

Title: Particle acceleration in plasma jets: from astrophysics to the laboratory

Principal Investigator: E. Paulo Alves, SLAC National Accelerator Laboratory

Co-investigators: Frederico Fiuza, SLAC National Accelerator Laboratory
Jonathan Zrake, Columbia University
Warren Mori, University of California Los Angeles

ALCC allocation:

Site: Argonne Leadership Computing Facility (ALCF)

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

Research Summary:

Astrophysical jets of magnetized plasma, emanating from supermassive black holes and rapidly rotating neutron stars, are among the most spectacular phenomena in the universe. From their bright radiation emission, which spans the entire electromagnetic spectrum, we infer that these are among the most powerful particles accelerators in the cosmos. Despite decades of observations and theoretical studies, the mechanisms behind these cosmic accelerators remain a long-standing mystery, which continues to fascinate physicists. Beyond its relevance to the understanding of the extreme universe, the study of particle acceleration in plasmas has a significant impact on inspiring the design of new laboratory accelerators for a variety of applications, from fusion to high energy physics research and medical imaging.

An important breakthrough in unveiling how jets may efficiently accelerate high-energy non-thermal particles through current-driven instabilities was achieved via unprecedented 3D *ab initio* massively particle-in-cell (PIC) simulations. This has opened the way for the exploration of this new acceleration mechanism. However, significantly more work is necessary in order to fully understand the range of parameters where this mechanism can operate, and how it may manifest in different astrophysical classes/regions that possess more complex magnetic field structures and plasma compositions.

Studying the dynamics of jets in subrelativistic regimes, including collisional effects, to determine the necessary conditions to observe the particle acceleration mechanisms associated with astrophysical jets in the laboratory will allow us to guide the design of future experiments aiming at probing the plasma processes and the acceleration mechanisms associated with the extreme cosmic accelerators.

Title: Towards Exascale Internal Combustion Engine Simulations with In-Situ Analysis

Principal Investigator: Muhsin Ameen, Argonne National Laboratory

Co-investigators: Saumil Patel (Argonne National Laboratory)

ALCC allocation:

Site: Argonne Leadership Computing Facility (ALCF)

Allocation: 630,000 node-hours on Theta

Research Summary:

To create and sustain leadership in the global transition to a clean energy economy, America needs sustainable transportation through advanced combustion engines with reduced emissions and improved efficiency. Engine manufacturers strive to improve the efficiency and reduce the emissions of internal combustion engines (ICEs) through modifications in engine design. Numerical simulations can play a crucial role by improving the understanding of the in-cylinder combustion process as well as the impact of design changes on the engine performance. Engine manufacturers currently use commercial codes that have inherent limitations in predictive capability and difficulties running at large scale.

The focus of this project is on developing the Nek5000, a leading high-order spectral element, open source code for accurate modeling of fluid turbulence, into an effective Exascale code for high-fidelity ICE simulations. This involves implementing the state-of-the-art submodels for combustion and spray as well as validating against several benchmark engine experiments. Integrating an in-situ analysis framework into Nek5000 will simultaneously reduce data storage requirements and expedite scientific discovery of turbulence, sprays, and combustion phenomena in ICEs.

The capability of the code to accurately capture the complex flow and heat transfer in an ICE with moving piston and valves will be validated using the TCC-III benchmark engine experiments. The implementation of the combustion and spray models into Nek5000 will be validated using benchmark spray and flame experiments. The knowledge gained will improve our understanding of ICEs and help in developing more efficient and cleaner engines.

Title: Accelerated Catalyst Discovery from First Principles Simulations and Machine Learning

Principal Investigator: Rajeev Surendran Assary, Argonne National Laboratory

Co-investigators: Mingxia Zhou (Argonne National Laboratory)
Hieu Doan (Argonne National Laboratory)
Ganesh Sivaraman (Argonne National Laboratory)
Seonah Kim (National Renewable Energy Laboratory)
Josh Schaidle (National Renewable Energy Laboratory)
Larry Curtiss (Argonne National Laboratory)
Alvaro Vasquez Mayagoitia (Argonne National Laboratory)

ALCC allocation:

Site: Argonne Leadership Computing Facility (ALCF)

Allocation: 240,000 node-hours on Theta

Research Summary:

Catalysts enable faster chemical transformation of all industrial processes essential for humankind. New and improved catalysts are desired for converting naturally abundant non-edible biomass and other waste streams to value added chemicals and transportation fuels. The crucial bottleneck associated with the catalyst discovery process is the cost and lack of *a priori* understanding of the properties of the catalyst when interacting with chemicals or reaction intermediates. In particular, chemical transformation such as deoxygenation (removal of oxygen atoms) and the carbon-carbon (C-C) bond formation reaction of bio oil components are essential to produce desired candidates for transportation fuels, are critical reactions that require efficient and durable catalysts. Here we propose to develop and utilize an automated computational approach that utilizing high fidelity first principles simulations coupled with machine learning (ML) to provide guidelines to the catalyst discovery challenges.

This project will utilize density functional theory (DFT) and ML methods that enable exhaustive searches for active catalyst facets (monometallic and bimetallic) and reveal active site motifs for deoxygenation and C-C bond formation reactions. Utilization of massively parallel computational architectures to generate first principle catalyst data along with an active learning approach would provide a unique toolset to search economically and synthetically accessible ‘catalyst subspace’ (Metal carbides, nitrides, and phosphides) that enable the discovery of most promising catalysts. This catalyst space and their selected alloys will be investigated and the guidelines and the design of catalyst candidates obtained from the investigation will be feed into the discovery/verification efforts of ChemCatBio consortium (<https://www.chemcatbio.org/>), which is part of the Energy Materials Network funded by the Energy Efficiency and Renewable Energy (EERE), Bio Energy Technology Office (BETO).

Title: Investigating energy-climate-biogeochemistry sensitivity with the Energy Exascale Earth System Model (E3SM)

Principal Investigator: Katherine Calvin, Pacific Northwest National Laboratory

Co-investigators: Dave Bader (Lawrence Livermore National Laboratory)
Susannah Burrows (Pacific Northwest National Laboratory)
Ruby Leung (Pacific Northwest National Laboratory)
Mathew Maltrud (Los Alamos National Laboratory)
Mark Taylor (Sandia National Laboratory)
Peter Thornton (Oak Ridge National Laboratory)

ALCC allocation:

Site: National Energy Research Scientific Computing Center (NERSC)

Allocation: 490,000 node-hours on Cori

Research Summary:

The Energy Exascale Earth System Model (E3SM) project has a fully functional Earth system model (E3SM v1) that includes innovative dynamic biogeochemistry capabilities in its land and ocean components. This model is uniquely suited to address the critical science question: How does the biogeochemical cycle interact with other Earth system components to influence energy-sector decisions? E3SM goes beyond previous Earth system modeling efforts by including representations of the energy system via the Global Change Assessment Model (GCAM). The goal in this proposed work is to address two important aspects of the climate-biogeochemistry-energy linkage:

- Implications of different energy futures for the biogeochemical cycle through changes in land use, land cover, water availability, and extreme events, and
- Interactions between changes in both the Earth system and the energy sector.

These areas address critical points of uncertainty in our ability to predict the future evolution of the Earth's climate system.

Title: Fundamental turbulence processes governing pedestal structure in fusion devices

Principal Investigator: Jeff Candy, General Atomics

Co-investigators: Walter Guttenfelder (Princeton Plasma Physics Laboratory)
Arash Ashourvan (Princeton Plasma Physics Laboratory)
Emily Belli (General Atomics)
Alessandro Marinoni (Massachusetts Institute of Technology)
Chris Holland (University of California, San Diego)
George Fann (Oak Ridge National Laboratory)
Igor Sfiligoi (San Diego Supercomputing Center)

ALCC allocation:

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 300,000 node-hours on Summit

Research Summary:

H-mode regime is the most promising operating scenario envisioned to achieve nuclear fusion in a tokamak configuration. An H-mode features a narrow region near the plasma boundary – an edge transport barrier (ETB) – where a strongly-sheared plasma flow reduces the turbulence intensity. This low turbulence causes development of steep gradients -- also known as pedestals -- in temperature and density profiles, which substantially increase the core density and temperature. Pedestals are thus key to achieving the confinement necessary to sustain thermonuclear fusion. A variety of H-modes have been measured that exhibit differing pedestal structure, turbulent fluctuations, and relative strengths of particle to heat transport. The existence of these different pedestal characteristics, each with potential benefits for reactors, motivates the study of pedestal structure across a variety of regimes.

Reduced models to predict the pedestal pressure have been developed and validated on conventional tokamaks. A successful pedestal structure model (EPED) for ELMy H-mode uses two constraints: (1) an MHD peeling-ballooning instability limit and (2) a transport limit that sets the pressure gradient. The transport limit is currently modeled via a proxy for the kinetic ballooning mode (KBM) turbulence threshold. However, a number of scenarios exist where EPED does not work well. This makes it imperative to develop transport models based on first-principles simulations that work in the H-mode pedestal. Construction of such model first requires covering the relevant parameter space using first-principles electromagnetic gyrokinetic simulation.

Despite years of successful gyrokinetic validation in the core of experiments, the community has only recently started to model the pedestal of experiments with true gyrokinetic fidelity. Capturing the cross-scale coupling from ion to electron scales requires multiscale simulations. Few multiscale simulations have been reported in the core region as they require leadership-class computational resources. No such simulations have been performed in the edge pedestal region. Anticipating many of these challenges, CGYRO was developed with a sonic rotation capability and the Sugama collision operator suitable for extreme edge plasma conditions.

This project proposes to exploit the leadership-scale edge-plasma turbulence simulation capabilities of CGYRO to accurately simulate the pedestal region of high confinement plasmas. Simulation data results will be made available to the community as part of a gyrokinetic database hosted by the AToM SCiDAC-4 project.

Title: Cryospheric Physics in E3SM: Impacts of Antarctic Ice Shelf Melting, Southern Ocean Resolution, and Sea Ice Coupling on Global Climate

Principal Investigator: Darin Comeau, Los Alamos National Laboratory

Co-investigators: Xylar Asay-Davis (Los Alamos National Laboratory)
Matthew Hoffman (Los Alamos National Laboratory)
Wuyin Lin (Brookhaven National Laboratory)
Mark Petersen (Los Alamos National Laboratory)
Stephen Price (Los Alamos National Laboratory)
Andrew Roberts (Los Alamos National Laboratory)

ALCC allocation:

Site: National Energy Research Scientific Computing Center (NERSC)

Allocation: 360,000 node-hours on Cori

Research Summary:

Antarctica and the Southern Ocean are critical regions of the global climate system. Their response to climate change is among the least understood in the climate system, but the potential global ramifications of their responses are among the most severe. Although the Southern Ocean covers less than a third of the global ocean, it is the primary sink for carbon and heat uptake from the atmosphere. Interactions between the ocean, fringing ice shelves from the Antarctic ice sheet, and sea ice produce some of the coldest waters on the planet, which play a crucial role in driving global ocean circulation. The Antarctic ice sheet holds many tens of meters of sea-level equivalent, and ice shelves directly exposed to the ocean are integral to restraining the grounded ice behind them. Severe melting of ice shelves could cause collapse of parts of the ice sheet, leading to rapid sea level rise that may displace a large part of the world's low-lying coastal population, as well as threaten vulnerable infrastructure.

Despite its importance, the Antarctic remains poorly represented in most Earth system models. Projecting future sea level rise from Antarctica using a global, coupled Earth system model that realistically simulates Antarctic and Southern Ocean processes is one of three primary science drivers for the Department of Energy's Energy Exascale Earth System Model (E3SM). This proposed suite of E3SM simulations will be the first climate simulations of the recent historical period and hypothetical future climate conditions to include the circulation of ocean currents and melting beneath the floating ice shelves surrounding Antarctica.

The project aims to simulate over 2000 years of the fully coupled (ocean, atmosphere, land, sea ice) Antarctic system at moderate resolution to investigate the impacts on global climate predictions from representing ocean circulation beneath ice shelves and improved sea ice physics. In addition, ~200 years of partially coupled (ocean and sea ice only) simulations with high resolution, regional refinement around Antarctica will demonstrate the value in better resolving key features in the Antarctic region. The fully coupled simulations follow two independent tracks: 1) exploring the impact of prognostic ice shelf melt fluxes on the Antarctic and the global climate system through standard climate model scenarios; 2) focusing on improving the representation of sea ice in E3SM, through better representation of ocean-ice-atmosphere coupling and sea ice physics. Ensembles of simulations will be critical to exploring the space of natural variability in the simulated climate. Together these sets of simulation will position E3SM and DOE at the forefront of modeling the Antarctic region, with broader benefits for the simulation of polar climate beyond the immediate focus of this research.

Title: Nuclear Physics from the Standard Model

Principal Investigator: William Detmold, Massachusetts Institute of Technology

Co-investigators: Phiala Shanahan, Massachusetts Institute of Technology
David Murphy, Massachusetts Institute of Technology
Michael Wagman, Massachusetts Institute of Technology
Zohreh Davoudi, University of Maryland
Assumpta Parreño, University of Barcelona
Frank Winter, Jefferson Laboratory
Kenneth Roche, Pacific Northwest National Laboratory

ALCC allocation:

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 320,000 node-hours on Summit

Research Summary:

Nuclei make up the majority of the visible matter in the universe; knowledge of nuclear structure and interactions thus underlies our understanding of the world we live in. Our knowledge of nuclei is derived from a century of innovative experiments which revealed nuclei, formed of protons and neutrons, at the cores of atoms. Over the last decades, the underlying structure of the protons and neutrons in terms of more fundamental constituents, quarks and gluons, has also been unveiled. Our theoretical understanding of the emergence of nuclei from these quarks and gluons—which are, to the best of our understanding, the fundamental degrees of freedom of nature—lags this impressive phenomenology.

New and more precise experiments are exploring the frontiers of nuclear physics and deepening our theoretical understanding of nuclear structure and interactions. At the same time, a host of current and future experiments are seeking to probe the limits of known physics using nuclei as targets. These experiments require detailed knowledge of nuclear structure to maximize their impact. Providing first-principles calculations of the properties and interactions of light nuclei from the Standard Model of particle physics addresses these missions. The NPLQCD Collaboration has pioneered efforts to extend the scope of LQCD to nuclei and has undertaken the first calculations of the structure and interactions of light nuclei directly from the Standard Model. LQCD will be used to compute the properties of light nuclei from first principles with close-to-physical values of the quark masses.

Title: Abacus2020: N-body Simulations for Precision Cosmology with DESI

Principal Investigator: Daniel Eisenstein, Harvard University

Co-investigators: Philip Pinto, University of Arizona
Sownak Bose, Harvard University
Lehman Garrison, Harvard University
Nina Maksimova, Harvard University
Boryana Hadzhiyska, Harvard University

ALCC allocation:

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 300,000 node-hours on Summit

Research Summary:

ABACUS 2020 will produce an enormous suite of cosmological N-body simulations in support of the large-scale structure modeling requirements of the analyses of the Dark Energy Spectroscopic Investigation (DESI). DESI is a next-generation cosmological project that will create the world's largest galaxy redshift survey. From the clustering of these maps, the baryon acoustic oscillation signature can be used to measure the cosmic expansion history to 0.3% accuracy, with which we can study the evolution of dark energy. Further, the redshift-space distortions of the galaxy clustering can be used to measure the large-scale velocities and compare them to the predictions of general relativity, thereby probing the question of whether the accelerating cosmic expansion might be due to modifications of gravity.

The cosmological tests enabled by DESI data depend upon careful modeling of the non-linear evolution of large-scale structure. For this, high-accuracy N-body simulations are a critical ingredient. Investigating a range of cosmologies will train accurate interpolation in the cold-dark matter model space and provide an opportunity to conduct mock analyses on blinded data sets as a validation of the DESI Collaboration's techniques. The planned suite will meet the stated requirements of the DESI Collaboration and provide an additional basis for robust interpolation in the cosmological model space.

Title: Semileptonic B- and D-meson form factors with high precision

Principal Investigator: Aida El-Khadra, University of Illinois, Urbana-Champaign

Co-investigators: Carleton DeTar (University of Utah)
Elvira Gámiz (University of Granada)
Steven Gottlieb (Indiana University)
Andreas Kronfeld (Fermilab)
John Laiho (Syracuse University)
Doug Toussaint (University of Arizona)
Ruth Van de Water (Fermilab)

ALCC allocation:

Site: Argonne Leadership Computing Facility (ALCF)
Allocation: 400,000 node-hours on Theta
Site: National Energy Research Scientific Computing Center (NERSC)
Allocation: 330,000 node-hours on Cori

Research Summary:

Many ongoing and upcoming experiments in particle physics aim, through precision measurements, to observe phenomena inconsistent with the Standard model. To be successful, comparably precise theoretical calculations are necessary, even when the challenge of nonperturbative quantum chromodynamics (QCD) is involved. Fortunately, an established method exists to meet this demand: lattice QCD. The Fermilab Lattice and MILC Collaborations have been especially successful in carrying out calculations required for quark-flavor experiments, such as reducing the total theoretical uncertainty on leptonic decay amplitudes down to the sub-percent level with all sources of error under control.

A next, and more crucial step, is to improve the precision of semileptonic decay amplitudes to this level. They play a central role in many phenomenological applications. For example, recent measurements give hints of (previously unobserved) lepton-flavor violation. The Belle II experiment at the KEK Laboratory in Japan and the LHCb experiment at the CERN Laboratory in Switzerland are pursuing new measurements that will result in unprecedented precision for many relevant observables, including probes of lepton-flavor universality. Their interpretation as tests of the Standard Model or discovery of new physics requires theoretical calculations more precise than those done to date. This project targets semileptonic decays with a B or D meson in the initial states and a pion or kaon in the final state, for which the transition matrix elements—known as form factors—for the vector, scalar, and tensor currents will be computed. Leadership-class computational resources are essential to reduce the theoretical uncertainties to the target level.

Title: Large-scale kinetic simulations of particle acceleration in laser-driven shocks

Principal Investigator: Frederico Fiuza, SLAC National Accelerator Laboratory

Co-investigators: Anna Grassi (Stanford University)
Warren Mori (University of California, Los Angeles)

ALCC allocation:

Site: National Energy Research Scientific Computing Center (NERSC)

Allocation: 395,000 node-hours on Cori

Research Summary:

What is the origin of cosmic rays? What are their dominant acceleration mechanisms? Can these mechanisms be reproduced in the laboratory? These longstanding scientific questions, closely tied to extreme plasma physics processes, continue to fascinate physicists. Beyond relevance to understanding of the extreme universe, the study of particle acceleration can offer a significant impact on applications from fusion plasma diagnostics to medical imaging.

Particle acceleration in high-energy-density (HED) plasma environments is associated with processes such as collisionless shocks, magnetic reconnection, and turbulence. It involves a complex and highly nonlinear interplay between the dynamics of flows, magnetic fields, and non-thermal particle distributions, which operate at different scales. A transformative advance in our understanding of these requires a combination of fully kinetic simulations and controlled laboratory experiments.

This project will use 3D *ab initio* massively parallel particle-in-cell (PIC) simulations to explore the physics of particle acceleration and magnetic field dynamics in HED plasmas. Particular focus will be given to the study of particle injection and acceleration in collisionless shocks using conditions from ongoing high-power laser experiments at the National Ignition Facility (NIF).

PIC codes provide a self-consistent description of the interplay between plasma particles and the electromagnetic fields they produce. Using unprecedented large-scale 2D and 3D simulations we aim to identify the processes responsible for the onset of particle acceleration, characterize their efficiency, their dependence on the plasma conditions, and their signatures in terms of fields and non-thermal particles. We will apply this understanding to design and guide the interpretation of ongoing HED laser experiments at the National Ignition Facility (NIF), where, through appropriate scaling laws, these processes can be directly probed and connected with astrophysical plasmas.

This work will have a critical impact in the understanding of particle acceleration in high energy density plasmas, which is important to DOE's program in Frontiers of Plasma Science as recently highlighted in the Fusion Energy Sciences "Frontiers of Plasma Science Workshops" report. Our results will support NIF experiments that could not succeed without fully kinetic 3D simulations, and will dramatically change our ability to understand and model magnetic field dynamics and particle injection in laboratory and astrophysical plasmas.

Title: Energetics of Collisionless Plasmas in the Laboratory and Space

Principal Investigator: Will Fox, Princeton Plasma Physics Laboratory

Co-investigators: Kai Germaschewski (University of New Hampshire)
Derek Schaeffer (Princeton University)
Sam Totorica (Princeton University)
Amitava Bhattacharjee (Princeton Plasma Physics Laboratory)

ALCC allocation:

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 400,000 node-hours on Summit

Research Summary:

Understanding the processes which convert energy between different forms—magnetic energy, flow energy, thermal energy, and accelerated particle populations—is a key unifying question of plasma astrophysics. The key processes—magnetic field generation by Weibel instability and the Biermann process; magnetic field destruction by magnetic reconnection; and flow thermalization by magnetized shocks—can potentially explain some of the most important questions in astrophysics, including acceleration of high energy particles such as cosmic rays, which are the highest-energy particles observed in the universe, millions of times more energetic than can be produced at earth-bound particle accelerators. Recent simulations, in combination with laboratory experiments at national laser facilities (NIF and OMEGA), have demonstrated significant advances toward understanding these important problems. One example is the first time ever generation of high-Mach number magnetized shocks in the laboratory, believed to be the key mechanism for cosmic ray particle acceleration.

Magnetic reconnection is ubiquitous in fusion, space and astrophysical plasmas, playing a central role in liberating stored magnetic energy in phenomena as diverse as sawtooth crashes in fusion plasmas, magnetospheric substorms, and solar flares. The Weibel instability is of great interest as one of the few known processes that generates magnetic field *de novo* in laboratory and astrophysical plasmas, enabling formation of collisionless shocks driven by astrophysical explosions such as supernova remnants and gamma ray bursts. Both problems are grand-challenge, multi-scale problems, where macroscopic flows and fields are coupled to microscale current sheets (in the reconnection case) and shock transition layers (for collisionless shocks). Both processes (reconnection and shocks) have been proposed to accelerate cosmic rays.

This project will build off previous work to study (1) fast magnetic reconnection by kinetic-scale plasmoid instabilities in laser-produced plasmas; (2) particle acceleration by collisionless shocks in laser-plasmas; and (3) simulations of coupled shocks and reconnection applied to the Earth's magnetosphere.

The simulations use the Plasma Simulation Code (PSC), a particle-in-cell code which has GPU capability and has been benchmarked on Summit. Fully-kinetic plasma simulations conducted on leadership-scale machines such as Summit can provide 1-1 matching of plasma parameters with experiments. Results will be compared with ongoing experiments on magnetic reconnection, Weibel instability, and collisionless shocks in laser-produced plasmas conducted at the Omega-EP facility at the University of Rochester and the National Ignition Facility at LLNL.

Title: Large Scale Numerical Simulations of Polymer Nanocomposites

Principal Investigator: Gary S. Grest, Sandia National Laboratories

Co-investigators: Shengfeng Cheng (Virginia Polytechnic Institute and State University)
Sanat Kumar (Columbia University)
Dvora Perahia (Clemson University)
Michael Rubinstein (Duke University)

ALCC allocation:

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 175,000 node-hours on Summit

Site: National Energy Research Scientific Computing Center (NERSC)

Allocation: 270,000 node-hours on Cori

Research Summary:

The potential of nanoparticles has been long realized; however, integrating them in different devices remains a challenge. Multi-million atom molecular dynamics simulations will probe ways to overcome the major barriers to integrating nanoparticles into a range of advanced devices, controllably dispersing and organizing them within durable matrices while retaining their unique properties. Zooming into the interaction regions between the nanoparticles and the surrounding polymeric matrix will provide the needed molecular level understanding of the assembly process.

A specific goal is to resolve the factors that control the assembly of multiple nanoparticles in the bulk and in thin films. As producing well-dispersed polymer nanocomposites is a major hurdle for making new materials, studies will focus on identifying the parameters that control nanoparticle dispersion. One approach for dispersing nanoparticles in a polymer matrix is to separately optimize the enthalpic and entropic components by end-grafting a mixture of short and long chains to the nanoparticle. The short chains are to screen the enthalpic interaction between nanoparticles while the long chains mix with the polymeric matrix, producing enhanced mechanical and thermal stability. As many polymer films are made by first dissolving the nanoparticles and polymer in a solvent and evaporating the solvent, modeling the effect of solvent evaporation rate on the dry polymer film is of interest.

Motivated by the numerous potential applications of polymer nanocomposites with unique new properties, large-scale molecular dynamics simulations will be used to define the interaction of nanoparticles and control their miscibility in polymer melts and their self-assembly that will serve as building blocks for new materials and devices. Access to petascale computational resources will enable capture of the essential time and length scales that dominate the physics of nanoparticle assembly and integration into polymer matrices.

Title: Prediction of Engine Knock in a Gasoline Direct Injection (DI) Engine

Principal Investigator: Ronald O. Grover, Jr., General Motors

Co-investigators: Jian Gao (General Motors)
Muniappan Anbarasu (General Motors)
Yangbing Zeng (General Motors)
Charles Finney (Oak Ridge National Laboratory)
Wael Elwasif (Oak Ridge National Laboratory)
Dean Edwards (Oak Ridge National Laboratory)
Russell Whitesides (Lawrence Livermore National Laboratory)
Keith Richards (Convergent Science, Inc.)

ALCC allocation:

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 260,000 node-hours on Summit

Research Summary:

Regulations and market pressures are driving the transportation industry to develop new designs and operating strategies for light-duty gasoline engines capable of higher efficiency, greater performance, and lower pollutant emissions. Down-sized boosted direct-injection gasoline engines and high compression ratio strategies are limited by engine knock at the high-speed, high-load conditions. Engine knock occurs when the end-gas in front of the spark-ignited, propagating flame is heated and compressed to its auto-ignition point resulting in uncontrolled combustion and high cylinder pressures which can damage the engine. Several factors, including turbulent mixing and flame propagation, create cycle-to-cycle and cylinder-to-cylinder variability in end-gas conditions and knock propensity. As a result, engine operation under knocking conditions is highly variable. Available computing resources have traditionally limited the effectiveness of predictive knock modeling efforts due to the need for high-fidelity models based on first-principles physics and chemistry with sufficient spatial and temporal detail to resolve small regions of high knock propensity within the end-gas where auto-ignition first occurs. This project proposes to develop a high-fidelity engine model based on the commercial CONVERGE™ CFD software tool to more accurately predict occurrence of knock at a high-speed, high-load operating point.

An LES turbulence model with high spatial resolution will be used to resolve turbulence and capture small-scale regions of high knock propensity within the end-gas. The Zero-RK GPU-enabled chemical kinetics solvers developed by Lawrence Livermore National Laboratory (LLNL) will be used to enable detailed chemistry to capture the initiating auto-ignition reactions. High temporal resolution will enable evaluation of the actual peak cylinder pressures during the high-frequency pressure oscillations associated with knock. A statistics-based parallel perturbation method (previously shown to be effective at capturing cyclic variability at dilute conditions) will be used to perform a combination of parallel and serial engine cycle simulations to develop a statistical sampling of the resulting cyclic variability based on 300 cycles. The accuracy of the model predictions will be assessed through comparison to experimental measurements from a single-cylinder boosted spark-ignited (SI) engine. Advances made in the proposed project are anticipated to help guide the transportation industry in meeting consumer needs while achieving performance, fuel efficiency, and emissions goals.

Title: Extending the Discovery Space of Cosmological Surveys

Principal Investigator: Salman Habib, Argonne National Laboratory

Co-investigators: Katrin Heitmann (Argonne National Laboratory)
Nicholas Frontiere (Argonne National Laboratory)
Andrew Hearin (Argonne National Laboratory)
Esteban Rangel (Argonne National Laboratory)
Thomas Uram (Argonne National Laboratory)

ALCC allocation:

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 290,000 node-hours on Summit

Research Summary:

Large-area surveys of the sky reaching ever deeper into the observable Universe are providing us with a remarkable composite data set, covering different frequency ranges as well as different ways of probing the dynamics and matter content of the Universe. The level of statistical control is quite remarkable with spatio-temporal summary statistics being measured at the percent level or even smaller. The drive to improve how well we can measure and understand the Universe is strongly motivated by the fact that the simple model nicely describing our current observations, Lambda-CDM, is widely expected to break down, but in ways that are hard to predict. Because Lambda-CDM is so successful, looking for variations from its predictions requires serious control over a host of systematic effects, both observational and modeling-related, and extension of the reach of cosmological probes into domains where the model has so far not been strongly tested.

Cosmological simulations are currently the only known way to provide the level of detailed predictive and modeling power that the above program requires. This project aims to carry out a state-of-the-art gigaparsec (Gpc)-scale cosmological simulation that will allow the modeling of galaxy statistics relevant for all current and near-future surveys without compromises arising from limitations on simulation volume or on the mass resolution of the simulation. The simulation will be run with the GPU-accelerated HACC code, of which we are the main developers, and which has recently been proven on Summit, out to the full scale of the system.

Along with the main simulation capability, we have implemented a major suite of galaxy modeling and analysis tools that are currently unique. These tools allow us to create highly realistic synthetic galaxy catalogs and sky maps using a combination of forward simulation and sophisticated empirical modeling. The results obtained with this project will allow precision cosmology probes such as galaxy clustering (in real and velocity space) and weak gravitational lensing to go to the small length scales that have up to now resisted a robust analysis, despite being extremely information-rich.

The results of our program will strongly impact all DOE-supported cosmological surveys – CMB-S4, DES, DESI, eBOSS, LSST, and SPT. In addition, we are members of the science team for the very recently approved NASA mission SPHEREx, the first all-sky spectral survey of the Universe. It is exciting to contemplate a future in which the combination of data from SPHEREx and the ground-based surveys can result in a major transformation in our knowledge of the Universe. This ALCC project will provide unique leverage in bringing this future very much closer.

Title: The next leap forward in LSST sky simulations

Principal Investigator: Katrin Heitmann, Argonne National Laboratory

Co-investigators: Rachel Mandelbaum (Carnegie Mellon University)
Chris Walter (Duke University)
James Chiang (SLAC National Accelerator Laboratory)
Salman Habib (Argonne National Laboratory)
Michael Jarvis (University of Pennsylvania)
Antonio Villarreal (Argonne National Laboratory)
Scott Daniel (University of Washington)
Tom Uram (Argonne National Laboratory)
Joshua Meyers (Lawrence Livermore National Laboratory)

ALCC allocation:

Site: Argonne Leadership Computing Facility (ALCF)

Allocation: 400,000 node-hours on Theta

Research Summary:

Understanding the nature of dark energy is at the core of DOE's High Energy Physics mission. The Large Synoptic Survey Telescope (LSST), a ten-year imaging survey starting in 2022, is a cornerstone of the U.S. cosmological community's efforts to understand dark energy, the cause of the accelerated expansion rate of the Universe. The Dark Energy Science Collaboration (DESC), an international collaboration with more than 800 members, was formed specifically to prepare to study dark energy with LSST. A key part of the preparation for LSST data arrival involves making sure that we can properly correct for systematic uncertainties due to astrophysical and observational nuisances, and therefore fully take advantage of the exquisite statistical constraining power of the data to measure dark energy with percent-level precision. Building on established simulation infrastructure, planned upgrades in features and performance are expected to enable us to reach the ~1000 square degree scale, a factor of ~3 beyond what is currently possible.

The resulting dataset and associated databases produced will be documented and made available (on request) to the wider LSST science community, thus enabling a broader range of science cases. The infrastructure improvements that are developed for this simulation campaign are part of the longer path to science readiness undertaken by the LSST DESC.

Title: Epistatic signatures for Climatic Adaptation in Plant Genomes

Principal Investigator: Daniel Jacobson, Oak Ridge National Laboratory

Co-investigators: Wayne Joubert (Oak Ridge National Laboratory)
Joao Gabriel Felipe Machado Gazolla (Oak Ridge National Laboratory)
David Kainer (Oak Ridge National Laboratory)
Jared Streich (Oak Ridge National Laboratory)
Michael Garvin (Oak Ridge National Laboratory)
Jonathon Romero (University of Tennessee)
Ashley Cliff (University of Tennessee)

ALCC allocation:

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 400,000 node-hours on Summit

Research Summary:

The Department of Energy has launched ambitious research agendas to accelerate the development of domestic, renewable alternatives to liquid fuels. Fuels converted from cellulosic biomass offer one alternative to conventional energy sources, and supplementing fossil fuels with ethanol or other biofuels derived from bioenergy crops would benefit the economic growth and energy security of the United States. To develop cost-effective bioenergy crops, biological and technological barriers need to be overcome in areas such as biomass yield (dry tons/ha), biomass deconstruction (which ultimately enables access to a plant's sugars) and sustainability (including plant-microbiome associations and responses to biotic and abiotic stress).

Short rotation woody crops, such as *Populus* species and hybrids, are expected to be an important renewable feedstock for biofuels due to their relatively low delivered cost. Phenotypes associated with sustainability are also of high importance to DOE's biofuels programs. Feedstocks like *Populus trichocarpa* and model plants like *Arabidopsis thaliana* are derived from complex genetic systems that are composed of pleiotropic and epistatic functional networks of interacting molecules and macro-molecules. The subsequent phenotypes, including a plant's adaptation to abiotic conditions and the interactions with its phytobiome, are the result of orchestrated, hierarchical, heterogeneous collections of expressed genomic variants regulated by and related to biotic and abiotic signals. The effects of these variants can be viewed as the result of historic selective pressure and current environmental as well as epigenetic interactions, and, as such, their co-occurrence can be seen as genome-wide associations in a number of different manners.

This project aims to use two different algorithmic approaches to find the genome-wide associations that form the complex, epistatic architectures responsible for these traits. The goal is to build on population genomics and climatypes information in order to better discover the complex genetic architectures responsible for abiotic adaptation in *Populus trichocarpa* and *Arabidopsis thaliana*. These will form important components of systems biology models that will further our understanding of these two complex biological systems. Understanding the molecular mechanisms of these adaptations is important to the DOE-funded Plant Microbial Interface Scientific Focus Area, Feedstock Genomics program and the Bioenergy Research Centers. Understanding these adaptations in the model plant *Arabidopsis thaliana* is of interest to the entire plant research community.

Title: First-principles design of thermodynamically stable 2D quantum devices

Principal Investigator: Jacek Jakowski, Oak Ridge National Laboratory

Co-investigators: Jerzy Bernholc (North Carolina State University)
Mina Yoon (Oak Ridge National Laboratory)
David Lingerfelt (Oak Ridge National Laboratory)

ALCC allocation:

Site: Oak Ridge Leadership Computing Facility (OLCF)
Allocation: 200,000 node-hours on Summit

Research Summary:

It is widely believed that Moore's Law shrinkage of computer circuitry is approaching physical limits, beyond which conventional silicon-based integrated circuits will cease to function reliably. Further extensions of the classical silicon circuitry roadmap appear impossible because of quantum effects, such as undesired current leakage due to quantum tunneling, or unreliable doping that will depend on the atomic positions of the dopants relative to line edges. The beyond-Moore's law advances in computer devices will thus depend on harnessing quantum effects to increase the speed of computing, decrease the energy cost of computation per elementary operation, and guarantee the integrity and reliability of computational results.

This project aims to simulate quantum systems and circuits for quantum information processing and computing well beyond the scales possible with current silicon technology. We focus on three classes of 2D quantum systems that will be needed in future computers and circuits: (i) single photon emitters, necessary to communicate between various parts of quantum computers, (ii) topological electrides and topological quantum dots that can be used to design high temperature qubits, and (iii) high-speed ballistic tunneling transistors. In photonic quantum computing devices, the generation and manipulation of quantum information are based on the utilization of single photon emitters. Semiconductor quantum dots have emerged as promising candidates because they allow for a quantum interface between stationary spin qubits and single photons that handle communications.

Eventually, the above elements could be used to form quantum computers manipulating quasiparticles called anyons, and high-speed interfaces between classical and quantum parts of future generations of computers, which will use "quantum accelerators" to enable novel, currently unimaginable types of computation.

Title: Study of a Disrupting Plasma in ITER

Principal Investigator: Stephen Jardin, Princeton Plasma Physics Laboratory

Co-investigators: Nate Ferraro, Princeton Plasma Physics Laboratory
Brendan Lyons, General Atomics

ALCC allocation:

Site: National Energy Research Scientific Computing Center (NERSC)

Allocation: 395,000 node-hours on Cori

Research Summary:

The ITER device, now being constructed in Cadarache, France, is a monumental step towards the realization of fusion power. The \$40+ B experiment, being financed by 7 international governments including the U.S., is the same scale as a follow-on fusion power plant. Success in this endeavor will be a turning point in the quest for a sustainable, carbon free, and safe source of energy for the planet.

The ITER tokamak plasma will be a donut-shaped fully ionized gas of Deuterium and Tritium which will carry about 15 MA of electrical current (\mathbf{J}). In normal operation, this current flows in a direction largely parallel to the 6 T magnetic field (\mathbf{B}) creating a nearly force-free equilibrium configuration. However, if the plasma suffers a disruptive plasma instability, the plasma current is transferred to nearby conductors (the vacuum vessel) where it is no longer nearly parallel to the magnetic field. This will cause large $\mathbf{J} \times \mathbf{B}$ forces on the vessel and other conducting components, potentially damaging the machine.

A disruption is an abrupt termination of a tokamak discharge due to large scale plasma instabilities. Leadership-class computational resources are required to generate a realistic simulation capability of a disrupting tokamak plasma and to apply this to ITER.

Title: Higher order cumulants of net-charge fluctuations

Principal Investigator: Frithjof Karsch, Brookhaven National Laboratory and Bielefeld University

Co-investigators: Swagato Mukherjee, Brookhaven National Laboratory
Alexei Bazavov, Michigan State University
Heng-Tong Ding, Central China Normal University, China
Peter Petreczky, Brookhaven National Laboratory
Patrick Steinbrecher, Brookhaven National Laboratory

ALCC allocation:

Site: National Energy Research Scientific Computing Center (NERSC)

Allocation: 800,000 node-hours on Cori

Research Summary:

A primary goal of studying the phase diagram of strongly interacting matter at high temperature and/or densities is to establish whether or not a true second order phase transition exists, at which the well-established rapid but smooth crossover transition at low net baryon-number density turns into a first order phase transition with discontinuities in basic thermodynamic observables. The existence of a second order phase transition point, the so-called critical point, will give rise to large fluctuations of conserved charges and their higher moments. Experimental evidence for such large fluctuations is searched for in heavy ion experiments, in particular with the next Beam Energy Scan (BES-II) performed at the Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory during 2019-2021.

Current lattice QCD results strongly disfavor a critical point at small values of the baryon chemical potential and temperature larger than about 150 MeV, so calculations will be focused on lower temperatures and aim at results for high order cumulants. These results will not only enhance our fundamental knowledge about strong-interaction matter but also guide the experimental explorations of the QCD phase diagram now continuing with BES-II at RHIC.

Title: Scaling LHC proton-proton collision simulations and Machine Learning for the ATLAS experiment

Principal Investigator: Eric Lancon, Brookhaven National Laboratory

Co-investigators: Douglas Benjamin (Argonne National Laboratory)
Abid Malik (Brookhaven National Laboratory)
Paolo Calafiura (Lawrence Berkeley National Laboratory)

ALCC allocation:

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 60,000 node-hours on Summit

Site: National Energy Research Scientific Computing Center (NERSC)

Allocation: 400,000 node-hours on Cori

Research Summary:

The Standard Model (SM) of particle physics is being tested with ever increasing precision in proton-proton collisions by the ATLAS experiment at the Large Hadron Collider (LHC). The ATLAS experiment is one of the major tools being used to search for physics beyond the SM at the Energy Frontier. The research focus is transitioning from quick searches for new particles, driven by large increases in beam energy or integrated luminosity, to precision measurements of SM predictions, as increases become more gradual. These measurements will guide both experimentalists and theorists in the search for physics beyond the SM. The precision measurements performed at the LHC need to be compared with precision model simulations which can be computationally demanding. DOE supercomputers are becoming a necessity to ensure the US physics program at the LHC meets its scientific goals. The computing and data challenges will increase by an order of magnitude with the High-Luminosity LHC looming on the horizon.

A key challenge facing LHC over the next decade is the prohibitive cost of computing to generate enough simulated event samples for its diverse physics program. Machine learning methods have the potential to achieve accurate simulations at a fraction of the cost of existing full physics-based simulations. However, producing high quality machine learning models also requires large computational resources. The ALCC has been vital in driving optimization and development of internal software frameworks and workflow engines within the collaboration to use supercomputers more efficiently. This new award will drive additional software optimization efforts to improve efficiency and increase production.

Title: Molecule-Informed Continuum Direct Numerical Simulation Method in Non-Equilibrium Combustion

Principal Investigator: Myoungkyu Lee, Sandia National Laboratories

Co-investigators: Jacqueline H. Chen (Sandia National Laboratories)
Michael A. Gallis (Sandia National Laboratories)

ALCC allocation:

Site: Oak Ridge Leadership Computing Facility

Allocation: 250,000 node-hours on Summit

Research Summary:

All transport phenomena originate from nonequilibrium of selected states. If everything is in equilibrium, a state will not change without any external forcing source. Considering nonequilibrium of all states renders most engineering and scientific problems too complicated to be solved. Hence, we adopt the concept of quasi-equilibrium to simplify most macroscopic problems. Study of reactive flows takes advantage of the quasi-equilibrium assumption. For example, we assume the energy state of molecules is in local equilibrium. This assumption is valid when the characteristic length and time scales of given reactive flow is much longer than the scales of molecular movements. However, such an assumption is no longer valid if the system includes certain phenomena which change the property of reactive flow extremely rapidly, e.g., shock, strong instability, etc.

This project aims to focus on the effect of thermal nonequilibrium in chemically reactive flows. The reactive flows with extreme thermal and chemical gradients induced by the Kelvin-Helmholtz instability will be studied for hydrogen/oxygen chemistry. Two primary simulation methods will be used: Direct Simulation Monte Carlo (DSMC) and Direct Numerical Simulation (DNS). DNS will represent flows with quasi-equilibrium in the continuum regime, but is not suitable to study the effect of the nonequilibrium process. On the other hand, DSMC simulates discontinuous particles (or simulators), which represents parcels of molecules, and their stochastic movement, collision, and reactions describing nonequilibrium processes at the molecular level, but cannot be used to simulate the macroscopic problem because of the enormous computational power demands.

Bridging DSMC and DNS will circumvent deficiencies of each method. The project will start with DSMC to simulate flows with strong thermal and chemical nonequilibrium, and then will build the macroscopic reaction library. The chemical reaction rate is usually a function of a single temperature in quasi-equilibrium. Instead, the chemical reaction rate will be described as a function of multi-temperatures which represent the different energy modes of molecules; translational, rotational and vibrational modes. We will exploit the different reduced-order modeling techniques to build such a reaction library. Finally, the governing equations for the different temperatures and new reaction library will be implemented in the existing DNS code, to facilitate comparisons from solutions in equilibrium and in nonequilibrium.

DSMC is implemented by using SPARTA and DNS by using S3D-Legion. These codes were developed at Sandia National Laboratories and show excellent performance on modern heterogeneous high-performance computing systems. The scientific findings from this project will be presented in multiple journal articles and conferences.

Title: High-Fidelity Multi-Physics Simulation of Pebble bed cores

Principal Investigator: Elia Merzari, Argonne National Laboratory

Co-investigators: Paul Fischer (Argonne National Laboratory)
Misun Min (Argonne National Laboratory)
Derek Gaston (Idaho National Laboratory)
Paul Romano (Argonne National Laboratory)

ALCC allocation:

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 305,000 node-hours on Summit

Research Summary:

Nuclear power is a reliable, carbon-free resource capable of meeting the energy needs of our nation and the world. Amid an increase in cost of traditional nuclear reactor designs, in recent years advanced reactor concepts based on molten salt coolants have emerged as an attractive solution capable of greatly increasing inherent safety while reducing cost. An example is the fluoride-salt-cooled, high-temperature reactor (FHR) that can be deployed with robust safety, affordable cost, and flexible operation to accommodate the expansion of variable renewables.

Leadership-class computational resources are necessary to generate high-fidelity datasets for a reference FHR pebble bed core design. We will conduct high-fidelity multiphysics simulations with Nek5000, OpenMC and BISON, state-of-the-art codes for the high fidelity three-dimensional modeling of fluid flow, heat transfer and neutronics. In particular we will perform wall-resolved large eddy simulation (LES) of the flow in random pebble beds. The simulations will use geometric information available in the open literature and will generate vast datasets (e.g., first- and second-order statistics, correlations) that can be used to inform lower-fidelity models (porous media, systems codes and diffusion codes).

Title: Nuclear Energy Industry Validation of Nek5000: ALAIN and HYMERES

Principal Investigator: Aleksandr Obabko, Argonne National Laboratory

Co-investigators: Elia Merzari (Argonne National Laboratory)
Haomin Yuan (Argonne National Laboratory)
Kostas Karazis (Framatome)
Christopher Boyd (US Nuclear Regulatory Commission)

ALCC allocation:

Site: Argonne Leadership Computing Facility (ALCF)

Allocation: 340,000 node-hours on Theta

Research Summary:

Nuclear energy (NE) power promises to become a reliable, carbon-free resource capable of meeting the energy needs of our nation and the world. Numerical simulation has been an intrinsic part of nuclear engineering in research, design, and licensing of existing and proposed conventional nuclear power plants. Nuclear modeling and simulation tools available today, however, are mostly low dimensional, and empirically based, valid for conditions close to the original experiments, and in many cases represent only incremental improvements on decades-old legacy codes. The development, deployment, verification, and validation of higher-fidelity computational capabilities for analyzing, modeling, simulating, and predicting complex thermo-fluid phenomena will advance nuclear power capabilities by resolving technical, cost, safety, and licensing issues. In particular, Argonne National Laboratory's close collaboration with Framatome and the US Nuclear Regulatory Commission (NRC) in the area of high performance computing (HPC) and computational fluid dynamics (CFD) provides an opportunity for further advancement of modeling and simulations for a scalable carbon-free energy quest.

In close collaboration with Framatome and the NRC, this project aims to validate open-source Nek5000 code for two sets of experiments: the ALAIN loop flow test and the PANDA facility Geometries. Framatome designs, manufactures, and installs components, fuel, and instrumentation and control systems for nuclear plants. The second partner, the US NRC, has a final say on licensing of new nuclear power plants and extending the operations of old ones in the US. Recent collaboration with the NRC Office of Nuclear Regulatory Research on Nek5000 CFD application to a range of problems shows great potential for support of some of the NRC's confirmatory calculations, exploratory analysis of generic safety issues, and the extension of experimental data for developing nuclear system analysis code flow correlations, especially beneficial in specific regimes and geometries where the experimental data are hard to come by or are of insufficient resolution and accuracy. In particular, further Nek5000 validation work for PANDA facility geometries including the recent modifications of HYMERES-2 tests can be used to inform lower-fidelity fast-turn-around models from Reynolds-averaged Navier-Stokes (RANS) to reduced-order modeling approaches.

Title: Predictive Simulations of Phase Transitions in Dynamically Compressed Carbon Materials

Principal Investigator: Ivan Oleynik, University of South Florida

Co-investigators: Aidan Thompson (Sandia National Laboratories)
Mitchell Wood (Sandia National Laboratories)
Anatoly Belonoshko (Royal Institute of Technology)

ALCC allocation:

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 551,000 node-hours on Summit

Research Summary:

This project aims to achieve a transformative impact in the understanding of phase transitions in carbon materials by integrating predictive atomic-scale simulations and experiment in one synergistic effort. Fundamental mechanisms and kinetics of phase transitions in carbon materials under dynamic compression will be studied through combination of first principles density functional theory and quantum accurate SNAP interatomic potentials molecular dynamics simulations using a DOE leadership HPC resource. The goals are:

- (1) To perform very accurate first-principles molecular dynamics (FP-MD) simulations of free energies, and phase diagrams within a wide range of pressures and temperatures;
- (2) To construct an extensive first-principles database of materials properties at high pressures and temperatures for training SNAP interatomic potentials and devise quantum accurate SNAP interatomic potentials for large-scale MD simulations of carbon materials at extreme temperatures and pressures;
- (3) To conduct MD simulations of phase transitions hand-in-hand with an ultrafast compression driver, in which the length and time scales of experiment and MD are completely matched.

The ultimate goal is to uncover fundamental atomic-scale mechanisms of carbon solid-liquid and solid-solid phase transitions from MD simulations.

Title: Hadron Structure from Lattice QCD

Principal Investigator: Konstantinos Orginos, William & Mary

Co-investigators: David Richards (Jefferson Laboratory)
Frank Winter (Jefferson Laboratory)
Balint Joo (Jefferson Laboratory)
Raza Sufian (Jefferson Laboratory)
Savvas Zafeiropoulos (Heidelberg)
Jianwei Qiu (Jefferson Laboratory)
Anatoly Radyushkin (Jefferson Laboratory/Old Dominion University)

ALCC allocation:

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 238,000 node-hours on Summit

Research Summary:

In the last few years, a major achievement in hadron structure has been the development of new lattice QCD methods that allow for direct computations. This is a groundbreaking development as it allows us for the first time to determine the full longitudinal momentum fraction dependence of the parton distribution functions (PDFs), and thus opens up a new window for the theoretical study of the structure of fundamental building blocks of matter such as the pion and the nucleon. Experimentally hadron structure studies are a central part of DOE's nuclear physics programs both with current experimental facilities such as the 12 GeV upgrade of Jefferson Lab, as well as at the future electron-ion collider (EIC). Furthermore, parton distribution functions of the nucleon are essential input in the LHC era of the search for new physics and for discovering new heavy particles.

This project will perform detailed numerical studies of hadronic structure to provide essential theoretical support to the current (JLab 12 GeV, BNL) and future (Electron Ion Collider) experimental programs in hadronic physics, goals set forward in the recent Nuclear Physics Long Range plan. Using these new methods to calculate the pion and the nucleon PDFs, as well as the pion distribution amplitude (DA), will reveal vital information about the quark and gluon degrees of freedom in these hadronic states. This investigation has the same physics goals as the approved 12 GeV experimental program at Jefferson Lab. Furthermore, the approved experiments E-03-012, E-00-002, at Jefferson Lab (12 GeV) are targeting the determination of the nucleon structure functions at large-x, precisely the region where our computational methods are more effective.

Petascale computers can make a transformational impact on hadronic physics. Without them, projects such those described here would not be possible. The USQCD collaboration has achieved important results with unprecedented precision in both high energy and nuclear physics by performing challenging numerical calculations that were previously not possible. Members of USQCD also continue to play leading roles in the development of algorithms and community codes. This project aims to build on these achievements to enable further major advances in hadronic physics.

Title: Low energy neutrino-nucleus interactions

Principal Investigator: Saori Pastore, Washington University in St. Louis

Co-investigators: Maria Piarulli (Washington University in St. Louis)
Robert Wiringa (Argonne National Laboratory)
Lorenzo Andreoli (University of Trento, Italy)
Noemi Rocco (Argonne National Laboratory)
Ewing Lusk (Argonne National Laboratory)
Alessandro Lovato (Argonne National Laboratory)

ALCC allocation:

Site: Argonne Leadership Computing Facility (ALCF)

Allocation: 390,000 node-hours on Theta

Research Summary:

Current and planned US experimental programs will benefit from better understanding of nuclear structure, reactions and properties emerging from nucleonic dynamics. This project aims to investigate, from the microscopic perspective, a kinematical region for which microscopic calculations are still lacking. The objective is to study electroweak reactions in light nuclei at low values of energy and momentum transfer. Using Quantum Monte Carlo (QMC) computational methods to solve the nuclear many-body problem, and phenomenological and chiral effective field theory (chiral EFT) to describe nuclear many-body interactions and currents, we will perform calculations of beta decay and neutrinoless-double-beta decay matrix elements, and low energy/momentum response functions induced by electrons and neutrinos, and muon-capture in light nuclei.

We anticipate delivering studies of these decay matrix elements in systems less than mass number $A \sim 13$, and electroweak responses and muon-capture for nuclei between $A = 4$ and 12. Calculations are based on the most recent implementations of many-body chiral EFT potentials and associated electroweak currents. These many-body operators have been derived from a chiral EFT that uses pions, nucleons and deltas as fundamental degrees of freedom. The derivation of interactions and currents within the same EFT formulation ensures a consistent description of short- and intermediate-range dynamics in both many-body interaction and current operators.

One of the objectives is to validate the axial currents against available experimental data. To this end, we propose to perform a detailed analysis of beta-decay matrix elements in light nuclei, transitions for which data are abundant and known, in most cases, with great accuracy. Having constrained fully the many-body axial current operators by validating the model against available data for beta-decay, we will proceed to study neutrinoless-double-beta decay matrix elements in light nuclei.

Understanding the relevant dynamics and the role of short-range physics contributing to neutrinoless-double-beta decay processes requires study of decay mediated by the exchange of a neutrino. This decay can involve values of momentum transfer of the order of a few hundred MeV. To test the validity of the axial currents in this kinematical region we propose to study low-momentum responses and muon capture in light nuclei for which data are available.

Title: Advances in Artificial Intelligence to Improve Sensory Perception

Principal Investigator: Robert Patton, Oak Ridge National Laboratory

Co-investigators: Thomas Potok (Oak Ridge National Laboratory)
Steven Young (Oak Ridge National Laboratory)
Catherine Schuman (Oak Ridge National Laboratory)
Zachariah Tyree (General Motors)
Ajay Deshpande (General Motors)

ALCC allocation:

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 200,000 node-hours on Summit

Research Summary:

Sensor data have several characteristics that differ from traditional machine learning data sets. Consequently, existing machine learning methods have not been optimized to deal with them. Sensor data sets contain both spatial and temporal features whose resolution, location, and significance change over time or need to be tracked over time. Currently, these data sets are typically transformed and reduced in a way as to make them tractable for machine learning algorithms. Much of the detail in the data is lost in this transformation, although there are likely important insights to be gained from the original data. In order to take advantage of the sensor data and enhance the perception capability from these sensors, goals include:

1. Develop new computational approaches primarily based on convolutional and spiking neural network techniques to transform the original, non-reduced data sets in a more meaningful way and to understand deeper, spatiotemporal relationships in sensor data.
2. Demonstrate the viability of new computational approaches for automatically learning a neural network for the task of object detection from sensor data. This will be performed using publicly available sensor datasets.

In order to accomplish these tasks, significant exploration of convolutional and spiking neural network architectures will be needed, requiring the use of high-performance computers to allow for large-scale architecture exploration. The key capability of this work is twofold: 1) leverage more of the available sensor data in the analysis, and 2) push advanced perception techniques closer to the sensor. We will use evolutionary optimization methods to design optimal architectures for convolutional neural networks (CNN) and spiking neural networks (SNN) to create a high-performance perception system.

Title: Integrating HPC molecular simulation with neutron scattering to study complex biological systems

Principal Investigator: Loukas Petridis, Oak Ridge National Laboratory

Co-investigators: Julie Mitchell (Oak Ridge National Laboratory)
Jeremy Smith (Oak Ridge National Laboratory)
Omar Demerdash (Oak Ridge National Laboratory)
Utsab Shrestha (Oak Ridge National Laboratory)

ALCC allocation:

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 380,000 node-hours on Summit

Research Summary:

Cells use dynamic, flexible structures to accomplish many tasks. Understanding cell function requires probing the shape and conformation of biological complexes and assemblies whose structure is neither rigid nor static. Many of the cellular molecular machines are multicomponent systems that transiently change state. Others are intrinsically disordered proteins (IDPs) a class of proteins that plays an important role in signaling pathways. A major challenge is characterizing the structural flexibility of these proteins to correlate with their function, and dysfunction. HPC simulation methods offer a predictive understanding to address this challenge.

Adequately sampling the conformational dynamics of flexible bio-systems cannot presently be realized with standard all-atom simulations. Therefore, to accelerate conformational sampling, atomistic simulations will be augmented by AdaptiveMD-enhanced sampling methods that prevent the system from being trapped in local free energy minima.

The goal is to develop computational pipelines for neutron science in advance of the Spallation Neutron Source Second Target Station (STS). Techniques that enable integration of neutron structural and dynamical information across multiple resolutions (and from other experimental techniques) will be key to providing on-line visualization and analysis of the structure and dynamics of complex biological systems at STS. Advanced tools that integrate theory, molecular dynamics simulations and scattering experiments will enable kinetic analysis of conformational transitions in complexes, of transient macromolecular assemblies and structures, and of the mechanisms of molecular recognition both *in vitro* and *in vivo*, providing new ways to understand how cellular processes are coordinated in space and time.

To test this combined neutron scattering/HPC simulation approach, we will investigate PeSCL7, a transcription factor that regulates abiotic stress response in the *Populus* bioenergy feedstock, and the flexible N-terminal of cellulose synthase, a critical part of the plant cellulose synthesis machinery. One goal is to reproduce, without any fitting, the small angle neutron scattering (SANS) intensity of the two IDPs. Once the simulations have been validated by comparison to SANS experiments, they will provide an atomic-detail description of the heterogeneous conformational ensemble of the IDPs. The thermodynamics and kinetics of the IDP processes will be calculated.

The computational models will be employed to develop a conceptual design for a small angle neutron scattering instrument at the Second Target Station. The present work will demonstrate how integration of SANS and MD-based modeling can transform research of flexible bio-systems and be applicable to many other challenging systems.

Title: Neutrino flux, energy deposition and radiological studies for the DUNE-LBNF beamline

Principal Investigator: Igor Rakhno, Fermi National Accelerator Laboratory

Co-investigators: Nikolai Mokhov (Fermi National Accelerator Laboratory)
Sergei Striganov (Fermi National Accelerator Laboratory)
Igor Tropin (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 (called DUNE-LBNF project) are under development at Fermilab since early 2010s. This activity represents a convergence of a substantial fraction of the worldwide neutrino physics community.

The primary science objectives of DUNE are to carry out a comprehensive investigation of neutrino oscillations.

The LBNF will provide a powerful 1.2-MW proton beam utilizing an upgrade of the Fermilab accelerator complex which is expected to be completed by 2025 at the latest. Later, it is planned to be upgraded to 2.4 MW. The LBNF will provide neutrino fluxes and detector infrastructure at the near site (Fermilab) and far site (Sanford Underground Research Facility).

The neutrino beamline, which utilizes a target and horn systems, decay pipe, hadron absorber and other systems, is a core component of the LBNF. At present, the beamline is under development at Fermilab. Its design considers experience from previous neutrino projects, but various optimization studies must be performed to account for both unique scientific and budget constraints. The most realistic and accurate predictions of the energy deposition in the beamline components, radiation environment and neutrino beam parameters at the far site detector can only be done by means of Monte Carlo modeling. The precise modeling of high energy particle beam transport and interactions with matter takes substantial computer time. The latter is especially pronounced when dealing with spatially extended systems and quantities in a relatively small phase-space like the hadron absorber core or a far site detector. The planned optimization calculations will include the material, geometry and magnetic fields in the target assembly and horns because these all have a significant impact on the pion/kaon production and affect the neutrino spectra. Component lifetime goals (e.g., target, horns and beam windows can and should be replaced during the facility lifetime) affect the design and require the optimization calculations as well. Availability of substantial supercomputer resources is an essential prerequisite for the successful neutrino beamline design and optimization.

Title: High-fidelity Simulations of Rotating Detonation Engines

Principal Investigator: Venkat Raman, University of Michigan

Co-investigators: Venkat Tangirala (General Electric Global Research Center)
Sarah Monahan (General Electric Global Research Center)
Peter Strakey (DoE National Energy Technology Laboratory)
Donald Ferguson (DoE National Energy Technology Laboratory)

ALCC allocation:

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 130,000 node-hours on Summit

Research Summary:

Gas turbines form the core of energy conversion devices used for aviation propulsion and stationary power generation. Traditional gas turbines, operating under a Brayton cycle, are approaching the limit of opportunities to improve thermal efficiency. Applying continuous detonations as the mode of combustion is a disruptive technological approach that can overcome this challenge, yielding decreased entropy generation and pressure rise. Thermal efficiency of such devices can reach the 67% target that DOE is seeking, while reducing fuel consumption by 10-15%. One implementation of such detonation combustion is the rotating detonation engine (RDE), which has the potential to revolutionize gas turbine combustion technology by enabling high-efficiency power generation engines. The focus of this study is to develop highly detailed simulations of practical RDE configurations to identify spurious sources of efficiency losses.

Emerging RDE designs overwhelmingly inject separate fuel and air streams, and use turbulence to mix these reactant gases molecularly. The goal is to create a sufficiently homogeneous mixture to stably maintain the detonation process. However, spatial and temporal variations in the flow field lead to incomplete mixing that can adversely affect the detonation process, and in turn the efficiency of the RDE system. Further, the impact of the mixing process depends on the reactivity of the fuel. Here, a multi-pronged approach is followed to understand the interaction of turbulence, mixing and shock-driven chemical reactions. Highly detailed full-scale simulations of two different RDE configurations are used to study the relative role of turbulent mixing and detonation process on the stability of the RDE system. Simultaneously, full-scale RDE simulations of complex geometries are conducted to understand the impact of injector design and geometry on mixing. In particular, performance measured in terms of RDE size is obtained.

A third challenge concerns the long-time behavior of RDEs. Full-system simulations often can track only a few cycles of the detonation wave. To accelerate these computations, a novel machine-learning approach is used. The canonical flow simulations, which resolve all the details of the turbulence-chemistry interaction, are used to train a reduced-order model (ROM) that is computationally inexpensive and costs roughly 10% of the full chemistry model. Using this ROM, very long-time simulations that cover more than 50 detonation cycles are attempted for the first time.

In order to ensure direct relevance to the DOE mission, this project is a collaboration between academia (UM), industry (GE Global Research), and national lab (DOE NETL). If successful, these studies can revolutionize gas turbine research and development.

Title: Machine Learning Directed Adaptive Multiscale Simulations to Model RAS-RAF Cancer Initiation Pathway

Principal Investigator: Arvind Ramanathan, Argonne National Laboratory

Co-investigators: Jim Glosli, Helgi Ingolfsson (Lawrence Livermore National Laboratory)
 Gnana Gnanakaran (Los Alamos National Laboratory)
 Tim Carpenter, Harsh Bhatia (Lawrence Livermore National Laboratory)
 Andrew Stephen (Frederick National Laboratory for Cancer Research)
 Christopher Neale, Cesar Lopez (Los Alamos National Laboratory)
 Tomas Opplestrup, Francesco Di Natale (Lawrence Livermore National Laboratory)
 Gautham Dharuman, Timo Bremer (Lawrence Livermore National Laboratory)
 Angel Garcia (Los Alamos National Laboratory)
 Thomas Turbyville (Frederick National Laboratory for Cancer Research)
 Xiaohua Zhang (Lawrence Livermore National Laboratory)
 Dwight Nissley (Frederick National Laboratory for Cancer Research)
 Felice Lightstone, Fred Streitz (Lawrence Livermore National Laboratory)
 Nick Hengartner (Los Alamos National Laboratory)
 Liam Stanton (San Jose State University/ Lawrence Livermore National Laboratory)
 Mike Surh (Lawrence Livermore National Laboratory)

ALCC allocation:

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 400,000 node-hours on Summit

Research Summary:

At the heart of quantifying and predicting multiscale phenomena is the use of high-fidelity simulations to accurately and reliably represent the behavior of a given system. Such simulations take up 45-60% of time on existing supercomputing resources annually. Scaling of the software is currently achieved through hardware optimization and/or algorithmic innovations. Emerging Exascale architectures pose additional challenges, requiring new approaches to optimize their overall throughput. One target is refining search-and-sample techniques. This work is part of the Joint Design of Advanced Computing Solutions for Cancer (JDACS4C) program established by the U.S. Department of Energy (DOE) and the National Cancer Institute (NCI) of the National Institutes of Health.

The complexity of these techniques can be greatly reduced by (i) exploiting the intrinsic low dimensionality of the simulations to reduce the number of parameters tracked, (ii) while simultaneously using machine learning techniques to steer the search. The overall goal is to develop a massively parallel, machine learning driven, adaptive multiscale simulation (MLD-AMS) framework that enables bridging time- and length-scales using macro-level models, while simultaneously maintaining molecular-scale detail.

The cancer initiation pathway mediated by the RAS and RAF oncogenes is a complex biological phenomenon. Nearly a third of all cancers diagnosed in the US are driven by mutations in RAS genes and their protein products, whose aberrant behavior accounts for a particularly high percentage of pancreatic (~95%), colorectal (~45%), and lung (~35%) cancers. We aim to extend the framework to decipher complex spatiotemporal aspects of initial triggering of RAS-RAF signaling near the membrane. Comprehensively mapping the conformational states of RAS and RAF and their dependence on local membrane composition, will provide quantitative insights into the RAS-RAF complex, which can provide rationale for new drug design strategies. A data-driven approach interleaving multi-scale simulations will have transformative impacts in many areas (biology, materials science, climate sciences, fluid dynamics, nuclear fusion, etc.) on the high-fidelity simulation capabilities that drive fundamental scientific discoveries.

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

Principal Investigator: Emily Shemon, Argonne National Laboratory

Co-investigators: Yiqi Yu (Argonne National Laboratory)
Dillon Shaver (Argonne National Laboratory)
Taek K. Kim (Argonne National Laboratory)
Florent Heidet (Argonne National Laboratory)
Alex Levinsky (Westinghouse Electric Company)

ALCC allocation:

Site: Argonne Leadership Computing Facility (ALCF)

Allocation: 880,000 node-hours on Theta

Research Summary:

Under a multi-laboratory effort, DOE is designing the VTR, a proposed experimental nuclear reactor based on fast spectrum sodium-cooled technology. Upon operation, the facility will be used to assess the performance of irradiated materials such as accident-tolerant fuel for inclusion in U.S. advanced nuclear reactor technology. Design for the fuel assemblies is relatively mature; however, the control assembly design has not been finalized due to competing parameters which have not yet been optimized. Multiphysics simulation will be applied to analyze peak temperatures in both the control assembly and fuel assembly for a specific control assembly design. Designs for core components having a significant impact on the flow through the VTR assemblies have also not been finalized. The pressure drop induced by the required core components is a key design parameter for which uncertainty must be minimized. Accurately predicting pressure drop across complex component designs requires detailed simulations to properly resolve important flow features. High fidelity simulation will be applied to yield pressure drop data which will directly impact design decisions for these components and the overall plant design.

High fidelity multiphysics simulations will also be performed for the lead-cooled fast reactor (LFR) being developed by Westinghouse Electric Company, a world leader in the development and commercialization of nuclear power plants. Known as a leader in thermal, light water reactor technology, the company is expanding its portfolio with advanced designs including non-water coolant and fast spectrum systems due to the potential efficiency and safety of such systems.

Among other advantages, the Westinghouse LFR can be built at reduced cost, more efficiently utilize natural resources, and operate at higher thermal efficiencies than existing commercial light water reactors. Aligning national laboratory R&D with industry needs is a key priority, as indicated by a variety of funding opportunity announcements (DOE-NE's GAIN Gateway for Accelerated Innovation in Nuclear program, ARPA-E's MEITNER Modeling-Enhanced Innovations Trailblazing Nuclear Energy Reinvigoration program) and changes in DOE-NE programmatic direction. Therefore, a series of high fidelity multiphysics calculations will be performed to compute a set of hot channel factors (HCF) for this industry LFR design using DOE-NEAMS codes. HCFs are critical safety factors that account for how uncertainties in physical parameters, manufacturing tolerances, modeling approximations, or experiments impact the worst case scenario reactor temperatures. There are no historical HCF datasets available for LFRs, so the results of the proposed calculations will be very informative for the reactor technology industry; including Westinghouse, which has already identified priority hot channel factors to be assessed here.

Title: Supercomputing for automotive high-temperature alloy design Consortium/End-Station

Principal Investigator: Dongwon Shin, Oak Ridge National Laboratory

Co-investigators: Bala Radhakrishnan (Oak Ridge National Laboratory)
Xiaohua Hu (Oak Ridge National Laboratory)
James Morris (Oak Ridge National Laboratory)
Rishi Pillai (Oak Ridge National Laboratory)
Adri van Duin (The Pennsylvania State University)
Sangkeun Lee (Oak Ridge National Laboratory)
James Haynes (Oak Ridge National Laboratory)

ALCC allocation:

Site: Argonne Leadership Computing Facility (ALCF)

Allocation: 350,000 node-hours on Theta

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 391,000 node-hours on Summit

Research Summary:

ORNL recently designed a new family of advanced higher-temperature Al-Cu alloys for automotive cylinder heads at an unprecedented pace (4 years from alloy concept to full engine testing). The accelerated development strategy significantly benefited from massively parallel atomistic calculations and microstructure simulations performed on at Oak Ridge Leadership Computing Facilities. DOE/EERE-VTO intends to extend the recent success by exploiting a suite of state-of-the-art ICME software packages (i.e., density functional theory, CALPHAD, phase-field, crystal plasticity, reactive force field, and machine learning) to accelerate the design of various classes of high-temperature alloys (i.e., Al-, Fe-, Ni- and Ti-alloys). This project will heavily exploit the GPU-capability of OLCF's Summit for atomistic simulations (VASP and LAMMPS) and microstructure evolution software (ORNL's in-house code), which have proven to be highly efficient. ORNL's world-class alloy design expertise and advanced characterization across multiple DOE national laboratories will actively support the proposed HPC ICME tasks. The outcome of this project is anticipated to fill critical knowledge gaps and shorten the timeframe from prototype automotive high-temperature alloys to their real-world deployment.

Title: Predictive Modeling and Machine Learning for Functional Nanoporous Materials Consortium/End-Station

Principal Investigator: J. Ilja Siepmann, University of Minnesota

Co-investigators: Evgenii Fetisov (Pacific Northwest National Laboratory)
Jason Goodpaster (University of Minnesota)
Chris Knight (Argonne National Laboratory)
Christopher Mundy (Pacific Northwest National Laboratory)
Yongchul Chung (Pusan National University)

ALCC allocation:

Site: Argonne Leadership Computing Facility (ALCF)
Allocation: 620,000 node-hours on Theta
Site: National Energy Research Scientific Computing Center (NERSC)
Allocation: 200,000 node-hours on Cori

Research Summary:

An interdisciplinary, collaborative team will use predictive hierarchical modeling and machine learning to accelerate the discovery and design of materials for a variety of energy-related applications and to advance scientific and technological capabilities with innovative discoveries. The research objectives of this ALCC proposal are aligned with the goals of the DOE-funded Nanoporous Materials Genome Center (NMGC). Research focus is on four topics:

- ***Hierarchical screening and machine learning for adsorption in nanoporous materials*** aimed at the discovery of materials with superior performance, including the development of transferable force fields for aluminosilicates and aluminophosphate zeotypes enabling future hierarchical screening studies, and the development of machine-learned models for the prediction of mixture isotherms from single-component isotherms.
- ***First principles Monte Carlo calculations for adsorption systems*** aimed at improving characterization of catalytic acid sites through modeling of reactive adsorption of propene and amines in acidic zeolites.
- ***First principles simulations for nanoparticle assembly*** aimed at understanding hydration forces and providing a molecular-scale view of the double layer that govern nanoparticle assembly and synthesis.
- ***Development of periodic wave function in density functional theory embedding methods*** aimed at improving accuracy and efficiency of predictive modeling for the selective oxidation of alkanes to alcohols in metal-organic frameworks.

Work enabled by prior ALCC allocations has led to the award of one US patent for the ethanol purification.

Title: AI-Enabled Computational Cancer Phenotyping for Precision Oncology

Principal Investigators: Georgia Tourassi, John Gounley, and Hong-Jun Yoon,
Oak Ridge National Laboratory

Co-investigators: Jacob Hinkle (Oak Ridge National Laboratory)
Mohammed Alawad (Oak Ridge National Laboratory)
Folami Alamudun (Oak Ridge National Laboratory)

ALCC allocation:

Site: Oak Ridge Leadership Computing Facility (OLCF)
Allocation: 200,000 node-hours on Summit

Research Summary:

Cancer is a global public health concern, with 1,735,350 new cases and more than 609,640 cancer deaths estimated in the US annually. With advancing age as the most important risk factor, cancer incidence continues to increase amidst a growing older population. The Precision Medicine Initiative in Oncology was established in 2015 to accelerate cancer research advances. Developing cancer phenotypic algorithms to provide measurable clinical and biological markers of the disease is a critical step for effective patient stratification to support the vision of Precision Medicine. Computational phenotyping is the process of deriving patient phenotypes from clinical data using computer-executable algorithms. The overarching goal of this project is to develop scalable computational solutions for large-scale accurate cancer phenotyping from unstructured clinical text data to accelerate precision medicine cancer advances. The key challenges addressed here are the development, scaling, and deployment of deep learning algorithms to enable and support reliable and comprehensive cancer phenotyping from population-level national cancer registry data. Our project will pursue promising but computationally expensive artificial intelligence (AI) approaches for accurate cancer phenotypic characterization at a population scale. NCI has partnered with DOE to unleash the power of supercomputers and artificial intelligence for modernizing the national cancer surveillance program. The collaboration is leveraging the capabilities of high-performance computing to support an advanced population-level cancer surveillance program and develop an integrated data-driven modeling framework. This work offers transformative implications for computational phenotyping, where scalable AI methods like deep learning hold substantial promise and provide novel approaches to key constraints in current implementations. By scaling these approaches to national-level volumes of clinical text reports, this work will enable the national cancer surveillance program achieve near-real-time cancer incidence reporting. Moreover, since clinical phenotypes are manifestations of the genome-exposome interaction, this work will help provide deep insights in the genome-phenome association, and ultimately lead to more effective cancer prevention, diagnosis, and treatment.

Title: Multiphase Flow in Shale

Principal Investigator: David Trebotich, Lawrence Berkeley National Laboratory

Co-investigators:

ALCC allocation:

Site: National Energy Research Scientific Computing Center (NERSC)

Allocation: 860,000 node-hours on Cori

Research Summary:

Capture and geological storage of CO₂ is among the most critical energy technologies of the next century. A major theme regarding the security of subsurface storage is the integrity of nanoporous (very low permeability) rock layers like shale that are critical for blocking upward flow of buoyant CO₂ during early phases of carbon storage. A key question is whether the pressure needed to inject commercial volumes of CO₂ into permeable sandstone and carbonate formations will cause mechanical failure that could provide leakage pathways that compromise the storage effort. Shale varies in mineralogical composition and pore structure and its mechanical response to pressurization varies accordingly. The role of geochemical processes in widening or healing fractures varies with mineralogy, pore structure, and invading fluid composition. The response of shales to perturbations from concurrent mechanical and chemical stressors is a coupled mesoscale behavior that has not been well understood or predictable. In this context, a new understanding of this coupled mesoscale behavior will also help to unravel phenomena associated with fracture-induced oil and gas extraction from shale, a problem which has significant impacts on our national and global economies, as well as the quest for energy independence in general.

Multiphase, reacting flow and transport modeling based on direct numerical simulation in realistic pore space is used to not only inform experiments *a priori*, but to also help interpret the experimental results, to interrogate properties of materials in the absence of experiments, and to generalize the results to the larger (porous-continuum) scales. The high performance simulation capability, Chombo-Crunch, was developed as part of the research of an Energy Frontier Research Center to address this problem. Chombo-Crunch has been used to validate a number of reactive transport experiments associated with carbon sequestration as well to interrogate nanoporous materials in the absence of experiments.

Leadership-class computational resources are required to simulate resolved multiphase flow in nanoporous Marcellus shale using Chombo-Crunch for the purpose of achieving multiphase breakthrough and estimating the effect of multiple fluid phases on bulk flow properties.

Title: Modeling the Response of Fusion Plasma Facing Components

Principal Investigator: Brian Wirth, Oak Ridge National Laboratory/University of Tennessee

Co-investigators: David Bernholdt (Oak Ridge National Laboratory)
Aiden Thompson (Sandia National Laboratories)
Wahyu Setyawan (Pacific Northwest National Laboratory)

ALCC allocation:

Site: Argonne Leadership Computing Facility (ALCF)

Allocation: 200,000 node-hours on Theta

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 250,000 node-hours on Summit

Research Summary:

The performance demands on plasma-facing components (PFCs), first wall and blanket systems of future fusion power plants are beyond the capability of current materials, which is one of the reasons that the United States National Academy of Engineering has ranked the quest for fusion as one of the top grand challenges for engineering in the 21st Century. Furthermore, it is clear that the plasma surface interactions (PSIs) occurring in the divertor and PFCs pose a critical scientific challenge that limits our ability to achieve electricity production from fusion. The realization of fusion as a practical, 21st Century energy source requires improved knowledge of plasma surface interactions and the materials engineering design of component systems to survive the incredibly extreme heat and particle flux exposure conditions of a fusion power plant. This project aims 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 including beryllium and neon. In particular, two tasks are envisioned to investigate the surface response following implantation of beryllium, helium and hydrogen into the tungsten divertor, and the underlying mechanisms controlling the reduced tritium permeation and retention observed in helium exposed tungsten.

