



**Title:** Laser-plasma instabilities driven by helical laser beam

**Principal Investigator:** Alexey Arefiev (UC San Diego)

**Co-investigators:** John Palastro (LLE, University of Rochester)

**ALCC Allocation:**

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

**Allocation(s):**

394,008 node-hours on Perlmutter-CPU

**Research Summary:**

Uninhibited propagation of laser beams through plasma over long distances is critical for the success of high-energy density (HED) experiments and inertial confinement fusion. However, this propagation can be adversely affected by multiple laser-plasma instabilities. The topology of plasma waves driven by helical laser beams offers a previously unexplored level of control over laser-plasma interactions. Unlike conventional beams, helical laser beams can exchange angular momentum with plasma and excite helical plasma waves. The helical topology of these plasma waves fundamentally changes their interactions with electrons and ions, altering the evolution and properties of instabilities, including growth rates, thresholds, and saturation. The research program of this project is structured to build in complexity, starting with the basic properties of warm helical plasma waves and progressing to laser-plasma instabilities in a single speckle. Due to the nature of their field structure, helical laser beams and helical plasma waves must be simulated in 3D. This project will employ 3D particle-in-cell (PIC) simulations to capture the relevant physics. The emerging capabilities to generate helical beams at high-energy laser facilities, such as OMEGA and NIF, underscore the importance of a timely examination of the impact that helical laser drivers could have on mitigating laser-plasma instabilities.

Our research into laser-plasma instabilities driven by helical beams is funded by two programs at the US Department of Energy (DOE): the High-Energy-Density Laboratory Plasmas program (DOE/SC/FES/HEDLP) and the Scientific Discovery Through Advanced Computing program (DOE/SC/SCiDAC). The specific area of interest for DOE that this project addresses is the “nonlinear optics of plasmas and laser-plasma interactions” for furthering fusion energy science. The SCiDAC project supporting this proposal aims at unlocking kinetic effects in inertial fusion energy-relevant simulations on Exascale supercomputers. The project will also provide training to graduate students in applying high-performance computing for research of laser-plasma interactions.



**Title:** **AI-Driven Multiscale Investigation of RAS-RAF Activation Lifecycle**

**Principal Investigator:** Fikret Aydin (Lawrence Livermore National Laboratory)

**Co-investigators:** Helgi Ingólfsson (Lawrence Livermore National Laboratory),  
Loïc Pottier (Lawrence Livermore National Laboratory),  
Konstantia Georgouli (Lawrence Livermore National Laboratory),  
Tommy Turbyville (Lawrence Livermore National Laboratory),

**ALCC Allocation:**

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

**Allocation(s):** 900,000 node-hours on Frontier

**Research Summary:**

More than 30% of all human cancers are driven by mutations in RAS proteins; with particularly high prevalence in lung, colon, and pancreatic cancers, such cancers have poor prognosis and high mortality. Despite advances in experimental techniques and structural biology, mechanistic underpinnings of RAS-driven cancers remain obfuscated due to a lack of sufficient resolution. Bridging experiments at the finest scale possible with simulations at the largest/longest scales attainable represents a holy grail for computational biology. Multiscale simulations are needed to create this bridge and assess the molecular mechanism of the growth signaling pathway, specifically RAS interaction with its downstream partner, RAF, and the key steps in the initiation of the signaling cascade.

As part of an ongoing NCI/DOE collaboration, Project ADMIRRAL (*AI-Driven Multiscale Investigation of RAS-RAF Activation Lifecycle*), we are developing a predictive multiscale model of RAS-RAF driven cancer initiation and growth that allows coupling experiments with multiresolution simulations to build a dynamic model of RAS-RAF biology in varying cellular membrane compositions. The combination of simulation with experiment generates a vast amount of data that, with the aid of high-performance analytics, establishes a detailed, predictive understanding of RAS-RAF activation—structural and dynamic insights that will open new possibilities for therapeutic intervention. Recently, we developed the **Multiscale Machine-learned Modeling Infrastructure (MuMMI)**, which represents a new paradigm of multiscale simulations and enables exploring large length- and time-scales (microns and milliseconds, respectively) using macro-level models, while simultaneously maintaining molecular-scale detail using a novel ML-based sampling framework.

Our data-driven approach that interleaves multiscale simulations using ML will have transformative impacts on high-fidelity simulation capabilities that drive fundamental scientific discoveries in the areas of biology, materials science, climate sciences, fluid dynamics, nuclear fusion, etc. This work leverages the DOE leadership to accelerate multiscale simulations using scalable and novel ML strategies.



**Title:** Hadronic contributions to the muon  $g-2$  from lattice QCD

**Principal Investigator:** Thomas Blum (Univ. of Connecticut)

**Co-investigators:** Alexei Bazavov (MSU),  
Peter Boyle (BNL),  
Carleton DeTar (Univ. of Utah),  
Aida El-Khadra (UIUC),  
Steven Gottlieb (Indiana Univ.),  
Taku Izubuchi (BNL),  
Luchang Jin (Univ. of Connecticut),  
Ethan Neil (Univ. of Colorado),  
Ruth Van de Water (FNAL)

**ALCC Allocation:**

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

**Allocation(s):** 3,846,000 node-hours on Frontier,  
1,000,000 node-hours 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 is being measured at Fermilab and calculated by theorists world-wide to fantastic accuracy in a high-stakes test of the Standard Model (SM) of Particle Physics. To test the SM to a degree that allows discovery of physics beyond our current understanding of Nature's laws, the theory errors on 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. These so-called hadronic contributions will be determined with improved precision in numerical simulations of Quantum Chromodynamics (QCD) known as lattice QCD.

The theoretical calculation and measurement of the magnetic moment of the muon comprise one of the highest priorities of the DOE's Office of High Energy Physics. Lattice QCD calculations from many groups, using different formulations and methods, agree with each other but differ with longer-standing data-driven calculations, which calls into question the latter's disagreement with the SM. To resolve the difference between lattice and data-driven theory values, the project aims to compute the hadronic contributions at the sub-percent level, and ultimately to reach the expected precision of the experiment, at the one-to-two-permille level.



**Title:** Decadal Simulations with the Simple Cloud-Resolving E3SM Atmosphere Model

**Principal Investigator:** Peter Caldwell (Lawrence Livermore National Laboratory)

**Co-investigators:** Ben Hillman (Sandia National Laboratory),  
Mark Taylor (Sandia National Laboratory)

**ALCC Allocation:**

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

**Allocation(s):** 1,600,000 node-hours on Frontier

**Research Summary:**

Computational expense has historically forced climate models to discretize the world so coarsely that important features like clouds and mountains are unresolved. This lack of resolution is a major source of uncertainty in climate predictions. Exascale computers free us from this constraint but require extensive modification of existing code. The Energy Exascale Earth System Model (E3SM) project spent 5 years rewriting its atmosphere model in C++ specifically for use on DOE's new exascale computers. [Caldwell et al., \(2021\)](#) demonstrates that this new model resolves away many of the problems that plague conventional climate models.

The goal of this project is to perform an unprecedented 20 year simulation with 3.25 km grid spacing globally (resulting in more than 7 billion total grid cells). This simulation will follow the well-established [AMIP protocol](#) of prescribing time-varying sea surface temperatures and other forcings based on historical observations. Multi-decadal storm-resolving simulations like this are critical for obtaining robust statistics of extreme events, which will be among the most impactful aspects of climate change. In addition, interannual variability in the climate system is large and important but hasn't been adequately evaluated in this new class of extreme-resolution models. By running a global cloud-resolving model following the same protocol as used for conventional climate models, this simulation will enable a quantitative evaluation of the benefits of km-scale global models relative to conventional climate models. In addition, this simulation will serve as the baseline against which a future-climate companion run (to be performed next year) can be evaluated. These simulations will be made publicly available and will be used by many researchers worldwide.



**Title:** Charged Macromolecules in the Nonequilibrium State

**Principal Investigator:** Jan Michael Carrillo (Oak Ridge National Laboratory)

**Co-investigators:** Yangyang Wang (Oak Ridge National Laboratory),  
Wei Ren Chen (Oak Ridge National Laboratory),  
Nick Hagerty (Oak Ridge National Laboratory),  
Monojoy Goswami (Oak Ridge National Laboratory),  
Elijah MacCarthy (Oak Ridge National Laboratory)

**ALCC Allocation:**

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

**Allocation(s):** 425,000 node-hours on Frontier

**Research Summary:**

Charged macromolecules play a crucial role in various technologies such as coatings, sensors, electrolytes, antifouling agents, and drug delivery systems. Despite the significant progress in understanding the equilibrium properties of charged macromolecules, their behavior in nonequilibrium state remains a difficult subject. This research addresses the nonequilibrium structures and dynamics of charged polymers under shear flow by using large-scale molecular dynamics simulations. By employing a newly developed coarse-grained dipolar solvent model and spectrum decomposition techniques, the study seeks to reveal how flow alters the conformations, interactions, and molecular motions of charged polymers. In collaboration with experimental researchers at the Oak Ridge National Laboratory, the project provides direct comparisons between theory, simulation, and small angle scattering experiments, with emphasis on analyzing the complex rheological behavior through the lens of space-time correlation functions. This research aligns with national priorities in pandemic readiness and advances computational tools for modeling charged macromolecules in a wide range of energy related applications.



**Title:** DNS of H<sub>2</sub>/NG blends towards fuel- and load-flexible industrial heat and power

**Principal Investigator:** Jacqueline Chen (Sandia National Laboratories)

**Co-investigators:** Ki Sung Jung (Sandia National Laboratories),  
Martin Rieth (Sandia National Laboratories),  
Cristian Lacey (Sandia National Laboratories)

**ALCC Allocation:**

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

**Allocation(s):** 2,100,000 node-hours on Frontier

**Research Summary:**

DOE's Industrial Heat Shot goal is to provide industrial heat with at least 85 percent reduction in GHG emissions by 2035. Exascale simulations can enable science and functional models to allow 0-100% hydrogen in natural gas (NG) fuel-flexible combined heat and power (CHP) combustion system design/operation and real-time, model-based control for fast load balancing and fuel-flexibility. To accelerate the commercial readiness of low-carbon fuel utilization, fundamental aspects of hydrogen combustion and/or low-carbon- input CHP need to be understood and modeled accurately to facilitate design or retrofit of industrial process heating equipment. Hydrogen's lower volumetric calorific value, its higher flame speed and adiabatic flame temperature, and differences in heating profiles result in unique combustion environments that need to be accommodated to avoid safety issues and to mitigate NO<sub>x</sub> emissions while ensuring ideal heating process environments with new fuels. Direct numerical simulation (DNS) for fuel- and operation- flexible CHP with reciprocating engines utilizing hydrogen/NG blends will shed light on NO emission formation, ignition dynamics, and flame-wall interactions for different fuel blend ratios at engine relevant conditions. The high reactivity and broad flammability limits of hydrogen translate into both increased thermal efficiency and higher heat fluxes to the piston and cylinder walls. The effect of preferential diffusion of hydrogen on the ignition dynamics of hydrogen/NG blends at engine-relevant condition will be elucidated, providing insights for strategies to mitigate knock, NO<sub>x</sub>, excessive wall heat fluxes and flashback. The unique DNS datasets will enable the development and validation of predictive physics- and data-based combustion and wall heat flux models for hydrogen/natural gas fuel blends at engine conditions. The DNS will also provide validation for a novel *on-the-fly* surrogate DNS approach to facilitate computational efficiency.



**Title:** **Building Digital Twin of a Model Host-Pathogen System for Enhancing Biopreparedness**

**Principal Investigator:** Margaret S. Cheung (Pacific Northwest National Laboratory/Environmental Molecular Sciences Laboratory)

**Co-investigators:** David Pollock (PNNL/University of Colorado Anschutz), James Evans (Pacific Northwest National Laboratory), Wei-jun Qian (Pacific Northwest National Laboratory), Amity Andersen (Pacific Northwest National Laboratory/Environmental Molecular Sciences 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):** 50,000 node-hours on Frontier, 100,000 node-hours on Aurora, 100,000 node-hours on Polaris, 200,000 node-hours on Perlmutter-CPU, 100,000 node-hours on Perlmutter-GPU

**Research Summary:**

The science of biopreparedness to counter biological threats hinges on understanding the fundamental principles and molecular mechanisms that lead to pathogenesis and disease transmission. The goal of this research is to address this challenge to create a powerful and user-friendly digital twin platform to elucidate the fundamental principles of how molecular interactions drive pathogen-host relationships and host shifts. Groundbreaking discoveries will be enabled by integrating a wide range of structural, genomics, proteomics, and other advanced “-omics” measurements, along with evolutionary and artificial intelligence predictions. It will be developed in the context of a tractable model system, the small, abundant, and accessible photosynthetic cyanobacteria and their constantly co-adapting viral pathogens, cyanophages. This will maintain the system’s applicability to real-world problems and techniques, but the overall focus will be on elucidating general principles that are system agnostic - detecting, assessing, surveilling molecular interaction, adaptation, and coevolution - and therefore extensible to any viral-host interaction.

The **impact** of the project will be to develop, implement, and test a platform to assess host-pathogen molecular interactions, adaptation to hosts and host shifts, and coevolution between hosts and pathogens. A successful project **outcome** will transform researchers’ ability to study any host-pathogen interaction, encourage diverse community contributions, and gain fundamental insights into how proteins adapt to new contexts relevant to DOE’s associated research in the broader biomanufacturing and bioeconomy.



**Title:** Huge Ensembles of Weather Extremes using the Fourier Forecasting Neural Network

**Principal Investigator:** William Collins (LBNL and University of California, Berkeley)

**Co-investigators:** Michael Pritchard (NVIDIA and the University of California, Irvine),

**ALCC Allocation:**

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

**Allocation(s):** 280,000 node-hours on Perlmutter-GPU

**Research Summary:**

Simulating Earth's climate at high resolution with high fidelity is a computational grand challenge. Reliable projections of how low-likelihood high-impact climate extremes are especially important given the outsize impacts these extremes have on society and the environment, and given measurements and modeling indicating these extremes are becoming more frequent and intense. It has proven very difficult to get adequate estimates of the risks and characteristics of these extremes using the climate record because of the highly localized and sporadic nature of these events. In this project, we overcome this sampling issue using 10,000-member ensembles of hindcasts of the historical record to attain adequate samples of these phenomena. . For the first time, we can now generate such massive ensembles using machine learning at five orders-of-magnitude less compute than traditional numerical simulations.

These huge ensembles will accelerate studies of the particularly intense heat waves, atmospheric rivers, and tropical cyclones and hurricanes. As the summers of 2023 and 2024 have shown, heat extremes have become far more prevalent, and the summer of 2024 is forecast to have an unusually active hurricane season. The huge forecasts enabled by this ALCC allocation will permit detailed assessment of the risks of the extreme phenomena actually experienced against the risks of the mildest to most severe variants of the same phenomena under identical meteorological conditions. This information and these simulations will be freely shared with the global climate community.





**Title:** Twist-3 GPDs of the pion and kaon

**Principal Investigator:** Martha Constantinou (Temple University)

**Co-investigators:** Andreas Metz (Temple University),  
Krzysztof Cichy (Adam Mickiewicz University),

**ALCC Allocation:**

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

**Allocation(s):** 206,400 node-hours on Perlmutter-CPU

**Research Summary:**

The theory of the strong interaction, Quantum Chromodynamics (QCD), binds permanently quarks into hadrons, giving them a rich and complex structure. Understanding their structure from first principles requires a systematic and quantifiable study of QCD, as fundamental properties of the hadronic matter arise from the Standard Model and QCD dynamics. At the hadronic energy scales, the solution of QCD is not amenable to perturbation theory. The only non-perturbative approach that captures the full QCD dynamics is discretizing the continuum theory on a 4-dimensional Euclidean lattice for numerical simulations on powerful computers. This approach, lattice QCD (LQCD), provides a rigorous framework for studying hadron structure non-perturbatively, starting directly from the underlying fundamental theory.

The internal properties of hadrons are studied using a set of essential quantities, the distribution functions, which can be accessed in a variety of high-energy processes. Parton distributions are classified into PDFs, GPD and TMDs, based on their dependence on the longitudinal parton momentum fraction  $x$ , the longitudinal momentum fraction  $\xi$  related to the longitudinal momentum transferred to the hadron, the momentum  $k_T$  transverse to the hadron direction, and the total momentum transferred to the hadron,  $t$ . In fact, GPDs and TMDs provide information on the 3-D tomography of hadrons, which remains a long-term goal of Nuclear Physics, and the science of JLab and the future EIC.

With this ALCC request for computer allocation we propose an extension to our research program to extract the pion and kaon GPD with focus the twist-3 case. The quasi-GPDs approach requires matrix elements with fast moving hadrons, and we will use two values for the momentum boost up to 2 GeV. Another novel aspect of the proposed research is a new method to optimize the computational cost by using an asymmetric frame to extract Lorentz-invariant amplitudes that are related to GPDs [PhysRevD.106.114512]. With this approach, we can access a few values of the momentum transferred squared,  $-t$ , within a single production. We intend to explore about eight values of  $-t$  up to 2 GeV<sup>2</sup>. All quantities will be calculated using an ensemble of  $N_f = 2+1+1$  at a pion mass of 260 MeV with lattice spacing  $a=0.093$  fm. The proposed study is unique and pioneering as these quantities have never been explored so far.



**Title: DNS of Buoyancy Driven Flows for Developing NN-informed High-fidelity Turbulence Closures**

**Principal Investigator:** Som Dutta (Utah State University)

**Co-investigators:** Paul Fischer (University of Illinois at Urbana-Champaign),  
Mauricio Tano-Retamales (Idaho National Laboratory),  
Izabela Gutowska (Oregon State 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 node-hours on Frontier,  
50,000 node-hours on Aurora,  
50,000 node-hours on Polaris,  
200,000 node-hours on Perlmutter-CPU

**Research Summary:**

Buoyancy modulated flows are ubiquitous in natural and the built environments. Accurately predicting the behavior of these flows have implications for energy systems, nuclear reactors, wildfire plumes, thermohaline flows causing accelerated melting of icebergs, etc. Turbulence models used for simulating these phenomena will be improved using machine-learning based surrogates, that will be developed using the data generated from direct numerical simulations (DNS) of buoyancy-driven plumes and currents. DNS resolve all the relevant physics of the flow, providing invaluable data and insight into the physics of the phenomena, while being computationally expensive. Thus, the allocated compute-time on the leadership-scale supercomputers will facilitate DNS of buoyancy-driven flows across the relevant parameter range. The high-fidelity simulations will be conducted using NekRS, a high-order incompressible computational fluid dynamics solver, that can efficiently utilize the GPU accelerated supercomputers. The tera-bytes of data generated from the DNS will be used to train Neural-Network and ML-based turbulence models, which are expected to predict the complex physics at relatively small computational cost.

The enhanced understanding of the physics and the ML-based models developed during this project will help improve design, optimization and safety of advanced nuclear reactors and energy systems, and can be extended to improve the fidelity of models used for predicting buoyancy-driven environmental flows.



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

**Principal Investigator:** Patrick Emami (National Renewable Energy Laboratory)

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

**ALCC Allocation:**

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

**Allocation(s):**  
150,000 node-hours 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. 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. This project will develop foundation models, which are large AI models trained on Internet-scale amounts of data using high-performance computing, that can be deployed within interactive dialogue systems commonly known as chatbots. Scientists will use these chatbots to answer questions about scientific literature, automate workflows, and construct scientific models and reason about ambiguous data. The first phase of the project will develop advances in cutting-edge AI topics such as multimodality, trustworthiness, and agent frameworks that augment foundation models with probabilistic reasoning abilities and connects them with scientific tools and verified knowledge. In the second phase, volunteer scientists will conduct dialogues with the chatbot about scientific modeling of turbulence, chemical reaction networks, and material degradation pathways. 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 an open-source family of foundation models called Theseus, in addition to open data and software. Additionally, the project aims to provide opportunities to researchers from diverse and historically under-represented populations and contribute to raising awareness of emerging challenges that lie at the intersection of AI and energy justice.



**Title:** Learning exchange-correlation functional in DFT from quantum many-body calculations

**Principal Investigator:** Vikram Gavini (University of Michigan)

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

**ALCC Allocation:**

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

**Allocation(s):** 250,000 node-hours 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 quantum Monte Carlo (QMC) and 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 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 work seeks to break this *size-accuracy barrier* through a data-driven approach to modeling the XC functional.

The first part of the proposed study will generate accurate QMC-based ground-state densities for jellium spheres in presence of Gaussian impurities to mimic various regimes of realistic solids. The study will also compute accurate QMC densities for various metallic and non-metallic solids. The second part of the study involves inverse DFT calculations on the QMC densities to obtain their corresponding XC potentials. Finally, various XC functional models will be developed, ranging from local to non-local forms, using the above densities and exact XC potentials as training data. The accuracy of the learnt XC functional models will be tested against widely used G2 and Sol62 databases, for thermochemistry and bulk properties, respectively. The proposed data-driven approach to modeling the XC functional offers a systematic route in alleviating known deficiencies of existing XC approximations—delocalization error, inability to capture strong correlations, inaccurate fundamental gap predictions, to name a few.

**Title:** Real-time extreme-scale inversion and prediction for tsunami early warning



**Principal Investigator:** Omar Ghattas (The University of Texas at Austin)

**Co-investigators:** Stefan Henneking (The University of Texas at Austin),  
Sreeram Venkat (The University of Texas at Austin),  
Milinda Fernando (The University of Texas at Austin)

**ALCC Allocation:**

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

**Allocation(s):** 100,000 node-hours on Perlmutter-GPU

**Research Summary:**

The overarching goal of this project is to create a GPU-accelerated mathematical and computational framework for real-time tsunami inversion and forecasting, and apply this framework to physics-based rapid response and early warning for tsunamis generated from the Cascadia subduction zone. Tsunamis generated from megathrust earthquakes are capable of massive destruction. Cascadia is thought to have a 30% probability of megathrust rupture (resulting in an up to magnitude 9 earthquake) within the next 50 years. Subduction zones are becoming increasingly instrumented. This project aims to employ data from seafloor acoustic pressure sensors in subduction zones, along with a forward model in the form of coupled acoustic-gravity wave equations, to infer spatiotemporal seafloor motion. The solution of this inverse problem then provides the boundary forcing to forward propagate the tsunamis toward populated regions along coastlines. The entire end-to-end “data-to-inference-to-prediction” computation is carried out in a Bayesian framework that rigorously accounts for uncertainties. The end-to-end computation is carried out in real time (order of seconds). Overcoming the above inverse problem challenges—extreme scale, uncertainty quantification, and real-time execution—is intractable with conventional algorithms. This work employs recently designed fast algorithms that exploit the structure of the inverse operator such that real-time execution, with high fidelity computational models, under uncertainty, can be achieved on a large GPU cluster such as Perlmutter-GPU.

Many problems within the DOE portfolio give rise to digital twins, a central component of which is data assimilation/inversion. Tsunami inversion has the same mathematical structure as a number of these problems, including inferring the source and transport of atmospheric contaminants from volcanos, greenhouse gases, or accidental or intentional releases of chemical, biological, nuclear, or other hazardous agents; and the real-time inference of seismic sources for treaty verification. The development of real-time, high-fidelity, uncertainty-equipped, GPU-efficient algorithms for digital twins will therefore have broad impact.

**Title:** Continually-Learning Foundation Models for Scientific Imaging



**Principal Investigator:** Jens Glaser (Oak Ridge National Laboratory)

**Co-investigators:** Irina Rish ( Université de Montréal)

**ALCC Allocation:**

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

**Allocation(s):** 710,000 node-hours on Frontier

**Research Summary:**

The rise of large, unsupervised pre-trained models, known as foundation models, has significantly enhanced AI's ability to generalize from minimal data and transfer learned knowledge across tasks. However, continually retraining these models as new data arrives is inefficient. This project focuses on continuously training existing models to adapt to new data without losing previously learned information. The research aims to leverage foundation models for scientific imaging in structural biology and brain imaging to facilitate automated scientific discovery. By initially training on a wide range of image data and then refining the models with specific scientific data, the project will improve the analysis of image data from neutron diffraction and MRI experiments.

This project will produce highly adaptable computer vision models, making them publicly accessible to the broader research community. By enabling real-time analysis of diverse scientific data, from neutron crystallography to high-resolution MRI brain scans, the research addresses the challenge of data scarcity and real-time processing demands. The project supports initiatives like dynamic nuclear polarization at Oak Ridge National Laboratory, accelerating structural discoveries in biology and materials science. Public access to these models will foster innovation and collaborative research, enhancing the ability to decode complex biological and material structures, and thereby enable imaging breakthroughs through artificial intelligence.



**Title:** **Highly Scalable Ab Initio Simulations of N-Doped Porous Materials for Carbon Capture**

**Principal Investigator:** Mark Gordon (Ames National Laboratory)

**Co-investigators:** Peng Xu (Ames National Laboratory),  
Buu Pham (Ames National Laboratory),  
Tosaporn Sattasathuchana (Ames National Laboratory),  
Sarom Leang (EP Analytics),

**ALCC Allocation:**

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

**Allocation(s):** 2,000,000 node-hours on Aurora

**Research Summary:**

This research aims to improve the carbon capture capabilities of nitrogen-doped porous materials, specifically nitrogen-assembly carbons (NACs), for carbon dioxide and methane in gas mixtures. The team will predict the nature and the magnitudes of the interactions between gas molecules and NAC surfaces using scalable quantum mechanical methods combined with Monte Carlo techniques. Building upon the Exascale Computing Project and utilizing Department of Energy computing facilities, the research will determine the role of nitrogen-containing species in NACs that exhibit selectivity for carbon dioxide and methane in environments such as flue gas, biogas, and gaseous products of carbon dioxide conversions. The outcomes of this research will guide the design and synthesis of NACs for carbon capture and storage, as well as chemical transformations. They will foster collaboration between experimental and theoretical scientists at Ames National Laboratory and beyond. The results will advance the development of carbon capture technologies essential for mitigating greenhouse gas emissions impacts on health, the economy, and the environment, aligning with the Department of Energy's mission to address energy and environmental challenges through science and technology.



**Title:** **Simulating large-scale long-lived neutron star remnants from binary neutron star mergers**

**Principal Investigator:** Ore Gottlieb (Flatiron Institute)

**Co-investigators:** Alexander Tchekhovskoy (Northwestern University),  
Francois Foucart (University of New Hampshire),  
Goni Halevi (Northwestern University),  
Danat Issa (Northwestern University),  
Kyle Parfrey (Princeton Plasma Physics Laboratory),  
Daniel Kasen (UC Berkeley),  
Brian Metzger (Columbia University),  
Nick Kaaz (Northwestern University),  
Carlos Palenzuela (University of the Balearic Islands)

**ALCC Allocation:**

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

**Allocation(s):** 400,000 node-hours 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:** High-fidelity WarpX simulations of stability and operational boundaries of advanced FRCs

**Principal Investigator:** Roelof Groenewald (TAE Technologies Inc.)

**Co-investigators:** Sean Dettrick (TAE Technologies Inc.),  
Elena Belova (Princeton Plasma Physics Laboratory),  
Calvin Lau (TAE Technologies Inc.)

**ALCC Allocation:**

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

**Allocation(s):** 65,000 node-hours on Perlmutter-GPU

**Research Summary:**

TAE Technologies, Inc, is a private company focused on the development of carbon-free baseload electricity production using fusion plasmas confined by magnetic fields in a field-reversed configuration (FRC). TAE's national lab scale experimental device, C-2W, which has achieved unprecedented steady-state FRC plasmas at record densities and temperatures, uses neutral beam injection (NBI) to sustain this new advanced FRC. In this project computer simulations of NBI into FRCs will be performed to gain deeper physics understanding of C-2W experimental results so that these advances can be carried forward to TAE's next generation machine. The particle-in-cell code, WarpX, which was developed by scientists at Lawrence Berkeley National Laboratory (LBNL) will be used for these simulations. In collaboration with LBNL, TAE Technologies has added new physics models to WarpX to access the time and space scales that are relevant to FRC plasmas. Simulations will reach milliseconds in simulation time while capturing the full plasma volume in 3D. The project will begin with validation of the new physics model in code-to-code comparisons with the HYM code developed by Princeton Plasma Physics Laboratory (PPPL), as well as comparisons to experimental results from C-2W. Then a first-of-its-kind kinetic study of global stability of FRCs with NBI will be performed to understand how NBI can stabilize FRC plasmas, and to find the operational boundaries and best operating points of steady-state FRCs.

This public-private partnership between TAE and the Department of Energy's LBNL and PPPL laboratories 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:** **Autonomy for DOE Simulations**

**Principal Investigator:** Allan Grosvenor (MSBAI — dba Microsurgeonbot Inc.)

**ALCC Allocation:**

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

**Allocation(s):** 100,000 node-hours on Frontier,  
30,000 node-hours on Aurora

**Research Summary:**

GURU, is an autonomous system that drives expert workflows in software, to remove the barriers that prohibit engineers from achieving the promise of simulation and High Performance Computing (HPC). MSBAI built and demonstrated GURU, reducing Virtual Prototype modeling and simulation setup time from hours to minutes for a range of engineering applications. This latest effort focuses on expanding the flexibility and usability of the system, by adding hybrid interaction, so users can type or speak requests, as well as gesturing and even sketching. Advancements to GURU's learning engine (the skills agent factory ), and new R&D will improve the speed and flexibility of generating new high quality, readily transportable, geometrical Computer Aided Design (CAD) files.

A broad range of industries, from energy, automotive, and aerospace, to telecommunications, medical devices, and consumer products, expect a third of their sales to come from new products, which means that \$30 trillion in revenues depends upon successful product development. In response to this critical need, organizations have increasingly adopted simulation as a standard practice, yet only 5% of these industries have fully integrated Artificial Intelligence (AI) into their operations. The setup of simulations remains a laborious process, potentially demanding months or even years for users to fully master. This challenge poses a barrier to achieving the Department of Energy's Exascale Computing Project (ECP) Vision. One of the longest-standing barriers to productivity in engineering utilization of computational physics simulations, as well as manufacturing, is the preparation and exchange of geometrical data. Data, quality, or functionality are routinely lost when converting, importing, or exporting CAD files which often causes engineers to lose enormous amounts of time and prohibits them from solving design and manufacturing challenges. Humanity has an opportunity to enter a new renaissance in creative and engineering design, as we bring the power of validated specialized modeling and simulation software together with HPC to scale up their utilization, driven by the latest AI methodologies to enable the mass adoption of these tools. The successful outcome of this research will achieve new scientific outcomes that demonstrate the latest potential for HPC, simulation, and AI to hyper-enable engineers to solve humanity's largest challenges, and the business implications of our success in this effort are represented by the \$30 trillion opportunity described above.



**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 (INT, University of Washington, Seattle)

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

**Allocation(s):** 540,000 node-hours on Frontier

**Research Summary:**

The standard model of elementary particle interactions has three flavors of neutrinos that are observed to transform into each other, a discovery that was awarded the 2015 Nobel prize in physics. This mixing implies that neutrinos have mass but does not inform us of the ordering of the three flavors by mass and whether the interactions of neutrinos and anti-neutrinos are the same. These two properties of neutrinos, the mass hierarchy and whether the interactions of neutrinos and anti-neutrinos are the same, which would imply that the combined charge-conjugation and parity (CP) is a good symmetry in the neutrino sector, will be measured in high precision neutrino oscillation experiments such as the Short Baseline Neutrino (SBN) and Deep Underground Neutrino Experiment (DUNE) centered at Fermilab. A large violation of the CP symmetry would satisfy a key condition for a process that could explain one of the most profound mysteries of the observed universe—why is there almost no antimatter today when all plausible initial starting points in the evolution of the universe after the big bang have equal amounts of matter and antimatter. To reach the design precision of these neutrino experiments requires knowing the strength of the interaction, encapsulated in the charged current neutrino-nucleus scattering cross section, and the incident neutrino flux to within a few percent. A key input in theoretical calculations of these is the isovector axial form factor of nucleons over the full energy range of the incoming neutrinos. The proposed calculations will reduce the current uncertainty of about ten percent in the axial vector form factors to the required few percent by addressing the least controlled systematic uncertainty through the proposed high statistics and better tuned simulations.



**Title:** **Developing biopreparedness models informed by heterogeneous data streams**

**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):** 130,000 node-hours on Frontier

**Research Summary:**

Large language models (LLMs) are poised to transform scientific inquiry and data-driven decision-making by facilitating the integration of vast amounts of data. Their ability to process complex concepts make them especially exciting in fields such as medicine. However, their broad use for precision medicine and public health surveillance is inhibited by two factors: (1) the siloed, tightly controlled, and sensitive nature of health data; and (2) the lack of organization around standardized and high-quality information retrieval. With the advent of generative AI, technological solutions to these challenges will be realized over the next few years. The creation of high fidelity synthetic datasets with real-world data presents a solution for enabling the exchange of large volumes of decentralized and sensitive health data. New developments in AI foundation models offer significant potential to help automate the standardization and harmonization of health information from unstructured clinical notes. This project aims to develop solutions that will enhance nationwide health data interoperability and result in an expanded high performance computing (HPC) infrastructure for the modeling of high-volume multimodal data at a global scale.

A framework for enabling the sharing of sensitive health data across the United States for the extraction of near-real-time disease information will be established through the completion of two aims. First, researchers will develop a novel multi-stage approach for producing synthetic datasets that are similar to real world data without sacrificing patient privacy. Second, researchers will conduct a deep interrogation into the foundation model paradigm, quantifying its promise as a solution for efficiently developing and deploying rapid extraction of disease phenotypes from electronic health data at population scale. The synthetic data and HPC workflows created by this project will provide a data-driven virtual health ecosystem that can be used to facilitate rapid response to threats to human health security; furthering the Advanced Scientific Computing Research's mission to develop and deploy computational tools for the advancement of science.



**Title:** Informing Forensics Investigations of Nuclear Materials

**Principal Investigator:** Sara Isbill (Oak Ridge National Laboratory)

**Co-investigators:** Ashley Shields (Oak Ridge National Laboratory),  
Jennifer Niedziela (Oak Ridge National Laboratory),  
Andrew Miskowiec (Oak Ridge National Laboratory)

**ALCC Allocation:**

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

**Allocation(s):** 200,000 node-hours on Frontier

**Research Summary:**

Researchers at ORNL 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 Nuclear Nonproliferation Research and Development program, 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. 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, allowing direct connections of chemical and physical changes in fuel cycle materials to the experimental observations. Additionally, determining the effect of process and environment 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. Atomistic modeling is an invaluable tool in nonproliferation materials research portfolio.

This project will connect density functional theory determinations of reaction products and lattice dynamics of carbon composites, 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 characterizations. Experimentally determining reaction kinetics for solid-state materials is challenging and an area that will greatly benefit from computational modeling of activation barriers. We will perform calculations on large simulation cells to capture experimentally relevant structural, vibrational, and kinetic properties with experimental efforts to understand fuel cycle-relevant materials for nuclear nonproliferation.



**Title:** Integrative Systems Biology For Epigenetics in the EPIC BRaVE Project

**Principal Investigator:** Daniel Jacobson (Oak Ridge National Laboratory)

**Co-investigators:** Veronica Melesse Vergara (Oak Ridge National Laboratory),  
Shawn Starkenburg (Los Alamos National Laboratory),  
Christina Steadman (Los Alamos National Laboratory)

**ALCC Allocation:**

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

**Allocation(s):** 600,000 node-hours on Frontier

**Research Summary:**

The EPIC BRaVE (Exascale Epigenomic Profiling of Pathogen Interactions with Chromatin) project aims to uncover how pathogens alter host genome structures to enhance infection and how hosts adapt their epigenetic landscapes in response. By using advanced single-cell epigenomics and high-performance computing (HPC) workflows, this research will analyze the changes in genome architecture and gene expression during infections. The approach will be applied to both plant and mammalian systems to identify specific epigenetic signatures linked to infection and resilience.

Key to this project is the hypothesis that resilient hosts show distinct epigenetic modifications compared to susceptible ones. The research will generate extensive multi-omic datasets, including single-cell Assay for Transposase-Accessible Chromatin using sequencing (scATAC-seq), single-cell Chromatin Immunoprecipitation sequencing (scChIP-seq), and single-cell Methylation sequencing (scMethylation-seq). Advanced computational methods will integrate these datasets to create detailed network models of host-pathogen interactions. By comparing these networks across different pathogens and hosts, the project aims to uncover common and unique epigenetic changes that signify infection processes and host defenses.

The expected impact of the EPIC BRaVE project is significant. By identifying epigenetic changes associated with infection and resilience, it will enhance our understanding of disease mechanisms and improve our ability to predict and respond to biological threats, including pandemics. The findings will support biosurveillance efforts by enabling early detection of pathogen-induced epigenetic alterations, potentially before clinical symptoms appear. This aligns with the Department of Energy's mission to advance national security and public health through scientific research. In summary, the EPIC BRaVE project aims to improve our understanding of host-pathogen interactions through epigenetic research, using advanced technologies and computational resources to address significant biosecurity and public health challenges.



**Title:** High-fidelity simulations of particle-laden non-equilibrium turbulent flows

**Principal Investigator:** Suhas Jain (Georgia Institute of Technology)

**Co-investigators:** Wen Wu (University of Mississippi),

**ALCC Allocation:**

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

**Allocation(s):** 3,000,000 node-hours on Frontier

**Research Summary:**

This project focuses on using high-fidelity predictive simulations to study the impact of particles on the performance of turbomachinery. There is growing evidence that the dispersed phases, such as dust and droplets, can alter the turbulent boundary layer and flow separation leading to aerodynamic degradation. However, the main focus of previous investigations has been on the role of heavy particles in causing erosion, fouling, and corrosion in these applications, and much less attention has been paid to characterizing their direct influence on the aero/hydrodynamic performance of these systems. Hence the team aims to perform high-fidelity numerical ‘experiments’ by adopting an Euler-Lagrange framework, to advance the capability to understand and predict the aerodynamic degradation of gas and wind turbines in multiphase environments. The planned specific objectives of the project include conducting wall-resolved large eddy simulations (WRLES) of particle-laden flow on two configurations: a NACA23012 airfoil and a NASA Rotor 37 compressor blade.

The project aims to develop a better understanding of how airborne contaminants, such as dust and droplets, affect the performance of gas and wind turbines. This research is directly relevant to the Department of Energy's (DOE) mission of ensuring the efficiency and reliability of energy technologies. The project focuses on addressing the lack of understanding regarding the impact of particle contamination on turbomachinery performance, by developing simulation tools that can accurately predict the efficiency of turbomachinery components in these environments, which are crucial for both power generation and energy storage. Current degradation models heavily rely on empirical estimations and limited data. This project aims to provide reliable data and knowledge on the interaction between particles and turbulence, which can inform engineering designs and operational schedules for gas and wind turbines. The project's findings are expected to have a significant impact on the design and operation of turbomachinery, leading to improved efficiency, reliability, and potentially extending the lifespan of these crucial energy technologies. This aligns with DOE's mission of promoting energy security and sustainability.



**Title:** Predicting heterogeneous photocatalysts using large-scale *ab initio* calculations

**Principal Investigator:** Felipe Jornada (Stanford University)

**Co-investigators:** Mauro Del Ben (Lawrence Berkeley National Laboratory)  
Tadashi Ogitsu (Lawrence Livermore National Laboratory)  
Alfredo Correa (Lawrence Livermore National Laboratory)  
Xavier Andrade (Lawrence Livermore National Laboratory)

**ALCC Allocation:**

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

**Allocation(s):** 140,000 node-hours on Frontier,  
100,000 node-hours on Aurora

**Research Summary:**

Mass-produced chemicals – from plastics to fertilizers – often involve chemical reactions that require intense heat and/or pressure, and hence significant amounts of energy. A promising alternative is to deliver energy to drive these reactions in a more controlled way in the form of light, in a phenomenon known as photocatalysis. Several light-driven reactions have been demonstrated with up to 100x improvement in their speed and efficiency compared to purely thermal processes. Still, the underlying mechanisms in photocatalysis are poorly understood, making it harder to systematically develop efficient light-driven reactions. This research will combine massively parallel computer simulations at the Frontier and Aurora supercomputers with modern, quantum-mechanical theories to understand photocatalysts with unprecedented accuracy and generate new design principles.

The calculations in this research will allow one to understand several key steps that dictate the processes of photocatalysis, including: 1) the absorption of light by a material (the photocatalyst); 2) the transfer of photon energy to the chemical reaction; and 3) the resulting evolution of the reaction with time. This research will study prototypical photocatalytic materials such as titanium dioxide and silver nanoparticles and derive connections between materials properties, such as their electronic properties and geometries, with photocatalytic efficiency. Such a connection is critical since materials that are good light absorbers are not necessarily good at driving chemical reactions. Our research has the potential to provide a standardized and systematic way of evaluating and engineering photocatalysts, which will greatly accelerate the development and deployment of this technology and address several of the DOE's missions, such as understanding and optimizing the energy flow at the nanoscale.





**Title:** **Scaling transformer neural networks for skillful and reliable weather, S2S, and Climate forecasting**

**Principal Investigator:** Rao Kotamarthi (Argonne National Laboratory)

**Co-investigators:** Sandeep Madireddy (Argonne National Laboratory)  
Troy Arcomano (Argonne National Laboratory)  
Timothy Smith (NOAA Physical Sciences Laboratory)  
Niraj Agarwal (NOAA Physical Sciences Laboratory)

**ALCC Allocation:**

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

**Allocation(s):**

250,000 node-hours on Polaris  
200,000 node-hours on Frontier

**Research Summary:**

The advent of foundation models for forecasting weather and climate marks a pivotal moment in scientific research, enabling a deeper understanding of complex dynamical systems across vast scales observed with complex Earth systems. These large multi-billion parameter models effectively integrate large-scale data and high-performance computing, and promise enhanced predictive accuracy for critical DOE science applications such as weather forecasting and climate simulation. While the proposed multi-scale foundation models are generally applicable across DOE science applications, we will demonstrate them with the application of weather/climate forecasting, leveraging our team's proven expertise and pioneering work in this domain. We propose research along the following five thrusts (a) develop neural architectures and data engineering pipelines for multi-scale modeling of geophysical flow systems across scales using large, heterogeneous datasets (b) refine pre-trained diffusion models on multi-scale physics simulations by ensuring their strict adherence to physical constraints at each scale (c) develop multi-billion parameter, multi-resolution diffusion models by integrating cutting-edge parallelism techniques and tokenization strategies to ensure scalable, efficient training and high accuracy, employing dynamic routing and novel scaling laws to manage diverse scales and enhance inference speeds (d) employ a Bayesian fine-tuning technique based on Low-Rank Adaptation, developing algorithms to quantify both epistemic and aleatory uncertainties, and implementing multi-scale uncertainty quantification strategies to improve the model's interpretability, reliability, and sample efficiency and (e) ensure the trustworthiness of our foundational models by developing strategies to measure and mitigate bias, address inaccuracies, and establish a robust risk mitigation plan for predictions across various scales.



**Title:** High-Fidelity Simulations of Helium-Air Mixing in High-Temperature Gas Reactor Cavities

**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):** 115,000 node-hours on Polaris

**Research Summary:**

Advanced reactors such as High Temperature Gas Reactors (HTGR) and Sodium Fast Reactors (SFR), are being developed by US companies for deployment in the late 2020s or early 2030s. One of the critical passive safety design tests for High Temperature Gas Reactors (HTGR) is the ability to dissipate decay heat safely during Depressurized Conduction Cooling (DCC) conditions. Following a break in the primary system of a HTGR, hot helium gas will be discharged from a high-pressure primary loop into the reactor cavities 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 vent paths. As helium continues to escape from the primary system, the primary loop pressure falls eventually to an equilibrium pressure. The cavity may then be filled with a well-mixed air-helium, stratified layers of air and helium, or only helium or air depending on the cavity vent paths. After the primary system pressure falls to an equilibrium pressure in the cavities, a helium-air mixture may enter the primary system through the break and potentially leading to catastrophic results. The amount of air entering the reactor is significantly governed by helium-air mixing and possible helium-air stratification inside the reactor cavity. It is necessary to check the design test by modeling the gas-mixing phenomena computationally using system level or detailed 3D CFD codes, and by conducting experiments using scaled models to validate the computational results. NekRS will be used to conduct high-fidelity (DNS and LES) simulations of the gas mixing and stratification process. Results from the simulations will be compared to experimental measurements. It is expected that the mixing of gases and high-speed jet could introduce a substantial degree of compressibility effect in the simulations. We will investigate the validity of low-Mach number assumption and the performance of NekRS as the speed of jet gradually is ramped up. Results from this work will be included in a V&V database for future validation of HTGR simulation codes and could be used to improve the sub-grid scale models in low-order codes.



**Title:** Coastal wetland carbon balance of the conterminous U.S.

**Principal Investigator:** Jinxun Liu (U.S. Geological Survey, Western Geographic Science Center)

**Co-investigators:** Christopher Soulard (U.S. Geological Survey),  
Benjamin Sleeter (U.S. Geological Survey)

**ALCC Allocation:**

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

**Allocation(s):** 30,000 node-hours on Perlmutter-CPU

**Research Summary:**

The conterminous United States (CONUS) contains approximately 40 million acres of diverse coastal wetlands. These wetlands and their associated watersheds are highly susceptible to climate change and human land-use changes, which can significantly impact their carbon storage and fluxes. However, there are substantial uncertainties in quantifying the carbon budget of coastal wetlands and watersheds, a critical component of national greenhouse gas accounting. This project addresses this gap by quantifying CONUS coastal wetland carbon changes from 1985 to the present (2024) and future projections from 2025 to 2100 at a high resolution of 250 meters using models and remote sensing data. The project will utilize remotely sensed monthly surface water extent data and annual land cover change data integrated with ecosystem models, a soil erosion model, and a land change model. This will enable the quantification of carbon sequestration, lateral carbon movement within the landscape, and carbon transfer across the terrestrial-aquatic continuum under the combined influences of land-use and climate change.

The project has two primary goals. (1) Comprehensive Assessment of CONUS Coastal Carbon Changes: This goal encompasses major ecosystem types (forests, shrublands, grasslands, agriculture, and wetlands), key carbon pools (live and dead woody biomass, ground litter, soil carbon), greenhouse gas fluxes (CO<sub>2</sub> and CH<sub>4</sub> emissions), and critical controlling processes such as climate, fire, harvesting, land-use and land-cover change, soil erosion, and sea level rise. (2) Quantifying Carbon Lateral Transport: This goal focuses on quantifying the lateral movement of carbon across the landscape, specifically from land to aquatic systems and from inland to coastal waters. This proposal aligns with the U.S. Department of Energy (DOE) Office of Biological and Environmental Research (BER) Earth and Environmental System Sciences program, which aims to strengthen the predictive capabilities of ecosystem and global-scale models using high-performance computing resources. Additionally, it directly addresses the National Multi-Agency Research and Development Priorities, specifically the focus area of "Tackling Climate Change: Nature-based Climate Solutions for Mitigation and Adaptation".



**Title:** Computational Systems for Tracking Ocean Carbon

**Principal Investigator:** Matthew Long ([C]Worthy)

**Co-investigators:** Alicia Karspeck ([C]Worthy),  
Scott Bachman ([C]Worthy)

**ALCC Allocation:**

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

**Allocation(s):** 492,000 node-hours on Perlmutter-CPU

**Research Summary:**

The Intergovernmental Panel on Climate Change emphasizes that industrial-scale carbon dioxide removal (CDR) will be necessary to meet the Paris Agreement's goal of keeping the rise in global temperature to below 2°C. The amount of CDR required is large — depending on the scenario, society must be removing as much as 10 billion tons of CO<sub>2</sub> from the atmosphere per year by mid-century. Ocean-based pathways are some of the most promising and scalable approaches to CDR. However, ocean pathways are open-system approaches, relying on natural biogeochemical processes outside of direct human control. Further, the vast and dynamic nature of the ocean, coupled with complex ocean ecosystem processes, requires specialized modeling tools to establish scientifically credible methods to quantify the safety and efficacy of ocean-based CDR. This project comprises a computational research plan to support the development and application of a numerical modeling system called C-Star, which will support site characterization, deployment design, and carbon accounting for ocean CDR. The computational resources on this project will be used for C-Star development, research into key controls on uncertainty, efficient computational methods, and simulations explicitly focused on supporting quantification of ocean CDR deployments.

This project will produce several key outcomes. The C-Star modeling system will be established as a standard for scientifically vetted, open-source tools tailored to the needs of the nascent ocean-based CDR industry. The targeted computational research will contribute to catalyzing and advancing scientific understanding of CDR pathways, clarifying the requirements for robust quantification, and establishing a prioritization of key factors driving uncertainty. Finally, by contributing scientific rigor and technological advancement, C-Star will build confidence in open system MRV, further catalyzing innovation and sustained investment to support a potentially critical industry for climate solutions.



**Title:** Scalable and Resilient Modeling for Federated Learning Systems and Applications

**Principal Investigator:** Xiaoyi Lu (University of California, Merced)

**Co-investigators:** Sheng Di (Argonne National Laboratory),  
Junjing Deng (Argonne National Laboratory),  
Yanfei Guo (Argonne National Laboratory),  
Kibaek Kim (Argonne National Laboratory),  
Guanpeng Li (University of Iowa)

**ALCC Allocation:**

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

**Allocation(s):** 15,360 node-hours on Frontier,  
8,960 node-hours on Aurora,  
35,840 node-hours on Polaris,  
76,800 node-hours on Perlmutter-CPU,  
130,720 node-hours on Perlmutter-GPU

**Research Summary:**

This project aims to harness substantial computational resources from the ASCR Leadership Computing Challenge (ALCC) to design and develop a scalable and resilient modeling and simulation framework for federated learning systems and applications. Leveraging ALCC computing allocations on Frontier, Aurora, Polaris, and Perlmutter, the project will conduct meticulous modeling and large-scale simulations essential for refining federated learning systems and understanding their dynamics in evolving real-world scenarios. Additionally, ALCC resources will facilitate the development and evaluation of advanced federated learning systems, workflows, and applications.

This project is expected to significantly impact mission-driven scientific applications and workflows of the Department of Energy (DOE). The project will pioneer advancements in federated learning by fostering more effective simulation and modeling techniques, addressing the critical challenges of scalability and resilience in distributed federated learning systems and the associated scientific workflows. The extensive computing resources provided by the ALCC are crucial, enabling the project to explore and expand the frontiers of current federated learning capabilities and applications.



**Title:** Global kilometer-scale atmospheric simulations with prognostic aerosol

**Principal Investigator:** Po-Lun Ma (Pacific Northwest National Laboratory)

**Co-investigators:** Jerome Fast (Pacific Northwest National Laboratory)

**ALCC Allocation:**

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

**Allocation(s):** 1,600,000 node-hours on Frontier

**Research Summary:**

The role of aerosol in the Earth system remains a major source of uncertainty in predicting the climate in the past, present, and future. The spatial resolution used in global models is insufficient for representing atmospheric features critical for aerosol and aerosol-cloud interactions (ACI). The objective of this research is to perform global kilometer-scale simulations to quantify aerosol effects on Earth's energy balance, on clouds and precipitation, and on weather and climate. This study uses the U.S. Department of Energy's Energy Exascale Earth System Model (E3SM) Atmospheric Model in C++ with the modal aerosol module. This new modeling capability is developed to run efficiently at global kilometer scale on Graphic Processing Units (GPUs). The global kilometer-scale simulations with detailed aerosol and ACI processes will provide more realistic simulations of aerosol and ACI, reducing the uncertainty of actionable predictions of the Earth system. The simulation data will provide new insights into the role of aerosol in the Earth system and serve as the basis for building a digital twin for the Earth system. Model simulation results will be evaluated against in-situ and remotely sensed observational data to ensure the fidelity.

This research leverages the Department of Energy's capability in (1) high-performance computing for generating and analyzing large model simulation datasets, (2) Earth system modeling across scales, and (3) machine learning techniques. This research improves the understanding and predictability of the Earth system by enabling aerosol-cloud interactions at global convection permitting scales.



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

**Principal Investigator:** Michael Mahoney (Lawrence Berkeley National Laboratory)

**Co-investigators:** Prasanna Balaprakash (Oak Ridge National Laboratory)  
Amir Gholami (University of California, Berkeley)  
Yaoqing Yang (Dartmouth)  
Wuyang Chen (Simon Fraser University)

**ALCC Allocation:**

**Site(s):** Oak Ridge National Laboratory (OLCF)  
National Energy Research Scientific Computing Center (NERSC)

**Allocation(s):**

500,000 node-hours on Frontier  
310,000 node-hours on Perlmutter-GPU

**Research Summary:**

Foundation models (FMs) have received interest in machine learning (ML) and Scientific ML (SciML) as a platform with which to train using extensive data, typically through large-scale self-supervision, that can be adapted to a variety of downstream tasks. It is natural to wonder whether it is possible to develop a broad-based scientific foundation model (SFM). Such a model would take advantage of the broad range of observational, experimental, and simulation data modalities available at the DOE, as well as the extraordinary high performance computing (HPC) capabilities available at the DOE.

There is currently no SFM that has been trained on a broad range of data across different scientific domains, space and time scales, to gain an understanding of multiple physical processes and their interactions in a complex scientific system. We have identified three main challenges that currently block the development of a general SFM: (1) lack of neural scaling, with respect to model/data/compute as well as spatio-temporal scaling; (2) lack of control on out-of-distribution (OOD) generalization; and (3) lack of broad-based multi-modal data for training.

In our proposed work, we will address each of these challenges, and in doing so develop SciGPT, an example of what we call a blueprint-SFM. We will address the above three challenges, and in doing so provide SciGPT, a broad-based blueprint-SFM model that (1) is applicable via transfer learning to multiple scientific domains and (2) provides a clear blueprint to develop a general SFM. Developing a general SFM is beyond the scope (budget/compute/data) of this project. However, our learnings will provide a clear path forward for more general investment in this area, both for a general SFM and for multiple domain-specific SciML models.



**Title:**            **Application of Exascale Computing to Regional-Scale Earthquake Hazard and Risk Assessments with a Focus on Emerging Energy Systems**

**Principal Investigator:**           David McCallen (Lawrence Berkeley National Laboratory)

**Co-investigators:**           Arben Pitarka (Lawrence Livermore National Laboratory),  
Rie Nakata (Lawrence Berkeley National Laboratory),  
Houjun Tang (Lawrence Berkeley National Laboratory)

**ALCC Allocation:**

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

**Allocation(s):** 1,100,000 node-hours on Frontier

**Research Summary:**

The DOE Exascale Computing Project (ECP) successfully developed the Earthquake SIMulation (EQSIM) application framework for regional scale, fault-to-structure simulations of earthquake processes. The EQSIM framework was developed specifically for GPU-accelerated systems and has substantially pushed the achievable fidelity and computational performance boundaries of regional-scale simulations for earthquake hazard and risk<sup>1</sup>. With the conclusion of the ECP, EQSIM is being pressed into application in the DOE development of advanced capabilities for risk assessments for critical energy systems. EQSIM models are under development for the San Francisco Bay Area, Southern California and the New Madrid region of the midwestern U.S. The earthquake simulation results from these regions will be shared in an open-access website at the Pacific Earthquake Engineering Research Center located at the University of California Berkeley to provide a unique database accessible by both the broad earthquake research and infrastructure design practitioner communities. In addition, the DOE Office of Science Advanced Scientific Computing Research (ASCR) Science Foundations for Energy Earthshots program is supporting the development of a *Hybrid Framework of Exascale Simulations, Observations and Deep Learning for System-Level Clean Energy Resilience and Risk Reduction*. This project, a combined effort of Caltech, LBNL and UCLA is exploring the application of deep learning to maximize the understanding obtainable through a synthesis of measured data and high-performance simulations. For all of these activities, access to the DOE's leadership computing and the Frontier system is essential, with the ultimate objectives of advancing our fundamental earthquake knowledge and providing the information necessary for developing the most resilient energy systems.

---

<sup>1</sup> McCallen, D.M., Pitarka, A., Tang, H., Pankajakshan, R., Petersson, N.A., Miah, M., Huang, J., 2024, Regional-scale fault-to-structure earthquake simulations with the EQSIM framework: Workflow maturation and computational performance on GPU-accelerated exascale platforms, *Earthquake Spectra*, early access, 2024, doi: 10.1177/87552930241246235.





**Title:** Exploring Exascale Quantum Chemical Methods for Transition Metal Chemistry

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

**Co-investigators:** **PNNL:** E. Aprà, N. Bauman, N. Govind, K. Kowalski, E. Mutlu, A. Panyala, H. Pathak. **UF:** A. A. Mata, R. Hennig, S. B. Trickey

**ALCC Allocation:**

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

**Allocation(s):** 313,442 node-hours on Frontier,  
145,628 node-hours on Aurora

**Research Summary:**

An under-appreciated reality of present-day digital electronics is that the electricity consumed for each digital operation is unsustainable. The power per operation of the next generation of microelectronic devices must be drastically smaller. Achieving that requires big advances that bridge current knowledge gaps. A plausible route forward is quantum technologies that manipulate electron spin rather than charge, as in current systems (for example, single-molecule spintronic devices). A key fundamental science challenge to that route is the limitation of current electronic structure methods to model transition metal physics and chemistry in diverse environments accurately. This is in stark contrast with other areas of chemistry and materials science for which contemporary electronic structure methods have provided invaluable predictive insight.

This project will use exascale-ready open-source software to exploit advanced, high-level quantum chemical methods to study transition metal systems with non-trivial electronic structure features, then use the data generated to train a machine-learning model that circumvents known limitations of the methods. The generated model will have high value for exploring the phenomena critical to emerging next-generation low-power, high-storage density technologies such as molecular-spin-based qubits for quantum information systems and molecular-complex spintronic components. We also expect the model to be valuable for exploring catalysts that promote non-natural transformations within living organisms with potential applications in bioremediation and carbon uptake. A collaboration with the University of Florida will test the generated model on qubits and spin-crossover.

This research is funded by the Center for Scalable Predictive methods for Excitations and Correlated phenomena (PNNL, FWP 70942) and the Center for Molecular Magnetic Quantum Materials (University of Florida, DE-SC0019330). A portion of the research is funded under the Laboratory Directed Research and Development Program at Pacific Northwest National Laboratory, a multiprogram national laboratory operated by Battelle for the U.S. Department of Energy.



**Title:** Exascale Computing for Energy Applications

**Principal Investigator:** Misun Min (Argonne National Laboratory)

**Co-investigators:** Paul Fischer (University of Illinois Urbana-Champaign / ANL),  
Elia Merzari (Penn State / 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):** 200,000 node-hours on Frontier,  
200,000 node-hours on Aurora,  
50,000 node-hours on Polaris,  
50,000 node-hours on Perlmutter-GPU

**Research Summary:**

Low-cost carbon-free energy is central to mitigating climate change and is a key driver in the overall mission of the Department of Energy. The goal of this project is to advance the DOE's simulation capabilities in important carbon-free energy sectors, including nuclear, fusion, and wind. In close collaboration with domain experts in each of these areas, we have identified important problems to be addressed with Argonne's multiphysics transport code, Nek5000/RS. These problems include (nuclear) simulations of fluid flow and heat removal through fuel, shield and reflector assemblies that have major impacts on operation and safety for sodium fast reactors; (fusion) high Reynolds/Hartmann-number magnetohydrodynamic (MHD) flows of liquid metal cooling that has been proposed in the blankets and diverters of fusion devices; (wind) high Reynolds-number boundary layer flows with sophisticated turbulence models for the atmosphere and for complex-geometry wind turbines.

The challenge problems will lead to enhanced simulation capabilities over a range of energy systems. All Nek5000/RS software is open source, which means that developments under this project will have a broad impact for challenging fluid-thermal/MHD applications, particularly those that require peta- or exascale resources. These problems will provide a vehicle for the test and development aspects of the project supported through DOE ASCR Applied Math Research and DOE's Nuclear Energy Advanced Modeling and Simulation program.



**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),  
Marius Millot (Lawrence Livermore National Laboratory),  
Sally Tracy (Carnegie Institution for Science)

**ALCC Allocation:**

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

**Allocation(s):** 900,000 node-hours on Frontier,  
600,000 node-hours on Aurora

**Research Summary:**

The historic December 5, 2022 experiment at Lawrence Livermore National Lab's National Ignition Facility reached fusion energy ignition for the first time, thus paving the way to future clean inertial fusion energy (IFE). This project addresses one of the priority research opportunities in IFE – the development of alternate ablative target materials, which are urgently sought to achieve high gain ( $G > 10$ ) by enhancing fuel compression in low-adiabat implosions.

The main objective of this ALCC project is to gain insight into the HED ablator physics of novel amorphous carbon (a-C) targets under compression by IFE drivers. Our scientific goals are: (1) to uncover the atomic-scale evolution of the a-C microstructure under shock compression, including potential phase transitions to diamond or metastable supercooled liquid carbon at high pressures and temperatures; (2) to accurately describe the complex states of carbon under conditions relevant to IFE and obtain high-quality EOS through a combination of experimental measurements and machine-learning molecular dynamics simulations. Our simulations at experimental time and length scales at DOE exascale supercomputers will guide experiments to observe predicted phenomena and validate our theoretical models.

This ALCC project will deliver key information on novel a-C ablators for target and implosion designs using a-C ablators. The transformative science impact of our work is in harnessing the unprecedented power of extreme-scale quantum-accurate MD simulations on exascale Frontier and Aurora to predict novel physical phenomena and guide experiments towards observing them.



**Title:** **A GPU-based multiscale ocean modeling system for real-time quantification of marine Carbon Dioxide Removal operations**

**Principal Investigator:** Raphael Ouillon (atdepth MRV)

**Co-investigators:** Carlos Munoz Royo (atdepth MRV), Thomas Peacock (atdepth MRV)

**ALCC Allocation:**

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

**Allocation(s):** 75,000 node-hours on Perlmutter-GPU

**Research Summary:**

The objective is to develop the Ocean Carbon Capture Model 1 (OCCM1), a GPU-based multi-scale simulation system that can capture near-field, regional, and global ocean and biogeochemical processes, with the goal of enabling scalable, cost-effective and high-certainty quantification of carbon dioxide removal in the ocean. At the core of OCCM1 will be two existing codes, (i) Oceananigans.jl, a cutting-edge, open-source fluid dynamics solver, and (ii) OceanBioME.jl, a state-of-the art biogeochemical modeling package that combines models for ecosystem dynamics, carbonate chemistry, individual-based models, and air-sea gas exchange. The project will implement core features to enable multiscale simulations of marine Carbon Dioxide Removal (mCDR) activities. The features will be integrated into OCCM1 to enable continuous, real-time simulations of pilot mCDR operations, characterizing the sources of uncertainty and demonstrating readiness for operational use.

Marine Carbon Dioxide Removal (mCDR) has the potential to scale rapidly and remove gigatons of CO<sub>2</sub> from the atmosphere, in an effort to mitigate climate change. Current approaches to quantifying carbon removal in the ocean via Measurement, Reporting and Verification (MRV), however, rely on legacy CPU-based ocean simulation codes that cannot resolve all the necessary scales, and are too slow and inefficient to offer a cost-effective and scalable solution. As a result, no science-based solution to quantify mCDR operations is currently available. This research project has the potential to underpin a paradigm shift in how numerical simulations can be used to quantify human activities in the ocean, by making real-time multiscale numerical simulations operational for both research and commercial applications. Applied to carbon removal, this can both accelerate scientific research on the risks and benefits of mCDR, and unlock cost-effective science-based MRV that enables high certainty ocean carbon credits.



**Title:** Atomic-scale design and characterization of sorbents for carbon capture

**Principal Investigator:** Jonathan Owens (GE Vernova Advanced Research Center)

**Co-investigators:** Bojun Feng (GE Vernova Advanced Research Center),  
Jerry Liu (GE Vernova Advanced Research Center),  
Xiaolei Shi (GE Vernova Advanced Research Center),  
Jinbo Cao (GE Vernova Advanced Research Center)

**ALCC Allocation:**

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

**Allocation(s):** 190,000 node-hours on Perlmutter-GPU

**Research Summary:**

High surface area, porous materials, like metal-organic frameworks (MOFs), zeolites, *etc.* are leading candidates for scalable and durable CO<sub>2</sub> sorbents and have shown great promise in this space. However, CO<sub>2</sub> molecules are present in ultra-dilute quantities in the atmosphere (~400 ppm), posing a challenge for sorbent materials, wherein they must selectively adsorb CO<sub>2</sub> over the other molecular species (e.g., H<sub>2</sub>O) with much higher concentrations. One prominent path to achieve this selectivity is through grafting chemical species that contain amine-groups that will selectively undergo chemical reactions with CO<sub>2</sub>, leaving other species relatively unaffected. There are large numbers of these potential amines that can be functionalized to a large variety of porous species, opening an enormous search space that is not explorable by experiment alone. In this project, we will leverage high accuracy first-principles methods, building upon the pipeline developed in our AY23 award, also on Perlmutter-GPU. These methods allow us to predict key quantities related to a hypothetical sorbent's performance, such as how much CO<sub>2</sub> it can adsorb, how quickly this process happens, and how structurally stable a material is predicted to be. In this allocation, however, we will be expanding the search space of our calculations to porous materials outside of MOFs, and additionally incorporating H<sub>2</sub>O and O<sub>2</sub> in our studies, to represent more realistic simulations of a sorbent's working environment. We will use the Leadership Computing Facility resources awarded on this project to perform these complex calculations.

Climate change is a major crisis of our time, and achieving a carbon negative future will require a multi-pronged approach, including electrification, clean energy generation, and carbon dioxide removal (CDR). CDR technologies are a key priority for the DOE and core component of the DOE's Carbon Negative Shot, the "Earthshot" related to CDR. This project will build upon our previously developed high-accuracy computational materials screening methods that are vital to CO<sub>2</sub> sorbent development, a task central to effective CDR technologies.



**Title:** High-fidelity combustor simulations to enable drop-in sustainable aviation fuels

**Principal Investigator:** Bruce Perry (National Renewable Energy Laboratory)

**Co-investigators:** Shashank Yellapantula (National Renewable Energy Laboratory),  
Sreejith Nadakkal Appukuttan (National Renewable Energy Laboratory),  
Marc Day (National Renewable Energy Laboratory),  
Mohammad Rahimi (National Renewable Energy Laboratory)

**ALCC Allocation:**

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

**Allocation(s):** 300,000 node-hours on Frontier,  
60,000 node-hours on Perlmutter-CPU

**Research Summary:**

Liquid jet fuel produced from biologically-derived sources, or Sustainable Aviation Fuels (SAF), are the most promising path to eliminating net carbon emissions in the aviation sector. However, due to rigorous certification requirements to ensure safety, the cost of testing can be prohibitively risky 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 candidates and de-risk scale-up. This research focuses on comparing a leading SAF candidate to petroleum-based Jet-A fuel in a NASA-funded open-source combustor design being studied at Georgia Tech for its potential to reduce pollutant emissions. The simulations will be performed with the Pele flow solvers, developed at NREL to efficiently utilize supercomputers. The project involves collaboration with experimental teams at Georgia Tech and NREL to provide combustor and fuel property measurements. The simulations will provide new insights into how SAF properties affect performance and pollutant formation in the latest aircraft combustor designs. They will also generate large data sets that can be used by outside researchers to train machine learning models for more rapid predictions in the future.

Recognizing the importance for decarbonizing the aviation sector, a federal-government-wide grand challenge has been established to expand production and use of SAF. Production of SAF is also a key component of the Department of Energy's Clean Fuels & Products Energy Earthshot. To achieve these goals, the end use of 100% drop-in SAF must be de-risked to enable certification and widespread adoption. This project is an important step toward these objectives. It will also provide insights into how SAF affects production of soot emissions, which play a significant role in the formation of contrails, another significant aspect of the climate impact of aviation.



**Title:** High-Fidelity Simulations of Open Fan Aeroacoustics Enabling Sustainable Aviation

**Principal Investigator:** Stephan Priebe (GE Aerospace Research)

**Co-investigators:** Trevor Wood (GE Aerospace Research),  
Kishore Ramakrishnan (GE Aerospace Research),  
Suryapratim Chakrabarti (GE Aerospace Research),  
Ramakrishnan Kannan (Oak Ridge National Laboratory)

**ALCC Allocation:**

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

**Allocation(s):** 477,000 node-hours on Frontier

**Research Summary:**

There is a societal need to eliminate use of fossil fuels associated with aircraft propulsion and achieve sustainable energy in flight. In 2021, GE Aerospace announced along with its CFM partner, Safran, the Revolutionary Innovation for Sustainable Engines (RISE) technology demonstration program to provide over 20% lower fuel consumption and CO<sub>2</sub> emissions compared to today's engines. Most of that benefit stems from the open fan architecture, in which the outer nacelle that encloses the fan in today's ducted turbofan engines is removed. Because open fans involve complex, fully 3D, and turbulent flow physics spanning a broad range of scales, from the smallest flow features near surfaces of highspeed airfoils to the coupled system interactions at an aircraft-level, the most computationally intensive simulation capability available today is required to understand open fan aerodynamics and acoustics. High-end, GPU-enabled supercomputing and application of Large Eddy Simulation (LES) at realistic conditions allows us to understand these complex flows and to optimize the open fan system aerodynamic performance as a steppingstone to sustainable flight while simultaneously reducing noise levels to meet future noise stringency goals.

In this ALCC program, we will execute wall-resolved LES with the aim of studying two important physical aspects of novel open fans: (1) broadband noise generated by the fan blade itself as well as through the interaction with a downstream stationary blade row; (2) the effect on the aerodynamic and aeroacoustic performance at high angle of attack as experienced by the rotating fan blades at takeoff. By applying machine learning techniques to the high-fidelity LES data, we will study improvements to lower-fidelity models. The research builds on prior work, which has investigated the aerodynamic and aeroacoustic performance of open fan blades for a range of high Reynolds numbers. This ALCC program will allow us to predict open fan noise generated by operation under angle of attack relevant to an open fan engine installed on an aircraft.



**Title:** Radiation studies for intensity frontier projects DUNE-LBNF, Mu2e and PIP-II

**Principal Investigator:** Igor Rakhno (Fermi National Accelerator Laboratory)

**Co-investigators:** Dali Georgobiani (Fermi National Accelerator Laboratory)

**ALCC Allocation:**

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

**Allocation(s):** 140,000 node-hours on Perlmutter-CPU

**Research Summary:**

The Deep Underground Neutrino Experiment and Long-Baseline Neutrino Facility (called DUNE-LBNF project) are under development at Fermilab since early 2010s. The activity represents a convergence of a substantial fraction of the worldwide neutrino physics community around the opportunity provided by the large investment planned by the U.S. Department of Energy (DOE). The primary scientific objective of DUNE is to carry out a comprehensive investigation of neutrino properties including conversion rules between neutrinos of different types. The LBNF will provide neutrino fluxes and detector infrastructure at the near site (Fermilab) and far site in South Dakota. Many of the LBNF project milestones are heavily dependent on thorough computer simulations. Significant work will be required that should address various components and issues such as current budget constraints. This work will be based on the existing optimized design with a 1.5-m graphite target and beam focusing systems consisting of three horns.

The Mu2e project at Fermilab – devoted to high precision measurements of muon-to-electron conversion without emission of neutrinos – is at its final stage of completion. The goal of the experiment is to identify physics beyond the Standard Model. Since the FY2022, this project has been identified as one of the top priorities for the lab’s mission. The beam commissioning has begun in FY2023 and will continue until end of 2024, the detector commissioning is planned to begin in FY2026. The transition of the project to operations requires studies of the radiation protection of the personnel and equipment, which involves large scale simulations of prompt and residual radiation levels in the beam lines and experimental enclosures.

The goal of the Proton Improvement Plan Phase II (PIP-II) project is to build a new superconducting linear accelerator which is expected to provide beam to the DUNE-LBNF and Mu2e experiments. It requires development of the final integrated radiological analysis and design which implies detailed and comprehensive computer modeling studies.

The planned work on studying properties of fundamental particles and discovering new physics laws aligns with the U.S. DOE mission to develop innovations in science and technology.





**Title:** Machine Learning Enabled Atomistic Simulation of Iron at Extreme Pressure

**Principal Investigator:** Robert Rudd (Lawrence Livermore National Laboratory)

**Co-investigators:** Federica Coppari (Lawrence Livermore National Laboratory)

**ALCC Allocation:**

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

**Allocation(s):** 150,000 node-hours on Polaris

**Research Summary:**

Iron is at the core of our planet and is thought to be at the core of countless exoplanets. Its behavior at the extreme pressures and temperatures of the planetary core determines much of the structure of the inner Earth. Its properties determine the size of the inner core, that region where the otherwise molten core is under such great pressure from gravity that the iron is solid. The transition from liquid to solid core is understood to affect the convective flows that give rise to the dynamo, generating the magnetic field that shields us from ionizing radiation from space, allowing life to form and thrive. Laboratory experiments that probe the properties of iron under such extreme conditions are difficult, but important strides have recently been made. Dynamic techniques, such as ramp compression, drive the iron to extreme conditions for a brief time, a few hundredths of a millionth of a second, and probe its properties. A direct determination of the temperature during that time has been elusive. Exciting new experiments using X-ray absorption measurements during laser-driven compression have produced excellent data but the analysis is incomplete. Additional atomic-scale information is needed to fully understand and characterize the competing effects of temperature and changes in microstructure induced by the rapid dynamic compression. Here we use advanced, machine-learned interatomic potentials in large molecular dynamics simulations of the iron ramp compression to provide that essential guidance. The machine-learned potential is very accurate, providing the accuracy of quantum-mechanical calculations at greatly reduced computational expense, allowing the multi-million atom simulations that are needed. These simulations still require extraordinary computational power, only possible with supercomputing resources.

Our simulations will provide a first-ever view into the changes to the microstructure of iron as it goes through phase transformation and extreme deformation, detail at the atomic level directly comparable to the experimental X-ray absorption data and allowing a deep understanding of iron's behavior under these extraordinary conditions. The outcome of the simulations will also more broadly impact the Department of Energy X-ray absorption experiments on other materials ramp compressed to high pressure, informing investigations of high-energy-density science and the exotic physics that occurs under extreme pressure.



**Title:** High-Fidelity CFD Enabling Advanced Nuclear Power

**Principal Investigator:** Dillon Shaver (Argonne National Laboratory)

**Co-investigators:** Haomin Yuan (Argonne National Laboratory),  
Aleksandr Obabko (Argonne National Laboratory),  
Elia Merzari (Pennsylvania State University),  
Jun Fang (Argonne National Laboratory)

**ALCC Allocation:**

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

**Allocation(s):** 200,000 node-hours on Frontier,  
100,000 node-hours on Aurora,  
50,000 node-hours on Polaris

**Research Summary:**

Advanced nuclear power relies on next-generation designs that go far beyond the currently operating light-water-cooled reactor (LWR) fleet. Every commercial nuclear power plant currently operating in the US is an LWR. While the use of water as a coolant has benefits – such as very well understood fluid properties and heat transfer behavior – there are more advanced designs which can achieve higher efficiency and greater safety margins. The advanced reactor designs proposed as the next generation of nuclear power require significant investigation into the proposed mechanisms of fluid flow and heat transfer to ensure that these claims are met to help achieve energy independence using low-carbon or carbon-free sources.

Computational Fluid Dynamics (CFD) is playing an increasingly important role in the design of advanced nuclear reactors. With the capabilities of the exascale computing era, full-core high-fidelity CFD simulations are now a reality offering greater insight into the fluid flow and heat transfer behavior of these reactors than ever before achievable. Using state-of-the-art computational tools developed by the Exascale Computing Project and the Nuclear Energy Advanced Modeling and Simulation Program on leadership-class computing facilities, CFD can provide a unique capability that has been consistently demonstrated to accurately predict turbulent fluid flow and heat transfer phenomena in a wide range of nuclear power applications.

High-fidelity simulation capabilities will be applied to analyze multiple advanced reactor designs and the generated data used to inform models implemented in fast-running tools in use by the nuclear industry, thereby accelerating the deployment of advanced nuclear power, including liquid metal cooled fast reactors, pebble bed reactors, and high-temperature gas-cooled reactors.



**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/ORNL),  
Monojoy Goswami (Oak Ridge National Laboratory),  
Arthur Ragauskas (The University of Tennessee, Knoxville),  
Hugh O'Neill (Oak Ridge National Laboratory)

**ALCC Allocation:**

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

**Allocation(s):** 48,384 node-hours on Frontier,  
145,920 node-hours on Perlmutter-GPU

**Research Summary:**

An essential component of the future circular bioeconomy will be the efficient deconstruction of plant matter to yield fuels and high value bioproducts. To achieve this a detailed multiscale understanding of the plant cell wall is required. However, no fundamental understanding of atomistic-scale interactions that drives multiscale organization (and the disassembly) of the large complex structures within plant cell walls structures exists at present. This project will develop both large-scale all-atom molecular dynamics and all-atom informed (bottom-up) coarse-grained models of the secondary plant-cell walls of three bioenergy species: Poplar, Pine, and Switchgrass. Additionally, we will develop models of two plant-cell wall derived polymer-blends (lignin-polyurethane mixtures and cellulose nanofiber impregnated polyethylene terephthalate matrices) to permit a comparative study of changes of the hierarchical interactions in naturally occurring matrices to those present in renewable-derived biomaterials.

We will develop models by leveraging the high-performance computing allocation provided by ALCC that will advance the conversion of abundant plant-biomass to renewable biofuels and performance of bio-derived materials. Indeed, by exploring the interplay of biopolymer reorganizations across several spatial (from Å to ~100nm) and temporal (femtosecond to microseconds) scales under both mechanical and chemical stresses and will aid in identifying critical molecular interactions responsible for the differences in the resistance of biomass sources to deconstruction and upgrading. Further, the modeling of lignin-polyurethane and cellulose-nanofiber composites will assist in providing molecular-based design principles to improve the properties of these renewable materials.



**Title:** High-fidelity modeling and next-generation surrogate models for floating offshore wind energy

**Principal Investigator:** Michael Sprague (National Renewable Energy Laboratory)

**Co-investigators:** Lawrence Cheung (Sandia National Laboratories),  
Stuart Slattery (Oak Ridge National Laboratory),  
Mark Taylor (Sandia National Laboratories),  
Ganesh Vijayakumar (National Renewable Energy Laboratory),  
Aaron Lattanzi (Lawrence Berkeley National Laboratory),  
Ann Almgren (Lawrence Berkeley National Laboratory)

**ALCC Allocation:**

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

**Allocation(s):** 475,000 node-hours on Frontier,  
35,000 node-hours on Perlmutter-CPU

**Research Summary:**

As part of a larger effort to decarbonize the grid, the U.S. Department of Energy's (DOE's) Floating Offshore Wind Shot seeks to reduce the levelized cost of energy (LCOE) of floating offshore wind by at least 70% by 2035. The Floating Offshore Wind Modeling and Simulation (FLOWMAS) Energy Earthshot Research Center (EERC), which supports this ALCC allocation, will enable the Wind Shot goal by addressing key science barriers to the profound engineering and technology advances needed. Building on years and tens of millions of dollars of DOE investments in exascale-class high-fidelity models and codes for climate and land-based wind, FLOWMAS researchers are creating a suite of high-fidelity codes for floating offshore wind energy that incorporate the microscale (floating turbines and farms), the mesoscale (regional weather dynamics), and the global/climate scale. FLOWMAS researchers will use results from high-fidelity simulations and from ongoing DOE-supported field campaigns to create data-driven surrogate models that are computationally efficient and able to explore parameter spaces and time horizons not accessible to high-fidelity models. Finally, the developed models will leverage exascale computing power to create a new understanding of the floating offshore wind energy system, including how climate change will impact offshore wind resources, the physics of floating wind farms and turbine wake dynamics, and the loads and dynamics of floating wind turbines in operational and nonoperational extreme events.



**Title:** Hierarchical, Scalable Green's Function Modeling of Chemistry at Interfaces

**Principal Investigator:** Mark van Schilfgaarde (National Renewable Energy Laboratory)

**Co-investigators:** Ross Larsen (National Renewable Energy Laboratory),  
Swagata Acharya (National Renewable Energy Laboratory),  
Mai-Anh Ha (National Renewable Energy Laboratory),  
Dimitar Pashov (King's College London)

**ALCC Allocation:**

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

**Allocation(s):** 610,000 node-hours on Frontier

**Research Summary:**

The ability to predict properties of chemical systems generally, without recourse to model or empirical parameters, is the “holy grail” of quantum chemistry. Since properties can be calculated only approximately, three large challenges present themselves: (1) high fidelity in the approximations made, (2) efficiency in obtaining a solution and scaling with system size, and (3) range of properties accessible to the theory. Quantum chemists have largely relied on wave-function ( $\psi$ ) methods for high fidelity and density-functional (DFT) methods when efficiency is needed. Green's function methods are a promising alternative (less developed in chemistry than in physics) because they lie between  $\psi$  and DFT methods. NREL has a CCS center funded by BES-CSGB to build a next-generation Green's function platform (Questchem, derived from the solid-state platform [Questaal](#)), with the aim of realizing accuracy approaching the best of the  $\psi$ -methods, while being far more efficient and scaling well with system size. If these goals are realized, many key outstanding questions involving surfaces and interfaces now out of reach can be addressed, including studies of catalysis, electrochemistry, and properties of 2D systems. Green's function methods have another key advantage: response functions form an integral part of the theory and yield information beyond ground-state properties. This is especially important since practical probes are often optical, e.g. fluorescence; absorption and Raman spectroscopies.

Green's function theory is usually implemented in the form of many-body perturbation theory. At the lowest level is the *GW* approximation. Questchem/Questaal has special, quasiparticle self-consistent form of *GW* (QSGW), which dramatically improves its fidelity compared to the usual (DFT-based) *GW*. It also largely eliminates the starting-point dependence (present in both  $\psi$  and GF methods) and crucially, errors are highly systematic, which clarifies which diagrams are still needed. One key finding of this project was to show that adding ladder diagrams to the polarizability dramatically the fidelity of excited states and optical properties. Even without the more efficient redesign now in progress, we have been able to pursue studies (particularly optics in 2D magnetic systems) that are not possible to accomplish --- at least with the same fidelity --- using any other available platform.



**Title:** Foundation Neuroscience AI Model-NeuroX

**Principal Investigator:** Shinjae Yoo (Brookhaven National Laboratory)

**Co-investigators:** Jiook Cha (Seoul National University),  
Thomas Flynn (Brookhaven 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):** 152,000 node-hours on Frontier,  
220,000 node-hours on Aurora,  
70,000 node-hours on Polaris,  
120,000 node-hours on Perlmutter-GPU

**Research Summary:**

The use of brain imaging and neuroscience in medicine is limited due to our current models' poor predictive power for assessing brain health and cognitive functions. The successful development of a foundational model has the potential to significantly advance brain imaging capabilities. This advancement can be realized by incorporating key features such as scalability, multimodality, expressiveness, compositionality, and an enhanced memory capacity. These improvements are in line with the technological advancements seen in Generative Pretrained Transformer (GPT) models, which have rapidly accelerated the evolution of conversational AI applications, including ChatGPT. This revolution is transforming our approaches to learning, communication, and decision-making, marking the beginning of a new era in technological innovation. With its successful deployment, our foundation model in neuroscience is anticipated to be revolutionary in various scientific disciplines, including neuroscience, medicine, and psychology. Designed for robust adaptability, the model will support a wide range of applications, such as predicting cognitive, psychological, clinical, and behavioral outcomes, and their potential developments over time. Enhanced task performance is expected to address major societal issues like depression, substance abuse, and suicide. Additionally, the model's influence is likely to permeate various commercial industries, enhancing communication, marketing strategies, Brain-Computer Interface (BCI) technologies, and brain decoding devices, leading to substantial progress in these domains. The proposed foundation model is also a generic spatio-temporal model and could be applicable to various DOE flagship science domains including climate, fusion, material science, particle physics, etc.



**Title:** High Fidelity Numerical Analysis on Flow and Heat Transfer behavior in Involute Plate Research Reactor to Support the Conversion Program

**Principal Investigator:** Yiqi Yu (Argonne National laboratory)

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

**ALCC Allocation:**

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

**Allocation(s):** 200,000 node-hours on Polaris,  
500,000 node-hours 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 elements share a similar configuration of coolant channel, which is of extremely thin and involute shape. Better understanding of flow behavior and heat transfer in these coolant channels is important for design of LEU fuel elements.

Reynolds-averaged Navier–Stokes (RANS) approaches have exposed both practicality and limitations on the prediction of turbulent flows. Discrepancies are seen in RANS simulations with different turbulence models, flow conditions and codes. Due to the lack of relevant experimental data, it is difficult to quantify the uncertainty introduced by these turbulence models. Performing high fidelity simulation (LES&DNS), uncovers behavior of turbulent flows with fewer modeling assumptions and can thus potentially provide more reliable benchmark for engineering predictions and thus more accurate thermal hydraulic safety analysis. Moreover, the data from the high fidelity simulations can be used to improve the coefficients used in system codes. LES data enrichment are essential for understanding the fundamental mechanism of flow and heat transfer behavior in an involute plate research reactor under various conditions.

The simulations in this project will be performed with both Nek5000 and NekRS. The 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 expands their application to the Research Reactor Conversion Program and attract more attention from either industry partner or university to NEAMS code as a surrogate for expensive commercial code.



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

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

**Co-investigators:** Franck Capello (Argonne National Laboratory)  
Bogdan Nicolae (Argonne National Laboratory)  
Paul Hovland (Argonne National Laboratory)

**ALCC Allocation:**

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

**Allocation(s):**

58,000 node-hours on Perlmutter-GPU

**Research Summary:**

Large foundation models for science will face the same challenges of state-of-the-art large language models (LLMs): the pre-training stage requires thousands of GPUs (graphic processing units) and takes a few months. This leads to a huge financial cost and increasing concerns about its environmental impact. Running such expensive pre-training on the exascale computing systems of DOE (Department of Energy) will also reduce the availability of computing resources for other mission-critical computing research. This project aims to enable energy-efficient training and inference of extreme-scale foundation models for science, reducing pre-training time by 10x and resource utilization (GPU-hours) by about 50x.

Using low-rank tensor-compressed computing, 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, this project will (1) develop theoretical understanding and novel optimization algorithms of low-rank tensor-compressed pre-training for large-scale foundation models, (2) accelerate pre-training via mixed-precision low-rank tensor optimization and customized tensorized automatic differentiation, (3) optimize GPU performance on leadership computing platforms for large-scale tensor-compressed pre-training. The pre-training framework will be evaluated via foundation models for science of various sizes.