ridge National Laboratory
HOURS ALCF: 350,000 Node-Hours
OLCF: 391,000 Node-Hours

PHYSICS

Low Energy Neutrino-Nucleus Interactions

PI Saori Pastore, Washington University in St. Louis
HOURS ALCF: 390,000 Node-Hours

Modeling the Response of Fusion Plasma Components

PI Brian Wirth, Oak Ridge National Laboratory
HOURS ALCF: 200,000 Node-Hours
OLCF: 250,000 Node-Hours

Neutrino Flux, Energy Deposition and Radiological Studies for the DUNE-LBNF Beamline

PI Igor Rakhno, Fermilab
HOURS ALCF: 450,000 Node-Hours

The Next Leap Forward in LSST Sky Simulations

PI Katrin Heitmann, Argonne National Laboratory
HOURS ALCF: 400,000 Node-Hours

Particle Acceleration in Plasma Jets: From Astrophysics to the Laboratory

PI Paulo E. Alves, SLAC National Accelerator Laboratory
HOURS ALCF: 1,000,000 Node-Hours

Semileptonic B- and D-Meson Form Factors with High Precision

PI Aida El-Kahdra, University of Illinois at Urbana-Champaign
HOURS ALCF: 400,000 Node-Hours
NERSC: 330,000 Node-Hours

ALCC 2020–2021

COMPUTER SCIENCE

Enabling Resilient and Portable Workflows from DOE's Experimental Facilities

PI Katie Antypas, Lawrence Berkeley National Laboratory
HOURS ALCF: 100,000 Node-Hours
OLCF: 20,000 Node-Hours

CHEMISTRY

Benchmarking Many-Body Perturbation Theory

PI Olle Heinonen, Argonne National Laboratory
HOURS ALCF: 100,000 Node-Hours

Interpretable Machine Learning Force Fields for Accurate Chemical Reactive Dynamic

PI Olexandr Isayev, Carnegie Mellon University
HOURS ALCF: 359,000 Node-Hours

Stochastic A Priori Dynamics for Complex Reactive Chemical Environments

PI Ahren Jasper, Argonne National Laboratory
HOURS ALCF: 100,000 Node-Hours

Understanding the Role of Hierarchical Correlations in Solution-Phase Chemical Separations

PI Lynda Soderholm, Argonne National Laboratory
HOURS ALCF: 100,000 Node-Hours

EARTH SCIENCE

Variable-Resolution Earth System Modeling of the Cryosphere with E3SM

PI Darin Comeau, Los Alamos National Laboratory
HOURS ALCF: 400,000 Node-Hours
NERSC: 500,000 Node-Hours

ENERGY TECHNOLOGIES

Automatic Building Energy Modeling

PI Joshua New, Oak Ridge National Laboratory
HOURS ALCF: 300,000 Node-Hours

MATERIALS SCIENCE

Data-Driven Molecular Engineering of Advanced Functional Materials

PI Jacqueline Cole, University of Cambridge
HOURS ALCF: 100,000 Node-Hours

High-Temperature Material Properties from First Principles

PI Mark Messner, Argonne National Laboratory
HOURS ALCF: 200,000 Node-Hours

Many-Body Perturbation Theory Meets Machine Learning to Discover Materials for Organic Photovoltaics

PI Noa Marom, Carnegie Mellon University
HOURS ALCF: 100,000 Node-Hours

Plasma Surface Interaction Modeling

PI Brian Wirth, University of Tennessee
HOURS ALCF: 318,000 Node-Hours
OLCF: 155,000 Node-Hours
NERSC: 30,000 Node-Hours

Predictive Modeling of Nanoporous Materials and Multiphase Systems

PI Joern Siepmann, University of Minnesota
HOURS ALCF: 220,000 Node-Hours

Supercomputing for Automotive High-Temperature Alloy Design

PI Dongwon Shin, Oak Ridge National Laboratory
HOURS ALCF: 100,000 Node-Hours

PHYSICS

Toward the Future: High-Fidelity Simulation for Next-Generation Nuclear Reactors

PI Yiqi Yu, Argonne National Laboratory
HOURS ALCF: 208,000 Node-Hours
OLCF: 300,000 Node-Hours

Reconstructing Neutrino Data with the MicroBooNE Liquid Argon Detector

PI Andrzej Szelc, University of Manchester
HOURS ALCF: 200,000 Node-Hours

Optimization Studies of the LBNF - PIP-II Complex for Megawatt Beams on Target

PI Igor Rakhno, Fermi National Accelerator Laboratory
HOURS ALCF: 450,000 Node-Hours

DNS Simulations of Coolant Flow in the High-Flux Isotope Reactor

PI Emilian Popov, Oak Ridge National Laboratory
HOURS ALCF: 220,000 Node-Hours

Chiral Nuclear Interactions from Nuclei to Nucleonic Matter

PI Maria Piarulli, Washington University in St. Louis
HOURS ALCF: 200,000 Node-Hours

Nucleon Axial Charge with All-Staggered Lattice QCD

PI Andreas Kronfeld, Fermi National Accelerator Laboratory
HOURS ALCF: 200,000 Node-Hours
NERSC: 870,000 Node-Hours

Distributed Large Wavefield Propagation and 3D Reconstruction Beyond the Depth of Focus Limit

PI Ming Du, Argonne National Laboratory
HOURS ALCF: 250,000 Node-Hours

Field-Reversed Configuration Stability and Transport

PI Sean Dettrick, TAE Technologies Inc.
HOURS ALCF: 64,000 Node-Hours

Improving Direct-Drive Laser Fusion Predictive Capability with a Simulation Database

PI Duc Cao, Laboratory for Laser Energetics
HOURS ALCF: 318,000 Node-Hours

High-Luminosity LHC Detector Upgrade Studies by the ATLAS and CMS Collaborations

PI Douglas Benjamin, Argonne National Laboratory
HOURS ALCF: 950,000 Node-Hours

Multiphase Flow Simulations of Reactor Flows

PI Igor Bolotnov, North Carolina State University
HOURS ALCF: 192,000 Node-Hours
NERSC: 250,000 Node-Hours

ALCF DATA SCIENCE PROGRAM

Advanced Materials Characterization with AI-Informed Computation

PI Marco Govoni, Argonne National Laboratory

Deep Learning at Scale for Multimessenger Astrophysics Through the NCSA-Argonne Collaboration

PI Eliu Huerta, University of Illinois at Urbana-Champaign

Developing High-Performance-Computing Applications for Liquid Argon Neutrino Detectors

PI Andrzej Szelc, University of Manchester

Dynamic Compressed Sensing for Real-Time Tomographic Reconstruction

PI Robert Hovden, University of Michigan

Machine Learning Magnetic Properties of van der Waals Heterostructures

PI Trevor Rhone, Harvard University

X-ray Microscopy of Extended 3D Objects: Scaling Towards the Future

PI Chris Jacobsen, Argonne National Laboratory and Northwestern University

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 and Oak Ridge National Laboratory

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

2020 DIRECTOR'S DISCRETIONARY

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

BIOLOGICAL SCIENCES

Agent-Based Model Called CityCOVID Capable of Tracking Detailed COVID-19 Transmission

PI Jonathan Ozik, Argonne National Laboratory

AI-Driven Integrative Biology for Accelerating Therapeutic Discovery Against SARS-CoV-2

PI Arvind Ramanathan, Thomas Brettin, Argonne National Laboratory

Compound Screening to Repurpose FDA-Approved Drugs Against SARS-CoV-2 Catalytic Enzymes

PI Albert Y. Lau, Johns Hopkins University

Machine Learning of Drug Binding and Toxicity Based on High-Throughput Free Energy Computations

PI Wei Jiang, Johns Hopkins University

Modeling Coronavirus

PI Zhangli Peng, University of Illinois at Chicago and University of Notre Dame

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

PI Ao Ma, University of Illinois at Chicago

CHEMISTRY

INtegrated and Scalable PredictIon of REsistanace (INSPIRE)

PI Peter Coveney, University College London

Ionic Liquid as a Potential Electrolyte Of High Performance Lithium Ion Battery

PI Zhengcheng Zhang, Argonne National Laboratory

COMPUTER SCIENCE

Collaboration with Cray on Interconnect-Related Studies

PI Sudheer Chunduri, Argonne National Laboratory

Distributed Relational Algebra at Scale

PI Sidharth Kumar, University of Alabama at Birmingham

Energy Efficient Tradeoff Among Execution Time, Power, and Resilience of Two ECP Applications

PI Xingfu Wu, Argonne National Laboratory

EARTH SCIENCE

Linking Climate to Water: Implementing a 4km Regional Climate Model with Hydrologic Model Coupling (WRF-Hydro)

PI Jiali Wang, Argonne National Laboratory

ENERGY TECHNOLOGIES

High-Fidelity CFD Simulations of Multi-Mode Combustion

PI Pinaki Pal, Argonne National Laboratory

ENGINEERING

Direct Numerical Simulation of Three-Dimensional Turbulence

PI Ramesh Balakrishnan, Argonne National Laboratory

High-Fidelity Simulation of Supersonic Turbulent Flow-Structure Interaction and Mixing

PI Ivan Bermejo Moreno, University of Southern California

MATERIALS SCIENCE

Metastable Phase Diagram of Material

PI Subramanian Sankaranarayanan, Argonne National Laboratory

Rational Design of Ultrastrong Composites

PI Hendrik Heinz, University of Colorado Boulder

Structure and Properties of Grain Boundaries in Materials for Energy Applications

PI Wissam Saidi, University of Pittsburgh

PHYSICS

Exploring Astrophysical Particle Acceleration in HED Laboratory Plasmas

PI Frederico Fiuza, SLAC National Accelerator Laboratory

Artificial Intelligence Assisted Safety Modeling and Analysis of Advanced Reactors

PI Rui Hu, Argonne National Laboratory

ALCF Publications

Researchers who use ALCF resources are major contributors to numerous publications that document their breakthrough science and engineering. The refereed journal articles and conference proceedings represent research ventures undertaken at the ALCF through programs supported by the U.S. Department of Energy and Argonne National Laboratory.

The publications are listed by their publication dates. An asterisk after a name indicates an ALCF author. ALCF publications are listed online at alcf.anl.gov/publications.

January

Adams, M. P., M. L. Adams, W. D. Hawkins, T. Smith, L. Rauchwerger, N. M. Amato, T. S. Bailey, R. D. Falgout, A. Kunen, and P. Brown. "Provably optimal parallel transport sweeps on semi-structured grids," *Journal of Computational Physics* (January 2020), Elsevier. doi: 10.1016/j.jcp.2020.109234

Alexander, F., A. Almgren, J. Bell, A. Bhattacharjee, J. Chen, P. Colella, D. Daniel, J. DeSlippe, L. Diachin, E. Draeger, A. Dubey, T. Dunning, T. Evans, I. Foster, M. Francois, T. Germann, M. Gordon, S. Habib, M. Halappanavar, S. Hamilton, W. Hart, Z. Huang, A. Hungerford, D. Kasen, P. R. C. Kent, T. Kolev, D. B. Kothe, A. Kronfeld, Y. Luo, P. Mackenzie, D. McCallen, B. Messer, S. Mniszewski, C. Oehmen, A. Perazzo, D. Perez, D. Richards, W. J. Rider, R. Rieben, K. Roche, A. Siegel, M. Sprague, C. Steefel, R. Stevens, M. Syamlal, M. Taylor, J. Turner, J.-L. Vay, A. F. Voter, T. L. Windus, and K. Yelick . "Exascale applications: skin in the game," *Philosophical Transactions of the Royal Society A* (January 2020), The Royal Society. doi: 10.1098/rsta.2019.0056

Ali, S., M. Du, M. F. Adams, B. Smith, and C. Jacobsen. "Comparison of distributed memory algorithms for X-ray wave propagation in inhomogeneous media," *Optics Express* (January 2020), OSA Publishing. doi: 10.1364/OE.400240

Bassenne, M., L. Fu, and A. Mani. "Time-Accurate and highly-Stable Explicit operators for stiff differential equations," *Journal of Computational Physics* (January 2020), Elsevier. doi: 10.1016/j.jcp.2020.109847

Bauberger, S., A. Freitas, and D. Wiegand. "TVID 2: evaluation of planar-type three-loop self-energy integrals with arbitrary masses," *Journal of High Energy Physics* (January 2020), Springer Nature. doi: 10.1007/jhep01(2020)024

Boughezal, R., A. Isgrò, and F. Petriello. "Next-to-leading power corrections to V+1 jet production in N-jettiness subtraction," *Physical Review D* (January 2020), APS. doi: 10.1103/physrevd.101.016005

Brockmeyer, L., E. Merzari, J. Solberg, and Y. Hassan. "One-way coupled simulation of FIV in a 7-pin wire-wrapped fuel pin bundle," *Nuclear Energy and Design* (January 2020), Elsevier, doi: 10.1016/j.nucengdes.2019.110367

Byna, S., M. S. Breitenfeld, B. Dong, Q. Koziol, E. Pourmal, D. Robinson, J. Soumagne, H. Tang, V. Vishwanath, and R. Warren. "ExaHDF5: Delivering Efficient Parallel I/O on Exascale Computing Systems," *Journal of Computer Science and Technology* (January 2020), Springer Nature. doi: 10.1007/s11390-020-9822-9

Chan, H., M. Cherukara, T. D. Loeffler, B. Narayanan, and S. K. R. S. Sankaranarayanan. "Machine Learning Enabled Autonomous Microstructural Characterization in 3D Samples," *Nature Communications* (January 2020), Springer Nature. doi: 10.1038/s41524-019-0267-z

Du, M., D. Gürsoy, and C. Jacobsen. "Near, far, wherever you are: simulations on the dose efficiency of holographic and ptychographic coherent imaging," *Journal of Applied Crystallography* (January 2020), International Union of Crystallography. doi: 10.1107/S1600576720005816

Eid, A., A. Eshein, Y. Li, R. Virk, D. Van Derway, D. Zhang, A. Taflove, and V. Backman. "Characterizing chromatin packing scaling in whole nuclei using interferometric microscopy," *Optics Letters* (January 2020), OSA Publishing. doi: 10.1364/OL.400231

Fodor, Z., K. Holland, J. Kuti, D. Nogradi, and C. H. Wong. "Case studies of near-conformal beta-functions," *Proceedings of Science* (January 2020), Sissa Medialab. doi: 10.22323/1.363.0121

Hannah, W. M., C. R. Jones, B. R. Hillman, M. R. Norman, D. C. Bader, M. A. Taylor, L. R. Leung, M. S. Pritchard, M. D. Branson, G. Lin, K. G. Pressel, and J. M. Lee. "Initial Results From the Super-Parameterized E3SM," *JAMES Journal of Advances in Modeling Earth Systems* (January 2020), AGU. doi: 10.1029/2019MS001863

Ho, P. J., B. J. Daurer, M. F. Hantke, J. Bielecki, A. Haddad, M. Bucher, G. Doumy, K. R. Ferguson, L. Flückiger, T. Gorkhover, Bi. Iwan, C. Knight, S. Moeller, T. Osipov, D. Ray, S. H. Southworth, M. Svenda, N. Timneanu, A. Ulmer, P. Walter, J. Hajdu, L. Young, F. R. N. C. Maia, and C. Bostedt. "The role of transient resonances for ultra-fast imaging of single sucrose nanoclusters," *Nature Communications* (January 2020), Springer Nature. doi: 10.1038/s41467-019-13905-9

Larsen, R., S. Meinel, S. Mukherjee, and P. Petreczky. "Excited bottomonia in quark-gluon plasma from lattice QCD," *Physics Letters B* (January 2020), Elsevier. doi: 10.1016/j.physletb.2019.135119

Loeffler, T. D., T. K. Patra, H. Chan, M. Cherukara, and S. K. R. S. Sankaranarayanan. "Active Learning the Potential Energy Landscape for Water Clusters from Sparse Training Data," The Journal of Physical Chemistry C (January 2020), ACS. doi: 10.1021/acs.jpcc.0c00047

Lozano-Duran, A., S. T. Bose, and P. Moin. "Prediction of trailing edge separation on the NASA Juncture Flow using wall-modeled LES," *AIAA Scitech 2020 Forum* (January 2020), Orlando, FL, AIAA. doi: 10.2514/6.2020-1776

Matsuno, K., and S. K. Lele. "Compressibility Effects in High Speed Turbulent Shear Layers – Revisited," *AIAA Scitech 2020 Forum* (January 2020), Orlando, FL, AIAA. doi: 10.2514/6.2020-0573

Milan, P. J., R. Torelli, B. Lusch, and G. Magnotti. "Data-Driven Model Reduction of Multiphase Flow in a Single-Hole Automotive Injector," *Atomization and Sprays* (January 2020), Begell House. doi: 10.1615/AtomizSpr.2020034830

Motamarri, P., S. Das, S. Rudraraju, K. Ghosh, D. Davydov, and V. Gavini. "DFT-FE – A massively parallel adaptive finite-element code for large-scale density functional theory calculations," *Computer Physics Communications* (January 2020), Elsevier. doi: 10.1016/j.cpc.2019.07.016

Nagakura, H., A. Burrows, D. Radice, and D. Vartanyan. "A systematic study of proto-neutron star convection in three-dimensional core-collapse supernova simulations," *Monthly Notices of the Royal Astronomical Society* (January 2020), Oxford University Press. doi: 10.1093/mnras/staa261

Nutschel, C., A. Fulton, O. Zimmermann, U. Schwaneberg, K.-E. Jaeger, and H. Gohlke. "Systematically Scrutinizing the Impact of Substitution Sites on Thermostability and Detergent Tolerance for Bacillus subtilis Lipase A," *Journal of Chemical Information and Modeling* (January 2020), ACS. doi: 10.1021/acs.jcim.9b00954

Permann, C. J., D. R. Gaston, D. Andrš, R. W. Carlsen, F. Kong, A. D. Lindsay, J. M. Miller, J. W. Peterson, A. E. Slaughter, R. H. Stogner, and R. C. Martineau. "MOOSE: Enabling massively parallel multiphysics simulation," *Physical Letters B* (January 2020), Elsevier. doi: 10.1016/j.physletb.2020.136029

Ponomarenko, N. S., O. Kokhan, P. R. Pokkuluri, K. L. Mulfort, and D. M. Tiede. "Examination of abiotic cofactor assembly in photosynthetic biomimetics: site-specific stereoselectivity in the conjugation of a ruthenium(II) tris(bipyridine) photosensitizer to a multi-heme protein," *Photosynthesis Research* (January 2020), Springer Nature. doi: 10.1007/s11120-019-00697-8

Ross, R. B., G. Amvrosiadis, P. Carns, C. D. Cranor, M. Dorier, K. Harms, G. Ganger, G. Gibson, S. K. Gutierrez, R. Latham, B. Robey, D. Robinson, B. Settlemyer, G. Shipman, S. Snyder, J. Soumagne, and Q. Zheng. "Mochi: Composing Data Services for High-Performance Computing Environments," *Journal of Computer Science and Technology* (January 2020), Springer Nature. doi: 10.1007/s11390-020-9802-0

Ryu, S., and J. Lee. "Evolution of the Deterministic Collapse Barrier of the Field Clusters as a Probe of Cosmology," *The Astrophysical Journal* (January 2020), IOP Science. doi: 10.3847/1538-4357/ab63ce

Shakya, G., S. E. Hoff, S. Wang, H. Heinz, X. Ding, and M. A. Borden. "Vaporizable endoskeletal droplets via tunable interfacial melting transitions," *Science Advances* (January 2020), AAAS. doi: 10.1126/sciadv.aaz7188

Somà, V., P. Navrátil, F. Raimondi, C. Barbieri, and T. Duguet. "Novel chiral Hamiltonian and observables in light and medium-mass nuclei," *Physical Review C* (January 2020), APS. doi: 10.1103/PhysRevC.101.014318

Sveinsson, H. A., A. Hafreager, R. K. Kalia, A. Nakano, P. Vashishta, and A. Malthe-Sørenssen. "Direct Atomic Simulations of Facet Formation and Equilibrium Shapes of SiC Nanoparticles," *Crystal Growth and Design* (January 2020), ACS. doi: 10.1021/acs.cgd.9b00612

Volpiani, P. S., M. Bernardini, and J. Larsson. "Effects of a nonadiabatic wall on hypersonic shock/boundary-layer interactions," *Physical Review Fluids* (January 2020), APS. doi: 10.1103/PhysRevFluids.5.014602

Wei, W., and E. A. Huerta. "Gravitational wave denoising of binary black hole mergers with deep learning," *Physics Letters B* (January 2020), Elsevier. doi: 10.1016/j.physletb.2019.135081

Yang, T. T., T. L. Tan, and W. A. Saidi. "High Activity toward the Hydrogen Evolution Reaction on the Edges of MoS2-Supported Platinum Nanoclusters Using Cluster Expansion and Electrochemical Modeling," *Chemistry of Materials* (January 2020), ACS. doi: 10.1021/acs.chemmater.9b05244

Zhdankin, V., D. A. Uzdensky, G. R. Werner, and M. C. Begelman. "Kinetic turbulence in shining pair plasma: intermittent beaming and thermalization by radiative cooling," *Monthly Notices of the Royal Astronomical Society* (January 2020), Oxford University Press. doi: 10.1093/mnras/staa284

Zimmer, M. D., and I. A. Bolotnov. "Exploring Two-Phase Flow Regime Transition Mechanisms Using High-Resolution Virtual Experiments," *Nuclear Science and Engineering* (January 2020), Taylor and Francis Group. doi: 10.1080/00295639.2020.1722543

February

Abramczyk, M., T. Blum, T. Izubuchi, C. Jung, M. Lin, A. Lytle, S. Ohta, and E. Shintani (RBC and UKQCD Collaborations). "Nucleon mass and isovector couplings in 2 + 1-flavor dynamical domain-wall lattice QCD near physical mass," *Physical Review D* (February 2020), APS. doi: 10.1103/PhysRevD.101.034510

Basak, S., A. Bazavov, C. Bernard, C. DeTar, L. Levkova, E. Freeland, S. Gottlieb, A. Torok, U. M. Heller, J. Laiho, J. Osborn, R. L. Sugar, D. Toussaint, R. S. Van de Water, and R. Zhou. "Lattice computation of the electromagnetic contributions to kaon and pion masses," *Physical Review D* (February 2020), APS. doi: 10.1103/PhysRevD.99.034503

Chu, W., Q. Zheng, O. V. Prezhdo, J. Zhao, and W. A. Saidi. "Low-frequency lattice phonons in halide perovskites explain high defect tolerance toward electron-hole recombination," *Science Advances* (February 2020), AAAS. doi: 10.1126/sciadv.aaw7453

Couch, S. M., M. L. Warren, and E. P. O'Connor. "Simulating Turbulence-Aided Neutrino-Driven Core-Collapse Supernova Explosions in One Dimension," *The Astrophysical Journal* (February 2020), IOP Science. doi: 10.3847/1538-4357/ab609e

Davies, C. T. H., C. DeTar, A. X. El-Khadra, E. Gámiz, S. Gottlieb, D. Hatton, A. S. Kronfeld, J. Laiho, G. P. Lepage, Y. Liu, P. B. Mackenzie, C. McNeile, E. T. Neil, T. Primer, J. N. Simone, D. Toussaint, R. S. Van de Water, and A. Vaquero (Fermilab Lattice, HPQCD, MILC Collaborations). "Hadronic-vacuum-polarization contribution to the muon's anomalous magnetic moment from four-flavor lattice QCD," *Physical Review D* (February 2020) APS. doi: 10.1103/PhysRevD.101.034512

Fairbanks, H. R., L. Jofre, G. Geraci, G. Laccarino, and A. Doostan. "Bi-fidelity approximation for uncertainty quantification and sensitivity analysis of irradiated particle-laden turbulence," *Journal of Computational Physics* (February 2020), Elsevier. doi: 10.1016/j.jcp.2019.108996

Hammond, K. D. “Parallel point defect identification in molecular dynamics simulations without post-processing: A compute and dump style for LAMMPS,” *Computer Physics Communications* (February 2020), Elsevier. doi: 10.1016/j.cpc.2019.106862

Hammond, K. D., D. Maroudas, and B. D. Wirth. “Theoretical Model of Helium Bubble Growth and Density in Plasma-Facing Metals,” *Scientific Reports* (February 2020), Springer Nature. doi: 10.1038/s41598-020-58581-8

Hong, S. “Quantum molecular dynamics simulations of chemical vapor deposition synthesis of MoS2 crystal assisted by H2 partial pressures,” *SPIE OPTO* (February 2020), San Francisco, California, SPIE, pp. 6. doi: 10.1117/12.2546359

Komperda, J., and F. Mashayek. “Filtered Density Function Implementation in a Discontinuous Spectral Element Method,” *Modeling and Simulation of Turbulent Mixing and Reaction* (February 2020), Springer Nature, pp. 169-180. doi: 10.1007/978-981-15-2643-5_7

Kumar, P., D. W. Kim, N. Rangnekar, H. Xu, E. O. Fetisov, S. Ghosh, H. Zhang, Q. Xiao, M. Shete, J. I. Siepmann, T. Dumitrica, B. McCool, M. Tsapatsis, and K. Andre Mkhoyan. “One-dimensional intergrowths in two-dimensional zeolite nanosheets and their effect on ultra-selective transport,” *Nature Materials* (February 2020), Springer Nature. doi: 10.1038/s41563-019-0581-3

Liu, J., A. Mohan, R. Kalia, A. Nakano, K. Nomura, P. Vashishta, and K. Yao. “Boltzmann machine modeling of layered MoS2 synthesis on a quantum annealer," *Computational Materials Science* (February 2020), Elsevier. doi: 10.1016/j.commatsci.2019.109429

Liu, X., R. Tom, X. Wang, C. Cook, B. Schatschneider, and N. Marom. “Pyrene-stabilized acenes as intermolecular singlet fission candidates: importance of exciton wave-function convergence,” *Journal of Physics: Condensed Matter* (February 2020), IOP Science. doi: 10.1088/1361-648x/ab699e

Stuke, A., C. Kunkel, D. Golze, M. Todorović, J.T. Margraf, K. Reuter, P Rinke, and H. Oberhofer. “Atomic structures and orbital energies of 61,489 crystal-forming organic molecules,” *Scientific Data* (February 2020), Springer Nature. doi: 10.6084/m9.figshare.11689347

Wei, W., E. A. Huerta, B. C Whitmore, J. C. Lee, S. Hannon, R. Chandar, D. A. Dale, K. L. Larson, D. A. Thilker, L. Ubeda, M. Boquien, M. Chevance, J. M. Diederik Kruijssen, A. Schruba, G. A. Blanc, Eand . Congiu. “Deep transfer learning for star cluster classification: I. application to the PHANGS–HST survey,” *Monthly Notices of the Royal Astronomical Society* (February 2020), Oxford University Press. doi: 10.1093/mnras/staa325

Withers, K. P., M. P. Moschetti, and E. M. Thompson. “A Machine Learning Approach to Developing Ground Motion Models From Simulated Ground Motions,” *Geophysical Research Letters* (February 2020), John Wiley and Sons. doi: 10.1029/2019GL086690

Yuan, H., M. A. Yildiz, E. Merzari, Y. Yu, A. Obabko, G. Botha, G. Busco, Y. A. Hassan, and D. T. Nguyen. “Spectral element applications in complex nuclear reactor geometries: Tet-to-hex meshing,” *Nuclear Engineering and Design* (February 2020), Elsevier. doi: 10.1016/j.nucengdes.2019.110422

Zhao, H., B. Wang, F. Liu, X. Yan, H. Wang, W. S. Leong, M. J. Stevens, P. Vashishta, A. Nakano, J. Kong, R. Kalia, and H. Wang. “Fluidic Flow Assisted Deterministic Folding of Van der Waals Materials,” *Advanced Functional Materials* (February 2020), John Wiley and Sons. doi: 10.1002/adfm.201908691

March

Cambareri, J. J., J. Fang, and I. A. Bolotnov. “Simulation scaling studies of reactor core two-phase flow using direct numerical simulation,” *Nuclear Engineering and Design* (March 2020), Elsevier. doi: 10.1016/j.nucengdes.2019.110435

Chen, H., J. M. Cole, G. B. G. Stenning, A. Yanguas-Gil, J. W. Elam, L. Stan, and Y. Gong. “Imaging Dye Aggregation in MK-2, N3, N749, and SQ-2 dye··· TiO2 Interfaces That Represent Dye-Sensitized Solar Cell Working Electrodes,” *Applied Energy Matter* (March 2020), ACS,. doi: 10.1021/acsaem.9b02001

Court, C. J., and J. M. Cole. “Magnetic and superconducting phase diagrams and transition temperatures predicted using text mining and machine learning,” *NPJ Computational Materials* (March 2020), Springer Nature. doi: 10.1038/s41524-020-0287-8

Dematties, D., G. K. Thiruvathukal, S. Rizzi, A. Wainselboim, and B. S. Zanutto,. “Towards High-End Scalability on Biologically-Inspired Computational Models,” *Computer Science: Faculty Publications and Other Works* (March 2020), vol. 36, Loyola University Chicago, pp. 497-506. doi: 10.3233/APC200077

Denoyelle, N., J. Tramm, K. Yoshli, S. Paramau, and P. Beckman. “NUMA-Aware Data Management for Neutron Cross Section Data in Continuous Energy Monte Carlo Neutron Transport Simulation,” *2020 Physics of Reactors Conference* (March 2020), Cambridge, UK, Nuclear Energy Group UCAM

Du, M., Y. S. G. Nashed, S. Kandel, D. Gürsoy, and C. Jacobsen. “Three dimensions, two microscopes, one code: Automatic differentiation for x-ray nanotomography beyond the depth of focus limit,” *Science Advances* (March 2020), AAAS. doi: 10.1126/sciadv.aay3700

Fang, J., M. K. Purser, C. Smith, R. Balakrishnan, I. A. Bolotnov, and K. E. Jansen. “Annular Flow Simulation Supported by Iterative In-Memory Mesh Adaptation,” *Nuclear Science and Engineering* (March 2020), Taylor and Francis Group. doi: 10.1080/00295639.2020.1743577

Feng, X., Y. Fu, and L.-C. Jin. “Lattice QCD calculation of the pion charge radius using a model-independent method,” *Physical Review D* (March 2020), APS. doi: 10.1103/PhysRevD.101.051502

Lai, J. K., E. Merzari, Y. A. Hassan, P. Fischer, and O. Marin. “Verification and validation of large eddy simulation with Nek5000 for cold leg mixing benchmark,” *Nuclear Engineering and Design* (March 2020), Elsevier. doi: 10.1016/j.nucengdes.2019.110427

Li, S., R. Paoli, and M. D’Mello. “Scalability of OpenFOAM Density-Based Solver with Runge-Kutta Temporal Discretization Scheme,” *Scientific Programming* (March 2020), Hindawi. doi: 10.1155/2020/9083620

Li, S., J. Qin, M. He, and R. Paoli. “Fast Evaluation of Aircraft Icing Severity Using Machine Learning Based on XGBoost,” *Aerospace* (March 2020), MDPI. doi: 10.3390/aerospace7040036

Liu, J., B. M. Sleeter, Z. Zhu, T. R. Loveland, T. Sohl, S. M. Howard, C. H. Key, T. Hawbaker, S. Liu, B. Reed, M. A. Cochrane, L. S. Heath, H. Jiang, D. T. Price, J. M. Chen, D. Zhou, N. B. Bliss, T. Wilson, J. Sherba, Q. Zhu, Y. Luo, and B. Poulter. “Critical land change information enhances the understanding of carbon balance in the United States,” *Global Change Biology* (March 2020), John Wiley and Sons. doi: 10.1111/gcb.15079

Mulligan, V. K. “The emerging role of computational design in peptide macrocycle drug discovery,” *Expert Opinion on Drug Discovery* (March 2020), Taylor and Francis Group. doi: 10.1080/17460441.2020.1751117

Shi, Y., O. V. Prezhdo, J. Zhao, and W. A. Saidi. “Iodine and Sulfur Vacancy Cooperation Promotes Ultrafast Charge Extraction at MAPbI3/MoS2 Interface,” *ACS Energy Letters* (March 2020), ACS. doi: 10.1021/acsenergylett.0c00485

Shriwise, P., J. Tramm, A. Davis, and P. Romano. “Towards CAD-Based Geometry Modelling with the Random Ray Method,” *2020 Physics of Reactors Conference* (March 2020), Cambridge, UK, Nuclear Energy Group UCAM

Tramm, J. R., A. R. Siegel, A. L. Lund, and P. K. Romano. “A Comparison of Stochastic Mesh Cell Volume Computation Strategies for the Random Ray Method of Neutral Particle Transport,” *2020 Physics of Reactors Conference* (March 2020), Cambridge, UK, Nuclear Energy Group UCAM

Yang, T. T., and W. A. Saidi. “Graphene Activation Explains the Enhanced Hydrogen Evolution on Graphene-Coated Molybdenum Carbide Electrocatalysts," *The Journal of Physical Chemistry Letters* (March 2020), ACS. doi: 10.1021/acs.jpclett.0c00615

April

Abhyankar, S., G. D. Betrie, D. A. Maldonado, L. C. McInnes, B. F. Smith, and H. Zhang. "PETSc DMNetwork: A Library for Scalable Network PDE-Based Multiphysics Simulations," *ACM Transactions on Mathematical Software* (April 2020), ACM. doi: 10.1145/3344587

Adams, C., M. Del Tutto, J. Asaadi, M. Bernstein, E. Church, R. Guenette, J. M. Rojas, H. Sullivan, and A. Tripathi. “Enhancing Neutrino Event Reconstruction with Pixel-Based 3D Readout for Liquid Argon Time Projection Chambers,” *Cornell University* (April 2020), IOP Science. doi: 10.1088/1748-0221/15/04/P04009

Andolina, C., P. Williamson, and W. Saidi. “Optimization and validation of a deep learning CuZr atomistic potential: Robust applications for crystalline and amorphous phases with near-DFT accuracy,” *The Journal of Chemical Physics* (April 2020), AIP. doi: 10.1063/5.0005347

Azevedo, G. M., J. L. Bez, P. Pavan, F. Z. Boito, and P. O. A. Navaux. “Requests sizes / O HPC applications on a supercomputer,” *2020: Proceedings of the 20th Regional School of High Performance Computing from Southern Brazil Undergraduate Forum* (April 2020), SBC. doi: 10.5753/eradrs.2020.10761

Bazavov, A., D. Bollweg, H.-T. Ding, P. Enns, J. Goswami, P. Hegde, O. Kaczmarek, F. Karsch, R. Larsen, S. Mukherjee, H. Ohno, P. Petreczky, C. Schmidt, S. Sharma, and P. Steinbrecher (HotQCD Collaboration). “Skewness, kurtosis, and the fifth and sixth order cumulants of net baryon-number distributions from lattice QCD confront high-statistics STAR data,” *Physical Review D* (April 2020), APS. doi: 10.1103/PhysRevD.101.074502

Blum, T., N. Christ, M. Hayakawa, T. Izubuchi, L. Jin, C. Jung, and C. Lehner. “Hadronic Light-by-Light Scattering Contribution to the Muon Anomalous Magnetic Moment from Lattice QCD,” *Physical Review Letters* (April 2020), APS. doi: 10.1103/PhysRevLett.124.132002

Chen, S. H., M. T. Young, J. Gounley, C. Stanley, and D. Bhowmik. “Distinct Structural Flexibility within SARS-CoV-2 Spike Protein Reveals Potential Therapeutic Targets,” *BIO Rx XIV* (April 2020), CSH. doi: 10.1101/2020.04.17.047548

Crespo, R., A. Arriaga, R. B. Wiringa, E. Cravo, A. Meccaa, and A. Deltuva. “Many-body effects in (p,pN) reactions within a unified approach,” *Physics Letters B* (April 2020), Elsevier. doi: 10.1016/j.physletb.2020.135355

Dharmawardhana, C. C., J. Zhou, M. Taylor, J. Miao, J. H. Perepezko, and H. Heinz. “Reactive modeling of Mo3Si oxidation and resulting silica morphology,” *Acta Materialia* (April 2020), Elsevier. doi: 10.1016/j.actamat.2020.01.048

Dematties, D., S. Rizzi, G. K. Thiruvathuka, M. D. Pérez, A. Wainselboim, and B. S. Zanutto. “A Computational Theory for the Emergence of Grammatical Categories in Cortical Dynamics,” *Frontiers in Neural Circuits* (April 2020), Frontiers Media. doi: 10.3389/fncir.2020.00012

Gandolfi, S., D. Lonardoni, A. Lovato, and M. Piarulli. “Atomic Nuclei From Quantum Monte Carlo Calculations With Chiral EFT Interactions,” *Frontiers in Physics* (April 2020), Frontiers Media. doi: 10.3389/fphy.2020.00117

Ho, P. J., C. Knight, and L. Young. “Extended x-ray emission times of clusters in intense x-ray pulses,” *Physical Review A* (April 2020), APS,. doi: 10.1103/PhysRevA.101.043413

Homerding, B., and J. Tramm. “Evaluating the Performance of the hipSYCL Toolchain for HPC Kernels on NVIDIA V100 GPUs,” *IWOCL '20: Proceedings of the International Workshop on OpenCL* (April 2020), ACM, pp. 1-7. doi: 10.1145/3388333.3388660

Iyer, K. P., J. D. Scheel, J. Schumacher, and K. R. Sreenivasan. “Classical 1/3 scaling of convection holds up to Ra = 1015,” *PNAS* (April 2020), National Academy of Sciences. doi: 10.1073/pnas.1922794117

Kaufmann, S., J. Simonis, S. Bacca, J. Billowes, M. L. Bissell, K. Blaum, B. Cheal, R. F. Garcia Ruiz, W. Gins, C. Gorges, G. Hagen, H. Heylen, A. Kanellakopoulos, S. Malbrunot-Ettenauer, M. Miorelli, R. Neugart, G. Neyens, W. Nörtershäuser, R. Sánchez, S. Sailer, A. Schwenk, T. Ratajczyk, L. V. Rodríguez, L. Wehner, and C. Wraith. “Charge Radius of the Short-Lived 68 Ni and Correlation with the Dipole Polarizability,” *Physical Review Letters* (April 2020), APS. doi: 10.1103/PhysRevLett.124.132502

Liu, Q., W. Jiang, M. J. Piernas-Muñoz, Y. Liu, Z. Yang, I. Bloom, T. L. Dzwiniel, Y. Li, K. Z. Pupek, and Z. Zhang. “Stabilized Electrode/Electrolyte Interphase by a Saturated Ionic Liquid Electrolyte for High-Voltage NMC532/Si-Graphite Cells,” *ACS Applied Materials and Interfaces* (April 2020), ACS. doi: 10.1021/acsami.0c06038

Maulik, R., A. Mohan, Be. Lusch, S. Madireddy, P. Balaprakash, and D. Livescu. “Time-series learning of latent-space dynamics for reduced-order model closure,” *Physica D: Nonlinear Phenomena* (April 2020), Elsevier. doi: 10.1016/j.physd.2020.132368

Morgan, H., R. T. Mills, and B. Smith. “Evaluation of PETSc on a Heterogeneous Architecture, the OLCF Summit System: Part 1: Vector Node Performance,” *OSTI.GOV* (April 2020), U.S. Department of Energy Office of Scientific and Technical Information. doi: 10.2172/1614879

Pastore, S., J. Carlson, S. Gandolfi, R. Schiavilla, and R. B. Wiringa. “Quasielastic lepton scattering and back-to-back nucleons in the short-time approximation,” *Physical Review C* (April 2020), APS. doi: 10.1103/PhysRevC.101.044612

Piarulli, M., I. Bombaci, D. Logoteta, A. Lovato, and R. B. Wiringa. “Benchmark calculations of pure neutron matter with realistic nucleon-nucleon interactions,” *Physical Review C* (April 2020), APS. doi: 10.1103/PhysRevC.101.045801

Renganathan, S. A., R. Maulik, and V. Rao. “Machine learning for nonintrusive model order reduction of the parametric inviscid transonic flow past an airfoil,” *Physics of Fluids* (April 2020), AIP. doi: 10.1063/1.5144661

Rosofsky, S. G., and E. A. Huerta. “Artificial neural network subgrid models of 2D compressible magnetohydrodynamic turbulence,” *Physical Review D* (April 2020), APS. doi: 10.1103/PhysRevD.101.084024

Saidi, W., W. Shadeed, and I. Castelli. “Machine-learning structural and electronic properties of metal halide perovskites using a hierarchical convolutional neural network,” *NPJ Computational Materials* (April 2020), Springer Nature. doi: 10.1038/s41524-020-0307-8

Schneider, A. S., E. O'Connor, E. Granqvist, A. Betranhandy, and S. M. Couch. “Equation of State and Progenitor Dependence of Stellar-Mass Black Hole Formation,” *The Astrophysical Journal* (April 2020), IOP Science. doi: 10.3847/1538-4357/ab8308

Sun, G., A. N. Alexandrova, and P. Sautet. “Structural Rearrangements of Subnanometer Cu Oxide Clusters Govern Catalytic Oxidation,” *ACS Catalysis* (April 2020), ACS. doi: 10.1021/acscatal.0c00824

Vorabbi, M., M. Gennari, P. Finelli, C. Giusti, and P. Navrátil. “Elastic Antiproton-Nucleus Scattering from Chiral Forces,” *Physical Review Letters* (April 2020), APS. doi: 10.1103/PhysRevLett.124.162501

Wang, X., R. Tom, X. Liu, D. Congreve, and N. Marom. "An energetics perspective on why there are so few triplet–triplet annihilation emitters," *Journal of Materials Chemistry C* (April 2020), Royal Science of Chemistry. doi: 10.1039/D0TC00044B

Wong, K., V. Zhdankin, D. A. Uzdensky, G. R. Werner, and M. C. Begelman. "First-principles Demonstration of Diffusive-Advective Particle Acceleration in Kinetic Simulations of Relativistic Plasma Turbulence," *The Astrophysical Journal Letters* (April 2020), IOP Science. doi: 10.3847/2041-8213/ab8122.

May

Aslangil, D., D. Livescu, and A. Banerjee. "Effects of Atwood and Reynolds numbers on the evolution of buoyancy-driven homogeneous variable-density turbulence," *Journal of Fluid Mechanics* (May 2020), Cambridge University Press. doi: 10.1017/jfm.2020.268

Benoit, A., V. Le Févre, P. Raghavan, Y. Robert, and H. Sun. "Design and Comparison of Resilient Scheduling Heuristics for Parallel Jobs," *2020 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)* (May 2020), New Orleans, LA, IEEE. doi: 10.1109/IPDPSW50202.2020.00099

Bertoni, C., J. Kwack, T. Applencourt, Y. Ghadar, B. Homerding, C. Knight, B. Videau, H. Zheng, V. Morozov, and S. Parker. "Performance Portability Evaluation of OpenCL Benchmarks across Intel and NVIDIA Platforms," *IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)* (May 2020), New Orleans, LA, IEEE. doi: 10.1109/IPDPSW50202.2020.00067

Choi, J., W. Wehde, R. Maulik, and S. Robinson. "What matters the most? Understanding individual tornado preparedness using machine learning," *Natural Hazards* (May 2020), Springer Nature. doi: 10.1007/s11069-020-04029-1

Feng, X., M. Gorchtein, L.-C. Jin, P.-X. Ma, and C.-Y. Seng. "First-Principles Calculation of Electroweak Box Diagrams from Lattice QCD," *Physical Review Letters* (May 2020), APS. doi: 10.1103/PhysRevLett.124.192002

Hager, R., C. S. Chang, N. M. Ferraro, and R. Nazikian. "Gyrokinetic understanding of the edge pedestal transport driven by resonant magnetic perturbations in a realistic divertor geometry," *Physics of Plasma* (May 2020), AIP. doi: 10.1063/1.5144445

Jin, Z., and H. Finkel. "Population Count on Intel® CPU, GPU and FPGA," *IEEE International Parallel and Distributed Processing Symposium Workshops* (May 2020), New Orleans, LA, IEEE. doi: 10.1109/IPDPSW50202.2020.00081

Jin, Z., V. Morozov, and H. Finkel. "A Case Study on the HACCmk Routine in SYCL on Integrated Graphics," *IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)* (May 2020), New Orleans, LA, IEEE. doi: 10.1109/IPDPSW50202.2020.00071

Kent, P. R. C., A. Annaberdiyev, A. Benali, M. C. Bennett, E. J. Landinez Borda, P. Doak, H. Hao, K. D. Jordan, J. T. Krogel, I. Kylanpaa, J. Lee, Y. Luo, F. D. Malone, C. A. Melton, L. Mitas, M. A. Morales, E. Neuscamman, F. A. Reboredo, B. Rubenstein, K. Saritas, S. Upadhyay, G. Wang, S. Zhang, and L. Zhao. "QMCPACK: Advances in the development, efficiency, and application of auxiliary field and real-space variational and diffusion quantum Monte Carlo," *The Journal of Chemical Physics* (May 2020), AIP. doi: 10.1063/5.0004860

Mukaddem, K. T., P. A. Chater, L. R. Devereux, O. K. Bahri, A. Jain, and J. M. Cole. "Dye-Anchoring Modes at the Dye···TiO2 Interface of N3- and N749-Sensitized Solar Cells Revealed by Glancing-Angle Pair Distribution Function Analysis," *The Journal of Physical Chemistry C* (May 2020), ACS. doi: 10.1021/acs.jpcc.0c02314

Rajak, P., K. Liu, A. Krishnamoorthy, R. K. Kalia, A. Nakano, K.-i. Nomura, S. C. Tiwari, and P. Vashishta. "Neural Network Molecular Dynamics at Scale," *2020 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)* (May 2020), New Orleans, LA, IEEE. doi: 10.1109/IPDPSW50202.2020.00167

Tom, R., T. Rose, I. Bier, H. O'Brien, Á. Vázquez-Mayagoitia, and N. Marom. "Genarris 2.0: A random structure generator for molecular crystals," *Computer Physics Communications* (May 2020), Elsevier. doi: 10.1016/j.cpc.2020.107170

Volpiani, P. S., P. S. Iyer, S. Pirozzoli, and J. Larsson. "Data-driven compressibility transformation for turbulent wall layers," *Physical Review Fluids* (May 2020), APS. doi: 10.1103/PhysRevFluids.5.052602

Wang, Y., and S. M. Day. "Effects of Off-Fault Inelasticity on Near-Fault Directivity Pulses," *Journal of Geophysical Research: Solid Earth* (May 2020), John Wiley and Sons. doi: 10.1029/2019JB019074

Wrobel, F., H. Park, C. Sohn, H.-W. Hsiao, J.-M. Zuo, H. Shin, H. N. Lee, P. Ganesh, A. Benali, P. R. C. Kent, O. Heinonen, and A. Bhattacharya. "Doped NiO: The mottness of a charge transfer insulator," *Physical Review B* (May 2020), APS. doi: 10.1103/PhysRevB.101.195128.

June

Bagchi, S., R. Kanungo, Y. K. Tanaka, H. Geissel, P. Doornenbal, W. Horiuchi, G. Hagen, T. Suzuki, N. Tsunoda, D. S. Ahn, H. Baba, K. Behr, F. Browne, S. Chen, M. L. Cortés, A. Estradé, N. Fukuda, M. Holl, K. Itahashi, N. Iwasa, G. R. Jansen, W. G. Jiang, S. Kaur, A. O. Macchiavelli, and S. Y. Matsumoto. "A two-neutron halo is unveiled in 29F," *Physical Review Letters* (June 2020), APS. doi: 10.1103/PhysRevLett.124.222504

Egger, A. T., L. Hörmann, A. Jeindl, M. Scherbela, V. Obersteiner, M. Todorović, P. Rinke, and O. T. Hofmann. "Charge Transfer into Organic Thin Films: A Deeper Insight through Machine-Learning-Assisted Structure Search," *Advanced Science* (June 2020), John Wiley and Sons. doi: 10.1002/advs.202000992

Feng, J., M. Acton, E. Baglietto, A. R. Kraus, and E. Merzari. "On the relevance of turbulent structures resolution for cross-flow in a helical-coil tube bundle," *Annals of Nuclear Energy* (June 2020), Elsevier. doi: 10.1016/j.anucene.2019.107298

Fiuza, F., G. F. Swadling, A. Grassi, H. G. Rinderknecht, D. P. Higginson, D. D. Ryutov, C. Bruulsema, R. P. Drake, S. Funk, S. Glenzer, G. Gregori, C. K. Li, B. B. Pollock, B. A. Remington, J. S. Ross, W. Rozmus, Y. Sakawa, A. Spitkovsky, S. Wilks, and H.-S. Park. "Electron acceleration in laboratory-produced turbulent collisionless shocks," *Nature Physics* (June 2020), Springer Nature. doi: 10.1038/s41567-020-0919-4

Fischer, P., M. Min, T. Rathnayake, S. Dutta, T. Kolev, V. Dobrev, J.-S. Camier, M. Kronbichler, T. Warburton, K. Swirydowicz, and J. Brown. "Scalability of high-performance PDE solvers," *Performance* (June 2020), SAGE Publications. doi: 10.1177/1094342020915762

Goc, K., S. Bose, and P. Moin. "Wall-Modeled Large Eddy Simulation of an Aircraft in Landing Configuration," *AIAA Aviation 2020 Forum* (June 2020), AIAA. doi: 10.2514/6.2020-3002

He, R., N. Li, B.-N. Lu, and D. Lee. "Superfluid condensate fraction and pairing wave function of the unitary Fermi gas," *Physical Review A* (June 2020), APS. doi: 10.1103/PhysRevA.101.063615

Hong, I., J. Ahn, H. Shin, H. Bae, H. Lee, A. Benali, and Y. Kwon. "Competition between Hückel's Rule and Jahn–Teller Distortion in Small Carbon Rings: A Quantum Monte Carlo Study," *The Journal of Physical Chemistry A* (June 2020), ACS. doi: 10.1021/acs.jpca.0c02577

Jin, Z. "The Rodinia Benchmark Suite in SYCL," *OSTI.GOV* (June 2020), U.S. Department of Energy Office of Scientific and Technical Information. doi: 10.2172/1631460

Kwack, J., C, Bertoni, B, Pham, and J. Larkin. "Performance of the RI-MP2 Fortran Kernel of GAMESS on GPUs via Directive-Based Offloading with Math Libraries," *International Workshop on Accelerator Programming Using Directives WACCPD 2019: Accelerator Programming Using Directives* (June 2020), Springer Nature, pp. 91-113. doi: 10.1007/978-3-030-49943-3_5

Lee, D. Y., W. Lin, W. and M. R. Petersen. "Wintertime Arctic Oscillation and North Atlantic Oscillation and their impacts on the Northern Hemisphere climate in E3SM," *Climate Dynamics* (June 2020), Springer Nature. doi: 10.1007/s00382-020-05316-0

Leman, J. K., B. D. Weitzner, and R. Bonne. "Macromolecular modeling and design in Rosetta: recent methods and frameworks," *Nature Methods* (June 2020), Springer Nature. doi: 10.1038/s41592-020-0848-2

Liu, Z., R. Lewis, R. Kettimuthu, K. Harms, P. H. Carns, N. Rao, I. Foster, and M. E. Papka. "Characterization and identification of HPC applications at leadership computing facility," *ICS '20: Proceedings of the 34th ACM International Conference on Supercomputing* (June 2020), Barcelona, Spain, ACM. doi: 10.1145/3392717.3392774

Rao, V., R. Maulik, E. Constantinescu, and M. Anitescu. "A Machine-Learning-Based Importance Sampling Method to Compute Rare Event Probabilities," *International Conference on Computational Science ICCS 2020: Computational Science – ICCS 2020* (June 2020), Springer Nature. doi: 10.1007/978-3-030-50433-5_14

Sahasrabudhe, D., and M. Berzins. "Improving Performance of the Hyper Iterative Solver for Uintah Combustion Codes on Manycore Architectures Using MPI Endpoints and Kernel Consolidation," *Lecture Notes in Computer Science* (June 2020), Springer Nature. doi: 10.1007/978-3-030-50371-0_13

Siriwardane, E. M. D., R. P. Joshi, N. Kumar, and D. Çakır. "Revealing the Formation Energy–Exfoliation Energy–Structure Correlation of MAB Phases Using Machine Learning and DFT," *ACS Applied Materials and Interfaces* (June 2020), ACS. doi: 10.1021/acsami.0c03536

Sivaraman, G., N. E. Jackson, B. Sanchez-Lengeling, Á. Vázquez-Mayagoitia, A. Aspuru-Guzik, V. Vishwanath, and J. J. de Pablo. "A machine learning workflow for molecular analysis: application to melting points," *Machine Learning: Science and Technology* (June 2020), IOP Science. doi: 10.1088/2632-2153/ab8aa3

Tandon, S., I. Marincic, H. Hoffmann, and E. Johnsen. "Enabling power-performance balance with transprecision calculations for extreme-scale computations of turbulent flows," *AIAA Aviation 2020* (June 2020), AIAA. doi: 10.2514/6.2020-2922

Zheng, W. "Research trend of large-scale supercomputers and applications from the TOP500 and Gordon Bell Prize," *Science China Information Sciences* (June 2020), Springer Nature. doi: 10.1007/s11432-020-2861-0

July

Bott, A. F. A., P. Tzeferacos, L. Chen, C. A. J. Palmer, A. Rigby, A. R. Bell, R. Bingham, A. Birkel, C. Graziani, D. H. Froula, J. Katz, M. Koenig, M. W. Kunz, C. Li, J. Meinecke, F. Miniati, R. Petrasso, H.-S. Park, B. A. Remington, B. Reville, J. S. Ross, D. Ryu, D. Ryutov, F. H. Séguin, T. G. White, A. A. Schekochihin, D. Q. Lamb, and G. Gregori. "Time-resolved turbulent dynamo in a laser plasma," *PNAS* (July 2020), National Academy of Sciences. doi: 10.1073/pnas.2015729118

Campbell, L. T., H. P. Freund, J. R. Henderson, B. W. J. McNeil, P. Traczykowski, and P. J. M. van der Slot. "Analysis of ultra-short bunches in free-electron lasers," *New Journal of Physics* (July 2020), IOP Science. doi: 10.1088/1367-2630/ab9850

Gao, X., I. Bermejo-Moreno, and J. Larsson. "Parametric numerical study of passive scalar mixing in shock turbulence interaction," *Journal of Fluid Mechanics* (July 2020), Cambridge University Press. doi: 10.1017/jfm.2020.292

Jeong, H., X. S. Asay-Davis, A. K. Turner, D. S. Comeau, S. F. Price, R. P. Abernathey, M. Veneziani, M. R. Petersen, M. J. Hoffman, M. R. Mazloff, and T. D. Ringler. "Impacts of Ice-Shelf Melting on Water-Mass Transformation in the Southern Ocean from E3SM Simulations," *Journal of Climate* (July 2020), AMS. doi: 10.1175/JCLI-D-19-0683.1

Khan, A., E. A. Huerta, and A. Das. "Physics-inspired deep learning to characterize the signal manifold of quasi-circular, spinning, non-precessing binary black hole mergers," *Physics Letters B* (July 2020), Elsevier. doi: 10.1016/j.physletb.2020.135628

Klawonn, A., M. Lanser, M. Uran, Ol. Rheinbach, S. Köhler, J. Schröder, L. Scheunemann, D. Brands, D. Balzani, A. Gandhi, G. Wellein, M. Wittmann, O. Schenk, and R. Janalík. "EXASTEEL: Towards a Virtual Laboratory for the Multiscale Simulation of Dual-Phase Steel Using High-Performance Computing," *Software for Exascale Computing* (July 2020), vol. SPPEXA 2016-2019, Springer, pp. 351-404. doi: 10.1007/978-3-030-47956-5_13

Lee, C.-W., J. A. Stewart, R. Dingreville, S. M. Foiles, and A. Schleife. "Multiscale simulations of electron and ion dynamics in self-irradiated silicon," *Physical Review B* (July 2020), APS. doi: 10.1103/PhysRevB.102.024107

Li, Y., A, Eshein, R, Virk, A. Eid, D. VanDerway, J. Frederick, R. Bleher, I. Szleifer, V. Dravid, and V. Backman. "Nanoscale Chromatin Imaging and Analysis (nano-ChIA) platform bridges 4-D chromatin organization with molecular function," *Microscopy and Microanalysis 2020 Proceedings* (July 2020), Microscopy Society of America. doi: 10.1017/S1431927620016785

Ma, H., M. Govoni, and G. Galli. "Quantum simulations of materials on near-term quantum computers," *NPJ* (July 2020), Springer Nature. doi: 10.1038/s41524-020-00353-z

Mallick, T., P. Balaprakash, E. Rask, and J. Macfarlane. "Graph-Partitioning-Based Diffusion Convolutional Recurrent Neural Network for Large-Scale Traffic Forecasting," *Machine Learning* (July 2020), SAGE Publications. doi: 10.1177/0361198120930010

Maulik, R., N. A. Garland, J. W. Burby, X.-Z. Tang, and P. Balaprakash. "Neural network representability of fully ionized plasma fluid model closures," *Physics of Plasma* (July 2020), AIP. doi: 10.1063/5.0006457

Mezzacappa, A., P. Marronetti, R. E. Landfield, E. J. Lentz, K. N. Yakunin, S. W. Bruenn, W. R. Hix, O. E. Bronson Messer, E. Endeve, J. M. Blondin, and J. A. Harris. "Gravitational-wave signal of a core-collapse supernova explosion of a 15 Solar mass star," *Physical Review D* (July 2020), APS. doi: 10.1103/PhysRevD.102.023027

Oliveira. M. J. T., N. Papior, Y. Pouillon, V. Blum, E. Artacho, D. Caliste, F. Corsetti, S. de Gironcoli, A. M. Elena, A. García, V. M. García-Suárez, L. Genovese, W. P. Huhn, G. Huhs, S. Kokott, E. Küçükbenli, A. H. Larsen, A. Lazzaro, I. V. Lebedeva, Y. Li, D. López-Durán, P. López-Tarifa, M. Lüders, M. A. L. Marques, J. Minar, S. Mohr, A. A. Mostofi, A. O'Cais, M. C. Payne, T. Ruh, D. G. A. Smith, J. M. Soler, D. A. Strubbe, N. Tancogne-Dejean, D. Tildesley, M. Torrent, and V. W.-z. Yu. "The CECAM electronic structure library and the modular software development paradigm featured," *Journal of Chemical Physics* (July 2020), AIP. doi: 10.1063/5.0012901

Oltchik, Y., and O. Schwartz. "Network Partitioning and Avoidable Contention," *SPAA* (July 2020), ACM, pp. 563-565. doi: 10.1145/3350755.3400242

Oqmhula, K., K. Hongo, R. Maezono, and T. Ichibha. "Ab Initio Evaluation of Complexation Energies for Cyclodextrin-Drug Inclusion Complexes," *ACS Omega* (July 2020), ACS. doi: 10.1021/acsomega.0c01059

Sivaraman, G., A. Na. Krishnamoorthy, M. Baur, C. Holm, M. Stan, G. Csányi, C. Benmore, and Á. Vázquez-Mayagoitia. "Machine-learned interatomic potentials by active learning: amorphous and liquid hafnium dioxide," *Computational Materials* (July 2020), Springer Nature. doi: 10.1038/s41524-020-00367-7

Wang, T., and A. Burrows. "Generalized Kompaneets formalism for inelastic neutrino-nucleon scattering in supernova simulations," *Physical Review D* (July 2020), APS. doi: 10.1103/PhysRevD.102.023017

Warren, M. L., S. M. Couch, P. O'Connor, and V. Morozova. "Constraining Properties of the Next Nearby Core-collapse Supernova with Multimessenger Signals," *The Astrophysical Journal* (July 2020) IOP Science. doi: 10.3847/1538-4357/ab97b7

Yao, Y., and Y. Kanai. "Temperature dependence of nuclear quantum effects on liquid water via artificial neural network model based on SCAN meta-GGA functional," *The Journal of Chemical Physics* (July 2020), AIP. doi: 10.1063/5.0012815

August

Abi, B., R. Acciarri, M. A. Acero, G. Adamov, D. Adams, M. Adinolfi, Z. Ahmad, J. Ahmed, T. Alion, S. Alonso-Monsalve et al. "Volume I. Introduction to DUNE," *Journal of Instrumentation* (August 2020) IOP Science. doi: 10.1088/1748-0221/15/08/T08008

Abi, B., R. Acciarri, M. A. Acero, G. Adamov, D. Adams, M. Adinolfi, Z. Ahmad, J. Ahmed, T. Alion, S. Alonso-Monsalve et al. "Volume III. DUNE far detector technical coordination," *Journal of Instrumentation* (August 2020) IOP Science. doi: 10.1088/1748-0221/15/08/T08009

Abi, B., R. Acciarri, M. A. Acero, G. Adamov, D. Adams, M. Adinolfi, Z. Ahmad, J. Ahmed, T. Alion, S. Alonso-Monsalve et al. "Volume IV. The DUNE far detector single-phase technology," *Journal of Instrumentation* (August 2020), IOP Science. doi: 10.1088/1748-0221/15/08/T08010

Cambareri, J. J., J. Fang, and I. A. Bolotnov. "Interface capturing simulations of bubble population effects in PWR subchannels," *Nuclear Engineering and Design* (August 2020), Elsevier. doi: 10.1016/j.nucengdes.2020.110709

Court, C. J., B. Yildirim, A. Jain, and J. M. Cole. "3-D Inorganic Crystal Structure Generation and Property Prediction via Representation Learning," *Journal of Chemical Information and Modeling* (August 2020), ACS. doi: 10.1021/acs.jcim.0c00464

De Souza, F. A. L., G. Sivaraman, M. Fyta, R. H. Scheicher, W. L. Scopel, and R. G. Amorim. "Electrically sensing Hachimoji DNA nucleotides through a hybrid graphene/h-BN nanopore," *Nanoscale* (August 2020). Royal Society of Chemistry. doi: 10.1039/D0NR04363J

Elliott, S. N., K. B. Moore, A. V. Copan, M. Keçeli, C. Cavallotti, Y. Georgievskii, H. F.Schaefer, and S. J. Klippenstein. "Automated theoretical chemical kinetics: Predicting the kinetics for the initial stages of pyrolysis," *Proceedings of the Combustion Institute* (August 2020), Elsevier. doi: 10.1016/j.proci.2020.06.019

Huang, S., and J. M. Cole. "A database of battery materials auto-generated using ChemDataExtractor," *Scientific Data* (August 2020), Springer Nature. doi: 10.1038/s41597-020-00602-2

King, G. B., L. Andreoli, S. Pastore, M. Piarulli, R. Schiavilla, R. B. Wiringa, J. Carlson, and S. Gandolfi. "Chiral effective field theory calculations of weak transitions in light nuclei," *Physical Review C* (August 2020), APS. doi: 10.1103/PhysRevC.102.025501

Lau, C. K., D. P. Fulton, J. Bao, Z. Lin, S. Dettrick, M. Binderbauer, T. Tajima, and L. Schmitz. "Electrostatic quasi-neutral formulation of global cross-separatrix particle simulation in field-reversed configuration geometry," *Physics of Plasma* (August 2020). doi: 10.1063/5.0012439

Marques, O., and A. Barker. "Training Efforts in the Exascale Computing Project," *Computing in Science and Engineering* (Augus 2020), IEEE. doi: 10.1109/MCSE.2020.3010596

Maulik, R., B. Lusch, and P. Balaprakash. "Non-autoregressive time-series methods for stable parametric reduced-order models," *Physics of Fluids* (August 2020), AIP. doi: 10.1063/5.0019884

Saini, N., I. A. Bolotnov. "Interface capturing simulations of droplet interaction with spacer grids under DFFB conditions," *Nuclear Engineering and Design* (August 2020),. doi: 10.1016/j.nucengdes.2020.110685

Wu, X., and A. Sharma. "Artefacts of finite span in vortex-induced vibration simulations," *Applied Ocean Research* (August 2020), Elsevier. doi: 10.1016/j.apor.2020.102265

Yakushin, I., K. Mehta, J. Chen, M. Wolf, I. Foster, S. Klasky, and T. Munson. "Feature-Preserving Lossy Compression for In Situ Data Analysis," *ICPP Workshops '20: 49th International Conference on Parallel Processing - ICPP: Workshops* (August 2020). doi: 10.1145/3409390.3409400

Yu, Y., E. R. Shemon, and T. K. Kim. "Evaluation of hot channel factor for sodium-cooled fast reactors with multi-physics toolkit," *Nuclear Engineering and Design* (August 2020), Elsevier. doi: 10.1016/j.nucengdes.2020.110704

Zhu, E., W. Xue, S. Wang, X. Yan, J. Zhou, Y. Liu, J. Cai, E. Liu, Q. Jia, X. Duan, Y. Li, H. Heinz, and Y. Huang. "Enhancement of oxygen reduction reaction activity by grain boundaries in platinum nanostructures," *Nano Research* (August 2020), Springer Nature. doi: 10.1007/s12274-020-3007-2

September

Georgakoudis, G., J. Doerfert, I. Laguna, and T. R. W. Scogland. "FAROS: A Framework to Analyze OpenMP Compilation Through Benchmarking and Compiler Optimization Analysis," *IWOMP 2020: OpenMP: Portable Multi-Level Parallelism on Modern Systems* (September 2020), Springer Nature. doi: 10.1007/978-3-030-58144-2_1

Gin, C., B. Lusch, S. L. Brunton, and J. N. Kutz. "Deep learning models for global coordinate transformations that linearise PDEs," *European Journal of Applied Mathematics* (September 2020), Cambridge University Press. doi: 10.1017/S0956792520000327

Huang, S., S. Hong, Y. Su, Y. Jiang, S. Fukushima, T. M. Gill, N. E. D. Yilmaz, S. C. Tiwari, K. Nomura, R. K. Kalia, A. Nakano, F. Shimojo, P. Vashishta, M. Chen, and X. Zheng. "Enhancing combustion performance of nano-Al/PVDF composites with ß-PVDF," *Combustion and Flame* (September 2020), Elsevier. doi: 10.1016/j.combustflame.2020.06.011

Huhn, W. P., B. Lange, V. W.-z. Yu, M. Yoon, and V. Blum. "GPU acceleration of all-electron electronic structure theory using localized numeric atom-centered basis functions," *Computer Physics Communications* (September 2020), Elsevier. doi: 10.1016/j.cpc.2020.107314

Jana, M. K., R. Song, H. Liu, D. R. Khanal, S. M. Janke, R. Zhao, C. Liu, Z. V. Vardeny, V. Blum, and D. B. Mitzi. "Organic-to-inorganic structural chirality transfer in a 2D hybrid perovskite and impact on Rashba-Dresselhaus spin-orbit coupling," *Nature Communications* (September 2020) Springer Nature. doi: 10.1038/s41467-020-18485-7

Kanwar, G., M. S. Albergo, D. Boyda, K. Cranmer, D. C. Hackett, S. Racanière, D. J. Rezende, and P. E. Shanahan. "Equivariant Flow-Based Sampling for Lattice Gauge Theory," *Physical Review Letters* (September 2020). doi: 10.1103/PhysRevLett.125.121601

Lao, K. U., Y. Yang, and R. DiStasio "On the Higher-Order Static Polarizabilities and Dispersion Coefficients of the Fullerenes: An Ab Initio Study," *Europe PMC* (September 2020), Elixir. doi: 10.26434/chemrxiv.12940652.v1

Lovato, A., J. Carlson, S. Gandolfi, N. Rocco, and R. Schiavilla. "Ab Initio Study of ($\nu\ell$, ℓ^-) and ($\bar{}\nu\ell$, ℓ^+) Inclusive Scattering in 12C : Confronting the MiniBooNE and T2K CCQE Data," *Physical Review X* (September 2020), APS. doi: 10.1103/PhysRevX.10.031068

Sharma, H., and E. Jennings. "Bayesian neural networks at scale: a performance analysis and pruning study," *The Journal of Supercomputing* (September 2020), Springer Nature. doi: 10.1007/s11227-020-03401-z

Vartanyan, D., and A. Burrows. "Gravitational Waves from Neutrino Emission Asymmetries in Core-collapse Supernovae," *The Astrophysical Journal* (September 2020), IOP Science. doi: 10.3847/1538-4357/abafac

Zhang, N., B. Toonen, X.-H. Sun, and W. Allcock. "Performance Modeling and Evaluation of a Production Disaggregated Memory System," *MEMSYS 2020: The International Symposium on Memory Systems* (September 2020), ACM. doi: 10.1145/3422575.3422795

Zhang, X., S. R. Stroberg, P. Navrátil, C. Gwak, J. A. Melendez, R. J. Furnstahl, and J. D. Holt. “Ab Initio Calculations of Low-Energy Nuclear Scattering Using Confining Potential Traps,” *Physical Review Letters* (September 2020), APS. doi: 10.1103/PhysRevLett.125.112503.

October

Ahn, J., I. Hong, G. Lee, H. Shin, A. Benali, and Y. Kwon. “Energetic Stability of Free-standing and Metal-Supported Borophenes: Quantum Monte Carlo and Density Functional Theory Calculations,” *The Journal of Physical Chemistry C* (October 2020), ACS Publications. doi: 10.1021/acs.jpcc.0c06883

Atkinson, M. C., W. H. Dickhoff, M. Piarulli, A. Rios, and R. B. Wiringa. “Reexamining the relation between the binding energy of finite nuclei and the equation of state of infinite nuclear matter,” *Physical Review C* (October 2020), APS. doi: 10.1103/PhysRevC.102.044333

Attar, A. R., H.-T. Chang, A. Britz, X. Zhang, M.-F. Lin, A. Krishnamoorthy, T. Linker, David Fritz, D. M. Neumark, R. K. Kalia, A. Nakano, P. Ajayan, P. Vashishta, U. Bergmann, and S. R. Leone. “Simultaneous Observation of Carrier-Specific Redistribution and Coherent Lattice Dynamics in 2H-MoTe2 with Femtosecond Core-Level Spectroscopy,” *ACS NANO* (October 2020), ACS Publications. doi: 10.1021/acsnano.0c06988

Di Renzo, M., L. Fu, and J. Urzay. “HTR solver: An open-source exascale-oriented task-based multi-GPU high-order code for hypersonic aerothermodynamics,” *Computer Physics Communications* (October 2020), Elsevier. doi: 10.1016/j.cpc.2020.107262

Fang, J., D. Shaver, and E. Merzari. “Modeling the Effect of Spacer Grid and Mixing Vanes in Coupled CFD Simulations of Small Modular Reactors,” *ASME 2020 Fluids Engineering Division Summer Meeting collocated with the ASME 2020 Heat Transfer Summer Conference and the ASME 2020 18th International* (October 2020), ASME. doi: 10.1115/FEDSM2020-20370

Gelman, S., P. A. Romero, and A. Gitter. “Neural networks to learn protein sequence-function relationships from deep mutational scanning data,” *Europe PMC* (October 2020), Elixir. doi: 10.1101/2020.10.25.353946

Hammond, K. D., B. F. Lee, I. V. Naeger, W. Widanagamaachchi, L.-T. Lo, D. Maroudas, and B. D. Wirth. “Corrigendum and Addendum: Helium flux effects on bubble growth and surface morphology in plasma-facing tungsten from large-scale molecular dynamics simulations,” *Nuclear Fusion* (October 2020), IOP Science. doi: 10.1088/1741-4326/abb2d3

He, Y., P. Balaprakash, and Y. Li. “FIdelity: Efficient Resilience Analysis Framework for Deep Learning Accelerators,” *2020 53rd Annual IEEE/ACM International Symposium on Microarchitecture* (October 2020), Athens, Greece, IEEE. doi: 10.1109/MICRO50266.2020.00033

Herman, B., F. Samsel, A. Bares, S. Johnson, G. Abram, and D. F. Keefe. “Printmaking, Puzzles, and Studio Closets: Using Artistic Metaphors to Reimagine the User Interface for Designing Immersive Visualizations,” *2020 IEEE VIS Arts Program* (October 2020), Salt Lake City, UT, IEEE. doi: 10.1109/visap51628.2020.00009

Hwang, Y., and M. Lee. “The mean logarithm emerges with self-similar energy balance,” *Journal of Fluid Mechanics* (October 2020), Cambridge University Press. doi: 10.1017/jfm.2020.730

Järvi, J., P. Rinke, and M. Todorović. “Detecting stable adsorbates of (1S)-camphor on Cu(111) with Bayesian optimization,” *Beilstein Journal of Nanotechnology* (October 2020), Beilstein-Institut. doi: 10.3762/bjnano.11.140

Jiang, Y., Y. Jiang, S. Deng, S. Hong, S. Tiwari, H. Chen, K.-i. Nomura, R. K. Kalia, A. Nakano, P. Vashishta, M. R. Zachariah, and X. Zheng. “Synergistically Chemical and Thermal Coupling between Graphene Oxide and Graphene Fluoride for Enhancing Aluminum Combustion,” *Applied Materials and Interfaces* (October 2020), ACS Publications. doi: 10.1021/acsami.9b20397

Kallala, H., F. Quéré, and H. Vincenti. “Techniques to generate intense isolated attosecond pulses from relativistic plasma mirrors,” *Physical Review Research* (October 2020), APS. doi: 10.1103/PhysRevResearch.2.043007

King, G., S. Pastore, and M. Piarulli. “Investigation of electroweak processes in light nuclei with Chiral Effective Field Theory,” *Bulletin of the American Physical Society* (October 2020), APS

Kononov, A., and A. Schleife. “Pre-equilibrium stopping and charge capture in proton-irradiated aluminum sheets,” *Physical Review B* (October 2020), APS. doi: 10.1103/PhysRevB.102.165401

Kraus, A., E. Merzari, T. Norddine, O. Marin, and S. Benhamadouche. “Direct Numerical Simulation of Fluid Flow in a 5x5 Square Rod Bundle Using Nek5000,” *International Conference on Nuclear Engineering* (October 2020), American Society of Mechanical Engineers. doi: 10.1115/ICONE2020-16643

Li, Y., L. Fu, and N. A. Adams. “A low-dissipation shock-capturing framework with flexible nonlinear dissipation control,” *Journal of Computational Physics* (October 2020), Elsevier. doi: 10.1016/j.jcp.2020.109960

Linker, T., S. Tiwari, S. Fukushima, R. K. Kalia, A. Krishnamoorthy, A. Nakano, K.-i. Nomura, K. Shimamura, F. Shimojo, and P. Vashishta. “Optically Induced Three-Stage Picosecond Amorphization in Low-Temperature SrTiO3,” *The Journal of Physical Chemistry Letters* (October 2020), ACS Publications. doi: 10.1021/acs.jpclett.0c02873

Liu, J., and B. Sleeter. “Simulated CONUS-wide 1-km resolution 1971-2015 ecosystem carbon variables from the IBIS model,” *USGS Science Data Catalog* (October 2020), U.S. Geological Survey. doi: 10.5066/P989HRGF

Lu, Q., C. Sohn, G. Hu, X. Gao, M. F. Chisholm, I. Kylänpää, J. T. Krogel, P. R. C. Kent, O. Heinonen, P. Ganesh, and H. N. Lee. “Metal–insulator transition tuned by oxygen vacancy migration across TiO2/VO2 interface,” *Scientific Reports* (October 2020), Springer Nature. doi: 10.1038/s41598-020-75695-1

Mallick, T., P. Balaprakash, E. Rask, and J. Macfarlane. “Graph-Partitioning-Based Diffusion Convolutional Recurrent Neural Network for Large-Scale Traffic Forecasting,” *Machine Learning* (July 2020), SAGE Publications. doi: 10.1177/0361198120930010

Malone, F. D., A. Benali, M. A. Morales, M. Caffarel, P. R. C. Kent, and L. Shulenburger. “Systematic comparison and cross-validation of fixed-node diffusion Monte Carlo and phaseless auxiliary-field quantum Monte Carlo in solids,” *Physical Review B* (October 2020), APS. doi: 10.1103/PhysRevB.102.161104

Maulik, R., K. Fukami, N. Ramachandra, K. Fukagata, and K. Taira. “Probabilistic neural networks for fluid flow surrogate modeling and data recovery,” *Physical Review Fluids* (October 2020), APS. doi: 10.1103/PhysRevFluids.5.104401

Morgan, H., P. Sanan, M. Knepley, and R. T. Mills. “Understanding performance variability in standard and pipelined parallel Krylov solvers,” *The International Journal of High Performance Computing Applications* (October 2020), SAGE Publications. doi: 10.1177/1094342020966835

Mulligan, V. K., C. S. Kang, M. R. Sawaya, S. Rettie, X. Li, I. Antselovich, T. W. Craven, A. M. Watkins, J. W. Labonte, F. DiMaio, T. O. Yeates, and D. Baker. “Computational design of mixed chirality peptide macrocycles with internal symmetry,” *Protein Science* (October 2020), John Wiley and Sons. doi: 10.1002/pro.3974

Oldfield, R. A., M. Wolf, and R. B. Brightwell. “ECP Capability Assessment Report for Software Technologies,” *OSTI.GOV* (October 2020), U.S. Department of Energy Office of Scientific and Technical Information. doi: 10.2172/1717885

Qian, Y., X. Li, A. R. Harutyunyan, G. Chen, Y. Rao, and H. Chen. “Herzberg–Teller Effect on the Vibrationally Resolved Absorption Spectra of Single-Crystalline Pentacene at Finite Temperatures,” *The Journal of Physical Chemistry A* (October 2020), ACS Publications. doi: 10.1021/acs.jpca.0c07896

Remolina González, J. D., K. Sharon, B. Reed, N. Li, G. Mahler, L. E. Bleem, M. Gladders, A. Niemiec, A. Acebron, and H. Child. "Efficient Mass Estimate at the Core of Strong Lensing Galaxy Clusters Using the Einstein Radius," *The Astrophysical Journal* (October 2020), IOP Science. doi: 10.3847/1538-4357/abb2a1

Schultz, W. C., L. Bildsten, and Y.-F. Jiang. "Convectively Driven 3D Turbulence in Massive Star Envelopes. I. A 1D Implementation of Diffusive Radiative Transport," *The Astrophysical Journal* (October 2020), IOP Science. doi: 10.3847/1538-4357/abb405

Schumann, J.-E., S. Toosi, and J. Larsson. "Assessment of Grid Anisotropy Effects on Large-Eddy-Simulation Models with Different Length Scales," *AIAAJ Aeronautics/Astronautics* (October 2020), AIAA. doi: 10.2514/1.J059576

Seylabi, E., D. Restrepo, R. Taborda, and D. Asimaki. "Deterministic ground motion simulations with shallow crust nonlinearity at Garner Valley in Southern California," *Earthquake Engineering and Structural Dynamics/Early View* (October 2020), John Wiley and Sons. doi: 10.1002/eqe.3360

Song, B., Y. Yang, M. Rabbani, T. T. Yang, K. He, X. Hu, Y. Yuan, P. Ghildiyal, V. P. Dravid, M. R. Zachariah, W. A. Saidi, Y. Liu, and R. Shahbazian-Yassar. "In Situ Oxidation Studies of High-Entropy Alloy Nanoparticles," *ACS Nano* (October 2020), ACS Publications. doi: 10.1021/acsnano.0c05250

Totorica, S. R., M. Hoshino, T. Abel, and F. Fiuza. "Nonthermal electron and ion acceleration by magnetic reconnection in large laser-driven plasmas," *Physics of Plasma* (October 2020), AIP. doi: 10.1063/5.0021169

Wang, G., A. Annaberdiyev, and L. Mitas. "Binding and excitations in SixHy molecular systems using quantum Monte Carlo," *The Journal of Chemical Physics* (October 2020), AIP. doi: 10.1063/5.0022814

Weerasinghe, A., L. Hu, K. D. Hammond, B. D. Wirth, and D. Maroudas. "Non-dilute helium-related defect interactions in the near-surface region of plasma-exposed tungsten," *Journal of Applied Physics* (October 2020), AIP. doi: 10.1063/5.0023356

November

Allen, B. S., M. A. Ezell, D. Jacobsen, C. Lueninghoener, P. Peltz, E. Roman, and J. Lowell Wofford. "Modernizing the HPC System Software Stack," *SIGHPC Systems Professionals Workshop* (November 2020), Atlanta, GA, SYSPROS

Allen-Dumas, M. R., A. N. Rose, J. R. New, O. A. Omitaomu, J. Yuan, M. L. Branstetter, L. M. Sylvester, M. B. Seals, T. M. Carvalhaes, M. B. Adams, M. S. Bhandari, S. S. Shrestha, J. Sanyal, A. S. Berres, C. P. Kolosna, K. S. Fu, and A. C. Kahl. "Impacts of the morphology of new neighborhoods on microclimate and building energy," *Renewable and Sustainable Energy Reviews* (November 2020), Elsevier. doi: 10.1016/j.rser.2020.110030

Benali, A., K. Gasperich, K. D. Jordan, T. Applencourt, Y. Luo, M. C. Bennett, J. T. Krogel, L. Shulenburger, P. R. C. Kent, P.-F. Loos, A. Scemama, and M. Caffarel. "Toward a systematic improvement of the fixed-node approximation in diffusion Monte Carlo for solids—A case study in diamond," *Journal of Chemical Physics* (November 2020), AIP. doi: 10.1063/5.0021036

Benali, A., H. Shin, and O. Heinonen. "Quantum Monte Carlo benchmarking of large noncovalent complexes in the L7 benchmark set," *The Journal of Chemical Physics* (November 2020), AIP. doi: 10.1063/5.0026275

Bertagna, L., O. Guba, M. A. Taylor, J. G. Foucar, J. M. Larkin, A. M. Bradley, S. Rajamanickam, and A. G. Salinger. "A performance-portable nonhydrostatic atmospheric dycore for the energy exascale earth system model running at cloud-resolving resolutions," *SC '20: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis* (November 2020), no. 92, Atlanta, GA, IEEE, pp. 1-14

Chase, S., B. J. Winjum, F. S. Tsung, K. G. Miller, D. Hinkel, and W. B. Mori. "Stimulated Raman backscatter in the kinetic regime of lasers with orbital angular momentum," *62nd Annual Meeting of the APS Division of Plasma Physics* (November 2020), APS

Cruz-Torres, R. D. Lonardoni, R. Weiss, M. Piarulli, N. Barnea, D. W. Higinbotham, E. Piasetzky, A. Schmidt, L. B. Weinstein, R. B. Wiringa, and O. Hen. "Many-body factorization and position–momentum equivalence of nuclear short-range correlations," *Nature Physics* (November 2020), Springer Nature. doi: 10.1038/s41567-020-01053-7

Finkel, H., A. McCaskey, T. Popoola, D. Lyakh, and J. Doerfert. "Really Embedding Domain-Specific Languages into C++," *2020 IEEE/ACM 6th Workshop on the LLVM Compiler Infrastructure in HPC (LLVM-HPC) and Workshop on Hierarchical Parallelism for Exascale Computing* (November 2020), Atlanta, GA, IEEE Xplore. doi: 10.1109/LLVMHPCHiPar51896.2020.00012

Freund, H. P., R. L. Ives, and W. D. Sessions. "High-Efficiency Inductive Output Tubes Using a Third Harmonic Drive on the Grid," *IEEE Transactions on Plasma Science* (November 2020), IEEE. doi: 10.1109/TPS.2020.3027202

Gupta, A., E. Huerta, Z. Zhao, and I. Moussa. "Deep Learning for Cardiologist-Level Myocardial Infarction Detection in Electrocardiograms," *EMBEC 2020: 8th European Medical and Biological Engineering Conference* (November 2020), Springer Nature. doi: 10.1007/978-3-030-64610-3_40

Harmon, K. J., K. Letchworth-Weaver, A. P. Gaiduk, F. Giberti, F. Gygi, M. K. Y. Chan, P. Fenter, and G. Galli. "Validating first-principles molecular dynamics calculations of oxide/water interfaces with x-ray reflectivity data," *Physical Review Materials* (November 2020), APS. doi: 10.1103/PhysRevMaterials.4.113805

Kang, Q., S. Lee, K. Hou, R. Ross, A. Agrawal, A. Choudhary, and W.-k. Liao. "Improving MPI Collective I/O for High Volume Non-Contiguous Requests with Intra-Node Aggregation," *IEEE Transactions on Parallel and Distributed Systems* (November 2020), IEEE. doi: 10.1109/TPDS.2020.3000458

King, G. B., L. Andreoli, S. Pastore, and M. Piarulli. "Weak Transitions in Light Nuclei," *Frontiers in Physics* (November 2020), Frontiers Institutional Membership Books. doi: 10.3389/fphy.2020.00363

König, S. A. Ekström, K. Hebeler, D. Lee, and A. Schwenk. "Eigenvector continuation as an efficient and accurate emulator for uncertainty quantification," *Physical Letters B* (November 2020), Elsevier. doi: 10.1016/j.physletb.2020.135814

Kruse, M., H. Finkel, and X. Wu. "Autotuning Search Space for Loop Transformations," *IEEE/ACM 6th Workshop on the LLVM Compiler Infrastructure in HPC (LLVM-HPC) and Workshop on Hierarchical Parallelism for Exascale Computing* (November 2020), Atlanta, GA, IEEE. doi: 10.1109/LLVMHPCHiPar51896.2020.00007

Lee, R., B. J. Winjum, F. S. Tsung, and W. B. Mori. "Theory and OSIRIS Particle-in-cell Simulations of Stimulated Raman Scattering in unmagnetized and magnetized Plasmas," *62nd Annual Meeting of the APS Division of Plasma Physics* (November 2020), APS

Li, J., X. Zhang, L. Han, Z. Ji, X. Dong, and C. Hu. "OKCM: improving parallel task scheduling in high-performance computing systems using online learning," *The Journal of Supercomputing* (November 2020), Springer. doi: 10.1007/s11227-020-03506-5

Marin, O., E. Constantinescu, and B. Smith. "A scalable matrix-free spectral element approach for unsteady PDE constrained optimization using PETSc/TAO," *Journal of Computational Science* (November 2020), Elsevier. doi: 10.1016/j.jocs.2020.101207

Maulik, R., T. Botsas, N. Ramachandra, L. R. Mason, and I. Pan. "Latent-space time evolution of non-intrusive reduced-order models using Gaussian process emulation," *Physica D: Nonlinear Phenomena* (November 2020), Elsevier. doi: 10.1016/j.physd.2020.132797

Maulik, R., R. Egele, B. Lusch, and P. Balaprakash. “Recurrent Neural Network Architecture Search for Geophysical Emulation,” *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis* (November 2020), vol. 1, Atlanta, GA, IEEE, pp. 89-102. doi: 10.1109/SC41405.2020.00012

Maulik, R., H. Sharma, S. Patel, B. Lusch, and E. Jennings. “A turbulent eddy-viscosity surrogate modeling framework for Reynolds-Averaged Navier-Stokes simulations,” *Computers and Fluids* (November 2020), Elsevier. doi: 10.1016/j.compfluid.2020.104777

Ouyang, K., M. Si, A. Hori, Z. Chen, and P. Balaji. “CAB-MPI: exploring interprocess work-stealing towards balanced MPI communication,” *SC '20: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis* (November 2020), no. 36, Atlanta, GA, IEEE, pp. 1-15. doi: 10.5555/3433701.3433748

Patel, T., Z. Liu, R. Kettimuthu, P. Rich, W. Allcock, and D. Tiwari. “Job Characteristics on Large-Scale Systems: Long-Term Analysis, Quantification, and Implications,” *2020 International Conference for High Performance Computing, Networking, Storage, and Analysis* (November 2020), IEEE. doi: 10.5555/3433701.3433812

Patra, T. K., T. D. Loefflerbc, and S. K. R. S. Sankaranarayanan. “Accelerating copolymer inverse design using monte carlo tree search,” *Nanoscale* (November 2020), The Royal Society of Chemistry. doi: 10.1039/D0NR06091G

Romano, P. K., S. P. Hamilton, R. O. Rahaman, A. Novak, E. Merzari, S. M. Harper, P. C. Shriwise, and T. M. Evans. “A Code-Agnostic Driver Application for Coupled Neutronics and Thermal-Hydraulic Simulations,” *Nuclear Science and Engineering* (November 2020), Taylor and Francis Group. doi: 10.1080/00295639.2020.1830620

Tovey, S., A. N. Krishnamoorthy, G. Sivaraman, J. Guo, C. Benmore, A. Heuer, and C. Holm. “DFT Accurate Interatomic Potential for Molten NaCl from Machine Learning,” *The Journal of Physical Chemistry C* (November 2020), ACS. doi: 10.1021/acs.jpcc.0c08870

Tsung, F. S., R. Lee, B. J. Winjum, and W. B. Mori. “Higher Dimensional Effects in Laser Plasma Interactions Relevant to Inertial Fusion,” *62nd Annual Meeting of the APS Division of Plasma Physics* (November 2020), APS

Winjum, B. J., R. Lee, F. S. Tsung, and W. B. Mori. “Simulations of Laser Plasma Interactions for Moderate to High Laser Intensities,” *62nd Annual Meeting of the APS Division of Plasma Physics* (November 2020), APS

Woodard, A. E., A. Trisovic, Z. Li, Y. Babuji, R. Chard, T. Skluzacek, B. Blaiszik, D. S. Katz, I. Foster, and K. Chard. “Real-time HEP analysis with funcX, a high-performance platform for function as a service,” *EPJ Web of Conferences* (November 2020), EDP Sciences. doi: 10.1051/epjconf/202024507046

Yu, V. W.-Z., C. Campos, W. Dawson, A. Garcia, V. Havu, B. Hourahine, W. P. Huhn, M. Jacquelin, W. Jia, M. Keceli, R. Laasner, Y. Li, L. Lin, J. Lu, J. Moussa, J. E. Roman, Á. Vázquez-Mayagoitia, C. Yang, and V. Blum. “ELSI — An open infrastructure for electronic structure solvers,” *Computer Physics Communications* (November 2020), Elsevier. doi: 10.1016/j.cpc.2020.107459

December

Bass, B., J. New, and W. Copeland. “Potential Energy, Demand, Emissions, and Cost Savings Distributions for Buildings in a Utility’s Service Area,” *Energies: Designing, Modeling and Optimizing Energy and Environmental Systems for Buildings* (December 2020), MDPI. doi: 10.3390/en14010132

Dunton, A. M., L. Jofre, G. Iaccarino, and A. Doostan. “Pass-efficient methods for compression of high-dimensional turbulent flow data,” *Journal of Computational Physics* (December 2020), Elsevier. doi: 10.1016/j.jcp.2020.109704

Gritsenko, D., and R. Paoli. “Theoretical Analysis of Fractional Viscoelastic Flow in Circular Pipes: Parametric Study,” *Applied Sciences* (December 2020), MDPI. doi: 10.3390/app10249080

Grosset, P., J. Pulido, and J. Ahrens. “Personalized In Situ Steering for Analysis and Visualization,” *ISAV* (December 2020), ACM. doi: 10.1145/3426462.3426463

Schwarz, N., S. Campbell, A. Hexemer, A. Mehta, and J. Thayer. “Enabling Scientific Discovery at Next-Generation Light Sources with Advanced AI and HPC,” *Smoky Mountains Computational Sciences and Engineering Conference* (December 2020), Springer, pp. 145-156. doi: 10.1007/978-3-030-63393-6_10

Schwartz, J., H. Zheng, M. Hanwell, Y. Jiang, and R. Hovden. “Dynamic compressed sensing for real-time tomographic reconstruction,” *Ultramicroscopy* (December 2020), Elsevier. doi: 10.1016/j.ultramic.2020.113122

Sobczyk, J. E., B. Acharya, S. Bacca, and G. Hagen. “Coulomb sum rule for 4He and 16O from coupled-cluster theory,” *Physical Review C* (December 2020), APS. doi: 10.1103/PhysRevC.102.064312

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.

Availability of this Report (ANL/ALCF-21/2)

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.
Alexandra, VA 22312

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

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

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831-0062

phone | 865.576.8401
fax | 865.576.5728
reports@adonis.osti.gov
www.osti.gov

ALCF Leadership: Michael E. Papka (Division Director), Jini Ramprakash (Deputy Division Director), Katherine Riley (Director of Science), Kalyan Kumaran (Director of Technology), Mark Fahey (Director of Operations), and Susan Coghlan (ALCF-X Project Director)

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

Design and production: Sandbox Studio, Chicago

Argonne NATIONAL LABORATORY

CONTACT
media@alcf.anl.gov
alcf.anl.gov

U.S. DEPARTMENT OF ENERGY

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