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

Principal Investigator: Fikret Aydin (Lawrence Livermore National Laboratory)

Co-investigators: Helgi Ingólfsson (Lawrence Livermore National Laboratory),

ALCC Allocation:

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

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

Research Summary:

Approximately one third of all human cancers are caused by mutations in RAS proteins, particularly in lung, colon, and pancreatic cancers, which have poor prognosis and high mortality. Despite advances in experimental techniques and structural biology, the underlying mechanisms of RAS-driven cancers remain unclear due to a lack of sufficient resolution. Bridging experiments at the finest scale possible with simulations at the largest/longest scales attainable represents a holy grail for computational biology. To create this bridge and understand the molecular mechanism of the growth signaling pathway, particularly the interactions between RAS and its downstream effector, RAF, and the key steps in the initiation of the signaling cascade, large multiscale simulations are necessary.

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. By combining experiments with simulations, a vast amount of data can be generated and analyzed with the help of high-performance analytics. This will give us a comprehensive, predictive view of RAS-RAF activation, with both structural and dynamic insights that may lead to development of new therapeutic interventions.

Our data-driven approach that interleaves multiscale simulations using machine learning 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 machine learning 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: Hadronic contributions to the muon g-2 from lattice QCD

Principal Investigator: Thomas Blum (University of Connecticut)

Co-investigators: Alexy Bazavov (Michigan State University)
Peter Boyle (Brookhaven National Laboratory),
Carleton DeTar (University of Utah),
Aida El-Khadra (University of Illinois UC),
Steven Gottlieb (Indiana University),
Taku Izubuchi (Brookhaven National Laboratory)
Luchang Jin (University of Connecticut),
Christoph Lehner (University of Regensburg)
Ethan Nei (University of Colorado)
Ruth Van de Water (Fermi National Accelerator Laboratory)

ALCC Allocation:

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

Allocation(s): 3,283,000 node-hours on Frontier
5,000 node-hours on Polaris

Research Summary:

The muon is an elementary particle identical to the ordinary electron except that it is about 200 times heavier. Its magnetic dipole moment is being measured at Fermilab and calculated by theorists world-wide to fantastic accuracy in a high-stakes test of the Standard Model (SM) of Particle Physics. To test the SM to a degree that allows discovery of physics beyond our current understanding of Nature's laws, the theory errors on the contributions to the magnetic moment from the cloud of virtual quarks, anti-quarks, and gluons surrounding the muon during its brief lifetime must be reduced. These so-called hadronic contributions will be determined with improved precision in numerical simulations of Quantum Chromodynamics (QCD) known as lattice QCD.

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

Title: Two-phase flow interface capturing simulations

Principal Investigator: Igor Bolotnov (North Carolina State University)

Co-investigators: Anna Iskhakova (North Carolina State University)
Nam Dinh (North Carolina State University)

ALCC Allocation:

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

Allocation(s): 200,000 node-hours on Theta,
300,000 node-hours on Perlmutter-CPU

Research Summary:

The planned ALCC project will capitalize on existing multiphase flow experience to utilize interface capturing methods and direct numerical simulation to perform state of the art large scale simulations two-phase flows. PHASTA code has a long history of HPC performance, and our group has been awarded 2014, 2016, 2018 and 2020 ALCC awards in the past advancing the mission of DOE. PHASTA is a finite-element based flow solver with level-set method for interface capturing approach.

Three major research projects would be supported by this allocation: (i) project on simulation of the two-phase flow separator design to support advanced boiling water nuclear reactor technology; (ii) two-phase boiling analysis for fundamental understanding of flow boiling in complex geometries to support advanced heat-exchanger designs; (iii) counter-current two-phase flow evaluation in complex geometries to support efficient carbon-capture technologies.

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

Title: Hydrogen for Carbon-Free Low Emissions Power Generation

Principal Investigator: Jacqueline Chen (Sandia National Laboratories)

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

ALCC Allocation:

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

Allocation(s): 1,150,000 node-hours on Frontier

Research Summary:

The use of hydrogen to replace fossil fuel is one of the most promising routes to de-carbonize power generation, transportation and industrial sectors, thus mitigating climate change. While progress has been made towards the use of hydrogen and hydrogen blends in gas turbine engines, challenges remain for their widespread adoption in large-scale power generation related to hydrogen production, storage, transportation and conversion. The proposed simulations will address challenges faced in the conversion of hydrogen in gas turbine applications that will deliver low cost, clean and carbon free electric power. Specifically, this work will focus on improving the fundamental understanding of the combustion behavior of hydrogen targeted at low-NOx operation with high performance, a critical need delineated in the DOE Hydrogen Program Plan1 and a goal of the Fossil Energy and Carbon Management Advanced Turbines Program to deliver carbon free electricity by 2035 and net-zero emissions by 2050. The advantages of hydrogen fueled gas turbines are that they provide carbon free dispatchable power offsetting non-dispatchable renewable power generation from wind and solar resources while utilizing existing infrastructure. While replacing natural gas with hydrogen eliminates carbon dioxide emissions, high combustion temperatures typically lead to nitrogen oxide (NO) which needs to be avoided due to adversarial health and greenhouse effects. The proposed study will address the issue of burning rate and nitric oxide emissions scaling with pressure and turbulence intensity by performing a series of direct numerical simulations in a turbulent shear layer configuration. It builds upon previous DNS work with ammonia/hydrogen/nitrogen mixtures in shear-driven turbulence which have demonstrated enhanced thermo-diffusive instabilities at higher pressures. The unique DNS database will be shared with the broader reacting flow modeling community in academia and industry to develop predictive models for the design, development and optimization of efficient low-NO_x hydrogen gas turbines.

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

Principal Investigator: William Collins (Lawrence Berkeley National Laboratory and University of California, Berkeley)

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

ALCC Allocation:

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

Allocation(s): 219,000 on Perlmutter-GPU

Research Summary:

Simulating Earth's climate at high resolution with high fidelity is a computational grand challenge – it not only requires tremendous compute, but also produces humongous volumes of data. Studying low-likelihood high-impact extreme weather and climate events in a warming world requires massive ensembles to capture long tails of multi-variate distributions. In combination, it is simply impossible to generate massive ensembles, of say 1000 members, using traditional numerical simulations of climate models at high resolution.

This project intends to bring the power of machine learning (ML) to replace traditional numerical simulations for short week-long hindcasts of massive ensembles, where ML has proven to be successful in terms of accuracy and fidelity, at five orders-of-magnitude lower computational cost than numerical methods. Because the ensembles are reproducible to machine precision, ML also provides a data compression mechanism to avoid storing the data produced from massive ensembles.

The machine learning algorithm is based on Fourier Neural Operators (FNO) and Transformers, proven to be efficient and powerful in modeling a wide range of chaotic dynamical systems, including turbulent flows and atmospheric dynamics. The algorithmic complexity and mathematical properties have been studied in detail.

What makes this project *special* is that until today generating 1,000- or 10,000-member ensembles of hindcasts was simply impossible because of prohibitive compute and data storage costs. Scientists have known, however, that such massive ensembles are needed to characterize high-impact but low-likelihood events. For the first time, we can now generate such massive ensembles using ML at five orders-of-magnitude less compute than traditional numerical simulations. Furthermore, this project will push the frontiers of HPC+AI for climate modeling, especially for characterizing extremes under climate change.

Title: Investigations of proton tomography: Chiral-odd GPDs

Principal Investigator: Martha Constantinou (Temple University)

Co-investigators: Krzysztof Cichy (Adam Mickiewicz University,),
Andreas Metz (Temple University),
Swagato Mukherjee (Brookhaven National Lab),

ALCC Allocation:

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

Allocation(s): 390,000 node-hours on Perlmutter-CPU

Research Summary:

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

This project directly aligns with the DOE's mission for high-risk, high-reward research. Also, lattice QCD utilizes large-scale simulations, and acquiring access to ASCR resources will allow one to obtain state-of-the-art results and optimize our algorithms in preparation for the exascale supercomputing era. The latter is among the six pillars of the mission of the U.S. Department of Energy, aiming at "generating the greatest possible intellectual and economic benefit". To this end, the project will engage young researchers in forefront research and HPC, ensuring their further development and establishing US leadership in advanced computing research. The scope of the project aligns with the scientific goals of the 3-D Quark-Gluon Tomography (QGT) Topical Collaboration in Nuclear Theory funded by the U.S. Department of Energy.

Title: Relativistic Quantum Dynamics in the Non-Equilibrium Regime

Principal Investigator: Eugene DePrince (Florida State University)

Co-investigators: Xiaosong Li (University of Washington),
Ed Valeev (Virginia Tech); Chao Yang (Lawrence Berkeley National Lab),

ALCC Allocation:

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

Allocation(s): 700,000 node-hours on Theta

Research Summary:

The goal of this project is to apply correlated, relativistic many-body quantum chemistry methods to challenging problems in chemistry where spin-dependent effects cannot be ignored. In particular, relativistic equation-of-motion coupled-cluster (EOM-CC) approaches will be applied to Fe(II)-based spin-crossover compounds and Ni(II)-based single-molecule magnets (SMMs), which have potential applications as components in novel magnetic materials and in quantum information science. For the spin-crossover complexes, ligand-field splittings and spin-transition temperatures will be derived from electron attachment (EA) EOM-CC calculations applied to the (N-1)-electron state of the Fe(II) complex (wherein the metal center has an easy-to-describe $3d^5$ configuration). For the SMMs, magnetic anisotropy parameters will be derived from double ionization potential (DIP) EOM-CC theory applied to the (N+2)-electron state of Ni(II) (wherein the metal center has a similarly easy-to-describe $3d^{10}$ configuration). Relativistic CC with up to (perturbative) triple excitations will also be applied to to lanthanide oxide molecules, the electronic structure of which can provide insights into larger lanthanide-containing SMMs.

Broadly, this research aims to address a growing need in the computational chemical sciences for accurate first-principles descriptions of the relativistic quantum dynamics of many-electron systems. This need is driven by emerging quantum technologies that are increasingly important in material design, as scientists and engineers seek to manipulate spins toward a variety of goals, including novel magnetic materials and quantum information science (using spin-based molecular qubits, for example). Fundamental to these scientific and technological applications are the correlated many-electron dynamics of systems driven far from equilibrium, the accurate and efficient description of which represents a grand outstanding challenge in computational chemistry, especially when considering spin-dependent processes (e.g., spin-coherence, spin-entanglement, intersystem crossing, etc.)

Title: Energy partition and particle acceleration in laser-driven laboratory magnetized shocks

Principal Investigator: Frederico Fiuza (SLAC National Accelerator Laboratory)

Co-investigators: Alexis Marret (SLAC National Accelerator Laboratory),

ALCC Allocation:

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

Allocation(s): 300,000 node-hours on Theta,
150,000 on Perlmutter-CPU

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 goal of this ASCR Leadership Computing Challenge proposal is to perform large-scale first-principles fully-kinetic simulations of magnetized collisionless shocks in HED laboratory plasmas that will be critical to support Discovery Science experiments at the National Ignition Facility that have been approved and are currently planned. The fundamental understanding of particle acceleration in plasmas 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 proposed simulations and experimental programs on NIF will also enable the important benchmark of widely used numerical plasma models in magnetized HED conditions of relevance to DOE programs.

Title: Earlier and cheaper fusion energy with stabilized magnetic mirrors

Principal Investigator: Manaure Francisquez (Princeton Plasma Physics Laboratory)

Co-investigators: James Juno (Princeton Plasma Physics Laboratory),
Cary Forest (University of Wisconsin – Madison),
Ian Able (University of Maryland - College Park),
Liang Wang (Princeton Plasma Physics Laboratory),
Maxwell Rosen (Princeton University)

ALCC Allocation:

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

Allocation(s): 93,000 node-hours on Perlmutter GPU

Research Summary:

The U.S. is exploring a novel approach to fusion energy using simpler and potentially cheaper, smaller, and faster-to-deploy machines called mirrors consisting of two coils aligned cylindrically. Mirrors were previously thought to be unstable, yet new ideas, data and technologies have found and demonstrated the ability to stably confine a plasma with a mirror. Some of these ideas are being explored in two new mirrors, one at the University of Maryland-College Park and another at the University of Wisconsin-Madison. The Maryland mirror aims to demonstrate stable operation by making the plasma spin at supersonic speeds, so that centrifugal forces keep the plasma confined and suppress instabilities. Wisconsin's mirror will leverage (subsonic) rotation to suppress instabilities as well, while also employing novel high temperature superconducting magnets that produce unprecedented magnetic fields within the mirror. This research project will perform computer simulations to characterize to what extent these strategies effectively stabilize the Maryland and Wisconsin mirrors.

This research has a twofold impact. First, we will provide the Maryland and Wisconsin teams with a prediction of whether their plasma will be stable for a given set of conditions (e.g. temperature) and parameters related to their mitigation strategy (e.g. rotation velocity). These simulations would then inform how to change those conditions and parameters to find more stable arrangements while still meeting their scientific and technological goals. Such simulations can designate the parameter space in which the machine can be safely operated and could suggest hardware alterations to improve stability and performance. Second, our high-fidelity simulations will reduce the uncertainty associated with future mirror-based machines aiming to attain net positive fusion energy production. This work intends to help today's new mirrors succeed at demonstrating stability in new fusion-relevant regimes, and to develop greater confidence in our predictions of future mirror-based fusion power plants. These outcomes are critical to the Advance Research Project's Agency-Energy (ARPA-E) that funded such mirrors, and indeed to the Department of Energy (DOE) as a whole, given their vision to shorten the timescale and cost of commercial fusion energy.

Title: Large scale simulations of materials for quantum information science

Principal Investigator: Giulia Galli (University of Chicago)

Co-investigators: Francois Gygi (University of California Davis),
Marco Govoni (Argonne National Laboratory),

ALCC Allocation:

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

Allocation(s): 400,000 node-hours on Theta,
200,000 node-hours on Polaris,
200,000 node-hours on Perlmutter-CPU,
200,000 node-hours on Perlmutter-GPU

Research Summary:

Our objective is to use quantum simulations to predict the properties of materials of interest for novel, optically addressable quantum platforms, including quantum sensors. We plan to simulate electronic excited state properties of heterogeneous materials, inclusive of defects and interfaces, using coupled first principles molecular dynamics and electronic structure methods beyond density functional theory (DFT), as implemented in the Qbox (http://qboxcode.org/) and WEST (http://west-code.org/) open-source codes. Both Qbox and WEST are optimized for and have been used on high performance DOE architectures. Qbox is also coupled with a suite of codes (SSAGES) for advanced sampling and WEST is coupled to both Qbox and Quantum Espresso (https://www.quantum-espresso.org/). The proposed simulations are planned within the Midwest Integrated Center for Computational Materials (MICCOM; http://miccom-center.org/), one of the computational materials science centers funded by DOE.

We will simulate point defects in wide band gap semiconductors for the realization of qubit and quantum sensors. The main deliverables are: (i) predictions of the structural and electronic properties of heterogeneous systems, to be compared with experiments in order to obtain an integrated mechanistic understanding of the interaction of defective, nanostructured and bulk materials; (ii) validated data for systems of interest for quantum technologies, computed on large-scale DOE platforms, which will serve as reference results.

Title: First-principles prediction of solute energetics at extended defects in Mg ternary alloys

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

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

ALCC Allocation:

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

Allocation(s): 210,000 node-hours on Frontier

Research Summary:

The project seeks to investigate solute energetics at dislocations and grain boundaries (GB) in Mg alloys using large-scale real-space density functional theory (DFT) calculations. Despite its high strength-to-weight ratio, Mg has low ductility due to insufficient density of dislocations required for <c> axis plastic deformation, and strong basal texture of standard wrought Mg alloys. This research effort builds on prior DFT studies of dislocation-solute and GB-solute energetics in binary Mg alloys by this team. The present effort extends the investigation to probe ductility enhancement in ternary Mg-X-Y (Ca-Zn) alloys, where increased non-basal (pyramidal and prismatic) slip activity and weakening of the basal texture is experimentally observed compared to the binary counterparts (Mg-Ca and Mg-Zn). Specifically, dislocation-solute interaction energetics in non-basal dislocation systems in Mg-Ca-Zn alloy will be investigated using explicit large-scale DFT calculations involving simulation sizes of up to 6,000 atoms. These energetics will be used as inputs to study key ductility enhancing/limiting mechanisms in these Mg ternary alloys. The second part of this work focuses on studying finite-temperature solute segregation free energy at random GBs in Mg, which involves computing solute-segregation and co-segregation free energies using large-scale DFT calculations on cell sizes of ~3,000 atoms. These ab-initio inputs will subsequently inform a free energy model to obtain equilibrium solute concentration at the GB, and to elucidate the importance of incorporating the vibrational entropic and co-segregation effects. The first-principles calculations of dislocation/GB-solute interaction energies planned in this study will require large system sizes containing up to 6000 atoms, which are inaccessible using plane-wave DFT codes. This effort will leverage the recent computational, algorithmic and high-performance computing advances in 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 e-.

Overall, the proposed first-principles informed modelling of solute segregation at dislocations and GBs will enable quantitative guidance on solute combinations, concentration ranges and thermomechanical processing regimes that simultaneously enhance both intrinsic ductility and weaken basal texture in ternary Mg alloys. 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), Zhenbang Dai (The University of Texas at Austin)

ALCC Allocation:

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

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

Research Summary:

The efficient use of renewable energy is an essential requirement to enable the sustainable development of society. To create and deploy more efficient energy systems, new materials with superior electrical, optical, and thermal properties are called for. In this broad context, atomic-scale computer simulations are acquiring an increasingly important role in accelerating materials design, discovery, and characterization in synergy with experimental efforts. This research program develops and utilizes advanced computational methods to examine diverse materials classes of immediate interest for solar photovoltaics.

In solar cells, an electric current is generated when the quanta of solar energy or photons are absorbed by a photoactive material, leading to the promotion of electrons to higher quantum energy levels. The mobile charge carriers generated in this process must be transported to the electrodes in order to obtain a photocurrent. This research will provide a detailed atomic-scale understanding of how charge is generated and transported across two classes of materials, halide perovskite and plasmonic ceramics. On the one hand, halide perovskites exhibit extraordinary light-to-electricity conversion efficiency, and emerged as potential additions to traditional silicon solar cells for achieving ultra-high performance photovoltaics. On the other hand, plasmonic ceramics have attracted significant interest in the area of photovoltaics due to enhanced light absorption by the surface plasmon resonance, which is a concentration of the electric field in a very narrow surface layer leading to increased photocurrent generation. This research focuses on a promising non-standard class of plasmonic materials, namely plasmonic ceramics such as transition metal carbides and nitrides, which exhibit excellent stability in harsh environments and high temperature conditions.

To investigate the optical and transport properties of halide perovskites and plasmonic ceramics, the team will employ the EPW code, an open-source software package for first-principles quantum-mechanical simulations of electron-phonon interactions and related temperature-dependent materials properties which is supported by the DOE Computational Materials Science program.

Title: Privacy enabled tumor classification for near real time population health analytics

Principal Investigator: Heidi Hanson (Oak Ridge National Laboratory)

Co-investigators: John Gounley (Oak Ridge National Laboratory), Hong-Jun Yoon (Oak Ridge National Laboratory)

ALCC Allocation:

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

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

Research Summary:

The National Cancer Institute (NCI) has partnered with the Department of Energy (DOE) on the MOSSAIC project with the aim of using artificial intelligence and high performance computing to improve cancer surveillance in the United States. The partnership couples the computational and methodological resources held by DOE with the data resources from NCI to accelerate and automate information extraction and classification across multiple modalities of cancer data. Computational approaches to precision health integrate and analyze voluminous amounts of health data with the intent to deploy adaptive tools that have real-world impacts at the population and individual levels. At the population level, precision public health aims to identify changes in disease incidence across sub-populations that can be used to further develop focused approaches to disease prevention and intervention at the individual level. At the individual level, precision medicine aims to improve response to treatment by developing patient-specific treatment plans tailored to each person's unique biologic and molecular traits.

The computational campaigns presented here will extend our current work in text and image classification to further precision public health and precision medicine. First, using OLCF's CITADEL environment, we will perform end-to-end privacy enabled training of large language models on clinical text to identify schemes which maximize data security in the federated learning environments while minimizing accuracy loss. Second, we will train deep learning models to identify local patterns of cellular organization in a tumor that can be used for abnormality detection and deep phenotyping. Our overarching goal is to create multi-modal computational approaches to tumor classification that are automated, non-biased, high-throughput, privacy-enabled, and capable of standardizing large amounts of historically messy electronic health data. Successful completion of the computational campaigns presented in this proposal will move us closer to the goal of creating widely deployable text and image classification systems that facilitate near real-time cancer reporting. The project holds potential for biopreparedness applications, precision public health, precision medicine, as well as all domains and applications utilizing sensitive unstructured text.

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

Principal Investigator: Dirk Hufnagel (Fermi National Accelerator Laboratory)

Co-investigators: Stephan Lammel (Fermi National Accelerator Laboratory),

ALCC Allocation:

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

Allocation(s): 20,000 node-hours on Aurora
50,000 node-hours on Polaris

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 ALCF Polaris/Aurora for the reconstruction of some of the Run3 data collected at the CMS detector at the LHC
- to develop and exercise better integration of HPC (LCF) resources into the CMS computing infrastructure.

Title: Informing Forensics Investigations of Nuclear Materials

Principal Investigator: Sara Isbill (Oak Ridge National Laboratory)

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

ALCC Allocation:

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

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

Research Summary:

Researchers at Oak Ridge National Laboratory actively support national and international efforts to prevent and deter the proliferation of nuclear weapons by developing tools and methods to detect nuclear materials and processes. This research requires the use of advanced experimental and computational resources to better understand process kinetics and environmental degradation of key fuel cycle materials.

Determining the effect of process conditions on the underlying crystal structure and reaction products of fuel cycle materials via atomistic computational modeling creates a direct connection between material chemical and physical changes to laboratory observations. The prediction of experimental observables, such as optical spectra, 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 ALCC project is to connect highly accurate density functional theory (DFT) and large-scale molecular dynamics (MD) determinations of reaction products and lattice dynamics of fuel cycle materials to ongoing experiments collecting vibrational spectra obtained from Raman, infrared, and neutron scattering. This project, a continuation of our 2021-2023 ALCC awards, will continue to bridge the gap between experimental observations, atomistic understanding of surface reactions, and the impact of defects to inform chemical characterization and forensic analysis.

Systems of interest to this research are reactive metals and actinide materials. Some systems and phenomena may be adequately modeled using classical MD simulations, but the high accuracy of DFT calculations may be necessary for reliable comparison to experimental observables, especially computations of lattice dynamics for comparison to optical vibrational spectra. Method selection between DFT and MD will balance efficiency and accuracy for experimental comparison. Additionally, machine learning may be used to provide DFT-level accuracy at the length and time scales of MD.

Advancements in material characterization benefit greatly from atomistic modeling to facilitate understanding of process and environmental effects on materials, and the use of highly scalable computational modeling, made possible through leadership-class computing resources, is a necessary complement to our 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 (Argonne National Laboratory),

ALCC Allocation:

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

Allocation(s): 540,000 node-hours on Theta,
170,000 node-hours on Polaris

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 proposed computation constitutes the simulation part of the current EERE VTO project, non-traditional electrolyte design from ionic liquid. This research is aimed at using large-scale, high-performance computing to assist discovery of novel battery electrolytes. The overall goal is to enable rational design of superior electrolytes for high voltage batteries. This study will focus on nontraditional electrolyte discovery from ionic liquids: a new entry to battery electrolytes. The 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: Continuum Limit Lattice Calculation of Direct CP-violation in Kaon Decays

Principal Investigator: Christopher Kelly (Computational Science Initiative, Brookhaven National Laboratory)

Co-investigators: Norman Christ (Physics Dept., Columbia University),
Masaaki Tomii (Physics Dept., University of Connecticut),
Amarjit Soni (Physics Dept., Brookhaven National Laboratory),
Tianle Wang (Computational Science Initiative, Brookhaven National Laboratory)

ALCC Allocation:

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

Allocation(s): 135,000 node-hours on Polaris

Research Summary:

CP-violation (CPV) is necessary to explain the asymmetry between the amount of matter and antimatter in the observable Universe. CPV is highly suppressed in the Standard Model of particle physics and the predicted amount appears insufficient to account for the measured amount of matter, hence a precise comparison between experiment and Standard Model theory is a promising path to uncovering new physics. In the late 1990's, impressive experiments at FNAL (DOE-funded) and CERN succeeded in measuring with 15% errors the quantity ϵ', which characterizes the elusive direct CPV in $K \to \pi\pi$ decay, a tiny, part-per-million effect that is highly sensitive to new physics. Unfortunately, a correspondingly precise theoretical calculation was not possible until recently due to large, low-energy non-perturbative effects intractable to traditional theory approaches. Lattice QCD is the only known method to compute these with controllable errors, by directly simulating the theory in a finite, discretized box using Monte Carlo techniques on supercomputers.

This research team, along with other members of the RBC & UKQCD collaborations, performed the first complete lattice calculation of the decay in 2015, and in 2020 published an improved result with significantly better control over the systematic errors. Both results agree with experiment within errors, but at present these are sizeable, O(39%), and completely dominated by systematic effects. Separate research is underway to develop strategies for addressing two of the largest systematic errors, those associated with isospin-breaking/electromagnetic effects, and the Wilson coefficients that encapsulate the high-energy weak-interaction physics. This research project aims to address the last of the three largest errors in the earlier work: the finite lattice spacing effects resulting from using a single, somewhat coarse lattice spacing. While these are estimated to be of O(15%), smaller than the other dominant errors, this estimate has substantial uncertainty and so it is vital to repeat the calculation with a second (finer) lattice spacing and perform a continuum extrapolation to reduce/remove this source of error. This requires significant computational resources only available through DOE leadership-class facilities.

Title: Learning Aerosol-Cloud Interactions Across Scales

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

Co-investigators: Balwinder Singh (Pacific Northwest National Laboratory), Meng Huang (Pacific Northwest National Laboratory)

ALCC Allocation:

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

Allocation(s): 384,000 node-hours on Perlmutter-CPU

Research Summary:

Aerosol-cloud interactions (ACI) is a major source of Earth system predictability due to the complexity of multiscale processes that are involved. To address this challenge, the Energy Exascale Earth System Model (E3SM) with Regionally Refined Meshes (RRMs) will be used to produce a simulation ensemble with the horizontal resolution of a few kilometers to improve the realism of ACI and account for the effects of aerosol-induced cloud adjustments typically neglected in Earth system models. A wide range of aerosol and cloud regimes will be included in the simulation ensemble to ensure that the parameter space is realistic and sufficiently large. The kilometer-scale E3SM-RRM simulations will be used to drive detailed aerosol and cloud microphysics models and to provide further information on aerosol mixing state, size distribution, as well as cloud droplet size distribution and microphysical processes. They will also be used to develop neural networks to emulate processes affecting ACI and unresolved adjustment and feedback processes. Results will be evaluated against in-situ and remotely sensed observational data to ensure the fidelity of the model simulations.

Developing scientifically robust representation of ACI in E3SM for running efficiently at kilometer- and climate scales on U.S. Department of Energy's (DOE's) high-performance computing systems is critical for providing accurate and actionable predictions of the integrated Earth system. This research directly contributes to the Earth System Model Development program area by developing and utilizing innovative and computationally advanced modeling capabilities to address a critical uncertainty source in Earth system modeling and leverages the DOE's capability in high-performance computing for generating and analyzing large model simulation datasets, Earth system modeling across scales, and machine learning techniques.

Title: Scaling Genomic Variant callers to Leadership-class systems: A collaboration between VA-MVP and DOE

Principal Investigator: Ravi Madduri (Argonne National Laboratory)

Co-investigators: Jennifer Huffman (Department of Veterans Affairs), Philip Tsao (Department of Veterans Affairs); Victoria Popic (Broad Institute)

ALCC Allocation:

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

Allocation(s): 210,000 node-hours on Aurora

Research Summary:

The human genome, which consists of 23 pairs of chromosomes containing over 3 billion nucleotide base pairs (A-T and C-G pairs) encodes the information for the formation of proteins that drive all life processes. Conditions like autism, hemophilia, schizophrenia, cardiovascular diseases, Huntington's disease, cancer and Alzheimer's, among many others, can be caused by malfunctioning proteins that are in turn related to the large-scale changes in the human genome referred to as "structural variations'' (SVs). SVs are difficult to identify due to their complex nature as well as some inherent limitations of genome sequencing and mapping techniques. This study will identify and quantify SVs from the 100,000 whole human genomes sequenced using Next Generation Sequencing (NGS) technique as part of the Million Veteran Program (MVP). NGS platforms create multiple copies of the 3 billion base pair long genome and split them into tens of millions of short pieces (referred to as "reads") for the ease of accurate identification of the bases. For this work, these randomly located reads will be rearranged for each of the genomes by "mapping" them to the human reference genome, and then identification and characterization of SVs will be done using AI-assisted and traditional methods. The Aurora supercomputer will be used to carry out these computationally intensive tasks as well as to generate synthetic genomes and train the AI-assisted models for improved accuracy.

The goals of this project align with the US Department of Energy's (DOE) mission to promote transformative growth in scientific research using supercomputers and AI. It builds upon the strategic partnership between the US Department of Veterans Affairs (VA) and DOE in leveraging DOE computing capabilities to enhance the health outcomes of veterans through the MVP-CHAMPION (Million Veteran Program Computational Health Analytics for Medical Precision to Improve Outcomes Now) project. By cataloging SVs from over hundreds of thousands of genomes available from various sequencing projects, this project will create one of the largest databases of SVs in the world. The inclusion of participants from diverse racial backgrounds further enhances the quality and representativeness of the data. Curated SVs from such a large and diverse population will enhance the statistical power of genome-wide association studies (GWAS), allowing for the identification of better correlation between large genomic variations and phenotypes that will greatly facilitate the early detection of variants responsible for various conditions, and allow quicker identification of therapeutic targets for precision medicine.

Title: E3SMv2-Seasonal-to-Multiyear Large Ensemble (SMYLE)

Principal Investigator: Gerald Meehl (NCAR)

Co-investigators: Stephen Yeager (NCAR),
Nan Rosenbloom (NCAR)

ALCC Allocation:

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

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

Research Summary:
There is growing societal demand for advance warning of near-term climate extremes that take place against the backdrop of ever-increasing anthropogenic climate change. The field of initialized climate prediction on timescales from subseasonal to decadal has grown to meet this demand. Multi-year Earth system prediction falls 'between the cracks' of traditional seasonal and decadal prediction efforts, offering great promise for new research. This project will use the Department of Energy (DOE) Energy Exascale Earth System Model, version 2.1 (E3SMv2.1) for exploration of Earth system predictability on multi-year timescales. The focus will be on evaluating model performance at predicting impactful climate anomalies that occurred in the past. To that end, we propose two 20 member large ensembles of 24-month E3SMv2.1 initialized hindcasts spanning start years 1970-2018 (Nov 1 and Feb 1 initializations). They will be examined to gain insight into model realism and potential for reliable predictions of future change that will offer a novel assessment of E3SM skill for Earth system predictability. An analogous large ensemble of multi-year initialized hindcasts produced using the Community Earth System Model (CESM2) has provided significant contributions to our understanding of climate predictability on multi-year timescales (Yeager et al., 2022).

This work will directly support the grand challenge of Earth system predictability, and result in a unique dataset created with E3SMv2, providing climate information with uncertainty estimates to support DOE's energy missions. The proposed simulations will facilitate a broad spectrum of multidisciplinary science by the E3SM community with excellent potential to deliver high-impact, actionable science. A key goal of this proposal is to lay the foundation for an extensible climate-scale prediction system using the E3SM model that can be built upon with successive model versions and capabilities.

Title: High-Fidelity Flow Data for Multiscale Bridging: Toward Industry Adoption

Principal Investigator: Elia Merzari (Pennsylvania State University)

Co-investigators: Igor Bolotnov (North Carolina State University),
Nam Dinh (North Carolina State University),
Paul Fischer (University of Illinois at Urbana-Champaign),
Misun Min (Argonne National Laboratory),
Milorad Dzodzo (Westinghouse)

ALCC Allocation:

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

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

Research Summary:
The US Department of Energy (DOE) established in 2020 a consortium focused on improving the modeling of thermal-fluids in advanced nuclear reactors. Over the course of the project, this consortium, in collaboration with universities, national laboratories, and industry stakeholders, aims to develop faster and more accurate models for complex physical phenomena in advanced reactors. The main challenges are the lack of available thermal-fluids data and inefficient utilization of existing experimental and simulation results. The consortium is addressing these challenges by bridging the gap between different scales of modeling and focusing on four key problems: heat transfer modeling, thermal striping of internals, mixing in large enclosures, and multiscale core modeling coupled to fuel performance. To support this effort, the consortium is creating a comprehensive database of direct numerical simulation and large eddy simulation results that leverage the leadership computing resources of DOE. The data generated is already contributing to improved modeling techniques and facilitating the transition of these methods from academia to industry. The final year of the project will involve simulating the entire primary system of a lead fast reactor. The anticipated outcomes include a better understanding of turbulent thermal mixing in complex geometries, advancements in multiscale simulations, and wider adoption of these techniques by the industry.

Title: Predicting performance degradation from rough surfaces using LES

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

Co-investigators: Suhas S. Jain (Center for Turbulence Research, Stanford University), Sanjeeb Bose (Cadence Design Systems)

ALCC Allocation:

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

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

Research Summary:

This project focuses on large-eddy simulations (LES) of flow over complex geometries with rough surfaces to assess the impact of roughness on aerodynamic performance degradation. Roughness can occur in systems, such as wind turbines, aircraft lifting surfaces, compressor blades, and turbine blades, through processes like ice accretion, erosive damage, molten particulate deposition, or due to limited precision in manufacturing. These critical components are subject to conditions where high Reynolds number boundary layers interact with surface roughness, which adversely affect the aerodynamic performance and operability of the devices. This study's outcomes will lead to more efficient and effective designs, enhancing the performance and safety of various systems that encounter rough surfaces. The three main objectives of this study include determining resolution requirements, evaluating roughness wall models, and creating a comprehensive database of wall-resolved roughness simulations in complex geometries. The study focuses on iced NACA23012 airfoils and the roughened NASA rotor 37 compressor rotor. Laser-scanned experimental data will be used for ice accretion roughness, while artificial roughness patterns will be applied to the compressor rotor. The project addresses the limitations of existing roughness modeling databases, especially in complex geometries. It also considers roughness parameters beyond height, such as slope and spacing, and aims to explore their influence on aerodynamic quantities.

By investigating the impact of roughness on aerodynamic performance, the study addresses broader applications relevant to the DOE. These applications include wind turbines, which play a crucial role in renewable energy generation. Understanding how roughness affects the aerodynamics of wind turbine blades can contribute to improving their efficiency and overall energy output. Furthermore, the study's exploration of roughness effects on turbomachinery components can have implications for energy systems that rely on turbines, such as power plants and jet engines. Enhanced understanding of the interaction between roughness and aerodynamic performance can lead to more efficient and reliable energy conversion processes. Additionally, the study's focus on realistic engineering systems and collaborations with teams at Boeing and NASA align with the DOE's emphasis on practical and impactful research. Therefore, by investigating the effects of roughness on aerodynamic degradation, the study contributes to the DOE's mission space by advancing knowledge and technology related to energy systems, renewable energy generation, and engineering design for improved performance and efficiency.

Title: Machine Learning-Enhanced Multiphase CFD for Carbon Capture Modeling

Principal Investigator: Jordan Musser (National Energy Technology Laboratory)

Co-investigators: Aytekin Gel (ALPEMI Consulting, LLC),
William Fullmer (National Energy Technology Laboratory)

ALCC Allocation:

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

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

Research Summary:

This project leverages machine learning (ML) to enhance simulation-based engineering models that analyze how solid particles move and interact with a carrier fluid. Scientists use these tools to evaluate designs and help troubleshoot advanced particle-based reactors; however, because of long simulation times, researchers commonly choose faster running, less accurate models to reduce the time to solution. MFIX-Exa, a newly developed software designed to efficiently run on modern GPU-accelerated supercomputers, will generate high-fidelity datasets for training ML derived surrogate models. The created ML models that characterize phenomena like the interaction force between particles and the fluid can be incorporated into the faster running simulation tools to improve accuracy. This effort supports the Department of Energy's mission by enabling scientists to rapidly evaluate novel gas-solid reactor designs for advanced CO2 capture technologies.

Title: High Energy Density Physics of Inertial Confinement Fusion Ablator Materials

Principal Investigator: Ivan Oleynik (University of South Florida)

Co-investigators: Mitchell Wood (Sandia National Laboratories),
Stan Moore (Sandia National Laboratories),
Rahulkumar Gayatri (NERSC),
Marius Millot (Lawrence Livermore National Laboratory),
Sally Tracy (Carnegie Institution of Washington)

ALCC Allocation:

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

Allocation(s): 1,500,000 node-hours on Frontier,
500,000 node-hours on Aurora

Research Summary:

The historic December 5, 2022 experiment at Lawrence Livermore National Lab's National Ignition Facility reached fusion energy ignition for the first time, thus paving the way to future clean inertial fusion energy (IFE). This project addresses one of the priority research opportunities in IFE – the development of alternate ablative target materials, which are urgently sought to achieve high gain by enhancing fuel compression in low-adiabat implosions. Amorphous carbon (a-C) is a novel ablator material for next-generation IFE capsules. The main objective is to gain insights into the high-energy-density physics of amorphous carbon targets under compression by IFE drivers. This ALCC project will perform machine learning molecular dynamics simulations at DOE exascale Frontier and Aurora supercomputers at experimental time and length scales to uncover complex response of a-C ablators under dynamic compression and guide experiments to observe predicted phenomena and validate our theoretical models.

This ALCC project will provide direct support to recently awarded project within DOE FES High Energy Density Laboratory Plasma (HEDLP) program (program manager – Kramer Akli) and LaserNetUS project, involving experiments at Omega EP laser facility at the Laboratory of Laser Energetics at the University of Rochester and European X-ray Free Electron Laser Facility. Our tightly integrated research, education, and collaborative activities, coupled with the use of the world's most powerful DOE supercomputers and cutting-edge experimental facilities, will allow us to achieve top-tier results while educating the next generation of scientists and engineers to advance national IFE program and lead the world in IFE research.

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

Principal Investigator: Jonathan Owens (GE Global Research)

Co-investigators: Bojun Feng (GE Global Research),
Jinbo Cao (GE Global Research),
Xiaolei Shi (GE Global Research)

ALCC Allocation:

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

Allocation(s): 200,000 node-hours on Perlmutter-GPU

Research Summary:

Metal-organic frameworks (MOFs) are porous, 3-dimensional materials, consisting of metal nodes connected via organic linkers. Owing to their large pore-size and high surface area, MOFs are highly tunable structures that are a leading candidate for carbon dioxide sorbents. The tunability of MOFs arises from altering the metal nodes and organic linkers, and it gives rise to a huge class of structures, of which over 90,000 have been synthesized and 500,000 have been theorized. However, not all MOFs have the same performance for CO_2 adsorption, and the presence of water or other atmospheric molecular species can enhance or impede this performance. The situation is further complicated by the fact that grafting small molecules within the pores (termed amine functionalization) can greatly enhance adsorption performance, meaning that not only do you have the vast landscape of potential MOFs, but each of these MOFs can be fitted with different amine species bonding at different structural sites. Taken together, this is a massive materials discovery and characterization challenge, impossible to solve with experiment alone. The objective of this allocation is to leverage Advanced Scientific Computing Resources (ASCR) to develop a multi-component computational pipeline aimed at discovery and screening of amine-functionalized MOFs, focused particularly on using high-accuracy, computationally expensive first principles methods. These methods will be employed to determine how much CO_2 a material can adsorb, how quickly this sorption process happens, and how robust these structures are to varying environmental conditions such as temperature and humidity.

Climate change is a major crisis of our time, and achieving a carbon negative future will require a multi-pronged approach, including electrification, clean energy generation, and carbon dioxide removal (CDR). CDR technologies are a key priority for the DOE and core component of the DOE's Carbon Negative Shot, the "Earthshot" related to CDR. This project will develop and perform high-accuracy computational materials screening methods that are vital to CO_2 sorbent development, a task central to effective CDR technologies.

Title: Probabilistic Comparative Modeling of Colorectal Cancer Screening Strategies

Principal Investigator: Jonathan Ozik (Argonne National Laboratory)

Co-investigators: Carolyn Rutter (Fred Hutchinson Cancer Center),
Iris Lansdorp-Vogelaar (Erasmus University Medical Center),
Karen Kuntz (University of Minnesota),
Fernando Alarid-Escudero (Stanford University)

ALCC Allocation:

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

Allocation(s): 150,000 node-hours on Theta,
10,000 node-hours on Polaris,
100,000 node-hours on Perlmutter-CPU

Research Summary:

Despite large increases in the uptake of screening in the past two decades, colorectal cancer (CRC) is still the second leading cause of cancer death in the US. This points to inadequate screening and treatment, 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 and describe CRC natural history using different underlying assumptions. Building on model calibration, comparison, and evaluation of screening efficacy that was accomplished in the 2022-2023 period, the project will extend analyses to extensions of the microsimulation models that incorporate the serrated pathway to colorectal cancer and examine the effectiveness of new blood-based screening approaches. 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.

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

Principal Investigator: Bruce Perry (National Renewable Energy Laboratory)

Co-investigators: Shashank Yellapantula (National Renewable Energy Laboratory), Lucas Esclapez (National Renewable Energy Laboratory), Sreejith Nadakkal Appukuttan (National Renewable Energy Laboratory), Marc Day (National Renewable Energy Laboratory)

ALCC Allocation:

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

Allocation(s): 320,000 node-hours on Frontier

Research Summary:

Liquid jet fuel from bio-derived sources, referred to as Sustainable Aviation Fuel (SAF), currently is the most promising avenue for decarbonizing the aviation sector. However, performance uncertainty and certification requirements represent major barriers to commercial adoption of SAF. Due to the high cost of testing, simulations must play a key role in de-risking SAF candidates. This project involves a series of high-fidelity simulations to characterize the performance and emissions of two SAFs, namely Alcohol-to-Jet (ATJ) fuel and Hydroprocessed Esters and Fatty Acids (HEFA), along with standard Jet-A fuel. These simulations will be based on an aircraft combustor architecture that is under development by GE Aviation in collaboration with Georgia Tech through a program to study supersonic commercial transport and 100% drop-in SAF use. The simulations will be performed with the Pele reacting flow solvers, developed under the Department of Energy's Exascale Computing Project (ECP). The project also involves tight collaboration with the fuels laboratory at the National Renewable Energy Laboratory, which is running experiments to characterize key fuel properties for the study. Three major outcomes are expected: validation of the software and modeling suite for drop-in SAF end use, comparative characterization of pollutant production and combustion of key fuel candidates, and assessment of the dominant properties affecting drop-in SAF performance.

Recognizing the importance for decarbonizing the aviation sector, a federal-government-wide grand challenge has been established to expand production and use of SAF. Production of SAF is also a key component of the Department of Energy's Clean Fuels & Products Energy Earthshot. To achieve these lofty goals, the end use of 100% drop-in SAF must be de-risked to enable certification of the new fuels and widespread commercial adoption. The simulations performed in this work will be a significant step toward achieving this goal, furthering scientific understanding of how fuel properties affect key combustion characteristics. Additionally, demonstrating and validating the open-source Pele codes for SAF simulations will pave the way for community use of the tool for SAF combustion explorations.

Title: Heavy quarks in QGP: lattice QCD inputs for sPHENIX and LHC

Principal Investigator: Peter Petreczky (Brookhaven National Laboratory)

Co-investigators: Swagato Mukherjee (Brookhaven National Laboratory)

ALCC Allocation:

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

Allocation(s): 943,000 node-hours on Frontier

Research Summary:

Heavy (charm and bottom) quarks and bound states of heavy quarks (quarkonia) are important tools in our understanding of Quark Gluon Plasma (QGP). Charm and bottom quark can help to understand the mechanism of parton energy loss in QGP and constrain the transport coefficients. Quarkonia probe collective properties of QGP at different length scales, in particular, the modification of inter-quark forces. Therefore, there is a large experimental effort focused on heavy quark probes in heavy-ion collisions. One of the goals of the new sPHENIX experiment at RHIC in BNL becoming operational in 2023, is the study of quarkonia and heavy-quark tagged jets. The STAR experiment at RHIC will be alsorunning during 2023-2025 with the study of quarkonium production being one of its focuses. Heavy-quark and quarkonium production is also an important part of the heavy-ion program at LHC experiments (ALICE, ATLAS,CMS and LHCb). In fact, one of the motivations for the upgrade of ALICE experiment was a better understanding of heavy-quark production. The theoretical understanding of the experimental results on heavy-quark and quarkonia production is limited because some QCD inputs for the dynamical models aiming to describe heavy-quark and quarkonium production are missing. The goal of this ALCC project is to perform lattice QCD calculations of important quantities urgently needed for the interpretation of the experimental results on heavy-quark and quarkonia production in heavy-ion collisions, including the heavy quark diffusion coefficient, the in-medium bottomonium masses and widths, and the heavy quark anti-quark potential at non-zero temperature. These key quantities will be calculated using lattices with large temporal extent together with novel calculational and analysis techniques, such as the gradient flow and the constrained spectral reconstruction of the lattice QCD correlation functions. The obtained results will largely reduce the uncertainties in the modeling of heavy quark and quarkonium production in heavy ion collisions allowing a much-improved interpretation of the experimental results.

Title: High-fidelity Simulations of Turbulence and SBLI in Transonic Fan Blades

Principal Investigator: Stephan Priebe (GE Research)

Co-investigators: Daniel Wilkin II (GE Aerospace),
Ramakrishnan Kannan (Oak Ridge National Lab)

ALCC Allocation:

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

Allocation(s): 290,000 node-hours on Frontier

Research Summary:

This project will execute high-fidelity simulations of the turbulent flow in transonic fan blades, which includes shock wave/boundary layer interaction (SBLI), possible flow separation, and turbulent wake mixing. The tip region of transonic blades is particularly challenging due to the presence of shock waves leading to SBLI, which are of high importance in aerospace applications, including turbomachinery aerodynamics since they can lead to significant aerodynamics losses and hence a reduction in efficiency and performance. SBLI is difficult to predict accurately using standard design tools due to shortcomings in the underlying models, which is why this ALCC project will apply machine-learning techniques to derive improved models trained on the high-fidelity simulation data.

There is a need to further improve the efficiency and fuel consumption associated with aircraft propulsion. The large fan at the front of modern turbofan engines is a key part to achieving these improvements, which are only possible with a comprehensive understanding of fan aerodynamics. In the realm of commercial aircraft propulsion alone, we estimate that the benefit of an improved physical understanding and prediction capability of transonic fans and SBLI would be of the order of 125 million gallons of jet fuel and 850,000 metric tons of greenhouse emissions for U.S. airlines every year.

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

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): 61,000 node-hours on Perlmutter-CPU

Research Summary:

The Deep Underground Neutrino Experiment and Long-Baseline Neutrino Facility (called DUNE-LBNF project) are under development at Fermilab since early 2010s. In 2016, the Critical Decision 1 (CD-1) was approved by the U.S. Department of Energy (DOE). The activity represents a convergence of a substantial fraction of the worldwide neutrino physics community around the opportunity provided by the large investment planned by the U.S. DOE. The primary scientific objective of DUNE is to carry out a comprehensive investigation of neutrino properties including conversion rules between neutrinos of different types. Also, a high-priority ancillary science program will be conducted. The LBNF will provide neutrino fluxes and detector infrastructure at the near site (Fermilab) and far site in South Dakota. Many of the LBNF project milestones are heavily dependent on thorough computer simulations including Beamline Final Design Review in Q4 2024 and sub-systems Final Design Reviews in 2023-2025. Significant work will be required that should address various components and issues such as current budget constraints. This work will be based on the existing optimized design with a 1.5-m graphite target and beam focusing system consisting of three horns. The Mu2e project at Fermilab – devoted to high precision measurements of muon-to-electron conversion without emission of neutrinos – is at its final stage of completion. The goal of the experiment is to identify new physics laws beyond the Standard Model. Since the FY2022, this project has been recognized as one of the top priorities for the lab's mission. The beam commissioning begins in FY2023, the detector is expected to be complete and commissioning is planned to begin in FY2026. The transition of the project to operations requires substantial computer modeling studies of the radiation protection of the personnel and equipment, which involves large scale simulations of prompt and residual radiation levels in the beam lines and experimental enclosures.

The goal of the Proton Improvement Plan Phase II (PIP-II) project is to build a new superconducting linear accelerator which is expected to provide beam to the DUNE-LBNF and Mu2e experiments. It requires development of the final integrated radiological analysis and design which implies detailed and comprehensive computer modeling studies. The latter will focus on multiple operational scenarios including six planned phases of commissioning. Credible beam accidents will be studied as well.

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

Title: Accelerating Deployment of Next-Generation Nuclear Power Using High-Fidelity CFD

Principal Investigator: Dillon Shaver (Argonne National Lab)

Co-investigators: Yuan Haomin (ANL), Aleksandr Obabko (ANL), Jun Fang (ANL)

Allocations site: Argonne Leadership Computing Facility (ALCF)

Allocation size: Aurora – 200k, Polaris – 100k, Theta – 200k node-hours

Research Summary:

Computational Fluid Dynamics (CFD) is playing increasingly important role in the design process of next-generation nuclear reactors. With the dawn of the Exascale computing era, full-core high-fidelity CFD simulations are now a reality. This offers unprecedented insight into complex flow and heat transfer behavior expected in many of the proposed designs. By leveraging high-fidelity simulations, we can generate more accurate models and implement them in the low-fidelity, fast-running design tools in use by the industry. This hi-2-lo concept is a driving force behind the DOE's ongoing Advanced Reactor Demonstration Program (ARDP) and is a key effort of the DOE's Nuclear Energy Advanced Modeling and Simulation (NEAMS) program. This project proposes to leverage data generated with high-fidelity CFD simulations to inform fast-running model development for design tools, aiding in the deployment of carbon-free energy on a commercial scale. It will consist of three primary components: ongoing support for the ARDP program, support for hi-2-lo efforts in NEAMS, and ongoing support for a safety validation benchmark with the U.S. Nuclear Regulatory Commission (NRC). All proposed simulations will use the spectral element code NekRS, which has a long history of development in the DOE NEAMS program and use on DOE leadership computing facilities.

ARDP demonstration projects are expected to yield a fully functional advanced nuclear reactor within 7 years. Simulations performed for both pebble beds and wire-wrapped fuel will enrich the available experimental data by evaluating the applicability of correlations – many of which were not directly intended for the design conditions. Simulations will further extend the knowledgebase by considering moderate flow rates in both designs, where natural circulation plays a significant role in heat transfer. This will fill an important gap that must be addressed.

To support validation of hi-2-lo efforts, a series of experiments have been recently conducted as part of an OECD/NEA international benchmark with an integral effect test facility for a High-Temperature Gas-cooled Reactor (HTGR). High-fidelity CFD simulations would enrich the available experimental data and provide valuable insights to the thermal fluid phenomena in a HTGR system.

Following the events of Fukushima Daiichi, A particular focus for nuclear reactor safety has been in hydrogen mitigation. In collaboration with the NRC, simulations will be performed of a large, complex stratification layer benchmark to provide validation for this phenomenon. This problem is a long-running transient and represents a major challenge in computing. The recent switch to NekRS on GPUs has yielded significant speed up, allowing for improvements in the scope of what can be simulated. This aspect of the project has the potential to enhance safety of currently operating reactors and continue broadening access to leadership computing.

Title: Integrative rate-kinetics modeling of diffusive limits to energy harvesting in bioenergetic membranes

Principal Investigator: Abhishek Singharoy (Arizona State University)

Co-investigators: David J. Hardy (University of Illinois at Urbana Champaign)

ALCC Allocation:

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

Allocation(s): 810,000 node-hours on Frontier,
50,000 node-hours on Perlmutter-GPU

Research Summary:

Photosynthesis is a fundamental process that sustains life on Earth by converting energy across extremely short timescales. This intricate process involves the coordination of thousands of proteins, posing a significant challenge for High-Performance Computing (HPC) to simulate efficient energy conversion. The diffusion of charge carriers within crowded protein domains determines the rate-limiting steps in photosynthesis and influences overall cellular performance. Our proposed approach involves molecular dynamics (MD) simulations of two bioenergetic systems, each associated with a Department of Energy (DOE) project. The first system focuses on the cyanobacterial iron deficiency-induced photosystem I (isiA-PSI) domain, examining the effects of external stress, such as iron deficiency, on the kinetics of charge carrier diffusion. The second system explores the supramolecular design principles of cellular energy metabolism at a systems level for the heliobacterial photosynthetic membrane. We employ a combination of coarse-grained (CG) MD simulations and all-atom simulations, guided by Reinforcement Learning, to study the long-time kinetics and electron hopping mechanisms within these systems. Additionally, we conduct a large-scale MD simulation to investigate protein function integration at the system level. Our bioenergetic simulations push the boundaries of HPC capabilities and provide valuable insights into optimizing bioenergy solutions. Moreover, we collaborate with experimental researchers at ASU to enhance the effectiveness of our approaches and contribute to the rational design of bioenergy systems.

Title: Grand-challenge predictive wind farm simulations

Principal Investigator: Michael Sprague (National Renewable Energy Laboratory)

Co-investigators: Jon Rood (National Renewable Energy Laboratory),
Philip Sakievich (Sandia National Laboratories),
Ashesh Sharma (National Renewable Energy Laboratory),
Ganesh Vijayakumar (National Renewable Energy Laboratory),
Lawrence Cheung (Sandia National Laboratories)

ALCC Allocation:

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

Allocation(s): 3,700,000 node-hours on Frontier

Research Summary:

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 and high-performance computing 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. The simulations performed under this ALCC allocation will use the Department of Energy (DOE) supported ExaWind suite of open-source codes and will address a grand challenge, which is the predictive simulation of a large wind farm comprised of modern megawatt-scale turbines with long and flexible blades. Such simulations will require the resolution of spatial scales spanning at least eight orders of magnitude, going from micrometer-scale boundary layers around turbine blades to the kilometer-scale wind farm domain. In addition to unlocking key insights into wind farm physics, we expect that the open-source ExaWind modeling and simulation environment, and the datasets generated from the ALCC allocation, will become the de facto standards for modeling best practices for simulating wind farm flows and will be used by the community to advance their engineering and operational tools. The simulations proposed here will be the crowning achievement of seven years of ExaWind development under the DOE Exascale Computing Project and the DOE Wind Energy Technologies Office.

Title: Quantum accurate large-scale atomistic simulations of advanced fusion reactor materials

Principal Investigator: Aidan Thompson (Sandia National Laboratories)

Co-investigators: Mary Alice Cusentino (Sandia National Laboratories),
Stan Moore (Sandia National Laboratories),
Mitchell Wood (Sandia National Laboratories),
Megan McCarthy (Sandia National Laboratories)

ALCC Allocation:

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

Allocation(s): 500,000 node-hours on Frontier,
600,000 node-hours on Aurora,
250,000 node-hours on Polaris,
250,000 node-hours on Perlmutter-GPU

Research Summary:

The goal of this project is to understand how plasma-facing metal alloys in fusion energy reactors degrade under bombardment by helium, hydrogen, and neutrons, by performing atomistic simulations of these materials at unprecedented scales and accuracy levels. The *FitSNAP* software will be used to develop *SNAP* high accuracy machine learning interatomic potentials describing implantation of hydrogen, helium, beryllium, and nitrogen. A similar approach will be used to describe neutron damage in tungsten-based refractory high entropy alloys. These simulations will help design new materials that can withstand the harsh fusion reactor environment. These novel simulations are achievable on DOE leadership class computing resources due to extensive *Kokkos*-based GPU optimization of the *SNAP* force kernel in the *LAMMPS* simulation code through the DOE Exascale Computing Project. The *FitSNAP*/*LAMMPS*/*Kokkos*/*SNAP* capability allows for quantum-accurate atomistic simulations of chemically complex materials at the plasma-material interface approaching realistic plasma environments.

Title: Simulation of flow and transport in desalination systems

Principal Investigator: David Trebotich (Lawrence Berkeley National Laboratory)

Co-investigators:

ALCC Allocation:

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

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

Research Summary:

In response to increasing water scarcity in drought-stricken areas, desalination has evolved into a viable alternative water supply. Semi-permeable seawater reverse osmosis membranes have become the dominant desalination technology outside of the Middle East. Besides engineering costs, there are performance costs associated with multiple physical processes involved in the technology. 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.

Chombo-Crunch is a production-capable application code that models high resolution flow, transport and reactions in arbitrarily complex geometries at unprecedented scale and resolution. Chombo-Crunch has successfully simulated multiphysics problems in a number of complex, heterogeneous geometries for applications including subsurface (pore scale reactive transport/BES EFRC), energy storage (lithium ion battery electrodes/DOE battery hub) and manufacturing (paper pressing/AMO HPC4Mfg, battery electrode slurry drying/EERE Roll-to-roll consortium). The Chombo-Crunch flow solver itself—based on the adaptive, finite volume Chombo software framework for PDE solutions—is robust for a wide range of flows, from low Reynolds number, Stokes flow to high Reynolds number turbulence. The code scales to full machine resources on OLCF Frontier as a result of recent performance portability and engineering under the ECP Subsurface application development project.

We propose 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). 5.0 micron resolution 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 under laminar flow conditions (16,384 x 8192 x 128 cells). The Reynolds number for the proposed device simulation is in the regime of transition to turbulence (Re=643). We have performed a preliminary turbulent flow study on the NERSC Perlmutter Nvidia A100 architecture using 4 boxes per MPI process under a previous ALCC award with our optimized, GPU-enabled code. With 9400 AMD nodes and 8 GPUs per node, Frontier provides capability for resolution of the turbulent flow boundary layer in the confines of a complex desalination device.

Title: Probing strong-field QED with Doppler-boosted lasers

Principal Investigator: Jean-Luc Vay (Lawrence Berkeley National Laboratory)

Co-investigators: Henri Vincenti (CEA Saclay, France),

ALCC Allocation:

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

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

Research Summary:

This project will use exascale-class simulations on Frontier to help design, guide and interpret novel laboratory experiments intended to probe Quantum ElectroDynamics (QED) effects occurring at very high electromagnetic field amplitudes, effects characterizing for instance the surroundings of extreme astrophysical objects like black holes or neutron stars. Simulations will be carried out with the massively-parallel, open-source, Particle-In-Cell (PIC) code WarpX (winner of the 2022 ACM Gordon Bell prize). This project will study a new strategy proposed by the team that consists in using optical devices called "plasma mirrors" to boost the intensity of a very short laser pulse by several orders of magnitude. Early theoretical and numerical estimates show that the new scheme could lead to electromagnetic field intensities high-enough to study QED effects in a strong field using existing PetaWatt-class lasers (e.g., BELLA, in the US or the ELI lasers in Europe).

QED is one of the most successful and well tested physical theories, but its strong-field regime is virtually unexplored in experiments. The "numerical experiments" that will be performed during this project will be key to guide experimental campaigns already planned at PetaWatt-class laser facilities in the USA and in France. This will be pivotal to validate decades-old theoretical predictions, potentially revealing new physics, and to improve our understanding of several extreme astrophysical scenarios.

Title: QMC-HAMM: High accuracy many-body data for multiscale models

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

Co-investigators: David Ceperley (University of Illinois at Urbana-Champaign)

ALCC Allocation:

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

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

Research Summary:

In this project, highly accuracy quantum Monte Carlo calculations will server as a base for two high-impact projects which will deepen our understanding of electronic materials. This project is part of the computational center, "QMC-HAMM: High accuracy multiscale models using quantum Monte Carlo."

This project is primarily focused on using first principles quantum Monte Carlo (QMC) calculations to provide highly accurate reference data on target materials. Quantum Monte Carlo offers explicit treatment of electron correlation for systems of electrons with up to around 1000 particles, and performs extremely well in parallel, taking advantage of the leadership-class resources available at NERSC. The project focuses on generating high quality first principles data and models, intended to be used in the development of larger length scale descriptions for materials.

This project links with efforts in computational materials science funded by the DOE such as the Midwest Integrated Center for Computational Materials at Argonne, the Materials Project at Lawrence Berkeley, Center for Computational Materials Science and Design at Brookhaven, and the Center for Predictive Simulation of Functional Materials at Oak Ridge. The first three of these centers focus on the development of effective and efficient theories for functional materials design, while the last one focuses on accurate but very microscopic physics. In this project, the very detailed and high accuracy QMC calculations will be linked to the effective and efficient theories.

In the current computational effort, two main applications will be considered, which are connected because the details at the microscopic affect the physics at the large scale. For the first project, simulations will be performed on liquid and solid hydrogen at pressures and temperatures similar to the interior of Jupiter. From these simulations, new understanding of the phase diagram of hydrogen will be attained. For the second project, detailed QMC calculations will determine how individual atomic layers of carbon, graphene, interact with each other to form corrugation, and how the electrons collectively interact to create new electronic phases of matter.

Title: Large Eddy Simulation on flow and heat transfer behavior in Involute Plate Research Reactor Supporting the Needs of the Materials Management and Minimization (M3) Reactor Conversion Program

Principal Investigator: Yiqi Yu (Argonne National laboratory)

Co-investigators: Cezary Bojanowski (Argonne National Laboratory), Aurelien Bergeron (Argonne National Laboratory), Jeremy Licht (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): 510,000 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. There are three research reactors in the world actively engaged in conversion that utilize involute shaped fuel elements: High Flux Isotope Reactor (HFIR) in USA; High Flux Reactor (RHF) in France and FRM II in Germany. These reactors share a similar configuration of coolant channel, which is of extremely thin thickness and involute shape. Better understanding of flow behavior and heat transfer mechanisms in these coolant channel 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.

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 assess the uncertainties of RANS models.

Title: 3D Transverse Spin Structure of the Proton: Probing Time-Reversal Symmetry

Principal Investigator: Yong Zhao (Argonne National Laboratory)

Co-investigators: Dennis Bollweg (Brookhaven National Laboratory),
Ian Cloët (Argonne National Laboratory),
Xiang Gao (Argonne National Laboratory),
Swagato Mukherjee (Brookhaven National Laboratory),
Qi Shi (Brookhaven National Laboratory)

ALCC Allocation:

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

Allocation(s): 1,570,000 node-hours on Frontier

Research Summary:

The objective of this research is the 3D transverse spin structure of the proton, which will test the fundamental property of time-reversal symmetry in strong interaction phenomena.

Using the Euclidean lattice formulation of quantum chromodynamics (QCD), the theory of strong interaction, the research team will carry out first-principles calculation of quark 3D momentum distributions in a proton, which are sensitive to the transverse spin of the quark or proton. The team will determine how a transversely polarized quark move within a proton with transverse spin, and verify the sign-change of Sivers and Boer-Mulders functions under the time-reversal transformation, which is a prominent prediction of QCD. The results from this calculation will impact the transverse-spin-asymmetry experiments at Fermilab, RHIC, Jefferson Lab 12 GeV upgrade, and the forthcoming Electron-Ion Collider.