Argonne NATIONAL LABORATORY

Argonne Leadership Computing Facility

Science

2022

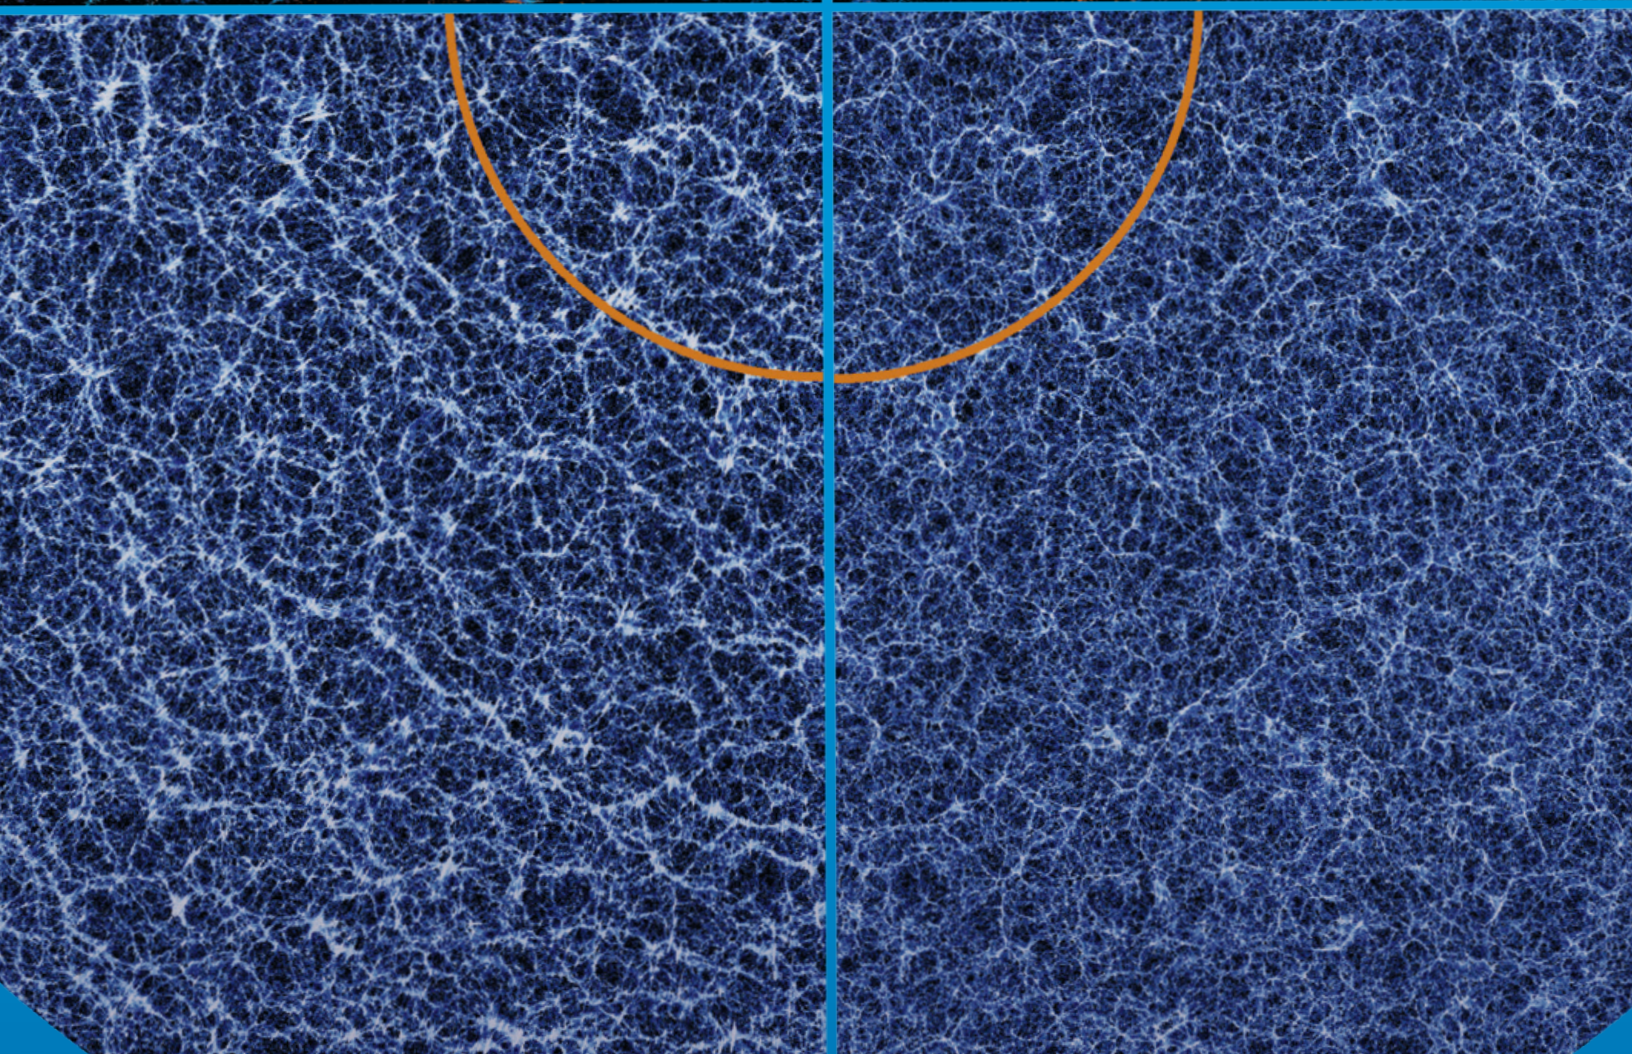

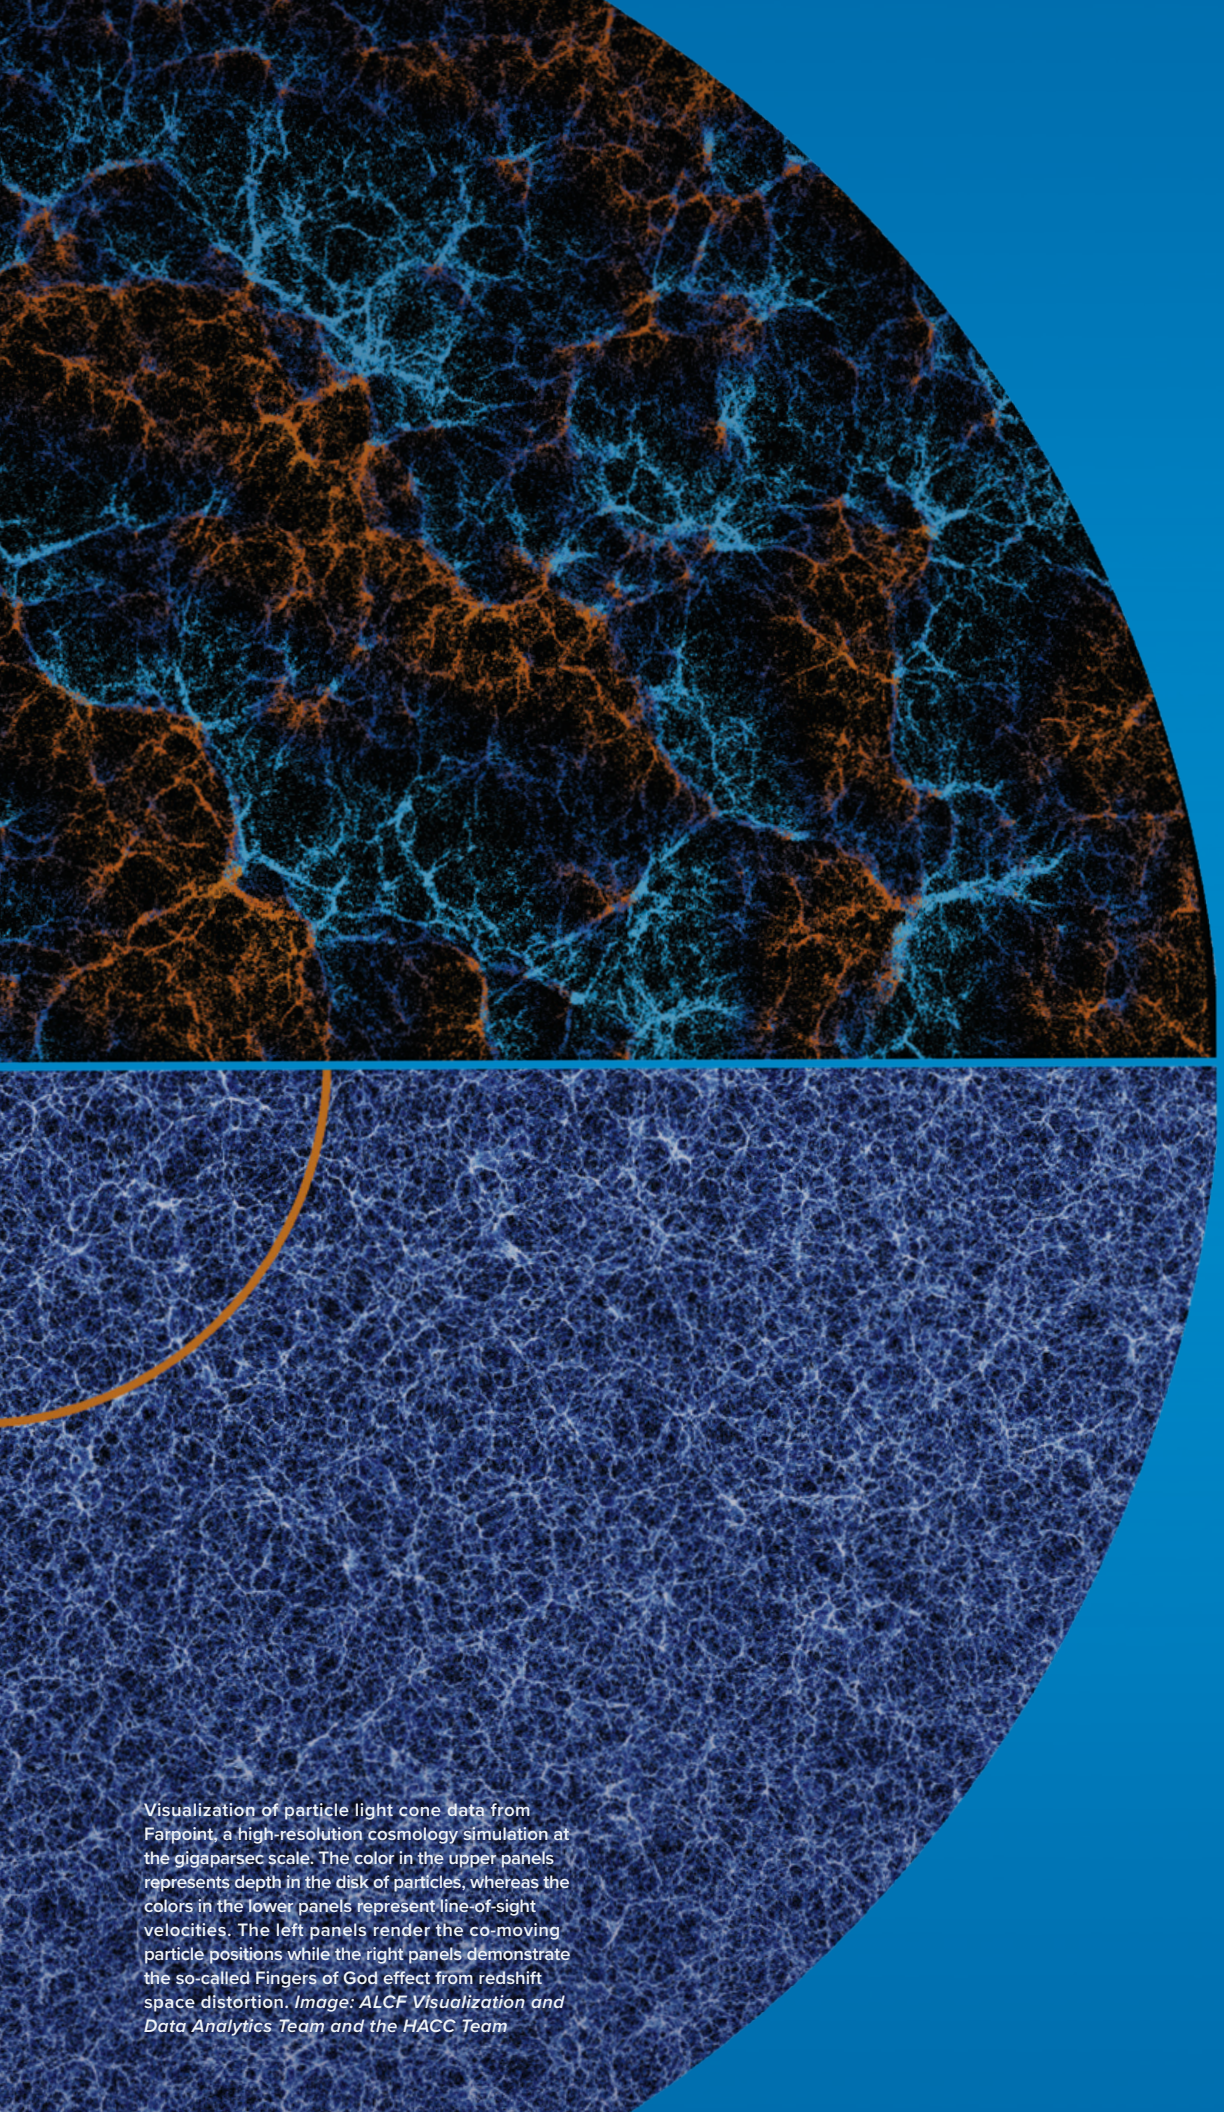

Visualization of particle light cone data from Farpoint, a high-resolution cosmology simulation at the gigaparsec scale. The color in the upper panels represents depth in the disk of particles, whereas the colors in the lower panels represent line-of-sight velocities. The left panels render the co-moving particle positions while the right panels demonstrate the so-called Fingers of God effect from redshift space distortion. *Image: ALCF Visualization and Data Analytics Team and the HACC Team*

CONTENTS

MICHAEL E. PAPKA
ALCF Director

BILL ALLCOCK
ALCF Director of Operations

SUSAN COGHLAN
ALCF-X Project Director

KALYAN KUMARAN
ALCF Director of Technology

JINI RAMPRAKASH
ALCF Deputy Director

KATHERINE RILEY
ALCF Director of Science

MESSAGE FROM ALCF LEADERSHIP

It's been a momentous year for the ALCF as we rolled out some brand-new computing resources to our user community with the deployment of Polaris and ALCF AI Testbed systems. Central to our ongoing efforts to build a data-centric leadership computing environment, Polaris and the novel AI Testbed accelerators offer advanced capabilities for research involving machine learning and big data analysis tasks—providing a playground to prepare for the future.

Polaris, a Hewlett Packard Enterprise system equipped with AMD CPUs (central processing units) and NVIDIA GPUs (graphics processing units), is our largest GPU-accelerated system to date. With a hybrid configuration that is similar to DOE's exascale supercomputers, Polaris is a key resource for advancing exascale readiness for Aurora and beyond. Teams participating in the Exascale Computing Project and the Aurora Early Science Program have begun tapping the system to scale up problem sizes and inform simulations and workloads planned for the larger exascale supercomputers. Polaris is also being used for a number of scientific research campaigns awarded time through DOE's INCITE and ALCC allocation programs.

In May, the first two ALCF AI Testbed platforms were opened up to the broader scientific computing community. We began accepting proposals from researchers interested in using these initial systems from Cerebras and SambaNova to help accelerate science. The testbed is allowing the ALCF and its user community to evaluate the usability and performance of machine learning-based applications running on the specialized systems. Our goal is to explore and understand how these accelerators can be integrated with current and future supercomputers to accelerate data-driven discoveries.

We also made significant progress in standing up Aurora over the past year. Much of the Intel-HPE system has been assembled in the new data center space with the installation of computer racks and several components, including the Sunspot test and development platform, the HPE Slingshot interconnect technology, and the DAOS storage system. The installation continues with the phased delivery of the Intel Ponte Vecchio GPUs and Sapphire Rapids CPUs that will provide Aurora's computational horsepower. The system is slated to be completed next year with an upgrade to Sapphire Rapids CPUs with high-bandwidth memory.

Work also continues behind the scenes to prepare scientific applications and software for the power and scale of Aurora. ALCF staff are playing a critical role in porting and optimizing codes, frameworks, and libraries to ensure their performance and functionality meets the needs of the HPC user community.

While these new systems represent the future of the ALCF, our primary production system, Theta, continued to enable innovative research in across a wide of disciplines. In 2022, Peter Coveney's INCITE project combined deep learning with physics-based simulations to speed up the search for promising new drugs to combat COVID-19 and other viruses; Rao Kotamarthi's ALCC project carried out studies to assess the risks of extreme weather events at local and regional levels across North America; and Logan Ward's ADSP project advanced their efforts to build an autonomous AI application that can help researchers identify better-performing battery materials.

In the following pages, you'll learn more about these scientific campaigns and many other projects that are using ALCF resources to advance the state of the art in their respective fields. Our annual Science Report also details how ALCF's powerful computing systems, talented staff, and diverse user community are helping to blaze new trails in using HPC and AI for science.

ARGONNE LEADERSHIP COMPUTING FACILITY

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

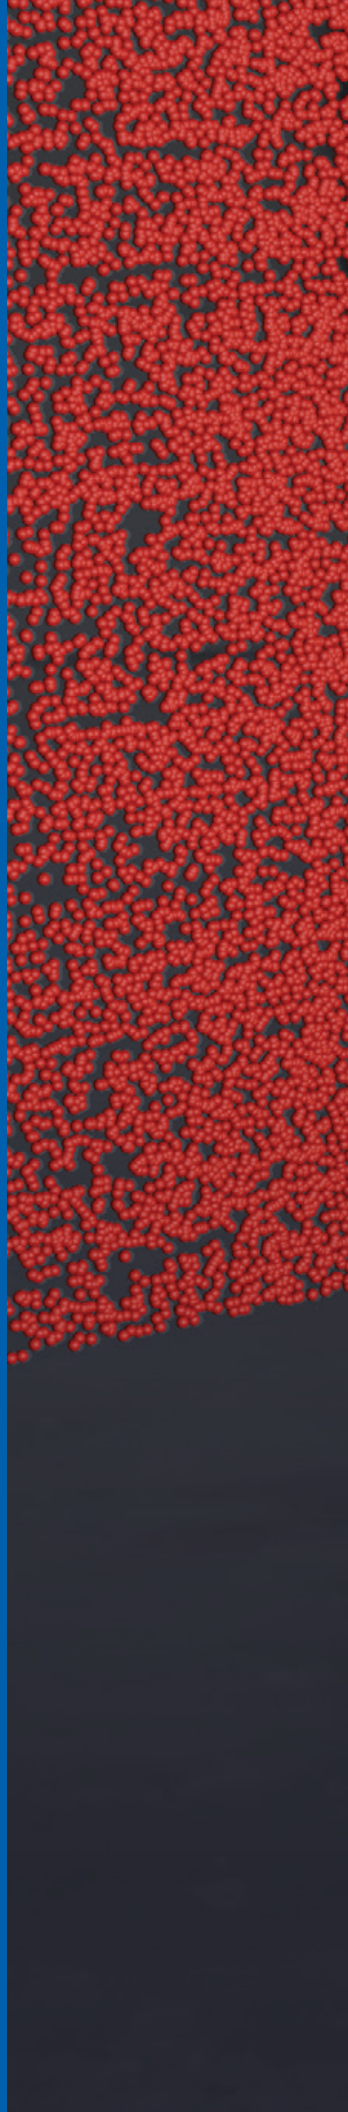

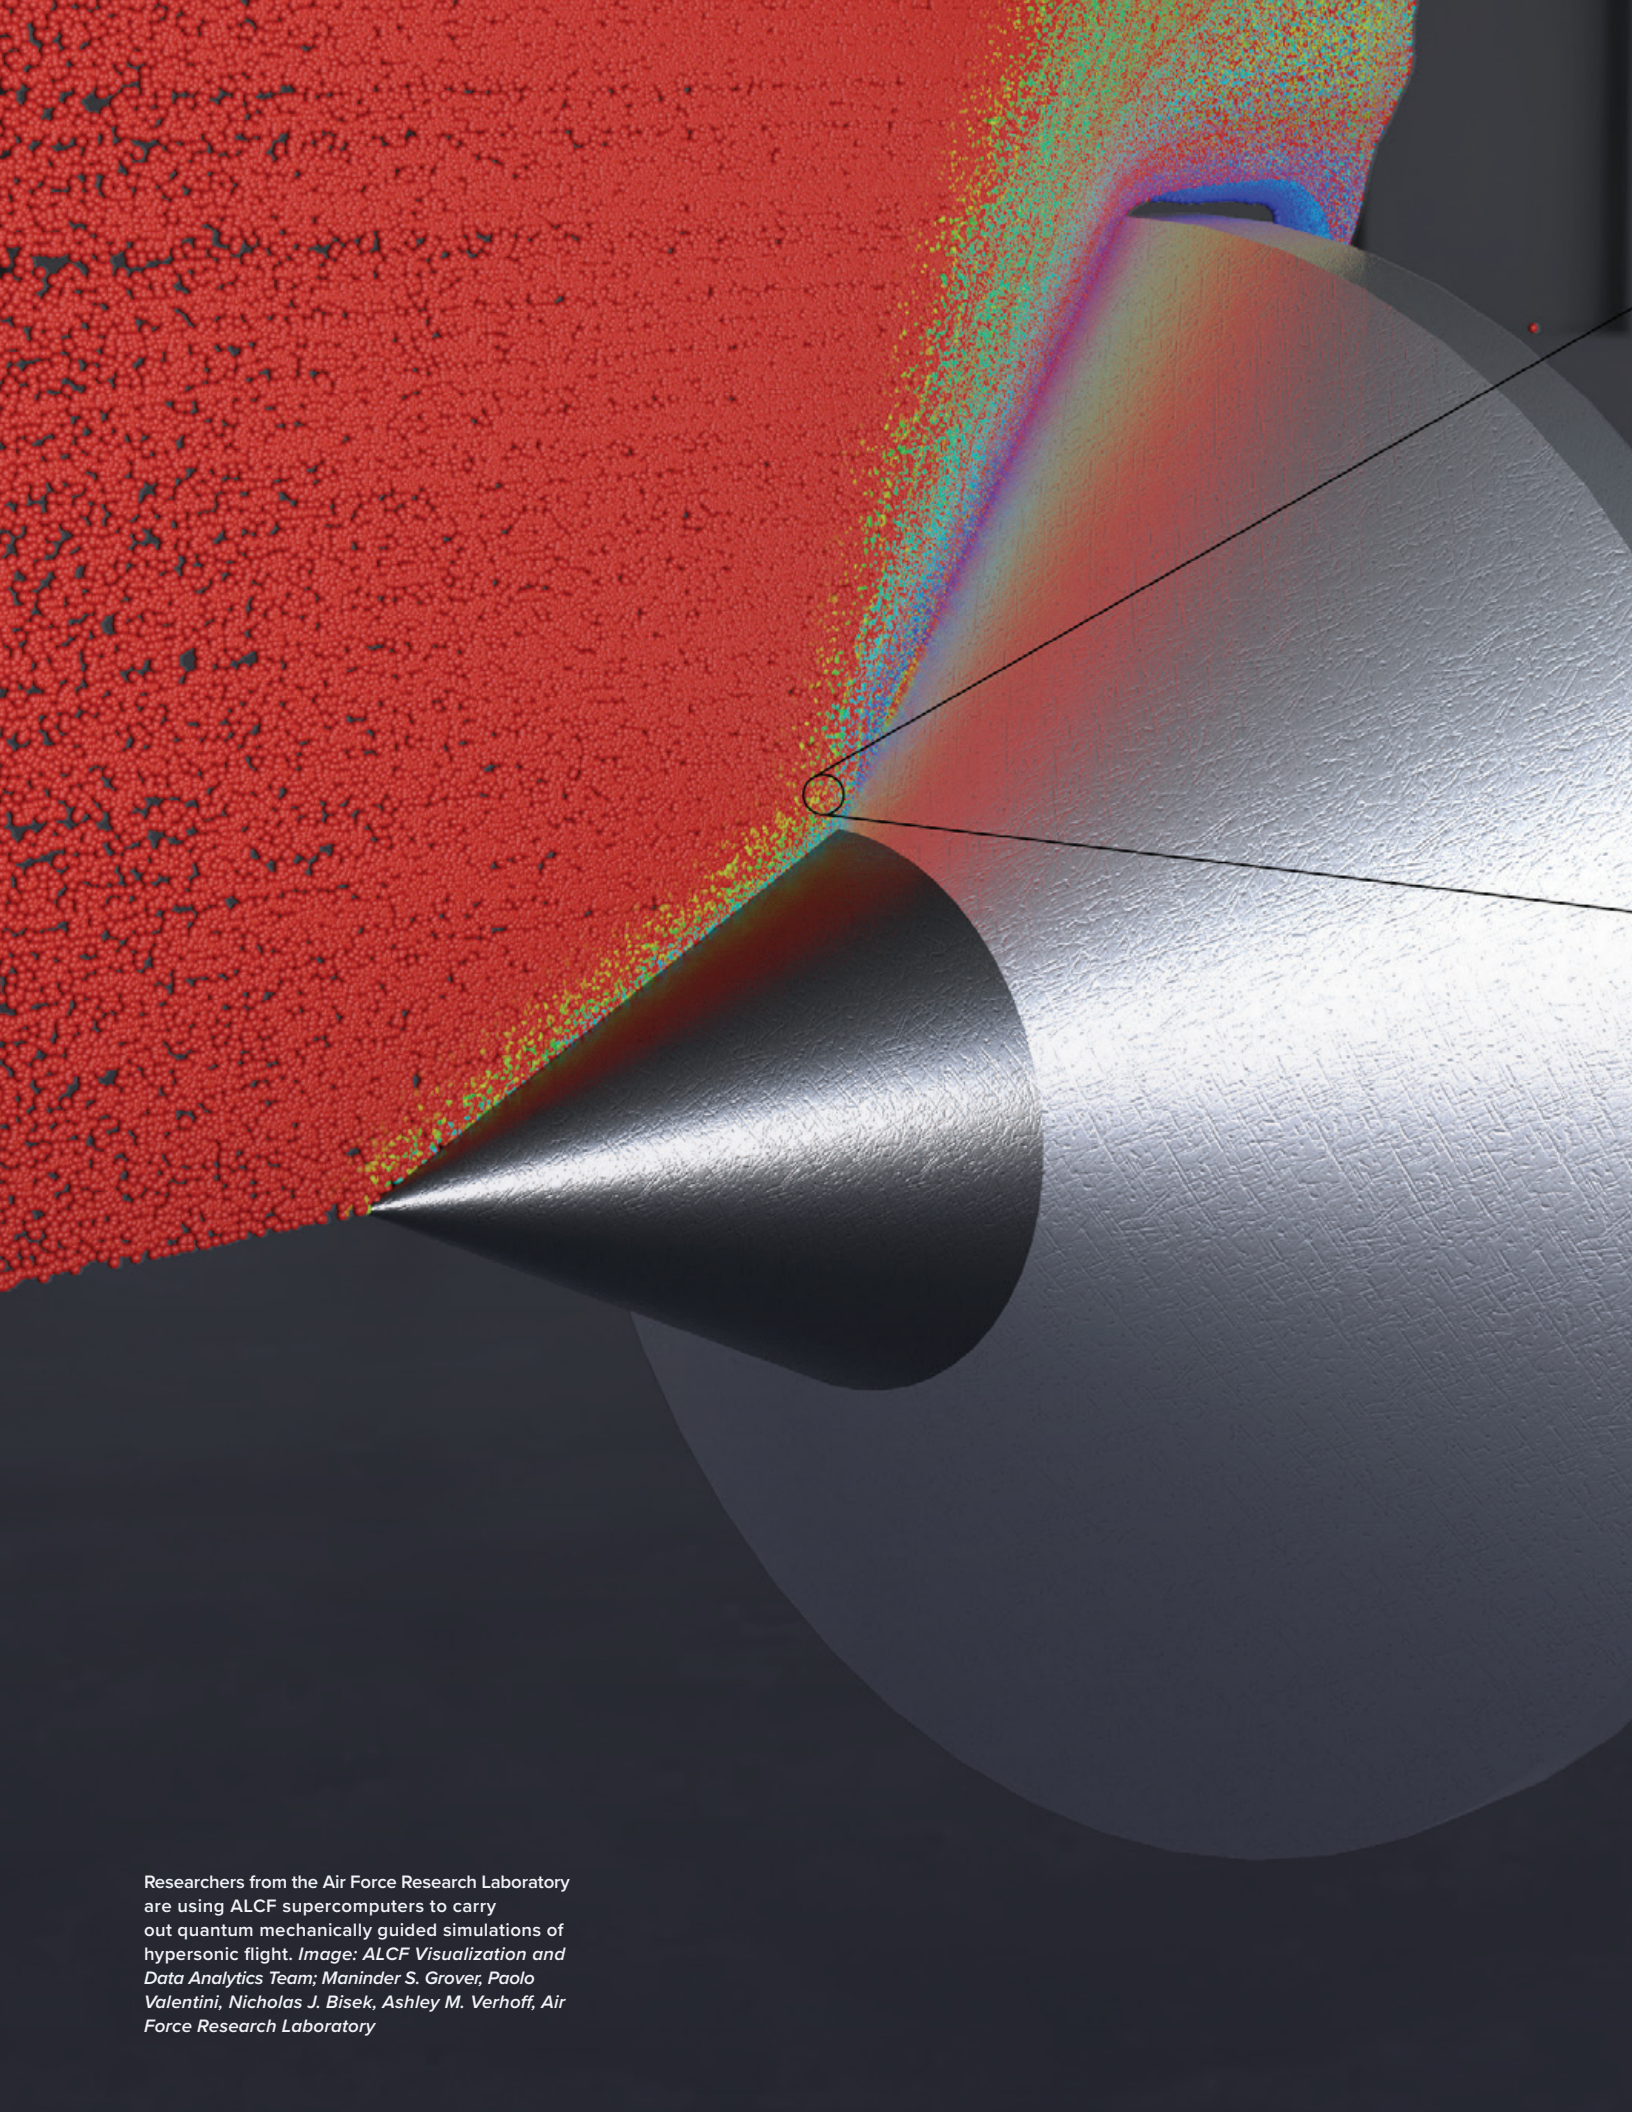

Researchers from the Air Force Research Laboratory are using ALCF supercomputers to carry out quantum mechanically guided simulations of hypersonic flight. *Image: ALCF Visualization and Data Analytics Team; Maninder S. Grover, Paolo Valentini, Nicholas J. Bisek, Ashley M. Verhoff, Air Force Research Laboratory*

About ALCF

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

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

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

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

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

ALCF Team

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

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

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

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

User support specialists provide technical 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 to external audiences the impactful research enabled by ALCF resources.

Supercomputing Resources

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

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

ALCF AI TESTBED

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

Groq

- Tensor Streaming Processor
- >26 billion transistors, 14 nm
- Chip-to-Chip interconnect
- GroqWare software stack, Onnx

SambaNova DataScale

- Reconfigurable Dataflow Unit
- >40 billion transistors, 7 nm
- RDU-Connect
- SambaFlow software stack, PyTorch

Habana Gaudi

- Tensor processing cores
- 7nm
- Integrated 100 GbE-based interconnect
- Synapse AI Software, PyTorch, Tensorflow

Graphcore MK1

- Intelligent Processing Unit (IPU)
- 1216 IPU tiles, 14 nm
- >23 billion transistors
- IPU-Links interconnect
- Poplar software stack, PyTorch, Tensorflow

Cerebras CS-2

- Wafer-Scale Engine
- >800,000 processing cores
- 2.6 trillion transistors, 7 nm
- SwarmX fabric
- Tensorflow, PyTorch

DATA STORAGE SYSTEMS

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

SYSTEM NAME	EAGLE	GRAND	THETA-FS0	SWIFT	TAPE STORAGE
File System	Lustre	Lustre	Lustre	Lustre	—
Storage System	HPE ClusterStor E1000	HPE ClusterStor E1000	HPE Sonexion L300	All NVMe Flash Storage Array	LTO6 and LTO8 Tape Technology
Usable Capacity	100 PB	100 PB	9 PB	123 TB	300 PB
Sustained Data Transfer Rate	650 GB/s	650 GB/s	240 GB/s	48 GB/s	—
Disk Drives	8,480	8,480	2,300	24	—

NETWORKING

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

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

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

JOINT LABORATORY FOR SYSTEM EVALUATION

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

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

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

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

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

Presque: Intel DAOS nodes for testing the Aurora storage system

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

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

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

ADVANCING HPC AND AI FOR SCIENCE

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

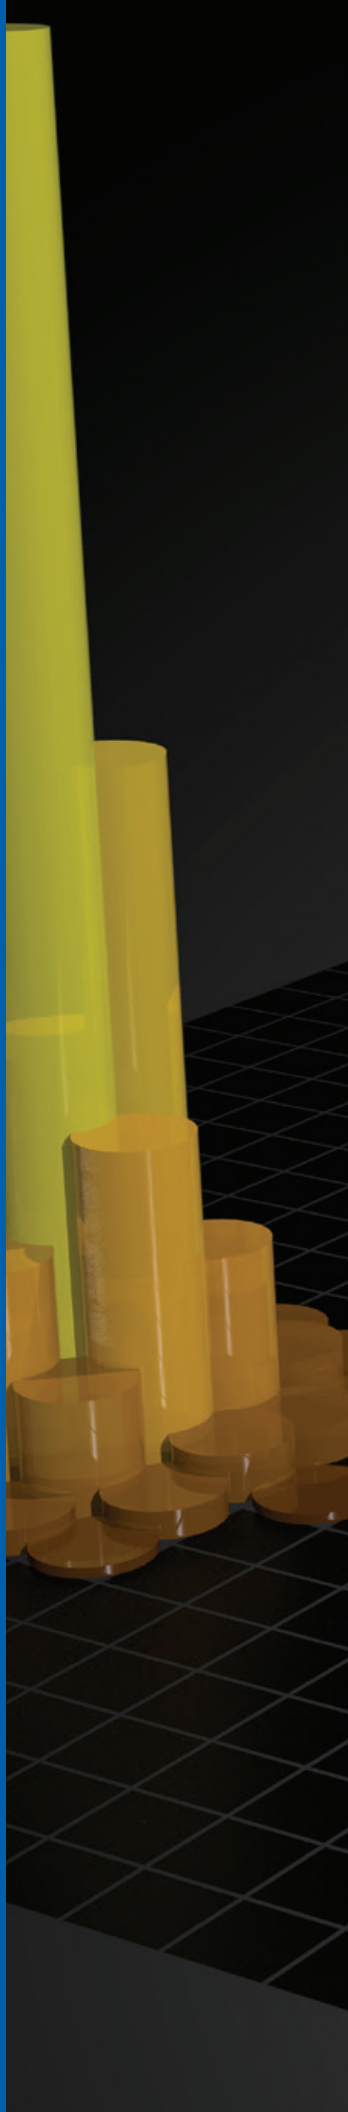

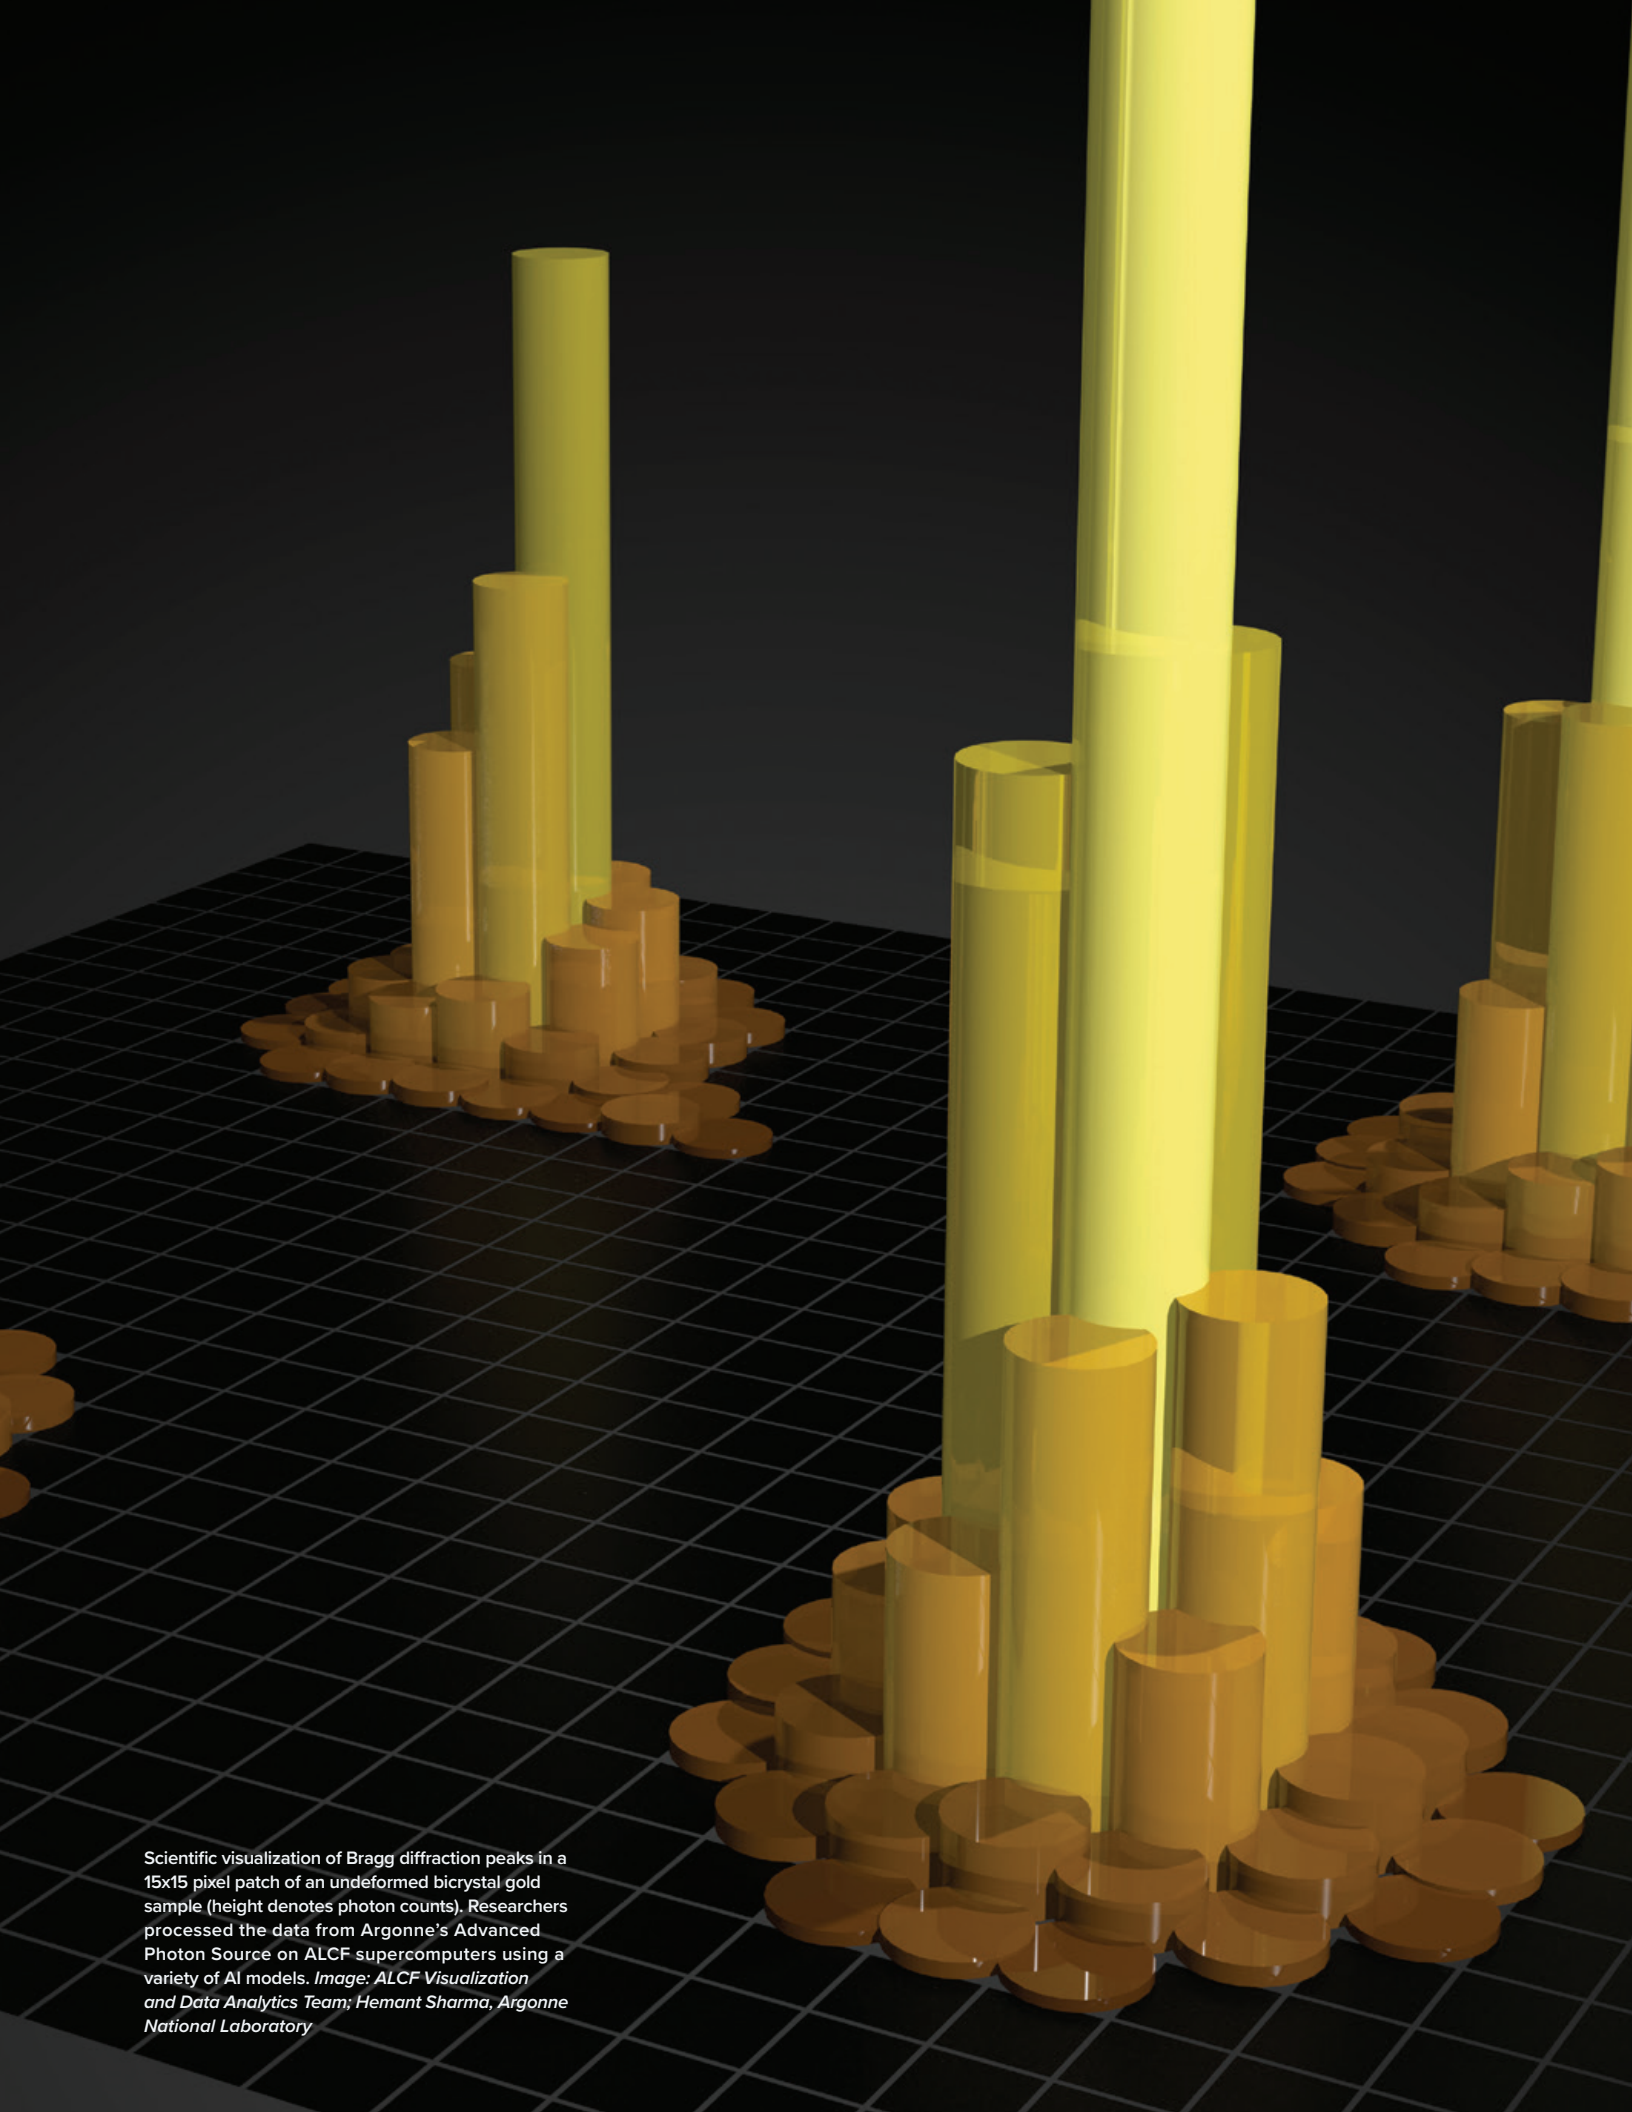

Scientific visualization of Bragg diffraction peaks in a 15x15 pixel patch of an undeformed bicrystal gold sample (height denotes photon counts). Researchers processed the data from Argonne's Advanced Photon Source on ALCF supercomputers using a variety of AI models. *Image: ALCF Visualization and Data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

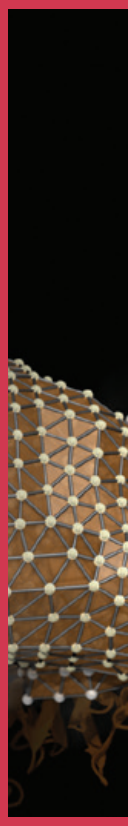

ALCF AI Testbed to Boost Data-Driven Science

The ALCF AI Testbed is enabling the facility and its user community to help determine the role of AI accelerators in scientific computing.

In May 2022, the ALCF AI Testbed was officially rolled out to the research community as the facility began accepting proposals for computing time on its Cerebras CS-2 and SambaNova DataScale systems.

A growing collection of some of the world's most advanced AI accelerators available for science, the ALCF AI Testbed also includes Graphcore, Groq, and Habana AI systems that will be made available to researchers in the near future.

The ALCF assembled and launched the AI Testbed to enable researchers to explore next-generation machine learning applications and workloads to advance the use of AI for science. The testbed platforms complement the ALCF's current and next-generation supercomputers to provide a state-of-the-art computing environment that supports pioneering research at the intersection of AI, big data, and HPC.

Offering architectural features designed specifically for AI and data-centric workloads, the AI Testbed systems are uniquely well-suited to handle the growing amount of scientific data produced by supercomputers, light sources, telescopes, particle accelerators, and other experimental tools and facilities.

The state-of-the-art accelerators are allowing researchers to explore novel workflows that combine AI methods with simulation and experimental science to accelerate the pace of discovery. Moreover, the testbed stands to significantly broaden data analysis and processing capabilities in the project workflows deployed at the ALCF beyond those supported by traditional CPU- and GPU-based machines.

To introduce researchers to using the AI accelerators for science, the ALCF hosted hands-on workshops for both the Cerebras and SambaNova systems. The two-day events covered system hardware, software, application porting, and best practices.

Prior to opening the ALCF AI Testbed up to the broader scientific community, Argonne researchers led several collaborative efforts to use the AI accelerators for a variety of data-centric studies. The following summaries provide a glimpse of some of the early success stories that have leveraged AI Testbed systems.

Using multiple supercomputers and the ALCF's Cerebras system, an Argonne-led team carried out an innovative study of the SARS-CoV-2 replication mechanism that was nominated for the Gordon Bell Special Prize for HPC-Based COVID-19 Research at SC21. *Image: Defne Gorgun, University of Illinois at Urbana-Champaign; Anda Trifan, University of Illinois at Urbana-Champaign/Argonne National Laboratory; and Arvind Ramanathan, Argonne National Laboratory*

Drug Discovery

A team of researchers leveraged the ALCF's Groq system to accelerate the process of searching through a vast number of small molecules to find promising antiviral drugs to fight COVID-19. With billions upon billions of potential drug candidates to sort through, the scientists needed a way to dramatically speed up their search. In tests on a large dataset of molecules, the team found they could achieve 20 million predictions, or inferences, per second, vastly reducing the time needed for each search from days to minutes. The most promising candidates were sent to a laboratory for further testing on human cells.

Edge Computing

To keep pace with the growing amount of data produced at DOE light source facilities, researchers are looking to machine learning methods to help with tasks such as data reduction and providing insights to steer future experiments. Using the ALCF's Cerebras and SambaNova systems, researchers demonstrated how specialized AI systems can be used to quickly train machine learning models through a geographically distributed workflow. To obtain actionable information in real-time, the team trained the models on the remote AI system and then deployed them on edge computing devices near the experimental data source.

Fusion Energy

As part of an effort to improve predictive capabilities for fusion energy research, researchers turned to the ALCF's Groq system to accelerate the performance of deep learning models used to investigate fusion control in real time. The Groq system's architecture ensured fixed, predictable compute times for a key phase of deep learning (inference) that would vary in duration if carried out on CPU- and GPU-driven machines. Ultimately, the researchers aim to develop a workflow that leverages AI and exascale computing power for training and inference tasks that will advance fusion energy research.

Neutrino Physics

To improve neutrino signal efficiency, scientists use image segmentation to tag each input pixel as one of three classes: cosmic-induced, neutrino-induced, or background noise. Deep learning has been a useful tool for accelerating this classic image segmentation task, but it has been limited by the image size that available GPU-based platforms can efficiently train on. Leveraging the ALCF's SambaNova system, researchers were able to improve this method to establish a new state-of-the art accuracy level of 90.23% using images at their original resolution without the need to downsample. Their work demonstrates capabilities that can be used to advance model quality for a variety of important and challenging image processing problems.

Polaris Kicks Off Exascale Era Computing at ALCF

The ALCF's most powerful system to date is helping researchers prepare for Aurora while supporting several large-scale science projects.

The ALCF officially deployed Polaris for scientific research in August, giving its user community a powerful new resource to prepare for science in the exascale era.

At 44 petaflops, the new supercomputer is the ALCF's most powerful system to date. Polaris is providing a platform for researchers to prepare codes and workloads for the upcoming Aurora exascale supercomputer while simultaneously supporting several large-scale research projects, including efforts focused on using AI for science.

At Argonne, Polaris will also be a key tool to advance efforts to integrate ALCF computing resources more closely with DOE experimental facilities. Initial work will be focused on coupling ALCF systems with the co-located Advanced Photon Source and Center for Nanoscale Materials, both DOE Office of Science user facilities.

Like Aurora, Polaris is a hybrid system powered by both central processing units (CPUs) and graphics processing units (GPUs). This similarity, along with some other shared technologies, will help ease the transition to using Aurora for scientific research ahead of its deployment.

The HPE Apollo Gen10+ based supercomputer is equipped with 560 AMD EPYC processors and 2,240 NVIDIA A100 Tensor Core GPUs. The Polaris software environment is equipped with the HPE Cray programming environment, HPE Performance Cluster Manager (HPCM) system software, and the ability to test programming models, such as OpenMP and SYCL, that will be available on Aurora and the next generation of DOE supercomputers. Polaris users will also benefit from NVIDIA's HPC software development kit—a suite of compilers, libraries, and tools for GPU code development.

Upon the deployment of Polaris, several research teams immediately began using the system to prepare for the arrival of Aurora via DOE's Exascale Computing Project (ECP) and the ALCF's Aurora Early Science Program (ESP). The system is also being leveraged by multiple projects awarded time through DOE's INCITE and ALCC allocation programs. The following summaries highlight some of the initial projects tapping Polaris to accelerate science.

Exascale Preparations

Researchers from DOE's Princeton Plasma Physics Laboratory are leading an Aurora ESP project that aims to use particle simulations of plasma, including impurities and the important magnetic field geometry at the edge, to predict behavior of ITER plasmas and to help guide future experimental parameters. The team is using Polaris to run simulations of the DIII-D National Fusion Facility's tokamak device with carbon impurities to inform their planned simulations of ITER plasmas with tungsten impurities on Aurora.

The NWChemEx team, supported by the ESP and ECP, will use Aurora to advance the understanding of chemical systems to aid in developing new methods for converting biomass into biofuels. In preparation for larger simulations on

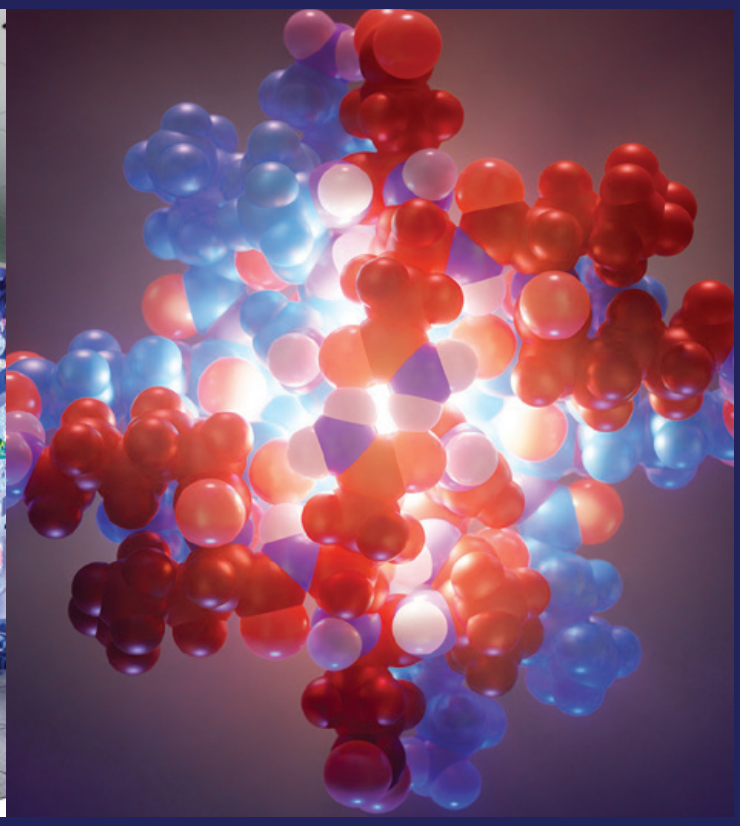

The ALCF's new Polaris system will support research across a wide range of scientific disciplines, including protein and peptide design (right: an artificially designed symmetric peptide macrocycle). ***Image: Vikram Mulligan, Flatiron Institute***

Aurora, the researchers are leveraging Polaris to carry out simulations to produce initial geometries appropriate for the conversion of propene to propanol within zeolites.

An ESP team from the Massachusetts Institute of Technology plans to use Aurora to couple machine learning and state-of-the-art physics simulations to determine possible interactions between nuclei and a large class of dark matter candidate particles. Their machine learning models had only been tested on small-scale problems so the team is using Polaris to scale up problem sizes to determine if they need to modify their machine learning models for their upcoming runs on Aurora.

Science Campaigns

With an INCITE award, researchers from the Flatiron Institute will use Polaris to aid in their efforts to reduce the computational and energetic costs of producing successful peptide macrocycle drugs or industrial enzymes with useful functions in medicine and manufacturing. The INCITE team will develop low-cost machine learning methods that can approximate the output of computationally expensive validation simulations, ultimately allowing users without access to large-scale resources to perform design and validation tasks on much more modest computing systems.

An Argonne-led ALCC team will leverage Polaris to perform high-fidelity computational fluid dynamics and multiphysics simulations to investigate fundamental flow phenomena in next-generation nuclear power reactors. The project's simulations will be carried out in collaboration with X-energy and TerraPower (recipients of awards from DOE's Advanced Reactor Demonstration Program) with the potential to accelerate the deployment of carbon-free energy solutions.

Building on years of work using Theta, a Princeton University-led team will use Polaris to advance their 3D simulations of core-collapse supernova. With their new INCITE allocation, the researchers will conduct not only a full suite of 3D simulations for the spectrum of progenitor stars but carry out these simulations for approximately five times the physical time possible with their previous INCITE allocations all the way to the asymptotic state.

Aurora: Teams Make Strides with Hardware and Software

Aurora began to take shape in the ALCF data center while staff continued efforts to prepare the exascale supercomputer for science on day one.

In 2022, the ALCF's Aurora exascale supercomputer officially began to take shape. The system's racks, along with several components, are now in place in the new data center at Argonne National Laboratory. The Intel-HPE supercomputer will soon be completed with the installation of Intel's state-of-the-art Ponte Vecchio GPUs and Sapphire Rapids CPUs with high-bandwidth memory.

The supporting infrastructure that surrounds the supercomputer is an engineering marvel in its own right. The Theory and Computing Sciences (TCS) building at Argonne was expanded with the addition of a new 80,000 square foot wing that includes the data center, office and meeting spaces, a public viewing area, and mechanical rooms to accommodate the power and cooling infrastructure.

The new area is equipped with 14 substations that provide 60 megawatts of capacity to power the data center that will house Aurora and future Argonne computing systems, as well as the building's day-to-day electricity needs. The cooling infrastructure for the data center pumps 44,000 gallons of water through a complex loop that connects cooling towers, chillers, heat exchangers, and other systems to ensure the water is at the perfect temperature, pressure, and purity as it passes through the Aurora blades.

Having a majority of the physical system and supporting infrastructure in place has allowed the Argonne-Intel-HPE team to test and fine-tune various components and ensure the power and cooling systems are ready for the massive system in advance of its deployment.

Behind the scenes, Argonne researchers are contributing to a broad range of activities to prepare Aurora for science on day one. To ensure key software is ready to run on the exascale system, scientists and developers participating in DOE's Exascale Computing Project (ECP) and the ALCF's Aurora Early Science Program (ESP) continue work to port and optimize dozens of scientific computing applications. With access to the Aurora software development kit and early hardware, the teams are working to improve the performance and functionality of various codes, frameworks, and libraries using the programming models that will be supported on Aurora.

Training events, such as workshops, hackathons, and webinars, have been a key mechanism for disseminating the latest details on hardware and software and helping developers ramp up for Aurora's arrival. In June, for example, the Argonne-Intel Center for Excellence hosted a multi-day workshop for ECP and ESP teams to provide updates on the system; share approaches and best practices for performance portability; and facilitate hands-on sessions with software tools, such as Intel VTune and Advisor, DPC++ and SYCL, and AI frameworks.

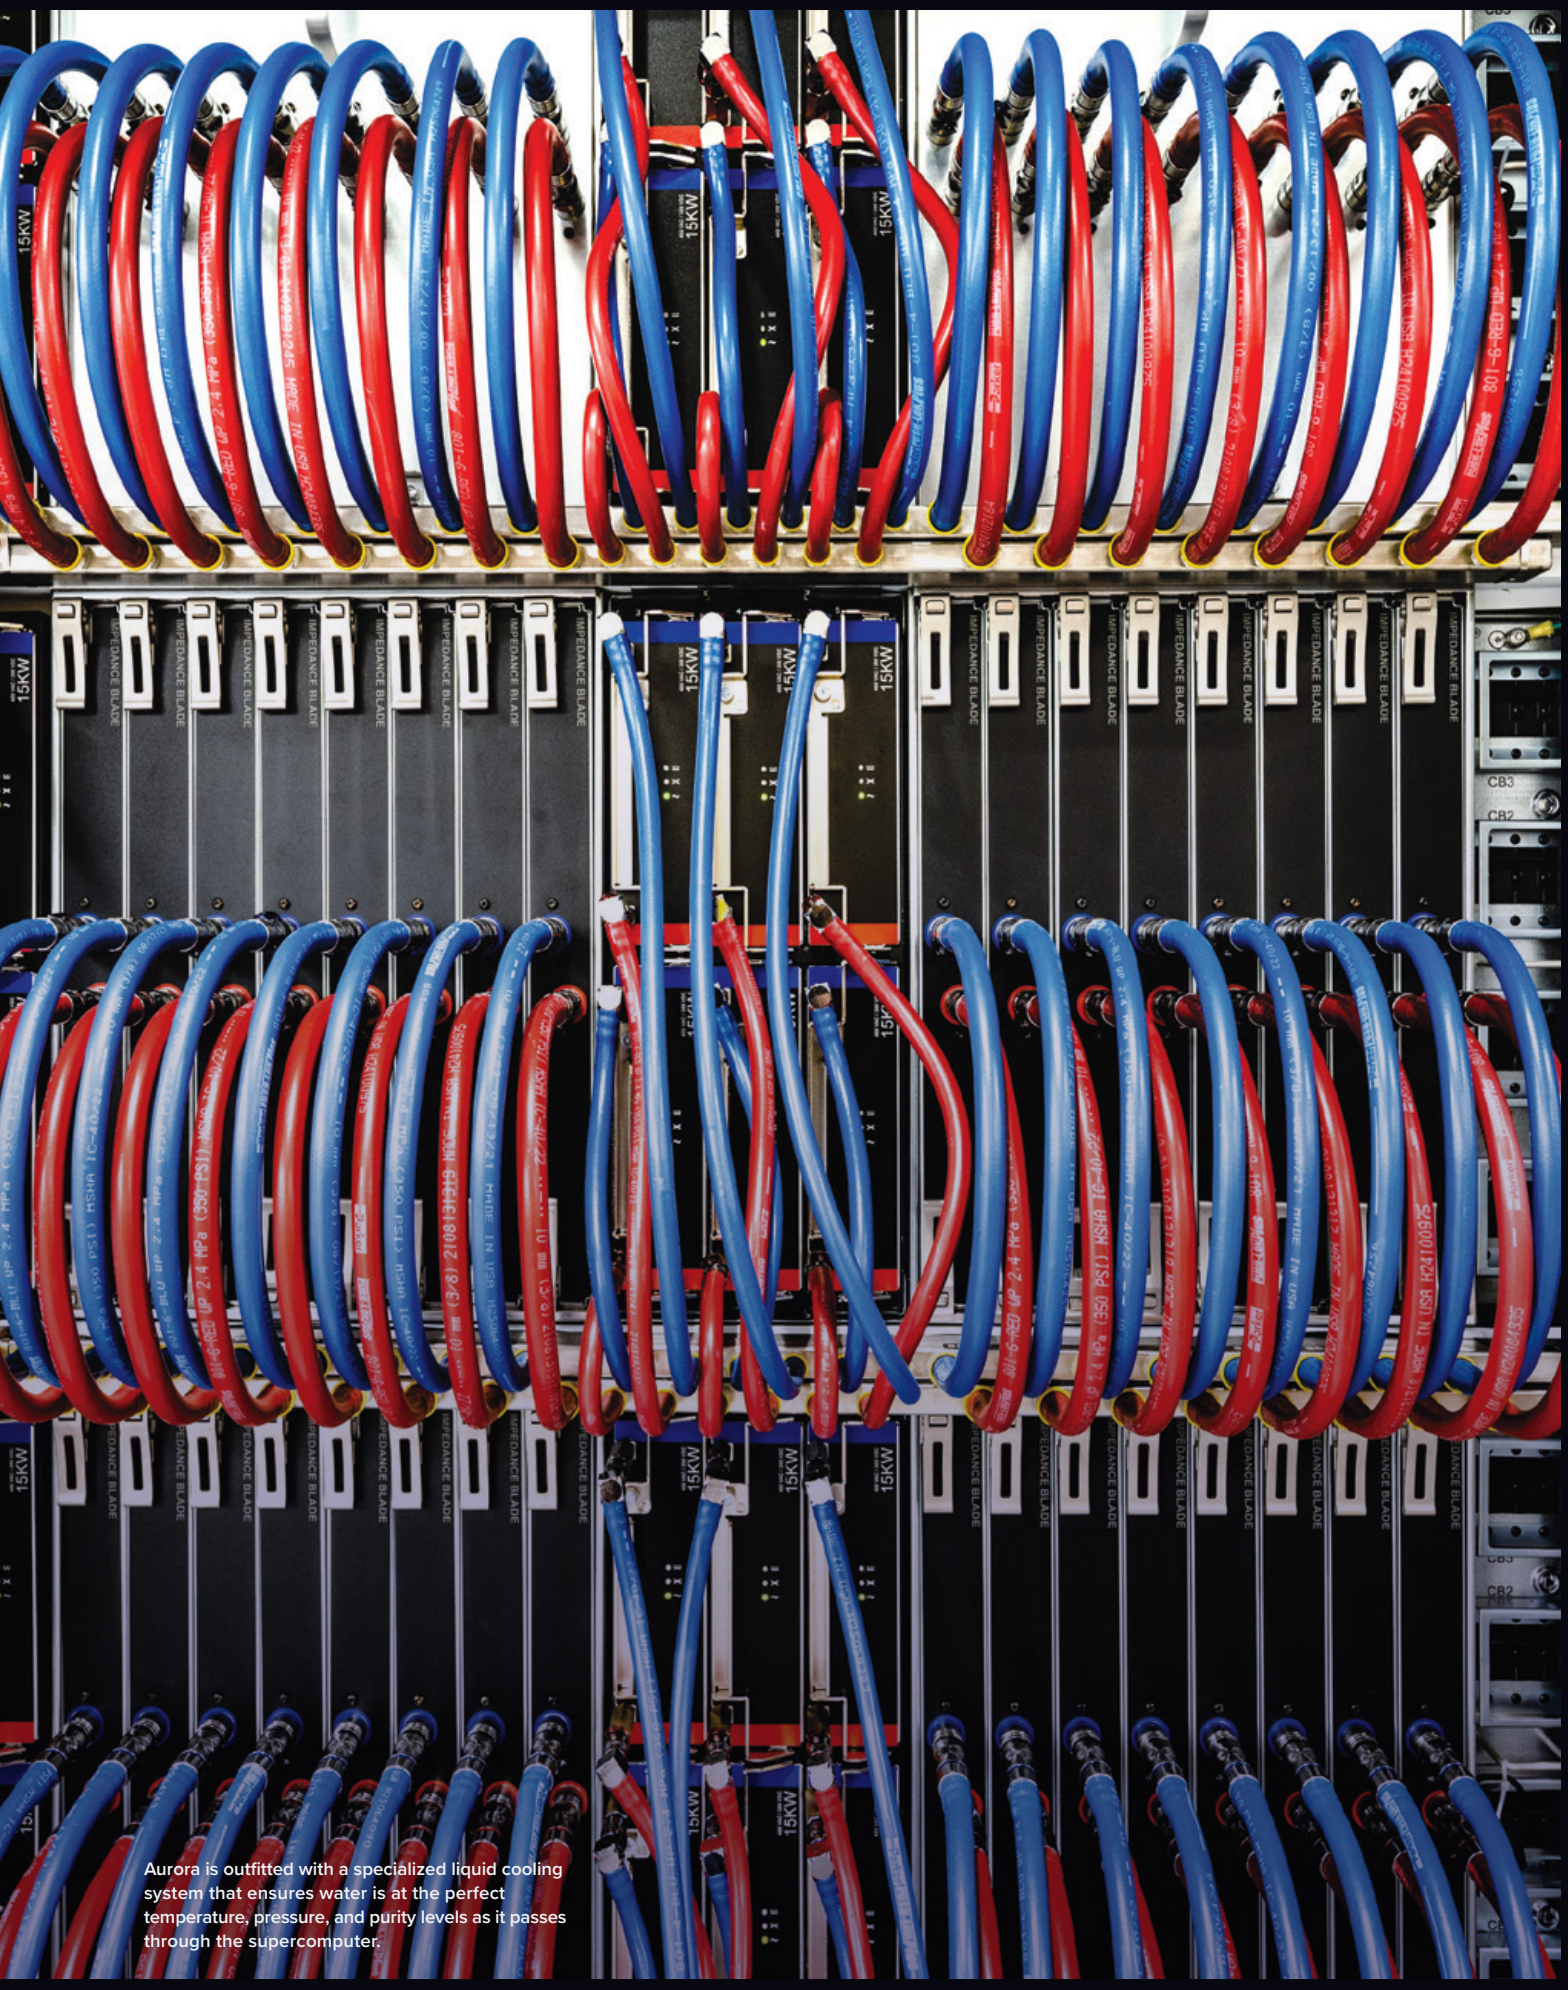

Aurora is outfitted with a specialized liquid cooling system that ensures water is at the perfect temperature, pressure, and purity levels as it passes through the supercomputer.

Exascale Expertise

ALCF's Bethany Lusch (left) and Murali Emani are working to prepare machine learning and data analysis capabilities for Aurora.

The process of planning and preparing for a new leadership-class supercomputer takes years of collaboration and coordination. It requires partnerships with vendors and the broader HPC community to test and develop various hardware and software components, validating that their performance and functionality meets the needs of the scientific computing community.

The following summaries provide a look at a few of the many ALCF staff efforts underway to ready the facility and its users for the exascale computing era.

Machine Learning and Data Science Tools

ALCF computer scientists Bethany Lusch and Murali Emani are working closely with Intel to prepare the oneDAL data analytics library and the scikit-learn machine learning package for Aurora. Designed to speed up big data analysis, oneDAL optimizes data ingestion along with algorithmic computation to increase throughput and scalability. Scikit-learn is a popular open-source Python package for machine learning and predictive data analysis tasks.

While Intel leads the efforts to write the software and develop the necessary interfaces, Lusch and Emani work closely with science teams, including projects supported by the Aurora Early Science Program, that use machine learning algorithms in their research and are interested in integrating their codes with these libraries for deployment on Aurora. This collaborative work helps ensure the machine learning tools will meet the needs of the scientific community.

When the work began, oneDAL and scikit-learn both entirely lacked GPU implementations. Lusch and Emani test the software on early hardware at Argonne's Joint Laboratory for System Evaluation. They meet regularly with Intel to provide input and feedback on advanced features, scaling across GPUs, scaling across nodes, APIs, and general performance.

The ALCF team's work also includes prioritizing support for the most widely or heavily used algorithms; enabling distributed implementation across multiple GPUs and at full scale on Aurora; facilitating interoperability with other data science libraries; and enabling the use oneDAL from multiple programming languages, such as Python and C++.

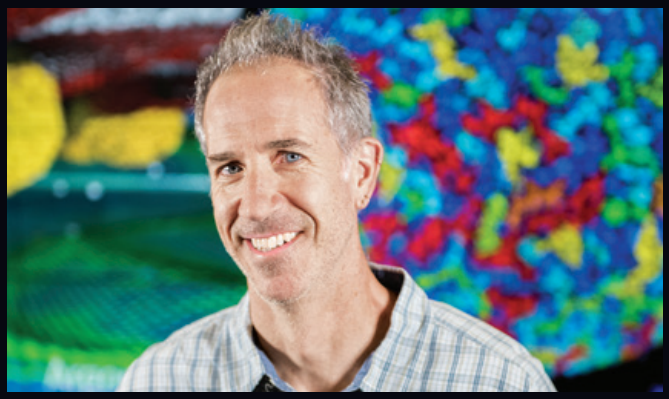

ALCF's Joseph Insley is helping to ready Aurora for scientific visualizations.

ALCF's Brice Videau is working to maximize the portability of HIP applications for exascale systems.

In Situ Visualization and Analysis

Joseph Insley, lead of the ALCF's Visualization and Data Analysis team, is helping to enable in situ visualization and analysis capabilities on Aurora. He and his team work closely with Intel's advanced rendering team, who are responsible for preparing the oneAPI Rendering Toolkit for Aurora. The toolkit is being ported to Intel GPUs to provide maximum rendering performance and reduce the need for data movement between different systems for different tasks (many standard visualization algorithms are implemented on CPUs).

The ALCF team is involved in several efforts related to the development of in situ frameworks, and collaborate with a number of DOE partners, including Lawrence Berkeley, Lawrence Livermore, Los Alamos, Oak Ridge, and Sandia National Laboratories, as well as corporate partners Kitware and Intelligent Light. Additionally, Insley and colleagues work with a number of application teams to instrument their codes for in situ visualization and analysis.

Insley is also actively involved in collaborating with teams from DOE's Exascale Computing Project (ECP). As part of ECP Software Technologies, the VTK-m and ALPINE projects aim to implement scientific visualization algorithms and in situ infrastructures on exascale supercomputers. He and his team work closely with these two projects, porting their algorithms and infrastructures to Aurora. They are also collaborating with the ECP ExaSky Application Development project, where they are preparing tools to visualize massive cosmology datasets generated on Aurora and other exascale machines.

HIP on Aurora

Brice Videau, a computer scientist at ALCF with years of experience in code generation, optimization, and auto-tuning high-performance computing, is leading efforts to maximize the portability of HIP applications being developed for the next generation of leadership computing systems, including Argonne's forthcoming exascale Aurora system.

Already established as the default programming model for AMD GPUs, many prominent applications are expected to adopt HIP as a programming model as they target deployment on upcoming exascale systems. Furthermore, HIP provides support for NVIDIA platforms, enabling HIP application portability across AMD and NVIDIA hardware. With HIP a likely programming model for numerous exascale codes, Videau's team wants to ensure that HIP applications are as portable as possible.

The HIPLZ project, to this end, aims to investigate the best way to enable native support for HIP applications on Aurora. Videau's team ultimately leverages Level Zero, a new API that Intel is developing for its GPUs, including the Ponte Vecchio accelerators that will drive Aurora.

As part of the team's development strategy, Videau and colleagues are building on the work of an earlier project that developed HIPCL, a library which enables HIP to run on top of OpenCL. Because Aurora will feature Level Zero as its primary low-level GPU interface, HIPLZ redirects the efforts of HIPCL in order to facilitate support for HIP applications on Level Zero.

ALCF's Kris Rowe (left) and Saumil Patel are working to bring the open-source OCCA library to Aurora.

ALCF's Corey Adams is helping to ensure deep learning frameworks are ready for Aurora.

OCCA Library

ALCF computational scientist Kris Rowe is leading collaborative efforts to bring OCCA—an open-source, vendor-neutral framework and library for parallel programming on diverse architectures—to Aurora.

Mission-critical computational science and engineering applications from the DOE and private industry rely on OCCA, which provides developers with transparency in the generation of raw backend code. NekRS, for example—a new computational fluid dynamics solver from the Nek5000 developers—is used simulate coolant flow inside of small modular reactors and design more efficient combustion engines.

Utilizing OpenMP and DPC++, Rowe worked with Saumil Patel of Argonne's Computational Sciences Division to establish initial benchmarks for various types of kernels before beginning development of the SYCL/DPC++ OCCA backend. Rowe and Patel meet biweekly with an Intel team to coordinate testing and performance aspects, and to drive the development of new features.

Programming models used by the OCCA backends continually evolve. Rowe currently surveys the latest DPC++, OpenMP, CUDA, and HIP specifications to identify common performance critical features—such as asynchronous memory transfers and work-group collectives—which have yet to propagate into the OCCA. Using this information, the team will propose extensions to the existing OCCA API for inclusion in future releases of the framework.

Deep Learning Frameworks

ALCF computer scientist Corey Adams is leading efforts to deploy advanced deep learning frameworks on Aurora, ensuring that key applications are fully performant on day one—that they run well and scale well in relatively bug-free implementations.

To this end, Corey and his colleagues selected a number of Argonne workloads that represent innovative AI-for-science approaches that will benefit from the Aurora architecture. In doing so, in order to grow application capabilities from a science perspective, they built on computer vision benchmarks established by Intel during the development of various deep learning and AI frameworks.

High-level frameworks in Python, such as TensorFlow and PyTorch, rely on Intel's deep neural network framework oneDNN for computationally intensive GPU processes such as convolution operations, the complex demands of which frustrate attempts at out-of-the-box performance. This necessitates extensive iterations of development and testing before an efficient kernel or source code can be produced.

Once optimal performance has been achieved on a single GPU, the Intel Collective Communications Library oneCCL helps deliver optimal performance on multiple GPUs by distributing optimized communications patterns to allocate parallel model training among arbitrarily many nodes. OneCCL and the synchronicity it encourages thereby enable tasks such as the uniform collection of gradients from a training iteration.

AURORA
EARLY SCIENCE PROGRAM
PROJECTS

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

Simulation Projects

Extending Moore's Law Computing with Quantum Monte Carlo

PI Anouar Benali, Argonne National Laboratory
SW QMCPACK

Extreme-Scale Cosmological Hydrodynamics

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

Extreme-Scale Unstructured Adaptive CFD

PI Kenneth Jansen, University of Colorado Boulder
SW PHASTA

High-Fidelity Simulation of Fusion Reactor Boundary Plasmas

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

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

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

Data Projects

Dark Sky Mining

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

Data Analytics and Machine Learning for Exascale Computational Fluid Dynamics

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

Exascale Computational Catalysis

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

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

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

Simulating and Learning in the ATLAS Detector at the Exascale

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

Learning Projects

Accelerated Deep Learning Discovery in Fusion Energy Science

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

Enabling Connectomics at Exascale to Facilitate Discoveries in Neuroscience

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

Machine Learning for Lattice Quantum Chromodynamics

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

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

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

Virtual Drug Response Prediction

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

SCIENCE

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

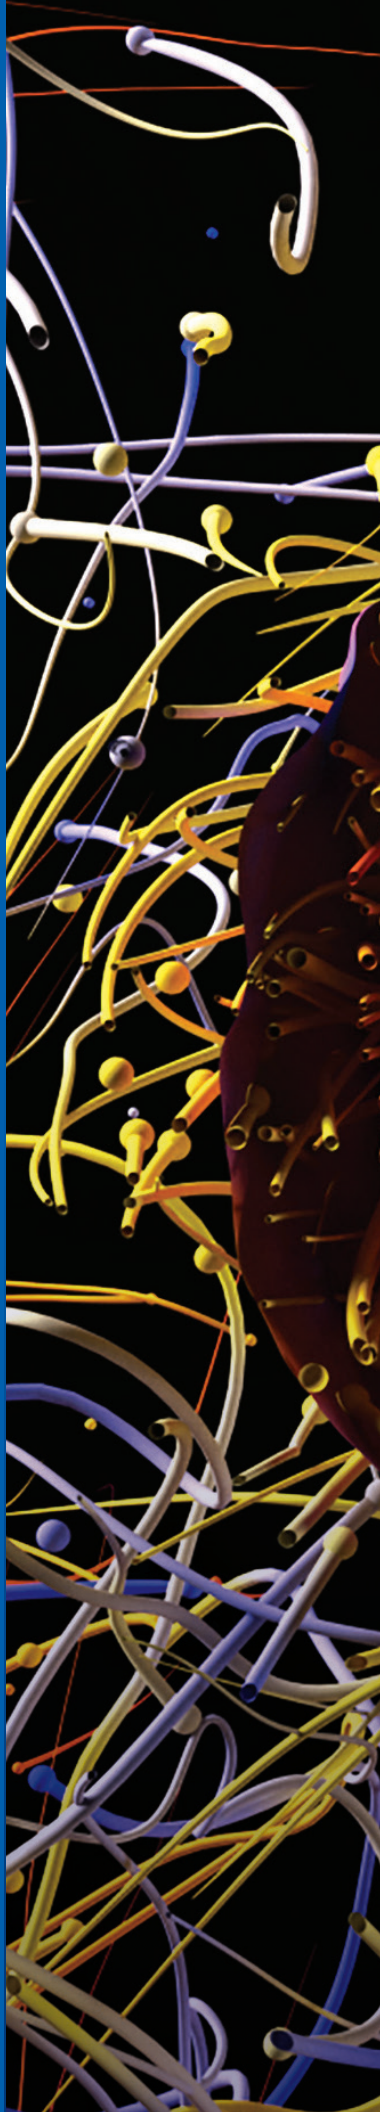

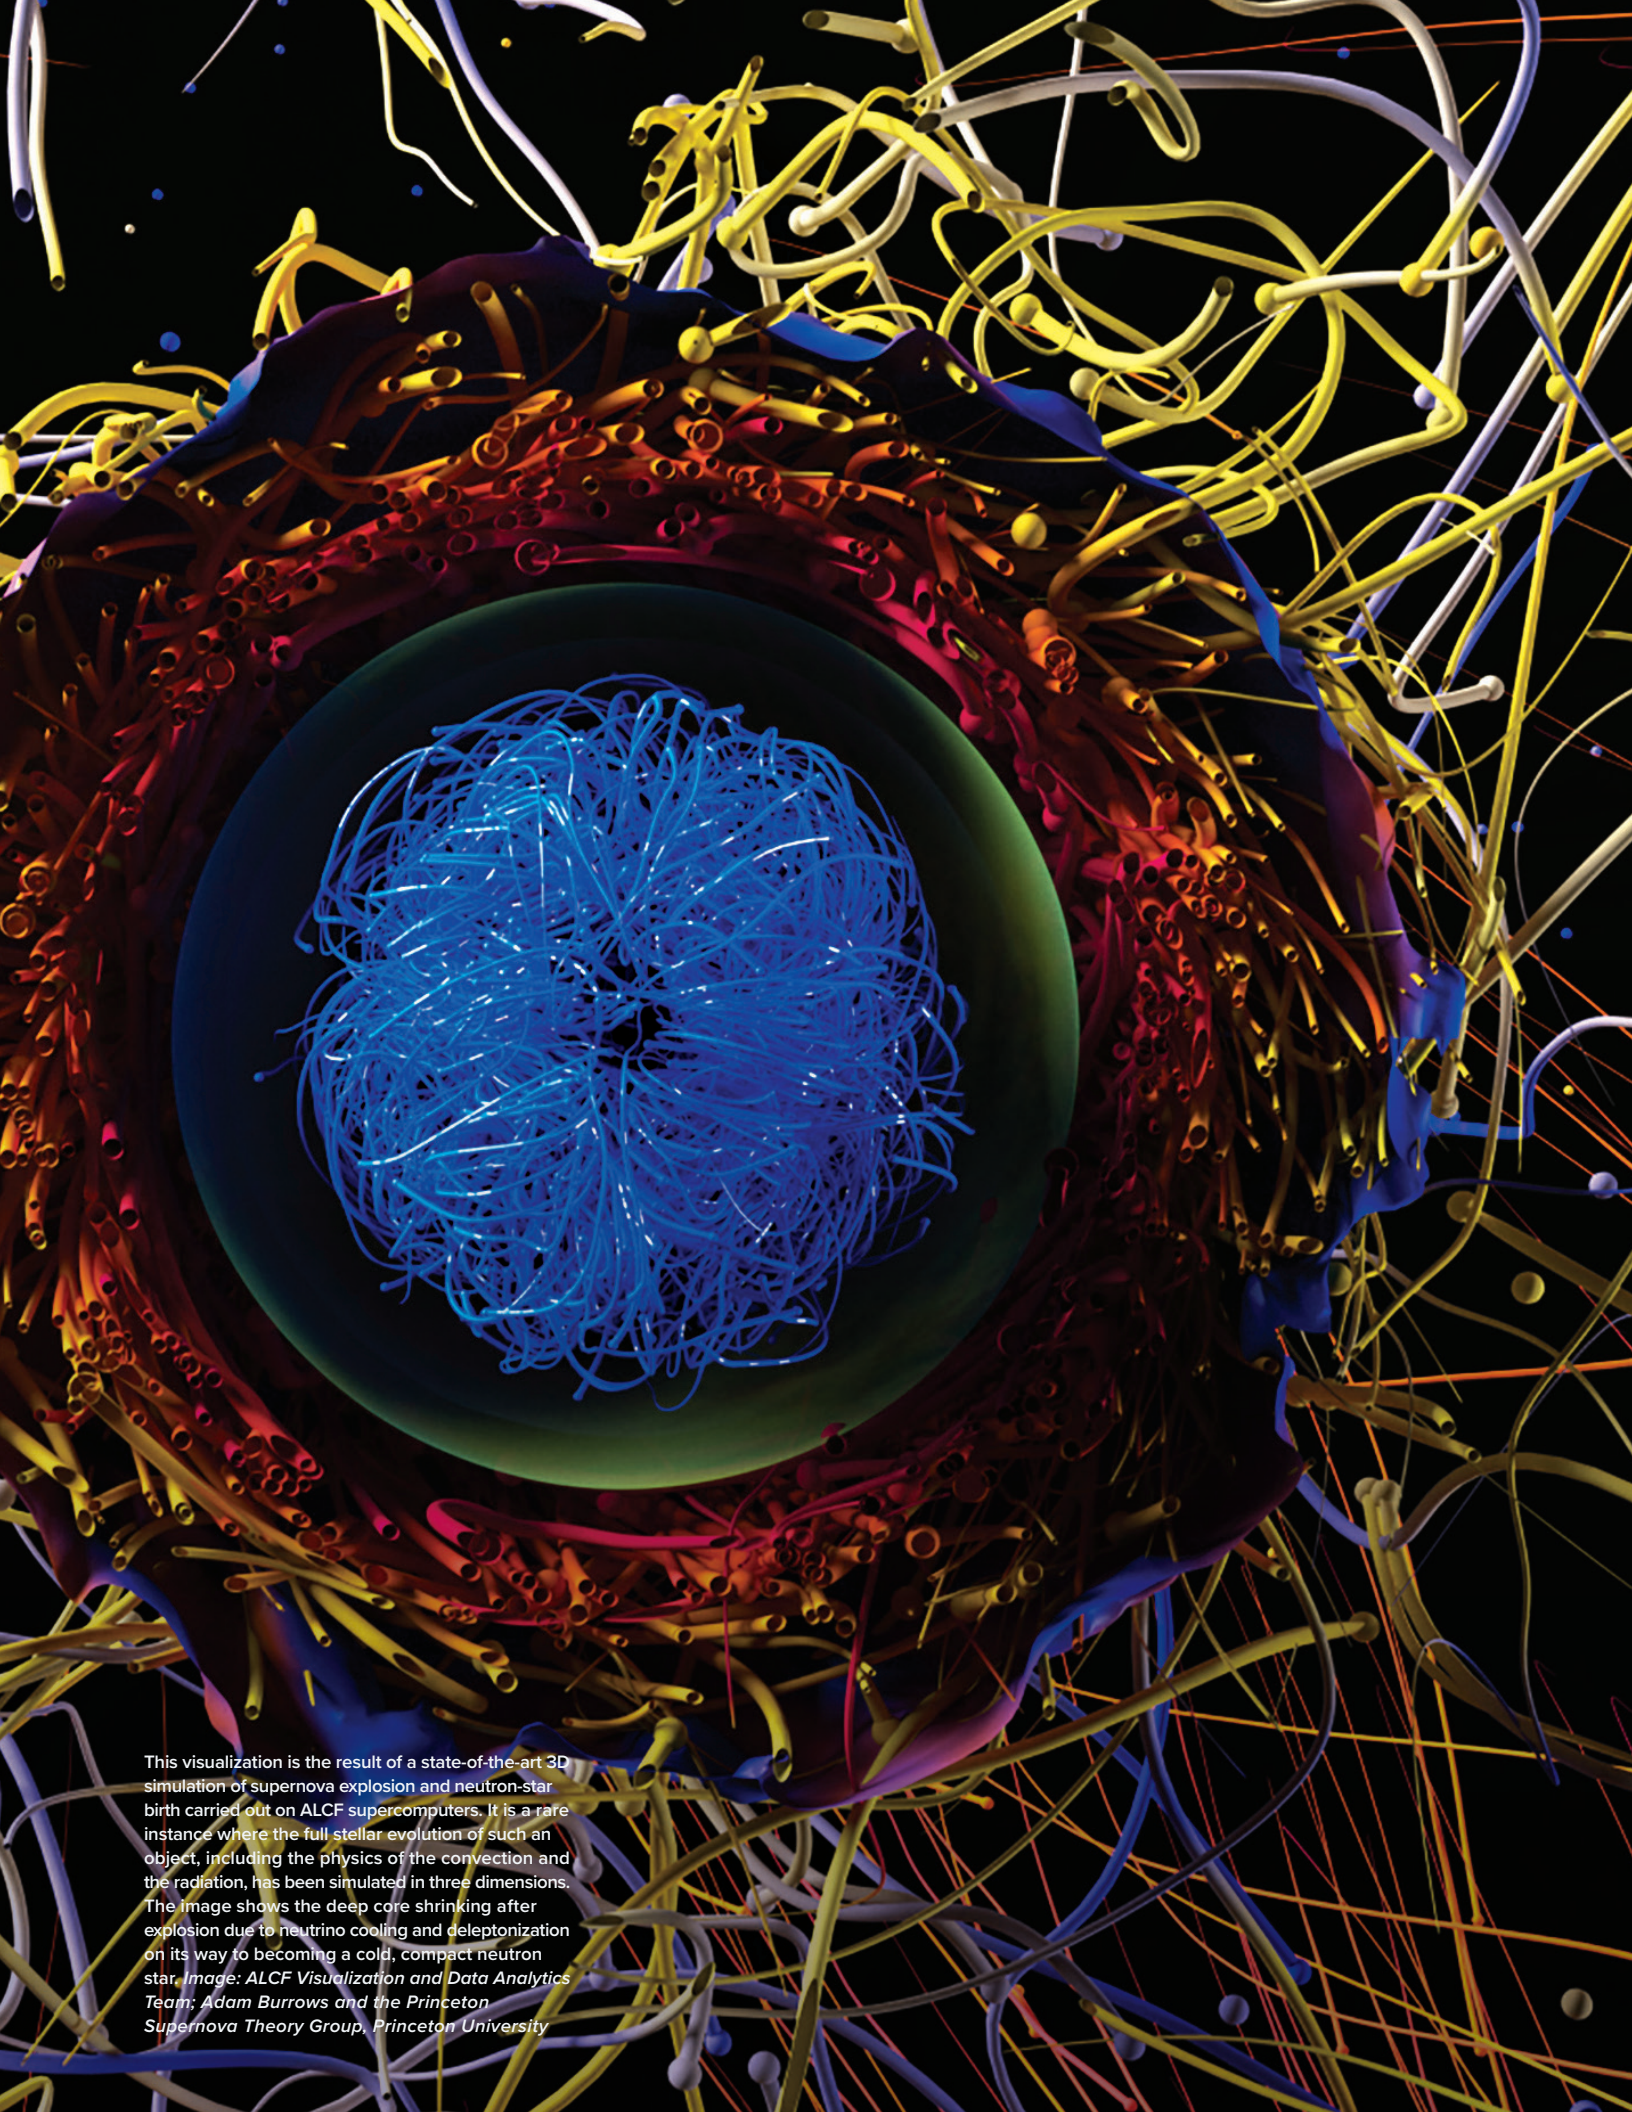

This visualization is the result of a state-of-the-art 3D simulation of supernova explosion and neutron-star birth carried out on ALCF supercomputers. It is a rare instance where the full stellar evolution of such an object, including the physics of the convection and the radiation, has been simulated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

Accessing ALCF Resources for Science

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

Researchers gain access to ALCF systems for computational science and engineering projects through competitive, peer-reviewed allocation programs supported by the DOE and Argonne.

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

INCITE/ALCC BY DOMAIN

2022 INCITE

17.8M NODE HOURS

A	Biological Sciences	3 %
B	Chemistry	19
C	Computer Science	–
D	Earth Science	1
E	Energy Technologies	4
F	Engineering	18
G	Materials Science	20
H	Physics	35

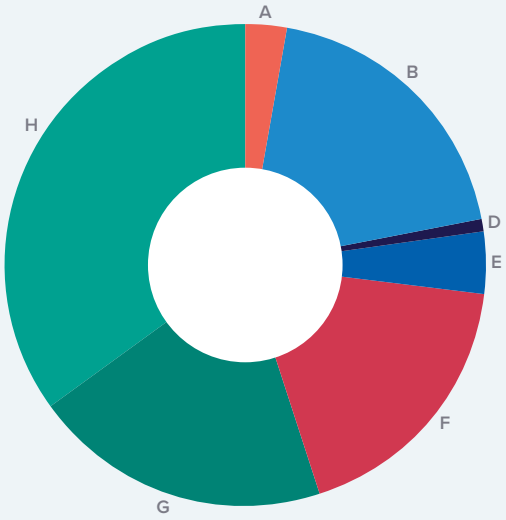

2022 ALCC

6.8M NODE HOURS

A	Biological Sciences	4 %
B	Chemistry	14
C	Computer Science	1
D	Earth Science	20
E	Energy Technologies	18
F	Engineering	6
G	Materials Science	11
H	Physics	26

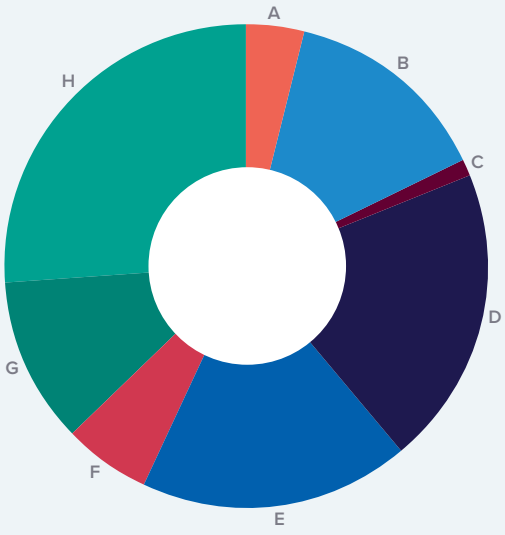

ALCC data are from calendar year 2022.

2022 Science Highlights

The ALCF user community accelerates scientific discovery across disciplines through research campaigns that deepen our understanding of the universe at all scales

ALCF users employ large-scale data analyses, learning processes, and simulations operating in tandem to conduct ambitious investigations and achieve scientific breakthroughs that would not otherwise be possible. From detailed atomic-level simulations to massive cosmological studies, researchers can probe extremely complex physical systems and processes that are too small or large, costly, or dangerous to study in a laboratory.

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

In the following pages, we present a selection of notable results from projects supported by ALCF allocation programs.

This year's highlights range from deep-learning-guided drug research to energy projects including battery and fusion investigations as well as sustainability studies.

Other efforts include work with ALCF resources to further integrate high-performance computing and experimental facilities such as Argonne's Advanced Photon Source. There are demonstrations of extreme-scale climate and earthquake models; simulations of oceans, turbulence, and combustion engines; and research into supernovae and other assorted astrophysical phenomena.

Biological Sciences | Simulation, Data, Learning

Combining Physics-Based Analysis with AI to Accelerate COVID-19 Drug Discovery

PI Tom Brettin and Hyunseung Yoo, Argonne National Laboratory
AWARD Director's Discretionary
SYSTEM Theta, Groq (ALCF AI Testbed)

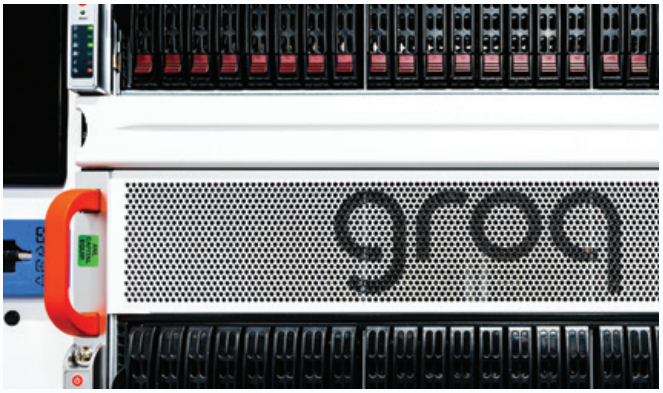

Argonne researchers used the ALCF AI Testbed's Groq system to accelerate the search for promising drug candidates to combat COVID-19. *Image: Argonne National Laboratory*

Since the beginning of the COVID-19 pandemic, scientists from Argonne National Laboratory have been applying AI and supercomputing resources to the search for drugs to combat the virus. These resources have been promising in vastly accelerating the drug discovery process. One of these efforts is using the ALCF AI Testbed's Groq system to accomplish within minutes searches which would otherwise have taken years.

CHALLENGE There are billions of possible small molecules to look through to see how well they would bind to different pockets on proteins of SARS-CoV-2. Each search can take days to complete. Physics-based computation is slower and less flexible than AI methods, while AI methods require a very large amount of training data. To overcome this, researchers are combining the two methods. Traditional GPU systems required reloading and context switching, which is inefficient, and this is addressed by the Groq system which avoids offloading.

APPROACH The researchers computed the binding affinity of small molecules with viral proteins, beginning with individual physics-based analysis on leadership-class supercomputers including ALCF's Theta system. They then used that data to train an AI algorithm to search for the lowest binding energy. To speed computations further, the team leveraged the ALCF AI Testbed's Groq system, which allowed them to make significantly more predictions per second. Using the GroqChip accelerator, the researchers were able to fit more models within the on-chip memory and use an end-to-end pipeline to avoid offloading quantization to the CPU.

RESULTS Because AI can quickly adapt to structures that it has never seen before, the team's efforts to combine physics-based methods with an AI algorithm resulted in the number of predictions run per second rising from 1,000 to 50,000. By employing the Groq system for tests on a large dataset of small molecules, the researchers found they could achieve 20 million predictions per second, vastly reducing the time needed for each search. The most promising drug candidates that were commercially available were sent to the University of Chicago's Howard T. Rickett's Laboratory for testing on human cells.

IMPACT Combining AI with physics-based analysis, the team was able to identify promising drug identification in a small fraction of the time typically required. Their approach resulted in a much more efficient route to drug discovery for COVID-19 and other life-threatening diseases.

Biological Sciences | Simulation, Data, Learning

COMPBIO2: Combining Deep Learning with Physics-Based Affinity Estimation 2

PI Peter Coveney, University College London
AWARD INCITE
SYSTEM Theta, ThetaGPU

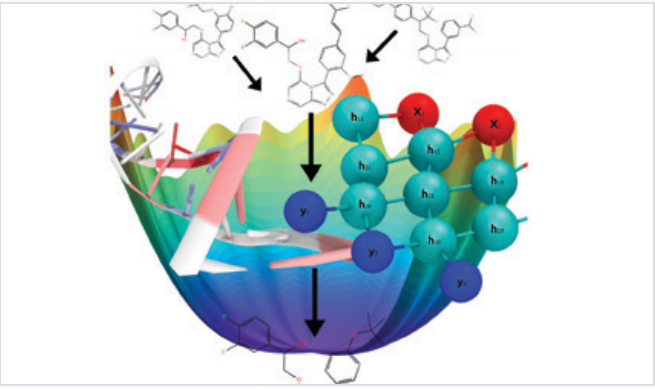

An artistic representation of the IMPECCABLE workflow that constructively combines physics-based simulations and machine learning approaches to accelerate the process of compound screening in drug discovery. *Image: Alex W. Wade, University College London*

Machine learning and physics-based methods combined can dramatically accelerate the process of drug discovery. With this INCITE project, researchers are developing computational methods to overcome severe bottlenecks in the pharmaceutical industry's current process and speed the search for drug candidates for SARS-CoV-2 and other viruses.

CHALLENGE Drug discovery requires exploring a vast chemical space of potential ligands to test binding potency with target proteins. *In silico* (computer simulation) physics-based methods play a key role in computational drug discovery but there is a very high computational cost when using relative binding free energy (RBFE) methods based on classical molecular dynamics (MD) which can lead to insufficient sampling of the chemical space (i.e., looking at too small of a subset of possibilities). The other, less computationally expensive method of finding drug binding potency is through deep learning, which requires a huge amount of training data. This project combines both to overcome limitations of each method.

APPROACH This project combines deep learning with physics-based methods and employs a range of ensemble simulation-based methods. The most accurate among them is called Thermodynamic Integration with Enhanced Sampling (TIES), which performs an ensemble of MD simulations at each alchemical intermediate state and then integrates the ensemble averaged energy derivative to obtain RBFEs. The pipeline using AI and simulation to look for drug candidates for COVID-19 is known as IMPECCABLE (Integrated Modeling PipelinE for Covid Cure by Assessing Better Leads). ALCF staff worked with the team to build NAMD code and helped implement the sampling enhancement method REST2 on NAMD CUDA.

RESULTS The team was able to simulate many more ligand transformations with greater accuracy, reproducibility, and precision than general simulation methodology. With TIES, the team replicated MD simulations in each free energy window five times with a different initial velocity allowing for statistically meaningful results and complete control of errors based on direct ensemble averages. They computed a broad selection of over 500 ligand transformations involving over 300 ligands binding to 14 different protein targets across a wide range of protein classes and found excellent correlation between experimental and prediction results. The systematic and extensive analysis of these transformations over such a large dataset is unprecedented.

IMPACT The team has demonstrated a method to overcome the limitations of physics-based drug simulations, allowing more accurate and reproducible results that can substantially reduce the time and cost of bringing new drugs to market. The team's method will also enable researchers to assess the resistance of COVID-19 protein variants and their impact on existing drugs and vaccines.

PUBLICATIONS

Bhati, A. P., and P. V. Coveney. "Large Scale Study of Ligand-Protein Relative Binding Free Energy Calculations: Actionable Predictions from Statistically Robust Protocols," *Journal of Chemical Theory and Computation* (March 2022), ACS Publications.

Wade A., A. P. Bhati, S. Wan, and P. V. Coveney. "Alchemical Free Energy Estimators and Molecular Dynamics Engines: Accuracy, Precision and Reproducibility," *Journal of Chemical Theory and Computation* (May 2022), ACS Publications.

Wan, S., A. P. Bhati, D. Wright, I. Wall, A. Graves, D. Green, and P. V. Coveney. "Ensemble Simulations and Experimental Free Energy Distributions: Evaluation and Characterization of Isoxazole Amides as SMYD3 Inhibitors," *Journal of Chemical Information and Modeling* (May 2022), ACS Publications.

Non-Synaptic Communication and Axonal Remodeling After Exposure to Cocaine

PI Bobby Kasthuri, University of Chicago
AWARD Director's Discretionary
SYSTEM Cooley

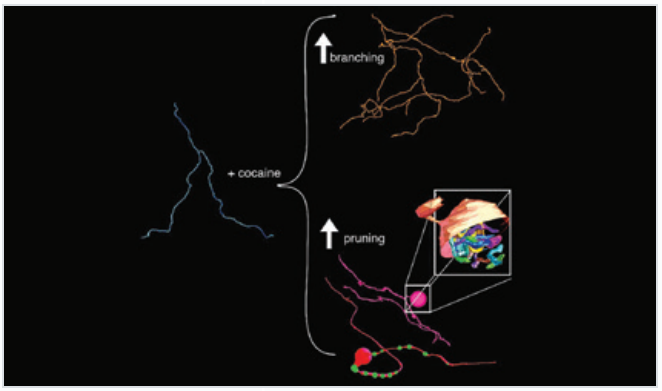

A connectomic analysis of the brains of mice treated with cocaine shows that dopamine axons undergo two major anatomical remodeling events: 1) axons increase the average number of branches they make (top image), 2) while simultaneously pruning, or removing, existing axons. *Image: Gregg Wildenberg, University of Chicago/Argonne National Laboratory*

Dopaminergic (DA) neurons exert profound influences on behavior including addiction. However, how DA axons communicate with target neurons and how those communications change with drug exposure remains poorly understood. A research team led by Argonne National Laboratory is leveraging experimental techniques and supercomputing resources to shed light on DA neuron structure, addiction, and the brain's ability to recover.

CHALLENGE Collecting and analyzing large volumes of brain tissue with electron microscopy (EM) has been a difficult and laborious task but recent advances in automized data collection and algorithms are helping to streamline the analysis of terabytes of data. Not much is known about the ultrastructural characteristics of potential DA synapses, where they are and how frequently they occur, so analyzing DA circuitry requires specific labeling.

APPROACH This project combines recent advances in serial EM (connectomics) and genetic labeling approaches for EM to examine dopamine axon wiring in mice. The team looked for structural changes in DA pathways resulting from brief exposure to cocaine to test the sensitivity of large-volume EM to these changes. The 1.5 terabytes of images collected with Argonne's serial electron microscope were digitally reassembled using the ALCF's Cooley visualization cluster. This process creates a 3D volume that allows researchers to identify and trace different anatomical features of the dopamine neurons, which, until recently, had proven something of a challenge.

RESULTS In a study published in *eLife*, the team detailed, for the first time, the specific changes that occur in the brains of mice exposed to cocaine. Their research revealed that individual DA axons have different varicosity types based on vesicle contents. They also discovered that days after brief exposure to cocaine, DA axons had extensively branched and formed blind-ended 'bulbs' filled with mitochondria that had more than doubled in length. In addition, the team found that DA varicosities lack features of classical synapses, they have a high degree of plasticity, and after exposure to cocaine they are extensively remodeled.

IMPACT This project reveals some of the physical basis of addiction in physical changes in dopamine circuits. Combining electron microscopy and digital reassembly allows physical changes in the brains of mice to be observed with a high degree of sensitivity in a large amount of brain matter. While the study has helped elucidate questions of form, function, and dynamics in the dopamine system, it also presents important new questions related to repeated exposure and addiction, as well as treatment and recovery.

PUBLICATIONS
Wildenberg, G., A. Sorokina, J. Koranda, A. Monical, C. Heer, M. Sheffield, X. Zhuang, D. McGehee, and B. Kasthuri. "Partial Connectomes of Labeled Dopaminergic Circuits Reveal Non-Synaptic Communication and Axonal Remodeling After Exposure to Cocaine," *eLife* (December 2021) eLife Sciences Publications Ltd..

Biological Sciences | Simulation, Data, Learning

Observing the SARS-CoV-2 Replication-Transcription Machinery in Action

PI Arvind Ramanathan, Argonne National Laboratory
AWARD Director's Discretionary
SYSTEM Theta, Cerebras (ALCF AI Testbed)

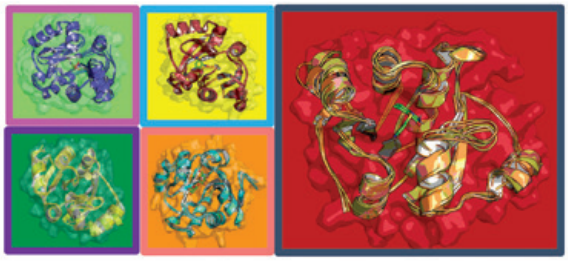

AI-driven MD simulations provide insights into how different ligands modulate the binding region of the viral ADP-ribose-1"-phosphatase protein. Ligands are shown in stick like representation and the protein is shown as a cartoon ensemble Note that each ligand has an effect on distinct regions of the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

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

CHALLENGE Collaborating closely with their colleagues at Argonne's Advanced Photon Source, the researchers leverage high-performance computers to examine the intricacies of SARS-CoV-2 self-reproduction. The coronavirus uses a precisely coordinated process known as the replication-transcription complex to reproduce at high speed when it invades a host's cells. Beyond eluding simple observation, the process represents a system of some 2 million atoms, making its study a tall order.

APPROACH Initial data for this work were obtained from cryo-electron microscopy, a technique that flash-freezes molecules and bombards them with electrons to generate 3D images. In performing their simulations, the team used a hierarchical AI framework running on the workflow engine Balsam to distribute their code across four leadership computing systems, one of which was Theta. Access to the ALCF AI Testbed enabled the team to train deep learning models that were coupled with the leadership systems. The approach built on the team's previous Gordon Bell-winning simulation of spike protein behavior in the SARS-CoV-2 virus.

RESULTS The team's innovative workflow bridges gaps in resolution and timescale that enable it to capture important dynamics in the SARS-CoV-2 replication transcription process. The results were captured in a paper that was a finalist for the Gordon Bell Special Prize for High Performance Computing-Based COVID-19 Research, as was another publication from the same team elucidating the virus's infiltration of the human immune system. The simulations performed provided information uncapturable by cryo-electron microscopes, meaningfully reconstructing viral motions that would otherwise remain incomprehensible.

IMPACT The team's research has grown our understanding of the SARS-CoV-2 virus and the COVID-19 disease. Elucidating the virus's behavior and mechanisms can help identify effective new treatments and therapies for combatting its spread and severity, and the research potentially could lead to the design of new generative models based on reinforcement learning for both small molecules and antibodies.

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* (April 2021), SAGE Publishing.

Trifan, A., D. Gorgun, Z. Li, A. Brace, M. Zvyagin, H. Ma, A. Clyde, D. Clark, M. Salim, D. J. Hardy, T. Burnley, L. Huang, J. McCalpin, M. Emani, H. Yoo, J. Yin, A. Tsaris, V. Subbiah, T. Raza, J. Liu, N. Trebesch, G. Wells, V. Mysore, T. Gibbs, J. Phillips, S. C. Chennubhotla, I. Foster, R. Stevens, A. Anandkumar, V. Vishwanath, J. E. Stone, E. Tajkhorshid, S. A. Harris, and A. Ramanathan. "Intelligent Resolution: Integrating Cryo-EM with AI-Driven Multi-Resolution Simulations to Observe the SARS-CoV-2 Replication-Transcription Machinery in Action," *International Journal of High Performance Computing Applications* (preprint), SAGE Publishing.

Chemistry | Simulation

Microscopic Insight into Transport Properties of Li-Battery Electrolytes

PI Wei Jiang, Argonne National Laboratory
AWARD ALCC
SYSTEM Theta

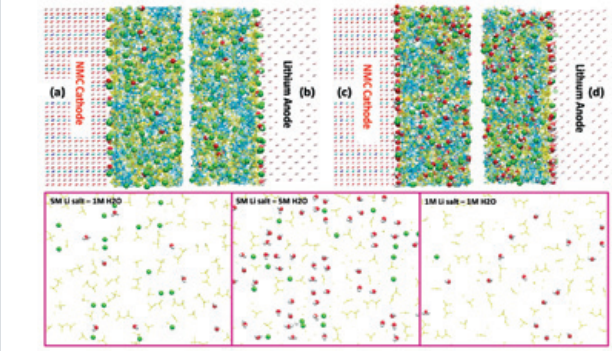

Simulated cluster of water (red) and lithium ions (green) at electrode surfaces under various water concentrations. Battery performance is retained when water molecules remain isolated. *Image: Wei Jiang, Argonne National Laboratory*

There is an increasing worldwide demand for high energy density batteries. The development of low-cost, safe, rechargeable batteries with high voltage, capacity, and rate capability is important for several applications, including electric vehicles. To aid in the discovery of novel battery materials, researchers from Argonne National Laboratory are using ALCF computing resources to enable the design of improved electrolytes for high-voltage lithium-ion batteries.

CHALLENGE One of the challenges in developing advanced batteries is the need for new insights into how to improve electrolyte behavior under increasingly extreme performance demands. However, at high potentials or high voltages, the state-of-the-art organic carbonate-based electrolytes tend to be oxidatively decomposed at the cathode surface causing gassing issues, low Coulombic efficiency, transition metal ion dissolution, and rapid capacity fade of the full cell. Moreover, these conventional electrolytes are extremely flammable due to their high vapor pressure and low flashpoint, which leads to potential safety issues especially when applied in electric vehicles. To address these issues, the Argonne team is exploring ionic liquids as a potential alternative due to their intrinsic physical properties, such as low vapor pressure, non-flammability, wide electrochemical window, and high ionic conductivity.

APPROACH This ALCC project is employing a joint experimental and computational approach that involves running the NAMD molecular dynamics code on Theta to explore the nanostructural organization at electrolyte-electrode interfaces, and to shed light on the transport properties and desolvation/solvation kinetics of charge carriers. The team is leveraging molecular dynamics methodologies, such as Hamiltonian Annealing and sampling enhanced free energy calculations, to provide insights into the underlying chemical mechanisms that give rise to favorable battery performance.

RESULTS In a study published in *ACS Applied Materials & Interfaces*, the team detailed the development of a new battery electrolyte that can hold 1,000 times more water than conventional electrolytes. Their simulations revealed that the material, composed of a lithium salt and an ionic liquid, could separate and bind up water, sequestering individual water molecules. They found the key is that the water molecules do not gather into “puddles” and thereby lose reactivity. The new electrolyte eliminates the critical moisture controls required for a state-of-the-art carbonate electrolyte, electrode, separator, and cell assembly, which would significantly reduce the cost of manufacturing of such batteries.

IMPACT As part of their work, the researchers identified a potential pathway for battery manufacturers to incorporate water in the fabrication process, allowing for lower cost, more environmentally friendly batteries. Ultimately, the team is working to develop a computer-aided electrolyte design protocol to enhance future high-throughput prediction and screening studies aimed at discovering new battery materials.

PUBLICATIONS

Liu, Q., W. Jiang, Z. Yang, and Z. Zhang. “An Environmentally Benign Electrolyte for High Energy Lithium Battery Materials,” *ACS Applied Materials & Interfaces* (November 2021), ACS Publications.

Computer Science | Data, Learning

Bridging Data Center AI Systems with Edge Computing for Actionable Information Retrieval

PI Zhengchun Liu, Argonne National Laboratory
AWARD Director's Discretionary
SYSTEM Cerebras, SambaNova (ALCF AI Testbed)

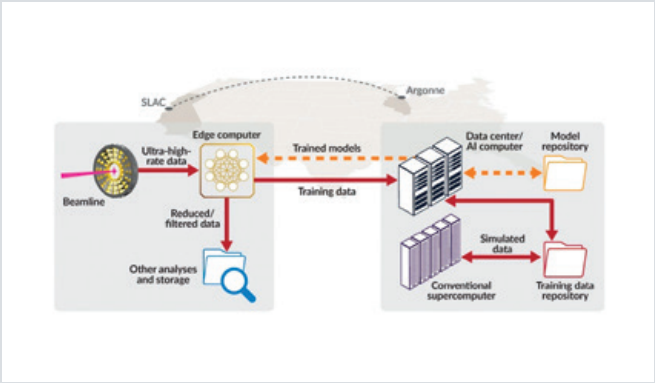

The actionable information architecture uses an edge computer co-located with the experimental apparatus for rapid reduction/filtering of high-velocity data; the ML model is trained on a remote data center AI computer (here, at the ALCF). Solid crimson arrows are data flows; dashed orange are models. *Image: Zhengchun Liu, Argonne National Laboratory*

The increased coherence and brilliance of DOE's next-generation x-ray light sources, such as the upgraded Advanced Photon Source and the Linac Coherent Light Source, will lead to increasingly large and rich datasets produced at high data rates. To keep pace with the deluge of scientific data, a team of researchers is looking to machine learning methods and edge computing to help with tasks such as data reduction and providing insights with sufficiently low latency to steer future experiments.

CHALLENGE The increasing complexity and velocity of scientific x-ray data means that extracting the desired physical information becomes a considerable computing challenge. Conventional analytical methods take too long to run on individual processors, particularly when information is needed rapidly. Machine learning (ML)-based surrogate models can provide an effective alternative, offering a means to approximate the results of an analytical method with high accuracy. One key challenge, however, is training the ML models with sufficient rapidity so that they can be deployed within useful timescales.

APPROACH With this project, the team explored how specialized data center AI (DCAI) systems and a geographically distributed workflow can help facilitate ML model training on remote DCAI systems and model deployment on edge devices. The researchers used funcX, Globus Flows, and Globus file transfer services to build a workflow to automatically train deep neural networks with given data or simulations using DCAI and supercomputing resources. To evaluate the automated workflow, the researchers compared two scenarios: one in which model training was performed remotely using the Cerebras or SambaNova AI accelerators at the ALCF AI Testbed, and one in which training was performed locally on an NVIDIA V100 GPU co-located with the experiment.

RESULTS The team's workflow achieved a turnaround time between initiating and model delivery to edge host of less than 151 seconds including all data movement overhead. As a comparison, it takes approximately 17 minutes (though no cost on data movement) when training the same model using one high-end GPU that is deployable within the data acquisition machine. The team's research proves the feasibility of using powerful yet remote DCAI systems to enable rapid (re)training of deep neural networks for production use on edge devices. Their work was recognized with the Best Paper Award at the XLOOP Workshop at SC21.

IMPACT This project demonstrated how specialized AI systems can be used to quickly train machine learning models through a geographically distributed workflow. The team's approach can help accelerate scientific discoveries at x-ray light sources by employing machine learning at the edge to facilitate data reduction, feature detection, and other data-intensive tasks.

PUBLICATIONS
Liu, Z., A. Ali, P. Kenesei, A. Miceli, H. Sharma, N. Schwarz, D. Trujillo, H. Yoo, R. Coffee, N. Layad, J. Thayer, R. Herbst, C. Yoon, and I. Foster. "Bridging Data Center AI systems with Edge Computing for Actionable Information Retrieval," *2021 3rd Annual Workshop on Extreme-Scale Experiment-in-the-Loop Computing* (XLOOP), (November 2021), IEEE.

Earth Science | Simulation

Convection Permitting Climate Scale Simulations for Extreme Event Modeling and Analysis

PI Rao Kotamarthi, Argonne National Laboratory
AWARD ALCC
SYSTEM Theta, ThetaGPU

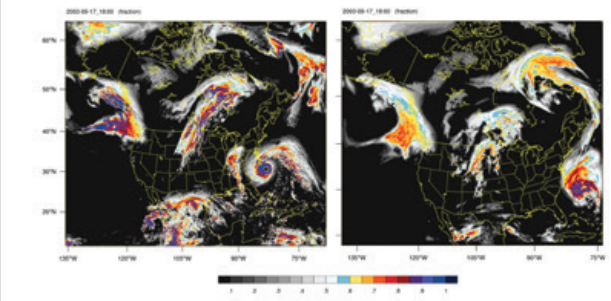

Mid-tropospheric cloud fraction (shaded; 0.05 interval) derived from (a) ERA5 driven 4 km convection-resolved simulation and (b) NCEP-R2 driven 12 km simulation on September 17, 2003. The detailed features of cloud patterns and a more realistic representation of a simulated hurricane are apparent in the 4 km simulations. *Image: V. Rao Kotamarthi, Chunyong Jung, and Jiali Wang, Argonne National Laboratory*

As greenhouse gas emissions cause the world to warm rapidly, the frequency and intensity of extreme events, such as wildfires, storms, heatwaves, floods, and droughts, have increased over the last several decades. To aid in the development of adaptation and mitigation strategies that address the impacts of climate change, researchers from Argonne National Laboratory are harnessing ALCF supercomputers to provide future projections of extreme weather events and assess their risks for local and regional levels across North America.

CHALLENGE The convection parameterizations used in most global or regional climate models introduce uncertainties in model projections. Making simulations at spatial scales where much of the convection is resolved (so-called convection-permitting, meaning there is no need for a parametrization) is found to lead to lower model bias when compared to observations for the recent past.

APPROACH By running the Weather Research and Forecasting (WRF) model on ALCF's Theta and ThetaGPU systems, the team is working to (a) build a large dataset of future projections of extreme weather events, (b) understand the environmental drivers that contribute to these extremes, and (c) estimate the risk from these events at regional and local scales with sufficient accuracy. To achieve these objectives, they are carrying out convection-permitting regional climate simulations at a few kilometer scales covering multiple decades and forming a set of ensembles that covers selected time slices in regions across North America. In addition, they are simulating selected past hurricanes and extratropical storms and projecting their responses in a changing climate using GPU-accelerated WRF on ThetaGPU, resulting in speedups of approximately 10x compared to the CPU version.

RESULTS The team has completed the current 20-year simulation at a spatial resolution of 4 km driven by ERA5 reanalysis data. The 4 km simulations yield more realistic representations of historical climate observations. For instance, it captured Hurricane Isabel's (2003) track, maximum sustained wind speed, and minimum sea level pressure in a very realistic manner that are nearly identical to observations, while a spatial resolution of 12 km simulation poorly reproduced the storm characteristics. The global climate model-driven simulations for multiple 20-year time slices are in progress.

IMPACT The team's work represents one of the first attempts to provide future projections of extreme events and assess their risks on a climate time scale using a spatial resolution that is actionable at local and regional levels. The resulting data and analyses will enable policymakers to develop effective adaptation and mitigation strategies, and pave the way for new climate change impact assessment studies.

Earth Science | Simulation

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

PI Christine Goulet, University of Southern California
AWARD INCITE
SYSTEM Theta

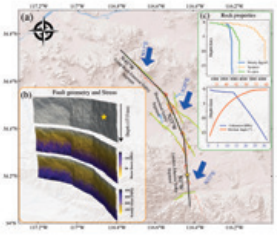

Simulation model setup (a) showing observational and model fault details with arrows indicating orientations of principal stresses, (b) fault-plane model and epicenter (star), and (c) rock properties. *Image: Christine Goulet and Yongfei Wang, University of Southern California*

Ground motions and fault displacements from earthquakes threaten infrastructure. Computer-based simulations can provide insight on those earthquake hazards to help develop mitigation strategies. Led by researchers from the Southern California Earthquake Center (SCEC), this INCITE project seeks to improve earthquake-modeling capabilities by integrating physics-based models into research software, validating ground motion simulation and hazard products, while improve code performance.

CHALLENGE Civil infrastructure is at risk of damage due to earthquake ground motions and fault displacement. In areas close to the fault (within 10–20 km), empirical models for both hazards are poorly constrained due to the paucity of data. This INCITE project is using DOE supercomputing resources to simulate ground motion and displacement using dynamic rupture codes that include complex physics to produce high-frequency simulations. The results are cross-validated with existing empirical data and used to extrapolate earthquake effects where no data exist. The researchers are looking to improve the accuracy of rupture simulations by adding features such as complex fault geometry and nonlinearity in the geological models.

APPROACH The team's work at the ALCF is largely driven by two open source codes, Waveqlab3D and SORD, which are used to simulate earthquake rupture dynamics and seismic wave propagation. The researchers are using Theta to look at fault displacement from dynamic rupture simulations for earthquakes with a magnitude greater than seven with different initial conditions. They are also generating broadband dynamic rupture sources with off-fault plasticity and geometrical fault complexity. With this work, the team is creating a database of slip-strike earthquake mechanisms, comparing broadband ground motions with empirical trends predicted by empirical ground motion models (GMM). They also are modeling ruptures across dipping faults, which have unique geometrical complexities. ALCF staff helped with scheduling jobs, and compiling and running SORD with Cray MPI.

RESULTS In one study, the INCITE team used Theta to simulate the 1992 Landers, CA earthquake by defining a preferred model that reproduced first-order fault displacement metrics including the on-fault partition of the total displacement. This has demonstrated that all the physics important to modeling the earthquake were properly captured, showing that dynamic rupture modeling can be used to assess fault-displacement hazard on a larger scale. The model reproduced intermediate and large-scale features including displacement, off-fault deformation, fault zone width and peak displacement.

IMPACT Physics-based dynamic rupture models and simulations inform earthquake hazard analyses and assessments, which help prepare for earthquakes and enable mitigation strategies that can save lives and alleviate the societal and economic damage caused by earthquakes.

PUBLICATIONS

Wang, Y., and C. Goulet. "Validation of Fault Displacements from Dynamic Rupture Simulations against the Observations from the 1992 Landers Earthquake," *Bulletin of the Seismological Society of America* (October 2021), SSA.

Earth Science | Simulation

Multiscale, Multiphysics Ensemble Simulations for Aerosol-Cloud Interactions

PI Po-Lun Ma, Pacific Northwest National Laboratory
AWARD ALCC
SYSTEM Theta

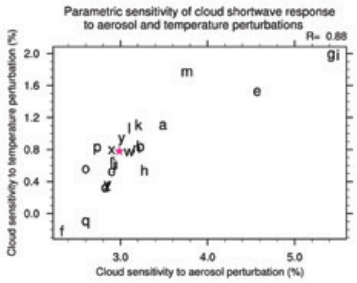

Scatter plot showing the sensitivity of cloud shortwave radiative effects to temperature and aerosol perturbations. Letters indicate the members of the perturbed physics ensemble. *Image: Po-Lun Ma, Pacific Northwest National Laboratory*

Climate change will have increasingly extreme effects in years to come. The U.S. Department of Energy is developing the Energy Exascale Earth System Model (E3SM) to predict the effects of climate change including extreme weather events, sea-level rise, coastal flooding, droughts, temperature changes, and change in atmospheric patterns to be able to respond to them. To help advance the development of E3SM, researchers from Pacific Northwest National Laboratory leading an ALCC project to increase our understanding of the role of aerosols and aerosol-cloud interactions (ACI) in the evolution of the Earth system.

CHALLENGE The processes affecting ACI are extremely complex and after decades of research ACI remains a source of uncertainty in climate predictions. This is largely due to coarse model resolution and incomplete representation of chemical and physical processes. To address these scientific challenges, researchers are developing a variable resolution capability for E3SM that will resolve atmospheric processes from the convection-permitting scale to the planetary scale.

APPROACH This project is leveraging ALCF's Theta supercomputer to carry out ensemble simulations with variable resolutions and perturbed ACI processes to assess the atmospheric characteristics and processes across scales; and provide insights into the robustness of ACI parameterization across scales. The PINACLES (Predicting INteractions of Aerosol and Clouds in Large Eddy Simulation) model is being used for simulations with large-eddy simulation (LES) resolution (e.g., 20 m horizontal grid spacing). The nonhydrostatic version of the E3SM with a regionally refined mesh is generating simulations with resolutions ranging between convection-permitting and mesoscales (e.g., ranging from 3 to 100 km horizontal grid spacing). The team's high-resolution model results will help researchers understand features at the subgrid scale of a coarse-resolution E3SM and provide guidance for parameterization improvements.

RESULTS The team's perturbation simulations testing model physics have shown that cloud feedbacks and ACI are sensitive to physical parameterization assumptions in microphysics, macrophysics, turbulence, and deep convections. Uncertainties in process representations in Earth system models contribute to uncertainties in cloud feedback and ACI, which are two major sources of uncertainty in climate prediction. This study quantifies the uncertainty and identifies the processes that affect cloud feedback and ACI, jointly or independently. The team's next step is to identify the physical mechanism that drives this correlation.

IMPACT The team's work aims to advance our understanding of the structural characteristics of ACI across scales, and therefore will provide unprecedented realism in climate predictions. Their development of improved aerosol and ACI modeling capabilities suitable for E3SM will help enable actionable predictions of Earth system variability and change, with an emphasis on the most critical scientific questions facing the nation.

Energy Technologies | Simulation, Data

Automatic Building Energy Modeling (AutoBEM)

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

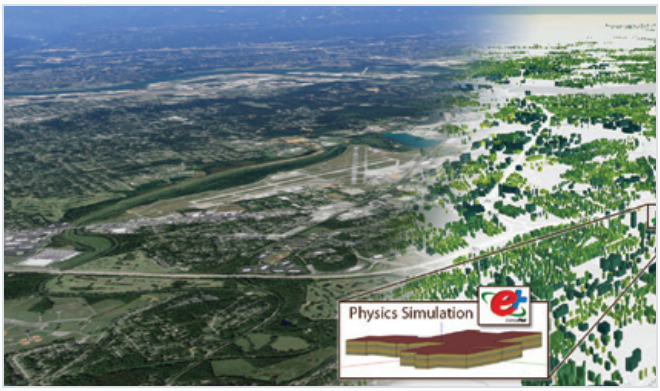

The Automatic Building Energy Modeling (AutoBEM) software is capable of detecting buildings, generating models, and simulating building energy use for every building in very large geographical areas. *Image: Joshua New, Oak Ridge National Laboratory*

Commercial and residential buildings consume nearly three-quarters of U.S. electricity. To help identify the potential energy, emissions, and cost reductions achievable with energy-efficient building technologies, researchers from Oak Ridge National Laboratory are using ALCF computing resources to model the energy use of every building in America.

CHALLENGE Urban-scale building energy modeling has grown significantly, allowing individual campuses and communities of buildings to be modeled. This has traditionally been limited to under 2,000 buildings, but the Oak Ridge team is advancing the capabilities of their AutoBEM suite to simulate all 125.7 million buildings in the United States. This is the first time that a nation-scale building modeling campaign has been attempted.

APPROACH The team expanded an existing HPC workflow to generate OpenStudio and EnergyPlus building energy models with building-specific geometry and characteristics synthesized from multiple datasets. In this way they generated unique models of almost every building in the United States. To manage the workload, the team divided the buildings into U.S. regions. They then took the files for each region and pre-processed them to be submitted to the queue on ALCF's Theta supercomputer. The output was over 80 terabytes of data, so the researchers employed a multistep process to enable postprocessing and analysis of the large dataset. The team's work at the ALCF continues with a 2022 INCITE award.

RESULTS In a study presented at the 2021 IEEE International Conference on Big Data, the researchers detailed their efforts to model 124.5 million buildings, or 98.9 percent of the nation's building stock. Using different iterations on reruns of processing, the building data yielded different degrees of success for different regions, with most regions having extremely high and consistent success. While the West North Central and Pacific regions (without California) had a much higher variation, this was likely due to the differences in the number of computationally complex buildings between regions. Moving forward the researchers are looking to make reliable predictions of processing times based on input parameters using artificial intelligence and comparing simulated data to measured energy use where the data are available.

IMPACT This project seeks to incorporate new data, methods, and stakeholder feedback to significantly improve crude, nation-scale building energy models toward investment-grade models. The team has made 122.9 million models freely available and is working with companies to analyze the models for business-relevant use cases while providing free data to help identify effective energy-savings towards a more sustainable and resilient built environment.

PUBLICATIONS

Berres, A. S, B. C. Bass, M. B. Adams, E. Garrison, J. R. New. "A Data-Driven Approach to Nation-Scale Building Energy Modeling," *IEEE International Conference on Big Data* (December 2021), IEEE.

New, J. R., B. Bass, M. Adams, and A. Berres. "Model America - Clark County (Vegas) extract from ORNL's AutoBEM (Version 1.1) [Dataset]," *Zenodo* (February 2021).

Deterministic Machine Learning for Latency-Critical Control Systems

PI Kyle Felker, Argonne National Laboratory
AWARD Director's Discretionary, Aurora Early Science Program
SYSTEM Groq (ALCF AI Testbed)

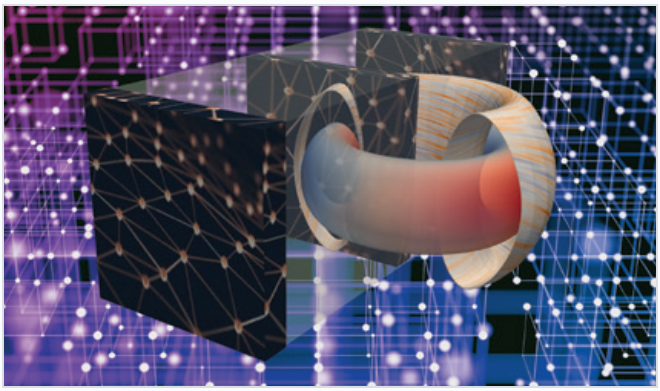

Researchers from Princeton Plasma Physics are developing the the Fusion Recurrent Neural Network (FRNN) code to predict disruptions in tokamak plasmas with unprecedented accuracy and speed. *Image: Eliot Feibush, Princeton Plasma Physics Laboratory*

As part of a project aimed at improving capabilities and mitigating large-scale disruptions in burning plasmas in tokamak systems such as ITER, Argonne researchers are working to leverage Groq AI platforms to accelerate deep learning-guided investigations into fusion control.

CHALLENGE The immense complexity of the equations that govern the behavior of tokamak fusion reactors creates myriad sources of potential instability that could prevent the fusion reaction from becoming self-sustaining. The team must incorporate deep learning models in real time; they have only a single millisecond to employ trained models to obtain a prediction to feed back to other fusion-reactor control systems.

The team intends to graduate use of the application to the prediction of instabilities in real-world plasma-discharge experiments being conducted as part of a path to viable fusion energy. Ultimately, the finalized workflow for the fusion application is to be one that connects Aurora to an AI machine—Aurora for training, and the AI machine for inference.

APPROACH The team uses ALCF supercomputing resources to train models, and through the ESP allocation is preparing to leverage Aurora for massively parallel training.

Given the significance of the inference problem, the team chose Groq as a platform on account of its architecture, which is based around the Tensor Streaming Processor (TSP). This unique architecture ensures fixed, predictable compute times for a key phase of deep learning (inference) that would vary in duration if carried out on CPU- and GPU-driven machines.

The TSP is designed to extend single core performance across multiple chips, can perform inference tasks quickly, and is deterministic with zero overhead for context switching.

RESULTS The researchers successfully produced benchmarks for a limited number of cases. As presented at the 2021 AI Summit, their results suggest that the TSP architecture might be able to outperform GPUs when running the tokamak application. Now afforded the ability to run ONNX files on the Groq TSP through a recently released compiler, the researchers are examining hundreds of different models as part of a comprehensive investigation into fusion control.

IMPACT Deepened understanding of the fusion control problem promises improved operational efficiency and safety, and stands to help bring fusion closer to becoming a viable source of energy. Real-time incorporation of deep learning models will help drive innovation in large-scale experiments performed at related facilities including particle accelerators and light sources.

Energy Technologies | Simulation

DNS and LES of Internal Combustion Engines to Understand Origins of CCV

PI Sibendu Som, Argonne National Laboratory
AWARD INCITE
SYSTEM Theta

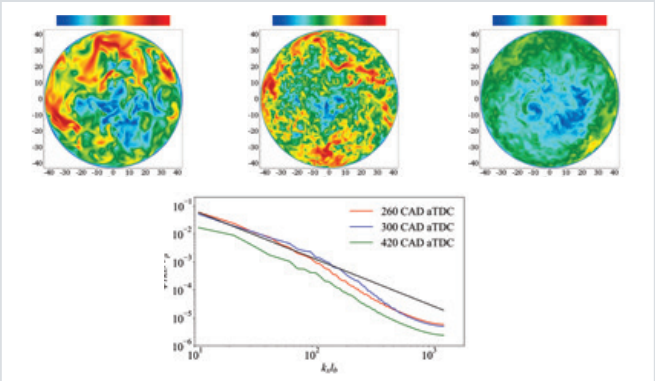

(Top) Distribution of the velocity field along a plane parallel to the piston showing the evolution of turbulent flow structures with increasing compression. (Bottom) 1D velocity spectra at a plane parallel to the piston showing the presence of more energy content in all the scales at higher compression. *Image: Muhsin Ameen, Argonne National Laboratory*

A key scientific challenge for internal combustion engine (ICE) is the understanding, modeling, and control of cycle-to-cycle variability (CCV) in engine performance, which can contribute to unevenness in the running of the engine, excessive engine noise and emissions, and potentially damaging engine knock. High-fidelity simulations can play a crucial role in providing insights into how in-cylinder flow, spray, and combustion processes contribute to CCV. With this project, researchers are using ALCF supercomputers to develop a detailed understanding of the origins of CCV in a modern laboratory-scale engine.

CHALLENGE Due to the extreme complexity of the physical processes and engine geometry, limited experimental data is available to enable scientific inquiry of CCV. In addition, the current state of simulations uses low-order methods and simple approximations to geometry (e.g., cylinders or boxes) which limit our understanding to cases that are not indicative of real engines. The Argonne team is working to overcome these obstacles by generating a high-fidelity numerical database of relevant engine conditions which exist in complex geometries that include spark plug, injector, valves, and bowl-in-piston features.

APPROACH To generate the database, the team is using ALCF computing resources to employ a hierarchy of one-of-a-kind direct numerical simulations (DNS) and high-fidelity large eddy simulation (LES) benchmarks that can be used for scientific discovery, model development, and validation. While the high-fidelity LES will provide "gold-standard" data to understand root causes for CCV, the DNS will provide "ground truth" data for portions of the cycles to improve heat-transfer and combustion models further. The team is carrying out a progression of cases with increasing complexity in physics and geometry for a laboratory-scale direct-injection spark-ignition test engine developed at Sandia National Laboratories.

RESULTS The team performed multi-cycle, wall-resolved LES of the Sandia DISI engine to understand relevant in-cylinder flow features that can affect CCV under lean/dilute conditions. This involved a combination of open-cycle engineering LES using the commercial CFD code, CONVERGE, and closed-cycle, wall-resolved high-fidelity LES using the high-order code, Nek5000. The distribution of energy across the different length scales was evaluated at different piston positions. They found energy content across all length scales increased with increasing compression. The team also used simulations to evaluate the correlations between turbulence distribution at the top dead center of compression with swirl and tumble ratios during the intake stroke.

IMPACT The team's datasets will be made publicly available so researchers across the world can access the data to pursue new insights. Ultimately, the knowledge gained from this project will improve our understanding of internal combustion engines and help in developing more efficient and cleaner engines.

PUBLICATIONS

Wu, S., M. M. Ameen, and S. S. Patel. "Investigating the Origins of Cyclic Variability in Internal Combustion Engines Using Wall-Resolved Large Eddy Simulation," *ASME Internal Combustion Engine Fall Conference* (October 2021), ASME.

Wu, S., M. M. Ameen, and S. S. Patel. "Investigation of Cycle-to-Cycle Variations in Internal Combustion Engine Using Proper Orthogonal Decomposition," *Flow, Turbulence and Combustion* (submitted), Springer Nature.

Ameen, M. M., S. S. Patel, J. D. Colmenares, and S. Wu. "Direct Numerical Simulation (DNS) and High-Fidelity Large-Eddy Simulation (LES) for Improved Prediction of In-Cylinder Flow and Combustion Processes," *2021 DOE VTO Annual Merit Review* (June 2021), U.S. Department of Energy.

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

PI Johan Larsson, University of Maryland
AWARD INCITE
SYSTEM Theta

Different aspects of shock-turbulence interactions. Background: shock-induced mixing shadowed by the velocity marking the shock. Foreground: oblique shock impinging on a turbulent boundary layer over a flexible panel. *Image: Ivan Bermejo-Moreno, University of Southern California*

High-speed air flows, such as those experienced by supersonic aircraft, create a thin boundary layer along a solid surface. If the boundary layer is hit by a shock wave, as forms when a plane approaches supersonic speeds and pushes on the surrounding air, the flow may separate from the surface. If the flow separates from the surface, the engine's performance degrades significantly. A team led by researchers from the University of Maryland used ALCF supercomputing resources to study supersonic turbulent boundary layers and their interaction with shock waves.

CHALLENGE With the aim of elucidating how shock/boundary-layer interactions (SBLI) are affected by crossflow (that is, flow along the shock and in a direction parallel to the wing span), researchers performed direct numerical simulations of skewed SBLI with crossflow so as to approximate the off-design operating conditions that can ultimately cause engine performance degradation and flameout.

APPROACH Leveraging the ALCF's Theta supercomputer, the team carried out SBLI simulations with the code Hybrid. Hybrid is a solver for the compressible Navier-Stokes equations that govern much of fluid dynamics.

RESULTS The team published their results in *Theoretical and Computational Fluid Dynamics and the AIAA Journal*. Their papers demonstrate that the size of a turbulent flow's separation bubble grows almost 50 percent larger in the presence of crossflow than in cases without crossflow.

IMPACT The ultimate goal of the team's research is to develop a predictive theory for SBLI physics and apply it so as to devise improved modeling techniques for surfaces in large-eddy simulations (LES). LES, as mathematical models for turbulence, are used in a broad range of applications throughout different fields of engineering.

The team expects their findings to help construct improved models that better predict the size of separation bubbles and account for the effect of crossflow for a range of engineering applications. This work is also significant in how it differs from other research in the fields. Historically, most studies of turbulent boundary layers focus on a specific canonical case: the case of incompressible flow, with constant density and viscosity, passing over rigid walls. This research, in contrast, examines the case of compressible flows non-uniform density and viscosity across their boundary layers.

PUBLICATIONS

Larsson, J., V. Kumar, N. Oberoi, M. Di Renzo, and S. Pirozzoli. "Large-Eddy Simulations of Idealized Shock/Boundary-Layer Interactions with Crossflow," *AIAA Journal* (2022), AIAA.

Di Renzo, M., N. Oberoi, J. Larsson, and S. Pirozzoli. "Crossflow Effects on Shock Wave/Turbulent Boundary Layer Interactions," *Theoretical and Computational Fluid Dynamics* (2022), Springer Nature.

Engineering | Simulation

Multiscale Bubble Breakup and Gas Transfer in Turbulent Oceanic Environments

PI Parviz Moin, Stanford University
AWARD ALCC
SYSTEM Theta

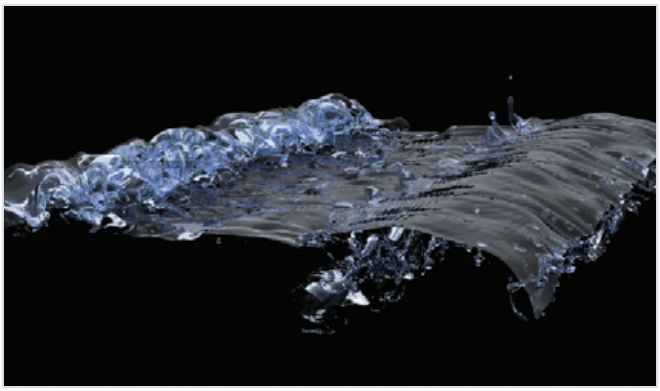

Wave surface generated in charLES run for verification of wave statistics. *Image: WH Ronald Chan, University of Colorado Boulder; Suhas S. Jain, Ali Mani, Shahab Mirjalili, Parviz Moin, and Javier Urzay, Stanford University*

This project, led by Stanford University researchers, uses computational fluid dynamics to focus on high-fidelity simulations of bubble breakup and gas dissolution in oceanic breaking waves that address energy and environmental challenges through the impact of oceans on weather/climate predictions and offshore wind technologies. The wave-breaking process gives rise to turbulent fluctuations that break up entrained air cavities in quick succession. A wide range of bubble sizes extending down to the Hinze scale—the critical length scale below which turbulent fragmentation ceases—is generated through a breakup cascade.

CHALLENGE The first objective of this work is to extend fundamental understanding of the bubble fragmentation process to sizes smaller than the Hinze scale, wherein bubble-formation mechanisms remain a subject of active research. The second objective is to quantify the bubble-induced gas dissolution that is enhanced by the breakup cascade. A major challenge for accomplishing these goals has been the inherent multiscale nature of the problem, which dramatically increases the computational cost and the complexity of the associated physical phenomena.

APPROACH The primary code used for this project is CTR-DIs3D, which is a structured, Cartesian-grid-based in-house solver developed at CTR that employs non-dissipative schemes for the simulation of incompressible and compressible two-phase flows. This solver implements the conservative diffuse-interface method to resolve the interface between two fluids and is designed to perform high-fidelity numerical simulations for the breakup and coalescence of bubbles in turbulent breaking waves, bubble acoustics, and transport of scalars in turbulent two-phase flows.

RESULTS In a paper published in *Journal of Computational Physics*—adopting a recently developed conservative diffuse-interface method for the simulation of compressible two-phase flows that conserves phase, momentum, and total energy—the researchers presented a set of numerical fluxes that results in an exact conservation of kinetic energy and approximate conservation of entropy in the absence of pressure work, viscosity, thermal diffusion effects, and time-discretization errors.

IMPACT Understanding sub-Hinze-scale bubble formation mechanisms and quantifying gas dissolution have significant impact on problems such as CO_2 transport in the carbon cycle. This project will yield deeper insights into oceanic bubble breakup and gas dissolution for the development of future predictive technologies for more energetic waves, as well as for associated phenomena such as hydroacoustics, radiation transfer, and scalar transport.

PUBLICATIONS

Jain, S. S., and P. Moin. "A Kinetic Energy- and Entropy-Preserving Scheme for Compressible Two-Phase Flows," *Journal of Computational Physics* (May 2022), Elsevier.

Engineering | Simulation, Learning

Next-Generation Nonwoven Materials Manufacturing

PI Ian Foster, Argonne National Laboratory
AWARD Director's Discretionary
SYSTEM Theta

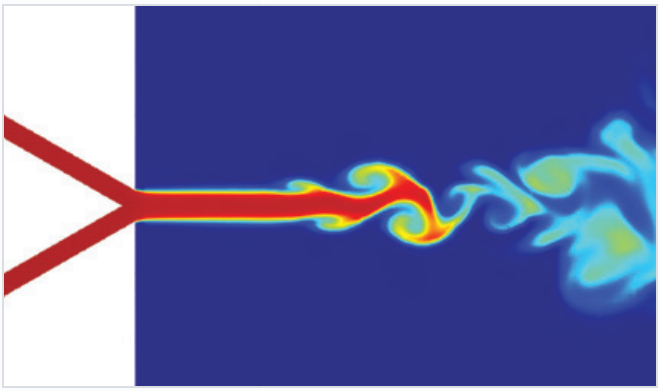

Computational fluid dynamics simulation of the melt-blowing process as modeled at Argonne. *Image: Argonne National Laboratory*

Melt-blown nonwoven materials form the basis of a range of filters, fabrics and insulation that includes critical items, such as N95 masks. Worldwide, manufacturers produce more than 300,000 tons of melt-blown nonwoven materials every year. To minimize the energy footprint of the process, researchers from 3M and Argonne National Laboratory are leveraging ALCF resources to model and optimize the production of melt-blown nonwoven materials.

CHALLENGE Combining simulations, data analysis, and machine learning (ML) to reduce the energy consumed in the melt-blowing production process—without compromising material quality—the researchers aim to reduce energy expenditures by 20 percent.

APPROACH The initial model for the melt-blowing process was developed through a series of simulation runs performed on the ALCF's Theta supercomputer with the computational fluid dynamics (CFD) software OpenFOAM. The researchers then scaled up the code using the CONVERGE CFD application. The optimization approach combines ML for simulation parameter selection with high-fidelity fluid dynamics simulations for modeling the process and nozzle geometry. For the models to closely match the company's melt-blowing process, 3M first performs experiments to establish a process map for a particular production setup. The process map must identify which combinations of processing choices yield acceptable non-woven materials. This stage is crucial for generalizing the process maps to other production setups and source polymers.

Once the simulations reproduce a 3M experimental process map with sufficient accuracy, they can be used to optimize processing choices for a given polymer. The researchers will then employ ML techniques to determine preferred processing parameters by simulating different configurations of settings, like air pressure and temperature, and various geometric parameters, like the angle of the nozzle or the diameter of the nozzle opening. Based on the simulations already generated, the ML algorithm suggests successive simulations with parameter settings iteratively refined to minimize energy consumption.

RESULTS While ML methods have thus far been used to refine and select successive experiments, the researchers intend to investigate the use of recurrent and physics-informed neural networks so as to predict fluid flow patterns, from which key metrics can be determined. The researchers are still working to identify and validate optimal process conditions in 3M facilities, but have successfully established a quantitative metric for improved parameter optimization. After simulation and validation are complete, the team will move on to industrial implementation.

IMPACT Reduction of energy expenditures by 20 percent in a market space as large as that of melt-blown nonwoven materials stands to carry global environmental and economic effects, including significantly smaller production costs and externalities.

Materials Science | Simulation, Learning

Autonomous Molecular Design for Redox Flow Batteries

PI Logan Ward, Argonne National Laboratory
AWARD ALCF Data Science Program
SYSTEM Theta, ThetaGPU

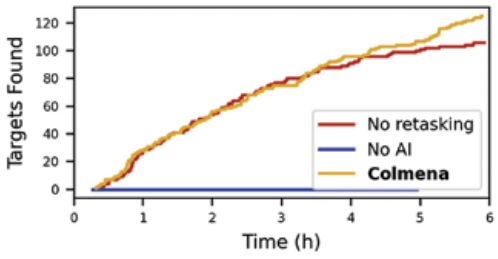

Number of high-performing redoxmer molecules found during a 256-node run on Theta. Colmena retasks HPC nodes between AI and simulation workloads to maximize the scientific output for a given run, finding over 120 target molecules in six hours. *Image: Logan Ward, Argonne National Laboratory*

Redox flow batteries are a promising energy storage technology for electrical grid storage, offering potential advantages in terms of cost, performance, and safety. To advance their development, researchers from Argonne National Laboratory are working to build an autonomous AI application for supercomputers that can select and perform the simulation and machine learning tasks needed to identify better-performing battery materials.

CHALLENGE Redox flow batteries work by storing energy in large tanks of electrolyte solutions, but they are currently limited by the performance of available electrolyte materials. With billions of potential candidate molecules, scientists need an improved method to speed the discovery of optimal materials for redox flow batteries. With this ADSP project, researchers are using a combination of quantum chemistry and machine learning methods to identify and optimize new electrolyte materials from within this intractable design space.

APPROACH The team's overarching goal is to develop an autonomous AI system that makes optimal use of supercomputing resources to design new molecules. To identify promising new battery materials, the team is employing active learning, which uses machine learning models trained on previous simulations to identify which new simulations are most crucial to finding an optimized molecule. As part of this effort, the researchers are also developing novel tools and frameworks that can help advance autonomous molecular design.

RESULTS Using active learning methods to couple AI and simulation tasks across Theta and ThetaGPU, the team has found hundreds of candidate molecules with promising properties for next-generation batteries. They are now working to fabricate the predicted molecules in a laboratory to facilitate experimental testing.

In addition, the team has made progress in developing new tools to enable their work. In a paper presented at the Workshop on Machine Learning in High Performance Computing Environments (MLHPC) at SC21, the team detailed the development of Colmena, an open-source Python framework that allows users to steer campaigns by providing just the implementations of individual tasks plus the logic used to choose which tasks to execute when. The team has also created a Function-as-a-Service (FaaS) system for executing complex, distributed computational campaigns that achieves performance parity with conventional workflow systems without any need for network configuration.

IMPACT The team's work to develop autonomous molecular design tools aims to accelerate the discovery of optimal energy storage materials for redox flow batteries. Ultimately, their work seeks to unlock new, more economical opportunities for grid storage and pave the way towards a greener energy future.

PUBLICATIONS

Ward, L., G. Sivaraman, J. G. Pauloski, Y. Babuji, R. Chard, N. Dandu, P. C. Redfern, R. S. Assary, K. Chard, L. A. Curtiss, R. Thakur, and I. Foster. "Colmena: Scalable Machine-Learning-Based Steering of Ensemble Simulations for High Performance Computing," *2021 IEEE/ACM Workshop on Machine Learning in High Performance Computing Environments (MLHPC)*, (November 2021), IEEE.

BraggNN: Fast X-ray Bragg Peak Analysis Using Deep Learning

PI Zhengchun Liu, Argonne National Laboratory
AWARD Director's Discretionary
SYSTEM ThetaGPU, Cerebras (ALCF AI Testbed)

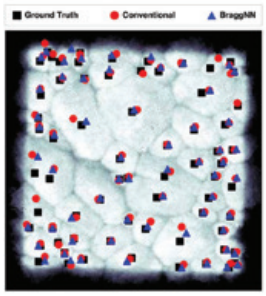

Comparison of *BraggNN*, pseudo-Voigt FF-HEDM, and NF-HEDM. Grain positions from NF-HEDM (black squares), pseudo-Voigt FF-HEDM (red circles) and *BraggNN* FF-HEDM (blue triangles) overlaid on the NF-HEDM confidence map. *Image: Antonino Miceli, Argonne National Laboratory*

A group of researchers from Argonne National Laboratory is harnessing AI to perform the challenging task of analyzing data from high-energy X-ray diffraction experiments. With a new neural network-based method called BraggNN, the Argonne team can more rapidly and precisely identify Bragg peaks—data points that indicate positions and orientations of tiny individual crystals—in a fraction of the time they used to.

CHALLENGE In recent years, a technique called high-energy diffraction microscopy (HEDM) has become an important tool for accurately characterizing complex materials with high resolution. Although HEDM has proven to be a great improvement over conventional techniques, it can also be expensive and time-consuming. It involves the collection of enormous datasets, analysis of millions of Bragg diffraction peaks and reconstruction of the sample using those peaks. The impending upgrade of Argonne's Advanced Photon Source (APS) is expected to vastly improve the speed of HEDM data acquisition to a minute or less. But the computing time to complete analysis Bragg peaks can extend to hours or weeks, even with the largest supercomputers. Not only do such delays slow research, but they also prevent the use of HEDM information to steer experiments.

APPROACH To address these challenges, the Argonne team developed BraggNN, a novel deep learning-based method that can accurately characterize Bragg diffraction peaks in HEDM images much more rapidly than conventional methods. The BraggNN network architecture comprises a series of convolutional neural network layers acting as feature extractors, followed by a series of fully connected layers that generate a regression prediction. The team used the ALCF's ThetaGPU system and Cerebras AI accelerator to train the BraggNN model with a collection of data containing diffraction peaks.

RESULTS To evaluate the performance of BraggNN, the researchers compared the results with the reconstructed position of grains in addition to diffraction peak positions. When applied to a real experimental dataset, they found the grain positions using BraggNN resulted in 15% smaller errors compared with those calculated using convention pseudo-Voigt fitting. Recent advances in deep learning method implementations and special-purpose model inference accelerators allow BraggNN to deliver significant performance improvements relative to the traditional method, running, for example, more than 200 times faster on a consumer-class GPU card with out-of-the-box software.

IMPACT The team has developed the first machine learned-based method for precisely characterizing Bragg diffraction peaks in HEDM images. Capable of precisely identifying Bragg peaks more quickly than conventional methods, the BraggNN method also stands to benefit high-resolution, high-throughput, and latency-sensitive applications, including real-time analysis and experiment steering.

PUBLICATIONS

Liu, Z., H. Sharma, J.-S. Park, P. Kenesei, A. Miceli, J. Almer, R. Kettimuthu, and I. Foster. "BraggNN: Fast X-ray Bragg Peak Analysis Using Deep Learning," *IUCrJ* (January 2022), International Union of Crystallography.

Materials Science | Simulation, Data

Dynamic Compressed Sensing for Real-Time Tomographic Reconstruction

PI Robert Hovden, University of Michigan
AWARD ALCF Data Science Program
SYSTEM Theta

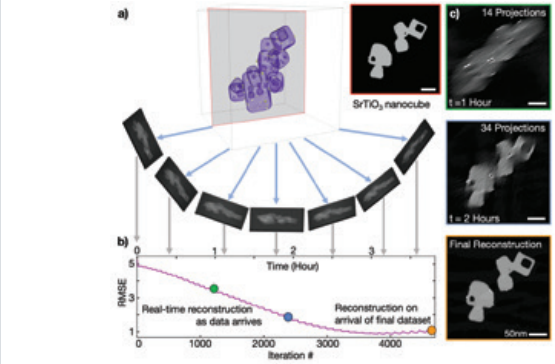

a) As the tomographic experiment progresses, projections are collected across an angular range. Measured projections are fed into the dynamic CS algorithm for 3D reconstruction. b) As the amount of data increases, the root mean square error (RMSE) decreases. c) 2D slices of the 3D reconstruction at various time stamps. *Image: Jonathan Schwartz, University of Michigan*

Electron and x-ray tomography allow researchers to perform 3D characterization of materials at the nano- and mesoscale, generating data that is critical to the development of nanomaterials for a wide range of applications, including solar cells and semiconductor devices. With an ALCF Data Science Program (ADSP) award, a University of Michigan-led research team is leveraging recent advancements in tomographic reconstruction algorithms, such as compressed sensing methods, to enhance and accelerate materials characterization research.

CHALLENGE Compressed sensing algorithms provide higher quality reconstructions, but they require substantially more computation time to complete, causing the rapidly expanding field of tomography to become critically bottlenecked by low throughput. To address these challenges and achieve real-time tomographic reconstruction using compressed sensing algorithms, the ADSP team has developed a dynamic framework that performs in-situ reconstruction simultaneously with data collection.

APPROACH With access to DOE supercomputing resources, the researchers are conducting comprehensive simulations for real-time electron tomography and developing reconstruction methods for through-focal tomography. The team is experimentally demonstrating their reconstruction workflow and methods on commercial scanning transmission electron microscopes and the ptychographic tomography instruments at Argonne's Advanced Photon Source.

RESULTS A paper published in *npj Computational Materials* introduced fused multi-modal spectroscopy, a technique offering high signal-to-noise ratio (SNR) recovery of nanomaterial chemistry by linking correlated information encoded within detector signals. The team's approach substantially improves SNRs for chemical maps by as much as 500 percent, reduces electron doses by an order of magnitude, and enables accurate measurement of local stoichiometry.

A paper accepted for published in *Nature Communications* demonstrated real-time tomography with dynamic 3D tomographic visualization to enable rapid interpretation of specimen structure immediately as data are collected with an electron microscope. The authors show volumetric interpretation can begin in less than 10 minutes and that a high-quality tomogram is available within 30 minutes. Real-time tomography is integrated into tomviz, an open-source and cross-platform 3D data analysis tool that contains intuitive graphical user interfaces (GUI), to enable any scientist to characterize biological and material structure in 3D.

IMPACT The team's approach will help advance materials characterization research by enabling real-time analysis of 3D specimens while an experiment progresses. By integrating their framework with an open-source 3D visualization and tomography software package, the team's techniques will be accessible to a wide range of researchers and enable new material characterizations across academia and industry.

PUBLICATIONS

Hovden, R., J. Schwartz, C. Harris, J. Pietryga, H. Zheng, P. Kumar, A. Visheratina, N. Kotov, B. Major, P. Avery, P. Erius, U. Ayachit, B. Geveci, D. Muller, A. Genova, Y. Jiang, and M. Hanwell. "Real-Time 3D Analysis During Tomographic Experiments using tomviz," *Nature Communications* (accepted), Springer Nature.

Large-Scale Simulations of Light-Activated Matter

PI Guilia Galli, University of Chicago and Argonne National Laboratory

AWARD INCITE

SYSTEM Theta

Machine learning can circumvent explicit calculation of certain material behavior to accelerate simulations of optical properties of complex materials at finite temperature. *Image: Argonne National Laboratory*

This INCITE project explores light-activated processes in materials and molecules using first principles, quantum mechanical calculations, and addresses two outstanding challenges: designing materials, namely solids and molecules, that are useful to develop sustainable energy sources, and promising systems for quantum technologies. To investigate these problems, the team is coupling first-principles atomistic simulations with computational spectroscopic techniques. Their studies are focused on oxides for photoelectrodes and low-power electronics, and on defective semiconductors for quantum sensing applications.

CHALLENGE One of the main challenges related to this project is that the materials of interest are hybrid, heterogeneous materials, and the predictions and design of their emergent behaviors require the ability to compute multiple properties of systems with hundreds or thousands of atoms. These properties encompass structural properties of interfaces and defective solids and electronic and vibrational spectra.

APPROACH The team is coupling different quantum mechanical methods, including first principles molecular dynamics to study structural properties at finite temperature, and the calculation of excited electronic states to investigate spectroscopic properties and response functions. Dynamical simulations are based on density functional theory (DFT) and carried out with the Qbox code (http://qboxcode.org/) and excited state calculations are based on post-DFT methods, namely many-body perturbation theory, and carried out with the WEST code (http://www.west-code.org/). Both open-source codes are developed within the DOE-funded Midwest Integrated Center for Computational Materials (MICCOM, http://miccom-center.org/). The team used ALCF's Theta supercomputer to calculate the properties of aqueous interfaces and defective solids.

RESULTS Using a data-driven approach based on machine learning, the team was able to simplify the solution of the quantum mechanical equations that describe how light is absorbed by a solid, liquid or molecule. The simplification led in turn to substantial computational savings and was applied to solid/liquid interfaces such as those found between water and a photoelectrode (a material that can help turn sunlight into electricity). The team also developed codes to accurately simulate how light is emitted by a spin-defect, and computed photolumiscence spectra that were validated within a collaboration with experimentalists at Argonne. Different levels of theory were tested and compared as well, leading to a general protocol to study light-activated processes in defects in semiconductors.

IMPACT The impact of the project is two-fold: the team developed and applied general first principles methods to understand how light interacts with heterogeneous solids; these methods can now be extended to study broader classes of problems. In addition, the team developed a robust protocol to compare theory and experiments that can be applied to diverse spin-defects in semiconductors and accelerate the prediction and engineering of their properties.

PUBLICATIONS

Yang, H., M. Govoni, A. Kundu, and G. Galli. "Combined First-Principles Calculations of Electron-Electron and Electron-Phonon Self-Energies in Condensed Systems," *Journal of Chemical Theory and Computation* (December 2021), ACS Publications,

Jin, Y., M. Govoni, G. Wolfowicz, S. E. Sullivan, F. J. Heremans, D. D. Awschalom, and G. Galli. "Photoluminescence Spectra of Point Defects in Semiconductors: Validation of First Principles Calculations," *Physical Review Materials* (August 2021), APS.

Dong, S., M. Govoni, and G. Galli. "Machine Learning Dielectric Screening for the Simulation of Excited State Properties of Molecules and Materials," *Chemical Science* (March 2021), Royal Society of Chemistry.

Materials Science | Simulation, Learning

Metastable Phase Diagrams for Materials

PI Subramanian Sankaranarayanan
Argonne National Laboratory
AWARD Director's Discretionary
SYSTEM Theta

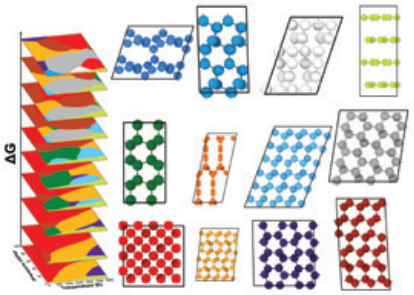

The final product of the machine learning algorithm: metastable phase diagrams for carbon. The colored regions indicate conditions at which carbon exists in certain metastable states (with similarly colored structures) that may yield useful material properties. *Image: Argonne National Laboratory*

Phase diagrams are an invaluable tool for materials synthesis, providing researchers with information on the phases of a material at any given thermodynamic condition (e.g., pressure, temperature, chemical composition). To extend their utility to a promising but mysterious class of materials, researchers from Argonne National Laboratory are using ALCF computing resources and machine learning to develop an automated workflow to construct phase diagrams for metastable materials.

CHALLENGE Conventional phase diagrams represent a reduced set of phases observed at distinct thermodynamic equilibria. In contrast, materials during their synthesis, operation, or processing, may not reach their thermodynamic equilibrium state, but instead remain trapped in a local (metastable) free energy minimum, that may exhibit desirable properties for various applications. Mapping these metastable phases and their thermodynamic behavior is therefore highly desirable, but it is a non-trivial and data-intensive task that is currently lacking due to the vast configurational landscape.

APPROACH Leveraging the ALCF's Theta supercomputer, Argonne researchers are creating an automated workflow that integrates first-principles physics and atomistic simulations with machine learning and high-performance computing to allow rapid exploration of the metastable phases of any given elemental composition. Their framework allows for the curation of metastable structures from published literature/databases and concurrently enables automatic discovery, identification, and exploration of the metastable phases of a material, and learns their equations of state through a deep neural network.

RESULTS Using carbon as a prototypical system, the team demonstrated automated metastable phase diagram construction to map hundreds of metastable states ranging from near equilibrium to far-from-equilibrium (400 meV/atom). The researchers incorporated the free energy calculations into a neural-network-based learning of the equations of state that allows for efficient construction of metastable phase diagrams. They used the phase diagrams to identify domains of relative stability and synthesizability of metastable materials. The team validated their metastable phase predictions with data from high-temperature, high-pressure experiments using a diamond anvil cell on graphite sample coupled with high-resolution transmission electron microscopy. Their approach is quite general, making it possible to extend the phase diagram construction to multi-component alloy systems. Their findings were detailed in a paper published in *Nature Communications*.

IMPACT The team's automated framework for constructing metastable phase diagrams lays the groundwork for computer-aided discovery and design of synthesizable metastable materials, which could help advance a range of applications including semiconductors, catalysts, and solar cells.

PUBLICATIONS

Srinivasan, S., R. Batra, D. Luo, T. Loeffler, S. Manna, H. Chan, L. Yang, W. Yang, J. Wen, P. Darancet, and S. K.R.S. Sankaranarayanan. (2022). "Machine learning the metastable phase diagram of covalently bonded carbon," *Nature Communications* (June 2022), Springer Nature.

Predicting Ion Transport Kinetics at Complex Interfaces for Energy Storage

PI Brandon Wood
Lawrence Livermore National Laboratory
AWARD INCITE
SYSTEM Theta

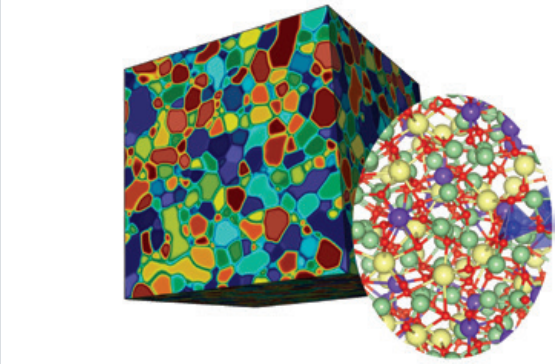

Predicted microstructure of a $Li_7La_3Zr_2O_{12}$ solid-state battery electrolyte showing diverse orientations of individually colored grains. The foreground image shows an atomic-scale representation of a disordered boundary region where lithium ions tend to congregate and limit battery performance. *Image: Brandon Wood, Lawrence Livermore National Laboratory*

Many of the most promising high-capacity, solid-state energy storage systems—including next-generation batteries and hydrogen storage materials—rely on the rapid transport of ions through an electrolyte. This process, however, can be severely inhibited by complex and disordered interfaces in various devices. To shed light on this phenomenon, researchers from Lawrence Livermore National Laboratory are employing a multiscale modeling framework to predict ion transport kinetics at solid-solid interfaces.

CHALLENGE Because buried solid-state interfaces are notoriously difficult to probe experimentally—particularly during device operation—computational simulations play a unique and increasingly visible role. However, the idealized models often used for simulation studies neglect the chemical and structural disorder that real interfaces exhibit, in large part due to computational cost and complexity. As a result, interfacial ion transport is an exceptionally challenging problem to tackle using conventional approaches, instead requiring extensive multiscale simulations capable of spanning broad ranges of time and length scales.

APPROACH Focusing on ceramic solid-state battery electrolytes and metal hydride hydrogen storage materials, this project integrates three sets of simulation capabilities to predict ion transport kinetics at interfaces. First, representative interface and disordered models are generated using large-scale ab initio molecular dynamics (AIMD) simulations, validated using computational spectroscopy. Second, these models are used as training sets for machine learning force fields able to span a much larger range of configuration space. Third, atomic-scale inputs are handed to a phase-field continuum model capable of simulating complex microstructures.

RESULTS In a study published in *The Journal of Chemical Physics*, the team employed machine learning to investigate $Li_7La_3Zr_2O_{12}$ (LLZO), a promising solid-state electrolyte material. The researchers developed and validated a machine learning potential for simulating crystalline, disordered, and amorphous LLZO systems across a wide range of conditions. Based on a neural network algorithm and trained using ab initio data, the machine learning potential was able to predict accurate structural and vibrational characteristics, elastic properties, and Li diffusivity of LLZO comparable to AIMD simulations. The team's study demonstrated how this approach can enable simulations of transitions between well-defined and disordered structures with quantum-level accuracy at speeds thousands of times faster than traditional *ab initio* methods.

IMPACT This project is harnessing supercomputing in combination with advanced machine learning techniques to simulate ion transport kinetics in novel energy storage systems with unprecedented accuracy. With a better understanding of how interfaces impact ion transport, researchers can develop rational engineering strategies for improving performance of advanced battery materials for grid and vehicular energy storage applications.

PUBLICATIONS

Kim, K., A. Dive, A. Grieder, N. Adelstein, S. Kang, L.F. Wan, and B.C. Wood. "Flexible Machine-Learning Interatomic Potential for Simulating Structural Disordering Behavior of $Li_7La_3Zr_2O_{12}$ Solid Electrolytes," *The Journal of Chemical Physics* (May 2022), AIP Publishing

Materials Science | Simulation

Towards Predictive Calculations of Functional and Quantum Materials

PI Paul Kent, Oak Ridge National Laboratory
AWARD INCITE
SYSTEM Theta

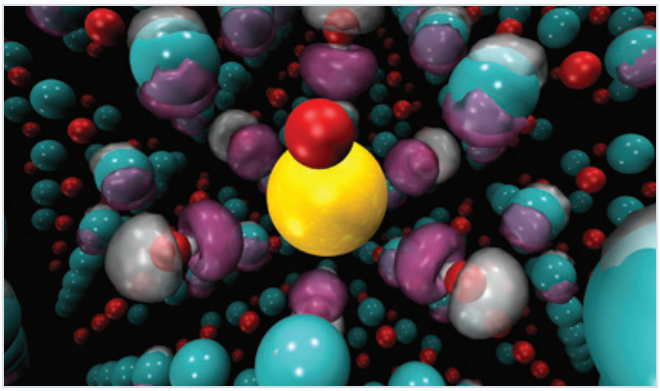

Diffusion Monte Carlo spin density difference between bulks of potassium-doped nickel oxideand pure nickel oxide, showing the effects of substituting a potassium atom (center atom) for a nickelatom on the spin density of the bulk. *Image: A. Benali, O. Heinonen, J. A. Insley, and H. Shin, Argonne National Laboratory*

Our understanding of and ability to predict and design functional and quantum materials is severely limited by the challenges in simulating properties that vary greatly with small changes in composition, pressure, strain, and doping. This project, led by Oak Ridge National Laboratory researchers, aims to predict and better understand the quantum-mechanical properties of novel functional materials.

CHALLENGE The couplings between electron charge, spin, and atomic positions are what yield functionality in functional materials, but their subtlety makes predictive simulations a grand challenge. Accurately predicting the properties and behavior of quantum materials—solids that can host novel quantum states and physical properties arising from uncommon quantum-mechanical effects—poses even greater difficulties, as the states of interest so strongly differ from existing materials theories. Understanding functional materials requires methods that permit systematic improvement, with known and potentially quantifiable approximations.

APPROACH Necessarily leveraging DOE supercomputing resources, the researchers utilize quantum Monte Carlo (QMC), as implemented in the open-source QMCPACK code. By directly solving the Schrödinger equation and by treating a system's electrons at a consistent and highly-accurate many-body level, QMC can be applied to general elements and materials, while employing very few approximations.

RESULTS As presented in a paper published in *The Journal of Chemical Physics*, the researchers developed an efficient energy-based method for structural optimization with stochastic electronic structure theories. Based on robust line-search energy minimization in reduced parameter space, the method exploits approximate Hessian information from a surrogate theory, such as density functional theory, also used to characterize potential energy surface, enabling simple but reliable maximization of statistical efficiency without sacrificing accuracy. Furthermore, this method is suitable for a broad range of applications and is easily generalized to any electronic structure method in which forces and stresses are under active development and implementation, as well as to deterministic approaches.

IMPACT Accurate and efficient means of optimizing optimization could shed light on a broad class of materials and molecules, showing high sensitivity of induced properties to structural variables. The materials studied are of outstanding fundamental scientific interest and present the potential for the development of new sensors and devices.

PUBLICATIONS

Tiihonen, J., P. R. C. Kent, and J. T. Krogel. "Surrogate Hessian Accelerated Structural Optimation for Stochastic Electronic Structure Theories," *The Journal of Chemical Physics* (February 2022), AIP Publishing.

Dumi, A., S. Upadhyay, L. Bernasconi, H. Shin, A. Benali, and K. D. Jordan. "The Binding of Atomic Hydrogen on Graphene from Density Functional Theory and Diffusion Monte Carlo Calculations," *The Journal of Chemical Physics* (April 2022), AIP Publishing.

Isaacs, E. B., H. Shin, A. Annaberdiyev, C. Wolverton, L. Mitas, A. Benali, and O. Heinonen. "Assessing the Accuracy of Compound Formation Energies with Quantum Monte Carlo," *Physical Review B* (June 2022), APS.

Materials Science | Simulation, Learning

Ultrafast Control of Functional Materials

PI Priya Vashishta, University of Southern California
AWARD INCITE
SYSTEM Theta

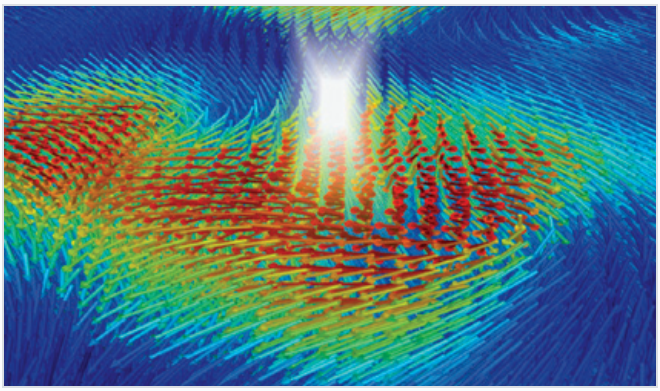

Photoinduced topological switching in $PbTiO_3$. *Image: Priya Vashishta, University of Southern California*

Artificial intelligence (AI) methods, such as machine learning and neural networks, have shown great potential in accelerating the discovery of new functional materials. With this INCITE project, researchers from the University of Southern California are combining advanced AI techniques with leadership-scale quantum dynamics simulations to extend the frontiers of computational materials science.

CHALLENGE Ferroelectric materials exhibit a rich range of complex polar topologies (i.e., topologies formed by spatial distribution of electric dipoles), but their study under far-from-equilibrium optical excitation has been largely unexplored because of the difficulty in quantum-mechanically modeling the large spatiotemporal scales involved. To study optical excitation at the scales where these topologies emerge, the researchers developed and performed multiscale excited-state neural network quantum molecular dynamics simulations that integrate quantum-mechanical description of electronic excitation and billion-atom machine learning molecular dynamics to describe ultrafast polarization control in an archetypal ferroelectric oxide, $PbTiO_3$.

APPROACH The INCITE team continues to expand the physics and chemistry that can be described by their two primary codes: QXMD, a non-adiabatic quantum molecular dynamics simulation engine; and RXMD, a reactive molecular dynamics simulation engine. They are also leveraging novel methods such as deep neural networks and reinforcement learning to advance their research into a variety of materials. In addition, the team is benefitting from synergy between their INCITE project and a related Aurora Early Science Program project, which is helping to further improve the performance of their simulation engines on Theta.

RESULTS As described in a paper published in *Science Advances*, far-from-equilibrium quantum simulations within picoseconds revealed a significant photo-induced change in the electronic energy landscape in prototypical ferroelectric material $PbTiO_3$, resulting in a crossover from ferroelectric to oxygen octahedral-tilting topological dynamics. The coupling and frustration of these dynamics, in turn, create topological defects in the form of polar strings similar to the Kibble-Zurek mechanism in cosmology. The demonstrated nexus of multiscale quantum simulation and machine learning will boost not only the emerging field of ferroelectric "topotronics" but also broader optoelectronic applications.

IMPACT The team's use of emerging AI techniques is accelerating efforts to identify promising material compositions and topological phases. Ultimately, this research will help inform the discovery and synthesis of new materials engineered for targeted applications, such as sensors, computer memory, and batteries.

PUBLICATIONS

Linker, T., K. Nomura, A. Aditya, S. Fukushima, R. K. Kalia, A. Krishnamoorthy, A. Nakano, P. Rajak, K. Shimamura, F. Shimojo, and P. Vashishta. "Exploring Far-from-Equilibrium Ultrafast Polarization Control Ferroelectric Oxides with Excited-State Neural Network Quantum Molecular Dynamics," *Science Advances* (March 2022), AAAS.

Physics | Simulation

Ab-initio Nuclear Structure and Nuclear Reactions

PI Gaute Hagen, Oak Ridge National Laboratory
AWARD INCITE
SYSTEM Theta

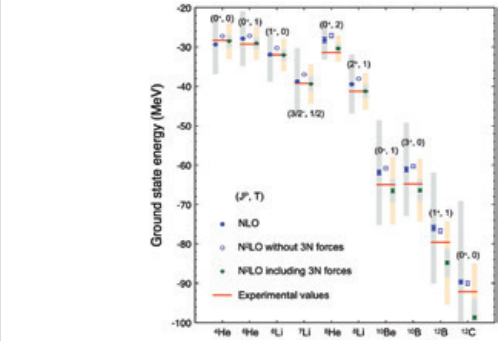

Calculated ground state energies in MeV using chiral NLO and N²LO interactions (blue and green symbols) in comparison with experimental values (red levels). *Image: Evgeny Epelbaum, Ruhr University Bochum*

Protons and atomic nuclei account for 99 percent of the visible mass in the universe, but how protons join together through astrophysical processes to form atomic nuclei is not well understood—nor are the resultant properties of those nuclei. This project, led by researchers from Oak Ridge National Laboratory, focuses on first-principles approaches to nuclear structure and reactions that apply interactions derived within effective theories of quantum chromodynamics (QCD). The calculations performed will make predictions for and guide new experiments at major DOE facilities, explain observed phenomena, and potentially propel the discovery of new laws of nature.

CHALLENGE Among the principal tasks of computational nuclear physics is the establishment of a reliable quantitative first-principles description of nuclear structure and reactions. It is presently believed that the most promising approach to achieve this goal combines chiral effective field theory (EFT) (for the description of nuclear reactions in concert with QCD symmetries and breaking patterns), with ab-initio few-body methods (for tackling the quantum-mechanical A-body problem). To address this challenge, the Low Energy Nuclear Physics International Collaboration (LENPIC) aims to develop a precise, accurate description of two- and three-nucleon interactions by pushing the EFT expansion to high-chiral orders and using these interactions to solve the structure and reactions of light nuclei.

APPROACH To perform calculations, the team used the no-core configuration interaction (NCCI) approach implemented in their Many Fermion Dynamics—nuclear (MFDn) code. For the largest runs they used almost the entire Theta machine in quad-flat mode, which gave the best performance and maximal memory, and allowed them to calculate properties of ground and excited states of light nuclei with robust theoretical error estimates.

RESULTS The team performed the first tests of novel chiral nucleon-nucleon potentials with consistent three-nucleon interactions. This demonstrates the importance of three-nucleon interactions and allows for a quantitative understanding of the theoretical uncertainties due to the chiral EFT expansion. The team also extended and tested a Bayesian statistical model that learns from the order-by-order EFT convergence pattern to account for correlated excitations. This enabled demonstrated agreement with experimental ground state energies as well as excitation energies. The results were compared to known experimental results to test consistent LENPIC chiral EFT interactions consisting of two- and three-nucleon interactions. The calculated results were consistent with the experimental results, confirming the validity of the approach.

IMPACT Establishing predictive power for nuclear physics carries ramifications for a wide range of fields, including astrophysics, medicine, nuclear power, and nuclear security.

PUBLICATIONS

Maris, P., E. Epelbaum, R. J. Furnstahl, J. Golak, K. Hebeler, T. Hüther, H. Kamada, H. Krebs, U.-G. Meißner, J. A. Melendez, A. Nogga, P. Reinert, R. Roth, R. Skibiński, V. Soloviov, K. Topolnicki, J. P. Vary, Y. Volkotrub, H. Witała, and T. Wolfgruber. "Light Nuclei with Semilocal Momentum-Space Regularized Chiral Interactions up to Third Order," *Physical Review C* (May 2021), APS.

Physics | Simulation, Learning

High-Fidelity Gyrokinetic Simulation of Tokamak and ITER Edge Physics

PI Choong-Seock Chang,
Princeton Plasma Physics Laboratory
AWARD INCITE, Aurora Early Science Program
SYSTEM Theta

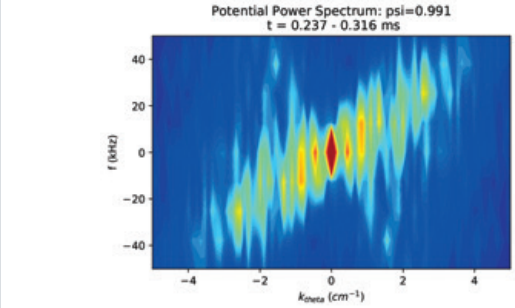

Frequency–wavenumber spectrum of nonlinear edge turbulence just before the L-H bifurcation. The pedestal profile is still evolving, and the turbulence has not reached its peak. *Image: C. S. Chang, Princeton Plasma Physics Laboratory*

This multi-year INCITE and ESP project seeks to advance our understanding of the edge plasma physics in fusion reactors, with a focus on ITER, and to assemble a burning plasma experiment that can demonstrate the scientific and technological feasibility of fusion.

CHALLENGE The INCITE team is performing studies on two high-priority challenges: (1) quantifying the narrowness of the heat-flux width on the ITER divertor material plates in the high-confinement mode (H-mode) operation during tenfold energy gain operation; and (2) understanding the basic physics behind the low-to-high mode L-H transition and pedestal formation at the edge, which is necessary to achieve a tenfold energy gain in ITER.

APPROACH The researchers are using the 5D gyrokinetic particle-in-cell code, XGC, on DOE leadership computing resources to address some of the most difficult plasma physics questions facing ITER. The team used this extreme-scale modeling code to solve kinetic equations for the tokamak edge by modeling plasma with a large number of particles. Predictions from XGC for the low-to-high mode transition revealed that a heavier-mass plasma can allow an easier transition to the high mode operation that is necessary for ITER to achieve its goal, in which a strongly sheared plasma flow suppresses edge turbulence. To be more specific, the L-H transition could occur more easily in a heavier plasma, due to a stronger mean plasma-flow effect, without the necessity for a significant turbulence energy transfer to the plasma-flow. A new electromagnetic solver algorithm has been developed for higher fidelity studies of the ITER edge plasma.

RESULTS The team's new electromagnetic solver results were described in a paper submitted to *Physics of Plasmas*. A simplified delta-f mixed-variable/pull-back electromagnetic simulation algorithm implemented in XGC for core plasma simulations was generalized to a total-f electromagnetic algorithm able to include—for the first time—the boundary plasma in diverted magnetic geometry with neutral particle recycling, turbulence, and neoclassical physics. It is confirmed that electromagnetic simulation is necessary for a higher-fidelity understanding of the tokamak edge physics in the presence of the divertor and magnetic separatrix.

IMPACT Establishing an accurate predictive formula for the exhaust heat-load width and low-to-high mode transition of future doughnut-shaped tokamak fusion reactors can help enable researchers to progress faster toward the goal of 0.5 GW of fusion power production from 50 MW of input power in ITER. More accurate understanding and prediction can also help inform more reliable designs for future fusion reactors, which currently suffer from the limitation imposed by exhaust heat-load width on the divertor plates and the required L-H transition power.

PUBLICATIONS

Hager, R., S.-H. Ku, A. Y. Sharma, C. S. Chang, and R. M. Churchill. "Electromagnetic Total-f Algorithm for Gyrokinetic Particle-in-Cell Simulations of Boundary Plasma in XGC," *Physics of Plasmas* (submitted), AIP Publishing.

Physics | Data, Learning

Model Quality Breakthrough for Cosmic Background Removal

PI Corey Adams, Argonne National Laboratory
AWARD Director's Discretionary
SYSTEM SambaNova (ALCF AI Testbed)

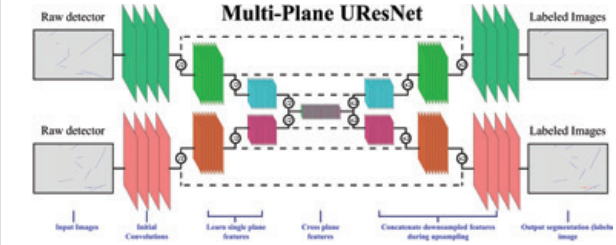

A representation of the multi-plane UResNet architecture. *Image: Argonne National Laboratory*

Billions of neutrinos pass through everything on Earth every second, but since they are weakly interacting particles, they are very difficult to detect. When looking at neutrino-triggered events on the Earth's surface, cosmic muons and other particles from the cosmic background predominate on the detectors, cause events, and make it difficult to record a single neutrino-induced event. Using the ALCF AI Testbed, researchers are advancing the use of deep neural networks to remove these background particles from detector images.

CHALLENGE Liquid argon time projection chambers (LArTPCs) for neutrino direction running near the Earth's surface experience backgrounds induced by cosmic interactions occurring at much higher rates than neutrino interactions. It is difficult to identify what pixels in images from LArTPCs are the result of cosmic particles, background noise, or neutrinos, so novel techniques using convolutional neural networks (CNNs) were devised in order to tag these pixels in the images. On GPU-based platforms, models had been restricted by the size of the images they were trained with, so the team turned to SambaNova's DataScale system to enable training on large, high-resolution images.

APPROACH The Argonne-led team used SambaNova's DataScale system to improve the accuracy of identifying pixels in images from LArTPCs. In an earlier effort, researchers developed a novel modification of UResNET architecture in order to distinguish cosmic pixels from neutrino pixels and segment events across all planes in images simulating data from the Short Baseline Near Detector neutrino detector. To avoid running out of memory, they had to downsample the images the neural network was trained on to 50 percent resolution, which resulted in lost information. The team overcame this problem using the SambaNova system's reconfigurable dataflow architecture to train CNNs on images beyond 50k x 50k.

RESULTS The team was able to ensure stable convergence and achieve better results compared to work carried out on GPU-based platforms. The results they obtained using larger images performed better by 6% mean intersection over union. The team also improved loss functions, including focal loss. In labeling the images, the researchers achieved a higher degree of accuracy in identifying pixels and demonstrated the ability to efficiently train models on gigapixel images.

IMPACT The team developed a novel AI method to remove backgrounds from LArTPC images for improved neutrino detection, providing an approach that can make it easier to study the elusive particles. Their technique is also applicable to other LArTPC detectors running at the surface including MicroBooNE, ICARUS, and ProtoDUNE. In addition, the team's research demonstrates capabilities that can be used to advance model quality for a variety of important and challenging image processing problems.

PUBLICATIONS

Acciarri, R., C. Adams, C. Andreopoulos, J. Asaadi, et al. "Cosmic Background Removal with Deep Neural Networks in SBND," *Frontiers in Artificial Intelligence* (August 2021), Frontiers Media.

Physics | Simulation, Learning

Neural Network Quantum States for Atomic Nuclei

PI Alessandro Lovato and Corey Adams, Argonne National Laboratory
AWARD Director's Discretionary
SYSTEM Theta, ThetaGPU

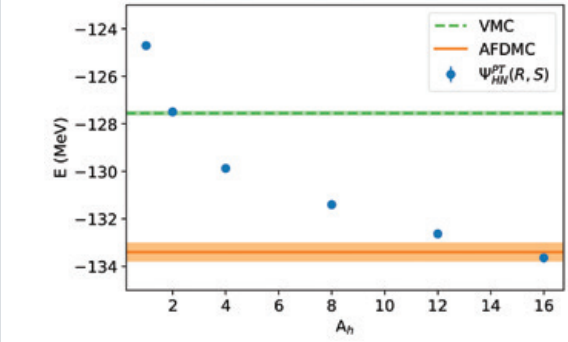

Ground-state energy of ^{16}O as a function of the number of hidden nucleons compared with conventional variational Monte Carlo and auxiliary-field diffusion Monte Carlo methods. *Image: Alessandro Lovato, Argonne National Laboratory*

The atomic nucleus is the small, dense region situated at the center of the atom, comprised of neutrons and protons bound together by the strongest known force. Understanding the structure of atomic nuclei starting from the interactions among their constituents is a highly non-trivial "many-body" problem. Its solution will help answering fundamental questions, ranging from the synthesis and stability of nuclei to the structure of compact astrophysical objects like neutron stars. In this project, researchers are developing novel solutions to the nuclear many-body problem using artificial neural networks.

CHALLENGE Solving atomic nuclei starting from the individual interactions between nucleons is extremely computationally complex. Using classical methods, these calculations require algorithms beyond the scale of current and future supercomputers. To solve this problem, the team is using neural network methods that compactly represent the quantum mechanical wave function of atomic nuclei, thus avoiding the problem of exponential scaling with the number of nucleons if these wave functions were computed classically.

APPROACH Using open-source TensorFlow and JAX frameworks, the team developed a novel artificial neural network representation of nuclear wave functions that by construction fulfills the symmetry of the problem, including the fermion one. In addition, dedicated optimization algorithms have been developed to train the network by efficiently minimizing the energy of the system. With nearly perfect weak scaling, the team's AI-based simulations can enable precise explorations into larger and more complex nuclei than currently possible, especially on larger supercomputers such as the ALCF's Polaris and Aurora systems.

RESULTS The team has published several studies on using neural networks to simulate atomic nuclei. In a paper published in *Physical Review Letters* the researchers demonstrated that an artificial neural network quantum state ansatz can model the ground-state wave function of light nuclei and approximately solve the nuclear many-body Schrodinger equation. In a *Few-Body Systems* journal paper they indicated how they have benchmarked the binding energies, point-nucleon densities and radii of helium and lithium nuclei with the highly-accurate hyperspherical harmonics method. In a paper submitted to *Physical Review Research* the team demonstrated generalizing the hidden-fermion family of neural network quantum states to include continuous and discrete degrees of freedom and solve the nuclear many-body Schrödinger equations in a systematically improvable fashion.

IMPACT The team is using neural networks to advance the understanding of atomic nuclei, allowing them to model and describe the short- and long-range dynamics of heavier nuclei than has previously been possible. The researchers have created a systematically improvable approach to solving the nuclear many-body problem, which will allow continuing progress. These methods are important in advancing condensed-matter physics, nuclear physics, and quantum chemistry.

PUBLICATIONS

Adams, C., G. Carleo, A. Lovato, and N. Rocco. "Variational Monte Carlo Calculations of A≤4 Nuclei with an Artificial Neural-Network Correlator Ansatz," *Physical Review Letters* (July 2021), APS.

Gnech, A., C. Adams, N. Brawand, G. Carleo, A. Lovato, and N. Rocco. "Nuclei with Up to *A*=6 Nucleons with Artificial Neural Network Wave Functions," *Few-Body Systems* (December 2021), Springer Nature.

Lovato, A., C. Adams, G. Carleo, and N. Rocco, "Hidden-Nucleons Neural-Network Quantum States For The Nuclear Many-Body Problem," *Physical Review Research* (accepted), APS.

Physics | Simulation

Next-Generation 3D Core Collapse Supernova Simulations

PI Adam Burrows, Princeton University
AWARD INCITE
SYSTEM Theta

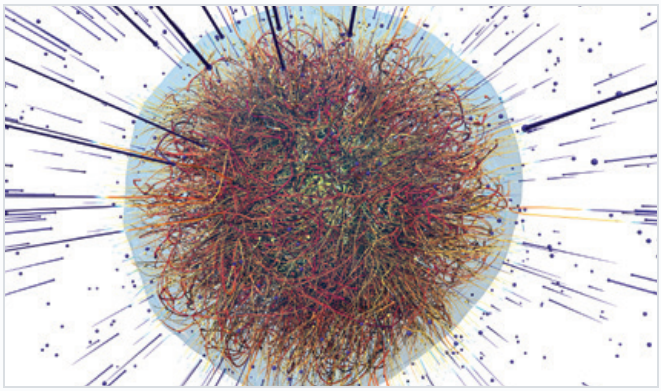

This image depicts the inner region where a proto-neutron star has been created due to the collapse of the core of a massive star. *Image: ALCF Visualization and Data Analysis Team; Adam Burrows, Princeton University*

Supernovae are the source of many chemical elements and have a profound effect on the composition of the universe. In order to understand how these events occur, researchers from Princeton University are using ALCF supercomputers to simulate the death of massive stars and birth of neutron stars and black holes in supernova explosions. The team is looking at turbulent 3D motions central to the explosion mechanisms and investigating which massive stars end their lives as supernovae, with what properties and why.

CHALLENGE The initial stages of supernovae explosions are mysterious due to the opacity of the cores of stars and direct observation is not possible. Conducting simulations in one direction—radially plus time—was challenging because the shockwaves from the exploding stars frequently stalled, moving outwards in mass but not in radius. The supercomputers at ALCF are capable of effectively computing the exponential complexity of simulations including highly complicated dynamics such as neutrino-matter interactions plus turbulence which allows modeling in 3D.

APPROACH To carry out their studies, the team uses the FORNAX code, which allows simultaneous simulations of hydrodynamics and radiative transfer. The researchers found that instead of solving radiative transfer for every grid point at each step, they could solve for radiative transport locally, which allowed simulations to go five time faster. 3D simulations with neutrinos and turbulence allow many more insights into star core collapse explosions. ALCF computing resources facilitate the possibility of deriving a comprehensive standard model of massive-star explosions to their final states.

RESULTS This project's 3D simulations have revealed many aspects of core collapse explosions and neutron star birth. In super-high energy and high density environments, many neutrinos are created, and these neutrinos drive a convective shell around the inner proto-neutron star just prior to explosion. This project has modeled the moments after the core bounce of the stalled shockwave and how the hydrodynamic stress and neutrino energy launch the explosion after a delay. The team also modeled the breaking of symmetry in the explosions, including element composition asymmetries and as well as modeling neutrino loss leading to neutron star formation and the possibility of induced pulsar spin from core collapse.

IMPACT With this project, researchers are seeking to create a standard model of star core collapse by investigating radiation, hydrodynamics, and neutrino-matter couplings. 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., and D. Vartanyan. "Core-Collapse Supernova Explosion Theory," *Nature* (January 2021), Springer Nature

Burrows, A., and M. Coleman, "The Character of Three-Dimensional Core-Collapse Simulation Results," *EPJ Web of Conferences* (February 2022), EDP Sciences.

Petascale Simulations of Kinetic Effects in IFE Plasmas

PI Frank Tsung, University of California, Los Angeles
AWARD INCITE
SYSTEM Theta

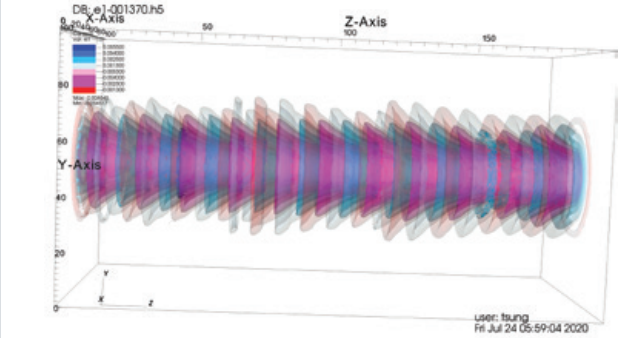

Isosurface of nonlinear plasma waves generated by laser plasma instabilities. The wavefronts bend due to a nonlinear frequency shift that causes the center waves to propagate slower than those at the edge, leading to a phenomenon called wave bowing. *Image: Frank Tsung, University of California, Los Angeles*

Inertial fusion energy (IFE) devices hold incredible promise as a source of clean and sustainable energy, but there are significant obstacles to obtaining and harnessing IFE in a controllable manner. Led by University of California, Los Angeles researchers, this project leverages DOE supercomputing resources to study laser plasma interactions (LPI) on meaningful spatial and temporal scales of particular relevance to various IFE scenarios.

CHALLENGE A comprehensive model of LPI is crucial to the success of any IFE scheme, but one so far remains elusive. The physics involved in these processes is complex and highly nonlinear, necessitating the use of computer models, such as fully explicit particle-in-cell (PIC) simulations. The ultimate goal is to construct a hierarchy of kinetic, fluid, and other reduced description approaches capable of modeling full spatial and temporal scales. Kinetic modeling has not yet yielded sufficient understanding across the array of scales necessary to make strong connections with more approximate models and experiments.

APPROACH Using the popular PIC code OSIRIS, the team is performing fully kinetic simulations that will help advance IFE research and experimentation. The researchers are also using the UCLA Parallel PIC Framework (UPIC), a unified environment for the rapid construction of new parallel PIC codes. UPIC contains a large collection of highly optimized libraries to rapidly build new parallel PIC codes from a reliable set of components and supports electrostatic, electromagnetic, and Darwin codes using spectral field solvers.

RESULTS The researchers are studying competition between multiple instabilities (stimulated Raman scattering (SRS) and two plasmon decay) as occurs under IFE-relevant conditions, so as to develop strategies to mitigate the instabilities through the use of bandwidth or external magnetic fields. 3D simulations enabled the first-ever observation of competition between forward and backward Raman scattering. The researchers also found that magnetic fields can interfere with the nonlinear frequency shift of SRS-driven electron plasma waves EPWs, and that for SRS dominated by the dynamics of this frequency shift, magnetic fields can thereby indirectly enhance the frequency resonance between the light and plasma waves involved in SRS.

IMPACT By developing our understanding of LPI, this work could help realize IFE devices as a source of clean and sustainable energy.

Physics | Data, Learning

Reproducible, Interpretable, and Physics-Inspired AI Models in Astrophysics

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

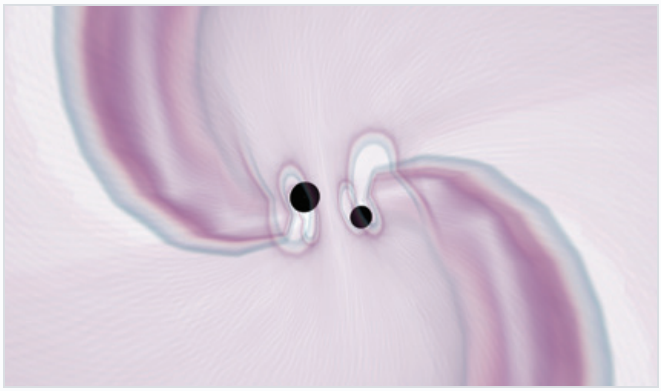

Scientific visualization of the binary black hole merger GW170814 produced on the ALCF's Theta supercomputer with the open source, numerical relativity, community software Einstein Toolkit. *Image: ALCF Visualization and Data Analysis Team*

Binary black holes in close orbit around each other emit gravitational waves that carry energy out of the system, shrinking their orbital separation until they eventually collide. Artificial intelligence (AI) methods can be used to find gravitational waves in experimental data, and characterize astrophysical properties of the compact objects that emit these cosmic messengers, providing key insights into physics, cosmology, and astrochemistry. To accelerate and increase the science reach of gravitational wave astrophysics, scientists from Argonne National Laboratory are developing novel tools at the interface of AI and high-performance computing.

CHALLENGE Gravitational wave detectors produce high velocity and high dimensional datasets with information about gravitational waves' sources. To extract new knowledge from these datasets in low latency and at scale, while other areas of research undergo similar transformations and compete for oversubscribed computing resources, new techniques to process the data more efficiently had to be developed.

APPROACH The team introduced an ensemble of AI models for gravitational wave detection. They designed a modified WaveNet architecture that takes in data from the twin LIGO (Laser Interferometer Gravitational-wave Observatory) detectors in Livingston, Louisiana, and Hanford, Washington. They optimized these models for accelerated interference using NVIDIA TensorRT and deployed them on the ALCF's ThetaGPU supercomputer. The team used the entire system to process a month of advanced LIGO data, then used time slides to quantify the performance of their AI ensemble with 5 years' worth of advanced LIGO data. The team also developed an AI model to learn and forecast the late-inspiral, merger, and ringdown of numerical relativity waveforms describing quasicircular, spinning, nonprecessing binary black hole mergers. They used ThetaGPU to train this model on 1.5 million waveforms.

RESULTS In a paper published in *Frontiers in Artificial Intelligence*, the team's interface-optimized AI ensemble for identifying mergers maintained the same sensitivity as traditional AI models, identifying all previously known binary black hole mergers, throughout August 2017, without misclassifications. Furthermore, they processed 5 years' worth of advanced LIGO data in a synthetically enhanced dataset and their AI ensemble reported an average of one misclassification for every month of searched data. Using ThetaGPU, the researchers achieved this three times faster than traditional AI models. In a paper published in *Physical Review D*, the team sampled a test set of 190,000 waveforms, finding that the average overlap, between ground truth signals and the AI predicted signals, is better than 99% across the entire parameter space. They combined visualization and accelerated computing to identify components of their AI model responsible for accurate predictions.

IMPACT This project provides tools to conduct AI-driven gravitational wave detection at scale, accelerating the identification and characterization of binary black hole mergers. This work has significance for many branches of physics and understanding the structure of the universe.

PUBLICATIONS

Khan, A., E. A. Huerta, and H. Zheng, "Interpretable AI forecasting for numerical relativity waveforms of quasicircular, spinning, nonprecessing binary black hole mergers," *Physical Review D* 105, 024024 (2022).

Pranshu, C, K. Asad, M. Tian, E. A. Huerta, and H. Zheng. "Inference-Optimized AI and High Performance Computing for Gravitational Wave Detection at Scale," *Frontiers in Artificial Intelligence* 5 (2022) 828672.

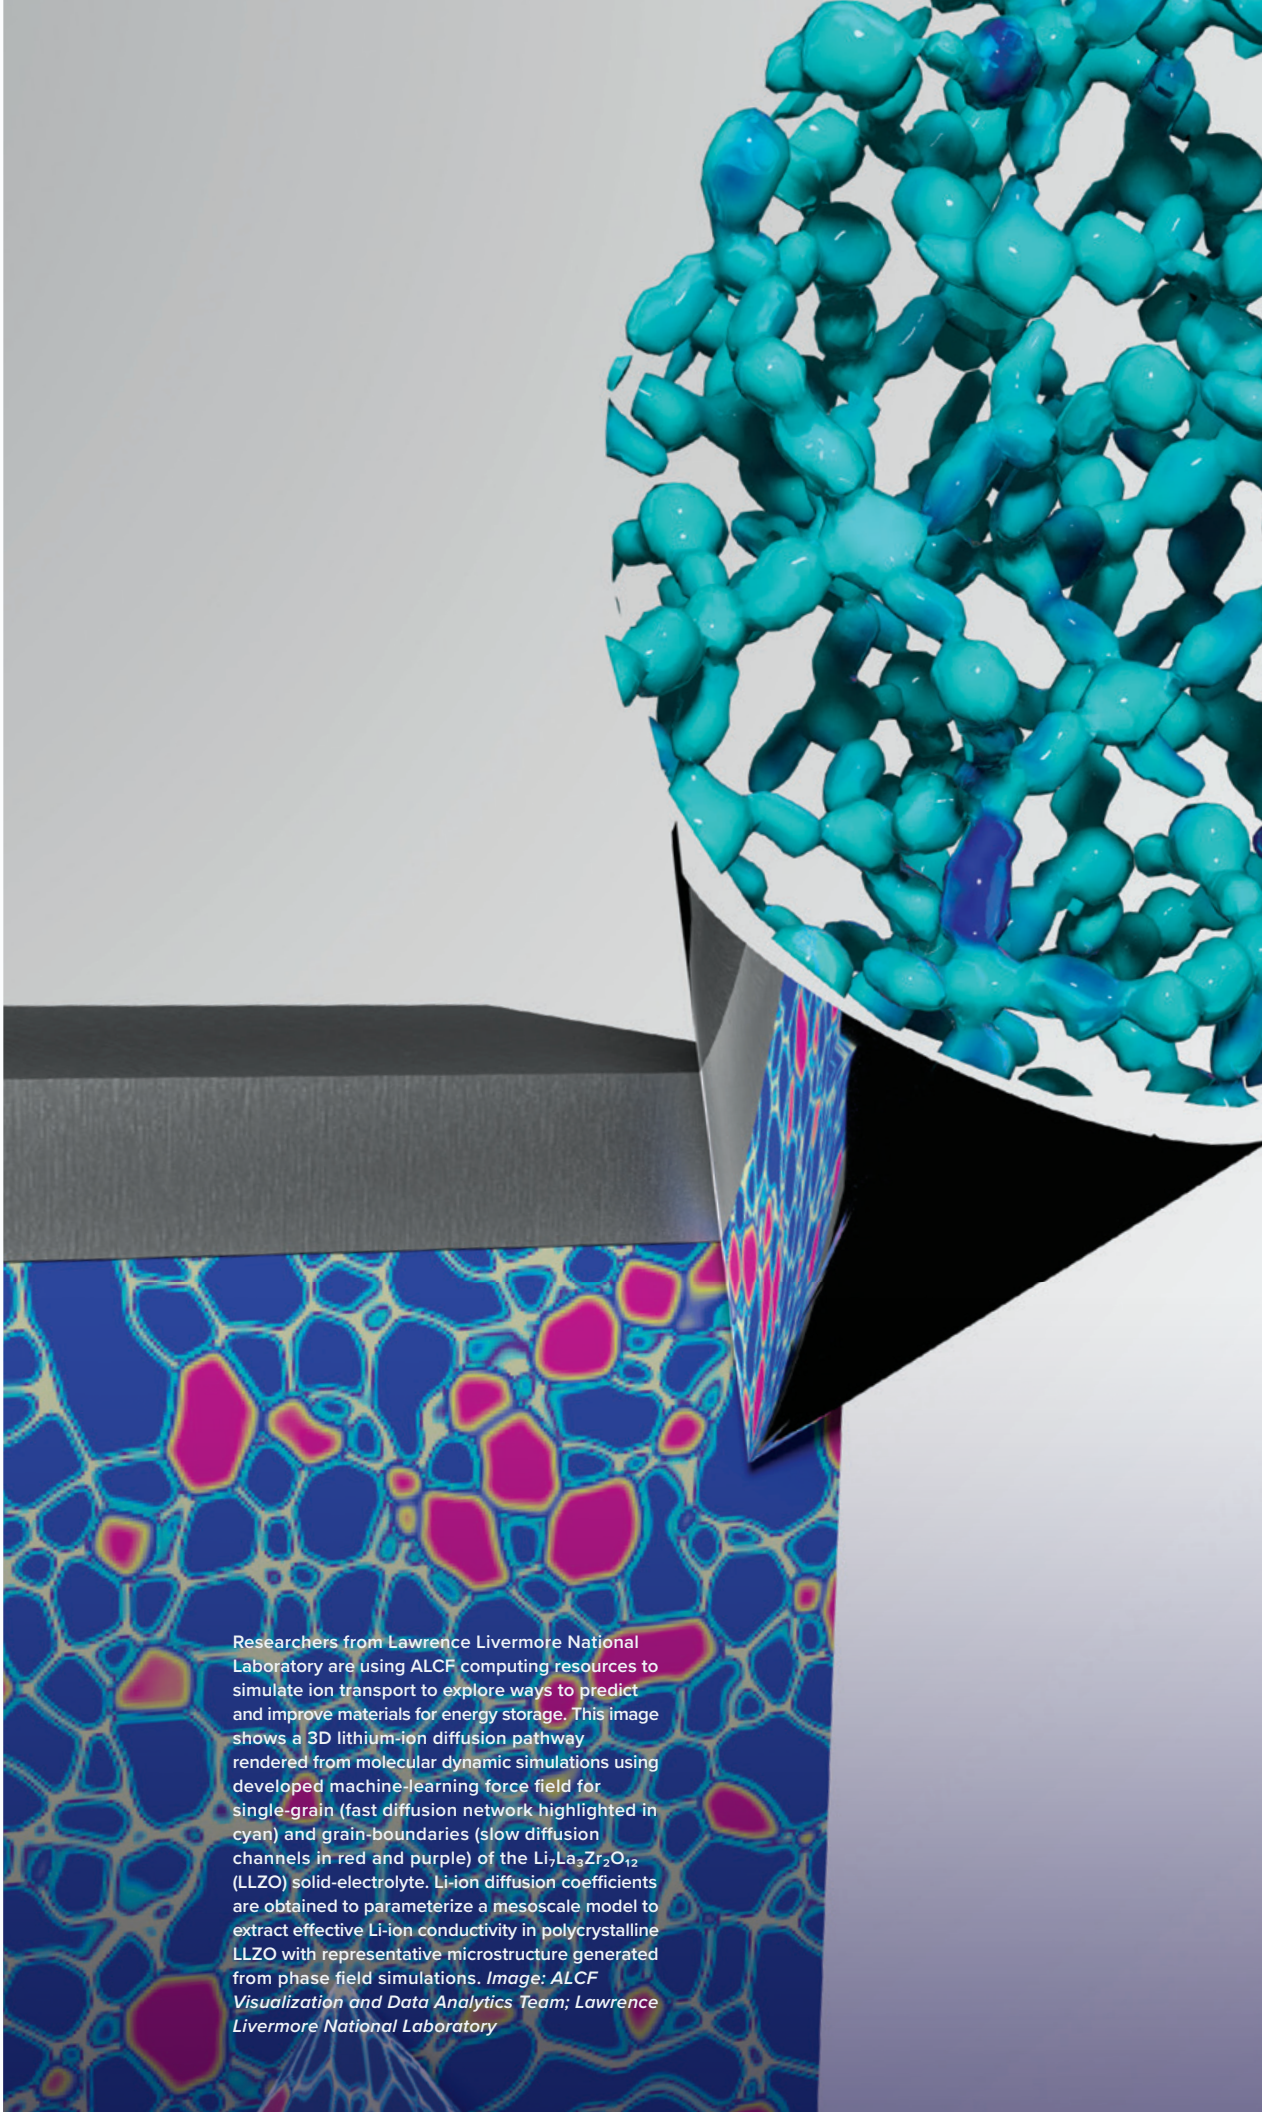

Researchers from Lawrence Livermore National Laboratory are using ALCF computing resources to simulate ion transport to explore ways to predict and improve materials for energy storage. This image shows a 3D lithium-ion diffusion pathway rendered from molecular dynamic simulations using developed machine-learning force field for single-grain (fast diffusion network highlighted in cyan) and grain-boundaries (slow diffusion channels in red and purple) of the $Li_7La_3Zr_2O_{12}$ (LLZO) solid-electrolyte. Li-ion diffusion coefficients are obtained to parameterize a mesoscale model to extract effective Li-ion conductivity in polycrystalline LLZO with representative microstructure generated from phase field simulations. *Image: ALCF Visualization and Data Analytics Team; Lawrence Livermore National Laboratory*

ALCF Projects

INCITE 2022

BIOLOGICAL SCIENCES

COMPBIO2: COMbining deep-learning with Physics-Based affinIty estimatiOn 2

PI Peter Coveney, University College London
HOURS ALCF: 100,000 Node-Hours

Design of Peptides and Proteins on Classical and Quantum Computing Hardware

PI Vikram Mulligan, Flatiron Institute
HOURS ALCF: 1,070,000 Node-Hours

CHEMISTRY

Heterogeneous Catalysis as a Collective Phenomenon Within Dynamic Ensembles of States

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

Implementing a Dispersive Interaction Database through Quantum Monte Carlo

PI Matteo Barborini, University of Luxembourg
HOURS ALCF: 1,000,000 Node-Hours

Nuclear Matter Dynamics in Real Time and the Heaviest Elements in Nature

PI Aurel Bulgac, University of Washington
HOURS ALCF: 200,000 Node-Hours
OLCF: 140,000 Node-Hours

EARTH SCIENCE

New Window into Tropical Meteorology with Global 1 km Atmosphere-Ocean Simulations

PI Inna Polichtchouk, European Centre for Medium-Range Weather Forecasts
HOURS ALCF: 50,000 Node-Hours
OLCF: 520,000 Node-Hours

ENERGY TECHNOLOGIES

Automatic Building Energy Modeling (AutoBEM)

PI Joshua New, Oak Ridge National Laboratory
HOURS ALCF: 800,000 Node-Hours

ENGINEERING

First Principles Simulation of Hypersonic Flight

PI Maninder Grover, Air Force Research Laboratory / University of Dayton Research Institute
HOURS ALCF: 1,000,000 Node-Hours

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

PI Johan Larsson, University of Maryland
HOURS ALCF: 1,750,000 Node-Hours

Online Machine Learning for Large-Scale Turbulent Simulations

PI Kenneth Jansen, University of Colorado Boulder
HOURS ALCF: 1,200,000 Node-Hours

MATERIALS SCIENCE

Large-Scale Simulations of Light-Activated Matter

PI Giulia Galli, University of Chicago and Argonne National Laboratory
HOURS ALCF: 600,000 Node-Hours
OLCF: 290,000 Node-Hours

Predicting Ion Transport Kinetics at Complex Interfaces for Energy Storage

PI Brandon Wood, Lawrence Livermore National Laboratory
HOURS ALCF: 1,120,000 Node-Hours

Towards Predictive Calculations of Functional and Quantum Materials

PI Paul Kent, Oak Ridge National Laboratory
HOURS ALCF: 1,000,000 Node-Hours
OLCF: 430,000 Node-Hours

Ultrafast Control of Functional Materials

PI Priya Vashishta, University of Southern California
HOURS ALCF: 1,700,000 Node-Hours

PHYSICS

Ab-initio Nuclear Structure and Nuclear Reactions

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

Approaching Exascale Models of Astrophysical Explosions

PI Michael Zingale, Stony Brook University
HOURS ALCF: 100,000 Node-Hours
OLCF: 590,000 Node-Hours

Electron Kinetic Plasma Physics of Black Hole Accretion Flows

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

First-Principles Simulations of Black Hole Accretion Flows and Coronae

PI Luca Comisso, Columbia 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: 1,440,000 Node-Hours

Internal Structure of Strong Interaction Nambu-Goldstone Bosons

PI Yong Zhao, Argonne National Laboratory
HOURS ALCF: 200,000 Node-Hours

Next-Generation 3D Core-Collapse Supernova

PI Adam Burrows, Princeton University
HOURS ALCF: 1,850,000 Node-Hours

ALCC 2022–2023

BIOLOGICAL SCIENCES

Probabilistic Comparative Modeling of Colorectal Cancer Screening Strategies

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

COMPUTER SCIENCE

A Multiscale Surrogate Model for Fracture Evolution Using DeepONet

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

CHEMISTRY

Microscopic Insight Into Transport Properties of Li-Battery Electrolytes

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

EARTH SCIENCE

A Climate Model Ensemble for Understanding Future Changes to Extreme

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

ENERGY TECHNOLOGIES

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

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

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

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

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

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

Predictive Simulations of Inertial Confinement Fusion Ablator Materials

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

ENGINEERING

HFIR DNS Simulation

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

MATERIALS SCIENCE

Computational Design of Novel Semiconductors for Power and Energy

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

PHYSICS

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

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

Energy Partition and Particle Acceleration in Laboratory Magnetized Shocks

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

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

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

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

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

Short Range Correlations from a Quantum Monte Carlo Perspective

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

The Spectrum and Structure of Hadrons

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

ALCC 2021–2022

BIOLOGICAL SCIENCES

Probabilistic Comparative Modeling of Colorectal Cancer Screening Strategies

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

CHEMISTRY

Microscopic Insight Into Transport Properties of Li-Battery Electrolytes

PI Wei Jiang, Argonne National Laboratory
HOURS ALCF: 1,032,000 Node-Hours

Multimodal Imaging with Intense X-ray Pulses

PI Phay Ho, Argonne National Laboratory
HOURS ALCF: 316,000 Node-Hours

EARTH SCIENCE

Advancing Multi-Year to Decadal Climate Prediction with High-Resolution E3SM and CESM

PI Ben Kirtman, University of Miami
HOURS ALCF: 1,000,000 Node-Hours

Improving the Representation of Mesoscale Convective Systems in Weather and Climate

PI Andreas Prein, National Center for Atmospheric Research
HOURS ALCF: 372,000 Node-Hours

Multi-Decadal, Climate-Scale Convection-Resolving Simulations for North America

PI V. Rao Kotamarthi, Argonne National Laboratory
HOURS ALCF: 300,000 Node-Hours

Multi-Scale Multi-Physics Ensemble Simulations for Aerosol-Cloud Interactions

PI Po-Lun Ma, Pacific Northwest National Laboratory
HOURS ALCF: 400,000 Node-Hours

ENERGY TECHNOLOGIES

Cavitation Dynamics in the Spallation Neutron Source Target

PI Eric Johnsen, University of Michigan
HOURS ALCF: 108,000 Node-Hours
OLCF: 4,000 Node-Hours

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

PI Dillon Shaver, Argonne National Laboratory
HOURS ALCF: 550,000 Node-Hours
OLCF: 80,000 Node-Hours

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

High-Fidelity Simulation of Flow and Heat Transfer Behavior to Support Conversion of Research Reactors with Involute Shaped Fuel Elements to Low Enriched Uranium

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

ENGINEERING

Multiscale Bubble Breakup and Gas Transfer in Turbulent Oceanic Environments

PI Parviz Moin, Stanford University
HOURS ALCF: 650,000 Node-Hours
NERSC: 500,000 Node-Hours

MATERIALS SCIENCE

Inverse Design of Multicomponent Oxide Catalysts with Generative Models and DFT

PI Rafael Gomez-Bombarelli, Massachusetts Institute of Technology
HOURS ALCF: 400,000 Node-Hours
OLCF: 50,000 Node-Hours
NERSC: 200,000 Node-Hours

Modeling of Polymeric Materials for Energy Storage Across Scales

PI Juan de Pablo, University of Chicago
HOURS ALCF: 100,000 Node-Hours

Response Functions of LaNiO2: Insights into High-Temperature Superconductivity

PI Gabriel Kotliar, Rutgers University
HOURS ALCF: 115,000 Node-Hours

PHYSICS

AI/Deep Learning Prediction & Real-Time Control for Fusion Energy Systems

PI William Tang, Princeton University
HOURS ALCF: 45,000 Node-Hours

LBNF - PIP-II Optimization Studies for Megawatt 120-GeV Beams on Target

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

Quantum Monte Carlo Calculations of Nuclei Up to 16O and Neutron matter

PI Alessandro Lovato, Argonne National Laboratory
HOURS ALCF: 632,000 Node-Hours

ALCF DATA SCIENCE PROGRAM

Autonomous Molecular Design for Redox Flow Batteries

PI Logan Ward, Argonne National Laboratory

Dynamic Compressed Sensing for Real-Time Tomographic Reconstruction

PI Robert Hovden, University of Michigan

Learning Optimal Image Representations for Current and Future Sky Surveys

PI George Stein, Lawrence Berkeley National Laboratory

Machine Learning for Data Reconstruction to Accelerate Physics Discoveries in Accelerator-Based Neutrino Oscillation Experiments

PI Marco Del Tutto, Fermi National Accelerator Laboratory

AURORA EARLY SCIENCE PROGRAM

Accelerated Deep Learning Discovery in Fusion Energy Science

PI William Tang, Princeton Plasma Physics Laboratory

Dark Sky Mining

PI Salman Habib, Argonne National Laboratory

Data Analytics and Machine Learning for Exascale Computational Fluid Dynamics

PI Kenneth Jansen, University of Colorado Boulder

Enabling Connectomics at Exascale to Facilitate Discoveries in Neuroscience

PI Nicola Ferrier, Argonne National Laboratory

Exascale Computational Catalysis

PI David Bross, Argonne National Laboratory

Extending Moore's Law Computing with Quantum Monte Carlo

PI Anouar Benali, Argonne National Laboratory

Extreme-Scale Cosmological Hydrodynamics

PI Katrin Heitmann, Argonne National Laboratory

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

PI Amanda Randles, Duke University

Extreme-Scale Unstructured Adaptive CFD

PI Kenneth Jansen, University of Colorado Boulder

High-Fidelity Simulation of Fusion Reactor Boundary Plasmas

PI C.S. Chang, Princeton Plasma Physics Laboratory

Machine Learning for Lattice Quantum Chromodynamics

PI William Detmold, Massachusetts Institute of Technology

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

PI Noa Marom, Carnegie Mellon University

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

PI Theresa Windus, Iowa State University and Ames Laboratory

Simulating and Learning in the ATLAS Detector at the Exascale

PI Walter Hopkins, Argonne National Laboratory

Virtual Drug Response Prediction

PI Rick Stevens, Argonne National Laboratory

DIRECTOR'S DISCRETIONARY

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

BIOLOGICAL SCIENCES

Combining Physics-Based Analysis with AI to Accelerate COVID-19 Drug Discovery

PI Tom Brettin and Hyunseung Yoo, Argonne National Laboratory

Large Ensemble Model of Single-Cell 3D Genome Structures

PI Jie Liang, University of Illinois at Chicago

Modeling the Coronavirus

PI Zhangli Peng, University of Illinois at Chicago

Non-Synaptic Communication and Axonal Remodeling After Exposure to Cocaine

PI Bobby Kasthuri, Argonne National Laboratory

Observing the SARS-CoV-2 Replication-Transcription Machinery in Action

PI Arvind Ramanathan, Argonne National Laboratory

Targeting Intrinsically Disordered Proteins Using Artificial Intelligence Driven Molecular Simulations

PI Arvind Ramanathan, Argonne National Laboratory

CHEMISTRY

Adsorptive CO_2 Removal from Dilute Sources

PI John Low, Argonne National Laboratory

Massively Parallel Electronic Stopping Simulations of High Energy Particles in Solvated DNA

PI Yosuke Kanai, University of North Carolina at Chapel Hill

Multimodal Imaging with Intense X-ray Pulses

PI Phay Ho, Argonne National Laboratory

Structure Elucidation for Nuclear Magnetic Resonance via Structured Prediction

PI Eric Michael Jonas, University of Chicago

COMPUTER SCIENCE

Bridging Data Center AI Systems with Edge Computing for Actionable Information Retrieval

PI Zhengchun Liu, Argonne National Laboratory

Data-Driven Molecular Engineering of Solar-Powered Windows

PI Jacqueline Cole, University of Cambridge

Training of Language Models on Large Quantities of Scientific Text

PI Ian Foster, Argonne National Laboratory

Unsupervised Analysis of Satellite Cloud Imagery

PI Ian Foster, Argonne National Laboratory

MATERIALS SCIENCE

BraggNN: Fast X-ray Bragg Peak Analysis Using Deep Learning

PI Zhengchun Liu, Argonne National Laboratory

Data-Driven Molecular Engineering of Solar-Powered Windows

PI Jacqueline Cole, University of Cambridge

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

PI Trevor David Rhone, Harvard University

Next-Generation Nonwoven Materials Manufacturing

PI Ian Foster, Argonne National Laboratory

Structure and Properties of Grain Boundaries in Materials for Energy Applications

PI Wissam A Saidi, University of Pittsburgh

PHYSICS

Deterministic Machine Learning for Latency-Critical Control Systems

PI Kyle Felker, Argonne National Laboratory; William Tang, Princeton Plasma Physics Laboratory

Model Quality Breakthrough for Cosmic Background Removal

PI Corey Adams, Argonne National Laboratory

Neural Network Quantum States for Atomic Nuclei

PI Alessandro Lovato, Argonne National Laboratory

Reproducible, Interpretable, and Physics-Inspired AI Models in Astrophysics

PI Eliu Huerta, Argonne National Laboratory

About the Argonne Leadership Computing Facility

Argonne's Leadership Computing Facility Division operates the Argonne Leadership Computing Facility (ALCF) as part of the U.S. Department of Energy's effort to provide leadership-class computing resources to the scientific community. The ALCF is supported by the DOE Office of Science, Advanced Scientific Computing Research (ASCR) program.

About Argonne National Laboratory

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

Availability of this Report (ANL/ALCF-22/3)

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

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

Design and production: Sandbox Studio, Chicago

Argonne NATIONAL LABORATORY

Argonne Leadership Computing Facility

CONTACT

media@alcf.anl.gov

alcf.anl.gov

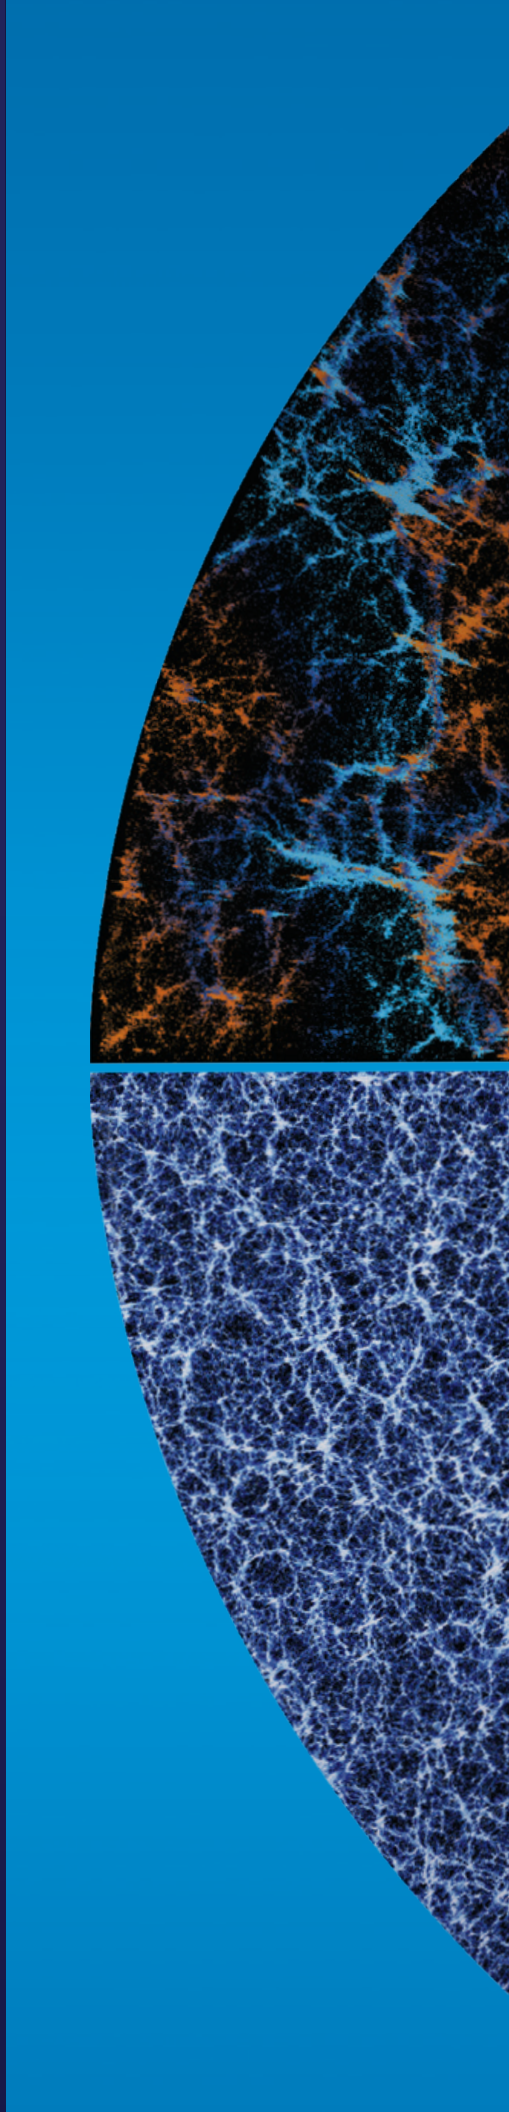

U.S. DEPARTMENT OF ENERGY

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