**Title:**     Large-Eddy Simulation Component of the Mesoscale Convective System Climate Model Development and Validation (CMDV-MCS) Project

**Principal Investigator:**     William Gustafson (Pacific Northwest National Laboratory)

**Co-Investigators:**     Steven Ghan (Pacific Northwest National Laboratory) | Jiwen Fan (Pacific Northwest National Laboratory) | Hailong Wang (Pacific Northwest National Laboratory)

**ALCC allocation:**
  **Site:**     Argonne Leadership Computing Facility (ALCF)
  **Allocation:**     74,000,000 core hours

## Research Summary

Deep convection as mesoscale convective systems (MCSs) is one of the most climatologically significant cloud forms. MCSs release latent heat that drives the general circulation; provide the dominant source of precipitation in many regions; redistribute water, clouds, and aerosol; and affect radiation through the atmosphere. Yet, climate models poorly simulate MCSs due to factors including the multiscale nature of convection. The Climate Model Development and Validation (CMDV) project focuses on improving the parameterization of MCSs in the Accelerated Climate Model for Energy, and the project is built upon a symbiotic hierarchy of models that span scales to address the range of processes necessary to tackle this problem. The computing allocation request specifically addresses the large-eddy simulation (LES) needs of the project during 2017. LES modeling of deep convection, particularly on the scale of MCSs, is extremely expensive. Therefore, the research aims to develop an alternative methodology using a smaller LES domain that encompasses a portion of a larger convective event, but that propagates with the convection to keep the domain aligned with the region of interest. Establishing the validity of this methodology requires a control simulation encompassing a much larger region for comparison. Carrying out a larger LES of a convective system and a series of smaller LES simulations will demonstrate the sensitivity of the results to the domain size and location with respect to the active convection.

**Title:** Investigating climate-biogeochemistry feedbacks with the Department of Energy's next-generation high-resolution Earth system model

**Principal Investigator:** Peter Thornton (Oak Ridge National Laboratory)

**Co-Investigators:** David Bader (Lawrence Livermore National Laboratory) | Ruby Leung (Pacific Northwest National Laboratory) | William Riley (Lawrence Berkeley National Laboratory) | Mark Taylor (Sandia National Laboratory) | Katherine Calvin (Pacific Northwest National Laboratory) | Forrest Hoffman (Oak Ridge National Laboratory) | Matthew Maltrud (Los Alamos National Laboratory)

**ALCC allocation:**
  **Site:** National Energy Research Scientific Computing Center (NERSC)
  **Allocation:** 60,000,000 core hours

## Research Summary

The Department of Energy's Accelerated Climate Model for Energy, Version 1 (ACME v1) includes unique capabilities in the representation of carbon cycle and other biogeochemical processes within its land and ocean components. Of particular note, ACME v1 is the first fully coupled Earth system model to include a representation of interacting carbon, nitrogen, and phosphorus cycles. The proposed research includes a series of experiments to evaluate the influence of carbon nutrient dynamics on fundamental climate system feedbacks under present-day and future climate scenarios. The research will explore the major feedback components driven by increasing carbon dioxide concentration and by changes in temperature and precipitation. It will also investigate the influence of model parametric and structural uncertainty using a comprehensive uncertainty quantification framework developed and tested with previous ALCC support. The model structural uncertainty investigation will focus on the representation of competition among plants and microbes for access to and acquisition of nitrogen and phosphorus resources. Those nutrients become available in the soil as organic matter (dead plant parts and byproducts of microbial growth and metabolism) is decomposed. The outcome of the study will be a first-ever assessment of the role of multiple nutrient limitations in the regulation of climate-carbon cycle feedbacks within a fully coupled Earth system model. The study may lead to lead to significant improvements in our ability to make robust predictions of future climate states under specified scenarios for greenhouse gas forcing.

**Title:**  Understanding the Role of Ice shelf-Ocean Interactions in a Changing Global Climate

**Principal Investigator:**  Mark Petersen (Los Alamos National Laboratory)

**Co-Investigators:**  Stephen Price (LANL) | Todd Ringler (LANL) | Jones Philip (LANL) | Luke Van Roekel (LANL) | Jeremy Fyke (LANL) | Xylar Asay-Davis (LANL) | Phillip Wolfram (LANL)

**ALCC allocation:**
  **Site:**  Argonne Leadership Computing Facility (ALCF)
    **Allocation:**  25,000,000 core hours
  **Site:**  Oak Ridge Leadership Computing Facility (OLCF)
    **Allocation:**  2,000,000 core hours
  **Site:**  National Energy Research Scientific Computing Center (NERSC)
    **Allocation:**  60,000,000 core hours

## Research Summary

Projecting the rate and probability of future sea-level rise is a primary science driver for the Department of Energy's (DOE) Accelerated Climate Model for Energy (ACME) project. ACME is focusing on better understanding the potential for rapid sea-level rise from the Antarctic Ice Sheet (AIS) where both observations and modeling indicate that ice sheet mass loss is highly sensitive to changes in oceanic melting of floating ice shelves. To understand this sensitivity and its potential impact on future sea level rise, ACME scientists have introduced new, variable-resolution ocean and sea-ice models for simulating the key processes responsible for delivering warm waters to the base of Antarctic ice shelves. The proposed research includes a suite of simulations of varying resolution and complexity designed to: (1) validate the ACME ocean models in global, coupled configurations with ocean circulation in ice shelf cavities; (2) explore the influence of global climate on sub-ice shelf melting (and vice versa); (4) explore the impact of anticipated changes in climate on sub-ice shelf melting. The simulations will produce a greatly improved understanding of the coupled interactions between ice sheets, oceans, and global climate, which is a necessary first step in being able to accurately project future sea-level change using dynamic ice sheet models coupled to Earth System Models. The research brings together the new ACME capability of modeling ice sheet-and-ocean interactions with DOE's expertise in climate modeling and high performance computing, and applies it to the pressing question of future ice shelf melting in Antarctica.

**Title:** Achieving the potential of ultraparameterization to advance understanding of low cloud / climate feedbacks

**Principal Investigator:** Michael Pritchard (University of California, Irvine)

**Co-Investigators:** Christopher Bretherton (University of Washington) | Balwinder Singh (Pacific Northwest National Laboratories) | Hossein Parishani (University of California, Irvine)

**ALCC allocation:**
  **Site:** National Energy Research Scientific Computing Center (NERSC)
  **Allocation:** 26,000,000 core hours

## Research Summary

Diverse low-lying clouds cover much of the world's oceans, where they act to cool the globe by reflecting sunlight that would otherwise warm the water. They depend on small turbulent eddies that pick moisture off the ocean surface which lift it into thin cloud sheets where the ocean is cool and patchy cumuli where it is warm. These motions are far too small to resolve with the grid of a global climate model, leading to large uncertainties in simulating the global distribution of low clouds. The multiscale challenge of simultaneously representing global and cloud-scale circulations underlies decades-long uncertainties about the effects of cloud feedbacks and changing aerosol emissions on climate change. A novel global model, UPCAM, sidesteps this uncertainty by explicit simulating the turbulent eddies and low clouds in a network of small interacting cloud-resolving models with very fine grids. Weather hindcast tests show UPCAM skillfully simulates the vertical and horizontal distribution of clouds. Parallelization advances and algorithmic acceleration now make UPCAM feasible for multiyear simulations with petascale computing resources. The proposed research will realize the potential of UPCAM by making 10-year long simulations in current and perturbed climates, and with and without current human contributions to aerosol emissions, providing an important new comparison with simulations by conventional climate models.

**Title:**     Non-Boussinesq effects on buoyancy-driven variable density turbulence

**Principal Investigator:**     Daniel Livescu (Los Alamos National Laboratory)

**Co-Investigators:**     Arindam Banerjee (Lehigh University) | Denis Aslangil (Lehigh University)

**ALCC allocation:**
  **Site:**          Argonne Leadership Computing Facility (ALCF)
  **Allocation:**     60,000,000 core hours

## Research Summary

The proposed research includes high-resolution Direct Numerical Simulations (DNS) to complete a study of the effects of high density-ratios on mixing. The proposed simulations aim to investigate the material mixing of two or more incompressible, miscible fluids with different densities as a consequence of stirring by turbulence created by buoyancy forces. The mixing of two or more miscible fluids with different densities (or molar masses) is of fundamental interest due to occurrence in atmospheric and oceanic flows, supernova formations, combustion and several engineering applications such as in inertial confinement fusion (ICF) capsules. The data will enable an understanding of the asymmetric behavior of mixing at higher density-ratios. In addition, transport equations for the moments of density gradient will provide unique information about the evolution of the high-order statistics of the scalar field and will allow estimation of the validity of the incompressible assumption in high-density ratio turbulence. All proposed DNS will be performed with the variable-density version of the peta-scale CFDNS code developed at Los Alamos National Laboratory. The database not only will be used to study the physics of the large-density ratio flows but also will provide fidelity data for verification/calibration of mix models and non-DNS codes, which is very important for the LANL predictive capability within the stewardship mission.

**Title:** Wall-Resolved Large Eddy Simulation of Transonic Shock-Induced Flow Separation

**Principal Investigator:** Mujeeb Malik (NASA Langley Ressearch Center)

**Co-Investigators:** Ali Uzun (National Institute of Aerospace)

**ALCC allocation:**
  **Site:** National Energy Research Scientific Computing Center (NERSC)
  **Allocation:** 84,000,000 core hours

## Research Summary

The Department of Energy and NASA's Green Aviation program share the strategic goals of reducing fuel usage and cutting carbon emissions. Reducing the drag on commercial aircraft is one way to meet these goals, but requires validated high-fidelity simulation techniques for drag prediction. One of the most challenging problems in flow simulation is the prediction of flow separation. Flow separation occurs when the thin boundary layer of slow moving air attached to a solid surface detaches from the surface in the presence of an adverse pressure gradient. The pressure gradient may be a result of the geometry of the surface, or the presence of a shock wave. The phenomenon commonly occurs in flows over transonic airfoils, helicopter rotors, turbomachinery blades and high-lift configurations, to name a few applications. Separation often leads to increased aerodynamic drag, stall and reduced performance. The proposed research will perform fluid dynamic simulations as part of a project concerned with aerodynamic drag prediction on transport aircraft wings. Also, virtual tests will be carried out to gather detailed information regarding the physics of flow separation. The simulation results will help in the improvement of existing turbulence models, which can be routinely used to predict drag and improve aerodynamic designs. The proposed work thus has a strong potential to provide the critical knowledge needed in the development of fuel-efficient air transportation technologies in which aerodynamic drag reduction offers significant performance and energy-savings benefits.

**Title:** High-Fidelity Simulations of Gas Turbine Stages for Model Development using Machine Learning

**Principal Investigator:** Gregory Laskowski (General Electric Aviation)

**Co-Investigators:** Richard Sandberg, (University of Melbourne) | Vittorio Michelassi (General Electric Oil&Gas/Aviation) | Richard Pichler (University of Melbourne) | Jack Weatheritt (University of Melbourne) | Gustavo Ledezma (General Electric Global Research)

**ALCC allocation:**
  **Site:** Oak Ridge Leadership Computing Facility (OLCF)
   **Allocation:** 30,000,000 core hours

**Research Summary**

The purpose of the research is to leverage recent advances in the fidelity of Direct Numerical Simulations (DNS) and Large Eddy Simultion (LES) tools to improve understanding and modeling of the complex unsteady physics that occur in the high-turbulence environment of gas turbines. Utilizing the HiPSTAR code that has been optimized for HPCs in previous programs, the research aims to identify opportunities to increase turbine aero-thermal efficiency by 2-4% and extend hot-gas-path durability, translating into combined cycle efficiency gains of 0.4%-0.8%. Using novel machine learning tools the data will also be used to develop affordable turbulence simulation methodologies that embody a step-change in predictive accuracy.

**Title:** Steady State Calibration of a Diesel Engine in CFD using a GPU-based Chemistry Solver and Conjugate Heat Transfer

**Principal Investigator:** Ronald Grover (General Motors)

**Co-Investigators:** Jian Gao (General Motors) | Venkatesh Gopalakrishnan (General Motors) | Ramachandra Diwakar (General Motors) | Wael Elwasif (Oak Ridge National Laboratory) | K. Dean Edwards (Oak Ridge National Laboratory) | Charles Finney (Oak Ridge National Laboratory) | Russell Whitesides (Lawrence Livermore National Laboratory) | Keith Richards (Convergent Science, Inc.)

**ALCC allocation:**
  **Site:** Oak Ridge Leadership Computing Facility (OLCF)
   **Allocation:** 5200000 core hours

## Research Summary

The automotive industry has been increasingly challenged to invest heavily in innovative powertrain technologies. These advancements can result in a complex array of tunable control parameters to optimize engine performance over a range of operating conditions. Hence, the calibration process of a modern passenger car diesel engine consumes significant time, effort and resources, making it a bottleneck during the development process. Advanced modeling tools, such as CFD, are often used with the goal of streamlining portions of the calibration process. The usefulness of CFD simulations tools for in-cylinder engine combustion is often compromised by the computational overhead of detailed chemical kinetics and uncertainty in the combustion chamber wall temperatures. Specifically, traditional diesel engine CFD simulations consist of partial geometry sector mesh computations utilizing reduced order kinetics mechanisms, fixed spatially uniform wall temperature boundary conditions, and a prescribed solid body swirl velocity field prior to spray injection in lieu of computing air induction with valve motion. The proposed research seeks to leverage recent advancements in CFD to improve the accuracy CFD computations and accelerate engine calibration. First, higher order kinetics will be solved using a GPU-based chemical kinetics solver (leveraged in past ALCC awards CMB119 and CMB124). Next, full in-cylinder 3D spray, flow and combustion simulations will be undertaken considering conjugate heat transfer to predict temporally and spatially varying wall temperature boundary conditions. The results will be analyzed to compare differences in combustion and emissions (NOx, CO, UBHC, Smoke) with actual engine measurements.

**Title:**       Interfaces to Control Chemical Reactivity of Confined Liquids

**Principal Investigator:**       Teresa Head-Gordon (University of California Berkeley)

**Co-Investigators:**       Luis Ruiz Pestana (Lawrence Berkeley National Laboratory)

**ALCC allocation:**
  **Site:**       National Energy Research Scientific Computing Center (NERSC)
  **Allocation:**       100,000,000 core hours

## Research Summary

The ability to control or accelerate a chemical reaction process is of great importance in basic energy science applications such as hydrogen production, energy conversion in fuel cells or artificial photosynthesis technologies. A particularly promising strategy to guide chemical reactions, ubiquitous in biological systems, is the use of confinement and the resulting altered physical properties of the aqueous solvent. Inspired by the dramatic changes that water undergoes when placed under hydrophobic nanoconfinement, the theoretical work proposed here aims to develop the ability to explore, using ab initio molecular dynamics (AIMD), catalytic reactions through altered solvent properties by tuning the chemistry and geometry of confining surfaces. Specifically, we will investigate methanol oxidation in confined water, an important intermediate reaction during hydrogen production from biomass. This ambitious work requires the large-scale impact that ASCR Leadership Computing Challenge can bring to the fundamental problem of chemical reactivity and catalysis under nanoconfinement by using more sophisticated levels of theory and large scale deployment of these advanced ab initio models and techniques, recently implemented in the HPC software platform CP2K and associated python libraries. The physical accuracy granted by the advanced AIMD techniques that we propose will help resolve long unsettled controversies regarding the confined behavior of water and will allow us to develop comprehensive, reliable knowledge on how to control reactions using altered solvent properties through confinement interfaces.

**Title:**     Computational engineering of electron-vibration coupling mechanisms

**Principal Investigator:**     Marco Govoni (Argonne National Laboratory)

**Co-Investigators:**     Huihuo Zheng (Argonne National Laboratory) | Christopher Knight (Argonne National Laboratory) | Ryan McAvoy (University of Chicago) | He Ma (University of Chicago) | Giulia Galli (University of Chicago) | Francois Gygi (University of California Davis)

**ALCC allocation:**
  **Site:**     Argonne Leadership Computing Facility (ALCF)
    **Allocation:**     60,000,000 core hours
  **Site:**     National Energy Research Scientific Computing Center (NERSC)
    **Allocation:**     15,000,000 core hours

## Research Summary

Solid-state technologies, light-emitting or absorbing devices, and spintronics applications strongly depend on the dynamics of excited carriers at finite temperature. The best candidate materials for these applications are routinely screened by computationally engineering their electronic properties. However, calculations of electronic structures that explicitly include the coupling between electrons and lattice vibrations (electron-phonon coupling) are rarely performed mostly due to their complexity and to the lack of an efficient implementation. This project will develop a first-principles theoretical framework applicable to studying the electron-phonon scattering mechanisms in both technologically mature and candidate nanostructured materials, thus enabling the investigation of materials for energy and quantum applications. The scalable, integrated first principles algorithms employed in this project combined with the unique computational resources provided by Mira, Theta, and Cori will enable studies of unprecedented scope. Key outcomes of this project will be improved predictions of the properties of new energy materials such as nanostructured inorganic solar cells and thermoelectric devices that will help optimize future device performances.

**Title:** Imaging Transient Structures in Heterogeneous Nanoclusters in Intense X-ray Pulses

**Principal Investigator:** Phay Ho (Argonne National Laboratory)

**Co-Investigators:** Chris Knight (Argonne Leadership Computing Facility) | Linda Young (Argonne National Laboratory) | Christoph Bostedt (Argonne National Laboratory)

**ALCC allocation:**
  **Site:** Argonne Leadership Computing Facility (ALCF)
  **Allocation:** 68,000,000 core hours

## Research Summary

This project seeks to develop a fundamental understanding of the interaction of intense x-rays with complex nanoscale systems, and to examine the potential applications of these x-rays. Study of the ultrahigh intensity regime of x-ray interactions with matter began with the advent of the Linac Coherent Light Source (LCLS) x-ray free-electron laser (XFEL). The achieved x-ray intensities can strip electrons from atoms from the inside out and create transient states of matter as the x-ray pulse progresses through the target. The project goal is to understand the correlation of sample excitation with the ultimate spatial resolution and elemental contrast achievable with XFEL pulses in heterogeneous samples, and to examine the dynamical processes induced by intense x-rays at the most fundamental level. A hybrid quantum/classical simulation methodology will be used to investigate the interactions of intense XFEL pulses with nanometer-sized samples to understand the ultrafast dynamics and its connection to x-ray imaging. The work proposed is an important step forward in understanding high-brightness, coherent x-ray laser pulses and their interactions with matter. Thus, the goal of this project is to develop a quantitative and predictive understanding of x-ray matter interactions in nanosized heterogeneous systems at high intensities at the level of atoms and electrons, characterize complex phenomena that emerges (from attosecond—picosecond timescales), and thus guide the experimental strategies with multicolor pulses at current and future XFEL sources. Particular attention is paid to the scattering response of nanosized heterogeneous systems to examine the imaging contrast and to control the ultrafast transient dynamics by exploiting sample heterogeneity with two x-ray pulses. This topic that is largely unexplored so far, but the result will be widely applicable to many systems and essential in guiding future XFEL experiments, and in establishing the applied methodology as an effective large-scale computational tool for the new research frontier of ultrafast x-ray science.

**Title:** Predictive Modeling of Functional Nanoporous Materials, Nanoparticle Assembly, and Reactive Systems

**Principal Investigator:** J. Ilja Siepmann (University of Minnesota)

**Co-Investigators:** Laura Gagliardi (University of Minnesota) | Jason Goodpaster (University of Minnesota) | Chris Knight (Argonne National Laboratory) | Christopher Mundy (Pacific Northwest National Laboratory) | David Sholl (Georgia Institute of Technology) | Randall Snurr (Northwestern University) | Donald Truhlar (University of Minnesota) | Yongchul Chung (Pusan National University)

**ALCC allocation:**
**Site:** Argonne Leadership Computing Facility (ALCF)
**Allocation:** 130,000,000 core hours
**Site:** National Energy Research Scientific Computing Center (NERSC)
**Allocation:** 16,000,000 core hours

## Research Summary

This interdisciplinary, collaborative team will use predictive modeling to accelerate the discovery and design of materials for a variety of energy-related applications. The computational campaigns in this ALCC project are integral parts of this team's ongoing research efforts and are aligned with the Department of Energy's mission to advance scientific and technological capabilities with innovative discoveries. The systems/processes targeted in this project are of significant interest for the chemical, biorenewable, and petrochemical industries, and improving nanoporous materials for these applications has tremendous societal benefits. The project's research objectives largely support goals of the DOE-funded Nanoporous Materials Genome Center (NMGC). The research objectives also support BES's mission in the fundamental core programs of CGBS through the Chemical Theory and Computation Program and of MSE through the Synthesis and Processing program. In addition, some of the application projects are directly related to the separation tasks of a DOE EERE award and an NSF SI2 (Software Infrastructure for Sustained Innovation) award, and other application projects indirectly support the goals of four Energy Frontier Research Centers (EFRCs): Inorganometallic Catalyst Design Center, Center for Gas Separations Relevant to Clean Energy Technologies, Catalysis Center for Energy Innovation, and Center for Understanding and Control of Acid Gas-Induced Evolution of Materials for Energy.

**Title:** The Materials Project - Pathway to Piezoelectrics

**Principal Investigator:** Kristin A. Persson (Lawrence Berkeley National Lab)

**Co-Investigators:** Joseph Montoya (Lawrence Berkeley National Lab) | Patrick Huck (Lawrence Berkeley National Lab) | Shyam Dwaraknath (Lawrence Berkeley National Lab)

**ALCC allocation:**
  **Site:**         National Energy Research Scientific Computing Center (NERSC)
    **Allocation:**   66,000,000 core hours

## Research Summary

This project's goal is using high-throughput computations and best practices from modern data management to create an "open, collaborative, and data-rich ecosystem for accelerated materials design". The project uses automated high-throughput workflows to perform and analyze results of electronic structure simulations. In addition, by designing open and efficient user interfaces to this data, the project will enable the materials science community to derive new insights from large, high-quality datasets and leverage its collective talents to their maximal extent. The project will compute the elastic and piezoelectric tensor via density functional theory (DFT) for 18,500 potential piezoelectrics to determine which chemistries and structures are ripe for further analysis. The resources provided by ALCC will support an effort to compute, organize, and publicly disseminate elastic and piezoelectric tensors data such that the discovery of materials with targeted sets of these properties may be accelerated.

**Title:**         ECP Consortium for Exascale Computing

**Principal Investigator:**         Paul Messina (Argonne National Laboratory)

**Co-Investigators:**         Julia White (Oak Ridge National Laboratory) | Doug Kothe (Oak Ridge National Laboratory) | Rajeev Thakur (Argonne National Laboratory) | Jim Ang (Sandia National Laboratories) | Terri Quinn (Lawrence Livermore National Laboratory)

**ALCC allocation:**
  **Site:**         Argonne Leadership Computing Facility (ALCF)
    **Allocation:**         530,000,000 core hours
  **Site:**         Oak Ridge Leadership Computing Facility (OLCF)
    **Allocation:**         300,000,000 core hours
  **Site:**         National Energy Research Scientific Computing Center (NERSC)
    **Allocation:**         139,000,000 core hours

## Research Summary

The ALCC allocation will be used for research and development (R&D) activities of the Exascale Computing Project (ECP), which include the development of exascale-capable Department of Energy (DOE) mission critical applications, implementations of computing motifs that are widely used, and optimization of many components of the exascale software stack. The ECP will develop scientific and large-data applications that exploit the emerging exascale-era computational architectures.

The majority of the applications codes that will receive time for exascale development are mature petascale codes that address major DOE programs and are already in use at the various ASCR leadership computing facilities. These representative codes cover science themes in accelerator physics, additive manufacturing, astrophysics, chemistry, climate, combustion, cosmology, materials, multiphase flows and computational fluid dynamics, fusion, nuclear physics, nuclear reactors, precision medicine and genomics, and subsurface modeling.

The software technology projects represent programming models and runtimes, tools, mathematical and scientific libraries and frameworks, data management and workflows, data analytics and visualization, system software, and resilience and integrity. The projects are critical components in the software stack for future exascale systems and form the infrastructure underpinning all ECP applications. The ALCC allocation will be used for testing newly developed features in these projects, for performance testing, and for testing the software at large scale.

The ECP R&D projects requiring computing resources through the ALCC program were selected with input from DOE stakeholders (e.g., program managers and senior staff) and are those that are important and relevant to programmatic mission needs.

**Title:**    Quantum Monte Carlo computations of chemical systems

**Principal Investigator:**    Olle Heinonen (Argonne National Laboratory)

**Co-Investigators:**    Yongkyung Kwon (Konkuk University) | Anouar Benali (Argonne National Laboratory) | Hyeondeok Shin (Argonne National Laboratory)

**ALCC allocation:**
  **Site:**    Argonne Leadership Computing Facility (ALCF)
  **Allocation:**    5,000,000 core hours

## Research Summary

A very large part of economic activity across the world depends on chemical processes. Much of the knowledge of how to harness chemical processes has come form empirical knowledge -- trial and error. This is for example the case for catalysis, which embodies chemical reactions that drive economies around the globe. Recently, there have been great efforts towards designing chemical processes using a closed cycle of theory, modeling, and experiments. This has become possible not in a small part because of advances in computational methods and computer hardware that have led to the possibility to more accurately model atoms, molecules, and chemical processes. Nevertheless, the presently most used computational methods are either efficient but suffer from a lack of accuracy that is needed to model chemical systems, or are prohibitively expensive to use. The allocation supports research using a highly accurate approach to model some fundamental chemical systems and reactions. This allocation supports work to develop more accurate and efficient methods and study novel low-dimensional materials that are not accessible to other methods at a required accuracy. The method used is the Quantum Monte Carlo (QMC) method, which includes all-important -- but frequently elusive -- electron-electron interactions at the quantum mechanical level. Outcomes of this project will help advnace methods for more physcially accurate and predictive modeling of chemical processes.

**Title:** Accelerated Global Human SettlementDiscovery

**Principal Investigator:** Jibonananda Sanyal (Oak Ridge National Laboratory)

**Co-Investigators:** Lexie Yang (Oak Ridge National Laboratory) | Dalton Lunga (Oak Ridge National Laboratory) | Jiangye Yang (Oak Ridge National Laboratory) | Mark Coletti (Oak Ridge National Laboratory) | Jeanette Weaver (Oak Ridge National Laboratory) | Robert Stewart (Oak Ridge National Laboratory)

**ALCC allocation:**
  **Site:** Oak Ridge Leadership Computing Facility (OLCF)
   **Allocation:** 25,000,000 core hours

**Research Summary**

Understanding where people live is fundamental to understanding what people do and what their social needs are with respect to energy security; policy and urban development; resiliency; disaster and emergency response; intelligence and security; humanitarian support, as well as understanding the behavioral social dynamics. This allocation supports computational research driven by ongoing efforts where initial investigations have been conducted to explore the application of machine learning — including deep learning — toward global scale human settlement detection from high resolution satellite imagery. Several use cases including support for Federal Emergency Management Agency (FEMA)'s disaster response in the western United States, to meet specific needs of the intelligence community, and strategic resourcing for polio and malaria eradication in sub Saharan Africa by the Gates Foundation are benefiting from this work. Thus far, the approaches undertaken run well on single Graphics Processing Unit (GPU)s and small clusters, and are well positioned to exploit a resource such as OLCF's Titan. The high resolution determination of settlements necessitates being able to dynamically produce very large datasets with fine temporal resolutions to adequately capture changes and fluxes. This problem is multi-petascale computationally and multi-petabytes in imagery data size. The successful completion of this research will produce a repeatable machine learning and image-processing pipeline that runs at continental to global scales at high resolutions and for the purposes of project sponsors, produces high resolution human settlement maps that are unprecedented in resolution and accuracy.

**Title:** Improved Physics-Based Modeling of Ground Motion Seismic Hazard Using High Performance Computing

**Principal Investigator:** Thomas Jordan (University of Southern California)

**Co-Investigators:** Christine Goulet (University of Southern California) | Yifeng Cui (San Diego Supercomputer Center) | Philip Maechling (University of Southern California) | Kim Olsen (San Diego State University) | Alexander Breuer (San Diego Supercomputer Center) | Dawei Mu (San Diego Supercomputer Center)

**ALCC allocation:**
  **Site:**           National Energy Research Scientific Computing Center (NERSC)
    **Allocation:**    14,100,000 core hours

## Research Summary

The Southern California Earthquake Center (SCEC) conducts and coordinates fundamental and applied research on earthquake system science using southern California as its main natural laboratory. Currently, over 1000 earthquake professionals at over 70 research institutions are participating in SCEC projects, covering a wide range of disciplines. SCEC differs from other research centers in that while it fosters new scientific development and discoveries from individuals and collaborative groups, it also integrates research results and new technologies into broader impact computational products to improve seismic hazard assessment.

This project will bring SCEC ground motion simulations to the next generation of supercomputers through the use of two Xeon Phi based codes: EDGE, a new finite element code based on the discontinuous Galerkin Finite Element Method and AWP-ODC, a wave propagation code currently used by the research team in a large suite of workflows to model ground motions from earthquakes. This project supports two main science objectives: 1) study uncertainties in dynamic rupture simulations, and 2) increase the upper frequency at which we can perform ground motion simulations while integrating more realistic physics.

Together, these highly efficient codes will be used to improve our understanding of ground motions through the integration of better physics. The results will allow researchers to better quantify ground motion hazard and its uncertainties. Because hazard computations usually control seismic risk at return periods of design interest, this work will contribute to reduced seismic risk for society. The outcomes of the proposal will be improved methods for seismic simulation, increased predictive power for seismic risk, and increased safety from seismic hazards.

**Title:** Multiscale Simulations of Hematological Disorders

**Principal Investigator:** George Karniadakis

(Browns University)

**Co-Investigators:** Xuejin Li (Brown University) | Alireza Yazdani (Brown University) | Zhen Li (Brown University)

**ALCC allocation:**
  **Site:** Argonne Leadership Computing Facility (ALCF)
    **Allocation:** 20,000,000 core hours
  **Site:** Oak Ridge Leadership Computing Facility (OLCF)
    **Allocation:** 26,000,000 core hours

## Research Summary

Thoracic aoritic aneurysms is an enlargement of the main blood vessel that carries blood away from the heart. Thoracic aortic aneurysms and dissections (TAADs) are responsible for significant mortality and morbidity, with estimates ranging from 30,000 to 50,000 associated deaths per year in USA alone. Sickle cell anemia (SCA) is one of the most common inherited blood diseases with more than 270,000 new patients each year; 72,000 individuals are suffering from SCA in USA. This allocation supports research to investigate TAADs, and perform *in-silico* testing of new drugs using a multiscale numerical model for SCA developed by the research team. Using the numerical framework, this allocation will help researchers gain quantitative understanding of several other hematological disorders such as hereditary spherocytosis (HS) and hereditary elliptocytosis (HE), which are inherited blood diseases that result in anemia. More specifically, the project will examine *in-silico*, for the first time, the onset of aortic dissections and subsequent thrombus formation and growth in TAAD, as well as the intrinsic biological processes that govern SCA, HS and HE.

**Title:** Protein-Protein Recognition and HPC Infrastructure

**Principal Investigator:** Benoit Roux (University of Chicago)

**Co-Investigators:** Donghyuk Suh (University of Chicago) | Chris Chipot (University of Illinois in Urbana-Champaigne)

**ALCC allocation:**
  **Site:** Argonne Leadership Computing Facility (ALCF)
  **Allocation:** 80,000,000 core hours

## Research Summary

A living cell is made of many different smaller pieces including macromolecules called proteins. Proteins are built from atoms and many work as nano-particle sized bio-machinery for the cell. Living cells can be viewed as a collection of these proteins and other macromolecules carrying out their well-defined tasks in the cellular enviornment. In order to perform their tasks, proteins often come together and form a complex. The ability to recognize their partner and form a tight complex is important. Disruption of complex formation due to mutation, often leads to loss of function and can lead to disease, such as cancer. Scientists therefore need to gain a deep understanding on protein recognition and assemblies at the atomic level.

To achieve the goal of understanding protein recognition and complex formation, this research team has developed powerful and scalable computational methodologies that are well-adapted for leadership computers. The team's method breaks-down the binding process into several physically meaningful steps and expresses the total binding free energy – the energy associated with complex-formation – as a sum of free energies associated with each of these smaller steps. This step-by-step framework provides a rational and tractable starting point to characterize protein-protein interaction quantitatively and explain how mutations and structural changes are associated with observed thermodynamic effects.

This allocation supports rigorouse testing of atomic scale (molecular dynamics) free energy simulation methods by considering a set of representative protein complexes of increasing size and complexity. The results will support determination of best practices and the building of a pipeline that automates the simulation setup according to the chosen protocol. This computational infrastructure will reduce human error and rapid simulation setup/execution in high performance computers. The research will result in great technical advances with cutting-edge computational methodologies, and generate new knowledge about an important class of biological systems.

**Title:**     Spin-Forbidden Catalysis on Metal-Sulfur Proteins

**Principal Investigator:**     Sergey Varganov (University of Nevada, Reno)

**Co-Investigators:**     Yuri Alexeev (Argonne National Laboratory) | Dmitri Fedorov (National Institute of Advanced Industrial Science and Technology, Japan)

**ALCC allocation:**
  **Site:**            Argonne Leadership Computing Facility (ALCF)
   **Allocation:**       42,000,000 core hours

## Research Summary

In chemistry, catalysts are ingredients in a reaction that speed up the reaction rate. In many energy and technology applications, chemical reactions are too slow for practical use and must be accelerated using catalysts. Modern energy applications require inexpensive catalysts for several "critical reactions" such as hydrogen oxidation/reduction, water splitting, and CO2 reduction to fuels. The development of catalysts based on earth-abundant and inexpensive first-row (first-row in the periodic table) transition metals (TM) is a critical challenge. Biological systems, such as metal sulfur proteins, have effectively used first-row TM to catalyze many important chemical reactions. Understanding how these TM based catalysts work and how their analogs can be made to work outside fragile biological systems is a major step toward developing future industrial catalysts.

This allocation supports computations to provide fundamental insight into the catalytic mechanisms used by biological systems and to help design new catalysts based on inexpensive and earth-abundant transition metals. Outcomes of this research will  be improved understanding for future development of effective, cost-saving catalysts for improved industrial processes and energy technology.

**Title:** Spin and Three-Dimensional Structure of the Nucleon

**Principal Investigator:** Huey-Wen Lin (Michigan State University)

**Co-Investigators:** Xiangdong Ji (University of Maryland) | Luchang Jin (Brookhaven National Laboratory) | Peng Sun (Michigan State University) | Yi-Bo Yang (Michigan State University) | Jian-Hui Zhang (Universit\"at Regensburg) | Jiunn-Wei Chen (National Taiwan University) | Tomomi Ishikawa (Shanghai Jiao Tong University)

**ALCC allocation:**
  **Site:**      National Energy Research Scientific Computing Center (NERSC)
   **Allocation:**      25,000,000 core hours

## Research Summary

We propose to study the spin and three-dimensional structure of nucleons, the particles (protons and neutrons) that compose the nuclei of all the elements in the periodic table. Since the discovery in the 1960s of quarks and gluons, the smaller particles that make up the nucleon, many experiments have expanded our knowledge of nucleon properties and internal structure. However, as we learn more, many questions remain: How do the constituents make up the nucleon? Where does the nucleon mass come from? What is the role of antiquarks and gluons? What is the transverse structure of nucleon? DOE has funded the 12-GeV upgrade at Jefferson Laboratory and planned a future electron-ion collider (EIC) in the US to fill in the gaps in our understanding. These large-scale experiments must proceed hand-in-hand with theoretical advances. A recently-formed special DOE topical collaboration entitled "Coordinated Theoretical Approach to Transverse Momentum Dependent Hadron Structure in QCD" is trying to bring together a wide range of theorists who traditionally use different approaches to study problems like this to form a universal understanding and advance the theory to reveal the final secrets of the nucleon. This project is an essential building block in the development of a pioneering new approach to calculating nucleon properties with well-understood uncertanties that is complementary to and necessary for DOE's experimental program.

**Title:** Nuclear Spectra with Chiral Forces

**Principal Investigator:** Alessandro Lovato (Argonne National Laboratory)

**Co-Investigators:** Ewing Lusk (Argonne National Laboratory) | Maria Piarulli (Argonne National Laboratory) | Steven Pieper (Argonne National Laboratory) | Robert Wiringa (Argonne National Laboratory)

**ALCC allocation:**
  **Site:** Argonne Leadership Computing Facility (ALCF)
   **Allocation:** 35,000,000 core hours

## Research Summary

Reliably modeling their structure of the particles that make up atomic nuclei, as well as their interactions with neutrinos, electrons, nucleons, and other nuclei is essential for DOE's large nuclear physics experiments at Argonne, FermiLab, Jefferson Lab, and Michigan State University. At present, it is not possible to derive the forces between nucleons from fundamental principles with sufficient accuracy. Instead, theoretical models use approximations of the true forces with parameters that are determined by fits to experimental data. This introduces uncertainties in the resulting predictions. To understand how large the uncertainty in a particular prediction might be, it's necessary to consider a range of realistic approximations to the true force.

Given a nuclear force, there are several methods to predict nuclear properties, such as their size and how tightly bound they are. This project uses quantum Monte Carlo, which is the most reliable for small nuclei. Over the last twenty years, we have used one particular force model and our quantum Monte Carlo program to make the most accurate available calculations for up to 12 nucleons. Recent work by another team provides systematic families of nuclear forces that could provide the range of options needed to understand the uncertainty in predictions made using these forces. We have been adapting their work for use with quantum Monte Carlo. We will use one of DOE's forefront computers, Argonne's Theta, to determine the parameters in these forces and then to model nuclei with them. These results, combined with those of our original model, will be important indicators of the reliability of theoretical predictions of nuclear properties, particularly those needed by DOE experiments. We will also provide the resulting forces to the nuclear physics community for use with other methods applicable to larger nuclei.

**Title:** Lattice QCD Equation of State for Moderately Large Baryon Densities: stretching the limits

**Principal Investigator:** Swagato Mukherjee (Brookhaven National Laboratory)

**Co-Investigators:** Alexei Bazavov (Michigan State University) | Frithjof Karsch (Brookhaven National Laboratory) | Edwin Laermann (Bielefeld University, Germany) | Peter Petreczky (Brookhaven National Laboratory) | Christian Schmidt (Bielefeld University, Germany) | Patrick Steinbrecher (Brookhaven National Laboratory)

**ALCC allocation:**
  **Site:** Oak Ridge Leadership Computing Facility (OLCF)
    **Allocation:** 32,000,000 core hours
  **Site:** National Energy Research Scientific Computing Center (NERSC)
    **Allocation:** 170,000,000 core hours

**Research Summary**

One of the central goals of the Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory is to probe the phase diagram of quarks and gluons, which are the building blocks of atomic nuclei. The familiar phase diagram for water indicates when we can expect to find steam, liquid water, or ice. The phase diagram for quarks and gluons is considerably more exotic, with a quark-gluon plasma phase, a color superconductor phase, and a phase that consists of a gas of quarks bound together by gluons.

Heavy-ion collisions at RHIC produce quarks and gluons in one of these three phases. By varying the collision energy, RHIC can map out the phase diagram, and potentially find an intriguing predicted critical point beyond which the quark-gluon plasma can smoothly transform into the gas phase. However, theoretical advances to support a comprehensive framework for modeling these heavy-ion collisions will be necessary to interpret the RHIC measurements.

This research will fill in critical theoretical inputs by performing the first state-of-the-art lattice QCD calculations to (i) extend the reach of the equation of state that describes nuclear matter to cover almost the entire beam energy range of the RHIC program; (ii) provide more stringent constraints on the location of the critical point in the phase diagram, vastly improving the present status; and (iii) establish the equilibrium baseline for certain experimentally-measured parameters for a large range of RHIC beam energies.

**Title:** Nucleon structure and electric dipole moments with physical chirally-symmetric quarks

**Principal Investigator:** Sergey Syritsyn (RIKEN BNL Research Center)

**Co-Investigators:** Michael Engelhardt (New Mexico State University, Las Cruces, NM, USA) | Jeremy Green (NIC, Deutsches Elektronen-Synchrotron, Germany) | Taku Izubuchi (RIKEN BNL Research Center, Upton, NY, USA) | Christos Kallidonis (Computation-based Science and Technology Research Center, The Cyprus Institute, Nicosia, Cyprus) | Meifeng Lin (Computational Science Initiative, Brookhaven National Laboratory, USA) | John Negele (Massachusetts Institute of Technology, Cambridge, MA, USA) | Hiroshi Ohki (RIKEN BNL Research Center, Upton, NY, USA) | Shigemi Ohta (RIKEN, Japan) | Andrew Pochinsky (Massachusetts Institute of Technology, Cambridge, MA, USA) | Jun-Sik Yoo (Department of Physics, Stony Brook University, Stony Brook, NY, USA)

**ALCC allocation:**
  **Site:** Argonne Leadership Computing Facility (ALCF)
   **Allocation:** 135,000,000 core hours

## Research Summary

This project will perform a collection of calculations using lattice QCD techniques to shed light on some of the most pressing open questions in nuclear physics. To address a puzzling discrepancy between measurements of normal hydrogen (an electron bound to a proton) and muonic hydrogen (a muon bound to a proton), we will calculate the charge radius of the proton with an uncertainty roughly 80% smaller than the measured discrepancy. To facilitate experimental searches for particles that may explain the inexplicably small masses of neutrinos or constitute Dark Matter, we will calculate nucleon properties that determine the strength of the interaction between nucleons in experimental detectors and these exotic particles. To explore possible reasons for the fact that there is significantly more matter than antimatter in our universe, we will calculate how various hypothetical explanations would affect the distribution of electric charge (specifically, the electric dipole moment or EDM) in protons and neutrons using a highly economical computing scheme that we have developed. These, and other related calculations, will be performed using software that was tested and optimized on the Bluegene/Q and used in other ALCC projects as well as unique tools developed through DOE's INCITE program and computer time at Brookhaven National Laboratory and the University of Edinburgh.

**Title:**　　　Large scale deep neural network optimization for neutrino physics

**Principal Investigator:**　　　Gabriel Perdue (Fermi National Accelerator Laboratory)

**Co-Investigators:**　　　Moulay Akhloufi (University of Moncton) | Adam Aurisano (University of Cincinnati) | Anushree Ghosh (Universidad Tecnica Federico Santa Maria) | Tomasz Golan (University of Wroclaw) | Alexander Himmel (Fermi National Accelerator Laboratory) | Evan Niner (Fermi National Accelerator Laboratory) | Jonathan Miller (Universidad Tecnica Federico Santa Maria) | Robert Patton (Oak Ridge National Laboratory) | Fernanda Psihas (Indiana University) | Alexander Radovic (College of William and Mary) | Dipak Rimal (University of Florida) | Marianette Wospakrik (University of Florida) | Barbara Yaeggy (Universidad Tecnica Federico Santa Maria) | Steven Young (Oak RIdge National Laboratory)

**ALCC allocation:**
　**Site:**　　　Oak Ridge Leadership Computing Facility (OLCF)
　　**Allocation:**　　　58,000,000 core hours

## Research Summary

Neutrinos are fundamental particles that may hold the key to understanding why there is a preference for matter over antimatter in the makeup of the universe today. The 2015 Nobel Prize in Physics was awarded for the discovery of neutrino oscillations, which shows that neutrinos have mass. We seek to establish whether neutrinos and antineutrinos oscillate differently and carry out a precision neutrino-nucleus cross section program.

The 2014 Particle Physics Project Prioritization Panel report highlights the physics of neutrino mass as one of the five Science Drivers for High Energy Physics (HEP). This proposal addresses this Driver with a synergistic neutrino cross section and oscillation physics program. A recent revolution in machine learning powered by new Graphics Processing Units and deep learning algorithms has propelled computers past humans in certain pattern recognition exercises, particularly in computer vision. Modern detectors are effectively imaging devices, and early results indicate that deep learning will significantly improve neutrino experiments by improving reconstruction efficiency and widening the set of accessible event topologies. Furthermore, the massive simulations available in HEP make this an excellent arena to study deep neural network optimization and subjects such as representation transfer and semantic segmentation in a physics context.

This proposal will analyze two complementary neutrino detectors, MINERvA and NOvA, and use deep learning to drive improvements in their physics reach. We will optimize deep neural network performance by conducting an enormous parameter search, using an evolutionary algorithm to evolve solutions using custom software: the Multi-node Evolutionary Neural Networks for Deep Learning (MENNDL) package. We will leverage common simulation software and neutrino interaction processes while contrasting different detectors to analyze network optimization. Addressing all of these questions is only possible with leadership computing.

**Title:** Snapshots into the exotic world of quarks and gluons

**Principal Investigator:** Robert Edwards (Thomas Jefferson National Accelerator Facility)

**Co-Investigators:** Raul Briceno (Thomas Jefferson National Accelerator Facility) | Jozef Dudek (College of William and Mary) | Balint Joo (Thomas Jefferson National Accelerator Facility) | Frank Winter (Thomas Jefferson National Accelerator Facility) | Christopher Thomas (DAMTP, Cambridge University, UK) | Mike Peardon (Trinity College, Dublin, Ireland) | Sinead Ryan (Trinity College, Dublin, Ireland) | David Wilson (Trinity College, Dublin, Ireland)

**ALCC allocation:**
  **Site:** Oak Ridge Leadership Computing Facility (OLCF)
   **Allocation:** 100,000,000 core hours

## Research Summary

Fundamental paritcles called quarks are held together by the strong nuclear force to form the protons and neutrons in the nucleus of every atom. However, the majority of the mass we observe in the visible universe comes not from these quarks but from the "glue" that binds the quarks through the strong nuclear force, described by the theory of Quantum Chromodynamics (QCD).

There are world-wide efforts to determine the spectrum of quarks that form subatomic particles and exploring the spectrum of QCD is a flagship project of the Thomas Jefferson National Accelerator Facility (Jefferson Lab). Mesons are subatomic particles made of one quark and one antiquark held together by the strong nuclear force. The excited meson spectrum is of particular interest because it may reveal the presence of exotic mesons that would directly signal presence of the "glue." The new GlueX experiment at Jefferson Lab has been built to discover and explore properties of exotic meson states.

This project aims to compute the excited-state meson spectrum of QCD. The goal of the computational research is to predict from theory and simulation the masses of possible exotic meson states in advance of the first experimental results from GlueX, which has commenced measurements in 2016. These calculations will guide future experimental searches and advance our understanding of matter and fundamental particle physics.

**Title:** Extreme-scale Simulations for Multi-Wavelength Cosmology Investigations

**Principal Investigator:** Katrin Heitmann (Argonne National Laboratory)

**Co-Investigators:** Salman Habib (Argonne National Laboratory) | Nicholas Frontiere (University of Chicago) | Adrian Pope (Argonne National Laboratory) | Jeffrey Emberson (Argonne National Laboratory) | Hal Finkel (Argonne National Laboratory) | Esteban Rangel (Northwestern University) | Thomas Uram (Argonne National Laboratory)

**ALCC allocation:**
**Site:** Argonne Leadership Computing Facility (ALCF)
  **Allocation:** 40,000,000 core hours
**Site:** Oak Ridge Leadership Computing Facility (OLCF)
  **Allocation:** 10,000,000 core hours
**Site:** National Energy Research Scientific Computing Center (NERSC)
  **Allocation:** 75,000,000 core hours

## Research Summary

During the last two decades, we have built a successful "Standard Model" of Cosmology which describes our Universe and its path from a hot fireball in its earliest moments to the sea of galaxies that we observe today. At the same time, this model poses deep puzzles that demand to be solved: What is the nature of dark matter and what is the origin of dark energy? Together, they account for 95% of the matter-energy content in the Universe but neither are currently understood.

Cosmology has entered a very exciting era where new and up-coming observations unprecedented in size and detail may find answers to these these deep fundamental questions about the nature and make-up of our Universe. In addition, multi-wavelength observations of the large-scale Universe can probe the nature of the inflationary phase in the very early Universe as well as determine the sum of neutrino masses.

Multi-wavelength simulations play a crucial role in answering these questions and providing predictions for different cosmological models and scenarios at an exquisite level of detail. In this project we will build upon some already available simulations and carry out new simulations, ranging from gravity-only to hydrodynamics, to build a suite of multi-wavelength, multi-cosmology synthetic sky maps. These simulations and resulting sky maps will be used to support the analysis of ongoing DOE-funded cosmological surveys, to prepare for upcoming data by creating data challenges, and to plan future observational surveys. This will advance our efforts to understand the nature of dark matter and dark energy.

**Title:** Simulating particle interactions and the resulting detector response at the LHC and Fermilab

**Principal Investigator:** John Taylor Childers (Argonne National Laboratory)

**Co-Investigators:** Thomas LeCompte (Argonne National Laboratory) | Doug Benjamin (Duke University) | Radja Boughezal (Argonne National Laboratory) | Paolo Calafiura (Lawrence Berkeley National Laboratory) | Stefan Hoeche (SLAC National Laboratory) | Burt Holzman (Fermi National Accelerator Laboratory) | Alexei Klimentov (Brookhaven National Laboratory) | Jim Kowalkowski (Fermi National Accelerator Laboratory) | Frank Petriello (Northwestern University) | Vakho Tsulaia (Lawrence Berkeley National Laboratory) | Craig Tull (Lawrence Berkeley National Laboratory) | Thomas Uram (Argonne National Laboratory) | Torre Wenaus (Brookhaven National Laboratory)

**ALCC allocation:**
  **Site:** Argonne Leadership Computing Facility (ALCF)
    **Allocation:** 58,000,000 core hours
  **Site:** Oak Ridge Leadership Computing Facility (OLCF)
    **Allocation:** 80,000,000 core hours
  **Site:** National Energy Research Scientific Computing Center (NERSC)
    **Allocation:** 50,000,000 core hours

## Research Summary

High Energy Physics studies the properties of matter, energy, space and time at the smallest possible distances. Scientists do this by building particle detectors of great size and complexity – in some cases weighing thousands of tons and having millions of readouts. Scientists then use these detectors by comparing computer simulations of what we would expect to see if different theories hold true with the actual data, and seeing which, if any, match. Progress can be made by improving either end of this comparison, and this effort aims to improve the quality of the computer simulation. Using supercomputers, we can produce more simulated events, and simulation events of better accuracy or greater complexity than we could otherwise.

This project supports a team of experimental and theoretical physicists working together to answer fundamental questions for High Energy Physics. We will create an end-station to advance computation on High Performance Computing resources in support of High Energy Physics experiments. The end-station environment will enable other particle physicists to use supercomputers more easily and quickly than they would on their own, to the benefit of the broader particle physics community. The project extends scientific research by producing simulated collision events for the ATLAS experiment at the Large Hadron Collider that are too complex for simulation on the Grid, by improving calculation of theoretical matrix elements used in comparing simulation with data, and by simulating and analyzing data of muon and neutrino experiments based as Fermilab and elsewhere.

**Title:**   Hadronic light-by-light scattering and vacuum polarization contributions to the muon anomalous magnetic moment from lattice QCD with chiral fermions

**Principal Investigator:**   Thomas Blum (University of Connecticut)

**Co-Investigators:**   Peter Boyle (University of Edinburgh) | Norman Christ (Columbia University) | Masashi Hayakawa (University of Nagoya) | Taku Izubuchi (BNL) | Luchang Jin (BNL) | Chulwoo Jung (BNL) | Andreas Juettner (University of Southampton) | Christoph Lehner (BNL) | Antonin Portelli (University of Edinburgh)

**ALCC allocation:**
  **Site:**   Argonne Leadership Computing Facility (ALCF)
    **Allocation:**   220,000,000 core hours

## Research Summary

The Standard Model of particle physics describes the fundamental particles and their interactions and, so far, its predictions have been extrodinarily accurate when tested by experiment. An important precision test of the Standard Model is the calculation and measurement of the magnetic moment of the muon, a heavy cousin of the electron. The measured value of the muon's magnetic moment is sensitive to the presnce of new particles not accounted for in the Standard Model, and new precision experiment at Fermi National Acclerator Laboratory seeks to measure this value to 140 parts per billion. Improvements in the Standard Model prediction of the muon's magnetic moment will enhance our ability to use this new measurement to search for new physics.

This project aims to complete computations of the leading and next-to-leading contributions to the muon's magnetic moment from Quantum Chromodyamics (QCD), the fundamental theory of strong nuclear force that describes interactions between quarks and gluons. These computations are necessary to improve the precision of the Standard Model prediciton for the muon's magnetic moment and enable this important test of the Standard Model.

**Title:**     High-fidelity Numerical Simulation of Wire-wrapped Fuel Assemblies

**Principal Investigator:**     Elia Merzari (Argonne National Laboratory)

**Co-Investigators:**     Aleksandr Obabko (Argonne National Laboratory) | Paul Fischer (Argonne National laboratory) | Brian Jackson (Terrapower) | Mariano Tarantino (ENEA) | Cody Permann (Idaho National laboratory) | Oana Marin (Argonne National Laboratory)

**ALCC allocation:**
  **Site:**     Argonne Leadership Computing Facility (ALCF)
  **Allocation:**     85,000,000 core hours

**Research Summary**

Recent interest in advanced fuel cycles and reactor fuel recycling has led to renewed interest in the development of liquid-metal-cooled fast reactors in the United States. Several private companies have expressed interest in this technology, with a focus on cores with longer life. The prediction of thermal performance in fuel assemblies is vital for evaluating overall reactor performance and safety, particularly as the assembly ages. In particular, deformations driven by thermal expansion and radiation damage is increasingly important as reactor designs look to longer assembly lifetimes. As assemblies deform, traditional heuristic modeling becomes less suitable for evaluating the thermal-hydraulic performance of a given assembly due complex flow patterns that may be introduced. This project has the goal of validating higher fidelity tools to provide accurate predictions of the flow field and heat transfer in wire-wrapped fuel assemblies, typical of sodium fast reactor designs. The data obtained with these simulations will be instrumental in further improving lower –fideility models and understanding the flow physics in this complex geometry.

**Title:** Numerical Simulation of Turbulent Flows in Advanced Steam Generators - Year 3

**Principal Investigator:** Aleksandr Obabko (Argonne National Laboratory)

**Co-Investigators:** Elia Merzari (Argonne National Laboratory) | Paul Fischer (Argonne National Laboratory) | Matthew Snyder (NuScale Power) | Yassin Hassan (Texas A&M University) | Robert Ferencz (Lawrence Livermore National Laboratory) | Jerome Solberg (Lawrence Livermore National Laboratory)

**ALCC allocation:**
  **Site:** Argonne Leadership Computing Facility (ALCF)
    **Allocation:** 50,000,000 core hours

## Research Summary

Safer nuclear energy power promises to become a reliable, carbon-free resource capable of meeting the nation's and the world's energy needs. Modeling and simulation traditionally used by the nuclear engineering community has been based on heuristic modeling. DOE programs such as NEAMS (DOE-NE), CESAR (DOE-ASCR), and exascale (ECP-ASCR) have invested in state-of-the-art tools for predictive, first-principles-based numerical modeling of turbulent flows and related phenomena. The NEAMS program has recently partnered with industry to address a topic of great interest and high impact: the numerical simulation of flow-induced vibrations in steam generators (SGs), which are an essential component of all pressurized water reactor (PWR) designs. The emphasis on reliable SGs has increased in recent years with the growing interest in small modular reactors (SMRs). Modeling and simulation is expected to play an important role in assessing the reliability of advanced steam generators, especially for what concerns complex phenomena such as flow induced vibration. This project, with its strong focus on validation, is a stepping stone toward achieving that goal.

**Title:**    CASL Gamma Heating Verification Analysis

**Principal Investigator:**    Tara Pandya (Oak Ridge National Laboratory)

**Co-Investigators:**    Kevin Clarno (Oak Ridge National Laboratory) | Ben Collins (Oak Ridge National Laboratory) | Thomas Evans (Oak Ridge National Laboratory) | Andrew Godfrey (Oak Ridge National Laboratory) | Steven Hamilton (Oak Ridge National Laboratory) | Seth Johnson (Oak Ridge National Laboratory) | Kang Seog Kim (Oak Ridge National Laboratory)

**ALCC allocation:**
  **Site:**    Oak Ridge Leadership Computing Facility (OLCF)
    **Allocation:**    40,000,000 core hours

**Research Summary**

The Consortium for Advanced Simulation of Light Water Reactors (CASL) is a U.S. Department of Energy (DOE) Innovation Hub, established in July 2010 for the modeling and simulation of nuclear reactors. In 2016, the VERA Core Simulator (VERA-CS) was used to successfully simulate and provide high fidelity solutions to quantify the uncertainty in predicting the onset of Chalk River Unidentified Deposits (CRUD) growth and therefore allowed for more optimal operating conditions for current U.S. reactors. In 2017, CASL will use VERA-CS coupled to Shift, a Monte Carlo neutronics solver, to verify and validate the gamma heating distribution for a wide variety of operational states and accident conditions. These calculations will allow for determining the errors in the lower-order approximations used in VERA-CS and will enable predictive simulation of these operational states and accident conditions. Along with this analysis, CASL will also perform high-fidelity vessel fluence calculations using excore modeling with Shift and VERA-CS.

**Title:**        Quantifying Systematic Error in Monte Carlo Simulations of Nuclear Reactors

**Principal Investigator:**        Paul Romano (Argonne National Laboratory)

**Co-Investigators:**        Christopher Perfetti (Oak Ridge National Laboratory) | Benoit Forget (Massachusetts Institute of Technology) | Andrew Siegel (Argonne National Laboratory)

**ALCC allocation:**
  **Site:**        National Energy Research Scientific Computing Center (NERSC)
   **Allocation:**        15,000,000 core hours

## Research Summary

The accurate simulation of neutral particle transport is of vital importance to a range of physical disciplines, including, for example, nuclear reactor analysis, medical imaging, radiation detection, fusion energy, and weapons design. While a variety of methods exist for solving the governing equations of particle transport, the Monte Carlo (MC) method is considered the most accurate because it requires very few approximations. When applied to fissioning systems (e.g., nuclear reactors) however, an error in the solution, known as an undersampling bias, can occur if too few particles are used. This project seeks to further our understanding of undersampling biases as they occur in MC simulations of large light-water reactors, which are the predominant reactor type in the U.S. The work will also apply to small modular reactors, a promising new class of reactors with key advantages in safety, economics, and manufacturability. Through large simulations of two reactor models using the OpenMC code, our work will quantify undersampling biases in local physical quantities, thus helping to establish best practices for large-scale MC simulations.

**Title:**     Elimination of Modeling Uncertainties through High-Fidelity Multiphysics
               Simulation to Improve Nuclear Reactor Safety and Economics

**Principal**              Emily Shemon (Argonne National Laboratory)
**Investigator:**

**Co-Investigators:**      Taek Kim (Argonne National Laboratory) | Yiqi Yu (Argonne
                           National Laboratory)

**ALCC allocation:**
  **Site:**                Argonne Leadership Computing Facility (ALCF)
    **Allocation:**        44,000,000 core hours

## Research Summary

This project will develop a series of high-fidelity multiphysics calculations of fast spectrum nuclear reactor configurations. These calculations are needed to eliminate modeling uncertainties in the calculation of important safety parameters, specifically "hot channel factors".  The use of high fidelity  multiphysics methods can reduce or eliminate modeling uncertainties due to more accurate geometry and physics representations of the system. Consequently, such simulations could dramatically improve reactor economics by allowing the reactor to operate at higher power while still maintaining safety margins from the design limit. In order to reduce the significant computational cost of using high-fidelity tools on explicitly modeled full-core reactor geometries, the project will also explore the use of high fidelity tools as a "zooming" tool. In the "zooming" concept, local regions of interest are represented with full fidelity and other regions are represented with much coarser geometry representation. This approach can reduce the overall calculation cost and make high fidelity tools much more computationally feasible for a larger array of problems. All calculations will be performed with the DOE NEAMS-developed SHARP Toolkit which includes the PROTEUS neutronics code and the Nek5000 thermal hydraulics/computational fluid dynamics code. This project has high relevance to DOE mission because it enables fast reactor designs that benefit the nuclear fuel cycle. The science goals of this project are directly requested by the DOE NEAMS program under the "SFR Grand Challenge Problem Workpackage."

**Title:** Portable Application Development for Next Generation Supercomputer Architectures

**Principal Investigator:** Tjerk Straatsma (Oak Ridge National Laboratory)

**Co-Investigators:** Katie Antypas (Lawrence Berkeley National Laboratory) | Timothy Williams (Argonne National Laboratory)

**ALCC allocation:**
  **Site:** Argonne Leadership Computing Facility (ALCF)
    **Allocation:** 20,000,000 core hours
  **Site:** Oak Ridge Leadership Computing Facility (OLCF)
    **Allocation:** 20,000,000 core hours
  **Site:** National Energy Research Scientific Computing Center (NERSC)
    **Allocation:** 20,000,000 core hours

## Research Summary

The ALCF, NERSC and OLCF supercomputing facilities have created early application readiness programs that prepare a range of application codes for the next-generation supercomputers on the path to exascale. These programs are: the Early Science Program (ESP) at ALCF, the NERSC Exascale Science Application Program (NESAP), and the Center for Accelerated Application Readiness (CAAR) at OLCF. The three facilities have launched these application readiness programs with approximately 40 applications teams participating. In order to explore portability issues and ensure the success of the application readiness program, it is crucial that application developers have allocations not only at their sponsored supercomputing facility, but also at facilities with alternate hardware architectures. This project will support the coordination and collaboration within the three facilities and provide computational resources to the application teams, with the overall objective of delivering applications that are high performing and architectural and performance portable.

**Title:**     Demonstration of the Scalability of Programming Environments By Simulating Multi-Scale Applications

**Principal Investigator:**     Robert Voigt (Leidos)

**Co-Investigators:**     Bala Balachandar (U. of Florida) | Jonathan Freund (U. of Illinois) | Karel MatouÅ¡ (U. of Notre Dame) | Lawrence Rauchwerger (Texas A & M U.) | Gianluca Iaccarino (Stanford U.) | Martin Berzins (U. of Utah)

**ALCC allocation:**
  **Site:**          Argonne Leadership Computing Facility (ALCF)
   **Allocation:**     110,000,000 core hours
  **Site:**          Oak Ridge Leadership Computing Facility (OLCF)
   **Allocation:**     46,800,000 core hours

## Research Summary

PSAAP II is a five-year program established in 2014 by the Advanced Simulation and Computing (ASC) program of the National Nuclear Security Administration to demonstrate predictive science in an extreme-scale computing environment. Five centers at U. Florida, U. of Illinois at Urbana-Champaign, Stanford, U. Notre Dame, Texas A&M, and U. Utah will focus on a diverse set of complex applications of importance to DOE.  This project will access ALCF and OLCF computing resources to demonstrate that the technologies and methodologies developed by these centers will effectively scale to extreme scale computing, simulating target applications with full physics and required resolution. The centers expect to gain insight into issues critical to the advancement to exascale computing, exposing areas where additional software research and development is needed on the path to exascale.

**Title:**  High fidelity gyrokinetic study of divertor heat-flux width and pedestal structure

**Principal Investigator:**  Choong-Seock Chang (Princeton Plasma Physics Laboratory)

**Co-Investigators:**  Stephen Abbott (ORNL) | Mark Adams (LBNL) | Ed D'Azevedo (ORNL) | Stephane Ethier (Princeton Plasma Physics Laboratory) | Robert Hager (PPPL) | Scott Klasky (ORNL) | Tuomas Koskela (NERSC) | Seung-Hoe Ku (PPPL) | Scott Parker (U. Colorado, Denver) | Mark Shephard (RPI) | Pat Worley (PHWorley Consulting)

**ALCC allocation:**
  **Site:**  Argonne Leadership Computing Facility (ALCF)
    **Allocation:**  80,000,000 core hours
  **Site:**  Oak Ridge Leadership Computing Facility (OLCF)
    **Allocation:**  100,000,000 core hours
  **Site:**  National Energy Research Scientific Computing Center (NERSC)
    **Allocation:**  89,900,000 core hours

## Research Summary

ITER is an international partnership to build the world's largest fusion experiment. The goal of ITER is to prove the feasibility of fusion as a large-scale carbon free energy source. One key challenge for ITER's succss is to build a device that can withstand the steady exhaust heat deposited by the hot plasma on a narrow strip of the divertor target plates. Accurate first-principles based simulations are needed to maximally inform ITER's operation and overcome the heat deposition challenges. The projections upon which ITER operation is currently based are extrapolations of existing experimental data or of simple reduced models and are made without a fundamental understanding of the underlying physical mechanisms, and are too far to be considered reliable. Realistic simulations at fundamental physics level can provide the needed insight into these complicated issues. The edge plasma is complicated by the presence of the magnetic X-point, being incontact with the material wall, and coupled to the neutral atoms and molecules resulting from that contact. The 5D gyrokinetic particle code XGC1 has been built to study such edge plasmas, but requires a large amount of CPU time on extreme-scale leadership class computers. This allocation will support the advancement of an ongoing study, using programming code XGC1 on Titan, Theta, and Cori, of the fundamental physics underlying the heat doposition on the divertor target plates and the related pedestal structure issue  in present tokamak devices and in ITER. Outcomes of this project will be higher confidence extrapolations of theory to practice, and increase the success likelihood of ITER.

**Title:**      Studying astrophysical particle acceleration in HED plasmas

**Principal Investigator:**      Frederico Fiuza (SLAC National Accelerator Laboratory)

**Co-Investigators:**      Paulo Alves (SLAC National Accelerator Laboratory) | Warren Mori (UCLA) | Tom Abel (Stanford University)

**ALCC allocation:**
  **Site:**      Argonne Leadership Computing Facility (ALCF)
   **Allocation:**      50,000,000 core hours

**Research Summary**

Beyond its relevance to the understanding of the extreme universe, the study of particle acceleration has a significant impact on generating new laboratory accelerators for a variety of applications, from diagnostics for fusion plasmas to medical imaging. A transformative advance in our understanding of the process of particle acceleration in plasmas requires a combination of fully kinetic simulations and controlled laboratory experiments where numerical findings can be tested. This project will perform unprecedented massively parallel particle-in-cell simulations to explore the physics of particle acceleration and magnetic field dynamics in High-Energy-Density (HED) plasmas, associated with collisionless shocks and magnetic reconnection, for conditions of high-power laser experiments already planned at the National Ignition Facility (NIF). The results of this project will support NIF experiments that will not succeed without 3D fully kinetic simulations on Mira and will dramatically change the ability of scientists to understand particle injection and acceleration in laboratory and astrophysical plasmas.

**Title:**        Integrated Simulation of Energetic Particles in Burning Plasmas

**Principal Investigator:**        Zhihong Lin (University of California, Irvine)

**Co-Investigators:**        Guoyong Fu (Princeton Plasma Physics Laboratory)

**ALCC allocation:**
  **Site:**        Oak Ridge Leadership Computing Facility (OLCF)
   **Allocation:**        50,000,000 core hours

## Research Summary

ITER is an international partnership to build the world's largest fusion experiment. The goal of ITER is to prove the feasibility of fusion as a carbon-free, abundant energy source. The confinement of energetic particles (EP) is a critical issue for burning plasma experiments since the ignition in ITER fusion reactor relies on the self-heating by energetic fusion products ($\alpha$-particles). Plasma confinement property in burning plasmas is one of the most uncertain issues when extrapolating from existing fusion devices to ITER. The fusion community has made impressive progress in developing comprehensive EP simulation codes and understanding key EP physics. In particular, the gyrokinetic turbulence simulation has been successfully established as a necessary paradigm shift in the energetic particle studies. Nonetheless, more coordinated efforts and advanced computing hardware and software are urgently needed for first-principles, integrated simulations incorporating multiple physical processes and disparate temporal-spatial scales to build the EP predictive capability for ITER.

The Integrated Simulation of Energetic Particles (ISEP) project aims to improve physics understanding of energetic particle (EP) confinement and EP interactions with burning thermal plasmas through large-scale simulations. This allocation supports ISEP development of a comprehensive predictive capability for EP physics and an EP module incorporating both first-principles simulation models and high fidelity reduced transport models to the fusion whole device modeling (WDM) project. This project will advance predictive capabilities of the fusion community and supports the success of ITER.

**Title:**     Modeling Helium-Hydrogen Plasma Mediated Tungsten Surface Response to Predict Fusion Plasma Facing Component Performance

**Principal Investigator:**     Brian Wirth (Oak Ridge National Laboratory)

**Co-Investigators:**     Blas Uberuaga (LANL) | Aidan Thompson (SNL) | David Bernholdt (ORNL)

**ALCC allocation:**
  **Site:**     Argonne Leadership Computing Facility (ALCF)
    **Allocation:**     98,000,000 core hours
  **Site:**     Oak Ridge Leadership Computing Facility (OLCF)
    **Allocation:**     75,000,000 core hours

**Research Summary**

The realization of fusion as a practical, 21st Century energy source requires improved knowledge of plasma surface interactions and the materials engineering design of component systems to survive the incredibly extreme heat and particle flux exposure conditions of a fusion power plant. The objective of this project is to further advance understanding of the response of tungsten, the proposed ITER diverter, to low energy, mixed H-He plasma exposure. In particular, two tasks are envisioned that investigate helium behavior and gas bubble aggregation kinetics in tungsten, and begin to address hydrogen retention in sub-surface gas bubbles. This allocation supports an extensive series of large-scale atomistic molecular dynamics simulations to investigate H diffusion and permeation through tungston surfaces containing a strained-network of over-pressurized helium bubbles, as well as an evaluation of different interatomic ptoentials to describe the bonding interactions in tungston-Helium-Hydrogen, and to enable new, higher fidelity potentials. This work will improve understanding of gas recycling and total tritium inventory in fusion plasma facing components.