

Title: AI Directed Adaptive Multiscale Simulations to Model RAS-RAF Cancer Initiation Pathway on Curved Membranes

Principal Investigator: Harsh Bhatia, Lawrence Livermore National Laboratory

Co-investigators: Helgi Ingolfsson (LLNL),
Arvind Ramanathan (ANL)

ALCC allocation:

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 740,000 node-hours on Summit

Research Summary:

Computational models can define the functional dynamics of complex systems in exceptional detail. However, phenomena of interest can be simulated either with high-resolution (microscopic) detail but at small spatiotemporal scales or with a coarser representation over large (macroscopic) length- and time-scales. Multiscale modeling and simulations, which aim to bridge this gap, are often key to reliable exploration of complex physical, chemical, and biological phenomena. We have previously demonstrated that unparalleled scaling of multiscale simulations can be achieved by (i) exploiting the intrinsic low dimensionality of the simulations to map the phase space of all configurations, while (ii) simultaneously using sophisticated machine learning (ML) techniques to computationally steer the exploration, which allows to (iii) judiciously utilize heterogeneous computational resources at scale. As part of an ongoing collaboration between the Department of Energy and the National Cancer Institute, we have developed Multiscale Machine-learned Modeling Infrastructure (MuMMI), which represents a new paradigm of multiscale simulations and enables exploring large length- and time-scales (microns and seconds, respectively) using macro-scale models, while simultaneously maintaining molecular-scale detail using a novel ML-based sampling framework. With the 2021-2022 ALCC allocation, this project will conduct the largest-of-its-kind simulation campaign on Summit – leveraging over 4.4M GPU hours to simulate over 20K coarse-grained and 5K all-atom simulations.

This project will study a complex biological phenomenon – the cancer initiation pathway mediated by the RAS and RAF oncogenes. RAS and RAF are important links in the cell's signaling system and nearly a third of all cancers diagnosed in the U.S. are driven by mutations in RAS genes and their protein products. Aberrant RAS behavior accounts for a particularly high percentage of pancreatic (~95%), colorectal (~45%), and lung (~35%) cancers. Whereas previous work focused on exploring the role of RAS and RAF lipid interactions and aggregation on an average plasma membrane model, this project will extend MuMMI to include variants of the RAS protein and decipher complex spatiotemporal aspects of initial triggering of RAS-RAF signaling with better controlled and more-relevant membrane conditions – by using a novel grand canonical scheme and titrating phosphatidylinositol lipid concentrations. Now equipped with three scales (continuum, coarse-grained, and all-atom) that model RAS-RAF dynamics at increasing complexity; these scales are coupled using ML-based decision making to automatically refine from coarser to finer, where needed. By comprehensively understanding the mechanisms of RAS-RAF interaction and their dependence on local membrane composition, we will provide quantitative insights into the RAS-RAF complex, which may provide rationale for new drug design strategies.

Title: Hadronic contributions to the muon $g-2$ from lattice QCD using chiral fermions

Principal Investigator: Thomas Blum, University of Connecticut

Co-investigators: Peter Boyle (BNL, University of Edinburgh),
Norman Christ (Columbia University),
Taku Izubuchi (BNL),
Luchang Jin (University of Connecticut),
Chulwoo Jung (BNL),
Christoph Lehner (Regensburg University, BNL)

ALCC allocation:

Site: Oak Ridge Leadership Computing Facility

Allocation: 610,000 node-hours on Summit

Research Summary:

The RBC/UKQCD collaboration is a collaborative effort involving physicists from several universities and laboratories in the United States and United Kingdom. The collaboration aims to compute hadronic contributions to the muon anomalous magnetic moment from first principles using lattice quantum chromodynamics (QCD). QCD is the fundamental theory of the interactions of quarks and gluons, and “lattice” refers to a discrete version of the theory that lives on a regular grid of points with a fixed spacing between them. These computations are crucial for a precision comparison between the Standard Model of particle physics and the new muon $g-2$ experiment at Fermilab (and a later one at J-PARC in Japan) and the previous one at Brookhaven National Laboratory (BNL). Fermilab recently announced their first results, and combined with BNL, they strengthen the current discrepancy with the Standard Model, which could point to new physical laws of Nature and/or never before observed particles. The experiment and theory value are top priorities for the DOE Office of High Energy Physics, and together they remain the community's best chance to discover new physics in decades.

The leading and next-leading hadronic contributions, which dominate the Standard Model theory errors, arise from the hadronic vacuum polarization (HVP) and hadronic light-by-light (HLbL) scattering, respectively. The research plans to continue calculations of these quantities in the framework of lattice QCD. The main goals of the HVP project are to improve the continuum limit with new calculations on a large r lattice and to reduce statistical errors that are intrinsic to Monte Carlo integration methods like lattice QCD. Lattice QCD results have now reached a similar precision as dispersive results but tensions between different lattice QCD calculations as well as between lattice QCD and dispersion theory have emerged. In a high-stakes test of the Standard Model, it is crucial to have precise, consistent results, using both methodologies, and the project calculations aim to clarify these issues. For the HLbL project the main goal is to reduce the dominant statistical errors. If the new calculations of the hadronic contributions uphold the Standard Model discrepancy with experiment, the impact on our understanding of Nature will be hard to overstate. In the alternative case where the Standard Model prevails, the impact on the field will still be significant.

Title: Achieving higher efficiency turbomachinery design via large eddy simulation

Principal Investigator: Sanjeeb Bose (Cascade Technologies/Stanford University)

Co-investigators: Adrian Lozano-Duran (MIT),
Parviz Moin (Stanford University)

ALCC allocation:

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 100,000 node-hours on Summit

Research Summary:

Recent advances in numerical methods, fast and efficient unstructured flow solvers, and developments in near-wall modeling have demonstrated that higher fidelity flow simulation techniques can tractably predict sensitive flow features that govern operability limits in realistic engineering flows. Leveraging high performance computing and modern architectures, turnaround times of less than a day are now possible for the prediction of complex flows, such as the maximum lift of a realistic aircraft model. This project aims to leverage these capabilities to aid prediction of two pacing items critical to the design of higher efficiency turbomachinery: predicting peak adiabatic efficiency/stall margins in compressor flows and heat transfer characteristics on turbine blades with active cooling. The accuracy requirements for both of these flow regimes are extremely stringent. Differences in adiabatic efficiencies of a few percent or limited operability in a compressor stage can render a new design infeasible, and inability to accurately characterize the heat transfer on turbine blades can limit overall thermal efficiencies (due to reduced operating temperatures/pressures). Two canonical problems with detailed experimental measurements have been selected to assess present predictive capabilities. For the prediction of adiabatic efficiencies and stall margins in compressor flows, wall-modeled and wall-resolved large-eddy simulations (LES) of the NASA Rotor 37 at 60% and 100% design speeds will be conducted. The selected off-design speed line has total pressure ratios consistent with novel energy storage devices (e.g., Brayton batteries) and the peak adiabatic efficiency approaches 95%. Similar wall-modeled and wall-resolved LES simulations will also be conducted to predict surface heat transfer on an effusion cooled turbine blade, experimentally characterized at Pennsylvania State University. The project includes tight collaboration from experimental groups at MIT (Gretizer and Spakovszky) and Penn State (Thole).

Title: Advancing Watershed System Science using Machine Learning and Extreme-Scale Simulation

Principal Investigator: Ethan Coon, Oak Ridge National Laboratory

Co-investigators: J. David Moulton (Los Alamos National Laboratory),
Scott Painter (Oak Ridge National Laboratory),
Carl Steefel (Lawrence Berkeley National Laboratory)

ALCC allocation:

Site: National Energy Research Scientific Computing Center (NERSC)

Allocation: 270,000 node-hours on Cori

Research Summary:

As Earth systems change places, increasing stress on the Nation's water resources, we can no longer rely on historical hydrological trends and simplified models to guide our understanding of the critical watershed functions that control downstream water supply and quality. Funded by a new Office of Science BER project, ExaSheds, this research aims to develop and test a hybrid Machine Learning (ML) and extreme-scale, process-based simulation workflow for predicting watershed function at hyperresolution (<100m) on extents approaching full river basins (e.g. the Upper Colorado River Basin). ML methods have been used to augment process modeling in several crucial ways that, for the first time, make watershed models at these scales sufficiently well-parameterized and performant to execute simulations on climate timescales (e.g. 100 years) to understand the impact of climate change on watersheds. The requested allocation will enable the continued development, training, and evaluation of this capability, and will result in hyperresolution simulations at climatological timescales on a full river basin – a first in the hydrologic modeling community.

The project involves developing a modeling capability that combines ML approaches and an existing High Performance Computing code, Amanzi – ATS, to simulate and predict watershed function including water quantity and quality in major river basins. Specifically, the research continues exploration of the use of ML and process-based models intertwined via:

1. High-resolution, downscaled precipitation and other meteorological forcing datasets that leverage Earth System Model projections (at coarse scales) and local information (at fine scales) to train an ML based downscaling approach. This downscaled dataset is then used to force high-resolution simulations.
2. Hybrid models that use Amanzi – ATS simulations as one type of input to an ML model to predict streamflow, effectively using the ML as a correction to the model.
3. ML-based alternatives to inverse models which use hindcast ensembles to efficiently calibrate the fields of parameters (e.g. permeability, porosity, soil structure) needed in high-resolution hydrologic models. The calibrated model is then used in projections.

Each of these methodological advances have been developed within ExaSheds on small simulation domains, and initial papers on these topics have been published or are in review. Supporting and evaluating these approaches at scale requires suites of simulations. Within the ExaSheds project, both the Gunnison River (within the Upper Colorado Basin) and the Neversink River (within the Delaware River Basin) have been targeted as crucial testbeds. A shared simulation plan ensures efficient use of resources to support each of the above scientific advances, and it will enable a first-of-a-kind simulation product predicting full river basin watershed function on climatological timescales at hyperresolution.

Title: Simulating Extreme Stellar Death: Magnetorotational Supernovae and Black Hole Formation

Principal Investigator: Sean Couch, Michigan State University

Co-investigators: Evan O'Connor (Stockholm University)

ALCC allocation:

Site: National Energy Research Scientific Computing Center (NERSC)

Allocation: 1,600,000 node-hours on Cori

Research Summary:

Core-collapse supernovae (CCSNe) are the most extreme laboratories for nuclear physics in the universe. Stellar core collapse and the violent explosions that follow give birth to neutron stars and black holes, and in the process synthesize most of the elements heavier than helium throughout the universe. The behavior of matter at supranuclear densities is crucial to the CCSN mechanism, as are strong and weak interactions. Beyond Standard Model behavior of neutrinos may also impact the CCSN mechanism. CCSNe play a key role in many aspects of astrophysics. After decades of research effort, recent progress in CCSN theory shows that the turbulence-aided delayed neutrino heating mechanism is successful at driving explosions in a broad range of progenitor stars. With this crucial step in our understanding, it is critical that we continue to refine our theory of the CCSN mechanism by making direct comparison to observations and by continuing to improve the fidelity of our simulations through the application of improved physics models and realism in the assumed initial conditions. A predictive model for the CCSN explosions is now within reach.

This project will carry out an investigation of CCSNe including the effects of rotation, magnetic fields, and progenitor asphericity. The comprehensive research program will consist of 3D magnetohydrodynamic (MHD) CCSN simulations with sophisticated multi-dimensional neutrino transport, the most realistic initial conditions ever adopted for the study of CCSNe, and an intensive comparison to observations. The objectives of this project will be achievable by leveraging the unique combination of skills in the project team, cutting-edge open-source software, and the Leadership-class resources available through the ALCC program.

This project will advance our understanding of the CCSN mechanism using realistic initial conditions, including rotation and magnetic fields. This project will develop and employ massive stellar progenitor models at the point of core-collapse that are evolved including rotation and magnetic fields. The critically important questions of how rotation and magnetic fields affect CCSN explosions and resulting observables across a range of initial progenitor structures will be directly addressed. The results will directly inform our understanding of the characteristics of newborn pulsars and magnetars, information that can be directly compared to observational data. This effort will also simulate proto-neutron star collapse to black hole in 3D including rotation and magnetic fields.

This project will address several critical questions: How do plausible rotation rates and magnetic field strengths influence the CCSN mechanism and stellar-mass black hole formation? What is the impact of realistic 3D progenitor structure including rotation and magnetic fields on the CCSN mechanism and observables? What are the resulting nucleosynthetic products from extreme stellar core collapse events? And what constraints can we place on the equation of state of nuclear matter through combining theory and observation of core-collapse supernovae? These questions are central to the scientific aims of the DOE Nuclear Physics program.

Title: Modeling of Polymeric Materials for Energy Storage Across Scales

Principal Investigator: Juan de Pablo, University of Chicago

Co-investigators: Elizabeth M.Y. Lee (University of Chicago),
Joshua Mysona (University of Chicago),
Ludwig Schneider (University of Chicago)

ALCC allocation:

Site: Argonne Leadership Computing Facility (ALCF)

Allocation: 45,000 node-hours on Theta GPU and early access to the Polaris Testbed

Research Summary:

Lithium-ion batteries are widely used for energy storage in electronic devices due to their high energy density. Current batteries, however, face challenges regarding long-term chemical stability and limited energy density, which are both necessary for energy-intensive applications. Addressing these challenges is key to developing next-generation energy storage technologies, and it has been identified as a priority by the Department of Energy (DOE). A promising strategy in lithium-ion battery research is to use lithium metal anode and liquid electrolytes to achieve higher energy densities, but the growth of dendrites continues to pose limits to the stability of such systems. A potential remedy for these problems is to use solid polymer electrolytes (SPEs) instead of liquid electrolytes.

This project uses a multi-scale simulation approach to gain insights into the morphology of two distinct interfaces (polymer/inorganic and polymer/polymer interfaces) of solid polymer electrolytes, and their influence on overall ion conductivity. The multi-scale approach is necessary because, in the lithium battery system, multiple distinct processes occur that are separated by orders of magnitude in both length and time, and they all contribute to overall battery performance. The first goal of the project is to identify chemical motifs in polymers that are thermodynamically stable in the vicinity of the inorganic solid phase, and that would enable fast lithium-ion transport across the interface. The second goal is to characterize the polymer/polymer interfaces that are found in microphase separated materials, including copolymers, using chemistries identified under the first goal. To connect each element of our multiscale approach, the individual pieces have to push the limits of what is currently possible with modern supercomputers. With modern supercomputers, state-of-the-art algorithms, and the combination of three different levels of description, it will be possible to simulate real engineering problems without neglecting fundamental effects that arise from the atomistic and quantum nature of the problem.

An important aspect of the project will be to develop new methods and augment software tools that will enable simulation of multiscale materials design problems, and which will be disseminated throughout the scientific and engineering communities.

Title: Precision Calculation of the Anomalous Magnetic Moment of the Muon with Highly Improved Staggered Quarks

Principal Investigator: Carleton DeTar, University of Utah

Co-investigators: Aida El Khadra (U Illinois U-C),
Steven Gottlieb (Indiana U),
Ethan Neil (U Colorado),
Ruth van de Water (Fermilab)

ALCC allocation:

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 300,000 node-hours on Summit

Research Summary:

A major goal of contemporary high energy physics is to discover new fundamental particles and interactions. Aside from looking directly for them in high energy accelerators, new physical processes will give rise to discrepancies between experimental measurements of various quantities and predictions based on the physics that we know. One such quantity of great current interest is the magnetic moment of the elementary particle called the “muon”. There is a significant discrepancy at the level of approximately four standard deviations between the experimental value and theoretical predictions. However, to claim evidence for new physics, one would need a still stronger disagreement. The discrepancy has led to a major effort to improve both the experimental measurement and the theoretical prediction. This project will sharpen our high-precision theoretical calculation of the anomalous magnetic moment of the muon using numerical lattice quantum chromodynamics (QCD) on Summit, the flagship computer at the Oak Ridge Leadership Computing Facility (OLCF). Timely results are needed for comparison with results of a new high-precision experiment (Muon g-2) currently underway at Fermilab.

This project falls squarely within the Department of Energy mission. The most recent Particle Physics Project Prioritization Panel (P5) report “Building for Discovery: Strategic Plan for U.S. Particle Physics in the Global Context” defines several science drivers including “Explore the Unknown: New Particles, Interactions, and Physical Principles.” One of the key opportunities highlighted in the report is study of the anomalous magnetic moment of the muon. First Fermilab Muon g-2 results were announced in April this year, and higher-precision results are expected very soon. The new result is still 4.2 standard deviations from the theoretical prediction based on the Standard Model of elementary particle physics. The Muon g-2 experimental group expects to improve the measured precision by a factor of four in the coming year or two. Another experiment is planned at JPARC in Japan. It is essential to obtain a commensurate improvement in the theoretical prediction to tell if the current tension between theory and experiment rises to a level of significance sufficient to consider it evidence for new physics beyond the Standard Model.

Title: Particle heating and acceleration in laboratory collisionless shocks

Principal Investigator: Frederico Fiuza, SLAC National Accelerator Laboratory

Co-investigators: Arno Vanthieghem (Stanford University)

ALCC allocation:

Site: National Energy Research Scientific Computing Center (NERSC)

Allocation: 820,000 node-hours on Cori

Research Summary:

Astrophysical collisionless shocks are among the most powerful particle accelerators in the Universe. Generated by violent interactions of supersonic plasma flows with the ambient medium, shock waves are observed to amplify magnetic fields and to accelerate electrons and ions to relativistic speeds. Recent developments in high-power lasers are now opening the opportunity to create high-energy-density (HED) plasma states in the laboratory that enable the experimental study of the physics of collisionless shocks in conditions relevant to astrophysical environments. The goal of this ASCR Leadership Computing Challenge project is to perform large-scale first-principles simulations of HED laboratory and astrophysical plasmas to identify the dominant mechanisms responsible for particle heating and particle acceleration in collisionless shocks. A transformative advance in our understanding of this long-standing scientific challenge requires the combination of three-dimensional, high dynamic range, fully kinetic simulations and controlled laboratory experiments where numerical findings can be tested and used to improve theoretical models, which is at the core of this project.

The simulations to be developed will be critical to design and support Discovery Science experiments at the National Ignition Facility (NIF) that will probe particle acceleration in shocks. The comparison between simulations of astrophysical and laboratory conditions will help identify and establish clearly the equivalence between the experimental findings and the physics governing astrophysical shocks. The fundamental understanding of particle acceleration in plasmas to be provided by this project is central to DOE's mission in Discovery Plasma Science. The results of this research are expected to have a significant impact on unveiling long-standing questions behind cosmic plasma accelerators, in advancing the understanding of interpenetrating HED plasmas, and in generating new ideas for efficient laboratory accelerators. Finally, the tight connection between the proposed simulations and experimental programs on NIF will also enable the important benchmark of widely used numerical plasma models in extreme HED conditions of relevance to DOE programs.

Title: Predicting the performance of heat-resistant alloys in extreme environments

Principal Investigator: Michael Gao, National Energy Technology Laboratory

Co-investigators: Laurent Capolungo (Los Alamos National Laboratory)

ALCC allocation:

Site: National Energy Research Scientific Computing Center (NERSC)

Allocation: 730,000 node-hours on Cori

Research Summary:

Cost-effective heat-resistant alloys that can withstand extreme mechanical and corrosive environments are crucial to enabling advanced energy systems. This project introduces a disruptive integrated strategy for both the design and lifetime assessments of structural metals used for power applications. The goal is to accelerate the discovery of cost-efficient alloys and other heat resistant materials capable of sustaining high temperatures, oxidizing environments, and with a lifetime above 100,000 hours. This project combines the strength of multi-national labs to predict creep performance of model high-temperature austenitic alloys using multiscale modeling approaches. The approach is to develop precise physics-based models of deformation and failure over many time and length scales, and then use these models to understand the relationship between failure and deformation to inform lifetime assessments in harsh fossil fuel environments. Additional effort is dedicated to the design and development of lightweight, low-cost, high-performance refractory high-entropy alloys for ultrahigh temperature applications over 1300°C such as gas turbines by predicting thermodynamic, kinetic, environmental, and mechanical properties of virtual alloys using multiscale modeling and machine learning.

The mission of the US. Department of Energy's Office of Fossil Energy High Performance Materials program is to develop materials suitable for extreme environments found in coal power generation to support existing and new plants. This project directly supports this mission by significantly reducing the time and lowering the cost to develop new high-performance heat-resistant alloys by using multiscale modeling and machine learning. Establishing the predictive capabilities of heat-resistant alloys subjected to real-world fossil power plant conditions will lay the foundation for life prediction for the existing and future fossil power fleet. The successful development of cost-effective refractory high-entropy alloys will not only improve the efficiency of gas turbines but also bring positive impacts on energy savings, carbon emissions, and the economy of wide arrays of industries and sectors.

Title: First principles calculations of dislocation core energetics in dilute Mg alloys

Principal Investigator: Vikram Gavini, University of Michigan

Co-investigators: Sambit Das (University of Michigan)

ALCC allocation:

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 130,000 node-hours on Summit

Research Summary:

Novel light-weight structural materials play an important role in reducing energy consumption and carbon footprint in automotive and aerospace sectors. It is expected that every 10% reduction in the weight of a vehicle will result in a 6-8% increase in the fuel efficiency. Magnesium (Mg), the lightest structural metal with a high strength to weight ratio, is an ideal candidate, but its technological usage has been severely limited by the lack of ductility of Mg. Ductility in Mg is governed, in part, by the energetics of dislocations—line defects in crystalline materials—as well as their interaction with solutes in Mg alloys. It has been suggested that the ductility of Mg can be enhanced by adding very small amounts of favourable transition-metal solute elements such as Yttrium, Zirconium (Zr) and Cerium (Ce). The objectives of this research are to conduct fundamental studies on energetics of pyramidal dislocations in Mg and their interaction with solutes in Mg alloys using large-scale Kohn-Sham density function theory (DFT) calculations. This research will leverage the significant computational, algorithmic and high-performance computing advances made via the development of the DFT-FE code—a massively parallel large-scale DFT code based on finite-element discretization. DFT-FE has enabled fast and scalable DFT calculations on large-scale systems with GPU acceleration. Over the past year, ongoing research efforts have focused on conducting large-scale Kohn-Sham DFT calculations on the Summit supercomputer to compute a sensitive parameter informing the ductility of Mg alloys—the energy difference of pyramidal I and II dislocations (ΔE_{I-II}). However, two critical DFT inputs still remain. The first research objective is to quantify the dependence of ΔE_{I-II} on macroscopic stresses (resulting from external loading), which can have an appreciable influence on the ductility predictions. The second research objective is focused on studying the dislocation-solute interaction energies between pyramidal dislocations and two emerging transition metal solutes, Zr and Ce.

Obtaining, accurately, the aforementioned energetics—energetics of Mg pyramidal dislocations and dislocation-solute interaction energetics for technologically important transition metal solutes—from large-scale DFT calculations will improve the quantitative accuracy of the ductility models for Mg. This will further enable a throughput search for the best solute combinations and the concentration ranges to significantly improve the ductility of Mg. This in turn has the potential for advancing light-weight multi-component Mg alloys, which has significant technological implications ranging from economic savings via improved fuel efficiency and reducing the carbon footprint in automotive and aerospace sectors.

Title: Inverse design of multicomponent oxide catalysts with generative models and DFT

Principal Investigator: Rafael Gomez-Bombarelli, MIT DMSE

Co-investigators: Yang-Shao Horn (MIT MechE)
Livia Giordano (MIT MechE),
Ben Blaiszik (ANL)

ALCC allocation:

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

Allocation: 50,000 node-hours on Summit
400,000 node-hours on Theta
200,000 node-hours on Cori

Research Summary:

Promising technologies such as solar fuels, fuel cells, electrolyzers, and metal-oxygen batteries are key for sustainable and independent energy generation and storage. However, they are being held back by a common challenge: the need to interconvert electricity and chemical energy in the formation and consumption of oxygen molecules. We lack efficient and cheap electrocatalytic materials for the so-called oxygen evolution reaction and oxygen reduction reaction (OER and ORR). Precious metals are effective but have limited practical applications because of their high cost and low abundance. Non-platinum group metal oxides such as Ni, Fe, and Co are active for OER and ORR in alkaline solutions, but are not stable in the acidic environment encountered in water electrolyzers or have low activity. Multi-component oxide materials can combine diverse earth-abundant elements to achieve the desired properties. Perovskites, for example, have an ABO_3 composition and have shown much promise as OER/ORR catalysts because they can be doped with a rich variety of A and B site cations, and these changes in elemental composition allow fine-tuning of the electrocatalytic properties. Despite the promise of these materials, and the success of atomistic simulations in explaining and predicting catalytic activity of surfaces, the large combinatorial space of substitutions hinders design of new materials. There are billions of 5-element combinations and, even with computers, they cannot be evaluated one by one in search for the best. This project will build an autonomous computational active-learning loop to perform global optimization of multicomponent oxide catalysts for OER/ORR. Combining machine learning and atomistic simulations through active learning, this project will explore and rank a very large composition space of multicomponent oxides according to their stability and activity. The predictions of these models will then be validated in the lab and scaled up through collaborations with academics and industry.

The project will utilize atomistic simulations based on quantum mechanics to rank candidate materials for their stability and their catalytic properties. Machine learning models will play two important roles. One in replacing the expensive simulations with fast surrogate functions that reproduce the accuracy at a fraction of the cost. Neural network architectures fine-tuned for materials will be trained on the continuously acquired data. The second use of machine learning is inspired in the recent successes in games like chess or go, where machines learn to play a game based on discrete decisions. In our case, the game is designing a material, and the discrete “moves” are selecting which elements to use in which positions in the material to get stable and active catalysts. By rapidly exhausting a large design space, this project will generate highly promising and diverse materials for OER/ORR that will then be validated experimentally and scaled up towards devices and products. The project will also produce a complete “map” of the chemical space of oxides for OER/ORR which can guide further material searches for other applications.

Title: Next-generation scalable deep learning for medical natural language processing

Principal Investigator: John Gounley, Oak Ridge National Laboratory

Co-investigators: Andrew Blanchard (ORNL),
Mayanka Chandra Shekar (ORNL),
Chris Forster (Nvidia),
Shang Gao (ORNL),
Josh Romero (Nvidia),
Junqi Yin (ORNL),
Hong-Jun Yoon (ORNL),
Todd Young (ORNL)

ALCC allocation:

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 130,000 node-hours on Summit

Research Summary:

The Joint Design of Advanced Computing Solutions for Cancer (JDACS4C) program was established as a partnership between the Department of Energy (DOE) and the National Cancer Institute (NCI) to accelerate cancer research by integrating high performance computing and deep learning approaches. Within JDACS4C, the Pilot 3 project is a collaboration between DOE and the NCI Surveillance, Epidemiology, and End Results (SEER) program focused on developing deep learning approaches to perform computational phenotyping, in which biological and clinical markers of cancer are extracted from population-level clinical text and medical image data. By learning high-level abstractions, deep learning models have the potential to significantly improve accuracy for patient phenotyping versus alternative automated methods. Developing these efficient and accurate phenotypic algorithms for cancer patient stratification is critical for advancing precision medicine in oncology. Deep learning models based on the Transformer self-attention mechanism, such as Bidirectional Encoder Representations from Transformers (BERT), have shown enormous promise for natural language processing and image processing. However, the applicability of generic pre-trained Transformers to computational phenotyping in the domain of clinical text and medical images is limited by the pretraining having been conducted on datasets and with learning tasks that are not representative of the clinical and biomedical domain. This limitation can be addressed by pre-training Transformer models from scratch for clinical text and medical images, which is computationally intensive but also highly scalable. The objective of this project is to use Summit at the Oak Ridge Leadership Computing Facility to improve computational phenotyping by pre-training a series of Transformer-based deep learning models from scratch with domain-specific datasets for clinical text and medical images.

This project will have transformative implications for computational phenotyping, where scalable deep learning methods hold substantial promise and provide novel approaches to addressing key constraints in current implementations. While contextual embedding approaches such as BERT are now the dominant state-of-the-art across a variety of natural language processing and image processing tasks, many latest advances have yet to be applied into the space of clinical, biomedical, and scientific data. Therefore, while the aims of the project focus on the DOE's partnership with NCI's population-level cancer surveillance program, the results of the project hold potential not just for the broader medical records space but also for text and image data in other domains, such as national security.

Title: Improving Shallow Clouds in a Multiscale Earth System Model

Principal Investigator: Walter Hannah, Lawrence Livermore National Laboratory

Co-investigators: Mark Taylor (Sandia National Labs),
David Bader (Lawrence Livermore National Lab),
L. Ruby Leung (Pacific Northwest National Labs),
Matthew Norman (Oak Ridge National Lab),
Kyle Pressel (Pacific Northwest National Labs)

ALCC allocation:

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 150,000 node-hours on Summit

Research Summary:

This research is in support of the Energy Exascale Earth System Model (E3SM) project, a multilaboratory project developing a leading-edge Earth system model to address U.S. Department of Energy (DOE) mission needs. E3SM development specifically targets current and future resources provided by DOE Leadership Computing Facilities. This is consistent with DOE's historical and ongoing role as the leader in adoption of disruptive new computational architectures to support a range of scientific endeavors. E3SM is a collaborative effort across eight national laboratories and twelve academic institutions. E3SM was publicly released in May 2018 and is now following an open source / open development model with all development work visible at <http://e3sm.org>.

With the appearance of machines like Summit, it has become apparent that Earth system models need to adapt to the coming changes in computational architecture, specifically GPUs. Most Earth system models involve frequent inter-process communication that scales with the problem size and makes it difficult to effectively leverage GPUs. One way to overcome this issue, while simultaneously improving the simulated clouds, is to replace the traditional cloud parameterizations with a fine-scale model embedded in each column of the global grid. This approach is known as a multiscale modelling framework (MMF) and can dramatically increase the workload of the model without changing the amount of the interprocess communication, which makes the MMF a good application for GPU acceleration.

After rigorous development and rewriting the code in C++, E3SM-MMF is the first Earth system model to effectively utilize Summit GPUs, allowing it to achieve throughput comparable to the standard E3SM. This makes E3SM-MMF much more attractive for longer climate experiments, but additional model limitations need to be addressed. The current schemes for microphysics and turbulence in E3SM-MMF are limited in their capacity to realistically represent clouds. The current microphysics scheme only predicts hydrometeor mass (i.e. single moment), which is problematic for simulating realistic cloud ice processes and aerosol-cloud interactions. The current turbulence scheme also lacks representation of higher order moments that are critically important for shallow convection. Both the microphysics and turbulence schemes impact the accuracy of climate sensitivity estimates with E3SM-MMF, which limits how the model can be used for climate experiments.

The goal of this ALCC allocation is to improve shallow clouds in E3SM-MMF by implementing new microphysics and turbulence schemes and evaluating their effects on the simulated climate. The planned schemes have been implemented and validated in other models, providing previous experience on which to build. The implementation will require special attention to performance concerns to ensure that E3SM-MMF continues to run with a reasonable throughput. After the implementation and testing is complete, decadal experiments will be conducted to assess the simulated climate and quantify the improvements.

Title: Multimodal Imaging with Intense X-ray Pulses

Principal Investigator: Phay Ho, Argonne National Laboratory

Co-investigators: Christopher Knight (Argonne National Laboratory), Adam Fouda (Argonne National Laboratory)

ALCC allocation:

Site: Argonne Leadership Computing Facility (ALCF)

Allocation: 316,000 node-hours on Theta

Research Summary:

Intense x-rays provide extraordinary spatial resolution and elemental specificity to enable tools that can resolve the dynamics of atoms and electrons. X-ray free-electron laser (XFEL) pulses, such as those recently available at the upgraded Linac Coherent Light Source (LCLS-II), Spring-8 Angstrom Compact Free Electron Laser (SACLA) facilities, the European-XFEL, Swiss-FEL and Korean PAL-XFEL provide unprecedented intensity pulses with femtosecond and shorter duration. These pulses can potentially enable the generation of real-time “movies” that can follow, in three dimensions, key ultrafast processes during chemical reactions in liquids, advanced materials, and biological systems. As new multimodal experimental techniques are being developed rapidly, concurrent theoretical development, as proposed here, will enable significant advances in imaging and the study of dynamical processes of materials. For 3D flash imaging with atomic resolution and elemental contrast, it is critical to fully understand the fundamental interactions of these intense x-ray pulses with nanosized heterogeneous systems (e.g., inner-shell ionization, relaxation, charge transfer and electron-ion recombination). A thorough understanding of XFEL pulse-mediated electronic transitions and the ensuing response of the environment serves not only to realize the LCLS Single Particle Imaging Initiative, but also guide the design of new, unique experiments and light source facilities. The planned research work is an important step forward in understanding high-brightness, high energy, coherent x-ray laser pulses and their interactions with materials, supporting the Department of Energy’s (DOE) mission to extend the scientific and technological strengths of the U.S.

This project will examine multimodal atomic imaging approaches enabled by the most intense femtosecond and attosecond XFEL pulses. In particular, this project will investigate and control resonant scattering processes for gain in signals and imaging resolution with intense femtosecond and attosecond x-ray pulses. To complement the scattering approach, the planned simulation campaigns will explore the potential of x-ray correlation methods by investigating fluorescence intensity correlations, based on the principle of Hanbury Brown and Twiss effect, for imaging structure and elemental contrast in heterogeneous systems. The challenges associated with tracking the motion of particles and evolution of electronic configurations are addressed with a novel Monte-Carlo/Molecular-Dynamics simulation algorithm implemented in the highly parallelized simulation code LAMMPS as required for the large multimillion particle systems under experimental study. The quantum nature of the initiating ionization process is accounted for by a Monte Carlo method to calculate probabilities of electronic transitions and thus track the transient electronic configurations explicitly. The freed electrons and ions are followed by classical particle trajectories using a molecular dynamics algorithm. The combination of this novel simulation method with leadership computational resources will facilitate the efficient investigation of the truly multiscale nature of these complicated processes in heterogeneous systems. The chosen xenon clusters and doped iron oxide nanoparticle systems are directly connected with the planned experimental efforts at current light facilities. The results from these simulations will provide predictions and new concepts to guide multimodal measurements using XFEL, and at the same time, maximize use of limited LCLS-II resources.

Title: E3SM Land Model (ELM) Biogeochemistry Perturbed Parameter Ensembles Project

Principal Investigator: Forrest M. Hoffman, Oak Ridge National Laboratory

Co-investigators: David M. Lawrence, National Center for Atmospheric Research;
Charles D. Koven, Lawrence Berkeley National Laboratory;
William J. Riley, Lawrence Berkeley National Laboratory;
James T. Randerson, University of California Irvine

ALCC allocation:

Site: National Energy Research Scientific Computing Center (NERSC)

Allocation: 550,000 node-hours on Cori

Research Summary:

This project will adopt a perturbed parameter ensemble (PPE) approach, employing hundreds of thousands of simulations, to develop an understanding of the importance of individual biogeochemical parameters in the Department of Energy's (DOE) Energy Exascale Earth System Model (E3SM) Land Model. These simulations will be used to quantify the impacts of model parameter uncertainties on changes in land carbon storage, carbon dioxide fertilization effects on land ecosystems and trends in water use efficiency in vegetation, Earth system feedbacks, extremes, and model skill. This effort provides a direct computational approach that uses the full land biogeochemical model instead of resorting to developing surrogate models that may not accurately emulate unusual cases. The project will apply a modified version of the International Land Model Benchmarking (ILAMB) package to quantitatively assess PPE results and identify optimal parameter combinations to reduce uncertainties in future fully coupled simulations and improve overall Earth system predictability, a primary mission of DOE's Biological and Environmental Research (BER) Earth and Environmental Systems Sciences Division (EESD).

Earth system models (ESMs) are designed to simulate the coupled multiscale, multiphysics processes associated with interactive dynamics, physics, chemistry, and biology across the land, ocean, sea ice, land ice, and atmosphere that drive the Earth's climate system. In addition to atmosphere and ocean physics and dynamics, today's ESMs simulate terrestrial and marine ecosystem processes, atmospheric chemistry, and human system interactions. Land surface models (LSMs) within coupled ESMs capture the mean state behavior of plants and soils over large spatial scales; however, process understanding limits the ability to reduce errors and biases when compared with observational data. LSMs rely on process representations that employ many often-uncertain parameters to approximate the evolution of carbon, water, and energy cycles. Data assimilation methods are used to calibrate and evaluate model accuracy and parameter uncertainty when measurements are available, but they do not provide insights into the sensitivity of model parameters or directly improve future predictions. Hence, using a PPE approach, the project will perturb about 50 biogeochemically relevant parameters around their nominal values, and run combinatorial ($\sim 50^3 = 125,000$) simulations to explore parameter uncertainty and their importance in predictions. Simulations will be run over an historical period (1850–2014), for comparison with observational data to gauge the divergence of terrestrial carbon variables based on combinations of parameter values.

Title: Microscopic Insight into transport properties of Li-battery electrolytes

Principal Investigator: Wei Jiang, Argonne National Laboratory

Co-investigators: Zhengcheng Zhang (Argonne National Laboratory)

ALCC allocation:

Site: Argonne Leadership Computing Facility (ALCF)

Allocation: 1,032,000 node-hours on Theta

Research Summary:

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

The proposed computation constitutes the simulation part of the current EERE VTO project, non-traditional electrolyte design from ionic liquid. This research is aimed at using large-scale, high-performance computing to assist discovery of novel battery electrolytes. The overall goal is to enable rational design of superior electrolytes for high voltage batteries. This study will focus on nontraditional electrolyte discovery from ionic liquids: a new entry to battery electrolytes. The nanostructural organization at electrolyte/electrode interface and wherein the transport properties and desolvation/solvation kinetics of charge carriers will be examined with advanced computational methodologies focused on exploring hierarchical length/time scales phenomena and the free energy landscape along with the desolvation/solvation process at electrode interface. The high-throughput capability will allow use of characterization approaches from simulation studies to link solution correlations with influences on lithium ion-transport behavior in electrolytes and enable the ability to seek multiscale structural attributes that allow facile and selective incorporation of the charge carrier while prohibiting the dissolution of cathodic transition-metal components.

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

Title: Cavitation dynamics in the Spallation Neutron Source Target

Principal Investigator: Eric Johnsen, University of Michigan

Co-investigators: Charlotte Barbier (ORNL)

ALCC allocation:

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

Allocation: 108,000 node-hours on Theta
4,000 node-hours on Summit

Research Summary:

The Spallation Neutron Source (SNS) at Oak Ridge National Laboratory produces the most intense neutron beam in the world. Neutron scattering can be used to uncover the molecular and magnetic structure of materials and has allowed ground-breaking discoveries in chemistry, biology, and materials science, among other fields. A high-energy pulsed proton beam impinges upon a flowing liquid mercury target, in which the spallation reaction occurs, thus producing neutrons. The pressure waves generated by this energy release reverberate within the target, ultimately giving rise to low-pressure regions in which the mercury cavitates (i.e., vaporizes). Such cavitation bubbles oscillate in response to the pressure field. If they reach a sufficiently large size, their collapse produces shock waves with the potential to cause structural damage. Although collapsing bubbles have been studied extensively in water under stationary conditions, cavitation in liquid metals such as mercury is far less well understood; for instance, the high density may produce stronger shocks, while the low vapor pressure may lead to smaller bubbles than in water. Additionally, the role of the flow (e.g., in certain regions of the SNS target) on the bubble dynamics is not well understood; for instance, the rotationality of the flow may lead to weaker shock waves, though perhaps they originate closer to the surface and thus may cause more damage. As higher neutron beam intensities are sought, cavitation damage is expected to become increasingly important and lifetime-limiting. Such cavitation induced damage is also problematic in liquid-metal-cooled reactors or thermal hydraulic systems during transient events such as steam explosions or water hammers. There is therefore a need to better understand cavitation dynamics in liquid metals and devise mitigation strategies.

The objective of this research is to simulate the collapse of individual gas bubbles in channel flow, in order to understand the role of confinement and shear on the bubble dynamics and resulting shock waves and to connect these phenomena to cavitation damage. In particular, the dependence of the magnitude of the maximum wall pressure (damage surrogate) on the stand-off distance of the bubble from the wall, on the local shear rate, and on the liquid will be determined. Additionally, the interaction of multiple bubbles with each other will be considered to determine the circumstances under which the emitted shocks are stronger or weaker than in the case of a single bubble. Finally, energy transfer between a macroscopic bubble and a small bubble will be investigated in the context of mitigation strategies for cavitation activity. In addition to new knowledge, these highly resolved simulations will provide data essential to developing data-driven models representing clouds of non-spherically collapsing bubbles. This research will be conducted using an in-house code for massively parallel direct simulations of the compressible Navier-Stokes partial differential equations in three dimensions for gas/liquid flows. The base spatial schemes are nominally non-dissipative (high-order accurate capturing at discontinuities, no dissipation in smooth regions) with numerical dissipation applied only where needed and explicit time-marching is used. This approach accurately and robustly handles discontinuities such as interfaces and shock waves.

Title: Advancing multi-year to decadal climate prediction with high-resolution E3SM and CESM

Principal Investigator: Ben Kirtman (Rosenstiel School of Marine and Atmospheric Science, U. of Miami)

Co-investigators: Leo Siqueira (U. Miams, RSMAS),
Gerald Meehl (NCAR, CATALYST),
Jadwiga Richter (NCAR, CATALYST)

ALCC allocation:

Site: Argonne Leadership Computing Facility (ALCF)

Allocation: 1,000,000 node-hours on Theta

Research Summary:

A team from the Rosenstiel School of Marine and Atmospheric Science (RSMAS) at the University of Miami and the Climate Change Prediction (CCP) group at the National Center for Atmospheric Research (NCAR) are jointly collaborating on this project through support from the Department of Energy Office of Science Biological and Environmental Research (BER) Regional and Global Model Analysis (RGMA) component of the Earth and Environmental System Modeling Program. The research will address the DOE/BER mission by exploring underlying mechanisms of predictability and quantifying interactions of climate processes in how they affect US extremes and understand the current and future impacts of these phenomena on regional and global climate. Several graduate students at RSMAS will learn how to run the high-resolution Energy Exascale Earth System Model (E3SM), and the NCAR team will collaborate with the CATALYST team to analyze the E3SM simulations.

The team hypothesizes that small-scale sea-surface temperature (SST) features made possible by high resolution ocean model component, along with improvements in atmospheric circulation and storm systems with higher resolution in the atmosphere, will affect decadal predictability and variability and the associated changes in extremes. The main objectives of the research are: (1) to investigate how the mesoscale features in the ocean feedback onto the representation of known modes of decadal variability (e.g., Atlantic Multidecadal Oscillation (AMO), Pacific Decadal Oscillation (PDO), among others), and diagnose how these known modes teleconnect to regional US extremes. The analysis will also include a detailed examination of subseasonal (e.g., Madden-Julian Oscillation, MJO) variability, seasonal (e.g., El Niño-Southern Oscillation, ENSO) variability, and how their interactions impact regional US extremes in terms of intensity and occurrence. Examples of extremes the team will investigate include flooding/drought, heat waves/cold spells, extreme wind events, and coastal flooding. Mainly, they seek to diagnose how the predictability of multi-year to decadal variability and the embedded extreme events are affected by oceanic mesoscale features and western boundary currents, and 2) quantify the role of volcanic eruptions in predictions of Pacific region SSTs and consequent global temperatures, with a focus on the regional details of the patterns of SSTs in verifying pattern correlations of observed and predicted SSTs that relate to decadal variability of globally averaged surface temperatures.

Initialized very high resolution global E3SM decadal predictions with observed volcanic emissions will be performed to test the above hypotheses and will comprise the most thorough set of simulations ever attempted for a decadal climate prediction project with the E3SM. Crucially, since these simulations are initialized predictions, they can make robust event-by-event comparisons with observations, which is particularly important when diagnosing extreme events. These simulations will be compared to lower resolution Community Earth System Model (CESM) simulations already conducted.

Title: Response functions of LaNiO₂: Insights into high-temperature superconductivity

Principal Investigator: Gabriel Kotliar, Rutgers University

Co-investigators: Corey Melnick (Brookhaven National Laboratory)

ALCC allocation:

Site: Argonne Leadership Computing Facility (ALCF)

Allocation: 115,000 node-hours on Theta

Research Summary:

One of the grand challenges in condensed matter physics is to understand materials wherein the correlations between electrons in *d* or *f* orbitals produce a rich tapestry of quantum. The response of a material to stimuli and the response functions which describe this provide invaluable insight into the precise mechanisms through which these phenomena arise. Accurate theoretical calculations of the response functions in strongly correlated materials are only now becoming possible thanks to advances in theory and computational power. In this project, the researchers will leverage these advances in order to compute the response functions in an important class of superconducting materials: the nickelates. In particular, they will compute the response functions in superconducting LaNiO₂ across a range of temperatures and electron concentrations. These response functions and their structure in this parameter space will be analyzed in order to gain fundamental insight into unconventional superconductivity.

Indeed, the structure of these response functions within the phase diagram will provide important insight into the superconducting glue in the nickelates, the fundamental interactions through which unconventional superconductivity arises in general, and the path towards room temperature superconductivity – a long-standing goal of the DOE and the scientific community, as high-temperature superconductivity would enable an incredibly efficient energy grid, new energy technologies, and more. The project will also illustrate the power of theoretical spectroscopy, wherein computational rather than experimental effort is used to probe a material's response to some stimulus. That is, it will provide unprecedented detail on the orbitals, the electronic structures, and the interactions through which high-temperature superconductivity manifests at a fraction of the cost. Finally, it will involve the continued development of open-source software which facilitates theoretical spectroscopy and material design.

Title: Characterizing Coastal Low-Level Jets and their Impact on Offshore Wind Farms

Principal Investigator: Jing Li, GE Research

Co-investigators: Shashank Yellapantula (National Renewable Energy Laboratory)

ALCC allocation:

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 240,000 node-hours on Summit

Research Summary:

The combined onshore and offshore wind will supply 20% of the US energy demand by 2030 as envisioned by the Department of Energy (DOE). The growth of US offshore wind energy, however, is hampered by a lack of understanding of the offshore wind conditions, the corresponding turbine behavior and the resulting farm performance. This project focuses on the numerical modeling of one such prevalent, yet poorly understood, offshore wind phenomenon—the coastal low-level jets (LLJ)—and its farm-level impact. This effort is aligned with DOE’s Clean Energy mission and has received funding from the National Offshore Wind Research and Development Consortium.

The work is a continuation of a 2020 ALCC award. The awarded allocation has allowed the team to employ a mesoscale-to-microscale coupling (MMC) technique required to model a period of LLJ occurring in the US Atlantic Coast wind energy Call Areas. With the 2021 allocation, this research will couple the Computational Fluid Dynamics (CFD) process of an LLJ with a turbine structural dynamic solver to elucidate LLJ’s impacts on wind turbines and wind farms. The expected results will be used to benchmark reduced-order models, as well as develop and demonstrate conceptual strategies to mitigate adverse LLJ impacts.

Accurately resolving the complex atmospheric turbulence and turbine wake requires the use of high-fidelity large-eddy simulations (LES) with high-end supercomputing. To leverage the leadership class computing resources, the team will use AMR-Wind—a massively-scalable, GPU-accelerated flow solver in the ExaWind software suite developed under DOE funding—to run these LES simulations. Coupled with OpenFAST, these simulations will generate a high-spatiotemporal-resolution description of wind farm flow, turbine power production and component loads under LLJ conditions. These results, together with an expanded predictive flow modeling capability, will accelerate the technology development in renewable energy, broaden the industry’s supercomputing applications, and ultimately contribute towards advancing DOE’s Clean Energy mission.

Title: Quantum Monte Carlo calculations of nuclei up to ^{16}O and neutron matter

Principal Investigator: Alessandro Lovato, Argonne National Laboratory

Co-investigators: Lorenzo Andreoli (Washington University),
Jason Bub (Washington University),
Garrett King (Washington University),
Saori Pastore (Washington University),
Maria Piarulli (Washington University),
Noemi Rocco (Fermi National Accelerator Laboratory),
Robert Wiringa (Argonne National Laboratory)

ALCC allocation:

Site: Argonne Leadership Computing Facility (ALCF)

Allocation: 632,000 node-hours on Theta

Research Summary:

Nuclear physics stands at the core of several exciting developments in science: state-of-the-art experimental facilities are probing the structure and reactions of stable and radioactive nuclei; the emerging gravitational-wave astronomy has opened a new window on the state of matter under extreme conditions; atomic nuclei are used in current and planned neutrino-oscillation and double-beta decay experiments, which will measure neutrino properties with unprecedented accuracy. The above efforts are accompanied by a remarkable progress of microscopic nuclear approaches that describe nuclei and nucleonic matter starting from the individual interactions among protons and neutrons. These nuclear forces are rooted in quantum chromodynamics, the fundamental theory of strong interactions, and are the main input of numerical methods that use high-performance computers to solve the nuclear Schrodinger equation with controlled approximation.

The overall goal of this research is to accurately describe atomic nuclei (including their spectra, densities, structure functions, transitions, low-energy scattering, and responses) while simultaneously predicting the equation of state of neutron matter. To achieve this objective, numerical quantum Monte Carlo methods will be extended to treat the ^{16}O nucleus, which has proven to be most effective in constraining three-nucleon potentials, relevant for the stability of neutron stars. In addition, novel computational methods will be used to study nuclear response densities induced by electrons and neutrinos scattering and gain more exclusive information about these processes. These studies will impact experimental programs on gravitational waves carried out by, e.g., the LIGO-Virgo collaboration, ongoing experiments that use electrons to study in medium nucleon-correlations carried out at Jefferson Lab, and neutrino experimental facilities, e.g., the Deep Underground Neutrino Experiment (DUNE).

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

Principal Investigator: Po-Lun Ma, Pacific Northwest National Laboratory

Co-investigators: Colleen Kaul (PNNL),
Kyle Pressel (PNNL)

ALCC allocation:

Site: Argonne Leadership Computing Facility (ALCF)

Allocation: 400,000 node-hours on Theta

Research Summary:

Accurate simulation of aerosol-cloud interactions (ACI) remains a major challenge for Earth system models (ESMs). The limited predictability is attributed to uncertainties and deficiencies in model resolution and process representations. To address these scientific challenges, the Energy Exascale Earth System Model (E3SM) version 4 that is currently under development will feature a variable resolution capability, resolving atmospheric processes from the convection-permitting scale to the planetary scale. The awarded computing resources will be used to generate ensemble simulations with variable resolutions and perturbed ACI processes to (1) assess the atmospheric characteristics and processes across scales; and (2) provide insights into the robustness of ACI parameterization across scales. A large-eddy simulations (LES) model and the nonhydrostatic version of the E3SM will be used to perform simulations of ACI at resolutions ranging from about 20 m horizontal grid spacing to about 100 km horizontal grid spacing. The high-resolution model results will be used to understand features at the subgrid scale of a coarse-resolution E3SM and provide guidance for parameterization improvements. Physical and chemical process representations affecting ACI (including the cloud and aerosol physics, chemistry, turbulence, and convection) will be perturbed to quantify the impacts of each process representation on atmospheric characteristics and ACI. The model simulations will cover a wide range of aerosol and cloud regimes to provide sufficient information for better generalizability of process representations.

Producing the large ensemble of simulations with variable resolutions and perturbed processes will provide realistic details on atmospheric characteristics affecting ACI. The simulation ensemble will be used to (1) identify emergent properties and controlling factors of ACI; (2) improve parameterizations suitable for multiple scales; (3) develop emulators for resolved and subgrid processes using novel machine learning techniques; and (4) validate the generalizability by synthetically reconstructing the high-resolution data from reduced-dimension representations using generative neural networks. This research fulfills the vision of Department of Energy Office of Biological and Environmental Research (BER) Earth and Environmental Systems Science Division (EESSD) to “develop an improved capability for Earth system prediction on seasonal to multidecadal time scales” and addresses EESSD’s three Scientific Grand Challenges: the Integrated Water Cycle, Drivers and Responses in the Earth System, and Data-Model Integration.

Title: Understanding the Role of Surface Energy in the Deformation of Metal Nanoparticles

Principal Investigator: Ashlie Martini, University of California Merced

Co-investigators: Dr. Tevis D. B. Jacobs (University of Pittsburgh)

ALCC allocation:

Site: National Energy Research Scientific Computing Center (NERSC)

Allocation: 300,000 node-hours on Cori

Research Summary:

Emerging energy technologies (such as fuel cells and metal-air batteries) require metal nanoparticles for their operation. By reducing nanoparticle size, the efficiency is increased while using less material, thus reducing the overall cost of these technologies. However, during use, the size and shape of nanoparticles changes spontaneously over time; even very small changes can dramatically degrade nanoparticle performance and consequently the efficiency and function of the application. Current theories describing nanoparticle deformation do not account for the important effect of surface energy and so cannot accurately describe nanoparticle behavior. The objectives of this research are to determine the processes and the factors that lead to permanent morphology changes, and to understand the effect of these changes on nanoparticle properties (especially mechanical properties). This research will use advanced computational models and cutting-edge experimental measurements to analyze the atomic rearrangement and changes in surface energies of nanoparticles under compression, revealing the fundamental mechanisms of nanoparticle deformation.

This research will provide new understanding into the previously unexplored role of surface energy in the deformation of small metal nanoparticles, enabling prediction of size-dependent nanoparticle properties. Such predictions are crucial for efforts to reduce cost and improve the performance of nanoparticle-based systems, enabling widespread use of energy technologies that are cleaner and more efficient.

Title: High-Fidelity Flow Data for Multiscale Bridging

Principal Investigator: Elia Merzari, Pennsylvania State University

Co-investigators: Igor Bolotnov (North Carolina State University),
Nam Dinh (North Carolina State University),
Paul Fischer (University of Illinois),
Misun Min (Argonne National Laboratory),
Rui Hu (Argonne National Laboratory),
Yang Liu (Argonne National Laboratory)

ALCC allocation:

Site: Oak Ridge Leadership Computing Facility (OLCF)
Allocation: 250,000 node-hours on Summit

Research Summary:

In 2020, the US Department of Energy (DOE), Office of Nuclear Energy, through the Integrated Research Projects (IRP) program, established a consortium focused on the research of thermal-fluids multiscale methods. This consortium represents a university component to the recently established Center Excellence for Thermal-Fluids Applications in Nuclear Energy that is fully integrated with the national laboratory efforts and its stakeholders. The consortium aims to deliver improved, fast-running models for complex physical phenomena relevant to advanced reactors.

The challenges in current thermal-fluids modeling of advanced reactors are dual. Primarily, there is a severe deficit of direct thermal-fluids data that can be applied to new, innovative nuclear system designs due to a lack of large-scale integral-effect test facilities for the wide range of scenarios and conditions relevant to these systems. Secondly, the results of experiments in advanced instrumented, small-scale, separate-effect experiments and high-fidelity numerical simulations that could be used to bridge the gap are inefficiently used. These conditions highlight the pressing needs for bridging this scale gap (multiscale bridging), enabling the use of high-fidelity simulation data to deliver improved fast running models. The consortium is focusing in particular on four challenge problems of great interest to the advanced reactor nuclear industry: (i) Flexible Modeling for Heat Transfer, (ii) Thermal Striping of Internals, (iii) Mixing in Large Enclosures and (iv) Multiscale Core Modeling Coupled to Fuel Performance. The work of the consortium covers numerical method development, high-fidelity numerical simulation and experimental research.

To support the goals of the project, an extensive high-fidelity fluid flow simulation database will be established. The planned computations will employ scale resolving techniques such as Direct Numerical Simulation and Large Eddy Simulation. The geometries simulated will span from experimental facilities involving jet mixing in large enclosures to more canonical geometries such as parallel plates in mixed convection and axial flow in simplified rod arrays. The database will be instrumental to the development of improved multiscale bridging methods and it will ultimately enable a paradigm shift in the way multiscale simulations are performed for nuclear engineering applications. Finally, the simulations will increase the understanding of turbulent thermal mixing in complex geometries, with a broad impact.

Title: Multiscale bubble breakup and gas transfer in turbulent oceanic environments

Principal Investigator: Parviz Moin, Center for Turbulence Research, Stanford University

Co-investigators: none

ALCC allocation:

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

Allocation: 650,000 node-hours on Theta
500,000 node-hours on Cori

Research Summary:

This is a computational fluid dynamics study focusing on high-fidelity simulations of bubble breakup and gas dissolution in oceanic breaking waves that address energy and environmental challenges through the impact of oceans on weather/climate predictions and offshore wind technologies. The wave-breaking process gives rise to turbulent fluctuations that break up entrained air cavities in quick succession. A wide range of bubble sizes extending down to the Hinze scale—the critical length scale below which turbulent fragmentation ceases—is generated through a breakup cascade. The first objective of this work is to extend fundamental understanding of the bubble fragmentation process to sizes smaller than the Hinze scale, at which bubble formation mechanisms remain a subject of active research. The second objective is to quantify the bubble-induced gas dissolution that is enhanced by the breakup cascade. A major challenge for accomplishing these goals has been the inherent multiscale nature of the problem, which dramatically increases the computational cost and the complexity of the associated physical phenomena. This project aims to resolve these issues by deploying in-house tools, solvers, and techniques that will readily leverage the capabilities of DOE supercomputing systems to obtain novel statistics and insights into the sub-Hinze-scale bubble population and accompanying gas dissolution.

Understanding sub-Hinze-scale bubble formation mechanisms and quantifying gas dissolution have significant impact on problems of practical importance, such as CO₂ transport in the carbon cycle through dissolution of soluble components of the entrapped bubbles, and light and sound scattering. In particular, maritime and climate studies may be informed by the temporal evolution of the bubble size distribution, which enables quantification of the total air–sea interfacial area and size-dependent effects like radiation scattering. This project will yield deeper insights into oceanic bubble breakup and gas dissolution for the development of future predictive technologies for more energetic waves and their associated multiphysics phenomena, such as hydroacoustics, radiation transfer, and scalar transport.

Title: Electromagnetic corrections to strong dynamics

Principal Investigator: Amy Nicholson, UNC Chapel Hill

Co-investigators: Kate Clark (NVIDIA),
Michele Della Morte (CP3-SDU),
Andrea Shindler (FRIB/MSU),
Andre Walker-Loud (LBNL),
Justus Tobias Tsang (CP3-SDU),
Henry Monge-Camacho (University of Costa Rica),
Ben Hoerz (LBNL),
Dean Howarth (LLNL), Pavlos Vranas (LLNL)

ALCC allocation:

Site: Oak Ridge Leadership Computing Facility (OLCF)
Allocation: 300,000 node-hours on Summit

Research Summary:

Lattice quantum chromodynamics (LQCD) is beginning to deliver theoretical results to be confronted with experiments at an unprecedented level of precision. This is a fundamental prerequisite to probe physics beyond the Standard Model (BSM). To control the precision of LQCD results at the order of, or below, 1% total accuracy, it becomes unavoidable to introduce Quantum Electrodynamics (QED) effects in LQCD calculations.

The goal of this project is to provide a robust and theoretically sound lattice QED formulation that allows control over all the systematics at the sub-percent level. A recently proposed method will be used to regulate infrared finite-volume corrections through introduction of a photon mass. With the time requested, controlled extrapolations to the physical point will be performed for a multitude of observables, including hadron spectra and charged pion scattering, with strong isospin breaking effects also included.

This research is closely aligned with the goals of the 2015 Long-Range Plan for Nuclear Physics, specifically in the area of Nuclear Structure and Reactions. This project focuses on a necessary step towards the understanding of nuclear interactions mediated by QCD and the Electroweak theory, and furthermore provides a testing ground toward calculations of important BSM decay modes, which will require novel non-perturbative techniques such as those being explored within this work: the massive photon method serves as a starting point for developing a massive neutrino method that can be used for neutrinoless double beta decay. All the codes developed for this project will be made publicly available.

Title: Investigation of Compressibility and Nonequilibrium Turbulence in Reacting Flows

Principal Investigator: Joseph C. Oefelein, Georgia Institute of Technology

Co-investigators: Kyle A. Schau (Georgia Institute of Technology),
Ramanan Sankaran (Oak Ridge National Laboratory),
Kareem Ahmed (University of Central Florida)

ALCC allocation:

Site: Oak Ridge Leadership Computing Facility

Allocation: 600,000 node-hours on Summit

Research Summary:

Turbulent reacting flows at high-pressure (e.g., supercritical conditions) or high-speed (e.g., supersonic and hypersonic conditions) involve significant levels of interscale nonequilibrium dynamics that are not well understood and not adequately accounted for in state-of-the-art models. The simultaneous presence of strong compressibility effects compound the problem due to the extreme complexity of the nonlinear coupling that is introduced between mechanical, thermal, and material properties. Flow conditions frequently involve shock waves, extreme thermophysical property gradients, thermal inhomogeneity, flow separation, high turbulence intensity, and exothermic heat release. Perturbations induced by interscale dynamics are generated by both turbulence-chemistry and shock-turbulence interactions. Nonequilibrium effects have a strong influence on these flows because chemical reactions, which require mixing at the molecular scale, can be delayed until the small scales in the velocity spectrum are energized. Similarly, heat release from chemical reactions moves turbulence away from equilibrium through the introduction of thermal energy at the viscous scales and its conversion into mechanical energy at larger scales. The timescales associated with the macroscopic fluid dynamics are short due to extreme mechanical, thermal, and material gradients across flames, shock waves, and boundary layers. The resulting interscale dynamics typically deviate significantly from classical theory, but can have a significant impact on critical flow processes.

This research will address the challenges outlined above by establishing an in-depth understanding of the nonlinear dynamical nature of turbulent reacting flows at the highly-compressible nonequilibrium conditions associated with advanced propulsion and power systems. The allocation will enable one-of-a-kind simulations that reveal the interscale dynamics associated with key multiscale-multiphysics processes at relevant flow conditions. This will be accomplished through a combination of Direct Numerical Simulation (DNS) and Large Eddy Simulation (LES) investigations guided by a unified treatment of turbulence, thermodynamics, and transport in both combustion-induced and shock-induced processes. Results will have a direct impact on development of the DOE, National Energy Technology Laboratory (NETL), University of Turbine Systems Research (UTSR) goal of removing environmental concerns over the future use of fossil fuels by developing revolutionary, near-zero-emission advanced turbine technologies. Project emphasis will be placed on understanding the underlying factors affecting combustion, aerodynamics, and heat transfer in advanced combustion turbines, supercritical carbon-dioxide power cycles, and pressure gain combustion devices.

Title: Applying Machine Learning to Reynolds Number Impact on High Pressure Turbine Flow

Principal Investigator: Michal Osusky, GE Research

Co-investigators: Rathakrishnan Bhaskaran (GE Research),
Gustavo Ledezma (GE Research),
Greg Sluyter (GE Aviation),
Sriram Shankaran (GE Aviation),
Ramakrishnan Kannan (ORNL)

ALCC allocation:

Site: Oak Ridge Leadership Computing Facility

Allocation: 256,000 node-hours on Summit

Research Summary:

This project will leverage Department of Energy leadership computing facilities to explore machine learning modelling capabilities for studying the impact of Reynolds number (Re) on flow within the hot gas path (HGP) of a high-pressure turbine, found, for example, at the core of an aircraft engine. Through their operating cycles, the turbines experience flows over a range of Re , affecting flow structures and behaviors critical to efficient turbine performance and longevity. The physics impacted by varying Re will be studied through high-fidelity Large Eddy Simulation (LES), which provide accurate numerical solutions for the flows of interest, providing engineering insights that will lead to novel turbine designs with improved robustness and efficiency. These include the ability to accurately predict the temperature fields around the turbine, as well as migration and evolution of turbulent flows. High Re flows, which require a significant increase in computational resources due to the increase in required simulation resolution, are still cost-prohibitive for use in design. To mitigate this, machine learning (ML) techniques provide a promising avenue for overcoming this hurdle. For example, the ability to correlate high-fidelity simulation results from lower Re analyses to a range of flows experienced by the turbine would lead to a significant advance in the ability to design robust, fuel efficient turbomachinery. This project will therefore focus on employing high-fidelity simulations to train and validate ML algorithms. The trained ML tools will be used for predicting the impact of Re on the flow in the turbine. The successful execution of this project will serve as a critical step towards leveraging high fidelity simulations across engine lines and flight conditions, employing simulations in lieu of testing (“computation wind tunnel”) in the design of competitive next-generation gas turbines.

Design and manufacture of globally-competitive gas turbines (used in aircraft engines and power generation) requires persistently higher and higher fidelity simulations to inform critical decision making. To enable this, the goal of this program is to advance the numerical prediction capability and the use of LES in the high-performance computing environment. This is especially important as the demands for improved efficiency and reduced cost come up against the limitations of current design tools. This project will take advantage of a peta-scale simulation of fine-scale turbulent flows coupled with the interaction effects of the larger complex system to guide design of the next generation of aircraft propulsion systems and power generation turbines. In the realm of commercial aircraft propulsion alone, the benefit of an improved physical understanding and predictive capability of HGP flow would be on the order of 70 million gallons of jet fuel saved for U.S. airlines every year, equivalent to a savings of 35 million US\$. In addition, greenhouse gas emissions would be reduced by 225,000 metric tons. From a durability perspective, improved HGP understanding has the potential to offset hundreds of millions of US\$ in maintenance costs.

Title: Probabilistic Comparative Modeling of Colorectal Cancer Screening Strategies

Principal Investigator: Jonathan Ozik, Argonne National Laboratory

Co-investigators: Carolyn Rutter (RAND Corporation),
Reinier Meester (Erasmus University Medical Center),
Karen Kuntz (University of Minnesota)

ALCC allocation:

Site: Argonne Leadership Computing Facility (ALCF)

Allocation: 160,000 node-hours on Theta

Research Summary:

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

This project will use leadership class computing resources to run comparative probabilistic sensitivity analyses (PSAs) of screening strategies across three CRC models. PSAs use samples from the distributions of uncertain model parameters to propagate uncertainties to model predictions. To date, formal PSAs addressing uncertainties in all relevant parameters, that is, in the natural history and screening parameters, have not been performed for complex cancer screening models. The goal of this project is to take this a step further by using a comparative framework that includes three state-of-the-art CRC models. Funded under the National Cancer Institute's (NCI) Cancer Intervention and Surveillance Modeling Network (CISNET) program, these CRC models were independently developed for the evaluation of interventions, with emphasis on screening, and describe CRC natural history using different underlying assumptions. These differences enable model comparisons that explore the impact of structural uncertainty, in addition to parameter uncertainty.

The comparative PSAs in this work will be used to generate cost-effectiveness analyses (CEAs) for complex interventions and to provide formalized assessments of both structural and parameter uncertainties across the three CRC models. For external parameters (i.e., costs, screening test characteristics and health-related quality of life), distributions will be derived from current systematic reviews. For the calibrated parameters, the joint posterior distributions for all three models will be obtained by applying sequential Bayesian calibration algorithms, using CRC registry and clinical trial data as calibration targets. The calibration algorithms will be implemented as HPC workflows using the EMEWS framework (emews.github.io) and the Swift/T workflow engine (swift-lang.org/Swift-T). EMEWS workflows have been run on leadership computing resources such as OLCF Summit, NERSC Cori and ALCF Theta. Additional EMEWS workflows will be used to sample from the joint posteriors and generate the comparative PSAs of screening strategy outcomes. This work integrates complex data, large-scale machine learning (ML) algorithms for uncertainty quantification (UQ), and simulation to advance HPC-enabled scientific discovery, supporting the DOE mission. The approaches also promote the adoption of HPC across other scientific domains as they are applicable beyond cancer or even microsimulation models to any computational modeling method (e.g., agent-based modeling) requiring heuristic/derivative free methods for UQ.

Title: Improving the Representation of Mesoscale Convective Systems in Weather and Climate Models

Principal Investigator: Andreas Prein, National Center for Atmospheric Research

Co-investigators: none

ALCC allocation:

Site: Argonne Leadership Computing Facility (ALCF)

Allocation: 372,000 node-hours on Theta

Large complexes of thunderstorms known as Mesoscale Convective Systems (MCSs) are dominating the weather and climate of equatorial regions and parts of the mid-latitudes during summer. The central U.S. is a hotspot of MCS activity during the warm season with up to 70 % of the total summertime rainfall produced by those systems. Furthermore, most flash flooding events are caused by MCSs such as the West Virginia flooding or the Baton Rouge flooding of 2016. Most weather and climate models have deficiencies in simulating MCSs, limiting their predictive skill on time scales from days to decades. Improving the simulation of MCSs in these models is of uttermost importance to realistically simulate the earth's energy and water cycle including the prediction of droughts and floods.

This project is tackling this challenge by using unique MCS observations from the Department of Energy's (DOE's) Atmospheric Radiation Measurement (ARM) sites in the U.S. central Great Plains and the Amazon rain forest to evaluate the simulation of MCSs in state-of-the-art atmospheric models. Turbulence (large-eddy) resolving simulations of well-observed MCSs at a leadership computing facility allow unprecedented insights into scale interactions of convective processes across multiple orders of magnitude and can answer fundamental questions about large-scale environmental impacts on MCS development and systematic model biases in current weather and climate predictions. Furthermore, this project will identify model configurations that allow the realistic simulation of MCS processes at affordable computational costs and develop novel cloud parameterizations that will improve coarse resolution global earth system models that are used for climate change assessments. The research supports DOE's Earth and Environmental Systems Sciences Division's mission to enhance the seasonal to multidecadal predictability of the Earth system and resilience of infrastructure against extreme weather and hydrological changes.

Title: Design of Novel Titanium Based Alloys for Additive Manufacturing using HPC-Aided Large-Scale Phase Field Simulations

Principal Investigator: Balasubramaniam Radhakrishnan, Oak Ridge National Laboratory

Co-investigators: Gorti Sarma (Oak Ridge National Laboratory),
Younggil Song (Oak Ridge National Laboratory)

ALCC allocation:

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 1,250,000 node-hours on Summit

Research Summary:

The objective of the project is to utilize HPC-aided simulations of the evolution of the underlying microstructure during solidification of titanium-based alloys during additive manufacturing (AM). This work is in support of a recently funded project under the High-Performance Computing for Manufacturing (HPC4Mfg) program at ORNL with collaboration from Raytheon Technologies Research Center. The basic science concept behind the work is that inducing a change from the typically observed columnar dendrites in AM-processed titanium alloys to equiaxed dendrites, i.e., columnar to equiaxed transition (CET), under AM solidification conditions consistently and reliably under the entire range of thermal conditions experienced during the manufacturing of a component with complex geometry will significantly reduce the mechanical anisotropy typical of AM titanium alloy components. The CET is significantly influenced by the ability to consistently achieve crystal nucleation in the bulk liquid ahead of the main advancing solid-liquid interface under the non-equilibrium AM solidification conditions. This effort involves performing high-fidelity, high-resolution, three-dimensional phase field simulations of the interplay between the destabilizing advancing solid-liquid interface that leads to columnar structures and the nucleation of solid in the undercooled liquid ahead of the advancing interface. The simulations will investigate the effect of alloy parameters including the solidification range, the solid-liquid interfacial energy, the strength of the interfacial energy anisotropy, the crystal orientations in the advancing front, etc. on the development of the solidification microstructure. The simulation results will be used to perform AM experiments using the model-based compositions and process conditions at the industry in order to optimize the alloy and process design.

The outcome of this project will make a significant impact on the design of alloys for additive manufacturing in general, and specifically for titanium-based alloys used in aerospace and automotive sectors. Currently, AM of titanium alloys is based mainly on compositions specifically developed for wrought alloys such as Ti-6Al-4V that pose many challenges both during the AM process as well as during service. This research effort will help create a new alloy design paradigm for AM. The research is directly relevant to the goal of the Advanced Manufacturing Office (AMO) of the U.S. Department of Energy to promote collaboration between the national laboratories and industries to investigate new manufacturing technologies that have critical impact on the national economy.

Title: LBNF - PIP-II Optimization studies for Megawatt 120-GeV beams on target

Principal Investigator: Igor Rakhno, Fermi National Accelerator Laboratory

Co-investigators: Nikolai Mokhov, Fermi National Accelerator Laboratory

ALCC allocation:

Site: Argonne Leadership Computing Facility (ALCF)

Allocation: 450,000 node-hours on Theta

Research Summary:

The Deep Underground Neutrino Experiment and Long-Baseline Neutrino Facility (DUNE-LBNF) have been under development at Fermilab since early 2010s. This international project represents a convergence of a substantial fraction of the worldwide neutrino physics community, provided by the large investment by the Department of Energy (DOE). The primary scientific objectives of DUNE are to carry out a comprehensive investigation of neutrino oscillations to test charge and parity (CP) violation in the lepton sector, determine the ordering of the neutrino masses, and to test the three-neutrino paradigm (electron, muon and tau neutrino). Independent measurements of the propagation of neutrinos and antineutrinos through matter will make it possible to observe neutrino transitions with the precision required to determine the CP-violating phase and the neutrino mass hierarchy. The LBNF will provide a 120-GeV proton beam on a neutrino production target utilizing a new 800-MeV superconducting linear accelerator which is expected to be completed in 2027 within Proton Improvement Plan-II project (PIP-II).

Many of the LBNF project milestones will heavily rely on computer simulations using the MARS15 computer code. The neutrino beamline, which utilizes a target and beam focusing horn systems, decay pipe, hadron absorber and other systems, is a core component of the LBNF. Such simulation studies are done by means of precise Monte Carlo modeling of radiation transport and interactions with matter in a broad energy region, utilizing the power and capabilities of the Fermilab's MARS15 computer code. Also, a detailed MARS model of the PIP-II facility—both accelerator beamline itself and infrastructure—have been developed. These resources are required for comprehensive Monte Carlo studies of radiation shielding, including both normal operation and accident scenarios. The model will be verified, and test simulations will be performed as a pre-requisite before performing numerous production simulations. This research will focus on significant further optimization and follow-up work which is required to address many components and issues such as current budget constraints. The performed simulations will leverage the substantial supercomputer resources which are essential to the success of the project.

Title: Predicting hot electrons for inertial confinement fusion

Principal Investigator: Chuang Ren, University of Rochester

Co-investigators: Riccardo Betti (University of Rochester)

ALCC allocation:

Site: National Energy Research Scientific Computing Center (NERSC)

Allocation: 800,000 node-hours on Cori

Research Summary:

Inertial confinement fusion aims to realize fusion, which powers the Sun and stars, as a clean and sustainable energy source for humankind. One major approach is direct drive, which uses high energy laser to compress a tiny target to achieve fusion ignition condition of high density (100's grams per cubic centimeter) and high temperature (100's million degrees). Compressing the target to high density requires fusion fuel to be kept at low temperatures before the compression starts. However, the laser can generate plasma waves through a process called laser-plasma instabilities. Electrons can be accelerated by the plasma waves to high energy, much like a surfer gaining speed riding a wave. These hot electrons can preheat the fuel to impede the compression. A predictive capability of hot electron generation is thus required to find successful experiment designs where the hot electrons are kept at a tolerable level. Such a capability is currently lacking and is identified as a critical need by the inertial confinement fusion community. This project uses a series of particle-in-cell simulations, combined with a hot electron database from experiments performed at the Laboratory for Laser Energetics, University of Rochester, to obtain a scaling law for hot electron generation. The simulations, limited by available computer resources, can provide a scaling law that includes only the most relevant physics. It will then be improved by comparing with the experimental data to account for any missing physics. Such a scaling law can be implemented inline in codes modeling target performance and help find a path to ignition, a scientific and engineering grand challenge.

If successful this work will significantly improve the fidelity of direct-drive implosion modeling, especially on the areal density and ignition likelihood. It will also explore the high-gain shock ignition scheme and potentially open new paths to ignition. It will expand an active knowledge base in laser plasma instabilities and have broad impact on indirect-drive inertial confinement fusion as well. It will also train new workforce in academia, national laboratory, and industry.

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

Principal Investigator: Dillon Shaver, Argonne National Laboratory

Co-investigators: Haomin Yuan (Argonne National Laboratory),
Aleks Obabko (Argonne National Laboratory),
Landon Brockmeyer (Argonne National Laboratory),
Yang Liu (Argonne National Laboratory)

ALCC allocation:

Site: Oak Ridge Leadership Computing Facility
Argonne Leadership Computing Facility

Allocation: 80,000 node-hours on Summit
550,000 node-hours on Theta

Research Summary:

This research will provide support in creating the next-generation of nuclear power reactor design tools needed by industry and regulators to enable the deployment of advanced nuclear power. By leveraging high-performance computing, simulations of fluid flow and heat transfer in novel nuclear power reactor cores will be performed. In particular, this work will perform high-fidelity simulations of fluid flow around wire-wrapped fuel pins for liquid metal cooled reactors, flow in bebble beds for both gas- and molten salt-cooled reactors, and analyze flow-induced vibration phenomena for structures such as spacer grids and mixing vanes in existing light water reactors (LWRs). The data generated from these simulations will then be analyzed using both traditional and machine-learning methods to inform fast-running design tools that can be used directly by the industry.

By focusing on designs similar to those being investigated as part of the recent DOE awards to the Advanced Reactor Demonstration Program (ARDP), this will provide insight into the complex phenomena associated with reactor designs most relevant to the industry. As a result of this work, the design tools available to industry will be made even more accurate, further enhancing the safety and reliability of current and next-generation nuclear power.

Title: Informing Forensics Investigations of Nuclear Materials

Principal Investigator: Ashley Shields, Oak Ridge National Laboratory

Co-investigators: Andrew Miskowiec (Oak Ridge National Laboratory),
Jennifer L. Niedziela (Oak Ridge National Laboratory),
Sara B. Isbill (Oak Ridge National Laboratory),
Erik Nykwest (Oak Ridge National Laboratory)

ALCC allocation:

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 368,000 node-hours on Summit

Research Summary:

Researchers at Oak Ridge National Laboratory actively support national and international nuclear forensics efforts by developing tools and methods to characterize nuclear material through ongoing research supported by the National Nuclear Security Administration. This research requires the use of advanced computational and experimental resources to fully elucidate process kinetics and environmental degradation of key fuel cycle materials, knowledge of which is necessary to provide a sound technical basis for monitoring nuclear fuel cycle activities.

Determining the effect of process conditions on the underlying crystal structure via electronic structure calculations directly connects chemical and physical changes to fuel cycle materials with laboratory observations. These calculations assist in this research in three key ways: (1) predicting relative stability of perturbed materials, (2) predicting physical and chemical changes due to thermodynamic perturbations, and (3) predicting spectroscopic observables. The prediction of optical spectroscopic observables is of particular interest, as optical vibrational spectroscopy is a key experimental technique for nondestructive chemical composition determination, relevant for nuclear forensics.

The goal of this project is to connect density functional theory (DFT) determinations of lattice dynamics under equilibrium and nonequilibrium conditions to experimental vibrational spectra obtained from Raman, infrared, and neutron scattering measurements. This project aims to bridge the gap between experimental observations and atomistic understanding of the influence of defects and other structural perturbations induced under fuel cycle-relevant conditions to inform chemical characterization and forensic analysis. Particular systems of interest to this research are uranium and defected carbon materials, where large simulation cells are required to simulate the parameters of interest. For the first time, the vibrational properties of defective uranium materials can be modeled with the accuracy afforded by DFT on large scale systems under real-world conditions. Large cell sizes are also required to capture the changing vibrational properties along forensics-relevant reaction pathways in uranium oxides. In addition, large simulation cells are needed to accurately model realistic defect concentrations and associated lattice dynamics between complex carbon-based materials. In each case, these high-fidelity comparisons support the interpretation of experimental measurements, and inform fundamental understanding of nuclear fuel cycle materials.

Title: Breaking the gene annotation bottleneck with structure-based machine learning

Principal Investigator: Jeffrey Skolnick, Georgia Tech Research Corporation

Co-investigators: Ada Sedova (ORNL),
Jianlin Cheng (University of Missouri-Columbia),
Jerry Parks (ORNL),
Mu Gao (Georgia Tech)

ALCC allocation:

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 50,000 node-hours on Summit

Research Summary:

This project aims to infer gene function from sequence via structural inference using advanced high-performance computing (HPC)-based deep learning methods and multiscale simulations. The ability to predict the function of a protein-coding gene from its sequence is a grand challenge in biology. With the advances in sequencing technology, the number of sequenced genomes is growing exponentially, leading to a new bottleneck in genome annotation: the ability to accurately infer an experimentally testable hypothesis for the majority of the protein coding genes in this massive dataset of sequences would revolutionize predictive and synthetic biology. Success could enable the synthesis of novel molecules and develop new approaches to the design of novel biomaterials. This technology gap was identified and characterized in a recent DOE BER workshop entitled “Breaking the Bottleneck of Genomes: Understanding Gene Function Across Taxa.” One of the important means of addressing this gap was the need for new approaches in computation. Inferring gene function from sequence computationally, for sequences that do not have close matches to proteins with known function (a substantial portion of the sequence space) is a challenging task. Existing technologies for gene function annotation are often based on evolutionary inference. If protein A has function X, and protein B is evolutionarily related to protein A, then protein B is inferred to have function X. This method fails when the query sequence has no close matches to known proteins. However, the structure of two proteins can be similar even with very low sequence identity matches (<15%), due to the interchangeable biophysical properties of certain amino acids. Currently being developed is a suite of novel deep-learning methods that use the protein’s structural information to advance proteome annotation, e.g., the first deep-learning algorithm SAdLSA for protein sequence alignment, as well as deep-learning approaches predicting the protein’s structure itself. Deep-learning approaches for predicting protein structure have recently achieved dramatic success in the recent CASP contests. The ability to match unknown proteins to existing crystallographic structures without the need for modeling the structure itself opens a vast avenue for novel functional annotation technology. In addition, the ability to identify remote homologs through deep-learning based sequence comparison, and accurately model the structure of an unknown protein could lead to the understanding of protein-mediated biochemical mechanisms such as enzyme catalysis. Finally, the prediction of protein-protein interactions and the structure of the resulting complex (PPIs) could help better understand the protein’s function as these complexes are often the functional units within a cell. Deep learning is also being used to address the PPI challenge.

DOE leadership HPC resources will be used to train large deep neural networks for prediction of structure, interaction, and function of proteins from their sequence, and apply these algorithms to entire proteomes of organisms with large numbers of proteins of unknown function. The use of massively parallel GPU-accelerated deep learning platforms will dramatically improve the accuracy of the models and enhance sampling and accuracy in molecular simulations to refine structures, predict stability and conformational flexibility, and model molecular interactions. This will allow an approach to this problem that is more powerful than was previously possible using academic resources. Success of this project will be important to one of DOE’s Office of Biological and Environmental Science’s primary missions, to translate nature’s genetic code into predictive models of biological function.

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

Principal Investigator: William Tang, Princeton University

Co-investigators: Zhihong Lin (University of California Irvine)

ALCC allocation:

Site: Argonne Leadership Computing Facility (ALCF)

Allocation: 45,000 node-hours on Theta

Research Summary:

An especially time-urgent and formidable problem facing the development of a fusion energy reactor is the need to reliably predict and avoid largescale major disruptions -- powerful instabilities that result in a loss of plasma confinement accompanied by rapid release of energy which if unmitigated can significantly damage the plasma facing components and harm the structural integrity of the whole device. Prediction and mitigation are key to major “tokamak systems” such as the EUROFUSION Joint European Torus (JET) experiment today and will be crucial for the ground breaking \$25B international burning plasma ITER device targeted for “first plasma” in 2026 -- and the exciting potential to exceed “breakeven” fusion power by a factor of 10 during the D-T (Deuterium-Tritium) experimental phase starting in 2036.

The research goals of this project is to develop realistic predictive plasma models of disruptions that are integrated within a modern plasma control system. This novel integrated modeling tool would ideally lead to a mature AI-enabled comprehensive control system for ITER and future reactors that feature integration with full pilot-plant system models. With engagement of the DOE leadership class high-performance computing (HPC) facilities managed by the Advanced Scientific Computing Research (ASCR) program, the practical impact of these studies is to deliver state-of-the-art AI-driven projections required to help cost-effectively plan, “steer,” & harvest the needed information from ITER to establish the foundations for a viable demonstration of a practical operating fusion reactor.

The proposed simulations will support the physics studies identified by the project “Machine Learning Application to Predictive Studies of Disruptions in Tokamaks” funded by the Fusion Energy Science (FES) program and the SciDAC (Scientific Discovery through Advanced Computing) project “Integrated Simulation of Energetic Particle in Burning Plasmas” jointly funded by FES and ASCR. In addition to the clear mission importance to the Fusion program, the critical reliance of HPC and AI/machine learning to the success of this project establishes a vital connection to the ASCR program and the DOE leadership computing facilities.

Title: QMC-HAMM: From the nanoscale to the mesoscale.

Principal Investigator: Lucas Wagner, University of Illinois at Urbana-Champaign

Co-investigators: David M Ceperley (UIUC)

ALCC allocation:

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 300,000 node-hours on Summit

Research Summary:

A major theme of condensed matter and materials physics is the relationship between the high-energy microscopic behavior of electrons and nuclei to the emergent relatively low-energy mesoscopic behavior of materials. Elucidating this relationship is a challenge, since the microscopic model requires advanced solution methods for many-body quantum mechanics, and many diverse behaviors can emerge at the mesoscopic scale. In complex materials, effective degrees of freedom phonons, spins, and electron-like excitations can interact in complex ways which are difficult to access experimentally.

The state of the art in creating mesoscopic models starting from the microscopic behavior is based on density functional theory (DFT) calculations. In recent years, modern machine learning techniques have been able to reproduce potential energy surfaces from standard DFT functionals to a very high accuracy; the accuracy potential energy surfaces can be limited by the underlying DFT data. In quantum materials such as twisted bilayer graphene, interactions between electronic excitations can be critical to their behavior, which DFT does not treat accurately. To resolve the above issues, it is necessary to move beyond density functional theory and to base mesoscopic models on more accurate microscopic calculations. In this project, quantum Monte Carlo calculations are used to develop high accuracy models of two materials systems: hydrogen at high pressure, and bilayer graphene.

This research is in support of the QMC-HAMM project, supported by the Department of Energy. The hydrogen data will support efforts to understand hydrogen at extreme conditions, including astrophysical observations and high temperature superconductivity in those materials. The graphene data will support efforts underway in many labs to understand the nature of electrons in twisted bilayer graphene. The data generated by both efforts will also provide a valuable reference data set for less accurate calculations, which will be used to improve the overall quality of materials modeling.

Title: Plasma Surface Interactions Modeling

Principal Investigator: Brian Wirth, University of Tennessee, Knoxville

Co-investigators: David Bernholdt (ORNL),
Aidan Thompson (SNL),
Karl Hammond (University of Missouri),
Wahyu Setyawan (PNNL),
Ilon Joseph (LLNL)

ALCC allocation:

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

Allocation: 1,100,000 node-hours on Summit
1,680,000 node-hours on Cori

Research Summary:

The realization of fusion as a practical, 21st Century energy source requires improved knowledge of plasma surface interactions (PSI) and the materials engineering design of the divertor and first wall systems to survive the incredibly extreme heat and particle flux exposure conditions of a fusion power plant. The objective of this project is to further advance understanding of the response of tungsten, the proposed ITER divertor, to low energy, mixed H-He plasma exposure in the presence of impurity atoms of beryllium and nitrogen. In particular, two tasks are envisioned that investigate the surface response and tritium retention and permeation following implantation of beryllium, helium and hydrogen into the tungsten divertor, as well as investigating the mechanical response of the plasma-damaged tungsten. These simulations will continue to expand our database to computationally benchmark and verify our continuum simulation approach to plasma surface interactions, and provide necessary knowledge to identify materials design strategies to effectively manage the high gas exposures in the fusion energy environment.

Title: Compact coherent X-ray sources based on plasma-based acceleration

Principal Investigator: Xinlu Xu, SLAC National Accelerator Laboratory

Co-investigators: Mark Hogan (SLAC National Accelerator Laboratory),
Warren Mori (UCLA)

ALCC allocation:

Site: National Energy Research Scientific Computing Center (NERSC)

Allocation: 800,000 node-hours on Cori

Research Summary:

Plasma-based acceleration (PBA) has been considered to be a compact and powerful source of high energy electrons that could be used to produce incoherent and coherent sources of X-rays. PBA has been shown to be able to accelerate charged particles over cm to meter distances with acceleration gradients more than three orders of magnitude larger than conventional accelerators. Incoherent X-rays and γ -rays are produced along the curved trajectories of the electrons inside the plasma or during the interactions between these electrons and separate lasers, magnets and solid targets. Coherent X-ray sources are far superior than incoherent sources on many aspects (e.g., spectrum brightness, photon flux and focusability) and will lead to numerous applications in science, industry and medicine. However, there is no feasible path to build such plasma-based compact coherent sources to date. A totally different level of understanding and control of the PBA is needed if coherent radiation is desired, including new ideas to inject and manipulate the electrons. The ultimate goal is to use the high-quality beam generated from PBA to drive an X-ray Free-Electron-Laser (XFEL). These high-quality beams could also permit FEL-like instabilities within the plasma wave wake window which could produce coherent radiation.

Developing early applications of PBA will support the ultimate goal of designing and building a next-generation plasma-based TeV-class collider. Indeed, the 2016 Advanced Accelerator Development Strategy Report recognized the importance of developing alternate applications on the way to a collider and acknowledged that “Likely early applications might be an X-ray Free Electron Laser (XFEL) and a gamma-ray source”.

This project will use fully nonlinear, particle-in-cell (PIC) simulations to explore PBA concepts to produce electron beams that can drive an XFEL and to study the XFEL process in regimes where conventional tools are not applicable. It will rely on the relativistic electromagnetic PIC code OSIRIS, which solves Maxwell’s equation and the motion of individual particles (plasma and beam particles) according to their relativistic equations of motion. The newly developed Maxwell solvers and particle pushers that permit simulations with the increased fidelity will be used to explore the generation of coherent X-rays with attosecond duration, 10s of TW power, multi-color structure and spatial-temporal coupled structures. These simulations will support and enhance the closely related FACET-II experiments and may form new proposals based on the simulation results. This project could lead to transformative results on the production of coherent x-rays from PBA and open up new research directions in FEL research.

Title: High fidelity simulation of flow and heat transfer behavior to support conversion of research reactors with involute shaped fuel elements to low enriched uranium

Principal Investigator: Yiqi Yu, Argonne National Laboratory

Co-investigators: Aurelien Bergeron (ANL),
Jeremy Richard Licht (ANL)

ALCC allocation:

Site: Argonne Leadership Computing Facility (ALCF)

Allocation: 500,000 node-hours on Theta

Research Summary:

The Materials Management and Minimization (M3) Reactor Conversion Program of the National Nuclear Security Administration (NNSA) is supporting the conversion of the research reactor 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. There are three research reactors in the world actively engaged in conversion that utilize involute shaped fuel elements: the Oak Ridge National Laboratory (ORNL) High Flux Isotope Reactor (HFIR) located in Tennessee, U.S.A.; the Laue-Langevin Institute (ILL) High Flux Reactor (RHF) located in Grenoble, France and; the Technical University of Munich (TUM) Research Neutron Source Heinz Maier-Leibnitz (FRM II) located in Garching, Germany. These reactors share a similar configuration of coolant channel, which is of extremely thin thickness and involute shape.

Better understanding of flow behavior and heat transfer mechanisms such as coolant mixing in the corner of the coolant channels is of great interest and importance for the design of LEU fuel elements since this is where minimum safety margins occur. It is also of fundamental interest due to the presence of turbulence-driven secondary flows that impact the rates of transfer of heat and momentum

In this project, Large Eddy Simulation (LES) and Direct Numerical Simulations (DNS), where minimal and no turbulence modeling is employed, respectively, will be performed with Nek5000 & NeRS to investigate local flow and heat transfer behavior and provide a benchmark for lower fidelity models as well. Nek5000 is a Gordon Bell and R&D 100 prize-winning code with high accuracy and demonstrated scaling to millions of processors. NekRS is a new version of Nek5000 that is targeting extreme-scale computers, including multicore and many-core platforms and graphics processing units (GPU). The two 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 (NEAS) program. Both codes include advanced algorithms, scalable iterative solvers and high-order discretization that enable scientists to efficiently simulate turbulence on the world's leadership-class supercomputers. These capabilities will reduce the time it takes to validate and certify new reactor designs, enabling an efficient conversion process. The performance of the codes have been extensively tested and demonstrated in previous ALCC projects.