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

Principal Investigator: Harsh Bhatia (Lawrence Livermore National Laboratory)

Co-investigators: Helgi Ingolfsson (Lawrence Livermore National Laboratory), Fikret Aydin (Lawrence Livermore National Laboratory)

ALCC Allocation:

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

Allocation(s): 350,000 node-hours on Summit, 70,000 node-hours on Frontier

Research Summary:

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

As part of an ongoing NCI/DOE collaboration, Project ADMIRRAL (AI-Driven Multiscale Investigation of RAS-RAF Activation Lifecycle), we are developing a predictive multiscale model of RAS-RAF driven cancer initiation and growth that allows coupling experiments with multiresolution simulations to build a dynamic model of RAS-RAF biology in varying cellular membrane compositions. The combination of simulation with experiment will generate a vast amount of data that, with the aid of high-performance analytics, will establish a detailed, predictive understanding of RAS-RAF activation — structural and dynamic insights that will open new possibilities for therapeutic intervention.

Computationally, this project will expand up the award-winning framework MuMMI (Multiscale Machine-learned Modeling Infrastructure), developed recently through previous NCI/DOE funding and ALCC allocations. MuMMI represents a new paradigm of multiscale simulations and enables exploring large length- and time-scales using macro-level models, while simultaneously maintaining molecular-scale detail using a novel machine learning (ML)-based sampling framework. This project will expand MuMMI capabilities to generate simulation ensembles that substantially expand the scope of multiscale simulations to accelerate different aspects of protein interactions along RAS-RAF activation pathway. Our data-driven approach that interleaves multiscale simulations using ML will have transformative impacts on high-fidelity simulation capabilities that drive fundamental scientific discoveries in the areas of biology, materials science, climate sciences, fluid dynamics, nuclear fusion, etc. This work leverages the DOE leadership to accelerate multiscale simulations using scalable and novel ML strategies. With the emerging hardware and software stacks, there is an opportunity to transform simulations to enable better utilization of extreme-scale parallelism while significantly reducing energy footprint and time-to-solution.

Title: Exploring pedestal structure via the electromagnetic gyrokinetic framework

Principal Investigator: Jeff Candy (General Atomics)

Co-investigators: Emily Belli (General Atomics)

ALCC Allocation:

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

Allocation(s): 70,000 node-hours on Summit,
60,000 node-hours on Perlmutter-GPU

Research Summary:

A one-year focused project is planned to exploit the leadership-scale edge-plasma turbulence simulation capabilities of CGYRO to explore tokamak pedestal structure, stability and transport. In the edge region of a fusion (tokamak) plasma, direct numerical simulation of turbulence is challenging due to strong density and temperature gradients and plasma cross-sectional shaping, as well as the proximity to electromagnetic stability thresholds. One particular edge configuration, the so-called H-mode, is the most promising scenario envisioned to achieve controlled nuclear fusion in tokamaks. In the transition from low to high confinement mode (L-H transition) a narrow region forms just inside the magnetic boundary where strong plasma flows can reduce the turbulence intensity. The reduced edge turbulence causes the development of steep gradients in temperature and density profiles (i.e., pedestals) which substantially increase the core density and temperature. The high density and temperature maximize the core fusion reaction rate and are thus beneficial for reactors. Capturing all the relevant space and time scales requires extremely high numerical resolution. CGYRO was designed to target these extreme resolutions. CGYRO optimizations include the use of GPUDirect MPI and have been further optimized during prior allocations on OLCF Summit. Understanding the roles of electromagnetic modes (KBM, MTM) in determining the pedestal structure is an important step to develop an improved reduced model for the pedestal. The planned work will have a very specific primary goal in mind: to reproduce some aspects of the pedestal structure observed in recent experiments. Secondary goals are to refine theoretical understanding of the pedestal, and to add to a formal and growing database of CGYRO turbulence results in order to calibrate TGLF (the GA turbulent transport model).

So that the results are available to the community, we will continue to make the simulation data available as part of a gyrokinetic database hosted by the Advanced Tokamak Modeling (AToM) SCiDAC-4 project (J. Candy, PI). CGYRO optimization and application is directly funded by AToM and thus time granted by ALCC is critical for this research. The work has a direct connection to U.S. DOE Office of Fusion Energy ITER and DOE Projects currently underway. In particular, this work will also directly support the U.S. DOE Edge Simulation Laboratory, a multi-institutional project jointly funded by both the DOE Office of Fusion Energy Sciences and the Office of Advanced Scientific Computing Research (ASCR).

Title: Deep learning-enabled ab initio simulation of heterogeneous aqueous systems

Principal Investigator: Roberto Car (Princeton University)

Co-investigators: Athanassios Z. Panagiotopoulos (Princeton University),
Annabella Selloni (Princeton University),
Xifan Wu (Temple University)

ALCC Allocation:

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

Allocation(s): 860,000 node-hours on Summit

Research Summary:

Chemistry in Solution and at Interfaces (CSI) is a Computational Chemical Sciences Center funded by the US Department of Energy (DOE) to study chemical phenomena in fluid and interfacial environments. These phenomena span a vast range of size and time scales, going from the local dynamics of the chemical bonds to the collective motions of the large assemblies of atoms behind macroscopic materials changes. These processes play an essential role in life and energy applications. To model them, CSI developed a suite of software tools based on artificial intelligence methods, such as machine learning, to learn the chemical interactions from costly quantum mechanical calculations on relatively small model systems, and deep neural networks to describe these interactions in large molecular environments. The codes are designed to run efficiently on graphic processing units (GPUs) and can harness the computational power of the huge arrays of GPUs available in the most advanced supercomputers of the national laboratories of the DOE, such as Summit at Oak Ridge National Laboratory.

The formidable computational power of Summit will be used in two flagship molecular dynamics simulation studies conducted at CSI. In these studies atomic trajectories in complex environments will be generated from interatomic potentials learned from quantum mechanics. In one project, the center will model ice nucleation processes on the surface of solid particles present in the atmosphere, a study that will provide microscopic information relevant for climate modelling. In the other project, the center will perform atomistic simulations of electrolyte solutions of varying ionic strength in contact with an electrode, a study relevant for better understanding chemical dynamics in batteries.

Title: Mitigating Climate Change Through Zero Carbon Fuels

Principal Investigator: Jacqueline Chen (Sandia National Laboratories)

Co-investigators: Martin Rieth (Sandia National Laboratories),
Alex Aiken (SLAC),
Seshu Yamajala (SLAC),
Elliott Slaughter (SLAC)

ALCC Allocation:

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

Allocation(s): 400,000 node-hours on Summit

Research Summary

Mitigating climate change while providing our nation's transportation and power generation needs are key to energy and environmental security. New e-fuels and hydrogen rich fuels provide a near zero-carbon alternative to fossil fuels for gas turbine engines for power generation and for compression ignition engines for marine shipping. A potential shift to hydrogen as a clean energy carrier is one of the most promising strategies to significantly reduce carbon dioxide emissions in the face of increasing energy demand. This is particularly relevant for large-scale power generation combined with pre-combustion carbon sequestration and large-scale energy-storage schemes, both relying on hydrogen obtained from fossil or renewable sources. The use of ammonia as a fuel for gas turbines is attractive due to its advantages over hydrogen from the perspective of transport and storage and has been explored in the past but was supplanted by conventional hydrocarbons due to its poor reactivity. However, recent studies with rich-lean fuel staging have shown some promising results. In principle gas turbines are fuel flexible; however, in reality, burner design becomes tuned to specific fuels, methane being the most common. Here, we consider mixtures of ammonia and hydrogen, as a carbon free fuel blend for gas turbines and marine shipping applications. Direct numerical simulation (DNS) results from our previous INCITE award in 2019 show that even in the presence of intense sheared turbulence, fast hydrogen diffusion in the ammonia blend in highly turbulent flames plays a crucial role in the enhancement of the combustion rate and in the prevention of extinction and blowoff relative to methane. More recently, DNS results from our 2020 INCITE allocation demonstrate the surprising robustness of combustion with ammonia/hydrogen blends at moderately elevated pressure (10 atm), again due to the intrinsic thermo-diffusive instability of hydrogen coupled with turbulent transport. Here, we propose DNS of turbulent premixed combustion with ammonia/hydrogen blends at 20 atm relevant to gas turbine operation to investigate the pressure scaling of ammonia/hydrogen blends and to understand the pathways for NO_x formation (NO_x, N_2O). The unique DNS data will be shared with the modeling community to develop predictive models for NO emissions and mitigation strategies at relevant conditions.

Title: Advancing Watershed System Science using ML and Process-based Simulation

Principal Investigator: Ethan Coon (Oak Ridge National Laboratory)

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

ALCC Allocation:

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

Allocation(s): 180,000 node-hours on Perlmutter-CPU

Research Summary:

Plentiful, clean water is a crucial resource required for life, yet our Nation's water resources are under threat from a changing climate. Predicting and understanding how stresses and extreme events effect water availability, flooding, and pollutants requires relevant, accurate, and timely simulations of the water cycle. To advance basic understanding and provide tools for water stakeholders, this project tackles the computationally demanding task of providing full-river basin predictions that are informed by processes at scales that resolve topography of the land surface and includes lateral flow and transport. Funded by an ongoing Office of Science, Biological and Environmental Research project, ExaSheds and enabled by these computing resources, this work will develop a Machine Learning (ML) informed simulation capability that addresses questions of how water quantity, temperature, and transport of chemical species will change under a changing climate and environment. This project will continue to develop, train, and evaluate this capability and will result in full-river basin simulations at ~100m resolution over climatological timescales – a first in the hydrologic modeling community.

This renewal proposal continues work in progress that combines ML approaches with Amanzi-ATS, a High-Performance Computing code that runs simulations of coupled surface and subsurface hydrologic flow, energy transport, and reactive transport of chemical species. The first year focused on the Upper Colorado Headwaters. In the coming year, this project will:

1. enable a few large-scale simulations of large river basins;
2. expand the use of surrogate and hybrid models and infer property-enabled surrogate models that are transferrable across basins without observations;
3. include energy transport for stream temperature prediction; and
4. begin to leverage a second version of Amanzi-ATS that is designed for exascale resources.

Each of these objectives has been under active development and prototyping on small-scale catchments that run on smaller computing resources. Using this larger allocation, the project will expand those prototypes to the full river basin; this critically integrates the impact of small-scale processes up to scales affecting water services for human use, and therefore can ultimately help inform decision-makers. In this way, the project will contribute to the Department of Energy Office of Science's mission to address environmental and energy challenges through transformative science.

2022 ASCR Leadership Computing Challenge Award

Title: Dynamics and decoherence of excited states in 2D systems for quantum technology

Principal Investigator: Mauro Del Ben (LBNL)

Co-investigators: Jack Deslippe (LBNL),
Felipe Jornada (Stanford University),
Diana Qiu (Yale University),
Steven G. Louie (UC Berkeley & LBNL),
Jeffrey Neaton (UC Berkeley & LBNL),

ALCC Allocation:

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

Allocation(s): 100,000 node-hours on Summit,
25,000 node-hours on Perlmutter-GPU

Research Summary:

Understanding and engineering defects that can be used for quantum information or to tune bulk states through quasiparticle interference phenomena is a grand challenge for materials science and engineering. A particularly interesting class of host systems for defects are low-dimensional or quasi-two-dimensional (quasi-2D) atomic layers, such as monolayer MoS_2. Localized defects in such systems may hybridize with the bulk band edges from the host quasi-2D systems, leading to potential applications in optoelectronics and quantum information science. Furthermore, new classes of defects, such as magnetic impurities, are promising candidates to realize dilute magnetic semiconductors, with potential applications in spintronics to utilize both charge and spin in transport, data storage and processing. For instance, recent theoretical works have shown that MoS_2 doped with transition metal magnetic impurities (such as Mn and V) can generate magnetic ordered states with a high Curie temperature, making them attractive candidates for spintronic devices. However, the fundamental understanding of such defects is complicated by the large structural complexity to describe them and the interplay between localized defects and delocalized bulk excitations. This project aims to study the quasiparticle, optical excitations and magnetic orderings associated with charged defects in 2D materials by studying charged chalcogen vacancies, transition metal vacancies, and anti-site defects in mono- and few-layers transition metal dichalcogenides (TMDs). These accurate calculations of defect quasiparticle wavefunctions will be used to study their evolution and decoherence employing the time-dependent GW formalism recently developed by our team, which is suitable for nonlinear and nonequilibrium dynamics. By further including electron-phonon interactions from first principles, we can study the physics of defect decoherence and motional narrowing as potential ways to engineer defect states for quantum applications. This study will be made possible at unprecedented computational scale and accuracy by this ALCC allocation and by using the massively parallel GPU-enabled BerkeleyGW software package.

Title: The spectrum and structure of hadrons

Principal Investigator: Robert Edwards (Jefferson Lab)

Co-investigators: Raul Briceno (Old Dominion University),
Chris Monahan (William and Mary),
Kostas Orginos (William and Mary)

ALCC Allocation:

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

Allocation(s): 300,000 node-hours on Polaris

Research Summary:

Using Summit at the Oak Ridge Leadership Computing Facility, we will compute the structure and spectrum of strongly-coupled hadronic states directly from the fundamental theory describing the interactions of quarks and gluons – the fundamental particles of nuclear matter. These calculations will provide essential theoretical support to the experimental program of the Thomas Jefferson National Accelerator Facility (Jefferson Lab) including the CLAS12 and GlueX experiments, and to the future Electron Ion Collider (EIC) at Brookhaven National Laboratory.

We have four goals. First, to compute the Bjorken x-dependent, isovector light-quark generalized parton distributions (GPDs) of the nucleon, in the continuum and physical quark-mass limits of lattice Quantum Chromo-Dynamics. Second, to provide a lattice determination of the flavor decomposition of the proton sea through isoscalar GPDs. Third, to calculate hadronic scattering amplitudes for scattering in the isospin 0, 1 and 2 channels and for KK scattering in the isospin 0 and 1 channels. Fourth, to obtain the first determination of radiative transitions for isovector and isoscalar mesons from lattice QCD.

Breakthrough advances in algorithms and computational implementations exploiting the capabilities of the Summit heterogeneous computing system have accelerated this project. The software has been developed largely under DOE SciDAC, the DOE Exascale Computing Project, and in collaboration with researchers at Nvidia. The project team is made up of both theoretical nuclear physicists, and computational scientists at the forefront of leadership-class computing.

Leadership class computing is critical for our goals, which will provide the ab initio answers to a question “essential for understanding the nature of visible matter” and central to the Department of Energy's experimental nuclear physics program: how do quarks and gluons form the wide range of hadronic bound states we observe in experiment?

Title: Energy partition and particle acceleration in laboratory magnetized shocks

Principal Investigator: Frederico Fiuza (SLAC National Accelerator Laboratory)

Co-investigators: Arno Vanthieghem

ALCC Allocation:

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

Allocation(s): 860,000 node-hours on Theta

Research Summary:

Astrophysical collisionless shocks are among the most powerful particle accelerators in the Universe. Generated by violent interactions of supersonic plasma flows with the ambient medium, shock waves are observed to amplify magnetic fields and to accelerate electrons and ions to highly relativistic speeds. Recent developments in laboratory high-energy-density (HED) laser-plasma experiments are now opening for the first time the opportunity to probe the microphysics and particle acceleration mechanisms of magnetized collisionless shocks in conditions relevant to high-energy astrophysical environments.

The objectives of this research are to perform large-scale first-principles simulations of magnetized collisionless shocks to address important longstanding questions associated with energy partition and particle acceleration: “What are the processes that mediate the onset of magnetic turbulence in the shock and does it decay?”, “What is the difference between electron and ion thermalization at the shock?”, “What is the injection process for particle acceleration at the shock front, and what is the efficiency?” A transformative advance in the ability to address these questions requires the combination of three-dimensional, high dynamic range, fully kinetic simulations and controlled laboratory experiments where numerical findings can be tested and used to improve theoretical models and the understanding of observations, which is at the focus of this project. The fundamental understanding of particle acceleration in plasmas to be provided by this project is central to DOE’s mission in Discovery Plasma Science. The results of this research are expected to have a significant impact on unveiling long-standing questions behind cosmic plasma accelerators, in advancing the understanding of interpenetrating magnetized HED plasmas, and in generating new ideas for efficient laboratory accelerators. Finally, the tight connection between the planned simulations and experimental programs on the National Ignition Facility will also enable the important benchmark of widely used numerical plasma models in magnetized HED conditions of relevance to DOE programs.

Title: Accelerated Discovery of Low-Cost Hydrogen-Resistant Alloys for Extreme ' Environments

Principal Investigator: Michael Gao, National Energy Technology Laboratory

ALCC Allocation:

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

Allocation(s): 250,000 node-hours on Perlmutter-CPU

Research Summary:

Hydrogen, even in a few parts per million by weight, can cause severe degradation to mechanical properties of metals and alloys, known as hydrogen embrittlement (HE). This includes loss of ductility and fracture toughness, increased fatigue crack growth rates, susceptibility to stress corrosion cracking, and hydrogen-induced cracking. This project aims to reveal the fundamental mechanisms of HE in metals and alloys, thereby allowing the rapid design of novel cost-competitive hydrogen-resistant high-performance alloys for high-pressure hydrogen environments. High-throughput multiscale computer modeling will be performed at NERSC to predict the thermodynamic, kinetic, mechanical, and environmental properties of steels and high entropy alloys in presence of hydrogen. Several state-of-the-art techniques will be engaged that bridge various length and time scales, from continuum to atomistic quantum mechanics calculations and machine learning. Most promising alloy compositions will be recommended for experimental validation that will be fed back to modeling for further refinement and improvement.

This project directly supports the mission of the U.S. Department of Energy's Office of Fossil Energy and Carbon Management (FECM) to develop technologies toward a net-zero carbon economy. Its success will bring transformational materials technology that can be safely used in a wide array of applications in presence of hydrogen. For example, hydrogen is recognized as an important fuel for gas turbine systems to reduce CO_2 emissions and increase efficiency, and hence holds the greatest potential in preventing harmful environmental emission. Hydrogen can also greatly reduce carbon footprint by replacing gasoline for automobiles as hydrogen-powered fuel cell electric vehicles emit only water and warm air. Because replacing fossil fuel with hydrogen requires that the materials for transporting, storing, or in contact with hydrogen (such as pipelines and compressors) are HE resistant, preventing HE of metals and alloys is vital to achieving a net-zero carbon economy.

Title: First-principles prediction of solute segregation at defects in Mg alloys

Principal Investigator: Vikram Gavini (University of Michigan)

Co-investigators: Sambit Das (University of Michigan),
Liang Qi (University of Michigan),

ALCC Allocation:

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

Allocation(s): 176,500 node-hours on Summit

Research Summary:

Novel light-weight structural materials play an important role in reducing energy consumption and carbon footprint in automotive and aerospace sectors. Despite its high strength-to-weight ratio, Mg has low ductility due to insufficient density of pyramidal <c+a> dislocations required for <c> axis plastic deformation, and strong basal texture of standard wrought Mg alloys. This has limited its utility as a structural metal in automotive and aerospace industries. The objectives of this work are to conduct first-principles calculations of the energetics of defects controlling the various ductility enhancing/limiting mechanisms in dilute Mg alloys. The first part of this research focuses on studying solute segregation at non-basal dislocation systems in Mg-Gd and Mg-Ca alloys. Specifically, the temperature and strain-rate effect of solute segregation will be studied on (a) activation barrier of <c+a> pyramidal II-I cross-slip mechanism, and (b) Portevin–Le Chatelier effect (serrated stress-strain curve) in non-basal dislocation systems (pyramidal and prismatic). The second part of the planned research focuses on studying solute segregation at random GBs in Mg using large-scale DFT calculations on realistic cell sizes of around 4,000 atoms. In particular, solute-segregation and co-segregation energies will be computed for Mg-X (Al, Ca, and Zn) and Mg-X-Y (Ca-Zn, and Al-Zn) systems, respectively. The first-principles inputs of dislocation/GB-solute interaction energies in the above studies will require large system sizes containing up to ~5000 atoms, which are inaccessible using state-of-the-art plane-wave DFT codes. The planned research will leverage the significant computational, algorithmic and high-performance computing advances made by the PIs recently in developing the DFT-FE code—a massively parallel large-scale real-space DFT code based on adaptive finite-element discretization for systems sizes up to 100,000 electrons.

Overall, the planned first-principles informed modelling of solute segregation at dislocations and random GBs will enable quantitative guidance on solute combinations, concentration ranges and thermomechanical processing regimes that simultaneously enhance both intrinsic ductility and weaken basal texture. This in turn has the potential for advancing light-weight multi-component Mg alloys, which has significant technological implications ranging from economic savings via improved fuel efficiency and reducing the carbon footprint in automotive and aerospace sectors.

Title: Computational design of novel semiconductors for power and energy applications

Principal Investigator: Feliciano Giustino (The University of Texas at Austin)

Co-investigators: Emmanouil Kioupakis (University of Michigan),
Joshua Leveillee (The University of Texas at Austin),
Jon Lafuente-Bartolome (The University of Texas at Austin)

ALCC Allocation:

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

Allocation(s): 883,000 node-hours on Theta

Research Summary:

One of the primary engineering challenges in developing new and more efficient energy systems is to design and deploy functional materials with tailored electric, thermal, mechanical, and optical properties, e.g. for solar photovoltaics, solid-state lighting, and power electronics. In this broad context, atomic-scale computer simulations are invaluable to complement the experimental characterization of new materials and to guide the discovery of new compounds. In this project, we aim to develop and leverage cutting-edge ab initio computational methods to investigate the temperature-dependent optical and transport properties of compound semiconductors, perovskites, and plasmonic ceramics for uses in power devices, solar cells, and light-emitting diodes.

Our long-standing expertise in the computation of electron-phonon interactions and related properties from first principles, together with the leadership-class computing resources provided by the ALCC program, will prove essential to undertake this project successfully. The new methodologies and computational workflows developed during the course of this project will be made available to the community through the open-source EPW code. The completion of this project will broaden the library of compounds available today for power and energy applications, and will push forward our overarching aim to introduce experimentally relevant temperature effects in the computer-aided design of new materials.

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

Principal Investigator: Steven Gottlieb (Indiana University)

Co-investigators: Alexei Bazavov (Michigan State U),
Carleton DeTar (U Utah),
Aida El-Khadra (U Illinois Urbana-Champaign),
Andreas Kronfeld (Fermilab),
Ethan Neil (U Colorado),
Ruth Van de Water (Fermilab)

ALCC Allocation:

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

Allocation(s): 1,000,000 node-hours on Summit,
100,000 node-hours on Polaris,
100,000 node-hours on Perlmutter-GPU

Research Summary:

The muon magnetic moment is one of the most precisely measured quantities in physics, but its value is not in good agreement with the theoretical prediction. On April 7, 2021, the Fermilab Muon g-2 (E989) experiment announced its first result in spectacular agreement with the Brookhaven National Laboratory E821 experiment done almost 20 years ago. In the next few years, the Fermilab experiment plans to increase its precision fourfold. Computers at Oak Ridge Leadership Computing Facility, Argonne Leadership Computing Facility, and the National Energy Research Scientific Computing Center will be used to refine the theoretical prediction based on a technique known as lattice Quantum Chromodynamics (QCD).

This calculation is both critical and very timely. It addresses a priority of the DOE Office of High Energy Physics. The most recent Particle Physics Project Prioritization Panel (P5) report “Building for Discovery: Strategic Plan for U.S. Particle Physics in the Global Context" defines several science drivers including “Explore the Unknown: New Particles, Interactions, and Physical Principles." One of the key opportunities highlighted in the report is study of the anomalous magnetic moment of the muon.

If the theoretical precision can be increased, and the difference between theory and experiment remains, it will be strong evidence for new particles or forces in Nature that are not part of the Standard Model of Elementary Particle and Nuclear Physics. This would be truly exciting.

Title: Improved Seismic Hazard Modeling Using Physics-based Simulations

Principal Investigator: Christine Goulet (Southern California Earthquake Center, University of Southern California)

Co-investigators: Yifeng Cui, Kim Olsen, Kyle Withers, Scott Callaghan

ALCC Allocation:

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

Allocation(s): 472,500 node-hours on Summit
783,000 node-hours on Cori

Research Summary:

Seismic risk reduction is a long-term endeavor that requires research programs focused on reducing uncertainties at various steps of the evaluation process, starting from the hazard. The reduction of uncertainties will lead to more effective engineering designs and reduce financial losses and human casualties due to earthquakes. Most types of U.S. infrastructure, from individual buildings to distributed infrastructure (energy, water, transportation) are subject to potential damage from earthquake ground motions and/or near-surface fault displacements. The timing and spatial distribution of ground motions also has a big impact on the damage of distributed systems, for which component failures can paralyze the whole system. To address these issues, the project will use physics-based simulations to model earthquake ruptures and wave propagation to the Earth's surface. The first step involves the validation of the models against recorded data, followed by the investigation of the impact of different types of models on the resulting ground motions and fault displacements. The proper quantification of the various cause and effects combinations can only be achieved through simulations for which both the input and output are known. This will help better quantify uncertainties in seismic hazard, which based upon specific local conditions can then be reduced. In addition to their potential for broader societal impacts, the simulations produced through these activities will stimulate years of scientific and engineering research, contributing to improved science and design innovation.

Title: Monte Carlo neutron transport for high burnup/high enrichment nuclear fuel

Principal Investigator: Steven Hamilton (Oak Ridge National Laboratory)

Co-investigators: Thomas Evans (Oak Ridge National Laboratory), Fausto Franceschini (Westinghouse)

ALCC Allocation:

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

Allocation(s): 250,000 node-hours on Summit

Research Summary:

The US nuclear industry is aggressively pursuing high burnup (HB) and high enriched (HE) fuel to increase sustainability of Light Water Reactors (LWRs). The current nuclear fleet is comprised of 94 LWRs that provide 20% of US electricity production. In addition, two new advanced LWRs, Westinghouse AP1000 units, will soon begin operation in the US. HB/HE fuel consists of Uranium-235 enrichments of up to 8 weight percent range compared to the current licensed limit of 5 weight percent. This increase in Uranium-235 content offers significant benefits for the current fleet, including increased operational flexibility and improved economic performance due to the longer operational cycles that can be achieved. Adopting HB/HE fuel would not only increase plant availability and reduce the need for refueling outages, but also reduce the number of spent fuel assemblies that need to be stored and disposed due to the higher discharge burnup.

Due to the lack of prior operating experience and the limited availability of experimental data for HB/HE fuel, highly accurate numerical simulations are needed to predict the operational behavior for reactors using these fuel designs. Such simulation capabilities can accelerate licensing and commercial deployment of advanced nuclear fuel by providing numerical benchmark data to reliably characterize the expected neutronic behavior of the fuel and support validation of the current industry tools and methods. This project will employ the Shift Monte Carlo neutron transport code, which is developed at Oak Ridge National Laboratory with support from both the DOE Nuclear Energy Advanced Modeling and Simulation (NEAMS) program and the ExaSMR project within the DOE Exascale Computing Project (ECP). It has been optimized for use on leadership class computing hardware such as the OLCF Summit supercomputer to perform Monte Carlo simulations of a full reactor core with isotopic depletion to model the evolution of the nuclear fuel over the duration of an operational cycle of a reactor. Shift has previously proved beneficial to industry to address validation challenges for advanced core designs with highly heterogeneous fuel loading patterns and heavy neutron burnable absorber loadings, confirming its highly predictive capabilities against actual start-up measurements from four Westinghouse AP1000 units. This excellent track record makes Shift uniquely positioned to address the challenges entailed in a transition to HB/HE fuel, where novel loading patterns with a high degree of heterogeneity and heavy burnable absorbers loading will be deployed to achieve optimum core economic and safety performance.

Title: Privacy-preserving Transformer models for clinical natural language processing

Principal Investigator: Heidi Hanson (Oak Ridge National Laboratory)

Co-investigators: Andrew Blanchard (ORNL),
Mayanka Chandra Shekar (ORNL),
John Gounley (ORNL),
Isaac Lyngaas (ORNL),
Adam Spannaus (ORNL),
Chris Stanley (ORNL),
Xiao Wang (ORNL),

ALCC Allocation:

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

Allocation(s): 150,000 node-hours on Summit,
30,000 node-hours on Frontier

Research Summary:

Public health surveillance systems are critical sources of actionable intelligence that can be used to protect the population against health threats. However, the establishment of well-curated, comprehensive, and population-wide data sets that can be used for real-time health surveillance is an unmet need in the United States (U.S.). Our overall vision is to create a set of tools that can be used by public health agencies and academic researchers to predict, detect, and respond to public health threats in near-real time. The MOSSAIC (Modeling Outcomes using Surveillance data and Scalable AI for Cancer) program has paved the way for realizing these goals for cancer, the second leading cause of death in the U.S. In the U.S. in 2020, there were an estimated 1.8 million new cancer diagnoses and 600 thousand cancer deaths. The MOSSAIC program belongs to the Joint Design of Advanced Computing Solutions for Cancer (JDACS4C), a program established as a partnership between the Department of Energy (DOE) and the National Cancer Institute (NCI) to accelerate cancer research by integrating high performance computing and deep learning approaches.

In our ongoing 2021–2022 ALCC project “Next-generation scalable deep learning for medical natural language processing”, we are using Summit to pre-train custom Transformer models on our clinical text corpus. As existing publicly available Transformer models underperform on clinical text, we are undertaking several improvements in our custom pretraining to handle long clinical documents and the various nuances of clinical text. A notable challenge of our current project is that, because the clinical text contains protected health information (PHI), these models must be trained securely using, e.g., the CITADEL capability on Summit at the Oak Ridge Leadership Computing Facility. As the trained models may memorize PHI from our data, they cannot be publicly released to the biomedical research community. The differential privacy and federated learning methods we have been exploring offer a solution to this problem and facilitate population-level health surveillance.

Title: Gyrokinetic Prediction of Burning Plasma Profiles Enabled by Surrogate Modeling

Principal Investigator: Nathan Howard (Massachusetts Institute of Technology)

Co-investigators: Pablo Rodriguez-Fernandez (MIT), Chris Holland (UCSD)

ALCC Allocation:

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

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

Research Summary:

Predicting performance of magnetic confinement fusion devices is key to the prompt realization of fusion as an energy source, as it provides an efficient means to explore operational regimes and optimize the design of fusion devices. However, to date, researchers have been unable to apply their highest fidelity physics models to predict fusion performance, as the extreme computation required has made it mostly intractable. As a result, next generation fusion reactors, such as ITER, were designed and planned only with simplified empirical models. This research seeks to reduce previous limitations by utilizing new scale-bridging techniques to find steady-state solutions for the kinetic profiles (density and temperature) in magnetic fusion devices while utilizing the most complete physics models available. Novel machine learning techniques have been coupled with physically accurate plasma turbulence calculations to achieve ~5x faster convergence (compared with traditional methods) to steady state solutions, with extremely high physics fidelity. This new approach will be used to predict the core kinetic profiles of next generation fusion devices such as ITER, SPARC, and candidate Fusion Pilot Plant (FPP) designs with unprecedented accuracy. These tools will allow researchers to explore the origin of well-known problems in fusion science, such as the favorable scaling of fusion energy confinement with fuel mass, and most importantly will enable the prediction of the fusion power generation in anticipated reactor scenarios.

Completion of this project will result a significant advancement in our ability to accurately predict fusion plasmas and will provide one of the first glances into whether favorable energy confinement properties with increased fuel mass (termed the “isotope effect”) are predicted by first-principles models to be observed in ITER, SPARC, Fusion Pilot Plant (FPP) candidates. This work may have significant implications for the design and operation of these devices, as their designs often assume that such favorable energy confinement properties will exist, based purely on empirical observations in current fusion devices. This research leverages leadership-class high performance computing, nonlinear plasma turbulence simulation, and cutting-edge machine learning to directly contribute to the Department of Energy Fusion Energy Sciences (DOE-FES) mission of advancing our predictive capabilities and furthering our understanding of burning plasma conditions that can lead to the realization of fusion energy.

Title: Using GPU to reconstruct LHC collisions recorded with the CMS detector

Principal Investigator: Dirk Hufnagel (Fermilab)

Co-investigators: Stephan Lammel (Fermilab)

ALCC Allocation:

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

Allocation(s): 50,000 node-hours on Summit

Research Summary:

The Compact Muon Solenoid, CMS, detector records high-energy proton---proton collisions of the Large Hadron Collider, LHC. The CMS collaboration operates the detector and analyzes the collisions to search for the fundamental constituents of matter, to precisely measure the forces between them, to identify new symmetries, and search for dark matter.

Position, timing, and energy deposits recorded by the detector are reconstructed back to particles generated in the collision and their momenta. These collisions are compared to collisions simulated with Monte Carlo techniques from the standard model, SM, extension of the SM, and other theoretical models to gain insight into nature at small scale or the conditions in our universe at about a picosecond age.

So far high-energy physics, HEP, experiments have used local and grid computing resources to reconstruct and simulate collisions. After the next upgrade of the LHC, the collisions will be so complex that grid resources alone will be insufficient. CMS are exploring high-performance computing, HPC, usage since a few years. HPC resources pose two challenges for CMS and HEP experiments in general:

1. The grid computing environment reflects the open collaborative approach in large HEP experiments. HPC resources, and especially leadership class facility HPC resources, have a tighter security model. This makes integration more difficult especially for our data intensive workflows where data needs to be transferred into and out of the site.
2. HEP software was written for CPUs. Migrating our software to use GPU started a few years ago. A number of algorithms have already been adapted and will be used in the processing of the upcoming data taking. For high-luminosity, HL-LHC, we anticipate simulation, reconstruction, and even some analysis workflows to fully harness the compute power provided by GPUs.

The objective of this proposal are two fold:

- to use GPU resources at OLCF Summit for the reconstruction of some of the Run3 (2022) data collected at the CMS detector at the LHC
- to develop and exercise better integration of HPC resources into the CMS computing infrastructure. This would be the first time USCMS has used the Power+nVidia architecture for production.

Title: Informing Forensics Investigations of Nuclear Materials

Principal Investigator: Sara Isbill (Oak Ridge National Laboratory)

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

ALCC Allocation:

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

Allocation(s): 175,000 node-hours on Summit

Research Summary:

Researchers at Oak Ridge National Laboratory actively support national and international efforts to prevent and deter the proliferation of nuclear weapons by developing tools and methods to detect nuclear material and processes. This research requires the use of advanced experimental and computational resources to more fully understand process kinetics and environmental degradation of key fuel cycle materials. Determining the effect of process conditions on the underlying crystal structure via electronic structure calculations allows direct connection of chemical and physical changes in fuel cycle materials to observations in the laboratory. These calculations can assist in all areas of this research in three key ways 1) predicting relative stability of perturbed materials, 2) predicting physical and chemical changes due to perturbations, and 3) predicting spectroscopic and other observables. The prediction of optical spectroscopic observables is of particular interest, as optical vibrational spectroscopy is the primary experimental technique used for non-destructive chemical composition determination, possessing sufficient fidelity to be of utility to nuclear forensics.

The chief aim of this project is to connect highly accurate density functional theory (DFT) determinations of lattice dynamics under equilibrium and non-equilibrium conditions to ongoing experiments collecting vibrational spectra obtained from Raman, infrared, and neutron scattering. This project will bridge the gap between experimental observations and an atomistic understanding of the impact of defects and other structural perturbations induced under fuel cycle relevant conditions to inform chemical characterization and forensic analysis. Systems of interest to this research are defective uranium and carbon materials and will extend upon previous work by calculating the stability and lattice dynamics of additional defects under equilibrium and non-equilibrium conditions.

Although calculating the lattice vibrations of large-scale defective materials at the level of DFT is computationally intensive, the high accuracy of DFT calculations is necessary for comparison to experimental observables, particularly vibrational spectra, and is a necessary complement to the growing nonproliferation materials research portfolio.

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

Principal Investigator: Wei Jiang (Argonne National Laboratory)

Co-investigators: Zhengcheng Zhang

ALCC Allocation:

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

Allocation(s): 500,000 node-hours on Theta

Research Summary:

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

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

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

Title: A Multiscale Surrogate Model for Fracture Evolution using DeepONet

Principal Investigator: George Karniadakis (Brown University)

Co-investigators: Somdatta Goswami

ALCC Allocation:

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

Allocation(s): 40,000 node-hours on Summit,
50,000 node-hours on Polaris,
60,000 node-hours on Perlmutter-CPU

Research Summary:
The research work is inspired by the increasing preference for data-driven models for creating atomistic-continuum multi-scale fracture analysis approaches employing machine learning techniques. Understanding fracture mechanics requires modeling of deformation and crack growth in material microstructures. The complex mechanisms involved in the process of crack growth operate on a variety of spatial scales, from the atomic scale of fracture initiation to the continuum scales of short crack propagation along with the microstructure deformation. Pure atomistic simulations using molecular dynamics (MD) can be used to investigate these processes. It is still very difficult to perform large-scale MD simulations of fracture propagation in larger domains that correspond to continuum length and time scales. Furthermore, the computationally demanding high-fidelity numerical solvers that can precisely anticipate the relevant parameters demand a fine-grained fracture resolution. Besides that, independent, expensive simulations must be run for every little change in the domain parameters and/or material properties. Partially getting around this restriction requires combining the continuum and atomistic models. The remaining portions of the domain are represented using continuum-scale constitutive relations, while critical sections with high stress intensity (crack tip) are modeled using MD at the atomic scale. The main objective of the research project is to develop a data-driven operator level surrogate model (DeepONet) to take the place of atomic scale modeling in a time-dependent multiscale system. As noted by DeepMind researcher Irina Higgins, “Once DeepONet is trained, it can be applied to new input functions, thus producing new results substantially faster than numerical solvers.”

This project will be the first of its kind to provide a unified fracture analysis model that can be used for failure analysis in any engineering material. It will also be a one-shot solution to determine the likely fracture path in order to prevent unexpected failures. By organizing materials in unique ways, based on predetermined failure paths, this research opens up the possibility of synthesizing a material with customized structural properties and the optimal topology and shape. It is directly related to DOE mission since this is a part of the PhILMs collaboratory is directing.

Title: Shock Turbulent Boundary Layer Interaction in Supercritical CO2 Flows

Principal Investigator: Sanjiva Lele (Stanford University)

Co-investigators: None

ALCC Allocation:

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

Allocation(s): 180,000 node-hours on Summit,

Research Summary:

Supercritical CO2 (sCO2) power cycles offer key advantages in energy conversion due to compactness, high thermal efficiency and lower life cycles cost from corrosion resistance. They can operate with a variety of heat sources, including sustainable sources, such as geothermal power, concentrated solar power and high temperature fuel cells. However, its development requires further advances in design and optimization of turbomachinery and heat exchangers. Improved fundamental understanding of the complex thermodynamic effects in turbulent fluid flows in supercritical regime which adversely affect turbomachinery and heat-exchanger performance is needed. This project studies fundamental turbulent flow behavior in sCO2 flows, with focus on shock turbulent boundary layer interaction (STBLI), a major contributor to performance losses.

Current state of the art in turbomachinery flow analysis and design is based on solving the Reynolds Averaged Navier-Stokes (RANS) equations. However, the existing RANS models are not aware of turbulent transport effects in thermodynamically complex fluids. This study is the first of its kind in generating a database of turbulent transport behavior under supercritical conditions for sCO2 turbomachinery relevant flows. The simulations use state-of-the-art parallel computing algorithms, high-order, low-dissipation and robust numerical methods. The simulation database from this work will be well documented and available for open science collaborations. Fundamental data on turbulent flow behavior in thermodynamically complex fluids may be of interest to wider applications in aerospace industry and energy technology.

Title: Terrestrial ecosystem carbon cycle of the conterminous U.S.

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

Co-investigators: Benjiamin Sleeter (USGS),
Zhiliang Zhu (USGS),
Qiuan Zhu (Hohai University China)

ALCC Allocation:

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

Allocation(s): 300,000 node-hours on Cori

Research Summary:

Terrestrial carbon (C) cycle is one important aspect of Earth's energy balance. Yet significant uncertainties exist in the quantification of regional carbon budget and its underlying driving forces. For the conterminous U.S. (CONUS), estimates of the annual land C sink during recent decades vary from 200 to 685 million metric tons per year, indicating serious discrepancies in data and methodology (inventory, modeling, and atmospheric inversion). This also indicates that land C sink derived as the residue of C budget of atmosphere and oceans could be inaccurate. Therefore, detailed independent estimates of land C sink in regional and global scale becomes a critically needed activity.

This research plans to quantify for the first time the terrestrial ecosystem C sequestration and greenhouse gas emission of the CONUS from 1971 to 2100 at 1-km resolution. Combining with newly available 30-m resolution land cover change data, 250-m remote sensing ecosystem productivity data, county level forest inventory data and county level agricultural census data, this research will apply three models: (1) the parallel Integrated Biosphere Simulator (IBIS) as the hourly carbon model; (2) the Unit Stream Power-based Erosion Deposition model (USPED), as the yearly C lateral redistribution model; and (3) the Land Use and Carbon Scenario Simulator (LUCAS) as the yearly land use and land cover change model for both baseline (1971-2020) and future scenarios (2021- 2100).

The goals include: 1. Assess terrestrial C dynamics of CONUS at 1-km by considering major ecosystem types (forest, shrub, grass, agriculture, wetland, and urban), major C pools (live and dead woody biomass, ground litter, soil carbon, grain and wood product pool), key GHG fluxes (CO2, CH4 and N2O emission), and major controlling processes (climate, wild fire, harvesting, erosion/deposition, land use and land cover change (LUCC), coastal sea level rise); 2. Quantify C sequestration on land as well as C transport across landscape and from land to aquatic system; 3. Use observational data from remote sensing, forest inventory and agricultural statistics to constrain model parameter and to evaluate model uncertainty; 4. Provide selected future LUCC and carbon simulations (2021-2100) under the IPCC Representative Concentration Pathway (RCPs).

This research potentially connects with DOE BER's Scientific Discovery through Advanced Computing (SciDAC) program and the Earth System Model Development and Analysis Program.

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

Principal Investigator: Zarija Lukić (Lawrence Berkeley National Laboratory)

Co-investigators: Solène Chabanier (Lawrence Berkeley National Laboratory),
Hyunbae Park (Lawrence Berkeley National Laboratory),
Jean Sexton (Lawrence Berkeley National Laboratory),
Nicholas Frontiere (Argonne National Laboratory),
JD Emberson (Argonne National Laboratory),
Salman Habib (Argonne National Laboratory),
Katrin Heitmann (Argonne National Laboratory),
Adrian Pope (Argonne National Laboratory),
Esteban Rangel (Argonne National Laboratory)

ALCC Allocation:

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

Allocation(s): 50,000 node-hours on Summit,
100,000 node-hours on Polaris,
50,000 node-hours on Perlmutter-GPU

Research Summary:

DOE-supported sky surveys such as Dark Energy Spectroscopic Instrument (DESI) and Rubin Observatory (a.k.a. LSST) require sophisticated hydrodynamical simulations in order to interpret the data and produce realistic mock skies on which working groups can test different methodologies. Combining in this proposal two exascale-ready cosmology codes – Nyx and HACC – allows us to model a wide range of redshifts and cosmological observables, from the epoch of reionization and vast regions of the intergalactic medium, to small-scale physics relevant for modern-day galaxies.

The ongoing DESI survey will observe 0.7 million quasars at z>2, more than three times as much as collected with previous surveys. This exquisite dataset will strongly improve several cosmological constraints, e.g., neutrino mass, the nature of dark matter, and the epoch of inflation. To interpret DESI, especially focusing on the DESI Year 1 Data Release of the Lyman alpha forest, scheduled for the summer of 2023, we will produce a set of high-resolution hydrodynamical simulations which are needed to accurately compute the large-scale clustering of the intergalactic medium within the cosmic web while capturing small-scale effects due to pressure broadening of gas or a finite mass of dark matter particles leading to free-streaming, at sufficient precision. At lower redshifts, the interpretation of the galaxy clustering and weak lensing signal from data collected by the LSST will be limited if we do not understand the effects of baryonic physics on small scales. Sophisticated hydrodynamical simulations which will be run under this project will include different subgrid/feedback models and will provide urgently needed predictions to fully exploit new information present in the LSST data.

Title: E3SMv2 Smoothed Biomass Burning Large Ensemble

Principal Investigator: Gerald Meehl (National Center for Atmospheric Research)

Co-investigators: John Fasullo (National Center for Atmospheric Research), Nan Rosenbloom (National Center for Atmospheric Research)

ALCC Allocation:

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

Allocation(s): 260,000 node-hours on Perlmutter-CPU

Research Summary:

A discontinuity in the interannual variability of prescribed biomass burning (BB) emissions in the Coupled Model Intercomparison Project (CMIP6) forcing database during the satellite era of wildfire monitoring (1997–2014) leads to spurious warming in the Northern Hemisphere extratropics in simulations with the Community Earth System Model version 2 (CESM2). The warming is specifically attributable to biomass burning variability from 40° to 70°N and arises from a net thinning of the cloud field and an associated increase in absorbed solar radiation. It is unknown whether other models respond in the same way. The objectives of this research are to explore the sensitivity of the Department of Energy's (DOE) Energy Exascale Earth System Model, version 2 (E3SMv2) to these biomass burning effects. We will conduct an abbreviated 20-member E3SMv2 large ensemble of historical+future simulations (1990-2080 with SSP370 forcing) using smoothed biomass burning in which prescribed BB emissions are homogenized and variability is removed. These abbreviated simulations will be compared against E3SMv2 simulations forced by biomass burning from the standard CMIP6 database.

This project supports the Department of Energy's grand challenge of actionable projections of Earth system variability and change, by using E3SMv2 to address the fidelity of model response to forcing in the satellite era by specifically exploring the sensitivity of E3SMv2 model simulations to biomass forcing.

Title: High-Fidelity Flow Data for Multiscale Bridging: Year 2

Principal Investigator: Elia Merzari (Pennsylvania State University)

Co-investigators: Igor Bolotnov (NCSU),
Nam Dinh (NCSU),
Paul Fischer (UIUC),
Misun Min (ANL)

ALCC Allocation:

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

Allocation(s): 250,000 node-hours on Summit

Research Summary:

In 2020, the US Department of Energy (DOE), Office of Nuclear Energy established a 3-year consortium, through the Integrated Research Projects (IRP) program, focused on the research of thermal-fluids multiscale methods. This consortium represents a university component to the recently established Center Excellence for Thermal-Fluids Applications in Nuclear Energy that is fully integrated with the national laboratory efforts and its stakeholders (e.g., industry, DOE programs, Nuclear Regulatory Commission).

The consortium aims to deliver improved, fast-running models for complex physical phenomena relevant to advanced reactors. The work of the consortium covers numerical method development, high-fidelity numerical simulation and experimental research. We are focusing in particular on four challenge problems of great interest to the advanced reactor nuclear industry: Flexible Modeling for Heat Transfer, Thermal Striping of Internals, Mixing in Large Enclosures and Multiscale Core Modeling Coupled to Fuel Performance.

To support the goals of the project, this proposal seeks to support the establishment of an extensive high resolution turbulence numerical database (Direct Numerical Simulation and Large Eddy Simulation). The geometries simulated span from experimental facilities involving jet mixing in large enclosures to more canonical geometries such as parallel plates in mixed convection and axial flow in simplified rod arrays. The datasets has already proven to be instrumental to the development of improved multiscale bridging methods. The planned work builds on the effort of a successful ALCC project awarded in 2021/2022. We envision that the work will fundamentally increase the understanding and modeling of turbulent thermal mixing in complex geometries and ultimately enable a paradigm shift in the way multiscale simulations are performed.

Title: Large-eddy simulations of dynamic stall in a boundary layer ingesting turbofan

Principal Investigator: Parviz Moin (Stanford University)

Co-investigators: Sanjeeb Bose (Cascade Technologie and Stanford University)

ALCC Allocation:

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

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

Research Summary:

Limiting the carbon footprint of air transport has been identified as one of critical needs in the coming decades for environmental sustainability. Boundary layer ingesting propulsion devices offer the potential to significantly improve fuel efficiency for air transport through two mechanisms. First, the reduced velocity fluid in the fuselage boundary layer is re-accelerated through the engine thereby reducing wake penalties. Second, integrating the nacelle into the fuselage with more aerodynamic shapes can reduce drag relative to contemporary nacelles that appear as bluff bodies hung underneath the wing. The operability limits of such devices can be limited by the onset of separation around the fan blades. The ingestion of the boundary layer, however, creates azimuthal asymmetry that leads to different stall mechanisms that are encountered by traditional turbofans. To this end, we plan to assess the ability of high-fidelity simulations to predict the operability limits (particularly the onset of a dynamic stall) in the NASA Boundary Layer Ingesting Turbofan (Cousins et al. 2017, Arend et al. 2020). Calculations will be conducted at several operating speeds and multiple grid resolutions and the accuracy of the results will be evaluated from comparisons against experimental data collected at NASA Glenn Research Center.

There have been numerous recent successes within our group in the ability for the high fidelity, large-eddy simulations for capturing stall characteristics in realistic geometries ranging from commercial aircraft models (Goc et al. 2021, Goc et al 2022) to axial compressor rotors (Bose et al, forthcoming). Simulations of the latter configuration have demonstrated the ability to characterize the onset of stall in an isolated sector and full wheel of a characteristic rotor, where the observed stall mechanism exhibits relative azimuthal symmetry. However, the simulations of the boundary layer ingesting turbofan incorporate new challenges into the prediction of the stall boundaries. Notably, the boundary layer ingesting turbofan configuration requires the ability to capture a dynamic stall that does not exhibit azimuthal symmetry at any operating condition due to the ingestion of the fuselage boundary layer and requires incorporation of a full stage (fan rotor and exit guide vanes).

The dynamic stall phenomena that is the focus of this study stresses some of the assumptions that are utilized in wall modeled LES approaches. The results of this study are expected to have spillover effects into other problems that experience similar dynamic stall events including aircrafts under maneuver or wind turbine applications.

Title: Predictive Simulations of Inertial Confinement Fusion Ablator Materials

Principal Investigator: Ivan Oleynik (University of South Florida)

Co-investigators: Aidan Thompson (Sandia National Laboratory),
Mitchell Wood (Sandia National Laboratory),
Stan Moore (Sandia National Laboratory),
Rahulkumar Gayatri (National Energy Research Scientific Computing Center)

ALCC Allocation:

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

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

Research Summary:

Very recent record breaking 1.3-megajoule (MJ) energy-yield experiment at DOE's National Ignition Facility (NIF) has brought practical realization of inertial confinement fusion (ICF) ignition within a reach. The ablator capsule converts the energy of high-power lasers into ablation pressure, which compresses and ignites the fusion fuel. To step up the energy yield, new alternatives to diamond ablator materials are urgently sought. This project will study the dynamic shock response of novel amorphous carbon ablator material in quantum-accurate billion atom molecular dynamics simulations at experimental time and length scales.

The extreme-scale MD simulations will deliver key information on physical properties of amorphous carbon including phase diagram, shock Hugoniot and the equation of state, the critical elements for the design of successful ICF ignition. They will also guide experiments at NIF and Omega EP laser facility at DOE's Laboratory for Laser Energetics at the University of Rochester.

Title: Probabilistic Comparative Modeling of Colorectal Cancer Screening Strategies

Principal Investigator: Jonathan Ozik (Argonne National Laboratory)

Co-investigators: Carolyn Rutter (RAND Corporation),
Iris Lansdorp-Vogelaar (Erasmus University Medical Center),
Karen Kuntz (University of Minnesota),

ALCC Allocation:

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

Allocation(s): 283,000 node-hours on Theta

Research Summary:

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

This project will use leadership class computing resources to run comparative probabilistic sensitivity analyses (PSAs) of screening strategies with three state-of-the-art CRC models. Funded under the National Cancer Institute's (NCI) Cancer Intervention and Surveillance Modeling Network (CISNET) program, these models were independently developed for the evaluation of interventions, with emphasis on screening, and describe CRC natural history using different underlying assumptions. Building on model calibration, comparison, and evaluation of screening efficacy that was accomplished in the 2021-2022 period, the project will extend analyses to extensions of the microsimulation models that will incorporate the serrated pathway to colorectal cancer. Research examining the serrated pathway has been thwarted by uncertainty in this process, which will be addressed by incorporating additional sources and combinations of uncertainty into simulations to examine their impacts on the projected results, including its effects on fecal immunochemical test (FIT) screening. The comparative PSAs in this work will be used to generate cost-effectiveness analyses for complex interventions and to provide formalized assessments of uncertainties across the three CRC models. This work integrates complex data, large-scale machine learning algorithms for uncertainty quantification, and simulation to advance HPC-enabled scientific discovery.

Title: Particle-in-cell simulations of beam-driven, field-reversed configuration plasmas

Principal Investigator: Jaeyoung Park (TAE Technologies, Inc.)

Co-investigators: Giovanni Lapenta (KU Leuven, Belgium)
Calvin Lau (TAE Technologies, Inc.)
Richard Magee (TAE Technologies, Inc.)

ALCC Allocation:

Site(s): Argonne Leadership Computing Facility (ALCF)
Allocation(s): 400,000 node-hours on Theta

Research Summary:

TAE Technologies combines accelerator physics and plasma physics to solve the challenge of fusion. The company has been developing an advanced, beam-driven Field-Reversed Configuration (FRC) for superior plasma confinement since 1998 and has been making significant strides toward net fusion power. The fifth-generation device, called C-2W, is currently in operation and has met its performance targets of total temperature greater than 3 keV, electron temperature greater than 500 eV and steady-state FRC operation. At present, TAE is in the design phase of the next device, called "Copernicus," which will achieve the plasma parameters necessary for net power production for deuterium-tritium fuel while running on hydrogen.

One of the critical questions currently facing the design of Copernicus is the global FRC stability of the plasma in the high density reactor regime. Since the fusion power output is proportional to the square of the density, high density operation is a major R&D focus towards developing compact, economical fusion power. Recent data from C-2W experiments indicate that there may have been a major breakthrough to overcome the long-standing FRC global stability limit with the use of neutral beam injection (NBI). As part of an ongoing investigation, we will utilize ALCF HPC resources to conduct first principles particle-in-cell (PIC) simulations to develop an understanding of this newly identified and highly impactful regime. TAE will collaborate with Prof. Lapenta at KU Leuven to apply the ECsim (Energy Conserving semi-implicit model) code to the NBI-stabilized FRC. ECsim is a massively parallel (tested up to 32,000 CPU cores) first-principles PIC code that has been successfully utilized for many challenging plasma physics problems such as magnetic reconnections in the Earth's magnetosphere and high-beta plasma systems for magnetic fusion.

We expect that successful ECsim simulations of high-density C-2W FRC plasmas will shed light on the role of energetic ions from NBI in global stabilization of the FRC. Knowledge of the underlying mechanisms of FRC stabilization in the current C-2W device will be applied to the next generation Copernicus device that is aimed at achieving a net power producing plasma condition. In addition, the project will demonstrate the use of public HPC resources to accelerate R&D in the private fusion sector, fostering future private-public ventures in the fusion industry.

Title: HFIR DNS simulations

Principal Investigator: Emilian Popov (ORNL)

Co-investigators: Igor Bolotnov (NCSU),

ALCC Allocation:

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

Allocation(s): 224,000 node-hours on Theta

Research Summary:

The High Flux Isotope Reactor (HFIR) at Oak Ridge National Laboratory (ORNL) is a source of thermal and cold neutrons for research projects throughout the US used to study phenomena in numerous scientific and engineering disciplines. As part of the US Department of Energy (DOE) National Nuclear Security Administration's (NNSA) initiative to reduce the enrichment of research and test reactors, a research project is underway to investigate the conversion of HFIR from a high enriched uranium (HEU) core to a low enriched uranium (LEU) core. Due to the complex channel geometry and the difficulty of performing LEU testing and experiments, data supporting this conversion is highly limited. Therefore, high-fidelity numerical data is required to verify and calibrate Reynolds-Averaged Navier-Stokes models.

Direct numerical simulation (DNS) of turbulent single- and two-phase flows at a leadership computing facility allows for users to attain unprecedented level of detail and can answer fundamental questions about the interaction and evolution of turbulence within complex geometries. The highly detailed simulation of all turbulent structures using a DNS approach will allow for the collection of statistical information relevant to turbulent flow parameters, such as the turbulent kinetic energy, k-ε model constants and the Prandtl number required to enforce the correct wall heat transfer.

Two subprojects will generate single-phase flow results for geometries relevant to HFIR, including (i) complex turbulent flow through the full radial span of a HFIR coolant channel and (ii) the simulation of turbulent flow through involute plates representing a portion of the HFIR coolant channel. The research code PHASTA will be utilized to study complex physical phenomena in unprecedented detail and allow for the collection of numerical data based on first-principal calculations. Major statistical parameters, such as the mean velocity profile, turbulent kinetic energy and secondary flow structure formation will be assessed to develop a new and more sophisticated understanding about single-phase flows in an involute geometry.

With the assistance of leadership computing facilities, detailed numerical data will be produced for the development of new closure laws to improve the prediction accuracy of computational fluid dynamics (CFD) models. These developments will help to capture turbulent flows in fine detail and facilitate the thermal-hydraulic design of HFIR and next generation energy systems.

Title: High-Fidelity Simulations of Turbulent Aeroacoustics Enabling Sustainable Aviation

Principal Investigator: Stephan Priebe (GE Research)

Co-investigators: Trevor Wood, Junsok Yi, Kishore Ramakrishnan (GE Research); and Daniel Tweedt (GE Aviation)

ALCC Allocation:

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

Allocation(s): 130,000 node-hours on Summit,
35,000 node-hours on Perlmutter-GPU,
50,000 node-hours on Frontier

Research Summary:

There is a societal need to eliminate use of fossil fuels for aircraft propulsion and achieve sustainable energy in flight. GE recently announced, along with its CFM partner, Safran, the Revolutionary Innovation for Sustainable Engines (RISE[1]) technology demonstration program to provide over 20% lower fuel consumption and CO_2 emissions compared to today's engines. Most of that benefit stems from the open fan architecture, in which the large fan at the front of the engine is unducted rather than being enclosed in an outer nacelle as typically found on today's engines. The RISE program also announced compatibility with sustainable aviation fuels for further CO_2 reductions, including hydrogen which enables a path to zero CO_2 emissions ultimately. However, with high volume per unit energy for hydrogen, aircraft range and capability (e.g., passenger load) could be significantly reduced for a given vehicle size and capability. Therefore, the path to sustainable flight starts with minimizing the energy needed for air transport, and the open fan architecture is well suited to maximize propulsive efficiency beyond what is achievable by turbofan propulsion systems.

This research aims to study two important aspects of open fan propulsion: (1) the effect of noise reduction technologies on the aerodynamic and acoustic performance; and (2) how the complex physics of flow separation, wake and turbulence prediction scale from subscale rig test to an aircraft in flight. These effects require the most advanced numerical simulations well beyond the capabilities of state-of-the-art modeling approaches, and the computational cost grows exponentially with the size of the model which is why past simulations have focused on smaller scale models, leaving the question of how our learnings from past research may translate to the "real thing". Supercomputing and application of Large Eddy Simulation (LES) allows us to extend prior research to realistic flight conditions to discover new ways of controlling turbulence and approach optimal performance for these complex flows.

[1]RISE is a registered trademark of CFM International, a 50/50 joint company between GE and Safran Aircraft Engines

Title: Modeling operating conditions in the US east coast offshore wind energy lease areas

Principal Investigator: Sara C Pryor (Cornell University)

Co-investigators: Rebecca J Barthelmie (Cornell University),
Tristan J Shepherd (Cornell University)

ALCC Allocation:

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

Allocation(s): 142,000 node-hours on Theta

Research Summary:

Recent research by the PI team suggests that if the current offshore wind energy lease areas along the US east coast are subject to deployments of wind turbines at the European average spacing they will meet 3% of national electricity supply. However, the associated levelized cost of energy (LCoE) is critically dependent on the nature of the operating conditions these wind turbines (WT) will be subject to during their anticipated deployment lifetimes of 30 years. Our research seeks to reduce uncertainty in 2 key aspects of likely operating conditions. It is highly responsive to research needs identified by the DoE EERE WETO. It will enable the innovations to advance U.S. wind systems, reduce the cost of electricity, and accelerate the deployment of wind power. We will use high-fidelity simulations with the Weather Research and Forecasting (WRF) model to:

Quantify hydroclimatic conditions associated with wind turbine blade leading edge erosion. WT blade failures and reduced aerodynamic efficiency due to erosion of the leading edge (LEE) is caused by collisions between the rapidly rotating blade and falling hydrometeors. The amount of kinetic energy transferred into the blade and the resulting material stress is dictated by the closing velocity between the hydrometeor and the blade (ƒ(wind speed)) and the hydroclimate (number, size and phase of hydrometeors). We will assess for the first time whether the Weather Research and Forecasting model can be used to predict highly-erosive events. Validation will be performed using in situ rain droplet and hail size distributions and NWS dual-polarization RADAR.

Coupled ocean-atmosphere-wave modeling for improved simulations of wind turbine wakes and dynamic loading. Most simulations of wind resources, wind turbine wakes, and wind turbine operating conditions offshore do not employ models that explicitly treat wind-wave-ocean coupling, even though waves and the ocean strongly interact with the overlying atmosphere and thus modify the wind profile. Additionally, wind-wave misalignment is a key source of structural loading. We will perform simulations with the Coupled-Ocean-Atmosphere- Wave-Sediment Transport (COAWST) Modeling System to improve quantification of wind speeds, turbulence conditions and wind turbine loading in the east coast offshore wind energy lease areas. Our results will inform optimal wind farm layouts, wind turbine selection and power production quantification along with associated uncertainties. Simulations will be led by the PI team but students from Cornell's M.Eng program will be embedded in each science use case to help in expanding the usage base of HPC.

Title: Optimization studies of LBNF neutrino beamline and hadron absorber complex

Principal Investigator: Igor Rakhno (Fermi National Accelerator Laboratory)

Co-investigators: Nikolai Mokhov (Fermi National Accelerator Laboratory),

ALCC Allocation:

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

Allocation(s): 450,000 node-hours on Cori

Research Summary:

The Deep Underground Neutrino Experiment and Long-Baseline Neutrino Facility (DUNE-LBNF project) are under development at Fermilab since early 2010s. At present, the work is being performed towards Critical Decision 2 (CD-2) by the U.S. Department of Energy (DOE). The project represents a convergence of a substantial fraction of the worldwide neutrino physics community around the opportunity provided by the large investment planned by the U.S. DOE. The primary scientific objectives of DUNE are to carry out a comprehensive investigation of neutrino oscillations to test charge and parity (CP) violation in the lepton sector, determine the ordering of the neutrino masses, and to test the three-neutrino paradigm (electron, muon and tau neutrino). Independent measurements of the propagation of neutrinos and antineutrinos through matter will make it possible to observe neutrino transitions with the precision required to determine the CP-violating phase and the neutrino mass hierarchy. The LBNF will provide neutrino fluxes and detector infrastructure at the near site (Fermilab) and far site (Sanford Underground Research Facility) in South Dakota.

Many of the LBNF project milestones are heavily dependent on simulations with MARS code: Beamline sessions of DOE Independent Project Review in January 2021, CD-2/3b in Q2_2022, Beamline Preliminary Design and CD3 in Q3_2022, and sub-systems Preliminary Design Reviews in 2021-2023.

The LBNF will provide a 120-GeV proton beam on a neutrino production target utilizing a new superconducting Linac which is expected to be completed in 2027. The neutrino beamline is a core component of the LBNF. It consists of a target and horn systems, decay pipe, hadron absorber and so on. As a result of numerous iterations, there exists an optimized beamline design with a 1.5-m graphite target and focusing system consisting of three horns. The design inherits experience from previous neutrino projects. Various energy deposition and radiological calculations have been performed for major components of the beamline. All such simulation studies are done by means of precise Monte Carlo modeling of radiation transport and interactions with matter in a broad energy region, utilizing the power and capabilities of the Fermilab's MARS code. Significant follow-up work will be required that should address many components and issues such as current budget constraints. In particular, realistic beam accident scenarios will be studied. Also, modeling of LBNF operation in an anti-neutrino mode with a focus on radiation-affected components is planned to be performed. Performance of the MARS code itself is planned to be improved. Performance of ray tracing on large spatial three-dimensional meshes will be addressed using contemporary tracking methods. Availability of substantial supercomputer resources is an essential pre-requisite for success of all the studies described above.

Title: Short Range Correlations from a Quantum Monte Carlo perspective

Principal Investigator: Noemi Rocco (Fermi National Accelerator Laboratory)

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

ALCC Allocation:

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

Allocation(s): 730,000 node-hours on Theta

Research Summary:

The experimental investigation of nuclear dynamics at short range has flourished over the last few years. Electron-scattering experiments conducted at Jefferson Lab and other facilities worldwide enable selecting kinematics where the role of short-range correlations in the target nucleus becomes dominant. In addition, atomic nuclei are used in current and planned neutrino-oscillation and double-beta decay experiments, which will measure neutrino properties with unprecedented accuracy. Similarly to the electron case, these experiments are also sensitive to short-range correlations.

In this proposal, we will pursue a multi-faceted strategy aimed at modeling short- and long range dynamics of nuclei providing reliable estimates of the associated theoretical uncertainty. On the one hand, we will improve the determination of the low energy constants describing the short-range component of the nucleon-nucleon interactions, assessing their (correlated) uncertainties. On the other hand, we will implement this class of interactions in the numerical quantum Monte Carlo methods used by our group to compute the binding energies and radii of selected light nuclei. Finally, novel computational methods will be developed to obtain the β-decay rates of A = 11 and A = 12 nuclei and study the response densities of ^{12}C and ^{16}O induced by electrons and neutrinos scattering and gain more exclusive information about these processes. These studies will greatly impact ongoing experiments that uses electron to study in medium nucleon-correlations carried out at Jefferson Lab, and neutrino experimental facilities, e.g., the Deep Underground Neutrino Experiment (DUNE).

Title: High-Fidelity CFD Simulations for Next Generation Nuclear Reactor Designs

Principal Investigator: Dillon Shaver (Argonne National Lab)

Co-investigators: Yuan Haomin (ANL),
Aleksandr Obabko (ANL),
Landon Brockmeyer (ANL),

ALCC Allocation:

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

Allocation(s): 400,000 node-hours on Summit,
400,000 node-hours on Theta,
100,000 node-hours on Polaris

Research Summary:

In this project we plan to use high-fidelity computational fluid dynamics (CFD) and multiphysics simulations to investigate fundamental flow phenomena in next generation nuclear power reactors. Efforts will focus on key areas relevant to the DOE's NEAMS program and will support recent awards by the Advanced Reactor Demonstration Program (ARDP), the Technology Commercialization Fund (TCF) and support for a safety benchmark validation in conjunction with the U.S. Nuclear Regulatory Commission (NRC). All simulations will use the spectral element codes Nek5000 and its GPU variant, NekRS. The spectral element method combines the speed and accuracy of spectral methods with the flexibility of the finite element method

Two proposed designs were awarded by the ARDP to build demonstration reactors. These are X-energy's Xe-100 and TerraPower's Natrium reactors. These designs use pebble bed and wire-wrapped fuel pin cores respectively. Traditionally, industry uses low fidelity models for quick turn-around. However, these next-generation concepts exhibit complex flow and heat transfer behaviors, requiring a new generation of design tools to capture the effects with greater accuracy. All results will be used to enhance the available experimental data for validation of design tools. The simulations will be performed in collaboration with X-energy and TerraPower to directly support the ARDP awards, with the potential to accelerate the deployment of carbon-free energy solutions.

In collaboration with Framatome, we will be developing a flow-induced vibration capability that can provide a new tool for industry to accelerate the design cycle by estimating vibrations in fuel bundles. Additionally, as part of an ongoing close collaboration with the NRC in the areas of verification, validation, and uncertainty quantification, there is an opportunity to advance the state-of-the-art in modeling and simulations for post-Fukushima safety concerns. We to perform first-of-a-kind multiscale model validation of anisotropic turbulent gaseous mixing, with a particular focus on the PANDA/HYMERES-2 safety benchmark. In a unique collaboration as part of this proposal, the NRC will directly leverage leadership class resources, broadening access to these facilities.

Title: Proteome-scale structural and function prediction with deep learning

Principal Investigator: Jeffrey Skolnick (Georgia Tech Research Corporation)

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

ALCC Allocation:

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

Allocation(s): 130,000 node-hours on Summit
50,000 node-hours on Perlmutter-CPU
20,000 node-hours on Perlmutter-GPU

Research Summary:

The ability to predict the function of a protein-coding gene from its sequence is a grand challenge in biology. Our team has been working on uncovering the functions of protein-coding genes via structural inference using advanced high-performance computing (HPC)-based deep learning methods and multiscale simulations. Advances in gene sequencing technology have boosted exponential growth in the number of protein-coding gene sequences, leading to a new bottleneck in genome annotation. The ability to accurately infer an experimentally testable hypothesis for a majority of putative proteins in this massive dataset of sequences would revolutionize predictive and synthetic biology. Our team is developing and applying novel protein structure based deep learning methods to advance proteome annotation on the genome scale on the Summit supercomputer.

In the next stage of our project, we will continue to develop and deploy our novel inference methods for proteins across genomes of interest to DOE. Furthermore, we will tackle the data challenge generated by these large-scale inferences by developing efficient, scalable analytics to find connections and correlations across datasets and exploit metabolomics data to help fill-in gaps in metabolic networks. Then, we will probe and share the wealth of information produced by these ground-breaking new methods made possible by deep learning and HPC that focus on function inference, structure modeling, and protein-protein interaction predictions. Our use of massively-parallel GPU-accelerated deep learning platforms dramatically improves the understanding of many biological molecular systems via predicting accurate structural models of the components and assemblies of the systems. These state-of-the-art computational resources enable us to address many fundamental biological questions in a more powerful and complete way than is possible using academic resources. Success of this project will be important to achieving one of DOE's Office of Biological and Environmental Science's primary missions: the translation of nature's genetic code into predictive models of biological function.

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

Principal Investigator: Micholas Dean Smith (University of Tennessee)

Co-investigators: Jeremy C. Smith (Oak Ridge National Laboratory),
Aris Tsaris (Oak Ridge National Laboratory),
Junqi Yin (Oak Ridge National Laboratory),
Allan Hicks (Oak Ridge National Laboratory)

ALCC Allocation:

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

Allocation(s): 300,000 node-hours on Perlmutter-GPU

Research Summary:

This ALCC application seeks to integrate two DOE User Facilities, the Oak Ridge Leadership Computing Facility and the Spallation Neutron Source, by performing the calculations in two DOE-funded projects. The first project addresses a major challenge in biology, both experimental and computational, which is characterizing the structural flexibility of biosystems. We address this challenge with HPC simulation methods to Adequately sample the conformational dynamics of flexible biosystems. Critically, the allocation will support the development of computational pipelines for neutron science in advance of the Spallation Neutron Source Second Target Station (STS). Development of high performance computing 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 to characterize the structural complexity of biological molecules relevant to DOE mission. 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.

The second project is related to the development of biorefineries, which convert whole plant biomass to fuels and chemicals. Their development is a pressing challenge as this would help address the need for renewable energy and reduce CO2 emissions. Here, we will apply state-of-the-art molecular dynamics simulation on the Perlmutter supercomputer and machine learning analysis of the simulations to obtain detailed knowledge about the fundamental molecular organization, interactions, and associations that occur during solvent-based pretreatments of lignocellulosic biomass and solvent-induced destabilization of bacterial membranes. Gaining a predictive understanding of these solvent effects will provide a road-map for optimizing industrial biomass conversion.

Title: Unlocking wind farm dynamics to secure a sustainable energy future

Principal Investigator: Michael Sprague (National Renewable Energy Laboratory)

Co-investigators: Michael J. Brazell (National Renewable Energy Laboratory), Lawrence C. Cheung (Sandia National Laboratory), Georgios Deskos (National Renewable Energy Laboratory), Paul Mullowney (National Renewable Energy Laboratory), Philip Sakievich (Sandia National Laboratory), Ganesh Vijayakumar (National Renewable Energy Laboratory)

ALCC Allocation:

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

Allocation(s): 800,000 node-hours on Summit

Research Summary:

The key to achieving low-cost, wide-scale deployment of wind energy is enabling a new understanding of, and ability to predict, the fundamental flow physics and coupled structural dynamics governing whole wind plant performance, including wake formation, turbine-turbine interactions through wakes, and the ocean-atmosphere environment. High-fidelity modeling (HFM) and high-performance computing (HPC) offer a path to drive significant reductions in the cost of wind energy by providing researchers and engineers with a virtual environment for understanding the physics driving wind plant performance and to explore technology innovations and new operational strategies with confidence.

This ALCC allocation will directly support research conducted as part of two closely coordinated Department of Energy (DOE) funded projects: the Exascale Predictive Wind Plant Flow Physics Modeling project, funded through the Exascale Computing Project (ECP), and the High-Fidelity Modeling project, funded through the Wind Energy Technologies Office (WETO). These projects share a common open-source software stack called ExaWind, which is composed of computational fluid dynamics and structural dynamics codes. Simulations performed under this ALCC allocation will enable scientific discovery in three areas, which are all part of the grand challenges in wind energy: 1. Improving the wind community's understanding of the science behind wind farm blockage effects and their impact on wind farm performance and economics; 2. Understanding the dynamics of turbine-wake formation/evolution and coupled aero-structural response under multiple turbulent atmospheric conditions; 3. Gaining novel insights into the dynamics of air-sea interactions and potential impacts on wind farm performance. Phenomena will be studied using the highest-fidelity simulations performed to-date that will significantly advance the state of the art in computational modeling of wind farm flows.

Title: Simulating Collapsar Accretion Disks, Outflows, and Nucleosynthesis

Principal Investigator: Alexander Tchekhovskoy (Northwestern University)

Co-investigators: Francois Foucart (U New Hampshire),
Dimitrios Giannios (Purdue),
Ore Gottlieb (Northwestern),
Daniel Kasen (Lawrence Berkeley Laboratory),
Matthew Liska (Havard),
Philipp Moesta (U Amsterdam),

ALCC Allocation:

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

Allocation(s): 400,000 node-hours on Summit,
40,000 node-hours on Perlmutter-GPU

Research Summary:

What is the origin of heavy r-process nucleosynthesis? Two cosmic cataclysms are the likely culprits. One is the neutron star mergers: the detection of their radioactive glow, or the kilonova, has established them as an important r-process factory. However, the other culprit remains poorly explored: dying massive stars, or collapsars, have similar conditions in their central engines to neutron star mergers and hence can generate heavy r-process elements. But can these elements escape out of the collapsing star? INCITE-2021 cycle enabled extremely high-resolution 3D general relativistic collapsar simulations that revealed a variety of transients – from spherical accretion shock to relativistic gamma-ray burst jets – naturally emerging from different structures of progenitor stars. These outflows shock the stellar envelope and may carry the r-process elements out of the star and enrich the universe. However, these simulations lacked nuclear synthesis calculations to probe the r-process elements and featured only a single model whose outflows escaped out of the star.

The planned research will address the production of r-process elements in different types of progenitors and outflows, and their ability to carry the heavy elements outside of the star. Additionally, we will calculate radiation estimates across the entire electromagnetic spectrum for the simulations of each of the outflows, for the first time (i) modeling the system in 3D at sufficiently high resolution to resolve the resulting jetted and broad outflows over groundbreaking 6 orders of magnitude in space and time; (ii) achieving unprecedentedly long durations of ~20 seconds needed for the outflows to escape from the star and reach the homologous expansion stage to infer the late evolution and emission; and (iii) including on-the-fly neutrino transport for computing the ejecta nuclear composition. The planned simulations will further define the explosive conditions under which heavy element nucleosynthesis may take place and will identify nuclear pathways of most relevance to the r-process. Calculations such as these are essential in helping experiments like the Facility for Rare Isotope Beams meet their science goal of explaining how the heavy elements come into being.

Title: Simulation of flow and transport in desalination systems

Principal Investigator: David Trebotich (Lawrence Berkeley National Laboratory)

Co-investigators: Sergi Molins (Lawrence Berkeley National Laboratory)

ALCC Allocation:

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

Allocation(s): 92,160 node-hours on Perlmutter-GPU

Research Summary:

Coupled flow, transport, and reaction processes lead to membrane fouling, seriously limiting the performance of reverse osmosis, nanofiltration, ultrafiltration, microfiltration and membrane distillation systems. Predictive models linking meso-scale reactive transport phenomena that occur on and within membranes to system-level desalination performance are needed to optimize brine treatment for beneficial reuse. Furthermore, three dimensional models of device scale systems with sufficient resolution of surface reactions have not been possible with current simulation capabilities.

We plan to simulate resolved flow, transport and reactions in 3D permeable membrane configurations for the design and optimization of desalination devices. The device domain length scale is on the order of 10 cm (8.16cm x 4.08cm x .06375cm). Resolution of 5.0 microns has been shown in previous proof-of-concept simulations to be sufficient to resolve viscous boundary layers of the flow and to capture concentration polarization of the reacting components (16,384 x 8192 x 128 cells). These simulations have been performed on Cori KNL using the full machine (8192 nodes, 524,288 cores); for 32^3 box size the load balancing and domain decomposition sweet spot is 1 box per core. The Reynolds number for the flow is in the transition to turbulence regime (Re=~600). On Perlmutter this same problem would require more nodes than currently available in the Phase 1 deployment if we applied the same load balancing of 1 box per core. However, we have shown in scaling studies on Perlmutter that we can fit 4 boxes per core with our optimized, GPU-enabled code on the Phase 1 architecture. This provides the load balancing we need for our ALCC allocation request with 4 GPUs, 1 MPI rank per GPU and 32 cores per MPI. We, therefore, request 92,160 node-hours on Perlmutter GPUs to perform 8 12-hr production runs on 960 Perlmutter nodes to model breakthrough of the reactive front. This allocation will allow us to evaluate permeable membrane configurations at device scale.

The underlying software development for this work is being performed under the Exascale Computing Project. Users for this particular application of Chombo-Crunch, besides those at LBNL, also include a faculty and student/post-doc team at San Diego State University through the Sustainable Horizons Institute (SHI) Sustainable Research Pathways (SRP) summer program at LBNL.

Title: A Climate Model Ensemble for Understanding Future Changes to Extreme Weather

Principal Investigator: Paul Ullrich (University of California, Davis)

Co-investigators: Sara Pryor (Cornell University),
Colin Zarzycki (Pennsylvania State University),
Stefan Rahimi-Esfarjani (Univeristy of California Los Angeles),
Melissa Bukovsky (National Center for Atmospheric Research),
Alan Rhoades (Lawrence Berkeley National Laboratory)

ALCC Allocation:

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

Allocation(s): 900,000 node-hours on Theta,
300,000 node-hours on Perlmutter-CPU

Research Summary:

An understanding of the effects of climate change on extreme weather and atmospheric hazards is essential to ascertain future socioeconomic and infrastructural impacts from these events. However, because extreme events are inherently rare and often localized, a large number of high-resolution simulations are required to ensure sufficient statistical fidelity when making statements regarding projected changes to these features. While large ensembles of climate model simulations have been invaluable in recent years for putting tight bounds on future changes in synoptic meteorology, the finest of these model simulations use a grid spacing of ~80km, which precludes their use for fine scale weather impacts from extreme features such as tropical cyclones or atmospheric rivers.

The awarded effort will produce the world's first high-resolution "medium" ensemble from a single global modeling system, using regional refinement in the Department of Energy's recently-released Energy Exascale Earth System Model (E3SM) version 2. Specifically, 10 simulations will be performed covering the period from 1950 through 2100 with a refinement region for E3SM that covers the conterminous United States (CONUS) at a grid spacing of approximately 22km. This grid spacing captures aspects of the large-scale structure, frequency and location of tropical cyclones, atmospheric rivers, wind storms, and winter storms, as well as quantify their changes in the future. These simulations will provide better estimates of mountain snowpack in the Western US, and will provide valuable feedback to the E3SM developers on model performance for extreme weather events. In addition, the most extreme 20 events will be simulated at high resolution using the Simple Cloud-Resolving E3SM Atmosphere Model (SCREAM) at 3.5km grid spacing and Weather Research and Forecasting (WRF) system at < 2 km grid spacing. This will enable us to investigate if phenomena generated by E3SM are as hazardous as historically-based events. Both the large ensemble simulations and downscaled simulations will further allow us to understand if historical extremes are close to the "worst case" for such extremes, or if there is risk for even more damaging extremes beyond the observational record.

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

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

Co-investigators: David Ceperley (UIUC),

ALCC Allocation:

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

Allocation(s): 277,000 node-hours on Summit,
83,000 node-hours on Perlmutter-CPU

Research Summary:
The variety of behavior in the world around us is almost all emergent from very simple ingredients—electrons and nuclei arranged in different ways, with motion described in exquisite accuracy by quantum mechanics. This statement applies to situations as diverse as the interior of Jupiter, which is made up of mostly hydrogen and helium atoms, to emerging electronic devices such as twisted bilayer graphene. At least in principle, the knowledge exists to use computer simulations to make a mapping between different arrangements of atoms to behavior; however, in practice it is too challenging to bridge the microscopic to the macroscopic directly, and multiphysics models are used. In these multiple length scale calculations, detailed calculations are done for small length scales, then coarse-grained for larger length scales. Ultimately the accuracy of coarse-grained models is constrained by the data they are based on. This project is focused on improving underlying data on which coarse-grained models are based, focused on twisted bilayer graphene and hydrogen at high pressure.

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

Title: Laser-plasma instability mitigation using broadband lasers

Principal Investigator: Han Wen (Laboratory for Laser Energetics, University of Rochester)

Co-investigators: Russell Follett (University of Rochester),
John Palastro (University of Rochester),

ALCC Allocation:

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

Allocation(s): 400,000 node-hours on Perlmutter-CPU

Research Summary:

Promising to deliver an almost endless supply of power with relatively low environmental impact and a nearly inexhaustible reserve of fuel, the realization of controlled fusion would have lasting impact both geopolitically and for the health of our planet. In the direct-drive inertial confinement fusion (ICF) approach, a major program within the DOE-NNSA portfolio, an ensemble of laser beams symmetrically illuminates a cryogenic target containing thermonuclear fuel. Achieving efficient coupling of laser energy to the fusion target is one of the most essential components of direct-drive ICF. This coupling, however, is often impeded by the excitation of laser-plasma instabilities (LPIs).

Present since the earliest days of ICF, LPIs continue to be a major limiting factor that prevents experiments at the OMEGA laser and the National Ignition Facility (NIF) from achieving their design objectives. Significant improvement on the ICF implosion performance can be made by mitigating the primary instabilities such as stimulated Raman scattering (SRS) and the two-plasmon decay (TPD), and reducing the hot electrons generated by these instabilities. One LPI mitigation strategy is to use broadband lasers as the ICF driver. There are different approaches to introduce bandwidth into a laser, which are categorized by various laser smoothing techniques such as smoothing by spectral dispersion (SSD), induced spatial incoherence (ISI), and the Fourth-generation Laser for Ultra-broadband eXperiments (FLUX) being developed at Laboratory for Laser Energetics (LLE).

Particle-in-cell (PIC) simulations, one of the most well-established kinetic techniques, offer a powerful tool for studying broadband-driven LPIs. Previous studies have been focused on LPIs driven by a time-varying single speckle or by multiple speckles statically distributed in space. In this project, we will conduct a series of 2D and 3D PIC simulations to model TPD and ignition-scale inflationary SRS driven collectively by multiple laser speckles in a broadband laser. These simulations will provide a deeper insight into the kinetic processes of LPIs and enable us to quantify the effectiveness of different laser smoothing techniques on LPI and hot electron mitigation.

Title: Automatic Histologic Diagnosis of Whole Slide Imaging at Scale

Principal Investigator: Hong-Jun Yoon (Oak Ridge National Laboratory)

Co-investigators: Sajal Dash (Oak Ridge National Laboratory),
Benjamín Hernández (Oak Ridge National Laboratory),
Aristeidis Tsaris (Oak Ridge National Laboratory)

ALCC Allocation:

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

Allocation(s): 300,000 node-hours on Summit

Research Summary:

Whole Slide Imaging (WSI), which refers to scanning conventional glass slides in pathology to produce digital slides, is the most recent cancer imaging modality. Applying machine learning (ML) models to the comprehension of WSI images is computationally challenging due to the gigapixel-level image resolution. We leverage HPC infrastructure to exploit optimal algorithms and training protocols for achieving AI/ML models for WSI classification.

The project is with National Cancer Institute and Emory University under the National Childhood Cancer Registry project. The collaboration utilizes the capabilities of HPC to support an advanced population-level cancer surveillance program and develop an integrated data-driven modeling framework.

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

Principal Investigator: Yiqi Yu (Argonne National laboratory)

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

ALCC Allocation:

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

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

Research Summary:

The Materials Management and Minimization (M3) Reactor Conversion Program of the National Nuclear Security Administration (NNSA) is supporting the conversion of the research reactor from Highly Enriched Uranium (HEU, 235U / U ≥ wt. 20%) fuel to Low Enriched Uranium (LEU, 235U / U < wt. 20%) fuel. Better understanding of flow behavior and heat transfer mechanisms in these coolant channels is of great interest and importance for the design of LEU fuel elements. As is well known, many of the most-widely used RANS approaches has exposed the limitations on the prediction of these turbulent flows. Due to the lack of relevant experimental data, it is difficult to quantify the uncertainty introduced by these turbulence models. By performing high fidelity simulation (LES&DNS), the behavior of turbulent flows involves fewer modeling assumptions and can thus potentially provide more reliable means for engineering predictions and thus provide more accurate thermal hydraulic safety analysis. Moreover, the database from the high fidelity simulation can be used to exam or improve the coefficients used in these system codes, such as RELAP5 and PLTEMP/ANL, which have been applied for reactor fuel conversion for a long time.

The proposed project consists of two parts. The first one focuses on using LES to reproduce the flow behavior and collect the turbulent statistics in an involute coolant channel for high Reynolds number. The Reynolds number in the research reactor can range from 50,000 to 100,000. Performing LES simulation with higher Reynolds number is of high importance to investigate the impact of the Reynolds number on the flow and heat transfer mechanism in involute plate research reactor. The second portion of the work focuses on using Large Eddy Simulation to provide benchmark for the highly simplified involute coolant channel. Hundreds of RANS simulations with all kinds of turbulence models have been performed with different commercial codes, such as STAR-CCM+, COMSOL, ANSYS-CFX under the cooperation of three institutions (ANL, ILL, TUM). Discrepancy are found between RANS simulations with different turbulence model, flow condition and codes. The LES results will be used for benchmarking and further investigating the modification of the RANS model. Since the geometry is highly simplified, more LES data are expected for understanding the fundamental mechanism of the flow and heat transfer behavior in an involute plate research reactor.