



Title: Integrated Exascale Computational Workflows for Accelerated Material Synthesis

Principal Investigator: Soumendu Bagchi (Oak Ridge National Laboratory)

Co-investigators: Panchapakesan Ganesh (Oak Ridge National Laboratory),
Ryan Morelock (Oak Ridge National Laboratory),
Ayana Ghosh (Oak Ridge National Laboratory),
Debangshu Mukherjee (Oak Ridge National Laboratory),
Emil Briggs (North Carolina State University),
Jerzy Bernholc (North Carolina State University)

ALCC Allocation:
Site(s): Oak Ridge Leadership Computing Facility (OLCF),
National Energy Research Scientific Computing Center (NERSC)

Allocation(s): 400,000 on Frontier,
150,000 on Perlmutter-GPU

Research Summary:

Traditional approaches to bridge atomic scale dynamics with experimental observations at the microstructural level often rely on phenomenological models or force-fields of the underlying physics, whose free parameters are in turn fitted to a small number of *intuition-driven* atomic scale simulations under limited number of thermodynamical and kinetic drivers. This approach becomes particularly cumbersome to study synthesis and characterization of materials with complex dependencies on local environment. This ALCC award will be used to deliver *exascale Bayesian workflows that couple automated high-throughput large-scale ab initio* (RMG (<https://rmgdft.sourceforge.net/>)) and classical (LAMMPS) simulations with a wide range of uncertainty quantification-driven active ensemble learning paradigms for *on-the-fly* force-field generation, tracking of structural transformation in molecular dynamics during targeted synthesis of multifunctional materials, and creation of digital-twins for synthesis to train AI/ML models and guide automated experimental synthesis (or manipulation) platforms.

At the Center for Nanophase Materials Sciences (CNMS) in ORNL, initially supported by the self-driving laboratories (INTERSECT) initiative, we have developed automated cross-platform workflows, as part of an ensemble driven materials workflow infrastructure that can efficiently exploit exascale resources autonomously. Within a networked computational framework incorporating ORNL's leadership theory, computing, experimental synthesis and characterization via advanced AI/ML approaches, our framework will enable theory-guided autonomous growth (or manipulation) of new thin-film heterostructures including transition metal chalcogenides with targeted multifunctional properties, while at the same time incorporating experimental feedback to improve simulation protocols. This integrated computational infrastructure strives to close the loop of autonomous discovery and experimental synthesis (or processing) of quantum materials thereby drastically reducing the time it takes to realize new multifunctional materials leveraging DOE's leadership class computing facilities.



Title: High-Precision Heterogeneous Catalysis by Quantum Monte Carlo

Principal Investigator: Michal Bajdich (SLAC National Accelerator Laboratory)

Co-investigators: Lubos Mitas (North Carolina State University)

ALCC Allocation:

Site(s): Argonne Leadership Computing Facility (ALCF)

Allocation(s): 1,000,000 on Aurora,
150,000 on Polaris

Research Summary:

This project aims to advance the accuracy and reliability of computer-based predictions for catalytic processes, which are vital for developing efficient and sustainable energy technologies. Leveraging state-of-the-art quantum Monte Carlo (QMC) methods and powerful supercomputing resources, the team will produce highly precise datasets to improve understanding and modeling of catalytic reactions, particularly where existing computational methods often fall short, or when the experiments are non-existent.

Project Impact:

Efficient catalysts are essential for addressing the global energy crisis and environmental challenges by speeding up chemical reactions involved in renewable energy processes and in carbon capture technologies. However, traditional computational tools used to predict catalytic reactions often struggle to achieve the necessary accuracy, particularly for complex materials that contain transition-metal such their surfaces, single-atom catalysts, and molecular complexes. This project will utilize state-of-the-art Quantum Monte Carlo (QMC) techniques in connection with this ASCR award to advance computational methods suited to modeling such challenging catalytic systems.

By employing QMC methods and running extensive simulations on supercomputers, including Aurora and Polaris, the project will generate benchmark-quality data that surpasses current computational standards as well as resolve real experimental catalytic problems. These results will help refine existing theoretical frameworks and significantly enhance the reliability of catalyst design. Furthermore, the team will make the generated datasets publicly available at our Catalysis-hub.org, facilitating broad scientific collaboration and accelerating the development of catalysts capable of efficiently converting carbon dioxide and producing renewable fuels. This work aligns closely with the Department of Energy's mission of fostering sustainable and clean-energy technologies.



Title: Hadronic Contributions to the Muon $g-2$ from Lattice QCD

Principal Investigator: Thomas Blum (University of Connecticut)

Co-investigators: Peter Boyle (BNL),
Carleton DeTar (University of Utah),
Aida El-Khadra (University of Illinois),
Steven Gottlieb (Indiana University),
Luchang Jin (University of Connecticut),
Christoph Lehner (University of Regensburg),
Ethan Neil (University of Colorado)

ALCC Allocation:
Site(s): Oak Ridge Leadership Computing Facility (OLCF),
Argonne Leadership Computing Facility (ALCF)

Allocation(s): 2,000,000 on Frontier,
2,000,000 on Aurora

Research Summary:

The muon is an elementary particle identical to the ordinary electron except that it is about 200 times heavier. Its magnetic dipole moment has been measured at Fermilab in an historic experiment (E989) to a precision of 1.27 parts-per-million. To test the limit of our current understanding of Nature's laws (aka, the standard model of particle physics) and fully leverage the DOE's investment in E989, the dominant theory errors arising from the contributions to the magnetic moment from the cloud of virtual quarks, anti-quarks, and gluons surrounding the muon during its brief lifetime must be reduced to this same level of precision. To reach this goal, these so-called hadronic contributions will be determined using numerical simulations of Quantum Chromodynamics (QCD) known as lattice QCD. QCD encapsulates the so-called strong force, one of the four fundamental forces of Nature comprising the standard model.



Title: Nuclear Thermal Hydraulics Simulations
Using Interface Capturing Approach

Principal Investigator: Igor Bolotnov (North Carolina State University)

ALCC Allocation:

Site(s): National Energy Research Scientific Computing Center (NERSC)

Allocation(s): 350,000 on Perlmutter-CPU

Research Summary:

The planned ALCC project will capitalize on existing multiphase flow experience to utilize interface capturing methods and direct numerical simulation to perform state of the art large scale simulations two-phase flows. PHASTA code has a long history of HPC performance, and our group has been awarded 5 ALCC awards in the past 10 years advancing the mission of DOE. PHASTA is a finite-element based flow solver with level-set method for interface capturing approach.

Three major research projects would be supported by this allocation: (i) project on simulation of complex liquid/steam two-phase flows through nuclear reactor steam separator geometries; (ii) project on bubble coalescence modeling to improve two-phase flow closure laws; (iii) two-phase boiling analysis to support the fundamental understanding of flow boiling.

All three subprojects are tightly aligned with DOE's mission statement: "is to ensure America's security and prosperity by addressing its energy, environmental and nuclear challenges through transformative science and technology solutions". They help resolve the existing challenges in predictive capabilities of two-phase flow and heat transfer. All of those are directly related to modern and future energy generation and transformation and involve HPC capabilities to demonstrate novel approaches of HPC applications to energy-related problems.



Title: Turbulence Database for Fusion Energy Science

Principal Investigator: Jeff Candy (General Atomics)

Co-investigators: Emily Belli (General Atomics),
Igor Sfiligoi (UCSD),
Chris Holland (UCSD),
Tom Neiser (General Atomics)

ALCC Allocation:

Site(s): Oak Ridge Leadership Computing Facility (OLCF),
Argonne Leadership Computing Facility (ALCF)

Allocation(s): 300,000 on Frontier,
300,000 on Aurora

Research Summary:

Accurate calculation of turbulent energy and particle flow is essential to the design and assessment of next-generation fusion pilot plants (FPPs). 5D gyrokinetic simulations have a proven track record of reproducing turbulent flows measured in existing experiments, providing high confidence for flow prediction in future reactor-scale devices. These high-fidelity simulations are fundamentally essential for both physics discovery and for calibrating reduced fusion engineering models. Industry-standard physics-based reduced models currently perform poorly in numerous critical operating regimes, motivating more systematic coverage of parameter space with direct simulation.

Leadership-scale CGYRO gyrokinetic simulations will be carried out to map the parametric dependence of plasma flows in a series of critical reactor operating regimes, particularly near the threshold of turbulence onset. This region is poorly-represented in existing simulation databases. The results will support DOE FES mission-critical needs for ITER and FPPs, and will facilitate AI/ML transport model development through expanded access to well-curated turbulence data. Results will directly support two US DOE SciDAC-5 FES Partnership Projects: *Surrogate Models for Accurate and Rapid Transport Solutions* (SMARTS) and *Frontiers in Leadership Gyrokinetic Simulation* (FRONTIERS). Results will be immediately available via the NERSC-CGYRO database (</global/cfs/cdirs/cgyrodb/gsharing>).



Title: **Massively Parallel Large Eddy Simulations for High-Efficient Gas Turbines Operating with Hydrogen and High Aerodynamic Loading**

Principal Investigator: Venkata Chaluvadi (GE Vernova Gas Power Engineering)

Co-investigators: Mark Flanagan (GE Vernova Gas Power Engineering)

ALCC Allocation:

Site(s): Oak Ridge Leadership Computing Facility (OLCF)

Allocation(s): 500,000 on Frontier

Research Summary:

The current computational proposal, which is part of a multi-year technology roadmap will focus on validating the in-house LES CFD solver against available high lift cascade data. This validation work will include various high aerodynamic loading vane and blade analyses and will focus on matching several aerodynamic loss sources: profile loss, trailing edge shedding loss, and end-wall secondary loss. Through previous OLCF awards, GE Vernova has been able to demonstrate preliminary validation of this flow phenomena for a limited set of conditions and geometries. This award will continue enhancing this validation and building flow understanding. After completion of the high aerodynamic loading LES CFD validation, the LES analysis system will be used to assess and optimize the performance of a GE Vernova gas turbine-representative stage 1 nozzle, positioned behind an advanced combustion system designed for operation with high hydrogen levels and axial fuel staging.

The GE Vernova Gas Power business is committed to developing environmentally friendly and efficient power generation equipment that leads in the energy transformation with reduced carbon emissions. To achieve the next generation global decarbonization goals, new gas turbine technologies that raise the aerothermal performance of Combined Cycle Power (CCP) plants and enable the use of zero-carbon hydrogen fuels need to be developed. Such breakthroughs in performance, combined with reductions in CCP cost per produced KWH, will help replace legacy coal-powered plants and provide real carbon-footprint reductions over the next 5-10 years. To achieve these goals, the adoption of hydrogen fuels for CCP plants will pose challenges to the interaction of the combustion system and the hot gas path (HGP). Future combustion systems operating on hydrogen are expected to employ higher levels of axial fuel staging (AFS) to manage nitrous oxides (NOx), which may lead to non-uniform turbine inlet conditions with high swirl and pockets of high turbulence. These non-uniform turbine inlet conditions may negatively impact component life and the operation of a highly efficient HGP. This research will improve the understanding of these complex flow interactions and lead to more efficient and cost-effective gas turbine designs.



Title: HyPowerGen

Principal Investigator: Jacqueline Chen (Sandia National Laboratories)

Co-investigators: Martin Rieth (Sandia National Laboratories),
Bruno Soriano (Sandia National Laboratories)

ALCC Allocation:

Site(s): Oak Ridge Leadership Computing Facility (OLCF)

Allocation(s): 638,000 on Frontier

Research Summary:

Hydrogen is expected to play a key role in providing secure, reliable and affordable clean energy co-fired with natural gas and as an energy carrier to power gas turbines and other hard-to-electrify sectors. Natural gas combined with hydrogen provides dispatchable electricity and industrial heat generation. Further development of hydrogen gas turbine technologies are needed to facilitate process integration and cost reduction to enable reliable fuel-flexible operation. The DOE Office of Fossil Energy and Carbon Management (FECM) is interested in research and development of novel strategies for fossil fuel and alternative fuel utilization for power and heat generation for U. S. industrial manufacturing.

To meet these objectives, direct numerical simulations (DNS) of a hydrogen micromix (MMX) gas turbine combustor will be performed to shed light on the mixing process, flame stabilization and emissions of hydrogen use. The micromix combustor concept relies on a large number of high-speed hydrogen injections into a crossflow of air. Injection of non-premixed hydrogen at high speeds mitigates flashback risks that can otherwise pose significant challenges with hydrogen, while also ensuring rapid mixing that reduces nitrogen oxide emissions. While this concept has been demonstrated with hydrogen for small gas turbines, challenges remain for medium- and large-sized gas turbines, mainly due to the high reactivity of hydrogen and thermo-diffusive instabilities that are amplified at higher pressures. Turbulent flame speed scaling for hydrogen-rich blends will be developed (e.g., from pure hydrogen to pure natural gas) based on DNS and existing experimental data from turbulent hydrogen jet flames at relevant operating conditions to industrial gas turbines. DNS of a single or array of micro-mix combustor(s) for select operating conditions will provide fundamental understanding of the flame stabilization behavior at low equivalence ratios and for different recirculation patterns that minimize residence times in hot regions, prevent flashback and minimize NO_x. The DNS will be performed in a simplified injector (MMX focus) configuration and validated against companion experimental data from industrial and academic partners. A set of such simulations will be performed to span atmospheric and elevated pressures as well as different momentum ratios (characterizing the interaction of the jet and crossflow). Ultimately, the DNS data will aid the design of fully hydrogen-capable combustors for carbon-free heat and power generation, as well as improve subgrid combustion models for large-eddy simulation (LES) of micromix conditions with hydrogen-rich fuels for computational fluid dynamics (CFD) design optimization.



Title: Real-Time Adaptive Disruption Forecasting at DIII-D

Principal Investigator: Ryan Coffee (SLAC National Accelerator Laboratory)

Co-investigators: David Rogers (ORNL-NCCS)

ALCC Allocation:

Site(s): Oak Ridge Leadership Computing Facility (OLCF)

Allocation(s): 350,000 on Frontier

Research Summary:

This project will extend an existing statically trained machine-learning based disruption prediction model for tokamak fusion reactors by leveraging a meta-learning method for fast optimization of the plasma state forecasting model and also leveraging an encoder/decoder model that accommodates a dynamic quantization scheme. The quantization will be optimized to explore the space of model and feature encoding. This exploration will inform decisions about reduced fidelity diagnostics that would remain sufficient for a reactor regime tokamak state prediction consistent for energy on the grid.

The demonstration of a real-time updating plasma state forecasting model for plasma disruption events at the DIII-D tokamak represents a critical technology development for fusion energy on the electric grid. Historical data acquisition systems archived data to storage rather stream it through HPC for automating real-time control system input. This project tightens the coupling between Exa-scale computing and the fusion reactor to the point of hosting key functions of a semi-autonomous control system for real-time plasma control. The results of this project will inform how a national effort for fusion energy will integrate with the computational resources available to both public and private efforts for a fusion energy pilot plant.

Transformers represent a generic neural network architecture that is capable of representing a complicated and even nonlinear functional systems within a latent space that is trained via self-supervised pre-training techniques. Recent work has shown that discrete tokenizers can be used to encode both image and video data as a fused latent representation. Doing so allows the multi-modal model to effectively leverage hidden correlations. To address the difficulty in forecasting plasma disruptions from the low-fidelity diagnostic signals for reactor regime within-shot conditions, we will adapt deployed models in real-time to the incoming data streams. This new paradigm creates data-driven optimization algorithms through meta-learning and have been shown to achieve strong generalization when they are trained on a sufficiently large number of tasks that require leadership computing resources. We will also use the allocation to explore the optimal use of reduced bit depth in training and encoding to optimize along the Pareto frontier for the conflicting goals of sufficiently high forecasting accuracy versus the inference latency which itself limits the control loop feedback bandwidth.



Title: Neutrino-nucleus Interactions from Lattice QCD

Principal Investigator: William Detmold (MIT)

Co-investigators: Phiala Shanahan (MIT),
Robert Perry (MIT),
Anthony Grebe (Fermilab),
Michael Wagman (Fermilab),
William Jay (Colorado State University),
Marc Illa (University of Washington),
Assumpta Parrreno (University of Barcelona),
Zohreh Davoudi (University of Maryland)

ALCC Allocation:

Site(s): Oak Ridge Leadership Computing Facility (OLCF)

Allocation(s): 970,000 on Frontier

Research Summary:

Neutrino physics is entering a new era of precision and discovery. The Deep Underground Neutrino Experiment (DUNE), currently under construction, is a flagship Department of Energy project that will use the world's most intense neutrino and antineutrino beams to shed light on the mysteries of neutrino masses and mixings. Furthermore, searches are ongoing for neutrinoless double- β (0nbb) decay, a hypothetical process that, if discovered, would reveal the fundamental nature of neutrinos and help explain the matter--antimatter asymmetry in the universe. These searches are tantalizingly close to the sensitivity needed to either rule out one possible scenario for the neutrino-mass hierarchy or to discovering the process. The construction of a next-generation tonne-scale 0nbb-decay experiment that aims to reach this definitive sensitivity is a high priority of the nuclear physics community. For both programs, the interpretation of upcoming experimental results will require controlled theoretical understanding of neutrino interactions with nuclei from first principles, and the goal of this project is to provide that vital input. In particular, this project will provide first-principles determinations of weak interactions in one- and two-nucleon systems using the technique of lattice QCD. To inform calculations of neutrino-nucleus cross sections for DUNE, the deuteron axial form factor will be calculated for the first time, along with associated nucleon and transition form-factors. To aid in interpretation of 0nbb-decay experiments, required nuclear matrix elements will also be determined for the first time. Machine learning (ML) will be used to optimize of the lattice QCD framework used for these calculations and will play an important role in achieving the proposal goals



Title: **Theseus: A Computational Science Foundation Model**

Principal Investigator: Patrick Emami (National Renewable Energy Laboratory)

Co-investigators: Sameera Horawalavithana (Pacific Northwest National Lab),
Jason Eisner (Johns Hopkins University),
Shaowu Pan (Rensselaer Polytechnic Institute)

ALCC Allocation:

Site(s): National Energy Research Scientific Computing Center (NERSC)

Allocation(s): 160,000 on Perlmutter-GPU

Research Summary:

Scientific discovery is a key driver of innovation. However, scientists are having trouble keeping up with the exploding volume of scientific information, particularly with the growth of data-intensive and compute-intensive science. Many of today's grand challenges expect scientific teams to reason about models, data, and background knowledge across multiple disciplines. The amount of uncertainty along the path to each new discovery or insight is rapidly growing. This presents a high barrier to entry, requiring large multidisciplinary teams that must organize their work around well-defined questions. Curiosity-driven research is becoming difficult.

Enthusiasm has grown around the idea of using artificial intelligence (AI) to accelerate scientific discovery. We're developing frameworks for science assistants, which are conversational AI systems that collaborate with scientists to construct scientific models and reason about ambiguous data. They also answer questions about scientific literature and help accelerate computational workflows by writing code. Assistants are powered by foundation models, which are AI models trained on Internet-scale amounts of data using high-performance computing. The first phase of the project will develop scientific benchmarks and advances in cutting-edge AI topics such as multimodality, uncertainty quantification, and reasoning frameworks that augment foundation models with probabilistic reasoning abilities and connects them with scientific tools and verified knowledge. In the second phase, we will connect scientists with prototypes of our science assistant to conduct dialogues about scientific modeling in areas such as computational fluid dynamics and materials characterization. Insights from these studies will inform future research towards AI that can act as responsive and knowledgeable science assistants. The main outcome will be the release of our assistant framework, Theseus, in addition to open data, software, and foundation models that demonstrate key improvements for scientific applications.



Title: Development of AI-powered Application for Advanced Nuclear Reactor Design

Principal Investigator: Jun Fang (Argonne National Laboratory)

Co-investigators: Aleksandr Obabko (Argonne National Laboratory),
Tingzhou Fei (Argonne National Laboratory)

ALCC Allocation:
Site(s): Argonne Leadership Computing Facility (ALCF)

Allocation(s): 100,000 on Aurora

Research Summary:

This project develops an artificial intelligence (AI)-powered application to accelerate and optimize the design of advanced nuclear reactors, with a focus on the Molten Salt Reactor (MSR), a promising technology known for its intrinsic safety and high thermal efficiency. By integrating high-fidelity Computational Fluid Dynamics (CFD) simulations with machine learning (ML), the project aims to automate the complex design process that traditionally requires expert input and time-consuming iterations. The resulting app will enable users—regardless of their background in CFD or ML—to explore optimized reactor configurations by simply inputting key design parameters. The system will feature two main components: a CFD module based on the NekRS code, and an AI/ML module that builds surrogate models from simulation data to guide design improvements.

The project is expected to significantly lower the barriers to advanced reactor design by reducing time, cost, and expertise requirements. It supports the U.S. Department of Energy’s mission to advance clean, secure, and sustainable energy solutions by enabling digital engineering tools that can accelerate innovation in nuclear energy. By demonstrating a novel workflow that bridges nuclear engineering and AI, the project lays the groundwork for a transformative shift in how reactors are conceptualized and built, potentially drawing interest from a wide range of technology vendors and stakeholders across the energy sector.



Title: **Ab Initio Simulations of Out-of-Equilibrium Heterogeneous Quantum Materials**

Principal Investigator: Panchapakesan Ganesh (Oak Ridge National Laboratory)

Co-investigators: Jacek Jakowski (Oak Ridge National Laboratory),
Liangbo Liang (Oak Ridge National Laboratory),
Wenchang Lu (North Carolina State University),
Emil Briggs (North Carolina State University),
Jerzy Bernholc (North Carolina State University)

ALCC Allocation:

Site(s): Oak Ridge Leadership Computing Facility (OLCF)

Allocation(s): 250,000 on Frontier

Research Summary:

This project will develop and apply state-of-the-art quantum dynamics approaches to perform large-scale *ab initio* simulations of heterogeneous quantum materials driven out of equilibrium. Our theoretical approaches range from long-time *ab initio* molecular dynamics simulations to quantum electronic dynamics simulations based on real-time time-dependent density functional theory (RT-TDDFT) of excited electrons that are coupled to ionic motion using perturbation theory and/or Ehrenfest dynamics. Heterogeneities such as defects and interfaces not only dominate the equilibrium properties of materials but also control their out-of-equilibrium properties and hence responses to external stimuli even at finite temperatures. Yet, to date, there is a lack of accurate quantum-mechanical approaches to describe out-of-equilibrium heterogeneous materials. This project will investigate quantum materials to: (1) understand and optimize the low-order nonlinear optical susceptibilities of chiral materials for advanced nonlinear optical applications; (2) determine the influence of thermally activated phonon modes on stability and lifetime of topologically protected edge states for room-temperature quantum applications; and (3) track the band population dynamics of excited charge carriers in quantum materials, including excited-state geometries, to aid in the interpretation of photoluminescence (PL) and time- and angle-resolved photoelectron spectroscopy (Tr-ARPES) measurements of emerging quantum materials. The simulations will use the open-source real-space multigrid (RMG) software suite (<https://github.com/RMGDFT/rmgdft>) for quantum simulations. This project will advance DOE's mission by targeting priority research goals "Advance Artificial Quantum-Coherent Systems with Unprecedented Functionality for QIS" and "Discover Novel Approaches for Quantum-to-Quantum Transduction", key to advancing new technologies for Microelectronics and Quantum Information Sciences (QIS). Expected participation from junior scientists, postdocs, and students will train the next generation of theoretical and computational quantum scientists.



Title: Transverse Momentum Dependent Helicity Distributions of the Nucleon

Principal Investigator: Xiang Gao (Brookhaven National Laboratory)

Co-investigators: Swagato Mukherjee (Brookhaven National Laboratory),
Fei Yao (Brookhaven National Laboratory),
Dennis Bollweg (Brookhaven National Laboratory),
Yong Zhao (Argonne National Laboratory),
Jinchen He (Argonne National Laboratory & University of Maryland),
Peter Boyle (Brookhaven National Laboratory)

ALCC Allocation:

Site(s): Oak Ridge Leadership Computing Facility (OLCF)

Allocation(s): 1,850,000 on Frontier

Research Summary:

Understanding the origin of proton spin is one of the most enduring mysteries in nuclear physics. Decades of experiments have revealed that the spins of the proton's constituent quarks account for only a small portion of its total spin, raising fundamental questions about the roles of gluons and quark motion. This project will perform cutting-edge simulations using lattice quantum chromodynamics (QCD), a theory that describes the strong force holding matter together, to map the proton's internal spin structure in three dimensions. The focus is on "helicity TMDs"—distributions that show how quark spin varies with momentum inside a spinning proton—providing new insight into how proton spin emerges from subatomic interactions.

By calculating these spin-dependent distributions from first principles, this project will deliver crucial nonperturbative input for interpreting upcoming results from next-generation experiments such as the Electron-Ion Collider (EIC). The simulations, powered by exascale computing on Frontier, will help constrain theoretical uncertainties in our understanding of nucleon structure and guide experimental efforts to probe the strong nuclear force. This work directly advances the U.S. Department of Energy's mission to uncover the fundamental building blocks of matter and supports key priorities outlined in recent long-range plans for nuclear science.



Title: Learning Correlation Functional in DFT from Quantum Many-body Calculations

Principal Investigator: Vikram Gavini (University of Michigan, Ann Arbor)

Co-investigators: Bikash Kanungo (University of Michigan, Ann Arbor),
Sambit Das (University of Michigan, Ann Arbor),
Paul Zimmerman (University of Michigan, Ann Arbor)

ALCC Allocation:

Site(s): Oak Ridge Leadership Computing Facility (OLCF)

Allocation(s): 290,000 on Frontier

Research Summary:

The quantum many body (QMB) methods and density functional theory (DFT) present two distinct *ab initio* methods for chemical and materials simulations. The QMB methods, such as configuration interaction (CI), provide quantum accuracy, but scale poorly with number of electrons. DFT scales favorably by reducing the many-electron wavefunction to an effective single-electron problem. However, DFT requires an exchange-correlation (XC) functional, which encapsulates the quantum many-electron interactions as a universal functional of the ground-state electron density. Existing approximations to the XC functional in DFT remain far from quantum accuracy. This teams' prior efforts in breaking this *size-accuracy barrier* involved a data-driven approach, where XC potentials corresponding to QMB electron densities were computed using inverse DFT and used to develop improved neural network based local/semi-local XC functionals. However, remedying notable deficiencies (e.g. errors in band gaps, charge-transfer, self-interaction errors) will require the use of non-local functionals. The generalized Kohn-Sham framework with exact exchange interactions provides a path forward, where hybrid functionals have been shown to substantially reduce the aforementioned errors.

In the proposed work, the data-driven approach will be extended to the generalized Kohn-Sham (GKS) framework, where the broad objective is to obtain exact correlation potentials corresponding to QMB densities and use this data to develop models for the correlation functional. The first part of the proposed work will generate the QMB based accurate electron densities for various molecules and jellium spheres. The second part of the proposed study entails using GKS-inverse DFT calculations on QMB densities to compute the exact correlation potentials. The final objective is to use the QMB densities, exact correlation potentials and energies to train models for correlation functionals and assess the accuracy of these models using a wide range of thermochemistry and solid-state benchmarks. The proposed study is expected to pave a systematic path for improved DFT functionals that can have an impact in many fields.



Title: **High-Throughput Multiscale Simulations for Designing Sustainable Copolymers**

Principal Investigator: Rafael Gomez-Bombarelli (MIT)

ALCC Allocation:

Site(s): National Energy Research Scientific Computing Center (NERSC)

Allocation(s): 243,750 on Perlmutter-CPU,
184,000 on Perlmutter-GPU

Research Summary:

Polymers are essential to the development of next-generation technologies for energy storage and industrial manufacturing, yet their immense design complexity, which compounds in copolymers made that consist of multiple monomer species, creates significant challenges. To that end, this project seeks to harness molecular simulations and machine learning to accelerate the design of two classes of advanced copolymers: solid polymer electrolytes for safer, higher-performing lithium batteries, and reprocessable thermoset plastics for high-performance industrial applications. Traditional lithium-ion battery materials suffer from limited stability, safety, and lifespan, hindering electric vehicle adoption and renewable energy integration. Similarly, high-performance thermosets, vital in the aerospace and automotive sectors, are difficult to recycle or repurpose, resulting in significant waste and inefficiency across the supply chain. To overcome these obstacles, high-throughput molecular simulations will be utilized to explore thousands of potential polymer structures, examining how variations in chemical composition, molecular architecture, and network topology affect key properties such as ionic conductivity and mechanical strength. Machine learning models, including graph-based neural networks and Bayesian active learning approaches, will be developed to rapidly predict material performance based on simulation data, thereby guiding experimental efforts toward the most promising candidates. This integrated approach enables a more systematic and accelerated exploration of the vast polymer design space, reducing reliance on slow and costly trial-and-error methods.

The anticipated impact of this work is closely aligned with several key priorities of the Department of Energy (DOE). The development of advanced solid polymer electrolytes will help speed the deployment of clean energy technologies, enhance energy security, and support the shift away from fossil fuels. Safer, longer-lasting batteries are crucial for expanding electric vehicle use and for providing reliable, efficient energy storage, which is necessary for greater adoption of renewable energy sources. At the same time, the project's focus on recyclable thermoset plastics advances the DOE's Strategy for Plastics Innovation by tackling plastic waste and enabling new recycling solutions. By integrating artificial intelligence with high-throughput simulations, the project aims to discover materials that improve resource efficiency and extend the useful life of critical industrial materials. These advances will help address major challenges in energy storage and materials management, while contributing to broader scientific progress in clean energy and advanced manufacturing.



Title: **Simulating Large-scale Long-lived Neutron Star Remnants from Binary Neutron Star Mergers**

Principal Investigator: Ore Gottlieb (Flatiron Institute)

Co-investigators: Brian Metzger (Columbia University),
Alexander Tchekhovskoy (Northwestern University),
Kyle Parfrey (Princeton University),
Francois Foucart (University of New Hampshire),
Carlos Palenzuela Luque (University of the Balearic Islands),
Daniel Kasen (UC Berkeley),
Elias Most (Caltech),
Danat Issa (Northwestern University),
Nick Kaaz (Northwestern University)

ALCC Allocation:

Site(s): Argonne Leadership Computing Facility (ALCF)

Allocation(s): 667,000 on Aurora

Research Summary:

The detection of the multi-messenger binary neutron star (NS) merger, GW170817, confirmed the long-standing prediction linking NS-NS mergers to short bursts of gamma-rays (GRBs) powered by jets, marking the dawn of the multi-messenger era. This watershed event has opened new avenues for studying the Universe's expansion rate and the NS equation of state. Additionally, the radioactive decay-powered kilonova emission from GW170817 has established NS-NS mergers as significant contributors to heavy element nucleosynthesis in the Universe. However, the central engine —whether a black hole (BH) or a NS — powering these jets remains elusive. This research aims to address this long-standing issue through first-principles simulations, focusing on the prospects of long-lived hypermassive neutron stars (HMNSs) as potential engines for short GRBs (sGRBs). The extended lifetime of HMNSs presents numerical challenges for advanced simulations over long timescales. Using state-of-the-art 3D general-relativistic magnetohydrodynamic simulations, this work will follow the outflows launched from an HMNS and compare them with those powered by BHs. This research has the potential to shed light on numerous mysteries: the central engine of sGRBs, the physical conditions at the collapse of the HMNS and those present at the time of BH formation, and provide means to distinguish between BH- and NS-powered jets.



Title: Validation of Stability Simulations and Predictive Scaling Toward Next Step FRC Fusion Device

Principal Investigator: Roelof Groenewald (TAE Technologies Inc.)

Co-investigators: Calvin Lau (TAE Technologies Inc.),
Sean Dettrick (TAE Technologies Inc.),
Toshiki Tajima (University of California, Irvine)

ALCC Allocation:

Site(s): Argonne Leadership Computing Facility (ALCF)

Allocation(s): 250,000 on Aurora

Research Summary:

TAE Technologies, Inc, is a private company focused on the development of carbon-free baseload electricity production using fusion plasmas confined in a field-reversed configuration (FRC), which the Office of Fusion Energy Science identified as a leading candidate for advanced fusion energy. TAE's national lab scale experimental device, C-2W, has achieved unprecedented steady-state FRC plasmas at record pressures using neutral beam injection (NBI) to sustain the configuration. Two landmark empirical results have emerged from TAE's experiments: (1) neutral beam current expands stable operational boundaries of FRCs, and (2) energy confinement increases with increasing electron temperature. The goal of this project is to build understanding of how these positive trends extend to fusion relevant conditions which is crucial for the success of FRC-based fusion reactors.

Building on a previous ALCC award where stable operating conditions for beam driven FRCs were identified, this project will use ALCC computing resources to identify desirable, globally stable, self-consistent, kinetic FRC states relevant to C-2W. These states will then be used in turbulent transport studies to understand energy confinement and predict how it scales to fusion relevant conditions. The particle-in-cell code, WarpX, which was developed at Lawrence Berkeley National Laboratory (LBNL) will be used for these simulations. In collaboration with LBNL, TAE added a hybrid-PIC algorithm to WarpX for the study of magnetically confined plasmas, enabling efficient simulations of macro-stability.

Recently TAE also added a semi-implicit electrostatic solver which enables large scale electrostatic micro-turbulence simulations. *In silico* demonstration that favorable energy confinement scaling with electron temperature continues to reactor conditions will reduce technical risk of TAE's fusion approach and greatly aid in securing continued funding for TAE's ultimate goal of building a viable fusion power plant.

This project will accelerate the development and commercialization of new carbon-free electricity technologies. The project leverages public investments by three major components of the Office of Science: the code was developed for the Accelerator R&D and Production (ARDAP) program, optimized for new computer architectures by the Exascale Computing Project (ECP), and enhanced with new physics models to so that it can address magnetic fusion plasmas within scope of the Fusion Energy Sciences (FES) program.



Title: Space Traffic Autonomy: Leadership Computing Advances
Hierarchical Planning R&D

Principal Investigator: Allan Grosvenor (Microsurgeonbot Inc.)

ALCC Allocation:

Site(s): Oak Ridge Leadership Computing Facility (OLCF),
Argonne Leadership Computing Facility (ALCF)

Allocation(s): 685,000 on Frontier,
685,000 on Aurora

Research Summary:

Picture a future where satellites dance through crowded orbits, effortlessly avoiding debris, and fusion reactors run like clockwork, delivering clean energy without a hitch. This project brings that vision to life with a groundbreaking AI designed for **mission-critical operations**. Unlike standard AI that falters under pressure, this system is a **hybrid intelligence marvel**—a layered network that learns, reasons, and adapts with human-like precision. It's powered by innovations like **Graph JEPA** to uncover hidden patterns, **scalable multi-agent reinforcement learning** to orchestrate complex tasks, **cross-domain knowledge transfer** to bridge insights across disciplines, and **neuro-symbolic integration** to fuse logic with instinct. In space **operations**, it harnesses **NEOWISE**'s infrared measurements, **CelesTrak**'s orbital data, and the **Universal Data Library**'s multi-sensor fusion to monitor thousands of satellites and debris, ensuring safe, real-time navigation. In fusion **operations**, it processes diagnostic data from DOE's DIII-D facility to detect subtle glitches before they become disasters, keeping reactors stable and efficient. Fueled by the raw power of Frontier and Aurora—two of the world's fastest supercomputers—this AI thrives in high-stakes **operations** where reliability is non-negotiable.

Anticipated Impact:

This isn't just science—it's a game-changer for U.S. leadership. In space **operations**, it safeguards orbits for satellites critical to communication, navigation, and national security—think sharper GPS and stronger defenses. In fusion **operations**, it enhances reactor reliability, supporting DOE's clean energy mission by the 2030s. Plus, it's driving high-tech jobs and could spark breakthroughs in fields like disaster management or autonomous tech. With a 98% success rate in predicting satellite maneuvers, this AI proves it can handle the toughest **operations**, cementing America's edge in innovation and safety.



Title: Axial-vector Form Factors for Neutrino-nucleon Scattering

Principal Investigator: Rajan Gupta (Los Alamos National Laboratory)

Co-investigators: Tanmoy Bhattacharya (Los Alamos National Laboratory),
Vincenzo Cirigliano (Institute for Nuclear Theory, University of Washington)

ALCC Allocation:
Site(s): Oak Ridge Leadership Computing Facility (OLCF)

Allocation(s): 550,000 on Frontier

Research Summary:

One of the profound mysteries of nature is the observed matter-antimatter asymmetry of the universe, i.e., the almost complete absence of anti-matter. One of the conditions for this to have arisen during the evolution of the universe from a symmetric state post inflation requires the violation of the combined symmetry of charge-conjugation and parity (called CP). CP violation of the required size can be either in the quark or neutrino sectors. To determine its size in the neutrino sector, the flagship experiment within the US called the Deep Underground Neutrino Experiment (DUNE) is being built. For the analysis of the data produced in DUNE to reach the precision needed to determine the size of CP violation in the neutrino mixing matrix, we need to know the cross-section for the interaction of neutrinos with nuclear targets, which in the case of DUNE is liquid Argon. A crucial input in this calculation is the axial-vector form factor—the energy dependent coupling with which neutrinos interact with a proton or a neutron. This can be calculated directly using large-scale simulations of the theory of quarks and gluons called quantum chromodynamics, whose bound states are neutrons and protons. Our calculation aims to calculate this axial form factor with the required precision.



Title: Privacy-Enhanced Simulation and Risk Evaluation for Collaborative Learning at Scale

Principal Investigator: Heidi Hanson (Oak Ridge National Laboratory)

Co-investigators: John Gounley (Oak Ridge National Laboratory)

ALCC Allocation:

Site(s): Oak Ridge Leadership Computing Facility (OLCF)

Allocation(s): 200,000 on Frontier

Research Summary:

AI-enabled advances in biotechnology and life sciences have expanded the biothreat landscape. Tools that enable near-real-time risk-assessment are necessary for an effective response to a catastrophic biothreat events. The ability to quickly learn from real-world data is crucial for rapid identification of emerging threats, decision support, and the evaluation of effective interventions and policies. Unfortunately, many data sources required for biopreparedness contain sensitive information and, consequently, remain siloed to ensure privacy. Existing approaches to learn across data silos, such as privacy-preserving federated learning, do provide robust guarantees against many privacy attacks but at a significant cost to utility. Moreover, federated training of large models requires substantial compute resources at every silo, which constrains the ability of many facilities to participate. Likewise, privacy protected data from each participating institution can exhibit unique patterns, creating vulnerability to mosaic attacks using public data. Addressing these challenges is a core goal of the Electronic Health Record informed Lagrangian method for preCision public Health (EHRlich) project, part of the Department of Energy's (DOE) Biopreparedness Research Virtual Environment (BRAVE) Initiative.

This project will use high-quality synthetic data produced by generative AI to reduce the utility costs of privacy guarantees while also enabling energy-efficient participation of lower-resourced facilities. The development of a generalizable suite of tools for synthetic data generation will enable broad use of population scale health data, solving a critical challenge in AI for health research. The successful completion of this project also takes advantage of world-class super computers and strong partnerships with the National Institute of Health to advance AI innovation.

This proposal would aim to address three questions via campaigns on Frontier at OLCF using CITADEL:

1. What is the risk of leaking sensitive information when training large-language models on a synthetic "twin" of private data?
2. What is the optimal approach for large-scale collaborative learning with sensitive data across heterogeneous systems?
3. What is the risk of leaking sensitive information when combining a synthetic 'twin' of private health data with synthetic trajectories of human mobility.



Title: **Cosmological Hydrodynamics for Multi-wavelength Cosmic Probes**

Principal Investigator: Katrin Heitmann (Argonne National Laboratory)

Co-investigators: Salman Habib (Argonne National Laboratory),
Nicholas Frontiere (Argonne National Laboratory)

ALCC Allocation:

Site(s): Oak Ridge Leadership Computing Facility (OLCF)

Allocation(s): 350,000 on Frontier

Research Summary:

Large-area surveys of the sky reaching ever deeper into the observable Universe provide a remarkable composite data set, covering different frequency ranges as well as different ways of probing the dynamics and matter content of the Universe. The level of statistical control is quite remarkable with spatio-temporal summary statistics being measured at the percent level or even smaller. Uncertainties due to lack of understanding of galaxy formation and the distribution of baryons affect most of the cosmological probes, and more importantly, prohibit from pushing cosmological analyses into the nonlinear regime where significant information will be harvested by the surveys. Understanding the impact of baryonic effects on cosmological probes across wavelengths is paramount to push the constraining power from ongoing and upcoming surveys to the next level. Additionally, understanding the detailed interplay of cosmological and hydrodynamics parameters, which is the aim of this project, is crucial. To reach this goal, a suite of simulations exploring both the impact of astrophysical and cosmological parameters on large-scale structure probes of cosmology will be carried out.

This project will enable breakthrough science by reducing the current modeling uncertainties, and by allowing systematics modeling using a much narrower range of uncertainty. The results of the program will strongly impact all DOE-supported cosmological surveys – the Dark Energy Spectroscopic Instrument (DESI), the Vera C. Rubin Observatory’s Legacy Survey of Space and Time (LSST), and the South Pole Telescope (SPT). In addition, we are members of the science team of the recently launched NASA mission SPHEREx, the first all-sky spectral survey of the Universe. It is exciting to contemplate a future in which the combination of data from space-based and the ground-based surveys can result in a major transformation in our knowledge of the Universe. This ALCC project will provide unique leverage in bringing this future very much closer.



Title: Informing Forensics Investigations of Nuclear Materials

Principal Investigator: Sara Isbill (Oak Ridge National Laboratory)

Co-investigators: Ashley Shields (Oak Ridge National Laboratory),
Andrew Miskowicz (Oak Ridge National Laboratory)

ALCC Allocation:
Site(s): Oak Ridge Leadership Computing Facility (OLCF)

Allocation(s): 300,000 on Frontier

Research Summary:

Researchers at Oak Ridge National Laboratory actively support national and international efforts to prevent and deter the proliferation of nuclear weapons by developing tools and methods to detect nuclear material and processes. Ongoing research efforts supported by the National Nuclear Security Administration's Office of Defense Nuclear Nonproliferation Research and Development leverage key chemical, structural, and isotopic signatures of special nuclear material production activities. This research requires the use of advanced experimental and computational resources to more fully understand process kinetics and environmental degradation of key fuel cycle materials to provide a sound technical basis for the detection and monitoring of nuclear fuel cycle and weapons development activities.

Optical vibrational spectroscopy is the primary experimental technique used for nondestructive chemical composition determination because it possesses sufficient fidelity to be of utility to nuclear forensics, and interpretation of the spectra is significantly aided by highly accurate computational models of the lattice dynamics that allow direct connections of chemical and physical changes in fuel cycle materials to the experimental observations. Additionally, determining the effect of process and environmental effects on the underlying crystal structure and reaction kinetics of fuel cycle-relevant materials can be nontrivial for complex reactions and benefits greatly from atomistic models. As such, atomistic modeling is an invaluable tool in our growing nonproliferation materials research portfolio.

The chief aim of this project is to connect highly accurate density functional theory determinations of reaction products and lattice dynamics of reactive metals and actinide materials with experiments, bridging the gap between experimental observations and their fundamental chemical or physical origins. Lattice dynamics calculations will be compared directly to ongoing experiments collecting vibrational spectra obtained from Raman, infrared, and neutron scattering spectroscopies, and reaction products will be compared to available characterization results. Experimentally determining reaction kinetics for solid-state materials is challenging and an area that will greatly benefit from computational modeling of activation barriers. Leadership-class computing resources allow the extensive number of calculations and large simulation cells required to capture experimentally relevant structural, vibrational, and kinetic properties to be done on a timeline commensurate with the accompanying experimental efforts. As a continuation of our previous ALCC awards from 2021 to present, we will use established workflows to ensure efficient use of the computational resources as we seek to advance understanding of fuel cycle-relevant materials in support of our nuclear nonproliferation mission.



Title: **High-fidelity Simulations of Bubble-laden Turbulent Flows with Surfactants**

Principal Investigator: Suhas Jain (Georgia Institute of Technology)

Co-investigators: Cyrus Aidun (Georgia Institute of Technology; Renewable Bioproducts Institute)

ALCC Allocation:

Site(s): Oak Ridge Leadership Computing Facility (OLCF),
Argonne Leadership Computing Facility (ALCF)

Allocation(s): 1,000,000 on Frontier,
1,000,000 on Aurora

Research Summary:

This computational fluid dynamics study investigates dense bubble-laden, high-Reynolds number turbulent flows with surfactants, relevant to sustainable energy and industrial applications like heat exchangers, nuclear reactors, electronics cooling, and petroleum processing. Surfactants notably impact bubble coalescence, breakup, and flow modulation, enhancing heat and mass transfer at interfaces. However, understanding bubble-turbulence interactions—especially under surfactant effects and near walls with shear—remains limited due to computational challenges and gaps in subgrid models. The project aims to advance this field through large-scale simulations, leveraging low-dissipation numerics and GPU-based supercomputing to study bubbly flows in decaying isotropic turbulence and in turbulent duct flows, with and without surfactants. Parameter sweeps will vary gas void fraction, Reynolds number, and surfactant concentration to assess effects on bubble behavior, turbulence, and engineering metrics like skin friction and foam stability.

This research has broad-reaching implications for various industries, including energy, chemical engineering, petroleum engineering, and in electronics thermal management. Studying bubbles in turbulent flows has transformative implications for the Department of Energy (DOE) and its commitment to advancing energy efficiency, sustainability, and safety. Understanding the behavior of bubbles in turbulent regimes enhances our ability to design and optimize energy systems where efficient heat and mass transfer are critical, such as nuclear reactors, advanced cooling systems, and renewable biofuel production facilities. Turbulent bubbly flows are also integral to improving the efficiency of fuel cells, gas pipelines, and large-scale reactors, where they enhance reaction rates and reduce drag, leading to significant energy savings and cost reductions. Furthermore, insights into bubble dynamics in turbulence help address challenges related to flow-induced vibrations and pressure fluctuations, which are key to maintaining the structural integrity and operational safety of pipelines and reactors.



Title: Microscopic Insight into Transport Properties of Li-battery Electrolytes

Principal Investigator: Wei Jiang (Argonne National Laboratory)

Co-investigators: Zhengcheng Zhang (Argonne National Laboratory)

ALCC Allocation:

Site(s): Argonne Leadership Computing Facility (ALCF)

Allocation(s): 600,000 on Aurora

Research Summary:

There is an increasing worldwide demand for high energy density batteries. The exploration of new Li-ion battery materials is an important focus of materials scientists and computational physicists and chemists throughout the world. The practical applications of Li-ion batteries and emerging alternatives may not be limited to portable electronic devices, and circumventing hurdles to their widespread adoption in electrical vehicle applications, requires new electrode materials and a fuller understanding of how the materials and the electrolyte chemistries behave. Computational prediction of ideal is a leading methodology in designing materials and electrolytes optimized for function, including those for Li-ion batteries.

The planned computation constitutes the simulation part of a current EERE VTO project. This research is aimed at using leadership computing to assist discovery of novel battery electrolytes. The overall goal is to enable rational design of superior electrolytes for high voltage batteries. This study will focus on nontraditional electrolyte discovery by diluents. The influence of structural perturbation at electrolyte molecules, such as fluoridation of cation ring, on nanostructural organization at electrolyte/electrode interface as well as the transport properties and desolvation/solvation kinetics of charge carriers will be examined with advanced computational methodologies, focused on exploring an optimal structure perturbation (synthesis) path to improve electrolyte performance in lithium-ion transport. The high-throughput capability will allow use of characterization approaches from simulation studies to link solution correlations with influences on lithium ion-transport behavior in electrolytes and enable the ability to seek multiscale structural attributes that allow facile and selective incorporation of the charge carrier while prohibiting the dissolution of cathodic transition-metal components.

This project is focused solely on computational methodologies that benefit from using exascale supercomputers, decreasing time to solution from months to days. Molecular dynamics methodologies such as sampling enhanced free energy calculations are ideal both for the research problems described here and the computer resources available for the allocation. Overall, this research will enable enhanced, fundamental understanding of how the charge carriers transport in hierarchical structuring of electrolytes and how simulation knowledge can be transferred to chemical synthesis and industrial environments. Advances enabled by this work will aid in the development of the US battery industry.



Title: **Integrated and Detailed Simulation of Combustor and Turbine Interaction in a Jet Engine.**

Principal Investigator: Dheeraj Kapilavai (GE Aerospace Research)

Co-investigators: Michal Osusky (GE Aerospace Research),
Eduardo Jourdan (GE Aerospace Research)

ALCC Allocation:

Site(s): Argonne Leadership Computing Facility (ALCF)

Allocation(s): 425,000 on Aurora

Research Summary:

The project aims to resolve the multi-scale spatio-temporal physics occurring in the continuous and integrated combustor and turbine path in a gas turbine engine. Understanding the physics plays a key role in bringing the overall efficiency of a gas turbine closer to the thermodynamic limits and in reducing emissions.

The scale of computationally resolving such physics and accommodating the dimensions of these components can only be realized through leadership class compute facilities. Aurora's compute capability when combined with advanced computational fluid dynamics (CFD) software will enable resolution of the detailed turbulent combustion physics. The insights from this project will enable enhanced efficiency and durability of future products. In gas turbine development, individual components are analyzed during development cycle but an ability to simulate two critical components of the hot gas path module will lay the foundation for more challenging multi-component simulations. The research proposed here would expand U.S. competitiveness in aerospace and computational capabilities by further pushing the application of computational techniques and leadership-class compute technologies to gas turbine systems. To overcome the complexity of modeling gas turbine, the project will combine modeling techniques as well as contrast combustion modeling strategies to understand trades between accuracy and simulation speed.



Title: Privacy-Preserving Federated Learning for Foundation Models

Principal Investigator: Kibaek Kim (Argonne National Laboratory)

Co-investigators: Thomas Flynn (Brookhaven National Laboratory),
Minseok Ryu (Arizona State University),
Olivera Kotevska (Oak Ridge National Laboratory),
Farzad Yousefian (Rutgers University - New Brunswick)

ALCC Allocation:
Site(s): Oak Ridge Leadership Computing Facility (OLCF),
Argonne Leadership Computing Facility (ALCF),
National Energy Research Scientific Computing Center (NERSC)

Allocation(s): 614,400 on Frontier,
307,200 on Aurora,
102,400 on Polaris,
102,400 on Perlmutter-GPU

Research Summary:

This project advances privacy-preserving federated learning (PPFL) to enable the training of large-scale foundation models (FMs) on sensitive, multimodal scientific data distributed across institutions. By leveraging the Department of Energy's (DOE) high-performance computing (HPC) facilities—including Frontier, Aurora, Polaris, and Perlmutter—the research team will train FMs in four key areas: extracting knowledge from scientific text, interpreting high-resolution imaging data from DOE light sources, forecasting building energy consumption using national building datasets, and modeling electric grid operations through graph-based learning. These models will be developed without centralizing data, preserving privacy while enabling collaborative AI development across national laboratories and universities.

The project supports DOE's mission by delivering AI capabilities that enhance energy resilience, scientific discovery, and secure collaboration. The PPFL framework will integrate scalable optimization and privacy-preserving mechanisms—such as adaptive compression, federated pruning, and differential privacy—to reduce communication costs and safeguard sensitive data. The outcomes will establish a foundation for collaboratively training foundation models on multiple DOE's exascale systems, contributing to breakthroughs in imaging science, biomedical science, and grid modernization, while setting a precedent for secure, multi-institutional research at leadership computing scale.



Title: **Pathfinding Integrated and Automatic Experimental Analyses for DIII-D Research**

Principal Investigator: Mark Kostuk (General Atomics)

Co-investigators: Sterling Smith (General Atomics),
Nick Tyler (National Energy Research Scientific Computing Center),
Christine Simpson (Argonne Leadership Computing Facility),
Severin Denk (General Atomics),
Torrin Bechtel Amara (General Atomics),
Akshay Deshpande (General Atomics)

ALCC Allocation:

Site(s): Argonne Leadership Computing Facility (ALCF),
National Energy Research Scientific Computing Center (NERSC)

Allocation(s): 15,000 on Polaris,
33,000 on Perlmutter-CPU,
2,000 on Perlmutter-GPU

Research Summary:

The DIII-D National Fusion Facility provides hundreds of scientists worldwide as well as the burgeoning private fusion energy industry with reliable access to experimental runtime on a research tokamak. Recently selected as a pathfinder for the US Department of Energy's Integrated Research Infrastructure (IRI), DIII-D is pioneering new ways of data analysis and simulation.

This ALCC award supports all researchers at DIII-D with automated, resilient data analysis and simulation at remote compute facilities like ALCF and NERSC which is then returned to the DIII-D control room fast enough in order to make actionable decisions for the on-going experiment. Specifically, both the plasma kinetic equilibria and the energetic particle heat-load deposition are calculated remotely, on-demand. Kinetic equilibria provide a very accurate self-consistent representation of the plasma state. Simulating the particle heat load deposition on the walls is crucial for protecting diagnostics, eliminating hot spots, and reducing impurities from carbon wall ablation. This data is used immediately to better inform the course of the experiment, for improved follow-on analyses that build upon this higher-fidelity foundation, as well as for populating historical databases that can be used for surrogate model creation by other researchers.

This effort is the first time that ALCC compute resources are being used in support of IRI workflow patterns for the DIII-D facility, and is directly aligned with the DOE-Advanced Scientific Computing Research's and Fusion Energy Science's mission to support fusion energy research and to continue building an Integrated Research Infrastructure that multiplies the scientific impact coming from the DIII-D User Facility.



Title: Causal Online Alignment for Reliable Foundation Models

Principal Investigator: Emmanouil Koukoumidis (Oumi PBC)

Co-investigators: Gokhan Tur (University of Illinois Urbana-Champaign),
Yulia Tsvetkov (University of Washington),
Georgia Gkioxari (Caltech),
Ruslan Salakhutdinov (Carnegie Mellon University)

ALCC Allocation:

Site(s): Argonne Leadership Computing Facility (ALCF),
National Energy Research Scientific Computing Center (NERSC)

Allocation(s): 100,000 on Polaris,
147,500 on Perlmutter-GPU

Research Summary:

This research addresses a critical challenge in DOE's mission to advance scientific discovery and intelligent automation with AI: developing multimodal foundation models that are not just powerful, but fundamentally reliable, i.e. models that can (a) accurately assess their own capabilities and limitations, providing explicit uncertainty estimation for their actions, (b) can reason about the real-world implications of their actions, particularly in scientific and mission-critical contexts, and finally (c) be transparent and able to clearly communicate their assumptions, limitations, and potential failure modes while being able to propose and leverage mitigation strategies such as human feedback, and external tools. In essence, this research will prevent foundation models from naively attempting tasks beyond their capabilities, especially when failure could have adverse consequences. It introduces a novel training framework that emphasizes cause-effect reasoning and transparent decision-making. The approach leverages large-scale AI feedback and high-performance computing infrastructure to develop foundation models that are introspective, transparent and ultimately more reliable.

The scientific and economic impact of this work is far-reaching as the arguably biggest barrier to the advancement of the sciences with AI and the deployment of AI models in most industry applications is not their power, potential, or cost, but their unreliability. The successful completion of this research will advance the fundamental understanding of foundation model reliability and transparency, while enabling the reliable deployment of AI in critical applications. It will greatly push the frontier through reliable AI systems in both scientific and industrial applications supporting DOE's mission of scientific leadership. The code, methodology, and research artifacts will be fully open, contributing to democratizing the use of HPC computing in AI research.



Title: **High-Fidelity Simulations of Helium-Air Mixing in HTGR Cavities 2.0**

Principal Investigator: Taehun Lee (The City College of New York)

Co-investigators: Saumil Patel (Argonne National Laboratory),
Masahiro Kawaji (The City College of New York)

ALCC Allocation:
Site(s): Argonne Leadership Computing Facility (ALCF)

Allocation(s): 70,000 on Polaris

Research Summary:

Advanced reactors such as High Temperature Gas Reactors (HTGR), are being developed by U.S. companies for deployment in the late 2020s or early 2030s. Their development requires extensive computational work involving accurate fluid dynamic simulations under normal operation and accident scenarios. One of the critical safety design tests for HTGRs is the ability to dissipate decay heat safely during Depressurized Conduction Cooling (DCC) conditions. In one accident scenario of HTGR operation, hot helium gas will be discharged from a high-pressure pipe into a cavity surrounding a Reactor Pressure Vessel (RPV) or steam generator. Hot helium jets will mix with the cold air in the cavity and this gas mixture may be vented through different locations in the reactor. The amount of air entering the reactor is significantly governed by helium-air mixing. Transmission of the gas mixture is a major concern as oxygen migration leads to the risk of exothermic (combustible) reactions. Therefore assessing the concentration levels of the species (helium, oxygen, nitrogen, etc.) become important to understanding the impacts of the accidental scenario.

The simulations in this project are a continuation of the work in previous ALCC 2024/25 project which used Nek5000 and NekRS for simulating helium-air mixing in HTGR cavities. This project will utilize the Polaris supercomputer to gather high resolution velocity, temperature and gas-concentration data and will provide a critical dataset to compare against the experiments already performed or currently underway at CCNY. The objective will be to build the dataset by including two additional design components to the simulations: (1) multiple cavities connected via vents and (2) variable location of the helium-jet entrance. Simulations that include both of these design features will help augment the dataset to help validate codes and improve our understanding of the heat transfer processes in HTGRs.



Title: Plasma Turbulence and Particle Acceleration in 3D Magnetic Reconnection

Principal Investigator: Xiaocan Li (Los Alamos National Laboratory)

Co-investigators: Fan Guo (Los Alamos National Laboratory)
Adam Stanier (Los Alamos National Laboratory)

ALCC Allocation:
Site(s): Oak Ridge Leadership Computing Facility (OLCF)

Allocation(s): 750,000 on Frontier

Research Summary:

Magnetic reconnection is a fundamental plasma process that restructures magnetic fields, converting magnetic energy into particle kinetic energy while accelerating charged particles to high energies. This multiscale phenomenon plays a central role in explosive events observed in laboratory plasmas, Earth's magnetosphere, the solar atmosphere, and astrophysical systems. However, accurately modeling reconnection remains a major challenge due to the vast separation between global system sizes and microscopic kinetic scales. This project advances understanding of reconnection by conducting three-dimensional, large-scale kinetic simulations using the Los Alamos National Laboratory Vector Particle-in-Cell (VPIC) code and its hybrid version, Hybrid-VPIC—both optimized for GPU supercomputers such as Frontier. By systematically varying plasma conditions and numerical setups, the research will reveal how reconnection-driven plasma turbulence develops across scales and how it governs the acceleration and transport of energetic particles.

The outcomes will address key open questions about how nature's most powerful particle accelerators operate, with implications for interpreting laboratory experiments, space missions such as the NASA Magnetospheric Multiscale (MMS) mission, and astrophysical studies. This work directly supports the U.S. Department of Energy Office of Science mission to expand fundamental knowledge in plasma physics and strengthens the scientific output of its two Collaborative User Facilities focused on reconnection research: the Wisconsin Plasma Physics Laboratory and the Facility for Laboratory Reconnection Experiments at Princeton.



Title: Gravitational Form Factors of the Nucleon and Pion

Principal Investigator: Keh-Fei Liu (University of Kentucky)

Co-investigators: Terrence Draper (University of Kentucky),
Frank Lee (The George Washington University),
Andrei Alexandru (The George Washington University)

ALCC Allocation:
Site(s): National Energy Research Scientific Computing Center (NERSC)

Allocation(s):
284,000 on Perlmutter-CPU,
193,000 on Perlmutter-GPU

Research Summary:

Where does the proton mass come from? Where does the proton spin come from? These are the primary goals to be addressed by experiments to be performed on the upcoming electron-ion collider (EIC). We will use the Monte Carlo calculation of lattice quantum chromodynamics (QCD) to address these questions. This entails calculations of the trace of the energy momentum tensor in the nucleon with chiral fermions. The quantum effect which leads to a trace anomaly is known to be responsible to the majority of the nucleon mass. We would like to calculate its form factor in order to obtain its spatial mass distribution from the Fourier transform. Similarly, we will calculate the topological charge distribution in the nucleon which is related to the quark spin contribution to the nucleon spin.



Title: Non-Equilibrium Rough-Wall Turbulence Database for Physics and Model Development

Principal Investigator: Adrian Lozano Duran (California Institute of Technology)

Co-investigators: Rong Ma (Massachusetts Institute of Technology)

ALCC Allocation:

Site(s): Oak Ridge Leadership Computing Facility (OLCF)

Allocation(s): 820,000 on Frontier

Research Summary:

In many real-world situations, fast-moving fluids—such as air or water—flow over rough surfaces. These flows are often affected by sudden changes in pressure or speed. Although such "non-equilibrium" conditions are common in engineering applications, they are still not well understood compared to simpler, more stable scenarios. This project aims to develop a publicly accessible database of turbulent flows over rough surfaces under non-equilibrium conditions, generated using high-fidelity simulations that fully resolve both the complex surface geometry and all relevant spatial and temporal scales of the flow. The resulting data will advance scientific understanding of flow over rough surfaces and support the development of machine learning–based modeling tools for predictive simulations.

Turbulent flows over rough surfaces are critical in many energy-related technologies, including ships, aircraft, pipelines, wind turbines, and gas turbines. Roughness caused by icing, erosion, or biofouling increases drag and reduces efficiency. Approximately one-quarter of the energy consumed by industry and commerce is lost to such effects. By improving our understanding of how surface roughness impacts turbulent flow, this project will help increase energy efficiency, reduce emissions, support renewable energy technologies, and improve predictive tools—contributing directly to the U.S. Department of Energy’s mission of promoting clean, efficient, and reliable energy systems.



Title: Scalable, Trustworthy, Energy-efficient Training of Graph Foundation Models for Material Design

Principal Investigator: Massimiliano Lupo Pasini (Oak Ridge National Laboratory)

Co-investigators: Kshitij Mehta (Oak Ridge National Laboratory),
Jong Youl Choi (Oak Ridge National Laboratory),
Pei Zhang (Oak Ridge National Laboratory),
Zach Fox (Oak Ridge National Laboratory),
Arindam Chowdhury (Oak Ridge National Laboratory)

ALCC Allocation:

Site(s): Oak Ridge Leadership Computing Facility (OLCF)

Allocation(s): 200,000 on Frontier

Research Summary:

This project aims to develop a scalable, trustworthy, and energy-efficient generative graph foundation model (GFM) tailored for atomistic materials design. By leveraging over 279 million atomistic structures from eight open-source datasets, the team will pre-train HydraGNN (<https://github.com/ORNL/HydraGNN>), a graph neural network architecture developed at ORNL, on chemically and structurally diverse compounds spanning much of the periodic table. Key innovations include the use of multi-task learning (MTL) to stabilize training on imbalanced, multi-source, multi-fidelity datasets, and the integration of scalable hyperparameter optimization (HPO) and energy telemetry tools to identify model architectures that balance accuracy with computational efficiency. Trustworthiness will be ensured via ensemble uncertainty quantification (UQ), which will assess the confidence and physical realism of generated atomic structures.

The resulting generative GFMs will be released open-source and integrated into a broader framework that combines generative and predictive models to propose, evaluate, and refine new material candidates. The predictive GFM ensemble (trained in prior work) will act as a grammar-checker for generated structures by assessing chemical and dynamical stability using surrogate models of formation energy and atomic forces. This unified AI framework will drastically reduce the data and energy needed to achieve accurate materials modeling, advancing U.S. Department of Energy (DOE) goals in materials discovery. The project will execute on DOE's OLCF-Frontier supercomputer using 200,000 node-hours, exploiting its graphics processing unit (GPU) infrastructure for large-scale training and model evaluation.



Title: **Compute for: SciGPT: Scalable Foundation Model for Scientific Machine Learning**

Principal Investigator: Michael Mahoney (Lawrence Berkeley National Laboratory)

Co-investigators: Prasanna Balaprakash (ORNL)

ALCC Allocation:

Site(s): Oak Ridge Leadership Computing Facility (OLCF),
National Energy Research Scientific Computing Center (NERSC)

Allocation(s): 376,000 on Frontier,
300,000 on Perlmutter-GPU

Research Summary:

Recent advances in machine learning have produced “foundation models”. These are large, general-purpose models trained on vast datasets using self-supervision and can be adapted to a wide range of tasks. While best known in natural language processing and computer vision, these models are beginning to impact scientific domains through the emerging class of scientific foundation models (SFMs). SFMs aim to learn from diverse scientific data—observational, experimental, and simulation-based—and support broad generalization across domains. The goal of this project is to develop **SciGPT**, a blueprint SFM that lays the groundwork for such cross-domain capabilities. Unlike existing scientific machine learning (SciML) models that are typically tailored to specific tasks within a single domain, SciGPT will be trained on a variety of scientific data types and scales. It will confront three central barriers to building truly general SFMs: (1) the absence of neural scaling laws that span model size, data complexity, and spatio-temporal resolution, (2) limited generalization to out-of-distribution inputs, and (3) the scarcity of unified, multimodal scientific datasets for training.

By directly addressing these limitations, SciGPT will provide a prototype demonstrating how a model can support transfer learning across scientific fields, adapt to new problems with minimal data, and reflect the complex interdependencies seen in real-world scientific systems. This effort will not only produce a functional model but also deliver architectural, methodological, and performance insights that clarify how to scale SFMs effectively. The work is aligned with the United States Department of Energy (DOE) mission to advance scientific discovery and innovation and will leverage the DOE’s high-performance computing infrastructure and rich data ecosystems, supporting long-term goals in several scientific domain areas.



Title: **A Codebook Language and Digital Twin Framework for Diffraction Data Analysis and Accelerator Operations**

Principal Investigator: Valerio Mariani (SLAC National Accelerator Laboratory)

Co-investigators: Jana Thayer (SLAC National Accelerator Laboratory),
David Rogers (Oak Ridge National Lab)

ALCC Allocation:
Site(s): Oak Ridge Leadership Computing Facility (OLCF),
National Energy Research Scientific Computing Center (NERSC)

Allocation(s): 150,000 on Frontier,
1,250 on Perlmutter-GPU

Research Summary:

The Linac Coherent Light Source (LCLS) at SLAC National Accelerator Laboratory generates unprecedented volumes of X-ray diffraction data through dozens of experimental techniques including serial femtosecond crystallography, coherent diffractive imaging, and single particle imaging. This project aims to create a universal visual language for X-ray diffraction data by training large-scale vision models that can represent any LCLS diffraction pattern through learned visual vocabularies. We will explore vector-quantized autoencoder models that create discrete visual codebooks, masked autoregressive models that can generate diffraction patterns, and hierarchical vision transformers that learn multi-scale representations from individual photon hits to complete datasets. These models will be trained on LCLS's extensive archives, creating robust representations that handle X-ray Free Electron Laser data and capture essential physical features across different experimental modalities. The project also includes digital twins for SLAC's particle accelerators, combining physics-based simulations with machine learning for real-time operational optimization.

This research will provide transformative tools for the broader light source community, enabling automated analysis of diffraction data at the unprecedented scales required by next-generation X-ray facilities. The universal visual language will accelerate scientific discovery by automatically identifying high-quality diffraction patterns and providing interpretable insights into measurement quality. These capabilities directly advance the Department of Energy (DOE)'s mission in materials science through rapid screening of quantum materials, accelerated structural biology research, and enhanced characterization of advanced materials and catalytic systems. The digital twin framework for accelerator operations will improve facility uptime and beam quality, maximizing scientific productivity across all user experiments. The computational demands of training these large vision models on multi-terabyte datasets, combined with the need for real-time data streaming and model training/distillation, require massively parallel Graphics Processing Unit (GPU) resources beyond SLAC's current capabilities. This ALCC allocation will utilize Oak Ridge National Laboratory (ORNL)'s Frontier and National Energy Research Scientific Computing Center (NERSC)'s Perlmutter systems to support both the intensive model development phase and the ongoing operational requirements for live data analysis and accelerator optimization.



Title: Multi-scale Modeling of Ion Beams from Hybrid Laser-plasma to RF Accelerator

Principal Investigator: Alexis Marret (SLAC National Accelerator Laboratory)

Co-investigators: Siegfried Glenzer (SLAC National Accelerator Laboratory),
Frederico Fiuza (Instituto Superior Tecnico)

ALCC Allocation:
Site(s): National Energy Research Scientific Computing Center (NERSC)

Allocation(s): 115,000 on Perlmutter-CPU

Research Summary:

The generation of well-controlled high-energy ion beams in a compact system is important for applications that range from radiography of dense plasmas to tumor therapy. Ion acceleration from laser-plasma interactions has attracted significant attention due to the possibility to exploit acceleration gradients in plasma that can be > 1000 times larger than those of conventional solid-state technology, enabling acceleration on very short distances compared to conventional accelerator. Laser-plasma experiments have shown the generation of high-charge, low-emittance ion beams with high energy but mostly with a broad energy, low-quality spectrum. On the other hand, accelerators based on radiofrequency (RF) cavities can achieve higher energies and precise control of the ion spectrum, albeit at the cost of significantly smaller acceleration gradients, and thus at much larger size and cost. The objective of this research is to perform for the first time multi-scale simulations of the coupling between a laser-plasma ion accelerator and a RF accelerator that could enable the generation of high-quality and high-energy ion beams in a compact system. Such a system would transform our ability to produce ion beams for a variety of applications that range from fusion plasmas, to material science, and medical therapy.

We will use the state-of-the-art particle-in-cell code OSIRIS to model ion acceleration for the full hybrid accelerator system, from the interaction of an intense laser with a solid density target (ion injection, micron scales) to the meter-scale transport and acceleration of the ion beams in the RF cavities. The full range of kinetic physics and space-charge effects will be self-consistently taken into account for the first time for the whole system. In order to achieve that we propose to use an adaptive mesh technique where the simulation resolution and time step is changed at different stages as a function of the ion beam propagation distance, which significantly reduces the computational workload. By varying the laser, plasma, and RF cavity parameters we aim to understand how to control space-charge effects and optimize this hybrid accelerator for the generation of high-quality, high-charge ion beams in a compact (meter-scale) system. The understanding of how to control laser-plasma interactions to produce high-quality ion beams in compact systems as well as the application of such beams are important to different areas of DOE, including the FES High-Energy-Density-Science, HEP, and NNSA programs. The results of this research are expected to have a significant impact on these programs by potentially uncovering a new route for the generation of high-quality and high-energy ion beams in compact systems and by developing the capability to do multi-scale end-to-end simulations of such systems.



Title: Multiphase Mixing Induced by Interface Breakup

Principal Investigator: Ryan McMullen (Sandia National Laboratories)

Co-investigators: Spencer Bryngelson (Georgia Institute of Technology)

ALCC Allocation:

Site(s): Oak Ridge Leadership Computing Facility (OLCF)

Allocation(s): 495,000 on Frontier

Research Summary:

Multiphase fluid flow is ubiquitous in both natural and engineering systems and has thus received considerable attention. While multiphase systems are often initially well-separated by a clearly delineated interface, less attention has been devoted to understanding how these interfaces break up and lead to the intermingling of the phases. This project aims to fill this gap in understanding by performing high-fidelity numerical simulations of multiphase interface breakup, focusing on the paradigmatic example of a liquid-gas interface breaking up due to strong vibration (having amplitude many times Earth's gravitational acceleration). Predicting the behavior of these vibrated multiphase fluid systems is challenging because of the rich variety of physical phenomena they exhibit: liquid ejection and droplet pinch-off, droplet impact and gas entrainment, bubble polydispersity, nonstationary multiphase turbulence, and rectified bubble motion induced by compressibility-vibration coupling. Moreover, there is currently a shortage of experimental data for these kinds of flows. Consequently, high-fidelity simulations are an essential tool for advancing our understanding of these systems. This project's objective is to understand the dominant physical mechanisms at play during the interface breakup process in multiphase fluid systems and how this process influences the long-term behavior of vibrated multiphase flows.

In support of the Goal 2, Strategic Objective 5 of the DOE Strategic Plan, Sandia National Laboratories is interested in predicting the behavior of multiphase fluid-structural systems subjected to vibration. Therefore, by achieving this objective, this project will advance basic scientific understanding for an important class of multiphase flow problems, support the DOE's mission, and provide a first-principles data set that would be instrumental in future efforts to develop reduced-order models of these flows.



Title: **Advanced Modeling for Two-Dimensional Materials and Transition Metal Oxides**

Principal Investigator: Daniel Mejia Rodriguez (Pacific Northwest National Laboratory)

Co-investigators: Ajay Panyala (Pacific Northwest National Laboratory)

ALCC Allocation:

Site(s): Oak Ridge Leadership Computing Facility (OLCF),
Argonne Leadership Computing Facility (ALCF)

Allocation(s): 275,000 on Frontier,
275,000 on Aurora

Research Summary:

Arguably, one of the more exciting developments in condensed-matter physics in the past few years is the discovery of moiré materials. A simple twist between two stacked sheets of graphene layers leads to superconductivity and other strongly correlated phenomena not seen in a single graphene layer. The twist angle induces a geometric interference pattern between the individual lattices known as the moiré pattern. The moiré pattern creates a periodic triangular superlattice which, for small twist angles, forms a unit cell with a typical lattice constant between tens of nanometers to hundreds of nanometers, orders of magnitude larger than that of single-layer graphene. It is also possible to create moiré patterns without any twisting of the layers. One such way is by stacking and aligning layers of different materials, like transition metal dichalcogenides (TMDs), which produces a moiré lattice because of the mismatch between the different crystal species. Remarkably, these moiré materials are highly tunable, with many parameters such as carrier densities, electronic band, and even band topology, controllable by experimental knobs, including external electrical field, relative interlayer twist angles, and dielectric environments. This level of control, rather unique to moiré materials, opens up exciting possibilities for the exploration of the interplay between correlation, delocalization, and topology and the potential discovery of emergent electronic phases arising from these interactions. The computational modeling and exploration of new moiré materials and how to tune their properties is rather challenging. This stems from the fact that there exist numerous combinations of twisting angles and two-dimensional materials that can be stacked together, and the immense moiré lattice constant resulting from such stacking.

We will use ExaChem, an exascale-ready computational chemistry tool, to obtain coupled-cluster-level corrections for the generation and fine-tuning of machine-learned interatomic potentials tailored for a diverse set of moiré materials. These interatomic potentials will be used to establish the mapping from moiré structures to moiré topology and, subsequently, to generate the “moiré kaleidoscope,” an open-access database of moiré materials part of the Navigating the Design Space of Heterostructures (HeteroFAM) project funded by the Computational Materials Science Program of the US Department of Energy (DOE).



Title: Training Multi-modal Models for HPC Code and Data

Principal Investigator: Harshitha Menon (Lawrence Livermore National Laboratory)

Co-investigators: Abhinav Bhatele (University of Maryland),
Tom Goldstein (University of Maryland),
Jonas Geiping (ELLIS Institute Tübingen),
William Godoy (Oak Ridge National Laboratory),
Arjun Guha (Northeastern University),
David Bau (Northeastern University)

ALCC Allocation:

Site(s): Oak Ridge Leadership Computing Facility (OLCF),
Argonne Leadership Computing Facility (ALCF),
National Energy Research Scientific Computing Center (NERSC)

Allocation(s): 300,000 on Frontier,
200,000 on Aurora,
200,000 on Perlmutter-GPU

Research Summary:

Large Language Models (LLMs) have become an important part of the toolchain for software development, but existing LLMs are not designed to handle the specialized tasks in High-Performance Computing (HPC), such as parallel code generation, performance optimization, and managing hardware heterogeneity. This project will train next-generation LLMs on Department of Energy (DOE) supercomputers, Aurora, Frontier, and Perlmutter, to support HPC software at scale. The models will be designed to handle very long code contexts and incorporate multiple sources of information, including performance traces and documentation. Beyond parallel code generation, the models will be trained to reason about performance and correctness of generated parallel code. Additionally, new attribution techniques will be developed to better understand how specific training examples influence a model's output, improving transparency and trustworthiness. Although broadly applicable, this work will focus on DOE's Extreme-scale Scientific Software Stack (E4S), a flagship software ecosystem developed under the Exascale Computing Project (ECP), to demonstrate how LLMs can accelerate software development across complex and heterogeneous systems.

The resulting multi-modal, performance-aware, explainable LLMs are expected to revolutionize scientific software development and boost HPC developer productivity by significantly reducing the manual effort required for porting, tuning, and maintaining software across existing and emerging hardware platforms. This work will enhance the long-term sustainability and usability of the DOE's software ecosystem, accelerate adoption of E4S across the scientific community, and advance Artificial Intelligence (AI) for HPC.



Title: **Toward Multi-scale Bridging for Thermal-hydraulic Transients:
High-Fidelity Simulations**

Principal Investigator: Elia Merzari (Penn State University)

Co-investigators: Igor Bolotnov (NCSU),
Misun Min (Argonne National Lab),
Paul Fischer (UIUC)

ALCC Allocation:

Site(s): Oak Ridge Leadership Computing Facility (OLCF)

Allocation(s): 250,000 on Frontier

Research Summary:

The U.S. Department of Energy (DOE) Office of Nuclear Energy has launched a consortium under the Integrated Research Projects (IRP) program to advance multiscale thermal-fluids modeling methods. This effort forms the university pillar of the newly established Center of Excellence for Thermal-Fluids Applications in Nuclear Energy, integrated with national labs and engaging key stakeholders, including industry, DOE programs, and the Nuclear Regulatory Commission. The goal is to develop rapid, high-accuracy computational models that capture complex phenomena in advanced nuclear reactors.

Current thermal-fluids modeling faces two key challenges: (1) a lack of integral-effect test data applicable to innovative reactor designs, and (2) underutilization of high-fidelity simulations and separate-effect experiments. Addressing these requires multiscale bridging methods that transform high-fidelity data into predictive, fast-running models.

Building on the initial phase, Phase II began in 2024 with a focus on transient phenomena, critical for understanding safety in advanced reactors. This phase targets scale-bridging techniques for complex, time-dependent thermal-fluid behavior, emphasizing three challenge areas: (i) Flexible Heat Transfer Modeling, (ii) Mixing in Large Enclosures, and (iii) Multiscale Core Modeling. The work spans numerical method development, high-fidelity simulations, and targeted experiments to generate essential validation data.

The project will produce a robust simulation database to support data-driven modeling approaches, fostering industry adoption of advanced modeling tools and accelerating deployment of next-generation reactor technologies.



Title: Predicting Flow Distortions in Serpentine Engine Inlets

Principal Investigator: Parviz Moin (Stanford University)

Co-investigators: Sanjeeb Bose (Cadence Design Systems, Stanford University)

ALCC Allocation:

Site(s): Oak Ridge Leadership Computing Facility (OLCF),
Argonne Leadership Computing Facility (ALCF)

Allocation(s): 1,600,000 on Frontier,
1,700,000 on Aurora

Research Summary:

Blended wing body configurations are emerging as a promising design for future aircraft, offering improved aerodynamic efficiency and a compact airframe that enhances the overall thrust-to-weight ratio. By integrating the engines within the fuselage, these designs also reduce radar cross-section and fuel consumption. However, such compact layouts require complex inlet systems, like serpentine ducts, which can ingest the large boundary layer that develops over the aircraft body. A major challenge with serpentine ducts is flow separation, induced by the high-curvature walls at duct bends. This can lead to severe consequences, including flow distortion, increased stress on turbine blades, reduced stall margin, and in extreme cases, engine failure. To this end, the present simulation campaign will perform direct numerical simulations (DNS) of a serpentine inlet (SD-2, from Burrows et al., PhD. Thesis, Georgia Tech, 2020). This campaign consists of two parts. First, we will simulate the SD-2 diffuser at three set points (Mach number/mass flow rate) to predict quantitative and qualitative trends in flow separation and distortion with incoming flow. As the time-varying flow distortion is of interest in designing these new engines, our simulations will also develop probability distribution maps of flow recovery and distortions as a function of Mach number for serpentine inlets. These maps will be compared with the experimental data. Second, a simulation campaign performing active flow control, with realistic zero net mass flux actuators (developed at Georgia Tech) will be pursued. The inclusion of zero net mass flux actuators was able to control the onset of flow separation in the experiments of Burrows et al., PhD. Thesis, Georgia Tech, 2020. These will help determine whether simulations can accurately predict the augmentation of the flow structures in the presence of active control strategies.



Title: **EMERGE: ExaEpi Calibration Runs and Surrogate Models**

Principal Investigator: Peter Nugent (Lawrence Berkeley National Laboratory)

ALCC Allocation:

Site(s): Oak Ridge Leadership Computing Facility (OLCF),
Argonne Leadership Computing Facility (ALCF),
National Energy Research Scientific Computing Center (NERSC)

Allocation(s): 425,000 on Frontier,
425,000 on Aurora,
100,000 on Perlmutter-GPU

Research Summary:

The EMERGE project (ExaEpi for Elucidating Multiscale Ecosystem Complexities for Robust, Generalized Epidemiology) is building a next-generation agent-based model (ABM) that captures how diseases spread both directly between people and through environmental pathways such as air, insects, and water. An ABM is like a computerized “mini-world” made up of many individual characters—called agents—that each follow a small set of rules. Every agent might represent a person, an animal, a car, or even a cell. The simulation lets these agents move around, bump into one another, make decisions, and change over time. By watching millions of simple agents interact with themselves and their environment, researchers can see how large-scale patterns—such as a traffic jam forming, a disease spreading, or an economy growing—emerge naturally. Thus an ABM starts with simple rules for individuals and ends up revealing complex behavior for the whole system. ABMs have been used to study a wide variety of phenomena spanning several different communicable human diseases, to business models involving production, selling and consuming, to even some nascent work on cancer and diabetes. However, their use for forecasting and control has been limited due to difficulties in calibrating them to the multitude of data streams available during an outbreak and quantifying the uncertainties of the model. Over the coming year, large ensembles of simulations using the exascale-capable ABM code ExaEpi will be carried out on DOE ASCR compute facilities to pinpoint the most influential variables for the COVID-19 outbreak. The refined set will train fast AI surrogate models that can be quickly calibrated to data such as hospitalizations, deaths, testing results, wastewater, etc. Applying the AI technique called reinforcement learning, EMERGE will create a foundation for a real-time decision system that tests a variety of intervention policies and their uncertainties.

By fusing massive computing power, diverse data streams, and AI, EMERGE is expected to deliver rapid forecasts with quantified uncertainties, giving public-health officials a clearer picture of how policy choices may play out days to years in the future for a variety of known and potential disease outbreaks. The work advances DOE’s mission by pushing the limits of exascale computing, data integration, and AI, while also strengthening national preparedness for emerging biological threats.



Title: **High Energy Density Physics of Novel Inertial Fusion Energy Ablator Materials**

Principal Investigator: Ivan Oleynik (University of South Florida)

Co-investigators: Mitchell Wood (Sandia National Laboratories),
Stan Moore (Sandia National Laboratories),
Jon Eggert (Lawrence Livermore National Laboratory),
Marius Millot (Lawrence Livermore National Laboratory),
Sally Tracy (Carnegie Institution for Science),
Aidan Thompson (Sandia National Laboratories)

ALCC Allocation:
Site(s): Oak Ridge Leadership Computing Facility (OLCF),
Argonne Leadership Computing Facility (ALCF),
National Energy Research Scientific Computing Center (NERSC)

Allocation(s): 1,000,000 on Frontier,
800,000 on Aurora,
100,000 on Perlmutter-GPU

Research Summary:

This project aims to advance the development of next-generation amorphous ablator materials for Inertial Fusion Energy (IFE) by reducing deuterium-tritium fuel/ablator mixing and lowering the fuel adiabat, thereby enabling multi-megajoule yields. The ablator, a fundamental component of IFE, converts deposited energy into rocket-like recoil that drives ablative fuel capsule implosion involving compression, heating, and ultimately, ignition of the fusion fuel.

This simulation campaign, leveraging billion-atom molecular dynamics (MD) simulations on DOE's Frontier and Aurora exascale supercomputers, seeks to uncover the potential of amorphous boron carbide (a-B₄C) as a high-performance IFE ablator material. We will generate critical data for high-gain IFE designs using a-B₄C targets by determining its equation of state across a broad range of pressure-temperature conditions, examining phase transformations along the shock Hugoniot, constraining its melting line, and investigating the emergence of shock front roughness at varying shock strengths.

Our simulations, uniquely coupled with dynamic compression experiments at DOE's high-power laser facilities, will provide simulation guidance that enables us to achieve top-tier results while fostering an inspiring intellectual environment for the graduate students and postdocs involved in this project.



Title: Probabilistic Comparative Modeling of Colorectal Cancer Screening Strategies

Principal Investigator: Jonathan Ozik (Argonne National Laboratory)

Co-investigators: Carolyn Rutter (Fred Hutchinson Cancer Center),
Iris Lansdorp-Vogelaar (Erasmus University Medical Center),
Fernando Alarid-Escudero (Stanford University),
Thomas Trikalinos (Brown University),
Iakovos Toumazis (The University of Texas MD Anderson Cancer Center)

ALCC Allocation:
Site(s): National Energy Research Scientific Computing Center (NERSC)

Allocation(s): 135,000 on Perlmutter-CPU

Research Summary:

Despite large increases in the uptake of screening in the past two decades, colorectal cancer (CRC) is still the second leading cause of cancer death in the US. This points to inadequate screening and treatment, gaps in care that need to be addressed. Technological advances are bringing new methods for risk-targeted screening and treatment and new screening modalities. There is a critical need to assess these potential improvements in terms of both their ability to reduce the burden of CRC and their associated costs. But it is not logistically or ethically feasible to conduct clinical trials of all possible interventions. Instead, computational models, in the form of natural history microsimulations, are used as *in silico* laboratories to evaluate the potential impact of changes in clinical practice and new policies on clinical and economic CRC outcomes. These models are based on information about underlying disease process, sensitivity and specificity of screening tests, and treatment effectiveness. There is uncertainty in both available data and the models, which describe unobservable processes. Due to these uncertainties, large-scale computation is required to provide robust evidence for effective screening approaches.

This project will use leadership class computing resources to run comparative probabilistic sensitivity analyses (PSAs) of screening strategies with three state-of-the-art CRC models. Funded under the National Cancer Institute's (NCI) Cancer Intervention and Surveillance Modeling Network (CISNET) program, these models were independently developed for the evaluation of interventions and describe CRC natural history using different underlying assumptions. Building on model calibration, comparison, and evaluation of screening efficacy that was accomplished in previous allocation periods, the project will extend analyses to extensions of the microsimulation models that incorporate the serrated pathway to colorectal cancer. The comparative PSAs in this work will be used to generate cost-effectiveness analyses for complex interventions and to provide formalized assessments of uncertainties across the three CRC models. Our prior work has also promoted the broader adoption of HPC in the application of microsimulation-based medical decision making for cancer and, as a result, this allocation expands the application of our HPC-scale approaches to two new cancers, bladder and lung.



Title: **High-fidelity Simulation of SAF End Use: Emissions and Operability Studies**

Principal Investigator: Bruce Perry (National Renewable Energy Laboratory)

Co-investigators: Jacqueline Chen (Sandia National Laboratories),
Martin Rieth (Sandia National Laboratories),
Bruno Souze Soriano (Sandia National Laboratories),
Sreejith Nadakkal Appukuttan (National Renewable Energy Laboratory),
Shashank Yellapantula (National Renewable Energy Laboratory),
Marc Day (National Renewable Energy Laboratory)

ALCC Allocation:

Site(s): Oak Ridge Leadership Computing Facility (OLCF),
Argonne Leadership Computing Facility (ALCF)

Allocation(s): 600,000 on Frontier,
300,000 on Aurora

Research Summary:

This project is a joint effort between Sandia National Laboratories (SNL) and the National Renewable Energy Laboratory (NREL) that will develop large scale computations for aviation combustors running with Synthetic Aviation Fuels (SAF). These simulations will be used to predict particulate emission formation in aircraft engines running on new fuels and investigate the performance and flame stability for these fuels in novel combustor designs, including lean premixed prevaporized combustors and trapped vortex combustors. A simulation pipeline will be established from simulations of a lab scale experiment at SNL used to validate models for the formation of particulate emissions to full combustor scale simulations in collaboration with industry and academic partners. High-performance computing capability through the ALCC program are essential in enabling simulations of configurations that are of interest to industry. The simulations will focus on Hydroprocessed Esters and Fatty Acids (HEFA), a novel fuel that must be blended with additives containing aromatic compounds to meet fuel certification requirements, to understand how these additives affect particulate formation and flame stability relative to conventional fuels.

Synthetic aviation fuels derived from biological feedstocks or electricity-driven processes provide an additional domestic source of fuel for the civil and defense aerospace industries. Therefore, developing and commercializing new synthetic fuels is well aligned with national energy priorities like increasing domestic energy production, maintaining energy security, keeping energy prices low/stable, and supporting economic growth in the agriculture sector. However, due to rigorous certification requirements to ensure safety, the cost of testing requirements can be a prohibitive risk for novel fuel production pathways. Aircraft engine combustor simulations can be used to predict performance based on properties measured with a small quantity of fuel to screen fuel candidates and de-risk scale-up, so the capabilities and new insight into new fuels and combustor designs developed in this work will help to hasten deployment of these technologies.



Title: Next Generation Turbomachinery Design Using Turbulence-resolving Large Eddy Simulation

Principal Investigator: Sylvain Pierre (GE Vernova Advanced Research)

Co-investigators: Balaji Jayaraman (GE Vernova Advanced Research)

ALCC Allocation:

Site(s): Oak Ridge Leadership Computing Facility (OLCF)

Allocation(s): 660,000 on Frontier

Research Summary:

Gas turbine designs today rely on insights from idealized low-order (quasi-2D) models and averaged turbulence formulations such as Reynolds-averaged Navier-Stokes (RANS), which are insufficient for next generation designs involving high-lift airfoil technology. Such high-lift airfoils are known to generate complex transonic flow separation dynamics that current RANS or (unsteady RANS) tools cannot predict accurately. High-fidelity (3D) modeling using large-eddy simulations (LES) of cooled airfoils at Reynolds numbers representative of heavy-duty power generation gas turbines are needed to generate insight into flow mixing, loss mechanisms and gas temperature migration in such designs. In an industry where small gains in efficiency are meaningful and increased reliability with aggressive designs are critical, the predictive accuracy of design tools matter.

Consequently, the turbomachinery community is starting to leverage high-fidelity LES for virtual testing, causality analysis, and ultimately design. This requires validating extreme fidelity wall-resolved LES models of cooled turbine airfoils under industry-relevant conditions against cascade test data to inspire confidence in their prognostic ability. Additionally, the computational barriers for industrializing LES at high Reynolds numbers (much higher than aircraft engines) must be overcome to support rapid design cycles. This project builds on prior work to transform industry design capabilities for next-generation products where higher fidelity GPU-accelerated tools can reduce reliance on physical testing, accelerate design cycles, and enable larger design space exploration with reduced uncertainty.



Title: **An Exascale Co-Design Framework for Spiking Graph Neural Networks**

Principal Investigator: Thomas Potok (Oak Ridge National Laboratory)

Co-investigators: Guojing Cong (Oak Ridge National Laboratory),
Robert Patton (Oak Ridge National Laboratory),
Ramakrishnan Kannan (Oak Ridge National Laboratory),
Chathika Gunaratne (Oak Ridge National Laboratory),
Mark Coletti (Oak Ridge National Laboratory),
Prasanna Date (Oak Ridge National Laboratory),
Ashish Gautam (Oak Ridge National Laboratory),
Seung-Hwan Lim (Oak Ridge National Laboratory),
Shruti Kulkarni (Oak Ridge National Laboratory)

Collaborators: Aditya Devarakonda (Wake Forest University)
Hao Lu (Oak Ridge National Laboratory)
Catherine (Katie) Schuman (University of Tennessee)
Gina C. Adam (George Washington University)
Maryam Parsa (George Mason University)

ALCC Allocation:

Site(s): Oak Ridge Leadership Computing Facility (OLCF)

Allocation(s): 420,000 on Frontier

Research Summary:

The Energy-efficient Novel Algorithms and Architectures for Graph Learning (ENGAGE) project is developing a new approach to building energy-efficient artificial intelligence (AI) that functions more like the human brain. The team is focused on creating next-generation graph-based AI models that can learn and make decisions using short bursts of electrical activity, or “spikes,” similar to the way real neurons communicate. These brain-inspired models are designed to consume significantly less energy than today’s deep learning systems, which, while powerful, require substantial computing resources. To design, test, and scale these models effectively, the team relies on high-performance computing (HPC) systems, which can simulate complex brain-like networks and explore a vast space of model designs far more quickly than conventional computers. By working closely with hardware designers, the project aims to construct these models in a way that makes them easier to run on specialized low-power computer chips, known as neuromorphic hardware. The initial phase of the project is focused on using high-performance computing resources to validate these ideas, with future efforts directed at adapting the models for real-world hardware.

This work aligns with the U.S. Department of Energy's (DOE) priority to co-design neuromorphic systems for energy-efficient computing. The anticipated impact includes enabling low-power, scalable AI solutions for DOE mission-critical applications such as base science, national security, autonomous vehicles, and smart infrastructure. In addition to advancing fundamental science, the project will support workforce development by engaging students through DOE-supported programs, thereby growing the next generation of AI and high-performance computing researchers.



Title: Targeting Coupled Fan Motor Assembly and Wind Tunnel Circuit Improvements

Principal Investigator: Russell Powers (Naval Air Warfare Center Aircraft Division)

Co-investigators: Craig Fernandes (Naval Air Warfare Center Aircraft Division),
Kyle Lukacovic (National Full Scale Aerodynamics Complex),
James Masters (Arnold Engineering Development Complex),
Joseph Sacco (National Aeronautics and Space Administration)

ALCC Allocation:

Site(s): Oak Ridge Leadership Computing Facility (OLCF)

Allocation(s): 380,000 on Frontier

Research Summary:

This research focuses on using transient wall-modeled large eddy simulation to predict coupled wind tunnel and wind tunnel fan motor performance. It will focus on the aerodynamic impacts due to proposed wind tunnel and fan motor modifications. The goal is to advance the state of the art in Digital Engineering by developing a massively parallel Large Eddy Simulation (LES) design framework for integrated simulations of an entire wind tunnel circuit including the fan motor assembly. This research will leverage established wall-modeled large eddy simulation tools and experience with the Fidelity CharLES flow solver from Cadence Design Systems.

The National Full-scale Aerodynamic Complex (NFAC) located at NASA Ames Research Center in Moffett Field, CA is a unique national asset used by the Department of Defense, NASA, industry and others for aerodynamic and acoustic testing. The facility consists of the world's two largest wind tunnel test sections that are both driven by a common six fan motor system. Test articles include large aircraft models, full-scale rotorcraft, and many other advanced technologies including the DoE SuperTruck program and wind turbine blades from the DoE Wind Energy Technology Office. Due to a facility mishap, the wind tunnel fan motor assembly is in a degraded state with only 12 of 15 legacy fan blades present on each of the 6 fan motors. The NFAC is currently undergoing a program for Performance Expansion and Restoration of the fan drive to replace each fan motor with 15 new blade designs. Predictions indicate new blades may improve maximum tunnel operating conditions while reducing the required tunnel power consumption for constant speed operations. Leadership class resources will improve confidence in the modeling assumptions while impacting design, risk, and program decisions for incorporation of the new fan blades at the NFAC. This research will enable improved performance and understanding of the wind tunnel operation that could impact a significant number of aircraft, helicopter, wind turbine, and clean energy transportation designs for both the commercial and defense sectors.



Title: **Highly-Resolved Simulations of Turbulence in a Transonic Fan at Flight Conditions**

Principal Investigator: Stephan Priebe (GE Aerospace Research)

Co-investigators: Daniel Wilkin II (GE Aerospace),
Arash Mousavi (GE Aerospace),
Paul Orkwis (GE Aerospace),
Luke D'Aquila (GE Aerospace Research)

ALCC Allocation:

Site(s): Oak Ridge Leadership Computing Facility (OLCF)

Allocation(s): 757,000 on Frontier

Research Summary:

There is a need to further improve the efficiency and fuel consumption associated with aircraft propulsion. The large fan blades at the front of modern turbofan engines are a key part to achieving these improvements, which are only possible with a comprehensive understanding of fan aerodynamics. Traditional computational simulations do not accurately predict the complex flow physics in transonic fans due to model shortcomings, thus requiring high-fidelity simulations to accurately capture these physics.

In this ALCC project, the team will execute high-fidelity simulations of the complex flow physics that occur in transonic fan blades at actual flight conditions, i.e., at the conditions that an engine on an aircraft would encounter in flight. The tip region of transonic blades is particularly challenging due to the presence of shock waves leading to a phenomenon known as shock wave/boundary layer interaction (SBLI), which can lead to significant aerodynamics losses and hence a reduction in efficiency and performance.

This work will investigate actual engine conditions which have a significantly richer turbulence field spanning a broader range of spatial and temporal scales than captured in prior simulations representative of wind tunnel tests. The simulations will enable physical insight into the transitional flow behavior and SBLI dynamics at flight conditions, and how these physics change from wind tunnel conditions at reduced scale to flight conditions at full scale, years in advance of when flight testing typically occurs. This will enable optimizing the design for steady and unsteady aerodynamics and aeromechanics, which is critical for guiding design of next-generation, fuel-efficient turbomachinery.

In the realm of commercial aircraft propulsion alone, we estimate that the benefit of an improved physical understanding and prediction capability of transonic fans and SBLI would be of the order of 130 million gallons of jet fuel saved for U.S. airlines every year.



Title: Kinetic Plasma Model Investigation of Z Pinch Physics at Fusion Conditions

Principal Investigator: Noah Reddell (Zap Energy Inc.)

Co-investigators: Uri Shumlak (Zap Energy Inc.)
Peter Stoltz (Zap Energy Inc.)
Daniel Crews (Zap Energy Inc.)
Eric Meier (Zap Energy Inc.)
Iman Datta (Zap Energy Inc.)
Steve Richardson (Zap Energy Inc.)

ALCC Allocation:

Site(s): Oak Ridge Leadership Computing Facility (OLCF),
Argonne Leadership Computing Facility (ALCF)

Allocation(s): 800,000 on Frontier,
200,000 on Aurora

Research Summary:

Plasma is an energetic state of matter with major implications for sustainable energy production, space travel, and foundational understanding of the universe, and yet it remains extraordinarily difficult to accurately simulate plasma behavior. Inadequate predictive modeling has been one of the major barriers to the advancement of fusion energy systems. To address this challenge, this project will advance and test state-of-the-art, high-fidelity kinetic simulation capabilities. Kinetic modeling captures the dynamics of plasmas far from local thermodynamic equilibrium. This regime is critically important for practical fusion technologies and is beyond the scope of more prevalent magnetohydrodynamic fluid descriptions. Specifically, this project will apply Vlasov kinetic modeling (an emerging alternative to the more frequently employed particle-in-cell method) to critical fusion plasma phenomena that are inherently three-dimensional. The simulation results and model performance will enhance understanding of Z-pinch fusion concepts and inform the viability of Vlasov kinetic modeling for broader applications in plasma science.

This project focuses on the study of Z pinches, a configuration for fusion plasma confinement and compression that directly drives a linear current to generate an encircling magnetic field that compresses plasma inwards. The configuration is simple and efficient, representing a potentially optimal fusion energy source, but is hindered by rapidly growing plasma instabilities. As such, phenomena of interest include instabilities, effects of compression, and electrode plasma sheath dynamics. Simulation of these phenomena is only recently possible on the most powerful supercomputers employing thousands of graphics processing units (GPUs). Improved models of Z-pinch dynamics will advance the domestic development of fusion energy and support ongoing R&D campaigns on Zap Energy's FuZE Z-pinch devices.



Title: **Enhancing APS-Enabled Research through Integrated Research Infrastructure**

Principal Investigator: Nicholas Schwarz (Argonne National Laboratory)

Co-investigators: Hannah Parraga (Argonne National Laboratory),
Ryan Chard (Argonne National Laboratory),
Thomas Uram (Argonne National Laboratory)

ALCC Allocation:
Site(s): Oak Ridge Leadership Computing Facility (OLCF),
Argonne Leadership Computing Facility (ALCF),
National Energy Research Scientific Computing Center (NERSC)

Allocation(s): 10,000 on Frontier,
150,000 on Polaris,
10,000 on Perlmutter-GPU

Research Summary:

The Advanced Photon Source (APS), located at Argonne National Laboratory, is a synchrotron light source funded by the U.S Department of Energy (DOE), Office of Science-Basic Energy Sciences (BES) to produce high-energy, high-brightness x-ray beams. As part of the facility's recent upgrade project, the APS has replaced the storage ring with a new advanced design, and is in the midst of commissioning new and enhanced beamlines offering transformative analytics for x-ray, scattering spectroscopy, and imaging. More than ever before, advanced computational approaches and technologies are essential to fully unlock the scientific potential of the APS. This project will deliver computational capabilities by leveraging Integrated Research Infrastructure (IRI) to couple APS instruments with supercomputers. Integrating this research infrastructure will be accomplished by running a set of scientific software and workflows using common Globus workflow infrastructure for reduction, reconstruction, multi-modal data utilization, and Artificial Intelligence (AI) training of massive x-ray scattering, spectroscopy, and imaging data in near real-time during experiments, and for post processing and refinement after experiments.

This project will help to enable groundbreaking photon science-enabled research at the upgraded APS, especially for the new and improved beamlines that are enabled by the facility's brighter, more intense x-rays. It will open a pathway for the critical data processing required for the facility's user community to conduct cutting-edge basic and applied research in the fields of materials science, biological and life science, physics, chemistry, environmental, geophysical, and planetary science, and innovative x-ray instrumentation. Moreover, this work will enlighten solutions around IRI by informing decisions regarding the management of resource scheduling across facilities, lead to better understanding and optimization tradeoffs between storage, network, and memory bandwidth and latency, and solidify a canonical set of workflows and workflow infrastructure for major classes of light source-enabled research.



Title: **Probing Lignocellulosic Biomass & Products via Large-Scale Molecular Simulation**

Principal Investigator: Micholas Smith (The University of Tennessee, Knoxville)

Co-investigators: Jeremy Smith (The University of Tennessee, Knoxville),
Monojoy Goswami (Oak Ridge National Laboratory)

ALCC Allocation:
Site(s): Oak Ridge Leadership Computing Facility (OLCF),
National Energy Research Scientific Computing Center (NERSC)

Allocation(s): 24,000 on Frontier,
1,000 on Perlmutter-CPU,
143,456 on Perlmutter-GPU

Research Summary:

An essential component toward the development of a robust domestic biofuels and bio-products industry requires the efficient deconstruction of plant matter into its constituent polymeric components (Fractionation). Changes to the organization of plant (secondary) cell walls undergoing deconstruction are expected to span multiple spatio-temporal scales. By making use of supercomputing resources at the Oak Ridge Leadership Computing Facility (OLCF) and National Energy Research Scientific Computing Center (NERSC), will enable large-scale molecular simulations to examine, in atomic detail, how differences in secondary cell wall composition impact secondary cell wall remodeling across multiple scales when exposed to chemical and physical stresses.

This ASCR ALCC project will develop and refine atomistic and coarse-grained models of the secondary cell wall of a Department of Energy relevant bio-energy crop, (*Populus trichocarpa*) to include interactions between biologically inorganic chemical species and three major cell wall polymers, lignin, cellulose, and hemicellulose. In addition to treating the inorganic component of the secondary cell wall, additional molecular models will also be constructed to explore how *intra-species* cell-wall variation impacts secondary cell-wall organization and reorganization under processing and fractionation relevant stresses.

An outcome of the simulations performed will be a window into how more realistic molecular scale models of plant secondary cell-walls respond to physical and chemical stresses. This will in turn yield a molecular-scaled accounting of how differences in cell wall composition modulate biomass fractionation processes. Insights drawn from these simulations may serve as to guide new fractionation technologies that effectively deconstruct biomass drawn from varied sources.



Title: Microscopic Fission Dynamics of Even-Even, Odd and Odd-Odd Nuclei

Principal Investigator: Ionel Stetcu (Los Alamos National Laboratory)

Co-investigators: Aurel Bulgac (University of Washington),
Ibrahim Abdurrahman (Los Alamos National Laboratory),
Kenneth Roche (University of Washington / AMD)

ALCC Allocation:

Site(s): Oak Ridge Leadership Computing Facility (OLCF)

Allocation(s): 528,000 on Frontier

Research Summary:

This team will employ the most advanced implementation of Density Functional Theory (DFT), extended to encompass superfluidity and time-dependent phenomena. The overarching goal is to enhance the understanding, description, and prediction of complex non-equilibrium quantum processes through cutting-edge microscopic methods and high-performance computing. This will significantly boost the predictive capability of the theoretical framework. The primary focus is to deepen our understanding of the intricate fission dynamics of odd and odd-odd nuclei and their relevance to national laboratories, the Facility for Rare Isotope Beams (FRIB), and the astrophysical origins of elements. A key additional objective is to extract insights from microscopic simulations to inform and refine models of prompt fission and gamma emission, critical components for transport codes like MCNP, developed and maintained at Los Alamos. While phenomenological models can successfully describe certain aspects of fission, their predictive range is often limited. Their reliability for exotic nuclei, such as those produced at FRIB or in astrophysical settings where empirical data are scarce, remains uncertain. A truly predictive and unified description of fission across the nuclear landscape is only achievable within a fully microscopic theoretical framework.

Building on prior work, the team will study the dynamics of a range of induced and spontaneous fission reactions—many involving even-even and odd nuclei (e.g., several U isotopes, ^{232}Th , Fm, Cf)—that have not yet been explored using time-dependent density functional theory. Theoretical results for odd and odd-odd nuclei are scarce, with only one known microscopic study. Using a fully time-dependent approach, which includes dissipation, will enable unique predictions of observables such as fragment excitation energy partition and saddle-to-scission time. The team will also examine how fission observables depend on the excitation energy of the initial compound nucleus, a currently unexplored aspect critical for nuclear data evaluations. Using a broad set of deformation parameters, time-dependent simulations will be performed to study fragment properties, including average charge and mass, kinetic energy, and excitation energy sharing. These findings will inform phenomenological models employed in nuclear data evaluations, thereby contributing to the generation of more reliable data for the scientific community.



Title: Enabling Exascale Discovery in Neutrino Science

Principal Investigator: Thomas Wester (University of Chicago)

Co-investigators: David Schmitz (University of Chicago)

ALCC Allocation:

Site(s): Argonne Leadership Computing Facility (ALCF)

Allocation(s): 150,000 on Aurora,
50,000 on Polaris

Research Summary:

Neutrinos are electrically neutral, weakly interacting fundamental particles with the surprising property that they change between different types, or “flavors”, as they travel. This flavor-changing phenomenon has been measured at different distances (baselines) over the past few decades, but several anomalous measurements at “short” baselines of about a kilometer remain in tension with theoretical predictions. If confirmed, these anomalies would imply a new fundamental particle and would radically change our understanding of particle physics. The Short-Baseline Neutrino (SBN) program, now operating at Fermilab, studies neutrinos using multiple advanced liquid argon time projection chamber (LArTPC) detectors placed in a neutrino beam and is poised to collect the world’s largest and most detailed data set of neutrino interactions to date. This data set will definitively resolve the short-baseline anomaly and simultaneously enable a rich program of neutrino-nucleus interaction research. The goal of this ALCC award is to integrate the high-performance computing (HPC) resources at the Argonne Leadership Computing Facility into the SBN research program to accelerate and enhance its scientific output. The HPC resources provided by the award will enable GPU-accelerated machine learning algorithms to be used at scale in the analysis of the many petabytes of data generated by the SBN experiments. This unique capability will drive the development of new algorithms by facilitating fast, iterative reprocessing of SBN data, and will enable the production of the large-scale simulated data sets necessary for making precision measurements.

This award will strengthen collaboration between national laboratories while showcasing the advantages of using HPC resources for a multi-stage particle physics research program. In addition to the scientific output of the SBN program, this award will result in a model for real-time data transfer between national labs, coordinated and scalable workflow execution, and multiple processing pipelines for real and simulated data sets. Further, the timely physics results and computational methods of the SBN program, enabled by this ALCC award, will inform the physics program and computational infrastructure of the next-generation flagship neutrino experiment: the Deep Underground Neutrino Experiment (DUNE).



Title: **LES&DNS Simulation on Flow and Heat Transfer Behavior in Involute Plate Research Reactor**

Principal Investigator: Yiqi Yu (Argonne National Laboratory)

Co-investigators: Aurelien Bergeron (Argonne National Laboratory),
Jeremy Licht (Argonne National Laboratory),
Cezary Bojanowski (Argonne National Laboratory)

ALCC Allocation:
Site(s): Argonne Leadership Computing Facility (ALCF),
National Energy Research Scientific Computing Center (NERSC)

Allocation(s): 150,000 on Polaris,
550,000 on Perlmutter-CPU

Research Summary:

The Materials Management and Minimization (M3) Reactor Conversion Program of the National Nuclear Security Administration (NNSA) is supporting the conversion of research reactors from Highly Enriched Uranium (HEU, $^{235}\text{U} / \text{U} \geq \text{wt. } 20\%$) fuel to Low-Enriched Uranium (LEU, $^{235}\text{U} / \text{U} < \text{wt. } 20\%$) fuel. The reactors that utilize involute-shaped fuel plates share a similar configuration of coolant channel, which is of extremely thin thickness and involute shape. Better understanding of flow behavior and heat transfer mechanisms in these coolant channels is of great interest and importance for the design of LEU fuel elements.

Reynolds-averaged Navier–Stokes (RANS) approaches have exposed both the practicability and the limitations on the prediction of the turbulent flows. Discrepancy are found between RANS simulations with different turbulence models, flow conditions and codes. Due to the lack of relevant experimental data, it is difficult to provide best practice guidelines (BPGs) for minimizing turbulence modeling uncertainties. By performing high fidelity simulations (LES & DNS), the behavior of turbulent flows can be investigated with fewer modeling assumptions, potentially providing more reliable benchmarks for engineering predictions and more accurate thermal hydraulic safety analyses. Moreover, the database generated from the high-fidelity simulations can be used to select or refine the coefficients used in the system codes. LES database enrichment is essential for understanding the fundamental mechanisms of the flow and heat transfer behavior in involute-shaped plate research reactors under various conditions.

The simulations in this project will be performed with both Nek5000(CPU) and NekRS(GPU). Both codes have won an R&D 100 Award. These codes are developed as part of the High-Order Methods for High-Performance Multi-physics Simulations project supported by the DOE Applied Math Research base program as well as a collaboration with Nuclear Energy Advanced Modeling and Simulation (NEAMS) program. This project can expand the application of Nek5000 and NekRS to Research Reactor Conversion Program and attract more attention from either industry partner or university to NEAMS code as a surrogate for expensive commercial code or separate engineering applications where experimental data is limited.



Title: **Tensor-Compressed Sustainable Pre-Training of Extreme-Scale Foundation Models**

Principal Investigator: Zheng Zhang (University of California, Santa Barbara)

Co-investigators: Franck Cappello (Argonne National Labs),
Bogdan Nicolae (Argonne National Labs),
Paul Hovland (Argonne National Labs)

ALCC Allocation:
Site(s): Argonne Leadership Computing Facility (ALCF),
National Energy Research Scientific Computing Center (NERSC)

Allocation(s): 600,000 on Aurora,
100,000 on Polaris,
80,000 on Perlmutter-GPU

Research Summary:

Large foundation models for science will face the same challenges of pre-training and inference as state-of-the-art large language models (LLMs). The time and energy needed on exascale computing systems will drastically limit the number of pre-training attempts (to only one) and capability to tune the model at this stage, leaving model corrections feasible only at post-training (alignment, fine-tuning) stage. The resource occupation needed for a single pre-training will reduce the availability of high-performance computing systems for mission-critical simulations and data analytics significantly. Thirdly, foundation models require significant resources to run inferences, which will limit their broad deployment for science. Reducing drastically the computing and memory cost of foundation models for science will have a critical impact on the feasibility, duration, and energy consumption of pre-training and inferences.

Leveraging the research team's prior research on tensor-compressed pre-training, this project will design and develop a memory- and computing-efficient pre-training framework and generate various resource-efficient foundation models based on it. To this end, the research team will explore three novel contributions: (1) theoretical foundation and novel optimization of low-rank tensor-compressed pre-training for large-scale foundation models, (2) training acceleration via mixed-precision low-rank tensor optimization and customized tensorized automatic differentiation, (3) graphic processing unit (GPU) optimization on leadership computing platforms for large-scale tensor-compressed pre-training on massive GPUs. This project aims to enable energy-efficient training and inference of extreme-scale foundation models for science, significantly reducing training time and energy cost.

Based on the Year-1 progress of the research team, this ALCC allocation will be used to support the Year-2 research plan of a 3-year research project. Main computing experiments include validation of the proposed methods and the comparison with various pre-training baseline methods on some public-domain LLMs and vision language models.