Type: New PW

Title: Continent-scale InSAR velocity fields for volcanology and tectonics

Principal Investigator: falk amelung (University of Miami)

Co-Investigators:

Field of Science: Tectonics; Volcanology and Mantle Geochemistry

Abstract:

With this Pathways allocation we will obtain InSAR velocity fields using Sentinel-1 data for both the Indonesian volcanic arc and the Tibetan Plateau. These data will be used (1) to test the hypothesis that in Indonesia the is a temporal correlation between volcanic unrest and precipitation with more unrest occurring during the rainy season because infiltrated rainwater weakens the rock and (2) to constrain models of continental deformation.

Type: New PW

Title: Planet-disk interaction in three dimensions

Principal Investigator: Jaehan Bae (Carnegie Institution of Washington)

Co-Investigators:

Field of Science: Stellar Astronomy and Astrophysics

Abstract:

Recent observations of circumstellar disks around young, forming stars have revealed a plethora of structures. One of the most exciting possibilities is the interaction between the disk and planets embedded therein. In order to improve our understanding of planet-disk interaction and help better interpret observations taken at unprecedented spatial resolution and sensitivity, we propose to carry out three-dimensional numerical simulations of planet-disk interaction.

Type: New PW
Title: Using Frontera GPU computing to accelerate developments in interpretable machine learning

Principal Investigator: David Benkeser (Emory University)
Co-Investigators:

Field of Science: Statistics and Probability

Abstract:
Machine learning for health care and public health decision making faces a significant tradeoff: accuracy vs. interpretability. Simple rules (e.g., based on logistic regression) for predicting outcomes or providing treatment recommendations are easy-to interpret, but often suffer in terms of performance. More complex rules, based on black-box algorithms like deep learning, may predict more accurately, but are difficult to interpret. We were recently funded by a three year grant from the National Science Foundation to develop methods that bridge this gap and provide machine learning that is both accurate and interpretable, making it amenable for deployment in clinical and public health settings. GPU computing could play a vital role in achieving the grant's objectives of generating fast and scalable implementations of our proposed algorithms.

Type: New PW
Title: Development of multiphase fluid-structure interaction methods

Principal Investigator: Amneet Pal Bhalla (San Diego State University)
Co-Investigators: KAUSTUBH KHEDKAR (San Diego State University)

Field of Science: Applied Mathematics

Abstract:
Applications involving fluid-structure interaction (FSI) are ubiquitous in natural and engineering processes that can range from bacterial swimming to the interaction of waves with ships. FSI also plays a vital role in new methodological approaches for modeling energy harvesting devices such as wave energy converters (WEC) and simulating 3D printing processes. Executing these applications at large-scale with the traditional body-fitted grid approach becomes prohibitively expensive due to the re-meshing requirement of the numerical scheme to conform to the deforming or moving body in the domain. On the other hand, the immersed boundary (IB) methods discretize the computational domain using Cartesian grids, and the dynamics of the immersed structure is resolved by modifying the underlying equations of motion. Thus, the computational cost involved in resolving FSI is substantially reduced by adopting the IB approach. In our group, we study the interaction of surface gravity waves with moving structures such as wave energy converter devices and naval ships, as well as model the stereo-lithography 3D printing process using state-of-the-art immersed boundary method and adaptive mesh refinement based software infrastructure named IBAMR, which is a National Science Foundation (NSF - OAC 1450327, OAC 1450374 and OAC 1931516) sponsored advanced cyberinfrastructure (CI) project. In the following sections, we give a brief review of our research work and the questions that we aim to address. Our proposed project is funded by NSF award 1931368.

Type: New PW

Title: Real-time High Resolution Ensemble Numerical Weather Forecasts Using SAR-FV3 for the Hydrometeorology Testbed - Ramping toward Exascale

Principal Investigator: Keith Brewster (University of Oklahoma)
Co-Investigators: Nathan Snook (University of Oklahoma); Tim Supinie (University of Oklahoma)

Field of Science: Meteorology

Abstract:
The Center for Analysis and Prediction of Storms (CAPS) at the University of Oklahoma (OU) has, for several years, been testing high-resolution numerical weather prediction (NWP) with support from, and in collaboration with, the National Oceanic and Atmospheric Administration National Weather Service (NOAA NWS) Testbeds. The NOAA Testbeds evaluate, both subjectively and objectively, new tools for improving forecasts of high-impact weather events across the Contiguous United States (CONUS). This project extends that work by ramping high-resolution NWP ensembles toward Exascale in support of the Hydrometeorology Testbed Winter Weather Experiment (WWE) 2020-21. CAPS plans to produce an experimental high resolution NWP ensemble in support of the WWE as part of a project to develop and extend NOAA's next-generation Unified Forecasting System (UFS) in a high-resolution regional weather configuration, the Standalone Regional Finite Volume-3 (SAR-FV3). As part of the WWE CAPS will develop and test ensemble post-processing tools, including machine learning, for improving forecasts of heavy snow and other high-impact winter weather. The project will extend CAPS prior successful execution of SAR-FV3 ensembles on large scale XSEDE resources (TACC Stampede2) and move toward the Exascale on Frontera by augmenting our previous ensemble with an additional member with 1-km grid spacing requiring approximately 27 times the computations of the current 3-km grid-spacing. Following this Pathways project, CAPS intends to apply for a LARC Frontera allocation for subsequent NOAA Testbed work that would free Stampede2 from running 13-member ensemble forecasts in a high priority queue needed for real-time forecast turnaround.

Type: New PW
Title: Regulation of cardiac small-conductance calcium-activated potassium channel

Principal Investigator: Nipavan Chiamvimonvat (University of California, Davis)
Co-Investigators: Ryan Woltz (University of California, Davis)

Field of Science: Biochemistry and Molecular Structure and Function

Abstract:
Our laboratory was the first to identify the critical roles of small conductance Ca2+-activated K+ channels (SK channels) in cardiac repolarization. Detailed understanding of the mechanistic underpinning of channel activation has overarching implications from the regulation of cardiac excitability, cardiac arrhythmias to drug targeting. The current research in the laboratory is supported by three NIH R01 grants with strong and longstanding collaborative efforts from well-established investigators in the field. SK2 channel is activated by Ca2+ binding to calmodulin (CaM), that are bound to the CaM binding domain (CaMBD) in the C-terminus of the channel. In addition, the presence of phosphatidylinositol 4,5-bisphosphate (PIP2) in the plasma membrane is required for the activation of SK2 channel. While CaM binding to SK2 channel and Ca2+ activation mechanism is supported based on a recent cryo-EM structure of a full SK4-CaM complex, the structure documenting SK2-CaM bound to PIP2 remains uncertain since only SK2 peptide was used. Specifically, the putative interacting amino acids from the crystal structure (Zhang, et al 2014) contradict the more recent cryo-EM full-channel structure (Lee, MacKinnon et al 2018). There are currently molecular dynamics (MD) simulations detailing activation mechanisms for KcsA, Kir, hERG, and Shaker K+ channels but no simulation exists for any members in the SK family.
Here, we propose to explore and compare potential PIP2 binding sites on the SK2-CaM complex and the possible SK2 channel activation with and without PIP2 in a plasma membrane using Umbrella Sampling (US) MD simulations. While extensive work has been performed locally on CPU clusters (docking PIP2 to a static model) and MD desktops (CG-MD with complex membrane simulations), all of which provide intriguing, but inconclusive, results. We propose to explore a full atom resolution simulation of the ~245,000 atom SK2-CaM-PIP2 system. There are two main objectives: 1) Generate an US-MD hSK2-CaM system to compare the binding energies of the 2 most likely PIP2 binding sites. 2) Identify differences in the PIP2 binding in different states of the SK2 channel with/without PIP2 in the membrane with a focus on identifying PIP2 possible binding mechanisms. We plan to continue local work described above to explore all possibilities and factors that could lead to SK2 activation by PIP2, however, a full atom resolution simulation is critical to identify specific binding orientations and energies. We equilibrated and tested systems using allocations from the Comet cluster at XSEDE. The US-MD will be equilibrated for 10 ns and productions runs are estimated to be at 20 ns on Frontera. The results will serve as a proof of concept to use US-MD to explore less established but potential binding sites of PIP2 not explored in this proposal. This second phase of the project will be accomplished by applying for an LRAC allocation once our scaling up procedure has been optimized and proven.

Type: New PW
Title: Development of Universal Interatomic Potentials Based on Graph Convolutional Neural Networks and First-Principles Calculations

Principal Investigator: Ju Li (Massachusetts Institute of Technology)
Co-Investigators: Qing-Jie Li (Massachusetts Institute of Technology)

Field of Science: Materials Research

Abstract:
Significantly extending the time-/length-scale while keeping quantum mechanical accuracy is a grand challenge in the field of molecular dynamics (MD) simulations. Recent advances in machine learning show promising progresses in accurately mapping the high-dimensional potential energy surface from quantum mechanical calculations to artificial neural network models. However, poor transferability remains as Achilles' heel of existing neural network models to enable general purpose simulations. A universal interatomic potential model capable of accurately describing arbitrary elements, structures, transformations, and chemical reactions is thus highly needed. In this proposed work, we will extend TeaNet (tensor embedded atom network), a graph convolution neural network architecture demonstrating robust transferability for elements 1-18, to all elements across the periodic table. Specifically, we propose to scale up TeaNet in terms of model and data parallelization such that it can efficiently handle training dataset size on the order of 10^6- 10^7. Meanwhile, residual learning, optimized graph convolutions, and active learning scheme will be implemented to enhance extreme short-range interactions, long-range electrostatic interactions and training data sampling, respectively. Such universal interatomic potential will be extensively validated and applied to simulating various systems such as advanced structural alloys, complex molten salts, high-performance battery materials and so forth.

Type: New PW

Title: Advancing understanding of aerosol-cloud feedback using the world's first global climate model with explicit boundary layer turbulence

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

Co-Investigators:

Field of Science: Atmospheric Sciences

Abstract:

Aerosols, meaning tiny particles suspended in the atmosphere, play a key role in cloud formation, as cloud droplets and ice particles are produced when water vapor condenses onto aerosols. When more aerosols are present clouds tend to have a larger number of smaller droplets, making them brighter and more effective in reflecting sunlight back to space. Thus increases in aerosol amount due to industrial activity can increase the brightness of clouds, resulting in a cooling effect on climate. The extent to which the global temperature increase from greenhouse warming has been offset by human-induced radiative forcing from aerosol-cloud-interactions (RFaci) is an important and unsolved problem in climate science.

One obstacle to progress on RFaci is the difficulty of performing computer simulations which explicitly represent cloud properties yet cover the whole earth, so that global climatic effects can be assessed. Cloud motions are turbulent and require models with grid points spaced a fraction of a kilometer apart, while global model grid spacing is typically tens to hundreds of kilometers. To bridge this scale gap the PIs have developed an ultraparameterized (UP) model, meaning a global model with coarse grid spacing in which each grid box contains a fine-scale cloud resolving model with a domain size much smaller than the grid box. The model is challenging both scientifically and computationally, and the project includes a concerted effort to improve computational efficiency to make simulations practical.

The research addresses several specific questions regarding RFaci. One question is why climate models tend to overestimate RFaci compared to estimates from satellites, in some cases by a factor of two. Comparisons between the UP model and satellite observations will be facilitated by a nudging methodology, in which external forcing is used to constrain the simulated weather patterns to match the days when the satellite observations were taken. The nudging minimizes differences between simulated and satellite-estimated RFaci due to incorrect simulation of large-scale circulation features, allowing attribution of differences to aerosol-cloud interactions.

The work has broader impacts due to the societal implications of high versus low RFaci: if the cooling effect of industrially-driven RFaci is large, the strength of greenhouse warming must be at the high end of current estimates in order to explain the warming seen over the past century. Likewise, if industrial RFaci cooling was small over the last century, the sensitivity of global temperature to greenhouse gas increase is likely to be on the lower end of its estimated range. RFaci is thus among the largest uncertainties in determining climate sensitivity and the severity of climate change impacts. In addition, software developed under the project is made available to the research community, in part through a version of the Community Earth System Model. The project provides support and training for a postdoctoral research scholar, thereby providing workforce development.

This award reflects NSF's statutory mission and has been deemed worthy of support through evaluation using the Foundation's intellectual merit and broader impacts review criteria.

Type: New PW
Title: First-principles study of many-body interactions and excited-state properties in two-dimensional systems

Principal Investigator: Diana Qiu (Yale University)
Co-Investigators:

Field of Science: Condensed Matter Physics

Abstract:
Two-dimensional (2D) materials are the subject of significant ongoing research for technological applications in electronics, optoelectronics, valleytronics, and energy. Under optical excitation, 2D materials in general exhibit strong electron-hole interactions due to the strong spatial variations in the screening of the Coulomb interaction, as well as reduced total screening and electronic confinement resulting from the dimensionality reduction. It is thus very important to accurately calculate and thoroughly understand the many-electron properties of 2D materials for developing devices based on such materials. Our group is actively focused on studying the effect of many-body interactions on the excited state properties of materials in reduced dimensions and complex functional materials. Particular emphasis is placed on engineering optical excitations of 2D materials by inducing and controlling structural/chemical defects, aiming to provide an understanding of fundamental processes such as charge/energy transfer between heterointerfaces and guidelines for designing 2D systems for technological applications. Our group is also engaged in developing high-performance first-principles quantum physics methods based on the GW-Bethe-Salpeter-Equation (BSE) approach to calculate many-electron interaction effects for a broad range of material systems with a current focus on novel efficient algorithms for solving the BSE including non-uniform spatial sampling schemes, finite momentum excitons, time-dependent approaches, and electron-hole interactions in core-level spectra. Two specific research topics will be investigated on the Frontera Pathway allocation. In the first project, we will use the GW-BSE method to study how optical excitations and valley-selectivity in monolayer transition metal dichalcogenides (TMDs) can be tuned by proximity effects from the adsorption of chiral molecules and induction of structural defects. In the second project, we will propose a new many-body downfolding approach for calculating exciton effects on core-level excitations and apply this new technique to the X-ray absorption spectra of monolayer TMDs, where exciton effects are expected to be large.

Type: New PW

Title: Understanding and Modeling the Role of Horizontal Heterogeneity on the Dynamics of the Nocturnal Boundary Layer Across Scales

Principal Investigator: Rob Stoll (University of Utah)
Co-Investigators:

Field of Science: Atmospheric Sciences

Abstract:
A critical component of weather and climate models is the determination of the exchange rate of momentum, heat, and moisture between the land surface and the atmosphere. These exchanges have a profound impact on the accuracy of model predictions. This project seeks to use high-resolution simulations of fluid flow over smooth and rough surfaces to study the fundamental processes that govern these exchanges with the ultimate goal of improving our ability to simulate weather and climate. Uniquely, this project focuses on these exchanges at night when cooling from the ground surface lowers the exchange rates and dampens atmospheric flow motions. This complicates the physics of atmospheric transport near the land surface.

Type: New PW
Title: Massively parallel simulations of binary black hole intermediate-mass-ratio inspirals : Computational contributions for LIGO/LISA

Principal Investigator: Hari Sundar (University of Utah)
Co-Investigators:

Field of Science: Gravitational Physics; Computer and Computation Research

Abstract:
Existing codes for Numerical relativity and relativistic magnetohydrodynamics do not scale well on modern heterogeneous clusters and this is a major impediment towards scientific progress in these areas. In recent work we have developed a highly scalable adaptive numerical relativity framework, Dendro-GR. The framework provides a flexible high-level interface for numerical relativity using problem description using a Symbolic interface in symbolic python. The problem is then automatically discretized and architecture specific code is automatically generated. We currently support distributed memory parallelism via MPI, shared memory parallelism via OpenMP, SIMD vectorization using `avx2` and `avx512` and support for Nvidia GPUs using CUDA. We have demonstrated the excellent scalability of our framework and the code generation capabilities on OLCF's Titan supercomputing using upto 8192 GPUs as well as TACC's Frontera via a DD allocation.

In this work, we are requesting a pathways allocation to verify and evaluate our code and code-generation framework on Frontera, in preparation for an Leadership Computing Resource Allocations (LRAC) allocation request. The LRAC proposal will focus on estimating waveforms for intermediate mass ratio mergers. Compared to existing codes for numerical relativity, DendroGR has high levels of adaptivity in both space and time, and is optimized for modern supercomputers, resulting in simulations that are 1000x faster compared to community codes like the Einstein Toolkit. With the pathways allocation, we wish to demonstrate both the accuracy of our methods and codes, as well as the speedup, especially for simulating large mass ratio mergers.

Type: New PW
Title: Cyberinfrastructure hybrid for scalable machine learning and visualization using genomic and geospatial datasets

Principal Investigator: Tyson Swetnam (University of Arizona)
Co-Investigators: Anne Thessen (Oregon State University); Arun Ross (Michigan State University); Remco Chang (Tufts University)

Field of Science: Ecological Studies

Abstract:
To mitigate the effects of climate change on public health and conservation, we need to better understand the dynamic interplay between biological processes and environmental effects. The state-of-the-art, which has led to many important discoveries, utilizes numerical or statistical models for making predictions or performing in silico experimentation, but these techniques struggle to capture the nonlinear response of natural systems. Machine learning (ML) methods are better able to cope with non-linearity and have been used successfully in biological applications, but several barriers still exist, including the opaque nature of the algorithm output and the absence of ML-ready data. Here, we propose to significantly advance technologies in ML and create a new interdisciplinary field, computational eco-genomics. We propose to do this by (a) designing ML techniques for encoding heterogeneous genomic and environmental data, and mapping them to multi-level phenotypic traits, (b) reducing the amount of necessary training data, and (c) developing interactive visualizations to better interpret ML models and their outputs.

We utilize next generation GPU nodes to create ML models and to process and visualize genomic data, as well as high resolution geospatial and 3D data from small unoccupied aerial systems (sUAS) i.e., drones.

Type: New PW
Title: Fast and Accurate Deep Learning on Distributed Systems

Principal Investigator: Zhao Zhang (University of Texas at Austin)
Co-Investigators: Yang You (University of California, Berkeley)

Field of Science: Computer and Computation Research

Abstract:
In the past three years, we observed that the training time of ResNet-50 dropped from 29 hours to 67.1 seconds. However, there is still a huge gap between High Performance Computing (HPC) and ML. On the one hand, we had powerful supercomputers that could execute roughly 10^18 floating point operations per second. On the other hand, we still can not even make full use of 10% of this computational power for some applications like ImageNet/ResNet-50. The reason is that supercomputers need an extremely high parallelism to reach their peak performance. However, the high parallelism leads to a bad convergence for DNN optimizers. To solve this problem, I proposed the LARS optimizer and LAMB optimizer. These new methods enable ImageNet training with ResNet-50 to scale on thousands of chips without losing accuracy. In fact, all the state-of-the-art ImageNet training speed records were made possible by LARS since December of 2017. LARS became an industry metric in MLperf v0.6. Nevertheless, these approaches can not satisfy the needs of future applications. The current optimization approaches still fail to fully utilize the increasing performance of supercomputers. The increasing parallelism requires future optimizers to continue to scale up. For example, a TPU Pod needs at least a batch size of 256K for an ImageNet-level dataset to reach its peak performance. On the other hand, researchers are collecting larger datasets by the auto-labeling technique. As a result, we will have enormous parallelism if we can put all the available data samples in the memory of a future supercomputer. One possibly interesting method is a variant of Gradient Descent (GD), which scales linearly with the dataset. The number of iterations of GD can be O(1). Moreover, the communication often brings a significant overhead to modern distributed systems. The number of network communication messages of GD can be O(1). However, the concern is that the SGD achieves a much better testing accuracy than the GD-variant optimizer. The GD-variant optimizer suffers a poor generalization performance on modern machine learning datasets. I would like to investigate various optimization techniques and invent new auto-learning techniques to improve the generalization performance of the GD-variant optimizer. Another direction is studying the second-order optimization technique, which can potentially improve the scaling efficiency.

Type: New LSCP

Title: LSCP Support for the SolFER DRIVE Center

Principal Investigator: William Daughton (Los Alamos National Laboratory)

Co-Investigators:

Field of Science: Solar Terrestrial Research

Abstract:

SolFER (Solar Flare Energy Release) is a multi-institution collaboration funded by NASA's DRIVE program to study the explosive release of magnetic energy in solar flares and the associated production of energetic particles. It brings together 12 institutions and > 50 researchers with expertise in observation, theory, numerical modeling, and computer science. This LSCP proposal seeks computational resources for the theory and modeling aspects of SolFER as well as access to an archival system to facilitate comparisons of simulations with observations. Our team employs a rich set of numerical tools (particle-in-cell, hybrid, mag- netohydrodynamics, and more) and has substantial experience in performing simulations on leadership supercomputing facilities. In the three-year project, we will explore several key aspects of solar flare energy release, including onset, the heating and acceleration of both ions and electrons, and energetic particle transport in the flaring region.

Type: New LSCP

Title: Computational Study of Astrophysical Plasmas

Principal Investigator: Kenichi Nishikawa (Alabama A&M University)
Co-Investigators: Nicholas MacDonald (Max-Planck-Institut fur Radioastronomie); Ioana Dutan (Astronomical Institute of the Romanian Academy); Christoph Koehn (Technical University of Denmark); Michael Watson (Fisk University); Martin Pohl (Universit¨at Potsdam); Yosuke Mizuno (Goethe Universitat Frankfurt am Main); Kouichi Hirotani (Academia Sinica, Institute of Astronomy & Astrophysics); Athina Meli (University of Liege); Jacek Niemiec (Institute of Nuclear Physics, Polish Academy of Sciences)

Field of Science: Extragalactic Astronomy and Cosmology

Abstract:
Astrophysical jets are dynamically magnetized plasma flows that are launched most likely in regions where the Poynting (magnetic field energy) flux dominates over the particle flux. We will investigate how Poynting flux of relativistic jets dissipates into kinetic energy to accelerate particles rapidly by studying the interaction of the particles with the immediate plasma environment on the microscopic scales. We will also investigate synthetic spectra emitted from relativistic jets, polarity images through radiative transfer, particle acceleration related to electric discharges and dynamics of black hole. Our research team employs new computational tools to investigate the important topics of magnetic reconnection, nonlinearly generated turbulence, and associated particle acceleration in collisionless astrophysical relativistic plasmas. Simulations of relativistic jets injected into ambient plasmas with helical magnetic fields have demonstrated possible signatures of magnetic reconnection and nonlinearly generated turbulence. Based on our promising and published results using Comet, Bridges and Pleiades, this research project will further explore relativistic particle acceleration associated with magnetic reconnection in jets with helical magnetic fields including nonlinearly generated turbulence and other research topics using Frontera. There is a need to explore these processes in more realistic environments, with the best computer power. Therefore, our aimed systematic study in a much more realistic astrophysical context will elucidate the important processes and mechanisms of reconnection and particle acceleration in active galactic nuclei and gamma-ray burst
jets. In this research effort, we will analyze the evolution of jets through the use of our extensive visualization tools which provide the location of reconnection and its associated phenomena such as particle acceleration. Our initial simulation results and published work show complicated structures of jet evolution due to combined kinetic and kink-like instabilities in the relativistic jets containing helical magnetic fields. The magnetic field structures generated by kinetic and kink-like instabilities determine where and when reconnection occurs. More computational power using Frontera will help our team to make further progresses with our simulation investigations. Often explosive in nature, magnetic reconnection enables rapid release of magnetic energy stored in the jets. Our overall goal is to integrate interdisciplinary scientific aspects to study basic plasma physics and theoretical astrophysics and to develop new numerical methods to simulate the microphysical processes responsible for reconnection, turbulence, and high-energy particle acceleration. At Alabama A&M University (AAMU) we will educate students through this research. We will create animations of the 3D evolution of relativistic jets with kinetic processes, and present them not only at scientific meetings but also use them for AAMU student outreach, public outreach activities, high school outreach, and presentations at the Space and Rocket Center, and Planetarium in Huntsville.