Title: Multiphase Simulations of Nuclear Reactor Flows

Principal Investigator: Igor Bolotnov (North Carolina State University)

Co-Investigators: Michel Rasquin (Cenaero and University of Colorado at Boulder), Kenneth Jansen (University of Colorado at Boulder)

ALCC allocation: Processor Hours
Site: Argonne Leadership Computing Facility (ALCF)
Allocation: 72,100,000 processor hours

Research Summary

The design of new nuclear reactors, and the safe, efficient operation of existing reactors, can benefit from fundamental understanding of the bubbly two-phase flows created as the water boils. The most accurate technique for these flows, Direct Numerical Simulation (DNS), captures all the length scales of turbulence in the flow. DNS of turbulent two-phase flows at a leadership computing facility allows achievement of unprecedented level of detail and can answer fundamental questions about the interaction between the complex and evolving interfaces of the bubbles and droplets and turbulence in the flow.

The detailed simulation of all turbulent structures using DNS approach as well as interface evolution in turbulent flows (in form of bubbles, droplets, wavy interfaces between liquid and steam in various flow regimes) will allow the collection of statistical information about two-phase flow parameters over an unprecedented range of conditions. Advanced data collection coupled with interface tracking will be used to process the large data sets.

Three subprojects will generate two-phase flow results highly relevant to nuclear reactor flows: (i) bubbly flow through pressurized water nuclear reactor fuel bundles with complex geometry spacers, (ii) complex two-phase flow regimes (slug and churn flow) as well as (iii) annular flows to the highest level of detail. Major parameters will be compared to experimental data for validation, while the higher order statistics produced by the simulations will bring new knowledge about two-phase flows, such as spectral information about gas/liquid interaction, bubbles and droplets contribution to the turbulence.

In this project, the complex physical phenomena will be captured in unprecedented detail. The level set interface tracking method and on-the-fly mesh adaptation will be used to improve locally the mesh resolution of phase interface without the need to store intermediate files. In-situ co-processing of the data will also be used to accelerate the analysis. Advanced analysis tools will allow tracking every bubble and collect information about their behavior. This approach will allow partitioning and recognizing different patterns and phase interactions for two-phase flows in complex geometries.

This data processing will generate information about vapor/liquid interaction, enabling the creation of new multiphase fluid models and closure laws that can be used in engineering simulations, allowing reliable application of multiphase computational fluid dynamics (M-CFD) to nuclear reactor systems. This will improve the safety margin predictions for existing reactors and facilitate the development and design of next generation systems.

Title: Computational engineering of defects in soft and hard materials for energy and quantum information applications

Principal Investigator: Marco Govoni (University of Chicago and Argonne National Laboratory)

Co-Investigators: Alex Gaiduk (University of Chicago), Hosung Seo (University of Chicago), Christopher Knight (Argonne National Laboratory), He Ma (University of Chicago), Giulia Galli (University of Chicago), Francois Gygi (University of California Davis)

ALCC allocation: Processor Hours

Site: Argonne Leadership Computing Facility (ALCF)

Allocation: 53,700,000 processor hours

Research Summary

Defects are ubiquitous in matter. Defects are often considered as imperfections to be avoided since they show unexpected and peculiar properties, which could degrade the entire material's performance even at very small concentration. However, the unusual properties of defects can be used to design new functionalities that are not present in their host materials.

This project supports the study of defects in water and solid materials to understand how the defects impact electronic properties. Specifically ,this project supports the application of large-scale quantum simulation methods to aqueous solutions (ions in water as defects) and solid-state defect quantum bits (qubits), pertaining to renewable energy applications and quantum computation. The study will employ *ab-initio* molecular dynamics simulations to compute ensemble averages and thermodynamic properties of the defects from atomic trajectories. The study will also use many body perturbation theory to compute accurate spectroscopy signatures of the defects. To enable these calculations, we developed highly scalable codes (Qbox and WEST) that are capable of tackling systems of unprecedented size (several thousands of electrons).

The specific aims of the project are: (i) to provide knowledge and computational tools to interpret the large body of current experiments on solar-powered fuel production from aqueous solutions; and (ii) to establish design rules to predict robust defect qubits in solid-state environments, in which coherent qubit control, strong qubit-lattice coupling, and device scalability could be all achieved simultaneously.

Due to the similarities in methodologies involved by defects in soft and hard materials, advances achieved in one system can directly be applied to many other cases, and solving computational and physical challenges that arise will lead to breakthroughs that carry over to other fields of science and engineering.

The outcome of this proposal will be an improved understanding of the role of defects in hard and soft matter, which is critical for the design of new functional materials.

Title: Extreme scale gyrokinetic particle simulations to complete the 2016 OFES National Theory/Simulation Performance Target and to study the fundamental Edge Physics

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

Co-Investigators: Scott Parker (University of Colorado at Boulder), Mark Adams (Lawrence Berkeley National Laboratory), Seung-Hoe Ku (Princeton Plasma Physics Laboratory), Robert Hager (Princeton Plasma Physics Laboratory), Mark Shephard (Rensselaer Polytechnic Institute), Scott Klasky (Oak Ridge National Laboratory), Pat Worley (Oak Ridge National Laboratory)

ALCC allocation: Processor Hours

Site: Argonne Leadership Computing Facility (ALCF)
Allocation: 100,000,000 processor hours
Site: Oak Ridge Leadership Computing Facility (OLCF)
Allocation: 75,000,000 processor hours

Research Summary

ITER is an international research project to build an experimental fusion reactor. The success of ITER would mark an important step towards creating a sustainable carbon-free source of fusion energy. A major goal of ITER is to demonstrate the generation of 500 MW of fusion power from a 50 MW input. One of the biggest challenges for ITER will be to control the fusion plasma. Fusion reactions generate massive amounts of heat. The heat from the fusion reaction in ITER will cause the plasma inside ITER to reach temperatures as hot as the Sun, 150 million degrees Celsius. ITER will use strong magnets to contain the extremely hot plasma yet edge plasma may still interact with the surface of ITER and create a turbulent environment where edge plasma behavior is difficult to predict and control. Controlling the edge plasma behavior is necessary to create a sustained, controlled, and confined fusion process. Edge plasma behavior is therefore critical to ITER's success.

First, ITER must be able to reliably withstand the steady exhaust heat deposited in an extremely narrow strip on the divertor target plates. Second, ITER edge plasma must be able to access physical conditions that create suitable edge plasma behavior known as the high confinement mode of operation (L-H transition). Here the edge plasma supports confinement by forming a barrier called the 'edge pedestal'. Third, instabilities that periodically destroy the edge plasma pedestal must be avoided or mitigated.

To optimize the chances for ITER's success in controlling the high energy edge plasma, this project supports computational simulations to understand edge plasma physics using highly accurate extreme-scale edge gyrokinetic code XGC1, developed for this purpose by collaborative SciDAC activities between computer and fusion scientists supported by DOE's ASCR and FES programs.

Title: Unraveling Silent Owl Flight to Develop Ultra-Quiet Energy Conversion Machines

Principal Investigator: Anupam Sharma (Iowa State University)

Co-Investigators: n/a

ALCC allocation: Processor Hours
Site: Argonne Leadership Computing Facility (ALCF)
Allocation: 25,000,000 processor hours

Research Summary

Acoustic emission (noise) from wind turbines is curtailing the growth of wind energy, which is currently the primary renewable energy source in the US and in the world. A majority of the noise radiated from wind turbines is generated aerodynamically – due to interaction of wind with blade surfaces. Aerodynamic sound (aeroacoustics) is an issue not just for wind turbines but also for aircraft, jet engines, combustion turbines used for electricity generation, cooling fans, and ventilation systems.

A solution to the problem of aerodynamic noise generation is available in nature but has not yet been leveraged to develop silent machines. The nocturnal owl is known to have a silent flight both when gliding and flapping. This has been known for decades, but the physical mechanisms enabling its silent flight are not well understood. Previous investigations have identified three feather features that are unique to the owl. Experimental investigations have demonstrated that these unique feather features are responsible for the owl's acoustic stealth. However, these experiments alone are unable to identify the reasons/mechanisms behind noise reduction. This is because it is nearly impossible to measure the flow with the spatial and temporal accuracy required to fully understand these mechanisms.

This project supports very high resolution simulations to bridge the scientific gap between experimental results and theoretical understanding. A systematic numerical investigation of the unique owl feather features is proposed to answer key questions that will help unravel the mystery behind owl's silent flight. The extremely high spatial and temporal resolution offered by high-fidelity numerical simulations will enable source diagnostics to identify how the unique feather features curb noise generation. The knowledge and understanding gained from these simulations can empower us to design nearly silent energy conversion- and various other engineering machines.

The simulations will use a well-established high-order accurate flow solver, FDL3DI, and an in-house acoustics solver. This computational framework has been verified against experimental data for a model aeroacoustics problem. The simulations results will be combined with ongoing experimental measurements at Virginia Tech to allow a comprehensive understanding of silent owl flight and provide a transformative jump towards developing nearly silent energy conversion machines.

Title: The Materials Project - Completing the Space of Elastic and Piezoelectric Tensors

Principal Investigator: Kristin Persson (Lawrence Berkeley National Laboratory)

Co-Investigators: Maarten De Jong (University of California, Berkeley), Patrick Huck (Lawrence Berkeley National Laboratory), Anubhav Jain (Lawrence Berkeley National Laboratory), Miao Liu (Lawrence Berkeley National Laboratory)

ALCC allocation: Processor Hours

Site: Argonne Leadership Computing Facility (ALCF)

Allocation: 36,000,000 processor hours

Research Summary

The Materials Project is an effort to collect, catalogue, and organize materials properties in a publicly available online infrastructure. In order to further the collection of materials properties, this project support automated high-throughput workflows that perform and analyze results of electronic structure simulations. In addition, by designing open and efficient user interfaces to our 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 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: 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), Christopher Knight (Argonne National Laboratory), Christopher Mundy (Pacific Northwest National Laboratory), Neeraj Rai (Mississippi State University), David Sholl (Georgia Institute of Technology), Randall Snurr (Northwestern University), Donald Truhlar (University of Minnesota), Yongchul Chung (Pusan National University)

ALCC allocation: Processor Hours
Site: Argonne Leadership Computing Facility (ALCF)
Allocation: 117,000,000 processor hours

Research Summary

This project supports work to develop a functional understanding of nanoporous materials important to industrial clean energy production. An interdisciplinary, collaborative team will use Argonne's Leadership Computing Facility to perform predictive modeling studies that will accelerate the discovery and design of nanoporous materials with tailored functions for a variety of energy-related applications. Specific applications include nanoporous membranes for separation of C8 aromatics, for second-generation biofuel production, for purification of diols as renewable feedstock compounds for high-value polymers, and for separation of light gases. The research will also tackle uncertainty quantification for large-scale screening studies, reactive equilibria in compressed vapors and in nanoporous confinement, hydration forces and ion distribution involved in nanoparticles assembly, and halide perovskites for photovoltaics. In addition, this research will contribute to the development of a computational infrastructure for screening and design of new heterogeneous catalysts and associated scaffolds.

Title: High-Intensity Multibunch Physics in the Fermilab Accelerator Complex

Principal Investigator: James Amundson (Fermilab)

Co-Investigators: Robert Ainsworth (Fermilab), Qiming Lu (Fermilab), Alexandru Macridin (Fermilab), Eric Stern (Fermilab)

ALCC allocation: Processor Hours
Site: Argonne Leadership Computing Facility (ALCF)
Allocation: 50,000,000 processor hours

Research Summary

The U.S. Fermilab accelerator complex is under transition to become the science facility with the world's highest intensity proton beams. In the near term, the complex will be upgraded to provide 700 kW proton beams to support the current and near-future generations of neutrino and muon experiments. In the medium term, the U.S. is dedicated to building a large-scale long-baseline neutrino experiment to pursue the physics of neutrino mass. An international collaboration has been formed to build the world's largest neutrino detector, the Deep Underground Neutrino Experiment (DUNE). The neutrino beam for DUNE will be provided by a Long Baseline Neutrino Facility to be built at Fermilab. The Fermilab PIP-II project will upgrade the accelerator complex to provide the 1000 kW proton beam to power LBNF and DUNE. The sum of these projects are the highest priority for the future of the U.S. experimental particle physics program. The upgrades to the Fermilab accelerator complex for both 700 kW in the short term and 1000 kW in the long term rely on using the Fermilab Recycler storage ring to accumulate beam for acceleration in the Main Injector. The procedure to accumulate beam, called "slip-stacking," involves slowly combining two 80-bunch batches of injected beam into one high intensity batch. The process is sensitive to the collective interactions between the 160 bunches involved.

This project supports simulations of the full 80-on-80 bunch slip stacking process in the Recycler in support of the accelerator upgrade process. This is a very large computational challenge. It will require the use of our accelerator simulation package Synergia, which has a unique multibunch simulation capability, and running on leadership-class computing resources. The results of the detailed simulations will allow Fermilab to reduce both costs and risks as it moves forward to support the U.S. particle physics program.

Title: Modeling of Intense X-ray Laser Dynamics in Nanoclusters

Principal Investigator: Phay Ho (Argonne National Laboratory)

Co-Investigators: Christopher Knight (Argonne National Laboratory), Christoph Bostedt (Argonne National Laboratory), Linda Young (Argonne National Laboratory)

ALCC allocation: Processor Hours

Site: Argonne Leadership Computing Facility (ALCF)

Allocation: 10,000,000 processor hours

Research Summary

X-rays are light waves with wavelengths on the size of atoms to biological cells. Because of their atom to cell sized wavelengths, X-rays have been an invaluable probe for science on the nano-scale. Key to using X-rays is the understanding of how X-rays interact with matter. This project supports the development of a fundamental understanding of the dynamical processes induced by intense x-rays in complex systems. The ultrahigh intensity regime of x-ray interactions with matter was initiated with the advent of the LCLS x-ray free-electron laser (XFEL). The intensities in both the soft and hard x-ray regime can strip electrons from atoms from the inside out and create transient states of matter as the x-ray pulse progresses through the target. Using nanoscale clusters, particular attention is paid to the ionization dynamics starting on the atomic level, the nanoplasma formation, and the implications of the rapidly changing sample to the ultrafast x-ray scattering process. A hybrid quantum/classical simulation methodology will be used to investigate the interactions of intense XFEL pulses with nanoscale samples to understand the ultrafast dynamics and its connection to x-ray imaging. This work is proposed is an important step forward in understanding high-brightness, coherent x-ray laser pulses and their interactions with matter. The completion of the proposed work will enhance capabilities 10,000-fold and establish the applied methodology as an effective large-scale computational tool for the new research frontier of ultrafast x-ray science.

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

Principal Investigator: Robert Voigt (Leidos)

Co-Investigators: S. Balachandar (University of Florida), Jonathan Freund (University of Illinois), Karel Matouš (University of Notre Dame), Lawrence Rauchwerger (Teas A&M University), Gianluca Iaccarino (Stanford University), Martin Berzins (University of Utah)

ALCC allocation: Processor Hours

Site: Argonne Leadership Computing Facility (ALCF)
Allocation: 151,000,000 processor hours
Site: Oak Ridge Leadership Computing Facility (OLCF)
Allocation: 40,000,000 processor hours

Research Summary

Future computing systems pose new challenges for current scientific software and applications. These systems will have extremely high node counts with heterogeneous cores. This project supports development of technology and methods to support effective extreme computing for future computing systems. The project brings together six different science application areas, each of which focuses on three dimensional problems spanning a large range of scales in space and time. The application areas includes turbulence, radiation transport, electric power generation with carbon capture, plasma-coupled combustion, and shock wave-processing of advanced reactive materials. The approaches for scaling codes on advanced architectures will be evaluated by demonstrating their effectiveness on the chosen applications; at the same time our understanding of some key areas of science and engineering will be advanced.

Title: Nuclear Fission: from more phenomenology and adjusted parameters to more fundamental theory and increased predictive power

Principal Investigator: Aurel Bulgac (University of Washington)

Co-Investigators: Ionel Stetcu (Los Alamos National Laboratory), Nicolas Schunck (Lawrence Livermore National Laboratory), Kenneth Roche (Pacific Northwest National Laboratory), Piotr Magierski (Warsaw University of Technology), Gabriel Wlazłowski (Warsaw University of Technology)

ALCC allocation: Processor Hours

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 40,000,000 processor hours

Research Summary

Nuclear fission is the process in which an atom's nucleus splits into smaller fragments including other smaller atoms. Fission is one of the most complex quantum phenomena in physics. The present description of nuclear fission is based largely on observation and not on a fundamental understanding of the physical processes. Still waiting for a microscopic description since 1939, nuclear fission is relevant to an impressive range of fields: basic energy needs, nuclear waste disposal, national security, stockpile stewardship, nuclear forensics, fundamental problems in the non-equilibrium dynamics of strongly interacting quantum many-body systems, nuclear reactions, understanding in astrophysics the nucleosynthesis of about half of the elements heavier than iron produced in the r-process in core-collapse supernovae and neutron star mergers. These questions can be answered only by a microscopic theory based on fundamental properties of nuclear forces, since important fission observables cannot be measured in the laboratory. The information content of the wave function of a fissioning nucleus is immense and it cannot be stored in a classical digital computer, which would be larger than the size of the Universe. The alternative is the Density Functional Theory, which has been implemented on leadership class supercomputers. This project supports real-time nuclear fission simulations will provide crucial and reliable information for basic science and energy needs, national security, stockpile stewardship, nuclear forensics, nuclear waste disposal, and other applications. This project will bring together the efforts at LANL, LLNL, PNNL, and universities to develop a fundamental understanding of nuclear fission.

Title: Characterization of Microbial Carbon Cycling in Terrestrial Ecosystems using Big Proteogenomics Data

Principal Investigator: Chongle Pan (Oak Ridge National Laboratory)

Co-Investigators: Jill Banfield (University of California Berkeley), Melanie Mayes (Oak Ridge National Laboratory)

ALCC allocation: Processor Hours

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 9,000,000 processor hours

Research Summary

Terrestrial microbial communities play critical roles in global carbon cycling, but are poorly represented in climate models. The activities of microbial communities can be understood by studying their genetics and their expression products using metagenomics (reconstruction of microbial genomes) and metaproteomics (analysis of the protein products of genes). The increased throughput of metagenomics and metaproteomics has enabled systematic sampling and measurement of a large number of soil microbial communities across space and time in a terrestrial ecosystem. However, a large amount of sequencing data and mass spectrometry data generated in these metagenomics and metaproteomics studies poses a tremendous computational challenge. This project supports using a set of scalable algorithms on Titan to perform high-resolution data analytics for these integrated omics experiments. The results will provide a mechanistic understanding of the carbon turnover processes by soil microbial communities and can be used to inform climate models.

Title: Delivering 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), Robert Jacob (Argonne National Laboratory), Mark Taylor (Sandia National Laboratory), Philip Jones (Los Alamos National Laboratory), Philip Rasch (Pacific Northwest National Laboratory), Dean Williams (Lawrence Livermore National Laboratory), Patrick Worley (Oak Ridge National Laboratory), Todd Ringler (Los Alamos National Laboratory), William Riley (Lawrence Berkeley National Laboratory)

ALCC allocation: Processor Hours

Site: Argonne Leadership Computing Facility (ALCF)
Allocation: 158,000,000 processor hours
Site: Oak Ridge Leadership Computing Facility (OLCF)
Allocation: 53,000,000 processor hours

Research Summary

The ACME (Accelerated Climate Modeling for Energy) project builds on the National Laboratory system's expertise in climate, high performance computing and data to create a next-generation Earth system model for use by the scientific community. The priority of this project is to improve predictive understanding of the Earth's climate system, with a focus on the coupled system dynamics linking the atmosphere, oceans, land, and ice. The first phase of the project uses focused model development and a carefully-configured set of simulation experiments to address three primary scientific questions:

1) How do the hydrological cycle and water resources interact with the climate system on local to global scales?
2) How do biogeochemical cycles interact with global climate change?
3) How do rapid changes in glaciers, snow cover, ice sheets, and other forms of frozen water (collectively referred to as the cryosphere) interact with the climate system?

This effort concentrates the efforts of climate, computational, software, and data experts across seven National Laboratories and other key partners to target a high-resolution coupled simulation capability over a nominal 80-year window centered on the present-day (approximately 1970-2050). In parallel to the implementation of the Earth system model, domain experts will continually update the atmospheric, ocean, and, and ice sub-models to ensure that these sub-models, and the coupled Earth system model, incorporate the most recent science. This effort will allow ACME to deliver a landmark capability in high-resolution Earth system modeling

Title: Ab initio modeling of the dynamical stability of HED plasmas: from fusion to astrophysics-

Principal Investigator: Frederico Fiuza (SLAC National Accelerator Laboratory)

Co-Investigators: Warren Mori (University of California Los Angeles), Eduardo Paulo Alves (SLAC National Accelerator Laboratory)

ALCC allocation: Processor Hours
Site: Argonne Leadership Computing Facility (ALCF)
Allocation: 60,000,000 processor hours

Research Summary

Fusion energy is regarded as a possible long-term energy solution for humanity, capable of providing the energy resources to drive economic growth and social development. Fusion reactions release so much energy, the fusion material becomes plasma. In the inertial confinement fusion (ICF) approach, a small pellet of fuel is hit with a high energy laser. The laser instigates a fusion reaction inside the pellet. The reaction creates a plasma that if sufficiently hot, would cause *ignition*, a chain reaction allowing the full pellet to undergo fusion and release vast amounts of energy. Ignition is the key to harnessing fusion energy in the ICF approach. However, controlling the high energy hot plasma is exceedingly difficult and to date, ICF technology has not been able to achieve ignition.

Dynamical stability of converging plasmas has been recognized for decades as one of the most significant limitations to creating a sustained fusion reaction ignition. Hydrodynamic instabilities lead to non-uniform compression and mixing of ablator material into the fuel, degrading the fusion yield. The numerical study of such instabilities in plasmas, which are also ubiquitous in astrophysical environments, such as supernovae, is typically conducted using fluid simulations. Very recently, experimental studies have started to shed light on the importance of kinetic effects, not captured in fluid simulations, in modifying the dynamics of compression and ignition of fusion plasmas. However, the kinetic modeling of such systems has been out of reach due to the outstanding challenges posed by the large difference in scales and the complexity of the models needed to capture all relevant processes.

This project supports first principles multi-dimensional studies of the dynamical stability of plasmas using kinetic simulations, thus capturing the intrinsic multi-scale physics associated with fusion plasmas. This project will take advantage of a suite of massively parallel kinetic codes with the goal of identifying the regimes for which kinetic effects impact the development of hydrodynamic instabilities. Outcomes of the project could ultimately lead to strategies to better control the dynamical stability of fusion plasmas and have an important impact in advancing DOE's clean energy agenda.

Title: Nuclear structure for tests of fundamental symmetries and astroparticle physics

Principal Investigator: Calvin Johnson (San Diego State University)

Co-Investigators: Wick Haxton (University of California Berkeley), Kenneth McElvain (University of California Berkeley), William Ormand (Lawrence Livermore National Laboratory), Esmond Ng (Lawrence Berkeley National Laboratory), Hongzhang Shan (Lawrence Berkeley National Laboratory), Sam Williams (Lawrence Berkeley National Laboratory), Chao Yang (Lawrence Berkeley National Laboratory)

ALCC allocation: Processor Hours

Site: Argonne Leadership Computing Facility (ALCF)
Allocation: 6,000,000 processor hours
Site: National Energy Research Scientific Computing Center (NERSC)
Allocation: 24,000,000 processor hours

Research Summary

What is the universe made of, and how did it get that way? Why is the universe made of matter and not anti-matter? What is the nature of the non-baryonic dark matter that permeates the galaxies? To answer these questions, scientists have turned to nuclei as microscopic experimental laboratories searching for: neutrinoless double-beta decay, where matter is produced from nothing; permanent electric dipole moments, which can only occur if the laws of physics for matter are different from those for anti-matter; and the terrestrial detection of non-baryonic dark matter. This project supports detailed and reliable calculations of the quantum wave functions of atomic nuclei. This will require some of the largest such calculations ever attempted, involving matrices of dimension in excess of 20 billion. Although the project will use an efficient code we have been developing which uses significantly less memory (roughly 1/10th) than other algorithms, we will still need significant computational resources. The results, however, will help provide robust interpretation of these crucial experiments illuminating the fundamental building blocks of our universe.

Title: An End-Station for Intensity and Energy Frontier Experiments and Calculations

Principal Investigator: Taylor Childers (Argonne National Laboratory)

Co-Investigators: Thomas LeCompte (Argonne National Laboratory), Paolo Calafiura (Lawrence Berkeley National Laboratory), Craig Tull (Lawrence Berkeley National Laboratory), Brendan Casey (Fermi National Accelerator Laboratory), Stefan Hoeche (SLAC National Laboratory), Alexei Klimentov (Brookhaven National Laboratory), Torre Wenaus (Brookhaven National Laboratory), Thomas Uram (Argonne National Laboratory), Venkat Vishwanath (Argonne National Laboratory), Elizabeth Sexton-Kennedy (Fermi National Accelerator Laboratory), Douglas Benjamin (Duke University), Radja Boughezal (Argonne National Laboratory), Frank Petriello (Argonne National Laboratory)

ALCC allocation: Processor Hours

Site: Argonne Leadership Computing Facility (ALCF)
Allocation: 93,500,000 processor hours
Site: National Energy Research Scientific Computing Center (NERSC)
Allocation: 13,000,000 processor 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 we propose here 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. By creating an end-station, we have created an environment where other particle physicists can use these supercomputers more easily and quickly than they would on their own, to the benefit of everyone. 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: Modeling Helium-Hydrogen Plasma Mediated Tungsten Surface Response to Predict Fusion Plasma Facing Component Performance in ITER

Principal Investigator: Brian Wirth (University of Tennessee)

Co-Investigators: David Bernholdt (Oak Ridge National Laboratory), Karl Hammond (University of Missouri), Rick Kurtz (Pacific Northwest National Laboratory), Blas Uberuaga (Los Alamos National Laboratory)

ALCC allocation: Processor Hours

Site: Argonne Leadership Computing Facility (ALCF)
Allocation: 70,000,000 processor hours
Site: Oak Ridge Leadership Computing Facility (OLCF)
Allocation: 25,000,000 processor hours

Research Summary

The purpose of this project is to develop high performance materials suitable for a future fusion power plant. The performance demands on plasma-facing components (PFCs), first wall and blanket systems of future fusion power plants are beyond the capability of current materials, which is one of the reasons that the United States National Academy of Engineering has ranked the quest for fusion as one of the top grand challenges for engineering in the 21st Century. Furthermore, it is clear that of this list, the plasma surface interactions (PSIs) occurring in the divertor and PFCs pose a critical scientific challenge that limits our ability to achieve electricity production from fusion. Gaining physical understanding and predictive modeling capability in this critical PSI area requires simultaneously addressing complex and diverse physics occurring over a wide range of lengths (Ångströms to meters) and times (femtoseconds to seconds, days to years), as well as integrating extensive physical processes across the plasma-surface-bulk materials boundaries. It is now well established that He and H plasma exposure with ion energies in the range of 10 to 100 eV and for implanted gas fluences around 1E25 per m2 can have profound effects on the metallic surface morphology and on the quantity of retained gas atoms. The extent to which both surface morphology and sub-surface defect creation and evolution processes driven by neutron-induced damage influence the diffusion, trapping and precipitation of hydrogen isotopes into gas bubbles is an outstanding question that impacts the tritium permeation, retention and near-surface saturation levels; and will definitely influence the operation of the ITER fusion experimental reactor. The outcome of this proposal will be a better understanding of the materials and properties necessary to serve as walls of a future fusion power plant.

Title: Computational Study of Cycle-to-Cycle Variation in Dual-Fuel Engines

Principal Investigator: Ravichandra Jupudi (General Electric Global Research)

Co-Investigators: Roy J. Primus (General Electric Global Research), Bhaskar Tamma (General Electric Global Research), Sibendu Som (Argonne National Laboratory), Janardhan Kodavasal (Argonne National Laboratory), Charles E.A. Finney (Oak Ridge National Laboratory), Miroslav K. Stoyanov (Oak Ridge National Laboratory), Sameera Wijeyakulasuriya (Convergent Science Inc), Eric Pomraning (Convergent Science Inc), Keith Richards (Convergent Science Inc), P. Kelly Senecal (Convergent Science Inc)

ALCC allocation: Processor Hours
Site: Argonne Leadership Computing Facility (ALCF)
Allocation: 25,000,000 processor hours

Research Summary

Premixed staged combustion (PSC), where fuel is burned in stages, is a new technology that has the potential to improve efficiency while reducing pollutant formation in combustion turbine design. Advanced engine designs from multiple gas turbine companies already incorporate this technology. However, a deeper understanding of the unique interactions between turbulence and combustion chemistry that occur in this reaction regime is needed to advance combustor technology and deliver higher gas turbine efficiency.

This project will use leadership-class computing resources at Oak Ridge National Laboratory in a multi-pronged effort to characterize this important combustion regime and accelerate its application to industrial gas turbine engines. The project is organized in three stages:

1. Assess the performance of current state-of-the-art modeling approaches for PSC
2. Explore poorly understood regions of the PSC state space
3. Apply learning and understanding to impact realistic gas turbine engines

In stage 1, existing modeling approaches will be evaluated to assess their performance in PSC regimes. This task will harness the vast computing power of Titan to compute detailed combustion simulations to answer fundamental questions about how turbulence and chemistry interact in canonical PSC configurations. Results from engineering models will be compared with the detailed simulations and with experiments to identify the most effective modeling approaches for the conditions in today's gas turbine engines. Stage 2 will use high-fidelity large-eddy simulations to explore combustion regimes that are relevant to where gas turbine engines are heading in the near future, including higher pressures, shorter residence times, and more intense turbulence. Finally, in stage 3, realistic engineering configurations for PSC will be explored using high performance computing. The results of these simulation will provide the fundamental understanding required to use premixed stage combustion effectively in future combustion turbine designs.

Title: Numerical Simulation of Turbulent Flows in Advanced Steam Generators - Year 2

Principal Investigator: Aleksandr Obabko (Argonne National Laboratory)

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

ALCC allocation: Processor Hours

Site: Argonne Leadership Computing Facility (ALCF)

Allocation: 80,000,000 processor hours

Research Summary

Failure of steam generators in pressurized water reactor (PWR) nuclear systems can lead to expensive plant shut-downs. Flow-induced vibrations in the steam generators was identified as a high impact problems of the Department of Energy's Nuclear Energy Modeling and Simulation program as a high-impact problem for nuclear industry. In addition to the impact on existing nuclear reactors, there is a growing interest in small modular reactors (SMRs). Many of the proposed SMR designs, from vendors such as NuScale, mPower, and Westinghouse, are based on an integral PWR concept. In all of these systems, flow-induced vibrations are an important limiting factor in the operation of heat exchangers and steam generators.

An advanced numerical simulation capability for modeling such phenomena will help improve the analysis and evaluation of different design variants in terms of vibrations and heat transfer performance, complementing expensive experimental tests and reducing their cost. Such a tool will have a broad utility for various nuclear technology applications, including other fluid-structure interaction problems, such as fluid elastic instability for rod bundles in cross flows. This tool will couple Lawrence Livermore's Diablo structural mechanics code with Argonne's Nek5000 computational fluid dynamics (CFD) code. Because the structural displacements are small, a one-way coupling, where the CFD results provide inputs to the structural code, can be used.

In collaboration with US-based NuScale Power, the research team will simulate a series to perform a series of single phase turbulent flows, including several experimental and mockup geometries of a helical coil steam generator. These will be compared to, and validated, using existing experiments at Texas A & M University.

Title: Fluctuations of conserved charges in high temperature QCD: towards the continuum

Principal Investigator: Swagato Mukherjee (Brookhaven National Laboratory)

Co-Investigators: Frithjof Karsch (Brookhaven National Laboratory), Edwin Laermann (Bielefeld University), Peter Petreczky (Brookhaven National Laboratory), Christian Schmidt (Bielefeld University), Sayantan Sharma (Brookhaven National Laboratory), Patrick Steinbrecher (Brookhaven National Laboratory)

ALCC allocation: Processor Hours

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 167,000,000 processor hours

Research Summary

Solids, liquids, and gasses are three examples of phases of matter. The phase of a substance, such as liquid water vs water ice, is determined by variables such as the temperature and pressure. A phase diagram is a graphic showing a substance's phase as a result of variables such as temperature and pressure. Phase diagrams can show surprisingly rich information about a substance and are often crucial steps in understanding a material and developing real-world applications.

Phase diagrams describe not only everyday matter, but also the more exotic "nuclear matter", matter inside nuclei. Nuclei are made of strongly interacting protons and neutrons, which are themselves composed of quarks, bound together by gluons. Just as ice melts into water with sufficiently high temperatures, so the protons and neutrons in nuclei "melt" into a remarkable phase of matter, the "quark gluon plasma" (QGP). Establishing the boundary between atomic nuclei and the QGP in the nuclear matter phase diagram is a central goal of the heavy-ion research programs at major experimental facilities such as the Relativistic Heavy Ion Collider (RHIC) at BNL and ALICE at CERN.

As the theory that describes quarks and gluons, Quantum Chromodynamics (QCD), is precisely known mathematically and can be studied on supercomputers, one might hope to predict the phase diagram of QCD directly. This can already be done for gluons alone, but determining the phase diagram of the real-world mixture of quarks and gluons requires new computational techniques. This ALCC project supports the investigation of a promising method for predicting the QCD phase diagram in the presence of both quarks and gluons. A special goal of this research is to attempt to find the conjectured "critical point" in the QCD phase diagram, the end point of the line along which nuclei and the QGP phase coexist . The results of this project are expected to help in the interpretation of experimental observations at RHIC, and to contribute to our understanding of the physics of quarks and gluons more generally.

Title: The Glue That Binds Us All

Principal Investigator: Robert Edwards (Jefferson Lab)

Co-Investigators: Raul Briceno (Old Dominion University), Jozef Dudek (Old Dominion University), Balint Joo (Jefferson Lab), Frank Winter (Jefferson Lab), David Wilson (University of Cambridge), Christopher Thomas (University of Cambridge), Michael Peardon (Trinity College, Dublin), Sinead Ryan (Trinity College, Dublin)

ALCC allocation: Processor Hours
Site: Oak Ridge Leadership Computing Facility (OLCF)
Allocation: 90,000,000 processor hours

Research Summary

Understanding mesons, a type of subatomic particle, is critical to advancing our knowledge of matter and fundamental physics. The determination of the meson spectrum is a world-wide effort and exploring the excited meson spectrum is a flagship project of the Thomas Jefferson National Accelerator Facility (Jefferson Lab) The excited meson spectrum is of particular interest because it may reveal important properties of "exotic" mesons. GlueX is a new experiment at Jefferson Lab built to discover and explore properties of exotic meson states. This project supports computations of the excited-state meson spectrum of Quantum Chromodynamics (QCD), the theory of strong interactions (a fundamental force of nature). 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 will commence measurements in 2016. These calculations will guide future experimental searches and advance our understanding of matter and fundamental physics.

Title: Multi-scale, high-resolution integrated terrestrial and lower atmosphere simulation of the contiguous United States (CONUS)

Principal Investigator: Reed Maxwell (Colorado School of Mines)

Co-Investigators: Laura Condon (Syracuse University), David Moulton (Los Alamos National Laboratory)

ALCC allocation: Processor Hours

Site: National Energy Research Scientific Computing Center (NERSC)

Allocation: 15,000,000 processor hours

Research Summary

Recent research shows that changes in climate affect groundwater and overland (surface) flow, land surface processes, vegetation, and the atmospheric boundary layer. These processes also feedback into climate and can affect local and regional climate across a range of scales and climatic conditions. Furthermore, the type and distribution of vegetation cover affects not only water runoff generation and overland flow processes, but also the land–atmosphere water and energy fluxes. This creates feedbacks on local and regional climate. While the full effects of these feedbacks across terrain ranging from plains to mountains is not well understood, results suggest that understanding and predicting hydrologic response to changing climate and forest cover requires explicit consideration of groundwater–land surface–atmosphere feedbacks within terrestrial modeling frameworks.

Developing high resolution continental scale models has been identified as a Grand Challenge in computational hydrology. This project will explore the computational challenges and scientific potential of integrated hydrology models at continental scales, including their coupling to land process models. The project will build upon an integrated hydrologic model of most of the continental US. The proposed simulations will provide the first continental scale simulations of the interactions between the atmosphere and terrestrial hydrology within a fully integrated framework. The simulations will use a phased approach to explore multi-scale, multi-physics treatments of Earth system modeling from the bedrock into the lower atmosphere. To adequately capture feedbacks between deeper subsurface flow, the land energy budget and the lower atmosphere, the models explicitly connects these systems. The model will be used to run three high-fidelity simulations that feature integrated subsurface and surface flow, land-surface water fluxes and energy fluxes. The three simulations will include a baseline scenario, a warming scenario, and a groundwater pumping scenario to diagnose climate and aquifer extraction impacts to the hydrologic cycle. Each of these simulations will capture a decade of real time for the system. These simulations will provide state-of-the science predictions of the current and future hydrology over North America.

Title: Muon g – 2 Hadronic Vacuum Polarization from Lattice QCD

Principal Investigator: John Laiho (Syracuse University)

Co-Investigators: Claude Bernard (Washington University), Carleton DeTar (University of Utah), Steven Gottlieb (Indiana University), Andreas Kronfeld (Fermilab), Paul Mackenzie (Fermilab), Urs Heller (American Physical Society), Robert Sugar (University of California, Santa Barbara), Doug Toussaint (University of Arizona), Ruth Van de Water (Fermilab), Peter Lepage (Cornell University), Christine Davies (University of Glasgow), Chris Monahan (Rutgers University), Aida El-Khadra (University of Illinois)

ALCC allocation: Processor Hours

Site: Argonne Leadership Computing Facility (ALCF)

Allocation: 66,000,000 processor hours

Research Summary

The Standard Model of particle physics has been enormously successful in describing the sub- atomic world of particle physics. A persistent exception to this success is the disagreement between the Standard Model prediction and measurements of the muon's anomalous magnetic moment. The muon is a fundamental particle, similar to the electron, and the "anomalous" part of the magnetic moment stems from quantum mechanical interactions between the muon and elementary particles of all kinds. If the disagreement holds up, it would signal the presence of new, otherwise unobserved physical phenomena. The disagreement between theory and experiment is not yet definitive, so a new experiment, aiming to reduce the measurement uncertainties fourfold, is being mounted at Fermilab. To receive a full return on this investment, the weakest parts of the theory have to be commensurately improved. This allocation supports using lattice QCD, which is acknowledged to be the only theoretical tool up to the task, to carry out the first ab initio calculations of the hadronic contributions to the anomalous magnetic moment. Uncertainties in these calculations (obtained in less sophisticated ways) are the main theoretical obstacle to establishing new physics in the muon's magnetic moment. This work will either underpin a major discovery or, should theory and experiment end up in accord, be a powerful constraint on new models of particle physics.

Title: Functional Mechanisms and Allostery in Energy-driven Molecular Machines in Membranes

Principal Investigator: Harel Weinstein (Weill Cornell Medical College of Cornell University)

Co-Investigators: n/a

ALCC allocation: Processor Hours
Site: Oak Ridge Leadership Computing Facility (OLCF)
Allocation: 15,500,000 processor hours

Research Summary

This allocation supports work to develop an understanding of biological processes by which molecular machines transport substrates and signals across membranes. The project is a key element in a research program that aims to determine the functional properties and modes of action of membrane proteins in the family of secondary transporters as prototypical biological machines of the cell. Our goal is to understand their functions under normal conditions, and under conditions impacted by interactions with their environments, including changes in membrane components and various forms of energy coupling, e.g., through various chemical gradients. Here we study computationally a specific member of the family of sodium-coupled symporters (Na+-powered substrate transporters across cell membranes) and the key components of its transport-assembly (i.e., the membrane and protein complex essential for transporter functions). The mechanistic study of this protein, the dopamine transporter DAT, at the unprecedented level of computational molecular dynamics presents the distinct advantage of addressing a system that is being intensively investigated experimentally due to its very high significance in biology, in human physiology, and in medicine. As a result the specific protein members of its functional environment are also well characterized and broadly investigated. Using advanced methods of computational simulation and multiscale analysis, the aim is to uncover detailed dynamic properties and energetics of these complex molecular machines under the various relevant conditions. Through extensive multi-scale, integrative computational approaches, carried out on unprecedented spatial/temporal scales, we seek to uncover thermodynamics-based insights into chemical environment and factors necessary to maintain tight control of two key physiological energy transduction processes: transporter-mediated reuptake and substrate efflux. The outcome of this proposal will be an improved understanding of the molecular mechanisms of energy-driven transport and signaling across membranes which are essential for the construction of predictive computational models of complex biological processes and their use in various bioengineering projects.

Title: Molecular Modeling of Hot Electron Transfer for Solar Energy Conversion

Principal Investigator: Hanning Chen (George Washington University)

Co-Investigators: Wei Jiang (Argonne National Laboratory)

ALCC allocation: Processor Hours
Site: Argonne Leadership Computing Facility (ALCF)
Allocation: 16,000,000 processor hours

Research Summary

The Sun radiates ~200,000 TW of power onto the Earth; about 1,000 times the yearly global power demand. Because of this bounty, technology that converts solar energy into a usable energy (e.g. electricity) stands out among renewable energy technologies. Solar cell technology is one such technology and takes energy irradiated from the sun and converts it into usable electrical energy. At the heart of solar cell technology is light-matter interactions; light irradiated from the sun interacts with matter to excite electrons. The excited electrons can then enter conduction bands, move in a current, and create electricity. The physics of light-matter interactions determine how efficient light energy can be converted into electrical energy, with current theory limiting the efficiency to 33%. However, a theory of 'hot-carrier solar cells' predicts that under special non-equilibrium circumstances, efficiencies of up to 66% could be achieved. Realizing 66% efficiencies could drastically push solar cell technology forward, yet in spite of its theoretical feasibility, experimental implementation is still in its infancy. This project supports development of a computational paradigm based on vibrationally resolved electron transfer theory to guide the experimental implementation and optimization of the emerging hot-carrier solar cells. In addition, this framework will be implemented into a highly parallel open-source molecular simulation package that is freely available to the scientific community for research and education. Outcomes of this project will advance our understanding of light-matter interactions in photovoltaics and could provide key insights to realizing next generation solar cell technology.

Title: High Performace Computing for Manufacturing

Principal Investigator: Peter Nugent (Lawrence Berkeley National Laboratory)

Co-Investigators: John Turner (Oak Ridge National Laboratory)

ALCC allocation: Processor Hours

Site: Oak Ridge Leadership Computing Facility (OLCF)
Allocation: 41,000,000 processor hours
Site: National Energy Research Scientific Computing Center (NERSC)
Allocation: 10,000,000 processor hours

Research Summary

Advanced modeling and simulation capabilities and large scale data analysis provide tools to address challenging problems in decision making, process and design optimization, and performance prediction that occur in manufacturing. Effective use of High Performance Computing (HPC) in manufacturing has the potential to improve quality, reduce energy costs, accelerate or eliminate expensive testing, and accelerate the time of adoption of new technologies. With the support of the Department of Energy's Advanced Manufacturing Office (AMO), Lawrence Livermore National Laboratory (LLNL), Oak Ridge National Laboratory (ORNL,) and Lawrence Berkeley National Laboratory (LBNL) have formed the HPC4Mfg program designed to allow partnerships with industry in bringing HPC into manufacturing.

This award supports the HPC4Mfg program, providing time on HPC resources to three National Laboratory/Industry partnerships designed as pilot programs for applying HPC to advanced manufacturing projects. These partnership projects are:

1. "Computational Design and Optimization of Ultra Low Power Device Architectures," a partnership between Global Foundries and LBNL, will use atomistic simulation to simulate annealing of amorphous oxide layers, a key step in the production of semiconductors. Improved understanding of this process will enable future electronics with reduced power usage.
2. "Massively Parallel Multi-Physics, Multi-Scale Large Eddy Simulations of a Fully Integrated Aircraft Engine Combustor and High Pressure Vane, a partnership between GE and ORNL, will use high-fidelity Computational Fluid Dynamics (CFD) simulations to perform integrated simulations of combustion aircraft engine combustors and turbines, with the goal of accelerating the design process.
3. "Integrated Predictive Tools for Customizing Microstructure and Material Properties of Additively Manufactured Aerospace Components," a partnership between United Technologies Research Center and ORNL, will simulate heat transfer in titanium alloys during additive manufacturing processes. Because thermal effects determine the materials properties of the final part, the results will allow prediction of materials properties, and optimization of the manufacturing processes.

Title: 61-pin wire-wrap turbulent conjugate-heat transfer: V&V for industry and SESAME

Principal Investigator: Elia Merzari (Argonne National Laboratory)

Co-Investigators: Elia Merzari (Argonne National Laboratory), Oana Marin (Argonne National Laboratory), Brian Jackson (TerraPower, LLC), Anthony Chang (Areva Federal Services), Kostas Karazis (Areva Federal Services), Mariano Tarantino (ENEA), Paul Fischer (University of Illinois at Urbana-Champaign)

ALCC allocation: Processor Hours

Site: Argonne Leadership Computing Facility (ALCF)

Allocation: 120,000,000 processor hours

Research Summary

The prediction of thermal performance in wire wrapped assemblies, such as fuel pins proposed for use in sodium-cooled nuclear reactors, is vital to evaluate overall reactor performance and safety. To capture the performance over the life of the reactor, these simulations need to be performed not only for the original geometry of the pins, but for the deformed end-of-life geometries. The Thermal/Hydraulic Computational Fluid Dynamics (CFD) modeling required to reach these goals requires not only simulation of the complex geometry of hexagonal packed fuel assemblies with helically wire wrapped pins, but high-fidelity turbulent flow modeling. This can be achieved using the Nek5000 large-eddy simulation code.

To ensure that these simulations can be used to make operational decisions, this allocation supports simulations for the geometries that can be validated against experiments. Four geometries will be simulated: an isothermal undeformed duct with an undeformed 61-pin geometry, an isothermal deformed duct with a deformed 61-pin geometry, a heated undeformed duct with an undeformed 61-pin geometry, and a heated deformed duct with an undeformed 61-pin geometry. These results will be compared against experiments performed at Texas A&M and Areva Incoprorated's Richland (WA) facility. The high-fidelity data will be shared with the 25 European and American partners in the SESAME (Simulations and Experiments for the Safety Assessment of MEtal cooled reactors) project.

Title: Discovering Optimal Deep Learning and Neuromorphic Network Structures using Evolutionary Approaches on High Performance Computers

Principal Investigator: Robert Patton (Oak Ridge National Laboratory)

Co-Investigators: Thomas Potok, Thomas Karnowski, Bobby Philip, Seung-Hwan Lim, Derek Rose, Steven Young, Catherine Schuman. (Oak Ridge National Laboratory)

ALCC allocation: Processor Hours
Site: Oak Ridge Leadership Computing Facility (OLCF)
Allocation: 75,000,000 processor hours

Research Summary

Machine Learning is an approach to computing that provides the computer a general algorithm with unset variables. Through trial and error, the computer 'learns' the best values for the variables in order for the algorithm to optimally solve the problem at hand. Deep Learning (DL) is one important type of machine learning. DL uses a layered network model for the algorithm. Although the machine learning process learns values for variables in the network algorithm, the size and shape of the network itself is pre-determined. These predetermined values, called "hyper-parameters", describe the network topology and can have significant impact on the accuracy of the final machine learning program. In order for DL algorithms to be effective at producing new scientific discovery, the automated and optimal configuration of these algorithms must be solved. This project supports efforts to solve the single largest barrier to effectively applying deep learning for scientific datasets, which is the heuristic approach to configuring hyper-parameters, (e.g., number of layers, kernel size,, etc.), which are vital to DL implementations.

This project will develop a scalable Evolutionary Algorithm (EA), primarily using OLCF's Titan supercomputer, which searches for both the optimal hyper-parameters and topology of a deep learning network. A tool such as this currently does not exist, and this proposed approach would provide the capability to:

1. Study interactions between parameters and their values as applied to different data
2. Reduce the expertise required for parameter tuning
3. Extend the application of deep learning algorithms to extreme scale data

Solving the hyper-parameter optimization problem will make DL a valuable tool for scientific discovery with DOE and beyond, especially as an "in situ" data analysis tool for scientific simulations. The project will also explore extending the scalable EA framework to optimize network topology and parameters for neuromorphic networks, or neural networks meant for hardware implementation. Similar to DL implementations, the effect of network topology and parameters on the performance of neuromorphic networks is largely unknown. Utilizing the scalable EA framework for neuromorphic networks will provide valuable information about the capabilities and characteristics of neuromorphic devices.

Title: First principles design of magnetic materials, models, and mechanisms

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

Co-Investigators: Elif Ertekin (University of Illinois at Urbana-Champaign)

ALCC allocation: Processor Hours
Site: Argonne Leadership Computing Facility (ALCF)
Allocation: 30,000,000 processor hours

Research Summary

This projects studies the role of magnetism at the nanoscale. Magnetism is a quantum effect visible at macroscale, and is used for many purposes technologically. In addition to the familiar permanent magnets, magnetism is important in new more efficient spin-based electronics, high temperature superconductivity, and heating and cooling systems using magnetic phase transitions. Magnetism arises from the combination of quantum effects and the interaction between electrons. Studying magnetism in materials requires us to calculate the properties of a manybody system of quantum particles at the nanoscale, and build models from the bottom up of the magnetic behavior. The more reliable these models, the more powerful the control over magnetism, which enables applications such as those listed above.

This project uses the Mira supercomputer and quantum Monte Carlo (QMC) techniques to directly simulate quantum electrons interacting with one another to create magnetism. The goal is to investigate to what extent the QMC techniques can improve our description of magnetism, with final goal of developing quantitative models of magnetism. The outcome of this project is an improved understanding of the behavior of magnetic materials and will make a step forward in computer-aided design of next-generation materials.

Title: Cosmic Frontier Computational End-Station

Principal Investigator: Katrin Heitmann (Argonne National Laboratory)

Co-Investigators: Salman Habib (Argonne National Laboratory), Zarija Lukic (Lawrence Berkeley National Laboratory)

ALCC allocation: Processor Hours

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 55,000,000 processor hours

Research Summary

The Cosmic Frontier effort within DOE High Energy Physics targets the physics of dark energy and dark matter, cosmological probes of neutrino physics, and the nature of primordial fluctuations -- some of the most important questions in all of physical science. Cosmological surveys are essential components of a major national and international research program aimed at understanding these questions. DOE-supported experiments include the Dark Energy Survey (DES) and the South Pole Telescope (SPT). Planned future surveys include the Dark Energy Spectroscopic Instrument (DESI), the Large Synoptic Survey Telescope (LSST), and SPT-3G. Two essential aspects of interpreting results from survey measurements are 1) the central role played by large-scale simulations and data analyses in enabling and interpreting observations, and 2) the significance of cross-correlation analyses across different types of measurements, necessary to control systematic errors. Motivated by these considerations, this project supports a computational end-station tasked with a set of simulation and analysis sub-projects. This computational end-station as an important step in a larger effort. In this step, the project brings together simulations and analyses tasks to enable some of the analysis of ongoing surveys as well as preparing for the next generation surveys. The proposed set of simulations includes some of the largest ever performed in their respective class.

Title: Hadronic light-by-light scattering contribution 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 (Brookhaven National Laboratory), Luchang Jin (Columbia University), Chulwoo Jung (Brookhaven National Laboratory), Andreas Juettner (University of Southampton), Christoph Lehner (Brookhaven National Laboratory), Antonin Portelli (University of Edinburgh)

ALCC allocation: Processor Hours
Site: Argonne Leadership Computing Facility (ALCF)
Allocation: 180,000,000 processor hours

Research Summary

A primary goal of physics is to describe everything in the universe through its most fundamental forces and particles. The current culmination of this goal is The Standard Model, which successfully describes electromagnetism, the weak force, the strong force and the particles they act on. The Standard Model is so successful that it has remained fundamentally unchanged for the past 50 years. Beyond Standard Model physics is an exciting area of research aiming to look for experiments that violate the Standard Model and require a new theory for the universe. Recently, experiments on the muon, a fundamental particle in the Standard Model, suggest the true value of its magnetic moment may be in disagreement with what is predicted by the Standard Model. However, the detected difference between theory and experiment is within the error tolerances of both the theory and experiment. Does the muon magnetic moment agree with predictions of the Standard Model, or is it evidence of a break in one of physics most important theories? To answer this question, a new experiment is underway that will measure the muon magnetic moment to a very high degree of accuracy. To compare the result to theory, computations must be made to determine the prediction of the Standard Model to the same level of accuracy. The goal of this project is to compute the hadronic contributions to the muon anomalous magnetic moment from first principles using lattice quantum chromodynamics (QCD), the fundamental theory in the Standard Model of the strong nuclear force that describes the interactions of quarks and gluons. The hadronic contributions represent the largest of the theory uncertainties, and their accurate determination is crucial to compare the Standard Model with the precise experimental measurement at Fermilab to potentially discover new physics.

.

Title: Exploring Higgs Compositeness Mechanism in the Era of the 14 TeV LHC

Principal Investigator: George Fleming (Yale University)

Co-Investigators: Richard Brower (Boston University), Anna Hasenfratz (University of Coloroado, Boulder), Joseph Kiskis (University of California, Davis), Ethan Neil (University of Colorado, Boulder), James Osborn (Argonne National Laboratory), Claudio Rebbi (Boston University), Pavlos Vranas (Lawrence Livermore National Laboratory)

ALCC allocation: Processor Hours

Site: Argonne Leadership Computing Facility (ALCF)

Allocation: 55,000,000 processor hours

Research Summary

A primary goal of physics is to describe the universe through its most fundamental forces and particles. The current culmination of this goal, The Standard Model, has withstood ~40 years of rigorous experimentation and successfully describes the electroweak force, the strong force and the particles they act on. The recent observation of the Higgs Boson, a particle predicted by the Standard Model to explain how fundamental particles have mass, was an exciting validation of the Standard Model. Yet the discovery of the Higgs also offers new opportunities to test the Standard Model and possibly discover new physics beyond the Standard Model. Is the Higgs truly a fundamental particle as the current Standard Model predicts, or is it a composite particle, composed instead of more fundamental particles? This project supports studies of the interactions of the Higgs boson and the electroweak W and Z gauge bosons in a theory in which Higgs boson, as well as the longitudinal spin components of the W and Z, are composite. A new interaction describes the forces among hypothesized constituents of these particles. The outcome of this project will be an exploration of a new theory and may lead to important discoveries on the nature of the Higgs, fundamental particles and forces of our universe.

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

Principal Investigator: Mujeeb Malik (NASA Langley Ressearch Center)

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

ALCC allocation: Processor Hours
Site: National Energy Research Scientific Computing Center (NERSC)
Allocation: 66,000,000 processor 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 drag on aircraft advances these goals, but requires validated high-fidelity simulations 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. This pressure gradient may be a result of the geometry of the surface, or the presence of a shock wave. This 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.

Such separated flows are generally difficult to simulate because they involve high Reynolds number turbulence. Much of the previous computational work in the field has employed lower-fidelity simulation techniques with mixed success. Because of the computational requirements, higher-fidelity techniques have mostly been able to simulate Reynolds numbers that are much lower than those found in practice. This study will perform wall-resolved high-fidelity simulations of a benchmark test case that is representative of a very common problem encountered in transonic flows.

The case to be simulated will be shock-induced boundary layer separation over the upper surface of a transonic airfoil. The proposed calculations will simulate the National Aeronautics and Space Administration (NASA) experiment conducted by Bachalo and Johnson. The use of high-performance computing will allow the simulations to be performed at a Reynolds number identical to the experiment. The wall-resolved simulations will generate a detailed dataset that can be used to obtain a better understanding of the problem, guide the development of new/improved wall models for use in the wall-modeled large eddy simulation (WMLES) of separated flows, and recalibrate existing Reynolds-Averaged Navier-Stokes (RANS) turbulence models.

Improved wall models will enable future simulations of practical high Reynolds number problems with significantly reduced computational cost. Correct predictions of flow separation in turn can be used to improve aerodynamic designs and reduce drag. These improvements can ultimately pave the way towards the development of aerodynamic designs with reduced viscous drag and hence aircraft fuel burn reduction. Improved turbulence modeling capabilities also have the ability to impact energy technologies such as wind energy.

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: Processor Hours

Site: Oak Ridge Leadership Computing Facility (OLCF)
Allocation: 20,000,000 processor hours
Site: National Energy Research Scientific Computing Center (NERSC)
Allocation: 10,000,000 processor hours

Research Summary

Supercomputers are powerful scientific and mathematical instruments capable of calculating quadrillions of operations in a single second. Science harnesses the compute speed of supercomputers to tackle some of science's biggest challenges including exploration of complex and important theories, simulation of otherwise impossible experiments on sub-atomic to cosmological scales, and rapid prototyping to speed time-to solution for industrial technology. Supercomputer gain their speed from massively parallel architectures and one of the largest challenges lies in programming application software that is capable of using these massively parallel systems.

Three science driven DOE supercomputing facilities; Argonne Leadership Computing Facility (ALCF), National Energy Research Scientific Computing Center (NERSC) and the Oak Ridge Leadership Computing Center (OLCF), are supporting efforts to prepare scientists and engineers for the next-generation supercomputers that are planned for these centers through early application readiness programs that will prepare a range of application codes. This allocation supports these three facilities as they collaborate on their application readiness programs for planned computer upgrades in 2016-2018. This allocation supports computational resources to enable the application readiness teams to carry out software development with the common objective of architecture and performance portability across the distinct architectures.

Title: Premixed staged combustion for increased efficiency in gas turbine engines

Principal Investigator: Jin Yan (GE Global Research Center)

Co-Investigators: Frank Ham (Cascade Technologies), Sanjeeb Bose (Cascade Technologies), Lee Shunn (Cascade Technologies), Babak Hejazi (Cascade Technologies), Shashank Yellapantula (GE Global Research Center), Shoreh Hajiloo (GE Power)

ALCC allocation: Processor Hours
Site: Oak Ridge Leadership Computing Facility (OLCF)
Allocation: 20,000,000 processor hours

Research Summary

Premixed stage combustion (PSC), where fuel is burned in stages, is a new technology that has the potential to improve efficiency while reducing pollutant formation in combustion turbine design. Advanced engine designs from multiple gas turbine companies already incorporate this technology. However, a deeper understanding of the unique interactions between turbulence and combustion chemistry that occur in this reaction regime is needed to advance combustor technology and deliver higher gas turbine efficiency.

This project will use leadership-class computing resources at Oak Ridge National Laboratory in a multi-pronged effort to characterize this important combustion regime and accelerate its application to industrial gas turbine engines. The project is organized in three stages:

1. Assess the performance of current state-of-the-art modeling approaches for PSC
2. Explore poorly understood regions of the PSC state space
3. Apply learning and understanding to impact realistic gas turbine engines

In stage 1, existing modeling approaches will be evaluated to assess their performance in PSC regimes. This task will harness the vast computing power of Titan to compute detailed combustion simulations to answer fundamental questions about how turbulence and chemistry interact in canonical PSC configurations. Results from engineering models will be compared with the detailed simulations and with experiments to identify the most effective modeling approaches for the conditions in today's gas turbine engines. In stage 2 will use high-fidelity large-eddy simulations to explore combustion regimes that are relevant to where gas turbine engines are heading in the near future, including higher pressures, shorter residence times, and more intense turbulence. Finally, in stage 3, realistic engineering configurations for PSC will be explored using high performance computing. The results of these simulation will provide the fundamental understanding required to use premixed stage combustion effectively in future combustion turbine designs.

Title: High-Throughput Screening and Machine Learning for Predicting Catalyst Structure and Designing Effective Catalysts

Principal Investigator: Efthimios Kaxiras (Harvard University)

Co-Investigators: Matthew Montemore (Harvard University), Wei Chen (Harvard University), Ekin Dogus Cubuk (Harvard University)

ALCC allocation: Processor Hours

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 17,500,000 processor hours

Research Summary

This project supports the design of new and efficient catalysts important to the clean and energy efficient production of industrial chemicals. Catalysts are used to make chemical transformations more efficient, and processes involving catalysts make up a significant fraction of energy use and economic output in industrialized countries. However, new catalysts are difficult to discover and design, for two main reasons: (1) catalysts and their interactions with molecules are very complex systems, making them hard to understand, and (2) the number of possible catalysts is very large, making it hard to screen all of the possibilities. In this project, we tackle both of these challenges for important systems. First, we will stitch together the results of thousands of calculations to understand an important catalyst, nanoporous gold. The reason for nanoporous gold's excellent catalytic performance remains unknown, making it difficult to engineer further improvements. Second, we will computationally test thousands of possible catalysts for the production of an important chemical, formaldehyde, of which millions of tons are produced every year. Using so-called machine learning, including artificial neural networks, we will extrapolate the results to further test tens of thousands of possible catalysts. We will then work with our collaborators to synthesize and test our most promising candidates. Overall, we will improve our understanding nanoporous gold, and attempt to discover a new catalyst for an industrially important reaction. The data we generate and the insight we gain will be useful for designing and screening catalysts for other reactions. The outcome of this proposal will be an improved understanding of catalytic materials important to clean energy technology.

Title: Simulations of laser experiments to study the origin of cosmic magnetic fields

Principal Investigator: Petros Tzeferacos (University of Chicago)

Co-Investigators: Don Lamb (University of Chicago), Milad Fatenejad (University of Chicago), Norbert Flocke (University of Chicago), Carlo Graziani (University of Chicago), Klaus Weide (University of Chicago), Michael Papka (Argonne National Laboratory), Katherine Riley (Argonne National Laboratory), Venkat Vishwanath (Argonne National Laboratory)

ALCC allocation: Processor Hours

Site: Argonne Leadership Computing Facility (ALCF)

Allocation: 60,000,000 processor hours

Research Summary

Magnetic fields are ubiquitous in the universe, but their origin is not fully understood. A variety of ways have been proposed in which seed magnetic fields could be created. It is widely thought that the larger values of the cosmic magnetic fields that we observe are a result of the amplification of these seed fields by the nonlinear turbulent dynamo mechanism. However, this mechanism has never been demonstrated in the laboratory. We have conceived experiments designed to demonstrate nonlinear turbulent dynamo in the laboratory, providing a comprehensive picture of the energy cascade in magnetized, turbulent plasmas. The experiments utilize high-intensity lasers at the Omega laser facility at the Laboratory for Laser Energetics, and the National Ignition Facility at the DOE national laboratory, Lawrence Livermore National Laboratory. We have been awarded shots on both – highly competitive – facilities. These powerful lasers allow us to produce strong turbulence and large magnetic Reynolds numbers required for the nonlinear turbulent dynamo mechanism to operate. The experiments are designed and interpreted using validated simulations done with FLASH, a highly capable radiation-MHD code we have developed. FLASH simulations are vital to ensure the experiments achieve the desired flow properties. They must be 3D to capture the turbulence, and they must resolve the small spatial scales at which amplification occurs. FLASH simulations are also critical in determining diagnostics' timing, since the signals last only a few nanoseconds. Finally, validated FLASH simulations are crucial to interpreting the results of the experiments. This allocation supports NIF experiments that will not succeed without 3D FLASH simulations on Mira.

Title: Chombo-Crunch: Modeling Pore Scale Reactive Transport Processes Associated with Carbon Sequestration

Principal Investigator: David Trebotich (Lawrence Berkeley National Laboratory)

Co-Investigators: Brian Van Straalen (Lawrence Berkeley National Laboratory), Sergi Molins (Lawrence Berkeley National Laboratory), Carl Steefel (Lawrence Berkeley National Laboratory), Greg Miller (Lawrence Berkeley National Laboratory)

ALCC allocation: Processor Hours

Site: National Energy Research Scientific Computing Center (NERSC)

Allocation: 40,000,000 processor hours

Research Summary

Capture and geological storage of CO_2 from coal and natural gas combustion is among the most critical energy technologies of the next century. However, the real and perceived risks of CO_2 storage remain a significant barrier to public acceptance. A major theme regarding the security of subsurface storage is the integrity of nanoporous (very low permeability) rock layers like shale that are critical for blocking upward flow of buoyant CO_2 during early phases of carbon storage. A key question is whether the pressure needed to inject commercial volumes of CO_2 into permeable sandstone and carbonate formations will cause mechanical failure leading to leakage pathways.

Shale varies in mineralogical composition and pore structure and its mechanical response to pressurization varies accordingly. The role of geochemical processes in widening or healing fractures varies with mineralogy, pore structure, and invading fluid composition. The response of shales to simultaneous mechanical and chemical changes, such as those seen during carbon injection, is not well understood. A new understanding of coupled mesoscale behavior can be obtained through advanced simulation.

Flow in subsurface porous media is typically modeled at the reservoir or field scale ($\sim$1 km), making use of bulk parameters and averaged quantities, which do not take into account geometric heterogeneity of the pore space. An alternative approach is pore scale modeling. Pore scale modeling resolves flow and transport in the micron scale of the pore, computing permeability and other properties such as reactivity and diffusivity. These parameters can then be upscaled for more accurate field scale models. For this upscaling to be effective, the pore scale must be simulated on a domain the size of a representative elemental volume (REV) of the field scale.

This project will use the code Chombo-Crunch, which is able to model pore scale processes at unprecedented scale and resolution, to simulate flow and transport in Marcellus shale near the scale of a REV ($\approx$1 mm), allowing accurate upscaling of permeability and reaction rates.

Title: Large scale electronic structure calculations of nanosystems

Principal Investigator: Lin-Wang Wang (Lawrence Berkeley National Laboratory)

Co-Investigators: Emily Carter (Princeton University), Laurent Bellaiche (University of Arkansas)

ALCC allocation: Processor Hours

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 30,000,000 processor hours

Research Summary

This project uses supercomputers at Oak Ridge Leadership Computing Facility to tackle some of the most challenging problems in nanoscience research. In particular, it will study the electron localization in nanosystems from colloidal quantum dots to piezoelectric materials, and to two dimensional materials. The understanding of their electronic structures will help to use these nanosystems to optical and electronic applications. It will also study the mechanical properties of bulk Al due to ultrafine grain formations, which is an important question for this widely used light weight metal. Finally, it will also develop a method to calculate the high order many body perturbation theory which can significantly increase our accuracy to predict the material properties by first principle calculations.

Title: Lateral organization and inter-leaflet coupling in complex biological membranes

Principal Investigator: Xiaolin Cheng (Oak Ridge National Lab)

Co-Investigators: Jeremy Smith (University of Tennessee/Oak Ridge National Lab), Wei Yang (Florida State University), Goundla Srinivas (North Carolina A&T State University)

ALCC allocation: Processor Hours
Site: Oak Ridge Leadership Computing Facility (OLCF)
Allocation: 20,500,000 processor hours

Research Summary

This project supports the study of cell membranes and understanding the fundamental biological principle of compartmentalization which governs the assembly of complex biological systems from fundamental building blocks. Realizing a predictive understanding of biological systems requires moving beyond individual genes and proteins, to functional, interacting, mesoscopic systems compartmentalized in space and time. Membranes based on lipid bilayers define biological compartments and provide them with functional interfaces. Understanding how membranes perform these roles is a grand challenge that underpins many problems central to human health and DOE-BER's missions. Cell membranes display remarkable organization. For example, in the transverse dimension they are compositionally asymmetric, while in the lateral dimension they are believed to contain nanoscopic domains ("lipid rafts") critical to their function. In this regard, questions arise about the interplay between lipid rafts and compositional asymmetry, including how asymmetry is maintained, whether rafts bridge the two halves (leaflets) of asymmetric bilayers, and if so, how. To help resolve these questions, we propose to develop and perform large-scale molecular dynamics (MD) simulations on ORNL's leadership class supercomputer, TITAN. Complementing a wide range of experiments currently being performed on exactly the same systems at ORNL, the proposed simulations will help obtain detailed knowledge of lipid rafts and the transverse coupling of these nanoscopic domains. This work forms an integral part of a greater effort comprising of a DOE-BER sponsored Adaptive Biosystems Imaging (ABI) project and a Biomembranes Initiative supported by ORNL's Laboratory Directed R&D (LDRD) fund, which is aimed at integrating experiment and leadership-class computation to synergistically derive information on cellular communication, organization, and function of microbial and plant systems at atomistic level of detail. The outcome of this work will be a predictive understanding of complex membrane (protein) systems important for host-microbe interactions, a major target of which is the NIH Human Microbiome Project.

Title: The Weak Structure of Light Nuclei

Principal Investigator: Martin Savage (Institute for Nuclear Theory)

Co-Investigators: Silas Beane (University of Washington), Emmanuel Chang (Institute for Nuclear Theory), Zohreh Davoudi (Massachusetts Institute of Technology), William Detmold (Massachusetts Institute of Technology), Kostas Orginos (College of William and Mary and Jefferson Laboratory), Brian Tiburzi (City College of New York and RIKEN-BNL Research Center), Frank Winter (Jefferson Laboratory)

ALCC allocation: Processor Hours

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 40,000,000 processor hours

Research Summary

One of the biggest mysteries in physics is the observed asymmetry between matter and antimatter in our universe. Generating such an asymmetry from interactions between fundamental particles requires the non-conservation of a few basic symmetries during non-equilibrium dynamics in the earliest moments of the universe. While such non-conservations are present in the known fundamental laws of physics, the levels are insufficient to produce the observed matter-antimatter asymmetry. There are a number of current and planned experiments that are designed to explore these features, including the expected violation of time-reversal symmetry, of the combined operations of charge-conjugation and parity, of baryon number and of lepton number, in an effort to identify and understand the additional sources of non-conservation that must be present. A new thrust in nuclear physics research, as detailed in the 2015 Nuclear Physics Long Range Plan, is a program to search for signatures of one such interaction by exploiting unique possibilities for the properties of neutrinos. Neutrinos are a fundamental particle, and it is hypothesized that neutrino's are their own antimatter particles. If this is true, it could indirectly explain the asymmetry of matter and antimatter in the universe. In particular, a major experimental program to search for the "neutrinoless double-beta decay" ($0\nu\beta\beta$) of nuclei is planned. Simply observing such decays would unambiguously prove that neutrinos are their own antiparticles and that lepton-number is not conserved, and thereby a mechanism to generate the matter-antimatter asymmetry. Crucial to maximizing the scientific impacts of the experimental program are theoretical benchmarks conducted on high end computing resources. These calculations support the design of these new experiments, and to determine the details of the violation, once it is observed, the strength of the $\beta\beta$ and $0\nu\beta\beta$ processes in relevant nuclei. There is an urgent need to quantify the uncertainties that are currently present in theoretical calculations of these processes, and then to systematically reduce them. This project supports first principles calculations of the strengths of processes related to those exploited in the experimental program. The results of these calculations will be used as inputs and benchmarks for the necessary nuclear calculations of $0\nu\beta\beta$-decay rates and support experimental research aimed at understanding one of physics biggest unanswered questions.

Title: First Principles Design and Analysis of Energy-Efficient NanoElectronic Switches

Principal Investigator: Sefa Dag (Globalfoundries)

Co-Investigators: n/a

ALCC allocation: Processor Hours
Site: Argonne Leadership Computing Facility (ALCF)
Allocation: 10,000,000 processor hours

Research Summary

In integrated circuit (IC) technology, decades of progress were driven by the continued miniaturization of transistor dimensions to yield greater circuit density and functionality at lower cost and power per function. Currently, however, the chip industry is facing a crisis. Although transistor scaling has provided for enhanced performance, it has also resulted in increased in power per unit area of a chip. This is manifested in today's typical Complementary Metal-Oxide-Semiconductor (CMOS) processor, which operates at around the power density of a nuclear reactor. The fundamental reason for the rapid rise in the power density is that the supply voltage (VDD) used to drive the transistors has not scaled respectively with transistor density. Simply, transistor dimensions have continued to shrink to minute scales, but the voltage used to operate these transistors has plateaued. Since the MOSFET is not an ideal switch, the off-state leakage current (IOFF) is non-zero and hence dissipates power even when it is supposed to be off. In 2011, Intel's 22 nm CMOS node is the 1st commercially available bulk-FinFET technology and opens a new era of 3D CMOS for low power mobile electronics and continuously driving CMOS scaling and Moore's Law. However, at 10 nm node and beyond, new design requirements are essential for switching capabilities of FinFET for optimal energy efficiency. This allocation supports a new solution in design enablement by using the power of computational simulations and high performance computing clusters for cost effective and energy efficient switches for FinFET device technology.

Title: CASL CRUD-Induced Power Shift (CIPS) Analysis

Principal Investigator: Kevin Clarno (Oak Ridge National Laboratory)

Co-Investigators: Thomas Evans (Oak Ridge National Laboratory), Steven Hamilton (Oak Ridge National Laboratory), Tara Pandya (Oak Ridge National Laboratory), Ben Collins (Oak Ridge National Laboratory), Greg Davidson (Oak Ridge National Laboratory), Andrew Godfrey (Oak Ridge National Laboratory)

ALCC allocation: Processor Hours
Site: Oak Ridge Leadership Computing Facility (OLCF)
Allocation: 75,000,000 processor hours

Research Summary

High-fidelity simulation of nuclear reactors allows improved safety, reliability, and economics. In response to these opportunities, the Department of Energy established the Consortium for Advanced Simulation of Light Water Reactors (CASL) Innovation Hub for the modeling and simulation of nuclear reactors. To allow high-fidelity simulation of nuclear reactors, CASL is developing, applying, and deploying the Virtual Environment for Reactor Applications (VERA), which incorporates science-based models, state-of-the-art numerical methods, modern computational science and engineering practices, uncertainty quantification (UQ), and validation against data from operating pressurized water reactors (PWRs), single-effect experiments, and integral tests. In 2015, the VERA Core Simulator was used to model an operational transient that occurred at the Watts Bar Nuclear Power Plant during cycle 7 (2005-6) that resulted in a dramatic shift of the axial power distribution and an extended outage costing tens of millions of dollars. To avoid these outages, many utilities operate reactors in an excessively conservative manner, costing the utility, thus consumer, millions of dollars.

This first-of-a-kind simulation provided a high-resolution simulation of the reactor and CRUD (Corrosion-Related Unidentified Deposits) growth at fidelity never seen before. These results show the potential of advanced simulation to predict the onset of a CRUD-Induced Power Shift (CIPS), allowing operators to avoid these conditions. In 2016, CASL will use the Shift Monte Carlo code to evaluate discrepancies in VERA to discover possible sources of error as simulations are performed though multiple cycles of multiple nuclear power plants. By quantifying the error, the uncertainty in the simulation can be quantified, improving its use as a decision-making tool.

This allocation supports the modeling of multiple cycles of four operating nuclear power reactors. These simulations will use an ultra-high-fidelity Monte Carlo neutronics component in VERA-CS to validate the standard solver. The Monte Carlo package will also be used to simulate gamma heating and do validation for VERA-CS depletion.

Title: Accurate Predictions of Properties of Energy Materials with High Throughput Hybrid Functional Calculations

Principal Investigator: Christopher Wolverton (Northwestern University)

Co-Investigators: n/a

ALCC allocation: Processor Hours

Site: National Energy Research Scientific Computing Center (NERSC)

Allocation: 45,000,000 processor hours

Research Summary

The goal of this project is to develop a high accuracy database of materials and properties which is open to the public. Density Functional Theory (DFT) was developed in the 1960s as a tool to efficiently compute the properties of materials directly from quantum mechanics. Many scientists have used this theory to predict materials behavior without needing to synthesize them in the laboratory. Recently, several groups have used DFT and high performance computing to create large databases of computed material properties, which have proven to be powerful tools in the design and discovery of new materials. However useful, the accuracy of DFT is limited by its core approximation: that the interaction between electrons can expressed as functions of the electron density. In recent years, research has shown that a new class of approximations, namely those developed by Heyd, Scuseria, and Ernzerhof (HSE), can improve the accuracy of DFT calculations, but at a significantly higher computational cost. This project will develop the tools necessary to employ hybrid functionals to predict the properties materials on a large scale. As the magnitude of improvement between conventional and HSE functionals is not systematic, we aim to develop evidence based rules to identify which materials for which using HSE functionals, so that we can utilize this computationally-expensive technique only when necessary. This project will develop these rules by using machine learning on a dataset of several thousand HSE calculations that we will create in this work. The project team will use these rules to selectively re-evaluate materials in the Open Quantum Materials Database (oqmd.org), and make the improved results freely available to the entire materials science community.

Title: Multiscale Gyrokinetic Simulation of Reactor Relevant Tokamak Discharges: Understanding the Implications of Cross-Scale Turbulence Coupling in ITER and Beyond

Principal Investigator: Christopher Holland (University of California San Diego)

Co-Investigators: Nathan Howard (Massachusetts Insitutute of Technology Plasma Science and Fusion Center), Jeff Candy (General Atomics)

ALCC allocation: Processor Hours

Site: Oak Ridge Leadership Computing Facility (OLCF)
Allocation: 50,000,000 processor hours
Site: National Energy Research Scientific Computing Center (NERSC)
Allocation: 60,000,000 processor hours

Research Summary

Nuclear fusion has the potential to provide a near limitless supply of usable clean energy. For this reason, scientists have devoted years of experimental, theoretical, and computational research to overcoming the significant challenges faced in containing and controlling high temperature plasmas that produce fusion energy. Microturbulence on scales one hundred times smaller than the plasma size can have macroscopic effects, requiring detailed multiscale physics simulations to develop theoretical understanding and predictive models of the turbulence. In tokamaks, (the leading approach to magnetic confinement based fusion devices), the turbulent transport of mass, momentum and energy by microturbulence is often the primary determinant of the level of total plasma confinement and performance achieved. It is therefore essential to be able to predict this microturbulence and its associated transport in reactor-relevant plasma regimes. In its simplest description, the plasma is made of two primary components, ions and electrons, both of which move along the confining magnetic fields in helical spiral orbits. Because the electrons are much lighter than the ions, the characteristic scale of their gyromotion is roughly sixty times smaller than that of the ions. Historically, research has focused on microturbulence occurring on the (relatively) larger ion gyroradius scales – either neglecting the electron-scale fluctuations or treating electrons as independent of the ion dynamics. However, this approach has been unable to explain some experimental observations, and recent work suggests that neglecting electron-scale dynamics and their self-consistent coupling with ion-scale physics in simulations of these regimes may lead to fundamentally flawed predictions. This project supports more realistic simulations of fusion plasmas in order to better understand multiscale plasma microturbulence. The study will use multiscale gyrokinetic simulations that resolve both ion- and electron-scale dynamics using realistic ion-to-electron mass ratio, experimentally-derived inputs, and electromagnetic effects to compare predicted transport in multiple channels directly against experiments performed on the DIII-D and Alcator C-Mod tokamaks. The results of this work will shed light on whether cross-scale coupling is likely play an important role in reactor-relevant regimes, helping to define the requirements for reliable prediction of ITER and future fusion reactors.

Title: Evaluating Turbomachinery Aerodynamic Cross-Coupling Coefficients Using CFD

Principal Investigator: Ravi Srinivasan (Dresser-Rand)

Co-Investigators: Mark Krzysztopik (Dresser-Rand), Paul Brown (Dresser-Rand), Alain Demeulenaere (Numeca, USA), David Gutzwiller (Numeca, USA)

ALCC allocation: Processor Hours

Site: Oak Ridge Leadership Computing Facility (OLCF)

Allocation: 12,000,000 processor hours

Research Summary

Compression of carbon dioxide plays a key role in two Department of Energy technology areas: carbon capture and sequestration, and power systems using super-critical CO_2 as a working fluid. Efficient and reliable compression of CO_2 requires re-evaluating existing rotary machinery concepts. Stable operation of rotating machinery is essential to avoiding expensive service operations, and must be considered in the design phase of complex turbomachinery systems. Currently, engineers utilize a set of empirically derived correlations to evaluate the stability regimes. However, these correlations are constrained to systems with large stability margins, specific geometry configurations, and process type. New technology developments, including high-pressure ratio, single-stage machines with higher mole weight working fluids and large-coefficient designs, are compelling engineers to reevaluate these correlations.

The goal of this project is to evaluate models used by industry for stability. The advanced simulation capabilities of Titan will be used to perform baseline full-annulus, time-accurate computational fluid dynamic (CFD) simulations of unshrouded centrifugal impellers with varying eccentricities (axis of rotation offset from machine axis) and whirl angular velocity. The results will be compared to aerodynamic cross-coupling coefficients from various empirical models such as the Alford equation, API-Wachel, and SwRI-Wachel. These methods were developed for axial machines or shrouded compressors. Direct comparison will highlight the strengths, deficiencies and applicability of these empirical models. The data obtained from these simulations will help in developing models for unshrouded compressor designs. This understanding will be published for the entire technical community, and will be used to accelerate the development of super-critical CO_2 turbomachinery for both carbon sequestration and advanced power generation.

Title: Molecular Dynamics Studies of Biomass Degradation in Biofuel Production

Principal Investigator: Klaus Schulten (University of Illinois at Urbana-Champaign)

Co-Investigators: Rafael Bernardi (University of Illinois at Urbana-Champaign)

ALCC allocation: Processor Hours
Site: Argonne Leadership Computing Facility (ALCF)
Allocation: 50,000,000 processor hours

Research Summary

Biofuels are a well-known carbon neutral alternative to fossil fuels. Cellulose is an important molecular building block of many plants but is resistant to easy conversion to biofuels. Development of an effective method of converting cellulose into usable fuel would make biofuel production significantly more efficient and cost competitive with fossil fuels. Our group recently reported that the very stable cellulosomal complex exhibits one of the most mechanically robust protein-protein interactions reported thus far, pointing towards new mechanically stable artificial multi-component biocatalysts for industrial applications, including production of second generation biofuels. Adopting the strategy which some bacteria are employing and use several enzymes synergistically in large enzymatic complexes, namely cellulosomes, offers a promising approach to reduce biofuel production cost. This project supports use of molecular dynamics tools to investigate the cellulosomes at the molecular level. The outcomes of this project will be an important step towards developing effective methods of cellulose breakdown and is a step along the path towards low cost biofuels.

Title: Modeling Partial Fuel Stratification in a Homogeneous Charge Compression Ignition Engine using General Purpose Graphical Processing Unit (GPGPU) enabled CFD tool

Principal Investigator: Ronald Grover (General Motors)

Co-Investigators: Seung Hwan Keum (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), Sameera Wijeyakulasuriya (Convergent Science), Eric Pomraning (Convergent Science)

ALCC allocation: Processor Hours
Site: Oak Ridge Leadership Computing Facility (OLCF)
Allocation: 16,000,000 processor hours

Research Summary

Reducing emissions from internal combustion engines while improving fuel economy is a challenging task, requiring new strategies such as lean burn engines. In lean burn engines, the air-fuel ratio is above stoichiometric, meaning that there is excess oxygen available relative to the mass of fuel. This approach offers lower combustion temperatures, and reduced formation of NOx, a major pollutant. One specific approach to lean combustion is Homogeneous Charge Compression Ignition (HCCI). Understanding the effects of thermal and mixture stratification is required to avoid pre-ignition of the fuel, more commonly known as "knock," and optimize HCCI engines for use.

Simulation of HCCI, and all other combustion engine configurations, requires capturing multiple complex physics, including fuel vaporization, turbulent mixing and heat transfer, and complex chemical kinetics, across a range of length and time scales. Modeling of the chemical kinetics is particularly challenging, since thousands of chemical species may be present during the combustion process. In many cases, surrogate fuels, which simplify the chemistry and physical behavior of real fuels to allow simulation, are used to study combustion.

The current work takes a two-step approach using CONVERGE, a commercially available Computational Fluid Dynamics (CFD) code, along with a customized user defined function (UDF) that take advantages of Titan's Graphical Processor Units (GPUs) to solve chemical kinetics. In the first phase of the work, detailed kinetics simulations will be used to evaluate publicly available surrogate components, which are mostly focused on matching ignition of homogeneous mixtures. In the second step, surrogate components will be validated against vaporization measurements to best model the fuel mixing process. These results will be used to capture the in-cylinder mixture stratification and combustion processes with sufficient accuracy to resolve the underlying physics leading to HCCI engine knock reduction.

Title: Adjoint Based Optimization via Large Eddy Simulation of a Fundamental Turbine Stator-Rotor

Principal Investigator: Qiqi Wang (Massachusetts Institute of Technology)

Co-Investigators: Gregory Laskowski (GE Aviation), Umesh Paliath (GE Global Research)

ALCC allocation: Processor Hours

Site: Argonne Leadership Computing Facility (ALCF)

Allocation: 15,000,000 processor hours

Research Summary

Numeric simulations allow the study of physical systems and engineering concepts at conditions and scales beyond the reach of traditional experiments. One approach to achieving higher fidelity is to build and calibrate computational models by solving an inverse problem to match available measurements and data. MIT is currently developing optimization-based computational tools for solving such data-intensive inverse problems. The optimization-based computational tools are designed to work together with Large Eddy Simulation (LES), a computational method for solving turbulent flows that depends on accurate modeling of energy transfer between turbulent length scales.

Turbulent flows appear in engineering applications ranging from wind energy to internal combustion engines. To demonstrate the applicability of these tools to engineering design, MIT has partnered with GE Aviation. In this project, MIT and GE will partner to use these tools to simulate flow over jet engine turbine components, using a fundamental vane/blade geometry described in the literature by Kopriva and Laskowski. Applying these high-fidelity simulation and optimization methods to jet engine design has the potential to significantly reduce fuel consumption, leading to billions of dollars in fuel savings and a commensurate decrease in carbon emissions.

Title: Predictive Simulations of Complex Flow in Wind Farms

Principal Investigator: Matthew Barone (Sandia National Laboratories)

Co-Investigators: Michael Sprague (National Renewable Energy Laboratory), Stefan Domino (Sandia National Laboratories), Matthew Churchfield (National Renewable Energy Laboratory)

ALCC allocation: Processor Hours

Site: National Energy Research Scientific Computing Center (NERSC)

Allocation: 10,700,000 processor hours

Research Summary

The U.S. Department of Energy's Wind Power Program is targeting significant reductions in the cost of wind energy through its Atmosphere to Electrons (A2e) program. This initiative, which includes researchers from Sandia National Labs and the National Renewable Energy Laboratory, seeks to improve fundamental understanding of the complex physics governing wind flow into and through wind plants. At the core of this program is a high-fidelity modeling initiative to develop a computational capability that is truly predictive for wind farm flow physics.

As part of this effort, the research team will use peta-scale computing to develop and demonstrate computational fluid dynamics (CFD) simulations of whole wind plants. These simulations will use large eddy simulations (LES), a high-fidelity technique for the simulation of turbulence that, due to its computational cost, is rarely used in simulation of engineering systems at this scale. These simulations will be performed using a modified version of Nalu, a Department of Energy open source CFD code. Changes to Nalu for these simulations will include the implementation and testing of a higher- order control-volume finite element (CVFEM) scheme and implementation of actuator line models for wind turbine blades.

The Nalu code will be used to simulate the flow through an actual offshore wind plant, at a scale 3 to 4 times that of previous simulations. The simulations will have the following objectives: 1) verification of the behavior of the LES model with respect to mesh resolution and order of the numerical scheme; 2) validation of the results through comparison with available power performance data for the wind plant; 3) quantification of the parallel efficiency and scalability of the simulation code and the dependence of these performance metrics on the numerical scheme.

The immediate impact of this project will be to demonstrate a next-generation wind farm LES capability, upon which a new, massively parallel, open source solver to support wind plant simulation research can be built. The long-term goal of this effort is to integrate high-fidelity simulations into the design and optimization of wind plants, ultimately decreasing the cost of wind energy.