Blue Waters
2017 Annual Report
Science Highlights

Space Science

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: GLCPC/275 Knh
PI: Dinshaw S. Balsara
University of Notre Dame
Space Science

SIMULATING TWO-FLUID MHD TURBULENCE IN STAR-FORMING MOLECULAR CLOUDS ON THE BLUE WATERS SYSTEM

Research Challenge

Larger simulations of Two-Fluid MHD turbulence in star-forming molecular clouds are needed to match the observations from the HAWC+ instrument. The simulations are extremely CPU-intensive, and only simulations with very limited resolution (5123 zones) are currently manageable. At the present resolution, we will be unable to match the observations from HAWC+. With current simulations, NASA's investment in HAWC+ will be in vain because the detailed match between simulations and observations will not be possible. The work on this newly funded grant will rectify this situation.

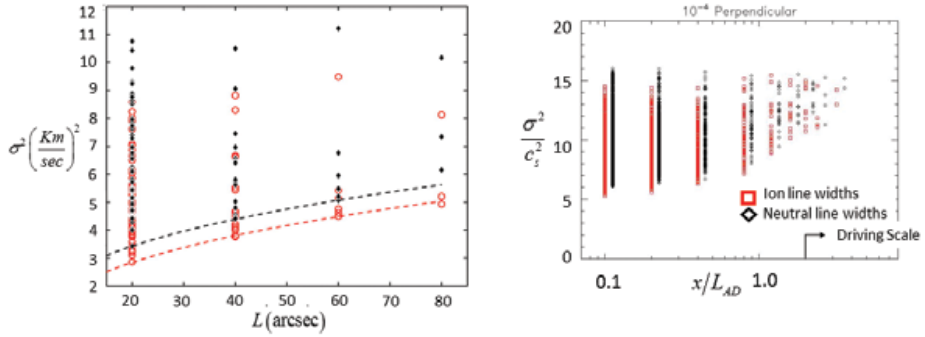

Left: Observed velocity dispersions as a function of length for the HCN molecule (black) and the HCO+ ion (red). The velocity dispersion in the ions is lower than that in the neutrals. Right: Shows simulated linewidth-size relationship from simulations.

Methods & Codes

The core MHD algorithms in our RIEMANN code are based on higher-order Godunov schemes. The research team has been on the forefront of the effort to develop high accuracy schemes for computational astrophysics in general and computational MHD in particular. Two-fluid methods have been described in the references.

Results & Impact

This project is newly funded and is in its initial stages where large-scale simulations have been planned and are ongoing on Blue Waters. Several papers have been published by our group using lower-resolution simulations. The new work will be a substantial improvement on our previous work in terms of resolution as well as in the details of input physics and accuracy of simulation code.

Why Blue Waters

This group has also simulated at petascale on Blue Waters via the Great Lakes Consortium for Petascale Computation. We are, therefore, extremely familiar with the Blue Waters system. The PIs are also funded via NSF grant DMS-1622457 to develop computational capabilities for turbulent simulations in computational astrophysics. This newly funded proposal will provide us with the impetus for developing petascale-ready simulation tools for astrophysical turbulence and making them freely available to the greater astrophysics community.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/1,750 Knh
PI: Tiziana Di Matteo
Carnegie Melon University
Space Science

GIANT BLACK HOLES AND TINY DISK GALAXIES AT THE COSMIC FRONTIER.

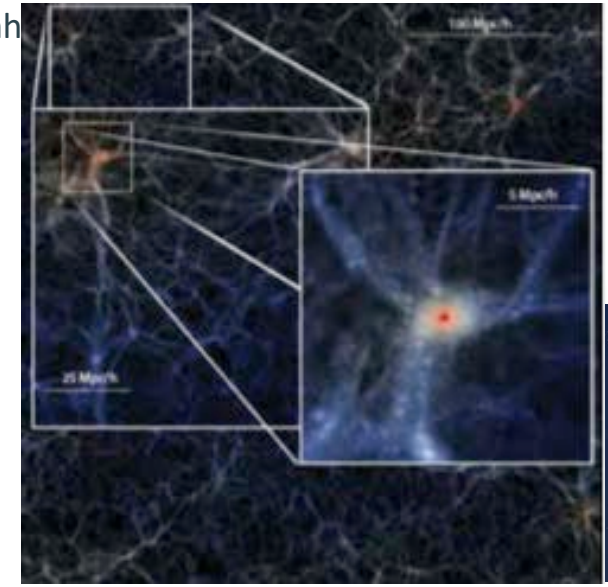

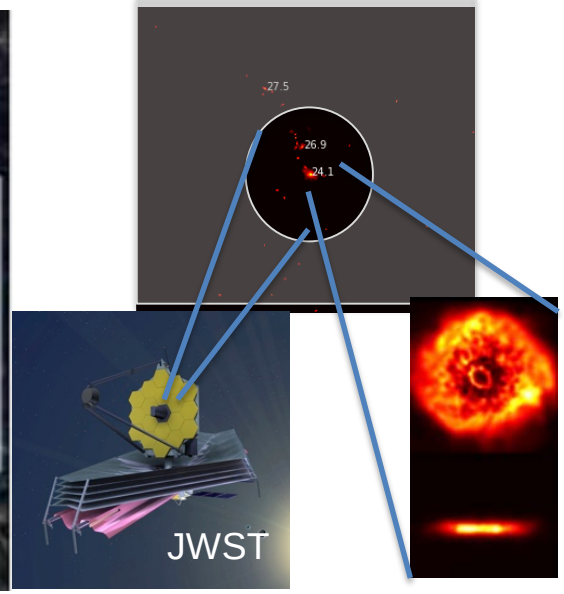

Clockwise from top left: the dark matter density in a slice through the BlueTides MassTracer simulation, a simulated Webb telescope field of view showing that the earliest massive galaxies will be visible, a Milk-Way like Galactic disk seen face and edge on, the Webb telescope, launching in 2020.

Research Challenge

The fundamental challenge is to understand the the first generation of galaxies and the elusive supermassive black holes that form in the first billion years of cosmic history. This is a pivotal time: the galaxies and black holes that form then shape and influence all future generations all the way to today. An extremely large volume of the universe needs to be simulated while at the same time extremely high resolution is required. The Blue Tides cosmological simulation run on Blue Waters is the first and only cosmological simulation of structure formation that has run on the full machine. The simulation is crucial for predicting what the Webb Telescope, successor to Hubble, will discover in this largely unexplored epoch.

Methods & Codes

- The simulations have been made possible with the use of the new cosmological hydrodynamic simulation code (MP)-Gadget which is massively parallel
- Computational methods used : Smoothed particle hydrodynamic method, hybrid tree and particle mesh/ N-body gravity solver

Results & Impact

- Simulations of structure formation have allowed the exploration of the conditions conducive to the growth of the earliest supermassive black holes
- They have predicted a population of tiny Milky Way-like galaxies to be observed by JWST.
- Simulations following the earliest phases of black hole critical growth have enabled new findings for the most massive black holes, on the importance of local tidal field.
- Dark matter-only simulations have addressed the question on where the most massive early-forming quasars are today.

Why Blue Waters

Large-scale uniform volume hydrodynamic simulation of the high-redshift universe is a problem that is perfectly suited to the largest modern petascale facilities such as Blue Waters, which has made simulations of unprecedented volumes and resolutions possible.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: GLCPC/450 Knh
PI: Thomas W. Jones
University of Minnesota
Space Science

TOWARD ROBUST MAGNETOHYDRODYNAMIC SIMULATIONS OF GALAXY CLUSTER FORMATION

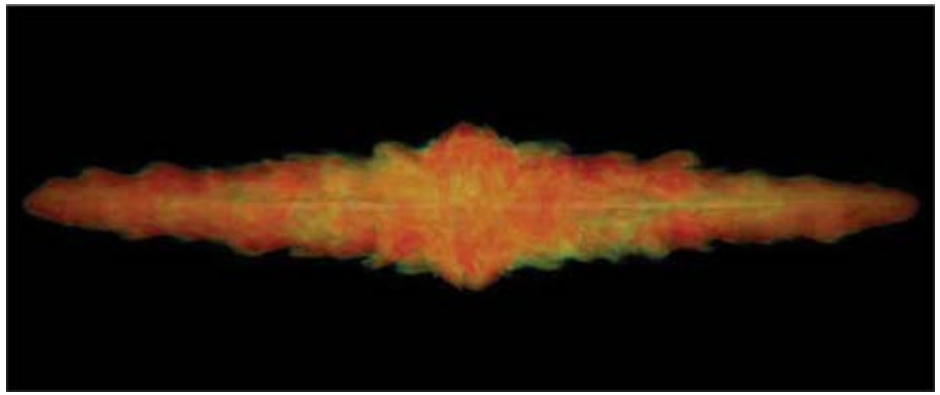

Volume rendering of the distribution of jet plasma resulting from 3D Mach 10, light, bipolar jets propagating through a galaxy cluster medium. The 1,728 x 576 x 576 cell MHD simulation was done with the WOMBAT MHD code on Blue Waters utilizing 2,196 MPI ranks. The volume shown spans approximately 280,000 light years, left to right.

Research Challenge

Galaxy clusters are massive with sizes of several million light years. Their formation is a critical diagnostic of cosmological theory. Understanding the physics of their matter is important for understanding the galaxy cluster's historical and dynamic analysis. Until now, simulations that could capture those small-scale interactions were beyond reach with existing software and high-performance computer systems. The current objective is to develop tools to resolve this problem for the upcoming generation of exascale systems while utilizing existing petasacle systems such as Blue Waters.

Methods & Codes

The team has developed an exceedingly high-performance and highly scalable magneto-hydrodynamics (MHD) cosmology code named "WOMBAT." The code optimizes local memory and vector performance and utilizes hybrid parallelization methods leveraging techniques pioneered at Cray that gain maximum "threading" performance within many-core nodes and MPI-RMA (message passing interface–remote-memory access) performance between nodes.

Results & Impact

The primary objective on Blue Waters for this allocation was to test and tune the technologies built into WOMBAT in order to prepare to address the cosmological problem. They demonstrated 75% scaling efficiency going from one Blue Waters XE6 node to 16,224 XE6 nodes with 16 floating point threads per node. A collaboration with CRAY Developers resulted in a new MPICH release with enhanced performance for hybrid parallel applications that will considerably benefit the broader high-performance community. They did complete very high-resolution WOMBAT MHD simulations (see the Figure).

Why Blue Waters

Blue Waters was essential for this project because using a system capable of sustained petaflop performance with very fast inter-node communications allows for proper WOMBAT performance testing. Blue Waters offers the unique opportunity to both test their strategy and to begin to carry out simulations that approach needed levels of physical fidelity.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/6,000 Knh
PI: Eric J. Lentz
University of Tennessee Knoxville
Space Science

EXPLORING THE NATURE OF EXPLODING MASSIVE STARS WITH HIGH RESOLUTION

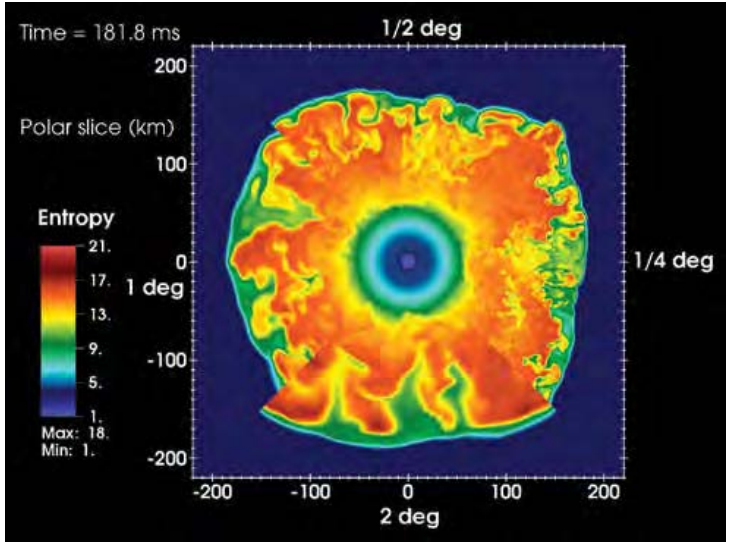

Slice in entropy through the 90° wedge models of various angular resolutions (2° to 1⁄4°) in the pre-supernova convective phase. As resolution increases, the number of fine structures increases, but the character of the models with regions of lower entropy inflow (green) and upwelling heated material (orange) remains.

Research Challenge

- Pre-collapse progenitors come in a wide variety, driven largely by variations in the initial mass, rotation, and composition
- During the explosion heating comes from neutrinos emitted from the neutron star forming at the center of the collapse, which requires energy resolved neutrino transport
- After the explosion begins the evolution of the nuclear isotopes in the ejecta requires a nuclear reaction network

Methods & Codes

- Chimera code
- Hydrodynamics, gravity, active nuclear reactions (burning)
- Dimensionally split piecewise parabolic finite volume scheme for hydrodynamics
- Global multipole expansion of Poisson equation for gravity
- 14 or 160 species in nuclear reaction network

Results & Impacts

- Measured the development of the cascade of turbulent energy to small scales
- Results indicate that the 1° models of our full-geometry 3D Blue Waters models should be sufficiently resolved and allay concerns that future gains in available capability must be thrown primarily into achieving better resolution
- Second set of simulations follows the collapse and explosion of a lower-mass star (9.6 solar masses) as a site for production of calcium-48 using an in-situ nuclear reaction network

Why Blue Waters

- Core-collapse supernova simulations are large, lengthy, and expensive, requiring 1000+ coupled nodes. Even a single 3D simulation can overwhelm the available allocations for a single project at other large sites, but with Blue Waters the team can perform about three simulations per year.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: Illinois/200 Knh
PI: Deborah Levin
University of Illinois at Urbana-Champaign
Space Science

MODELING PLASMA FLOWS WITH KINETIC APPROACHES USING HYBRID CPU-GPU COMPUTING

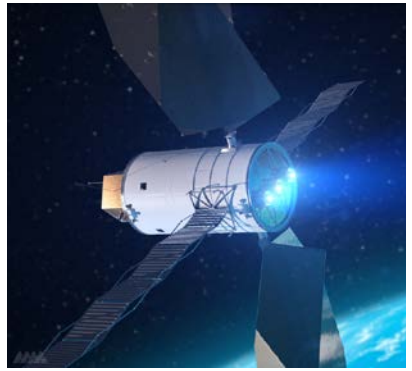

Ion thrusters propelling satellite.
Image Credit: NASA

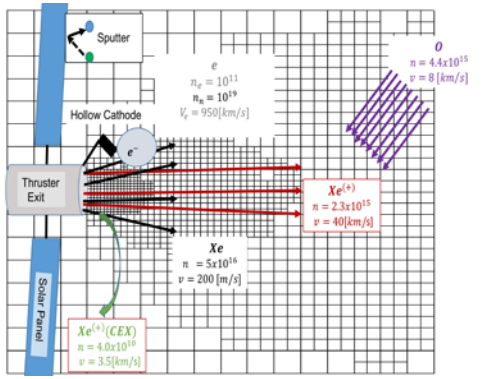

Schematic of thruster simulation experiment on Blue Waters.

Research Challenge

Ion thruster engines, which are used on spacecraft (satellites), create plasma plumes. We need to improve prediction of long-term effects of engine emissions on spacecraft surfaces to improve their efficiency and longevity.

Model the charge accumulation on the solar panel.

Model the energy with which ions bombard solar panels on satellites.

Model external neutralizer positioning which can decrease the volume of blow back ions.

Methods & Codes

Plasma modeling method based on modified DSMC (Direct Simulation Monte Carlo) code, CHAOS (Cuda-based Hybrid Approach to Octree Simulations).

For modeling the electric field, using AMR (Adaptive Mesh Refinement) involved single and multiprocessor stages.

Results & Impact

Satellites that are more efficient and last longer will save money, as billions of dollars are spent annually building and launching these craft.

A simulation shifting the position of the hollow cathode relative to the ion source exposed a weakness in traditional models used for simulation. This necessitates the CHAOS method for improved simulation accuracy, which also requires greater computation.

Why Blue Waters

It allowed testing on GPUs and scale prediction for 3D simulations.

Electrons are typically modeled as a fluid, and not as accurate as particle simulations in some aspects, which consume more computational power. About 1000 GPUs will be required for the electric field simulations, coupled with collisions and interactions with solar panels.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/1,800 Knh
PI: Yi-Hsin Liu
Department of Physics of Astronomy
Dartmouth College
Space Science

THREE-DIMENSIONAL NATURE OF COLLISIONLESS MAGNETIC RECONNECTION AT EARTH'S MAGNETOPAUSE

Research Challenge

The overall goal of this project is to develop an adequate understanding of the 3D nature of asymmetric magnetic reconnection, which will help to advance the modeling of space weather. To this end, the team aims to answer the question: Is there a simple principle that determines the orientation of the reconnection x-line (the line along which magnetic field lines reconnect) in such an asymmetric current sheet? The solution to this problem remains unclear according to the current understanding of magnetic reconnection.

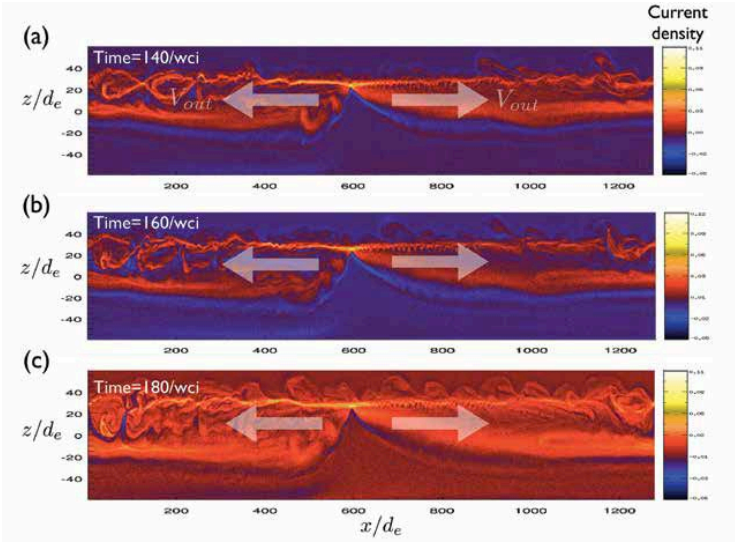

The current density on the reconnection plane in a sequence of time shows the self-generated turbulence during magnetic reconnection. The arrows indicate the reconnection outflows.

Methods & Codes

The particle-in-cell code VPIC is used for magnetic reconnection simulations, and the open-source package ParaView is used for visualization of the results. VPIC solves the relativistic Vlasov–Maxwell system of equations using an explicit charge-conserving approach. Charged particles are advanced, and then the current and charge density are accumulated on grid points to update electromagnetic fields.

Results & Impact

The team has studied the three-dimensional nature of the reconnection x-line. Knowledge gained from these simulations will enable space scientists to more accurately estimate the efficiency of flux transfer from solar wind to the Earth's magnetosphere. The work done here is also relevant to the study of dayside reconnection during the first phase of NASA's Magnetospheric Multiscale Mission.

Why Blue Waters

Because the x-line has a dimension down to electron scale, a fully kinetic description is necessary. Given the available computational capability, it has become possible to use a first-principle kinetic simulation to investigate the dynamics of the x-line in a reasonably large 3D system, which spans from electron kinetic (micro) scale to magnetohydrodynamics (macro) scale. A representative 3D run in this project simulates 2 trillion charged particles on 6 billion grids, generating hundreds of TBs of data for each run. Blue Waters not only provides the computational resource for the calculation but also the online storage for the output and restart files.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/ 3,200 Knh
PI: Warren B. Mori
University of California, Los Angeles
Space Science

TRANSFORMATIVE PETASCALE PARTICLE-IN-CELL SIMULATIONS

Research Challenge

The team focused on three key questions:

1. Can plasma-based acceleration be the basis of new compact accelerators?
2. Can laser plasma instabilities be controlled or harnessed in inertial fusion plasmas?
3. What are the collective processes responsible for the formation of shocks in collisionless plasmas?

Methods & Codes

Particle-In-Cell (PIC) codes employed:

- OSIRIS
- QuickPIC
- UPIC

These codes are locally developed by the UCLA simulation group and freely available.

Results & Impacts

Large-scale PIC simulations were performed and will impact the design of future experiments in plasma-based accelerators and inertial confinement fusion.

Why Blue Waters

Blue Waters provides the largest, time-tested, and stable supercomputing platform in the world. The combination of CPU and GPU nodes suits almost all research supercomputing needs and has provided a productive computational environment since the very beginning of the system.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/1,500 Knh
PI: Scott C. Noble
The University of Tulsa

Space Science

MINI-DISK DYNAMICS ABOUT SUPERMASSIVE BLACK HOLES BINARIES

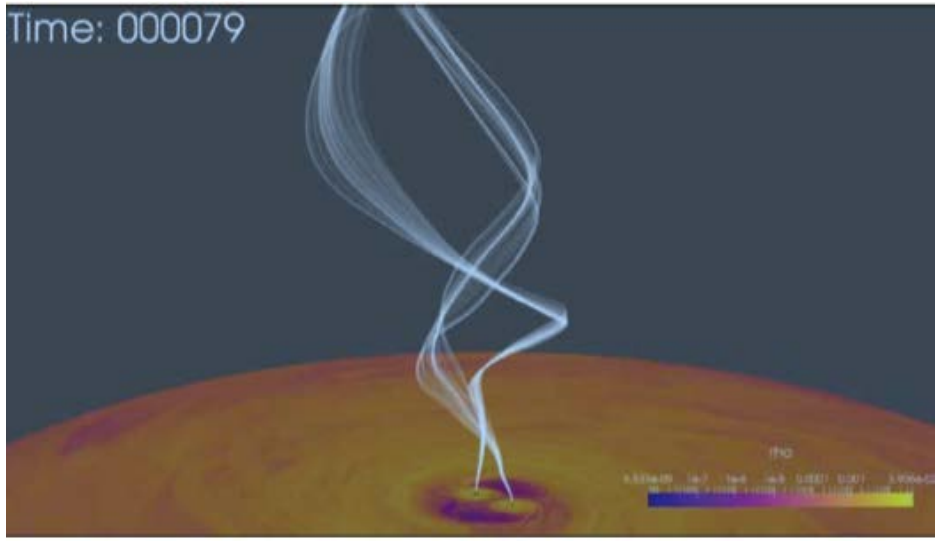

Magnetic field lines (white curves) emanating from two magnetized black holes from the 3D magnetized simulation. An equatorial slice of the density of the accreting gas is also shown (background).

Research Challenge

The research focus is on mini-disk simulations, which are first-of-a-kind comprehensive general relativistic magnetohydrodynamic (GRMHD) simulations of accreting supermassive black holes (SMBHBs). Including the mini-disks is critically important to understand the EM signatures. Mini-disks make up a large fraction of the total luminosity from the system, and give rise to the most variable emission, which is key for astronomers in characterizing and identifying SMBHBs.

Methods & Codes

- A flux-conservative, high- resolution, shock-capturing general relativistic magnetohydrodynamic (GRMHD) code, HARM3d, is used.
- GRMHD is written in a way so that any metric or coordinate system may be adopted, which accommodated the implementation of a novel, time-dependent, non-uniform gridding scheme to resolve the huge scale differences for features near and far away from black holes.

Results & Impacts

- 2D hydrodynamic simulations found significant mass exchange occurred between the two mini-disks. Also, binaries near merger will be bright and periodic at a time scale associated with the SMBHB's orbital period—a key to extracting information about an observed binary's orbit.
- The first full 3D MHD evolution simulation of mini-disks about black hole binaries in the relativistic regime, including the accretion from the circumbinary disk, will begin to address the importance of spiral shocks and the interaction between the mini-disks circumbinary gas.

Why Blue Waters

Blue Waters' capability and support has made the simulation of two orbital periods using 12.9 million floating-point-core-hours possible. The simulation used 600x160x640, or approximately 60 million cells, on about 2 million time steps, and took about one month to complete. The simulation is challenging because of the large dynamic range of time scales between the fast behavior near the black holes and the relatively slow orbital velocity of the binary.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/2,500 Knh
PI: Michael L. Norman
University of California, San Diego
Space Science

REALISTIC SIMULATIONS OF THE INTERGALACTIC MEDIUM: THE SEARCH FOR MISSING PHYSICS

Research Challenge

*More precise observations of the intergalactic medium (IGM) in the past decade have revealed a discrepancy with the well-established predictions of computational model*s. The discrepancy suggests that the standard model lacks some essential ingredient that we refer to simply as *"missing physics."* The project is addressing the issue of whether we are overlooking a key component of the mass-energy content of the universe.

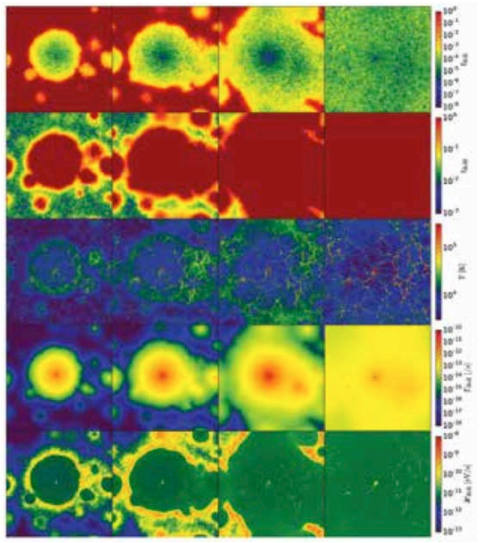

Redshift evolution of the thermal and ionization state of the intergalactic medium on a slice through the 80 Mpc simulation volume. Left to right: redshift snapshots at z=4, 3.5, 3, and 2. Top to bottom: He II fraction, He III fraction, gas temperature, He II photoionization rate, He II photoheating rate.

Methods & Codes

They have carried out a suite of the first fully coupled radiation hydrodynamic cosmological simulations that treat the quasars as a time-varying population of point sources. They have used the *multigroup flux-limited diffusion (MGFLD) branch of ENZO* developed by Dan Reynolds to perform the simulations. Using MGFLD, they accurately transport the hard UV radiation from quasars covering photon energies 54.4 eV to 500 eV and calculate its effects on the IGM self-consistently. *Results are analyzed using the open source yt toolkit.*

Results & Impacts

The MGFLD simulations' IGM temperature peaks around redshift 3, which is in agreement with observations, while the standard model peaks at a redshift of 3.5. The reason for the difference has to do with *the finite time it takes for an ionization front to propagate across the vast distances of intergalactic space*. This discovery helps *resolve one, but not all, of the discrepancies among observations and our earlier models*. The impact of these results is that all future models of the IGM must be revised to include this finite time propagation effect.

Why Blue Waters

The ability to do a suite of exploratory runs of this new type of simulation with excellent throughput, each of which is *quite computationally intensive, is only possible on a Blue Waters-scale system*. In this way they are able to home in on the model that best agrees with the high-precision observations.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/8,100 Knh
PI: Brian W. O'Shea
Michigan State University
Space Science

SIMULATING GALAXY FORMATION ACROSS COSMIC TIME

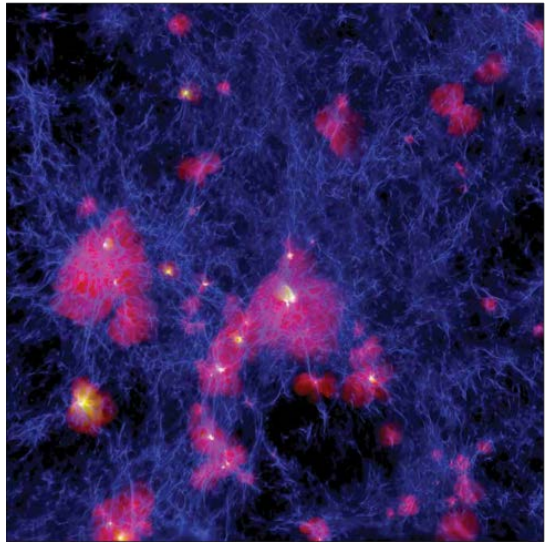

Composite volume rendering of several galaxies at z=8.6, approximately 590 million years after the Big Bang. The field of view is 250 kiloparsecs (approximately 800,000 light years) across. The blue color table shows density, red shows temperature, and green shows ionizing radiation.

Research Challenge

The goals are to understand *two critical issues in galaxy formation*: *the formation of the earliest generations of galaxies and their connections to the Milky Way* through hierarchical structure formation, and *how gas gets into and out of galaxies* and what it does while it is there. All of the calculations needed to study these problems require simulations with an extremely high dynamic range in space and time, complex physics (including radiation transport and non-equilibrium gas chemistry), and large simulation volumes.

Methods & Codes

Their simulation tool of choice is the *Enzo code*, an open-source and community-developed software platform for studying cosmological structure formation. Enzo allows them to include all the critical physical components needed to study galaxy formation—gravity, dark matter dynamics, fluid dynamics, the microphysics of plasmas, and prescriptions for star formation and feedback. *All analysis was done with the yt code*.

Results & Impacts

They find that *increasing the resolution by more than an order of magnitude results in the appearance of both spatial and chemical features that are seen in observations but not in previous models*. Similarly, they find in their idealized simulations that *galaxies can attain a dynamic equilibrium between cold gas condensing and falling into the galaxy, and the ejection of hot, metal-enriched gas into the circumgalactic medium*. This work changes understanding of the interface between the stellar and the diffuse plasma components.

Why Blue Waters

The simulations require extremely high spatial and temporal dynamic range, and also require complex physics; hence, in addition to *large simulation volumes* to model the many early galaxies, *huge numbers of cells* are required to accurately resolve the circumgalactic gas. Taken together, this requires the use of a supercomputer with *large memory and disk space, large computational resources, and an extremely high bandwidth, low- latency communication network* to enable significant scaling of the radiation transport code. *Blue Waters is the only machine available to the academic community that fits all of these requirements.*

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/3,730 Knh
PI: Christian D. Ott
California Institute of Technology
Space Science

3D GENERAL-RELATIVISTIC RADIATION-HYDRODYNAMIC SIMULATIONS OF CORE-COLLAPSE SUPERNOVAE

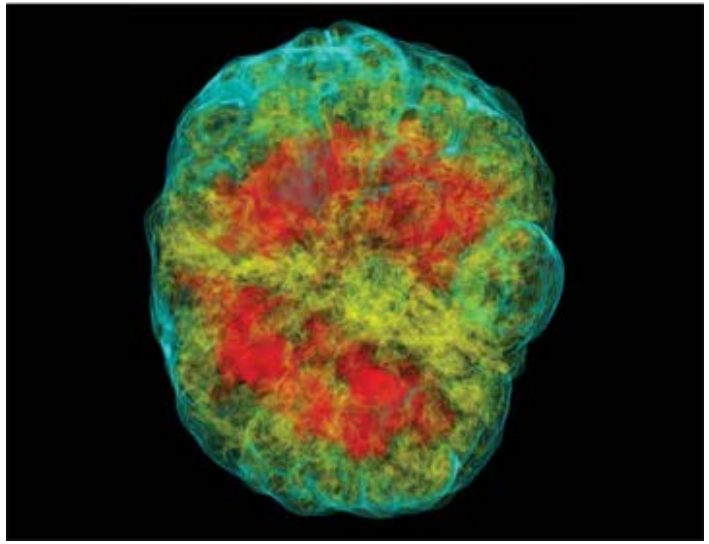

Volume rendering of the specific entropy distribution in a 3D core-collapse supernova at the onset of explosion. The region shown is 600 km^3. Red and yellow colors correspond to higher entropy regions while green, blue, and dark colors correspond to regions of lower entropy.

Research Challenge

- Core-collapse supernovae enrich the interstellar medium with nuclear products of stellar evolution
- Fundamentally a 3D problem involving turbulence, relativistic fluids, general relativity and neutrino radiation
- Truly scale resolving scales as small as 50m and up to 10,000 km

Methods & Codes

- Cactus / Zelmani code
- Finite-differences, finite-volume methods
- Relativistic hydrodynamics, neutrino radiation, general relativity
- Adaptive mesh refinement

Results & Impacts

- First set of full 3D core-collapse supernovae simulations
- Confirmed that “neutrino mechanism” in combination with turbulent convection is able to drive 3D core collapse supernovae
- Full 3D simulations are required since the explosion mechanism depends on the presence of large scale asymmetries in the flow that are suppressed in reduced simulations or incorrect in 2D simulations

Why Blue Waters

- While the team’s simulations typically use “only” 400 nodes, and other HPC systems could accommodate a single simulation, Blue Waters is the only U.S. resource that allows to carry out multiples of such simulations with high throughput. Without Blue Waters, this project would have taken many years to complete.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/2,140 Knh
PI: Nikolai Pogorelov
University of Alabama in Huntsville
Department of Space Science

MODELING PHYSICAL PROCESSES IN THE SOLAR WIND AND LOCAL INTERSTELLAR MEDIUM WITH A MULTISCALE FLUID–KINETIC SIMULATION SUITE

Research Challenge

Investigation of physical phenomena that start on the solar surface and result in the solar wind (SW) acceleration and propagation through interplanetary space toward the boundary of the heliosphere, where the SW interacts with the local interstellar medium (LISM). The simulations are data-driven and help interpret observations from such space missions as IBEX, New Horizons, Ulysses, Voyager, and a fleet of near-Earth spacecraft. Vector magnetogram data and STEREO observations are used to study the propagation of coronal mass ejections toward Earth, where they affect space weather.

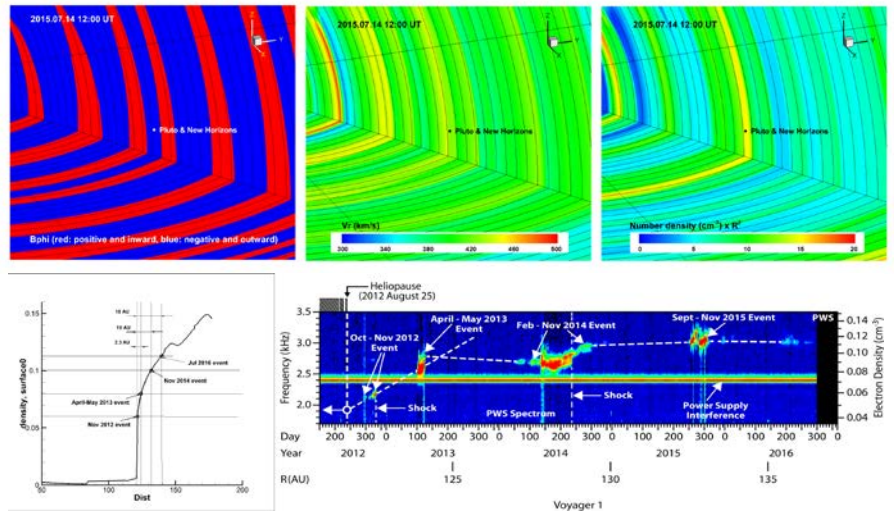

Top: Simulated interplanetary magnetic field direction. Bottom: The distribution of simulated interstellar medium plasma density (left) along the Voyager 1 trajectory and its comparison with the plasma wave events detected by the spacecraft beyond the heliopause (right).

Methods & Codes

Solving the ideal magnetohydrodynamics (MHD) equations coupled with the kinetic Boltzmann equation describing the transport of neutral atoms. In a less strict approach, the flow of atoms is modeled with a few systems of the Euler gas dynamic equations. All these are components of a Multi-Scale Fluid-Kinetic Simulation Suite (MS-FLUKSS)— an adaptive mesh refinement code built on the Chombo framework from LBNL.

Results & Impact

By addressing the basic physical phenomena occurring at the SW-LISM interface, this research is of importance for solar and heliospheric physics, physics of the interstellar medium, and plasma physics in general. The collaboration with the Blue Waters team promotes the application of adaptive technologies to contemporary plasma physics problems through the development of publicly available packages suitable for multiple applications. This is of particular importance because of the Parker Solar Probe to be launched this summer and become the major Heliophysics mission for the decades to come.

Why Blue Waters

First, neutral atoms are modeled kinetically, and the team needs of the order of 10^{12} particles in the Monte Carlo simulations. These simulations require particle splitting, multiple grids, and coupling with the MHD module. They also produce multiple data sets sometimes exceeding 1 terabyte, which require hybrid parallelization. Computational region sizes are very large, exceeding 7.5 x 10^{11} cubed astronomical units, as in the case of long-heliotail simulations necessary to explain the observed anisotropy in multi-TeV cosmic ray flux. Finally, very deep adaptive mesh refinement is necessary near magnetic reconnection sites.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/3,000 Knh
PI: Thomas Quinn
University of Washington
Space Science

UNIFIED MODELING OF GALAXY POPULATIONS IN CLUSTERS

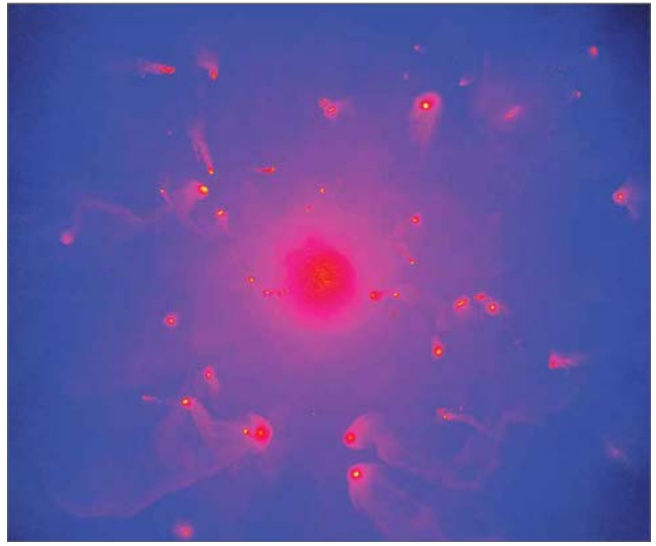

The gas density is shown for the central 3 megaparsecs in a simulation of a Virgo-size cluster (10^{14} solar masses). Note the lack of dense gas in the very center and the stripping of the dense gas out of the smaller galaxies.

Research Challenge

- Understanding the physical processes that occur in groups and clusters of galaxies is key to gaining insights into the evolution of baryons and galaxies across the age of the universe.
- Using clusters as cosmological probes requires understanding relationships between observables and the total mass of the cluster, which in turn requires detailed modeling of the gravitational/hydrodynamic processes using large simulations.

Methods & Codes

- The project uses the highly scalable N-body/hydrodynamics code, ChaNGa, to model formation and evolution of a population of galaxies in a Coma-sized galaxy cluster.
- ChaNGa code includes a well-constrained model for star formation and feedback, and improved implementation of supermassive black hole formation, growth, mergers, and feedback.

Results & Impact

- Prototype simulation of a cluster one-tenth the mass of the target cluster, but at the same mass resolution, completed.
- Preliminary analysis of the cluster indicates that the model continues to produce the observed stellar mass to halo mass even at 10^{14} solar masses.
- The central galaxy displays the typical morphology of the observed brightest cluster galaxies.

Why Blue Waters

- Mass resolutions on the order of 10^5 solar masses required to accurately follow star formation and galaxy morphology.
- Ten billion particles (on the order of 10^{15} solar masses) required to model galaxy clusters of interest.
- The number of particles and long-range nature of gravity requires a high-performance, low-latency network to perform the calculation.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/3,850 Knh
PI: Vadim Roytershteyn
Space Science Institute

Space Science

KINETIC SIMULATIONS OF LARGE-SCALE PLASMA TURBULENCE

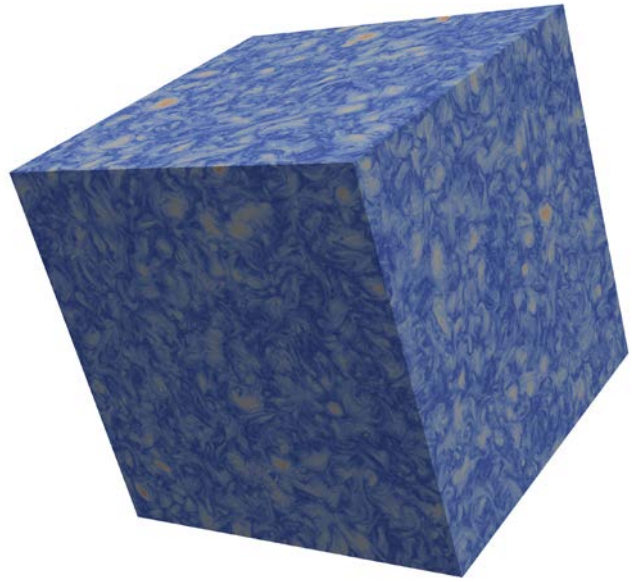

This figure illustrates the magnetic field in a large-scale hybrid-kinetic simulation of decaying turbulence.

Research Challenge

Plasma turbulence plays a significant role in the dynamics of many systems in the universe, from laboratory fusion experiments, to the Sun, and to astrophysical objects such as accretion disks. This project seeks a greater understanding of plasma turbulence by conducting simulations using codes that are capable of faithfully describing microscopic physical effects. This is important since plasma turbulence is a truly multiscale phenomenon, where the very nature of physical processes governing dynamics changes with scales.

Methods & Codes

The most complete description of the plasmas of interest is provided by the Vlasov–Maxwell equations, which we solve using the particle-in-cell technique. Plasma is represented as a collection of particles while equations describing electromagnetic fields are solved on a grid. Some of the simulations performed in this project were guided by and directly compared against observations conducted by the WIND spacecraft in the solar wind.

Results & Impacts

Current sheets in solar wind were studied using large-scale simulations. Simulation results were in good agreement with spacecraft observations, demonstrating the value of the employed modeling method and providing input to many future studies.

Hybrid-kinetic simulations of plasma turbulence decay were also performed. Preliminary comparison of the results with previous research, which was done using macroscopic approximations, shows significant differences that may indicate significance of microscopic effects.

Why Blue Waters

The simulations needed to address the scientific questions of this project solve a global problem that cannot be split into a series of smaller simulations. Typical large-scale simulations can simultaneously track upwards of a trillion particles in order to obtain reliable statistics. As such, they require large memory, fast on-node computation, and fast internode communications. For this reason, they require an HPC resource such as Blue Waters and cannot be conducted on cloud resources.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: Illinois/990 Knh
PI: Stuart L. Shapiro
University of Illinois at Urbana-Champaign
Space Science

MAGNETOROTATIONAL COLLAPSE OF SUPERMASSIVE STARS: BLACK HOLE FORMATION, GRAVITATIONAL WAVES, AND JETS

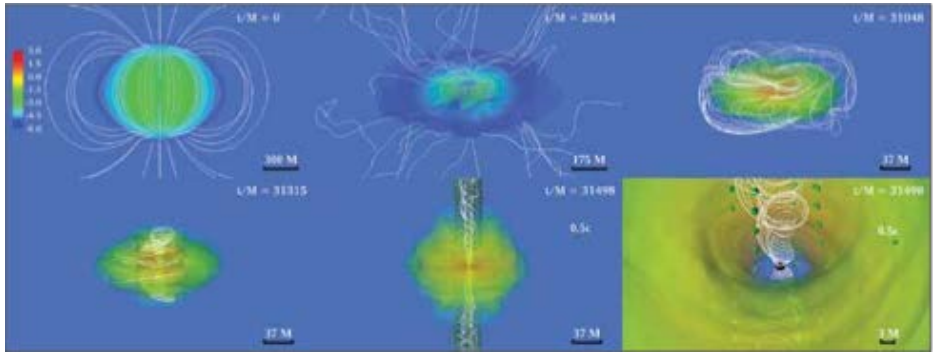

Rendering of the evolution of a magnetized, spinning SMS through collapse and subsequent black hole formation . The outcome is a spinning black hole immersed in a magnetized accretion disk which forms a collimated magnetic field and launches an incipient, relativistic jet emerging from the poles of the black hole.

Research Challenge

Supermassive black holes (SMBHs) reside at the center of most galaxies, including our own. The formation and growth mechanisms for these SMBHs is not clear since, in the early universe, the rate of growth by gas accretion is limited and there is little time (less than a billion years) to grow the oldest and most massive ones if the first black holes were born no more than ten to a few hundred times the mass of the sun. It is quite plausible, therefore, that SMBHs formed from the direct collapse of supemassive stars (SMSs). This project solves the equations of general relativity to follow the collapse of a rotating, magnetized SMS to a SMBH.

Methods & Codes

The team used magnetohydrodynamic simulations (MHD) in full general relativity (GR) to determine the gravitational fields (including the gravitational waves), the matter flow and the electromagnetic fields. The equations were solved using their state-of-the-art Illinois GRMHD code , built on the Cactus infrastructure and using the Carpet code, but employing the team's own algorithms for integratiing Einstein's equations and the equationsf relativistic MHD. The central algorithm for the field equations was the BSSN scheme, which the group helped to develop. Shapiro's undergraduate research team helped to create visualizations and movies with the VisIT software on Blue Waters.

Results & Impacts

In typical cases, following black hole formation, the team observed the formation of magnetically dominated regions above the black hole's poles where the magnetic field lines wound into a funnel, within which the plasma flows outward. This outflow constitutes an incipient jet. Further analysis of the process powering the jets allowed the team to estimate that, for observation times of ~10^4 s, FERMI and SWIFT could detect ultra-long gamma ray burst phenomena from these stars. The gravitational wave busts also could be detected by future space-based gravitational wave instruments. Simulations also crudely modeled the collapse of first generation stars to massive black holes, which could power some of the long gamma-ray bursts observed by FERMI and SWIFT satellites. The analysis thus represents a significant advance in ``multimessenger astronomy".

Why Blue Waters

By adding OpenMP support to the researchers' code and taking advantage of Blue Waters interconnect and processors, the code exhibits greater scalability and performance than on any other supercomputer used by the team. Another recent effort, building the code with the Intel compilers on Blue Waters, resulted in a 30% performance boost, making Blue Waters unique for tackling the astrophysical problems the research team wanted to address. The ability to store and retrieve the data on Blue Waters efficiently and access handy visualization tools to probe simulations in progress also made Blue Waters very desireable.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/5,220 Knh
PI: Alexander Tchekhovskoy
Northwestern University
Space Science

GPU ACCELERATED SIMULATIONS: BLACK HOLES, SPAGHETTIFIED STARS, AND TILTED DISKS

Research Challenge

- Numerical simulations of tidal disruption events of Sun-like stars by supermassive black holes are extremely challenging because it is difficult to numerically resolve the debris stream because it is very thin relative to the black hole.
- Tilted disk simulations require high resolution to properly resolve non-axisymmetric turbulence in the tilted disks.

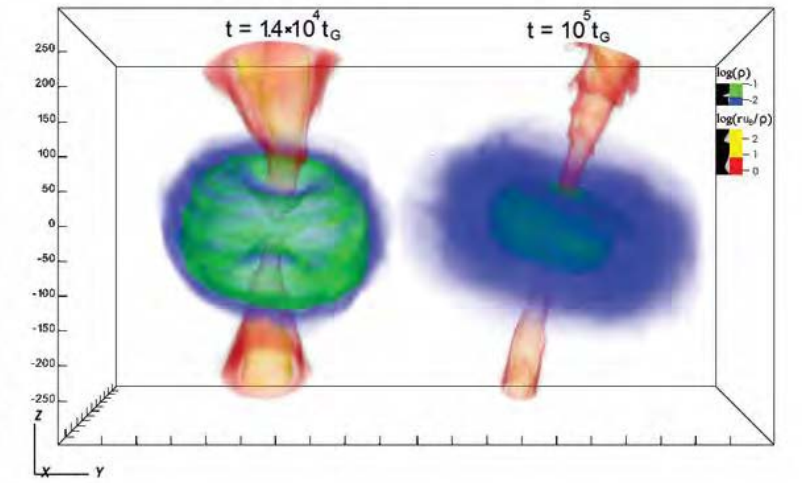

The PRAC allocation allowed the team to demonstrate for the first time that tilted precessing disks around Black Holes produce relativistic jets and that these jets precess together with the disk.

Methods & Codes

- H-AMR code
- general relativistic magnetohydrodynamics
- adaptive mesh refinement
- multi-GPU code

Results & Impacts

- Simulations establish for the first time that tilted accretion disks are capable of producing jets and that the jets undergo precession together with the accretion disk
- First demonstration that jets can be used as probes of disk precession
- Demonstration for the first time that the thin debris stream in typical tidal disruption events can efficiently circularize to form an accretion disk

Why Blue Waters

- These simulations require a high degree of parallelism as they run on hundreds to thousands of GPUs in parallel.
- On a CPU (central processing unit) cluster, the effective simulation cost would have been 500-million CPU core-hours, straining essentially any CPU-based cluster and requiring the largest GPU-based clusters such as Blue Waters

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/3,000 Knh
PI: Gabor Toth, **Co-PI:** Ward Manchester
University of Michigan
Space Science

ADVANCED SPACE WEATHER MODELING

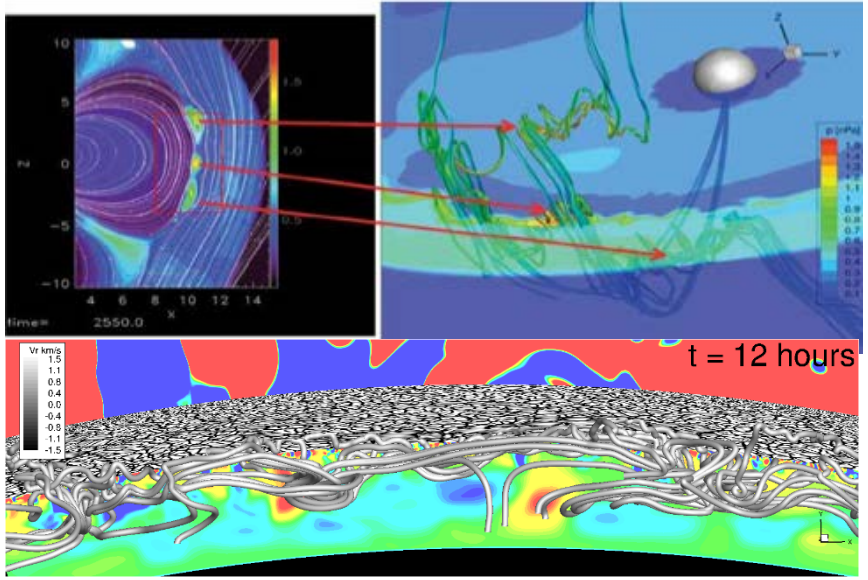

Top left: 2D cut of a simulation of Earth's magnetosphere showing pressure in color. Top right: 3D structure of a magnetic flux rope formed by reconnection. Bottom: Magnetic field emergence through the solar surface with radial velocity in colors. Lines show magnetic field structure.

Research Challenge

- Extreme space weather events are caused by the most energetic coronal mass ejections. These events can generate, among other effects, large-scale electric impulses that can melt transformers and cause cascading blackouts. Repair times for replacing the high-voltage transformers are estimated to be several months.
- Accurate modeling of magnetic storms requires prediction of the interplanetary magnetic field of coronal mass ejections and an accurate model for the reconnection process.

Methods & Codes

- Space Weather Modeling Framework (SWMF)
- Block-adaptive grids
- BATS-R-US used as the Spherical Wedge Active Region Model (SWARM)
- Magnetohydrodynamics with embedded particle-in-cell (MHD-EPIC) using iPIC3D embedded into BATS-R-US and coupled through the SWMF

Results & Impacts

- First three-dimensional global study of the solar wind magnetosphere interaction using a high-fidelity kinetic model for the magnetic reconnection
- Breakthrough advances in simulating flux emergence at active-region scale in spherical geometry
- Addresses the most salient questions of space weather

Why Blue Waters

- Blue Waters' capability allows simulating magnetic flux emergence from the convection zone into the corona to create active regions and coronal mass ejections
- Using Blue Waters allows modeling the global magnetosphere with embedded particle-in-cell model covering the magnetic reconnection sites

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

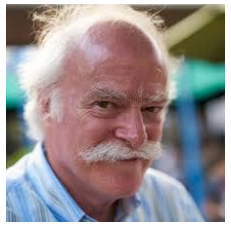

Allocation: NSF PRAC/3,000 Knh
PI: Paul R. Woodward
University of Minnesota

Space Science

3D SIMULATIONS OF I-PROCESS NUCLEOSYNTHESIS

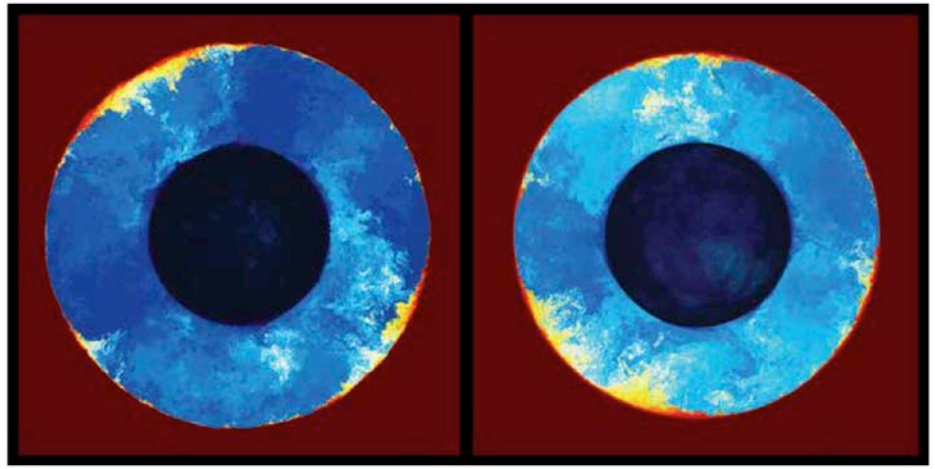

Snapshots of the distribution of ingested gas from above the convection zone generated by oxygen burning in model 25-solar-mass star. Left: 2.25x the luminosity of the 1D model. Right: 56.25x this luminosity to validate 1D descriptions of the ingestion phenomenon and to determine their dependence on parameters of the stellar context.

Research Challenge

The challenge to *the potential merger of multiple nuclear burning shells* in massive stars has involved *detailed studies of the ingestion process*, particularly in massive star contexts, as well as the aggressive development of a new simulation code. *Their results on i-process nucleosynthesis are important as inputs for the study of the chemical evolution of galaxies.* The new work with massive stars could have a large impact on the conditions just before those stars explode, and also on the injection of heavier elements from these explosions into the surrounding interstellar medium.

Methods & Codes

Work to date simulating hydrogen ingestion flashes exploits the *piecewise parabolic method (PPM)* coupled with the piecewise parabolic Boltzmann (PPB) moment-conserving advection scheme for the multifluid volume fraction. Their new code adds *a Level 3 AMR grid* for containing multiple nuclear burning shells and their respective convection zones in a single simulation. It is designed to *scale to 14,000 nodes* while running *roughly twice as fast as their older code per node* by exploiting 32-bit precision and GPU acceleration.

Results & Impacts

They are *producing a database of detailed simulations* that investigates the phenomenon of convective boundary mixing at unprecedented accuracy for convection zones that extend over ranges in radius of more than a factor of two. Global convection modes play an important role in these situations, making simulation difficult and costly. *Convective boundary mixing* plays an important role in stellar evolution. In particular, *in ingestion events* that they study, it *can have a dramatic impact on nucleosynthesis, which in turn affects galactic chemical evolution*

Why Blue Waters

They have carried out their simulations on Blue Waters because of its special ability to enable their simulation code to be run at *a sufficiently large scale* that their large computations can each be *completed in less than one week*. This allows their team to pose questions and get answers on a timescale that is conducive to productive thought and dynamic adjustment of their research direction.

Geosciences

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/240 Knh
PI: Marin Clark
University of Michigan
Geoscience

HIGH-RESOLUTION DIGITAL SURFACE MODELS OF THE 2016 Mw7.8 KAIKOURA EARTHQUAKE, NEW ZEALAND

Research Challenge

Regional assessment of landsliding and infrastructure damage in the aftermath of large earthquakes is a broad societal problem. Digital surface models derived from satellite data provide valuable information for quantitative assessment of surface changes after an event.

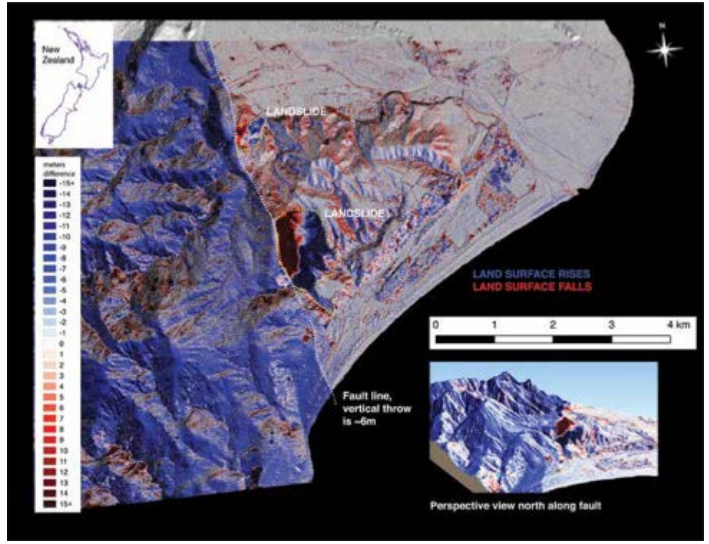

Elevation change from differenced digital surface models of the area near Waipapa Bay, northeastern South Island, New Zealand. Red regions—elevation drop in response to the earthquake. Blue regions—elevation gain. Landslides (yellow dots) are prominent. Satellite measurements of landslide volumes are identical to those made using UAV (unmanned aerial vehicle).

Methods & Codes

Surface Extraction with TIN-based Search-space Minimization (SETSM).

Results & Impacts

With rapid and reactive HPC access and an initial pre-event data set for areas of the globe that are likely, or have already been exposed to natural disasters, they have the ability to provide damage maps and cascading hazard monitoring in a space of hours to days.

This can potentially transform the way first responders are informed about the disaster zone.

Why Blue Waters

Blue Waters was necessary to complete the project in a timeframe useful to coordinate perishable field data collection and provide rapid feedback on event history.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: BW Professor/180 Knh
PI: Larry Di Girolamo
University of Illinois at Urbana–Champaign
Geoscience

THE TERRA DATA FUSION PROJECT

Research Challenge

The Terra satellite, launched in 1999, has collected more than 1 petabyte of data from its five instruments in a wide variety of formats and grid systems that are stored at different NASA centers. This project was launched to address two long-standing problems.

(1) How do we efficiently generate and deliver Terra data fusion products, that is, products derived by combining two or more Terra instruments?

(2) How do we facilitate the use of Terra data fusion products by the community in generating new products and knowledge and disseminate these through national data sharing services?

Methods & Codes

- Terra records from NASA were transferred to Blue Waters
- The team built software to perform whole-mission processing on Blue Waters to create basic fusion products
- Preparation began for mission-scale processing with help from the Blue Waters team

Results & Impacts

Researchers continue to the dataset in scientific studies. Current studies include how the EM radiation leaving Earth has changed over the Terra record, characterizing the global distribution of ice crystal roughness parameters that are important to our meteorological understanding of ice clouds, and to better characterize the size of liquid water droplets that make up clouds over the globe. An outreach video on Terra fusion was also created and can be found at: https://www.youtube.com/watch?v=C2uyjRGwwOs

Why Blue Waters

Key advantages of Blue Waters are local process of the data, with sharing and access that is global. Blue Waters provided a massively parallel system, one of the largest storage and bandwidth computing facilities and excellent sharing services.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: Illinois/200 Knh
PI: Marcelo H. Garcia
University of Illinois at Urbana-Champaign
Geoscience

LARGE-EDDY SIMULATION OF SEDIMENT TRANSPORT AND HYDRODYNAMICS AT RIVER BIFURCATIONS: USING A HIGHLY SCALABLE SPECTRAL ELEMENT-BASED CFD SOLVER

Research Challenge

- Fully understanding the Bulle Effect and related phenomena, such as secondary flows and vorticity-driven sediment transport.
- Resolve all the relevant turbulence eddies of the flow; this will not only provide an accurate description of the dynamics of the flow, additionally it will also help to model the sediment transport phenomena accurately when coupled with a Lagrangian particle model for the sediment.

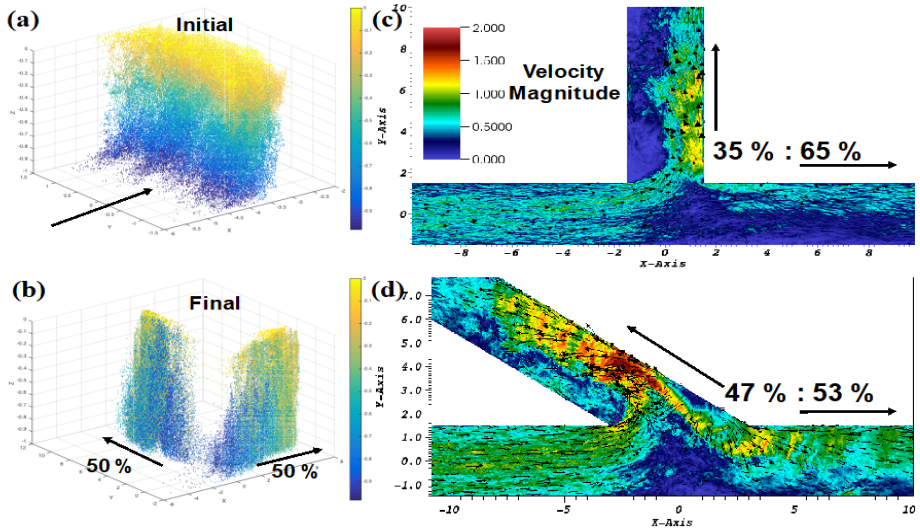

(a) Initial position of fine sediment uniformly distributed in the vertical. Sediment near the bottom is Bluer, and sediment near the top is Yellower. (b) Position of the particles once they have moved into the two channels. (c,d) Velocity Magnitude at 5 % height from the bottom, for flows in which, 35 % (90-degree) and 47 % (150-degree) of the flow is moving into the side/lateral channel.

Methods & Codes

- The computational fluid dynamic simulations use the open-source, highly scalable spectral element Navier-Stokes solver Nek5000.
- The maximum Reynolds number of the simulations are of the order of 10^4-10^5, which compares favorably with laboratory experiments.
- Simulations conducted with up to 240 million computational points, and 200,000 Lagrangian particles

Results & Impact

- This study provided new insights into the hydrodynamics and sediment transport at bifurcations, and shows that high-resolution LES can be used to study complex river-mechanics problems.
- Publications:

i. Dutta and Garcia, J. Hydraulic Eng., ASCE, 2018 a,b
ii. Dutta, Wang, Tassi, and Garcia, Earth Surface Processes Landforms, 2017
iii. Dutta, Fischer and Garcia, River Flow, 2016

Why Blue Waters

- This study pushes the limit of the scale at which eddy-resolving numerical simulations have been used to study complex multi-phase river mechanics problems, warranting the use of a computational resource that can provide sustained computing power at an unprecedented scale.
- Without access to Blue Waters completing the study within a realistic timeframe would be impossible.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

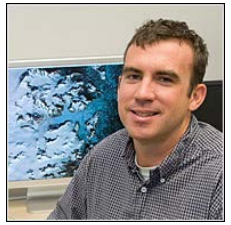

Allocation: Innovation & Exploration /1200 Knh
PI: Ian Howat
The Ohio State University
Geoscience

THE REFERENCE ELEVATION MODEL OF ANTARCTICA

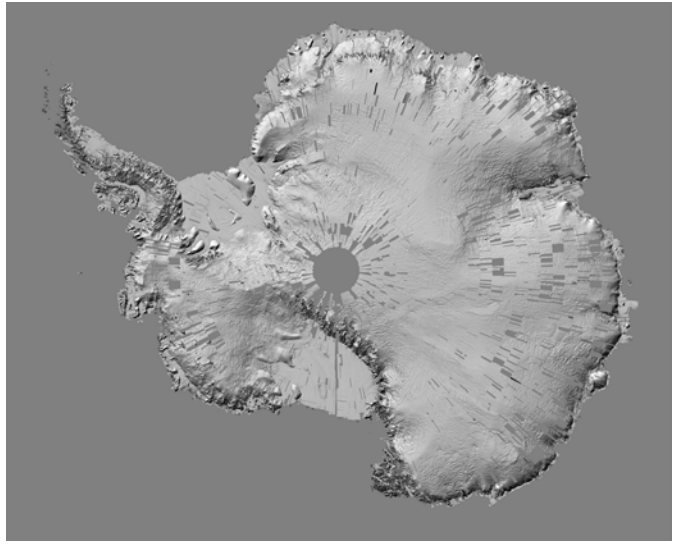

Hill shade representation of a preliminary Reference Elevation Model of Antarctica (REMA) consisting of over 500,000 individual 8-m resolution elevation models processed from satellite imagery on Blue Waters.

Research Challenge

The overall goal is to accurately map a high-resolution (8 m), high-precision (accuracy better than 1 m) reference surface for the continent of Antarctica

Accurate surface elevation is an essential dataset for glaciology, compromising vast sets of images and laser generated elevation data.

There is a large sub-meter set of stereo imagery held by the Polar Geospatial Center that is being utilized.

Methods & Codes

Fully automated algorithm (SETSM) which utilizes node parallelized OpenMP

SETSM is written in stand-alone C code with no external dependencies and requires no libraries, ensuring simple, multi-platform installation, support, and optimization.

SETSM is called from a single command line with the only required inputs being the filenames of the two stereo images and the RPC (Rich Photorealistic Content) file, typically provided in XML (eXtensible Markup Language) format

Results & Impact

Displayed Antarctica's ongoing rapid changes regarding overall ice coverage.

Estimated overall sea level change in the future based on findings

Useful for a wide range of applications beyond glaciology, ranging from geodynamics to multi-disciplinary logistics planning.

Why Blue Waters

Rapid throughput of thousands of individual jobs of highly variable and unpredictable wall times.

Impressive capacity for throughput of job volume

Backfill scheduling capabilities

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: BW Professor/250 Knh
PI: Sonia Lasher-Trapp
University of Illinois at Urbana-Champaign
Geoscience

UNTANGLING ENTRAINMENT AND PRECIPITATION IN CONVECTIVE CLOUDS

Research Challenge

A problem with meteorological models is that they often predict rain formation too early and in excessive amounts. This research focuses upon quantifying how much dry air is entrained by cumulus clouds and storms, that may limit their strength and the precipitation they produce. It requires high spatial resolution to represent small scale turbulence in the clouds and storms, but the resolution required for accuracy is itself an open question that the research team is investigating.

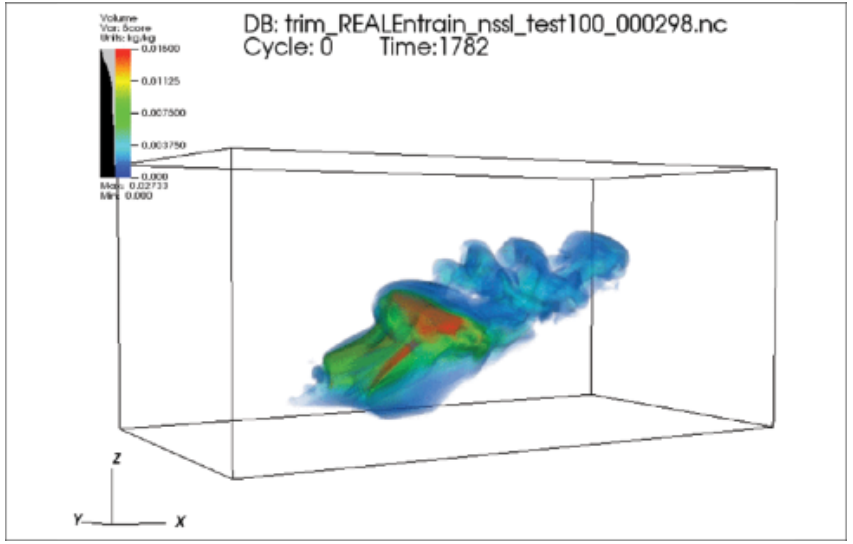

Visualization of the 3D core of a developing thunderstorm. Warm colors (yellow, red) denote areas of greater amounts of precipitation mass; cool colors (green, blue) denote areas with little water mass remaining.

Methods & Codes

The research team is using the NCAR CM1 model to simulate convective cumulus clouds and storms at high resolution using the Blue Waters system. They then apply their own codes to quantify the amount of dry air brought into the clouds/storms, and assess the effects upon the amount of precipitation produced. Comparison to observations and theoretical predictions provide confidence in the modeling results.

Results & Impacts

Atmospheric scientists are working to produce approximations of entrainment, as weather and climate models cannot represent all of the cloud motions. This research provides new quantitative knowledge for that effort.

The latest results show that closer spacing between storms, expected to decrease the effects of entrainment and encourage precipitation, can instead initially *delay and decrease* precipitation, as they compete for the warm moist air beneath their bases that fuel them. Identifying competing effects is critical for making progress upon this research problem.

Why Blue Waters

Blue Waters, with its large number of nodes, its high speed, and its large storage capability for high-resolution model output and analysis, allows us to push the spatial scale limit to include more turbulent motions much farther than in the past. Blue Waters staff have helped the research to learn new and practical ways to visualize the output for easier analysis.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

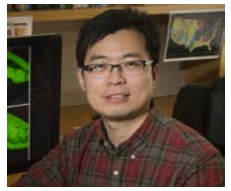

Allocation: NSF PRAC/880 Knh

PI: Lijun Liu
University of Illinois at Urbana-Champaign

Geoscience

UNDERSTANDING THE 4-D EVOLUTION OF THE SOLID EARTH USING GEODYNAMIC MODELS WITH DATA ASSIMILATION

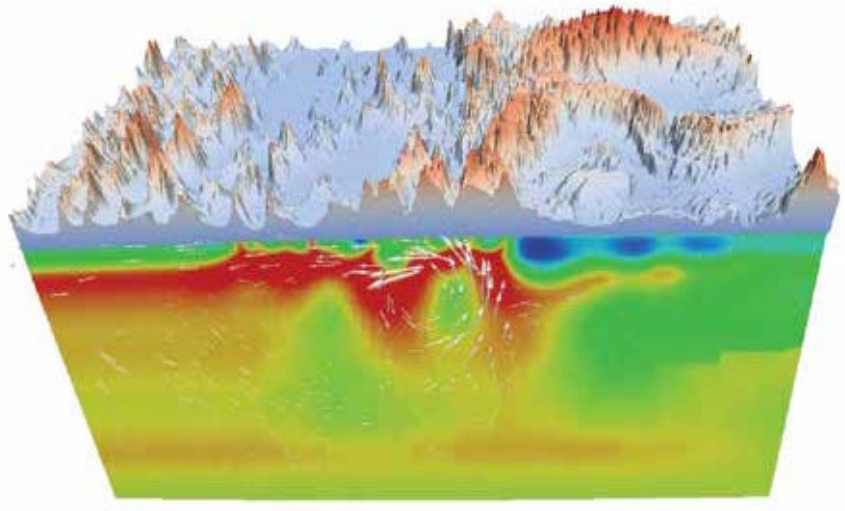

Lithosphere viscosity (color) and mantle flow (arrow) beneath the western United States. The viscosity is converted from electrical resistivity, and the flow is calculated from a geodynamic model.

Research Challenge

The goal of this project is to quantitatively understand current and past dynamic processes within the deep Earth, which are vital for explaining geological, geophysical, and geohazard observations. Such processes include:

- Heat source fueling of volcanoes far away from a subduction zone
- The influence of subduction on the mantle and the surface
- Earthquake- and volcano-generating deformation of continents

Methods & Codes

The team uses CitcomS, a community-based finite element code, to simulate a physics-based numerical model of the deep earth. This code, which can be run either forward or backward in time, uses various geodynamic modeling techniques that combine a variety of things such as mantle temperature and viscosity profiles, as well as tectonic plate motion.

Why Blue Waters

The enormous amount of data processing and computation makes Blue Waters the best platform. The CitcomS code has been designed and tested mostly on traditional supercomputers.

Results & Impacts

- Demonstrated that the electrical conductivity of rocks is an excellent proxy for the effective viscosity of the lithosphere, providing a practical way to “measure” the strength of the Earth’s rigid outer shell. (Liu & Hasterok, *Science*, 2016)
- Showed that the east-west topography difference of the Tibetan Plateau is due to the different crustal strength underneath. (Chen et al., *Nature Comm.*, 2017)
- Challenged the traditional view of intra-plate volcanism formation by demonstrating that Yellowstone-related volcanism was not caused by a deep-rooted mantle plume. (Leonard & Liu, *GRL.*, 2016)
- Found that the seismically fast upper-mantle structures beneath the southern Atlantic are mostly compositional anomalies, likely representing delaminated continental lithosphere. (Hu et al., *EPSL*, 2017)

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/9,696.484 Knh
PI: Paul Morin
University of Minnesota
Geoscience

ENHANCED DIGITAL ELEVATION MODEL FOR THE ARCTIC

Research Challenge

There is a lack of high-resolution, consistent, high-quality elevation data available for the Arctic. The team is able to construct digital elevation models (DEM) from sub-meter resolution stereo imagery from the National Geospatial Intelligence Agency. These data will be used by the Arctic research community to support activities that include sea level rise, ice mass balance, carbon cycling, coastal erosion, permafrost collapse and sustainable development.

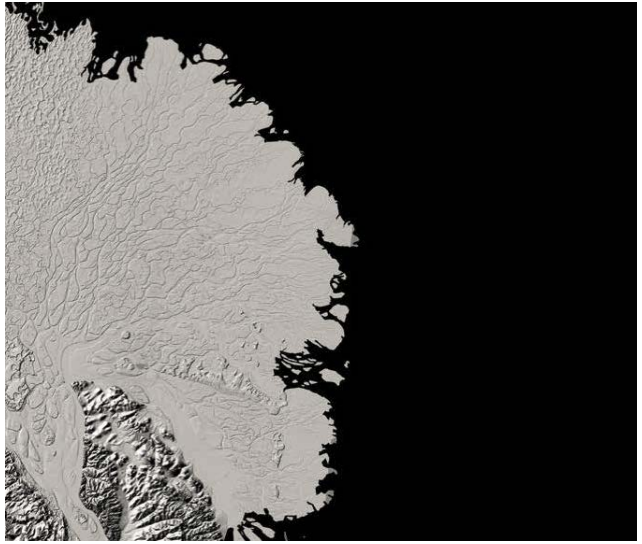

The Lena River is one of the three largest rivers flowing into the Arctic Ocean. This image shows the topography of the 100-km-wide delta as the river changes from being constrained by high relief to low relief.

Methods & Codes

The team developed the Surface Extraction from TIN-base Search-space Minimization (SETSM) code. It is a fully automated program for extracting DEMs from a stereo pair of source images.

Results & Impacts

The team produced 2-m posting DEMs of the Arctic with an average 4x repeat. All of these data have been released to the science community and the public through ArcticDEM.org. These data are being used by scientists, national geographic surveys, and regional and local governments for a broad range of applications.

Why Blue Waters

No other academic computer had the available capacity for this project. Blue Waters was able to execute the ArcticDEM workload without significantly impacting throughput of other projects.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/3,630 Knh
PI: Leigh Orf
University of Wisconsin–Madison
Geoscience

SIMULATING THE MOST DEVASTATING TORNADOES EMBEDDED WITHIN SUPERCELL THUNDERSTORMS

Research Challenge

Accurate forecasting of tornado behavior within supercell thunderstorms, including:

- Tornado Formation
- Tornado Maintenance
- Tornado Decay

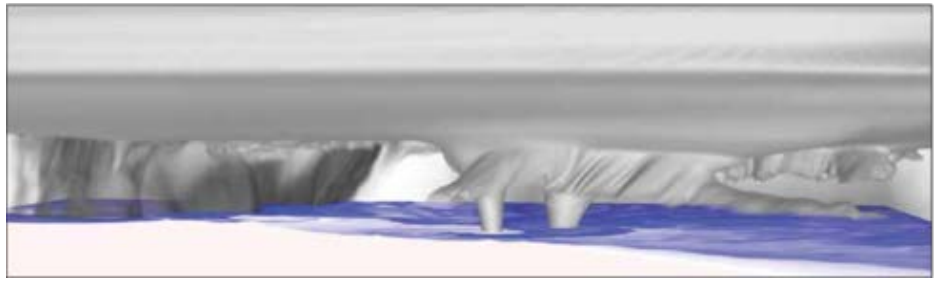

A short-lived anticyclonic tornado (left) adjacent to the long-lived cyclonic EF5 tornado (right), as seen in the volume-rendered cloud field of a 15-meter resolution simulation. The cloud field also reveals a wall cloud and tail cloud, features commonly observed in the field. Rain is visible as a dark grey field and is most prominent in the rear flank of the supercell (left half of image). The cold pool is represented by the surface buoyancy field where the coldest air is dark blue.

Methods & Codes

This research uses a modified version of the community climate model CM1, developed at the National Center for Atmospheric Research

The HDF5 library is used for I/O

Results & Impact

Our ability to forecast tornado behavior is limited by our current knowledge of how tornadoes form, are maintained, and decay in supercell thunderstorms

Tornadoes have extreme destructive capabilities, and a better understanding can help minimize the loss of life

Why Blue Waters

Simulations of the atmosphere, accurate enough to study tornadoes, require:

Very high resolution, large amounts of memory, inter-processor communication, and high I/O bandwidth

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/5,000 Knh
PI: Jamesina J. Simpson
University of Utah
Geoscience, Electrical Engineering

LOCATION-SPECIFIC SPACE WEATHER HAZARDS TO ELECTRIC POWER GRIDS CALCULATED ON A GLOBAL SCALE

Research Challenge

The research focus is to enhance our understanding of near-Earth electrodynamics associated with historically intense coronal mass ejection (CME) events. The models can simulate the effects of a planetary space weather event and provide location-specific information on possible hazards to societal infrastructure over both short (seconds) and long (hours) time spans.

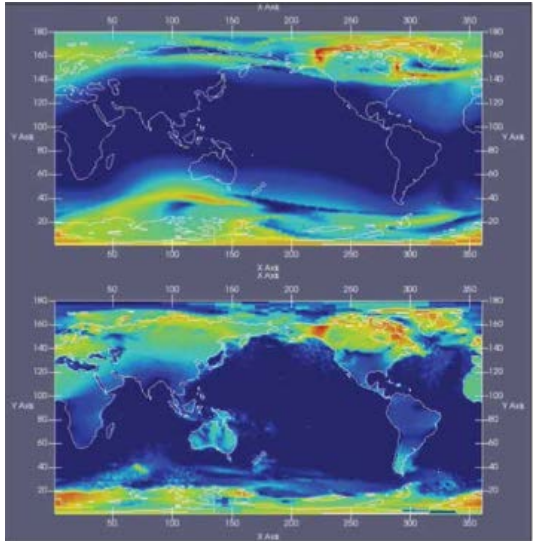

Snapshot of the electric field source amplitude versus position during the October 2003 Halloween geomagnetic storms as calculated by BATS-R-US model developed at the University of Michigan (top image) and the resulting surface-level electric field values calculated by the global FDTD model (bottom image).

Methods & Codes

The team developed their code via the finite difference time domain (FDTD) method to calculate the electromagnetic fields at the surface of the Earth. Their FDTD models are fully three-dimensional and solve for electromagnetic wave propagation through the ionosphere and into the ground and oceans. The Earth's topography is included in the analysis.

Results & Impacts

They were able to determine that space weather does not induce intense electric fields over a sufficiently large area along ocean-continent boundaries to pose a significant risk to power grids. It is expected that ocean-continental regions will only encounter regional hazards that are similar to those already encountered from lightning strikes, etc. However, islands and peninsulas may be at a higher risk in some areas.

Why Blue Waters

Using Blue Waters, they were able to develop higher grid resolutions than previously possible by over an order of magnitude on a global scale (1 cubic km *vs*. a box 40*40*5km) as well as higher resolutions on a localized scale. This permitted the investigation of unique space weather hazards at ocean-continent boundaries. It also opened up a wide variety of new applications, such as working with DARPA to develop a new electromagnetic system for geolocation, and working with the Office of Naval Research on detecting submerged objects in the oceans.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: Illinois/500 Knh
PI: Ryan L. Sriver
University of Illinois at Urbana–Champaign

Geoscience

IMPACT OF OCEAN COUPLING ON SIMULATED TROPICAL CYCLONE ACTIVITY IN THE HIGH-RESOLUTION COMMUNITY EARTH SYSTEM MODEL

Research Challenge

Tropical cyclones (TCs) are among the world's deadliest natural hazards. How TC activity will vary with the changing climate is a topic of great interest. Correct representation of air–sea interactions under TCs is important for simulating realistic storm intensities and track durations. Local feedbacks due to ocean mixing and surface fluxes can inhibit storm development and intensification, as well as influence larger-scale ocean and atmospheric circulations.

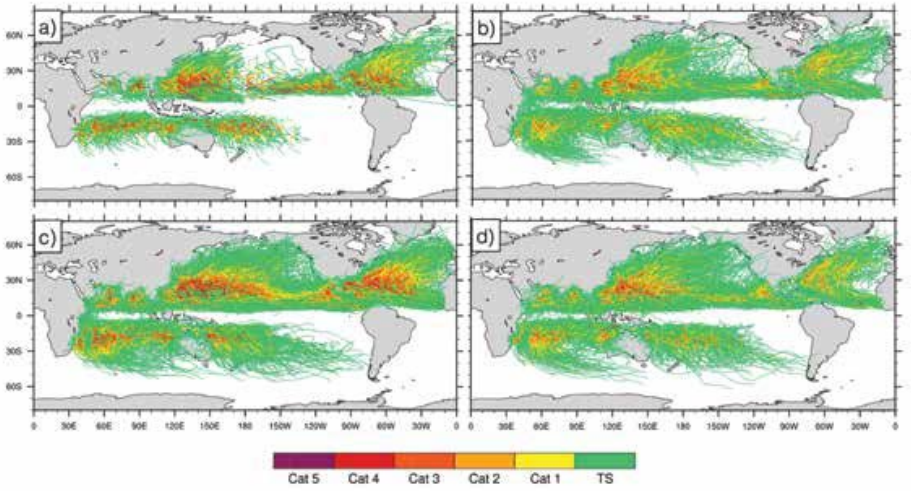

Global TC tracks accumulated over 30 years in the (a) observational best-track (1985–2014), (b) fully-coupled simulation, (c) atmosphere-only simulation, and (d) partially-coupled simulation.

Methods & Codes

In this project, the team assessed the impact of ocean coupling on simulated TC activity using a high-resolution configuration of the CESM (Community Earth System Model) with a 25-km resolution atmosphere. They performed three 30-year simulations in which the atmosphere model is configured with three different levels of ocean coupling. The models are configured to focus on ocean–atmosphere interactions associated with TCs. Each simulation is run under preindustrial climate conditions with an active carbon–nitrogen cycle.

Results & Impacts

The team found that TC number, geographical distributions, and intensity are sensitive to ocean coupling. Differences in TC characteristics are mainly attributed to model differences in local air–sea flux exchanges and large-scale climate conditions. This research enables fundamental advancement of our understanding about important physical processes related to TC dynamics, ocean mixing, ocean heat storage and transport, and global ocean–atmosphere circulations. It paves the way for more comprehensive coupled climate model experiments capable of linking extreme weather events with large-scale climate.

Why Blue Waters

Given the substantial computational expense of high-resolution Earth system models, it is difficult to apply these models to study tropical cyclones because of the necessary grid resolution (1/4 degree), model run length (multiple decades), and high frequency output (multiple times per model day). Blue Waters provides unique capabilities to handle the computational demand associated with running the model at ultra-high resolutions, including scalability to over 15,000 cores, high-frequency input and output, and post-processing and visualization of model results.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: BW Professor/240 Knh
PI: Robert J. Trapp
University of Illinois at Urbana–Champaign
Geoscience

PETASCALE MODELING OF CONVECTIVE STORMS UNDER CLIMATE CHANGE AND VARIABILITY

Research Challenge

- This research seeks to answer the basic question of how present-day extreme storm events might be altered by human-induced climate change
- Part of the challenge is that such storms—and especially the attendant tornadoes, hail, damaging "straight-line" winds, lightning, and localized flooding—have spatial scales that fall below the effective resolution of typical global models

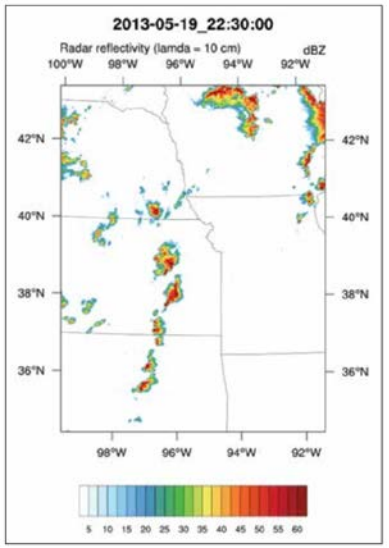

Simulated radar reflectivity from a high-resolution WRF model simulation of the May 19, 2013, outbreak of tornadoes, damaging wind, and hail. The radar reflectivity portrays the structure of the individual hail-producing storms in the CTRL simulation

Methods & Codes

Drawing on the success of previously reported work, they have further adapted the pseudo-global warming (PGW) methodology to investigate the impact of human-induced climate change on outbreaks of severe hail and on landfalling hurricanes. Modified atmospheric states drawn from GCM output were used to constrain WRF model simulations of these events at high resolution

Results & Impact

Exemplifying the hailstorm results are the simulations of the May 19, 2013, outbreak of tornadoes, damaging wind, and hail (see Figure)

Researchers are finding that the conditions under PGW promote a relatively more intense hurricane at landfall and lead to a much higher incidence of mesoscale vortices with tornadic potential

Why Blue Waters

Accurate simulations require very large geospatial domains that have fine grid point spacings and long-time integrations with high rates of model output. Moreover, quantifications of uncertainty require that such realizations be repeated over multiple experiments. The Blue Waters allocation is providing us with the resources needed to achieve this unprecedented level of climate simulation.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: Illinois/550 Knh
PI: Junshik Um
University of Illinois at Urbana-Champaign
Geoscience

IMPACTS OF ORIENTATION AND MORPHOLOGY OF SMALL ATMOSPHERIC ICE CRYSTALS ON *IN-SITU* AIRCRAFT MEASUREMENTS: SCATTERING CALCULATIONS

Light from our sun being scattered by ice crystals.

Research Challenge

They easiest way to identify something is by looking at it. However, there are many things that are very hard to see either because they are too small, or they blend in with their surroundings. Ice crystals in the atmosphere have both of these properties. In order to better understand ice crystals in our atmosphere, this research seeks to understand how light scatters off of non-spherical ice crystals. In other words, this research attempts to predict with better accuracy what our scientific instruments "see" when looking at ice crystals in the atmosphere.

Methods & Codes

This research simulates light scattering by non-spherical ice crystals using what is known as the Discrete Dipole Approximation (DDA). The DDA was implemented and run with a code called Amsterdam DDA.

Results & Impact

With these results, scientific instruments will possibly be able to "see" with greater accuracy the different types of small ice crystals in our atmosphere.

Why Blue Waters

The results of this work suggests that DDA may be a reasonably accurate model for predicting light scattering by atmospheric ice crystals. What's more, DDA has a natural scheme for algorithm parallelization. With the sheer computing volume of Blue Waters, the team is able to investigate a larger variety of scientific parameters in a shorter amount of time.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: Illinois/250 Knh
PI: Albert J. Valocchi
University of Illinois at Urbana-Champaign
Geoscience

PORE-SCALE SIMULATION OF MULTIPHASE FLOW IN POROUS MEDIA WITH APPLICATIONS TO GEOLOGICAL SEQUESTRATION OF CARBON DIOXIDE

Research Challenge

To understand the migration of multiple fluids within pore spaces in subsurface geological formations. This is critical for addressing important problems such as enhanced oil recovery, groundwater pollution from leaking tanks or pipelines, geothermal energy production, and geological sequestration of carbon dioxide (CO_2).

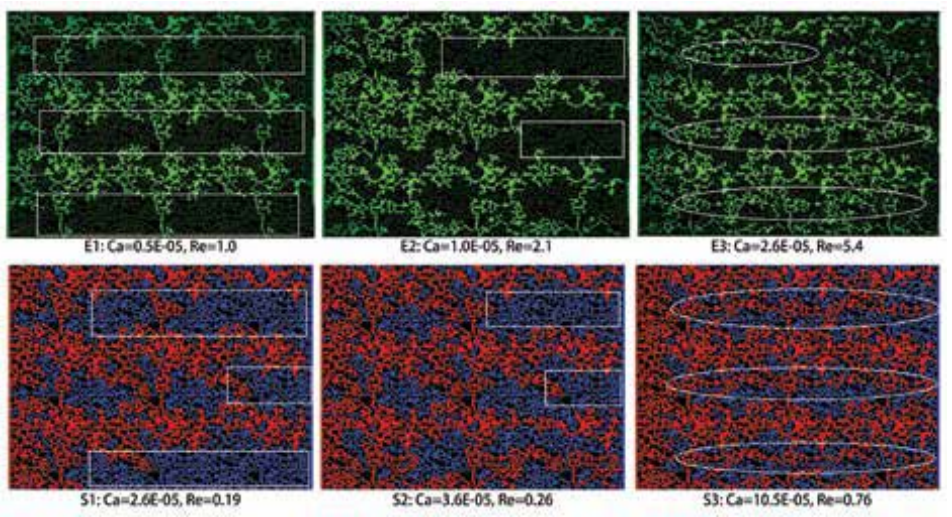

Comparison of CO_2 invasion patterns at approximately steady state between the simulations and experiments at different Ca. Red fluid represents liquid CO_2, blue fluid represents water, and solid grains are in black. E0, E1, E2, E3, and E4 denote the experimental results. S1, S2, and S3 denote the simulation results.

Methods & Codes

Lattice Boltzmann method (LB)

- Optimized a portable and scalable LB code based on a variant of the multiple relaxation time (MRT) color-fluid multiphase LB model.
- Employs a MPI-OpenMP/OpenACC hybrid programing model so that it can run on CPUs and GPUs.

Results & Impact

- Simulations and experiments of liquid CO_2 displacing water for reservoir pressure conditons were performed and compared for the first time.
- Development of secondary CO_2 pathways identified as a major discrepancy between experiment and simulation and further work determined responsible mechanisms.
- Published in *Journal of Contaminant Hydrology, 2017,* doi.org/10.1016/j.jconhyd.2017.09.005

Why Blue Waters

Blue Waters offers a large number of CPU and GPU nodes, which is essential to our research. A single iteration step in the LB simulation can be completed in a very short time, but a typical simulation requires millions of iterations. Therefore, LB simulation requires very low latency on message passing, which cloud resources cannot provide.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/6,600 Knh
PI: John Vidale (current PI)
Thomas Jordan (former PI)
University of Southern California
Geoscience

PHYSICS-BASED MODELING OF HIGH-FREQUENCY GROUND MOTIONS AND PROBABILISTIC SEISMIC HAZARD ANALYSIS

Research Challenge

The U.S. Geological Survey (USGS) currently uses empirical Probabilistic Seismic Hazard Analysis (PSHA) to promote seismic safety engineering and disaster preparedness across the United States, including California. SCEC's research goal is to develop physics-based seismic hazard models for California and elsewhere that are more accurate than the empirical USGS National Seismic Hazard Map Project standard models. The long-term goal is to extend physics-based PSHA across the full bandwidth needed for seismic building codes and other purposes.

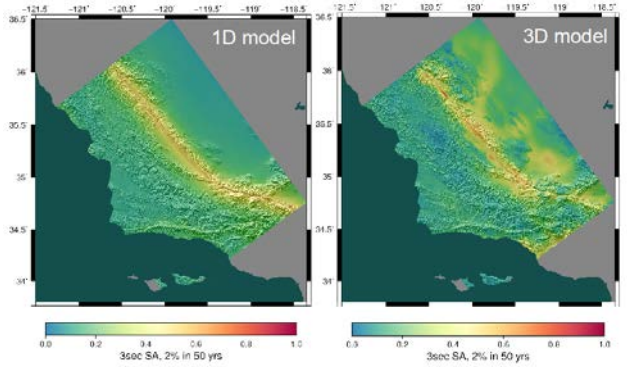

Seismic hazard maps for Central California from CyberShake Study 17.3 (left) a simple 1D seismic velocity model and (right) a 3D seismic velocity model were used by the CyberShake deterministic wave propagation simulations.

Methods & Codes

SCEC researchers added improved physics into their high-performance earthquake wave propagation software and improved the software performance on both CPUs and GPUs. They performed the first 4-Hz nonlinear magnitude 7.7 earthquake simulation using 4,200 GPUs, as well as CyberShake Study 17.3 that applied the CyberShake PSHA computational method to Central California for the first time. They Increased their use of third-party libraries including ADIOS, HDF5, and PnetCDF to scale up the I/O performance.

Results & Impact

CyberShake simulations are under review for use in new Los Angeles urban seismic hazard maps under development by the USGS. These new maps are being considered for the National Earthquake Hazards Reduction Program, the American Society of Civil Engineers 7–10 Seismic Provisions, and for the Los Angeles City building codes. SCEC's sustained work on Blue Waters is transforming and modernizing earthquake science and engineering, and thus represents a major contribution to Strategic Goal 1 of the NSF 2014–2018 Strategic Plan which is "Transform the Frontiers of Science and Engineering."

Why Blue Waters

New earthquake simulations that model advanced physics at high resolution require increasing amounts of computational, memory, and storage resources. Computational demands continue to grow because the calculations are not yet at full-resolution, because not all important physics have yet been included, and because individual earthquake simulations do not "solve" a problem when run just once or twice. Great uncertainty remains in the ground motions expected in future earthquakes, and urban society will be safer when better seismic hazard information is available about these critical scientific and public safety challenges.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: Illinois/300 Knh

PI: Matthew West
University of Illinois at Urbana-Champaign

Geoscience

3D PARTICLE-RESOLVED AEROSOL MODEL TO QUANTIFY AND REDUCE UNCERTAINTIES IN AEROSOL–ATMOSPHERE INTERACTIONS

Research Challenge

- The treatment of aerosol particles provides challenges in atmospheric modeling and simulation.
- Current models provide important insights, but do not resolve individual particles and their microscale interactions.
- This makes computation much cheaper, but it introduces unknown errors into model calculations. This has far-reaching consequences for the estimation of climate-relevant aerosol quantities, such as aerosols' ability to scatter and absorb sunlight as well as their ability to form clouds.

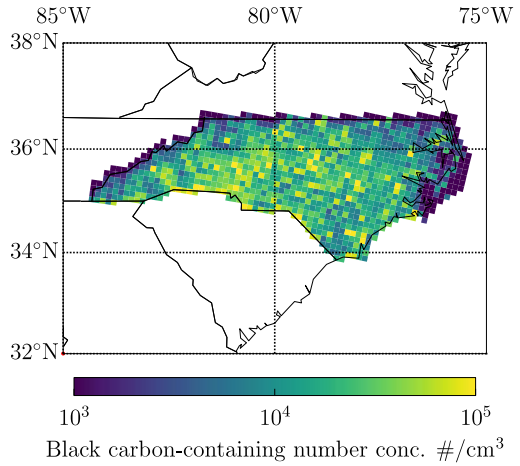

Horizontal distribution of simulated number concentration of black-carbon-containing particles near the surface.

Methods & Codes

The particle-resolved model PartMC-MOSAIC was coupled to the state-of-the-art 3D Weather Research and Forecast (WRF) model. These complement each other with the box model PartMC-MOSAIC handling the highly detailed aerosol model and the 3D regional WRF model capturing the transport of chemical species in the atmosphere. This next-generation model captures complex aerosol composition that current-generation models are unable to simulate.

Results & Impacts

- This research provides the first-ever particle-resolved aerosol simulation for a realistic, spatially resolved three-dimensional domain.
- On the order of 100 billion computational particles were tracked in this simulation.
- Future studies may be able to use this research to quantify how much individual source categories contribute to the pollution at a certain location.

Why Blue Waters

Cutting edge model formulations push both science and computing by combining the large-scale features of state-of-the-art 3D models with the process level physical representation of box models. Modeling 3D domains on the order of 100 billion tracked particles creates many computational challenges due to computationally intensive equations per particle and memory requirements to track high-dimensional particle composition. To simulate aerosols at both a high spatial and compositional resolution, tens of thousands of cores with fast interconnections among those cores, and sufficient memory per process are used.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/8,500 Knh
PI: Don Wuebbles
University of Illinois at Urbana–Champaign
Geoscience

HIGH-RESOLUTION EARTH SYSTEM MODELING USING BLUEWATERS' CAPABILITIES

Research Challenge

The team conducts multiple-century, high-resolution global climate model simulations to analyze past and projected changes in the earth's climate system. This work contributes to a larger body of ongoing research aimed at using high-resolution climate and weather forecast models to better understand high-impact events (such as tropical cyclones, storms, and extreme heat events) in the present, as well as future warmer scenarios.

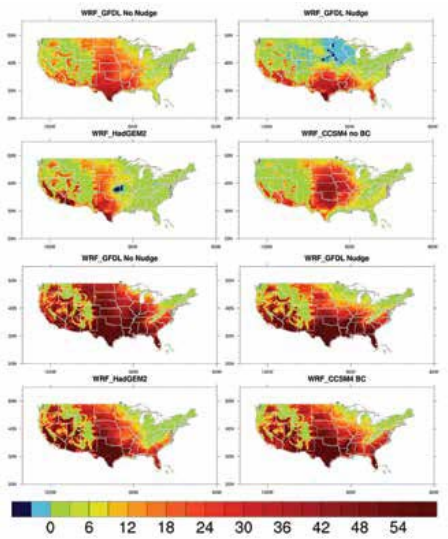

The change in the number of days that exceed 95°F between a moderate warming scenario and present day (top four panels) and a high warming scenario and present day (bottom four panels). The four panels show results from four different configurations of WRF using different boundary forcing (from either a global climate model or observations).

Methods & Codes

The team uses the Community Earth System Model (CESM) code, a coupled climate model for simulating the earth's climate system, to conduct fundamental research into the earth's past, present, and future climate states. Results from global climate model simulations can then be used to drive the Weather and Research Forecasting (WRF) code to examine the effects of climate change on smaller geographic regions.

Results & Impact

Simulations conducted for this project examined the frequency and strength of mid-latitude storms, as well as the the number of days that will exceed 95°F later this century. Insights gained from these simulations will help other scientists and policymakers understand the impacts that climate change will have on society and the surrounding ecosystem and develop plans for adaptation and mitigation.

Simulation results will be available to the wider scientific community for analysis by other science teams and organizations. They will also be used to inform the upcoming national and international climate assessments.

Why Blue Waters

First, CESM experiments begin with a multi-century long preindustrial control simulation. Then, multi-member ensembles of historical and future scenarios are run to quantify and reduce uncertainty. A high resolution (0.25°) atmosphere coupled with a 0.1°-horizontal ocean grid allows for full eddy-resolving ocean simulations. Using a 12-km grid, WRF simulations capture small-scale atmospheric processes like thunderstorms over the continental U.S. These high resolution simulations and analyses require petascale computing resources and cannot be completed without a system like Blue Waters.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: Illinois/500 Knh
PI: Donald J. Wuebbles
University of Illinois at Urbana–Champaign
Geoscience

Particulate Matter Prediction and Source Attribution for U.S. Air Quality Management in a Changing World

Research Challenge

The goal of this research is to determine the individual and combined impacts of global climate and emissions changes on U.S. air quality from the present to 2050 under multiple scenarios, quantifying pollution sources and assigning their attribution—natural vs. anthropogenic emissions, national vs. international agents, natural variations vs. climate changes—with associated probability and uncertainty. In summary, *this research presents a state-of-the-science approach for advancing quantitative knowledge of the impacts of global changes in climate and emissions on U.S. air quality.*

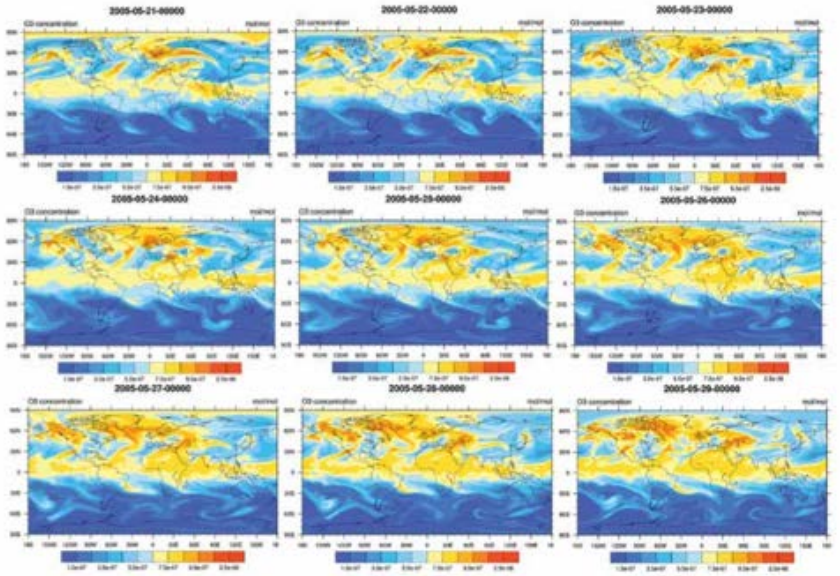

Evolution of daily ozone concentration for consecutive days in May 2005 from Cam5-chem global simulation.

Methods & Codes

The project uses the Community Earth System Model (CESM) in three phases:
(1) historical simulations for the period 1994–2013 to establish credibility;
(2) projections for 2041–2060 to quantify individual and combined impacts of global climate and emissions changes under multiple scenarios; and
(3) sensitivity analyses to determine future changes in pollution sources and their relative contributions.
The runs are based on CESM1.2 with fully coupled chemistry using CAM-5.

Results & Impact

Based on short-term global CAM5-chem simulations, the figure shows the global O_3 concentration for a period in May 2005. Previous studies have shown that trans-Pacific dust and aerosols contribute significantly to North American aerosol inflow, while export-related Chinese pollutants contributed 3–10% of annual mean sulfate concentration, 0.5-1.5% of ozone, and one or more day of noncompliance of ozone standards over many U.S. regions in 2006. A variety of additional modeling simulations like these are planned.

Why Blue Waters

The computational demand of the high-resolution global climate model used in this project is extensive. Additionally, the fully coupled model of the Earth's climate system with chemistry is computationally expensive. Blue Waters, with its petascale computational facility and storage capability for the output from the high-resolution model simulation, is essential for this project. Blue Waters staff have been critical in figuring out the various issues arising with the long-term fully coupled climate chemistry runs with CESM.

Physics and Engineering

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: GLCPC/290 Knh
PI: David Ackerman
Iowa State University
Physics & Engineering

EXPLORING CONFINEMENT VS. ORIENTATION EFFECTS IN RIGID AND SEMI-FLEXIBLE POLYMERS USING A MASSIVELY PARALLEL FRAMEWORK

Research Challenge

The research goal is to develop and utilize a highly parallel computational framework to model equilibrium structures of semi-flexible polymers efficiently in complex, confined, non-periodic geometries. The simulation system is designed to be capable of scaling to large numbers of processors for large, complex systems of semi-flexible polymer systems. These tools will enable the ultimate goal of studying confinement effects on semi-flexible chains.

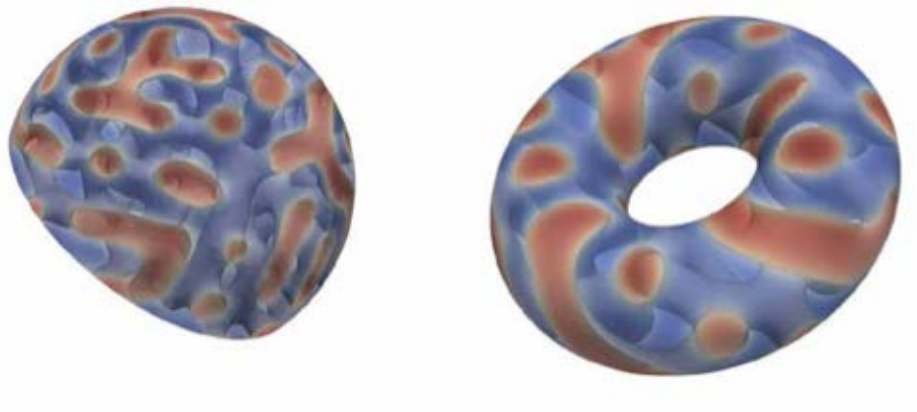

Example polymer microstructures generated with the framework developed in this project.

Methods & Codes

We built a custom finite element library from the Portable, Extensible Toolkit for Scientific Computation (PETSc). We use a simultaneous solution over the evolving chain contour space that enables efficient adaptivity and scaling. To model semi-flexible polymers, we use a two-level finite element method in which the entire spatial system is a finite-element volume mesh. Coupling the simulation spaces enables orientationally-dependent systems to be simulated.

Results & Impact

The development of a scalable finite element SCFT code utilized in a high-throughput study of polymer confinement, which will be used in several future studies.

The implementation of a simultaneous space–time finite-element method which is invaluable for a method aimed at enhancing scaling.

The development of a framework for modeling semi-flexible polymers which enables morphology studies of polymers with complex, orientation-dependent properties.

Why Blue Waters

The finite-element framework has been designed to model systems with large numbers of nodal points. At each finite-element nodal point there are 100 or more degrees of freedom to solve. For even the smallest problem, this results in billions of unknowns evolving under complex physical processes. The process requires significant, sustained computational resources. Utilizing a highly scalable framework and thousands of nodes makes these structure determinations feasible. The processing resources are not available outside the Blue Waters system.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: Illinois/250 Knh
PI: Jean Paul Allain
University of Illinois at Urbana-Champaign
Physics & Engineering

HARNESSING PETASCALE COMPUTING TO EXPLAIN FUNDAMENTAL MECHANISMS DRIVING NANOPATTERNING OF MULTICOMPONENT SURFACES BY DIRECTED IRRADIATION SYNTHESIS

Research Challenge

To learn more about the physical mechanisms of ion-induced changes on surfaces at measurements far smaller than what can be experimentally observed. With more knowledge, researchers hope to use ion beam techniques to create innovative nanomaterials in a single, scalable process step, which would be extremely useful in producing semi-conductors, biomaterials, thin films for energy production, and other advanced applications.

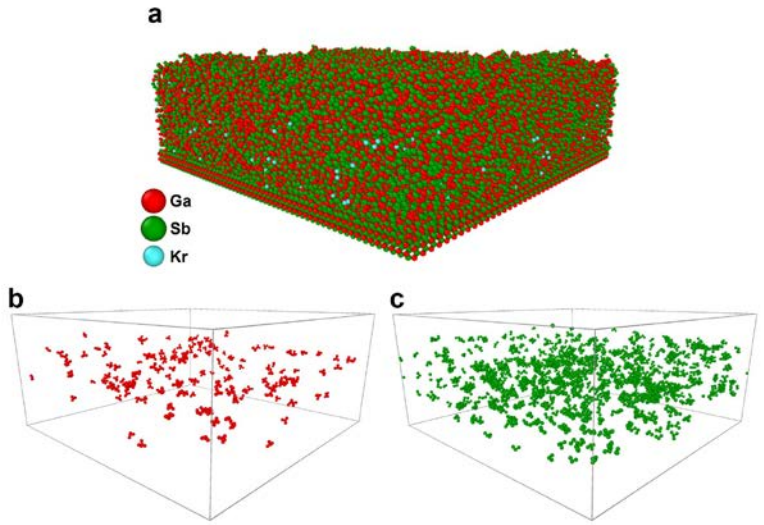

Snapshots of the GaSb surface after irradiation by 500 eV Kr^+ ions to a fluence of 7.5 × 10^{15} cm^{-2}—(a) the entire surface (b) only Ga atoms in clusters; (c) only Sb atoms in clusters.

Methods & Codes

Molecular dynamics simulations using LAMMPS for 500 eV Ne^+, Ar^+, and Kr^+ ion species incident on an initially pristine GaSb surface. The simulation cells were 25 × 25 nm^2. For each ion species, the GaSb surface was irradiated to a fluence of 7.5 × 10^{15} ions/cm^2, an experimentally relevant fluence at which compositional depth profile evolution has been observed.

Results & Impact

No compositional depth profile was observed, indicating more study with longer timescales is needed in this area. Formation of small clusters of gallium (Ga) or antimony (Sb) were seen, inducing compositional changes in the surface, which may provide “seed” structures to grow into larger-scale lateral compositional gradients. Ion-induced structural changes provide potential pathways for mechanisms at long temporal scales, such as diffusion, to create compositional depth and lateral profiles.

Why Blue Waters

This work required carrying out three molecular dynamics simulations in parallel, each of which contained around 150,000 to 200,000 atoms at any time, for over 230 million time steps each. The simulations also generated 300,000 output data files totaling 2.3 TB in size. Without Blue Waters, this would have required several years to complete using a conventional computing cluster — if it were even possible to do.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: Illinois/975 Knh
PI: Narayana R. Aluru
University of Illinois at Urbana-Champaign
Physics & Engineering

DETECTION OF ANTIBODY SUBCLASSES USING A NANOPOROUS SINGLE-LAYER GRAPHENE

Research Challenge

- Distinguishing the different subclasses of immunoglobulin G (IgG) antibodies could enable breakthrough advances in mapping the immune system and the human health management.
- This study performs molecular simulations (up to ~1,000,000 atoms) and a total aggregate simulation time of 2.7 microseconds (μs).

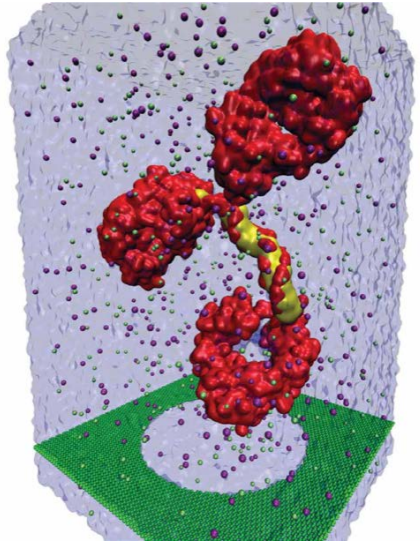

System consisting of IgG3 protein (red: chains; yellow: disulfide bonds), ions (pale green), and graphene sheet (green).

Methods & Codes

- Uses NAMD to do complex Molecular Dynamic simulations of antibody/nanopore structures
- Use Machine Learning to determine features and to cluster the ionic current and the dwell time data during multiple antibody translocation events

Results & Impact

- The study determined that an atomically thin graphene nanopore is capable of sensing and discriminating among different subclasses of IgG antibodies despite minor and subtle variations in atomic structure.
- The histogram of ionic current for each segment of IgG can provide high-resolution spatial detection of antibody segments.
- Parallel nanofluidic studies during IgG translocation reveal distinct water flux rates for IgG subclasses facilitate additional recognition mechanism.

Why Blue Waters

The size of the atomic structures being simulated (1,000,000 atoms) and the fact the NAMD molecular dynamics package scales almost linearly with the number of cores up to 1,000 requires a leadership class computing system such as Blue Waters.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/3,120 Knh
PI: Jerzy Bernholc
North Carolina State University
Physics & Engineering

PETAFLOPS SIMULATION AND DESIGN OF NANOSCALE MATERIALS AND DEVICES

Research Challenge

This project focuses on high-performance calculations for materials and devices, and on development of petascale methods for such simulations. The research challenges are the investigation of sensor configurations for detection of biologically important molecules, the analysis of high-performance dielectric materials that store and release energy electrostatically through polarization and depolarization, and the pursuit of atomically precise and bottom-up fabrication of graphene-based electronics.

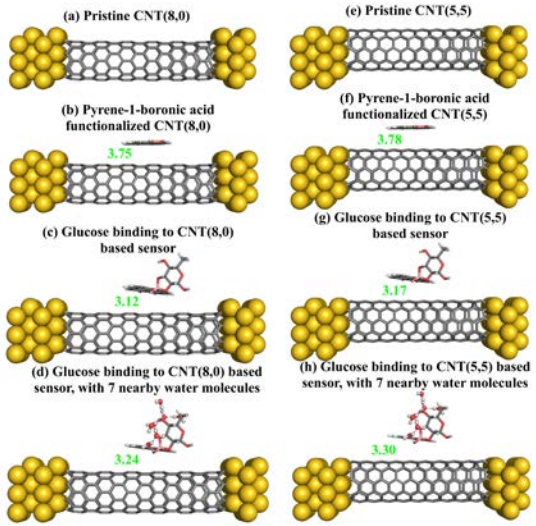

Glucose detection in the semiconducting nanotube CNT(8, 0)-based systems (a, b, c, d) and the metallic CNT(5, 5)-based systems (e, f, g, h). The closest distances between the pyrene-1-boronic acid molecule and the nanotube surface (in Å) are labeled in green.

Methods & Codes

The quantum transport and large-scale electronic structure calculations rely on the use of the RMG code developed by the team. Polymer simulations employ the LAMMPS code. Calculations that include van der Waals interactions use the PWSCF code.

Results & Impacts

The studied glucose sensor configuration and mechanism opened the path to the design of other nanotube sensors. The project found that blending a poly(arylene ether urea) (PEEU, K = 4.7) and an aromatic polythiourea (ArPTU, K = 4.4) leads to a compound that exhibits a very high dielectric constant, K = 7.5, while maintaining low dielectric loss (< 1%.) The team established how the bottom-up synthesis of a graphene nanoribbon can be controlled by charge injections from a scanning microscope (STM) tip.

Why Blue Waters

The conducted applications require very large compute and high-speed interconnect resources that are available on Blue Waters. Each project requires performing many runs to explore the various scientific issues, with a substantial amount of analysis between the runs. High availability and quick turnaround that Blue Waters provide are essential for success of the project.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: Illinois/280 Knh
PI: Huck Beng Chew
University of Illinois at Urbana-Champaign
Physics & Engineering

SCALABLE NANOPATTERNING OF GRAPHENE BY HYDROGEN- PLASMA ETCHING

Research Challenge

While there exists ample experimental evidence for the patterning of graphene by hydrogen-plasma treatment, the reported etching reactions and the resulting graphene nanostructures have been vastly different. The complete parameter space of substrate temperature, ion energy, and incident flux has not been systematically studied due to the cost limitations of plasma experiments.

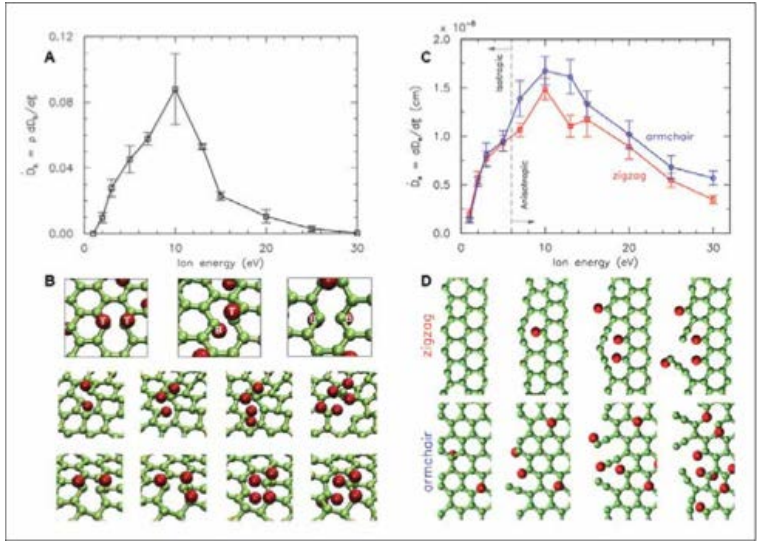

a) Steady state basal plane etching rate versus ion energy for monolayer graphene; b) Mechanism of etching showing the three possible configurations of damage nucleation and progression; c) Etching rates of the zigzag and armchair configurations versus ion energy; and d) Mechanism of etching for zigzag and armchair edges.

Methods & Codes

The team performed length-scale bridging by delineating the contributions of the edge and basal plane etching using ReaxFF- based molecular dynamics (MD) and linking these processes together via a mechanistic model. The simulations were performed with the C++-based open source LAMMPS code.

Results & Impact

Results demonstrate distinct ion energy regimes for isotropic versus anisotropic etching. These distinctive etching mechanisms, which are operative within narrow ion energy regimes, fully explain the differing plasma-graphene reactions observed experimentally.

Why Blue Waters

The Blue Waters computational capacities were necessary for several reasons. First, the complex chemistry and plasma surface interactions involved in the hydrogen etching of graphene require the use of fully reactive MD potential, allowing for potential reactions between the Si, O, C, and H species at each MD time step. Second, the impact dynamics of impinging H atoms on graphene requires the use of a small time step (0.15 fs), which further increases computational cost. Third, studying the edges of the multilayer graphene presents a wide range of possible configurations, as the edges can be partially or completely covered by a graphene layer. Finally, because of the random process of H deposition, a large number of simulation runs are required to obtain statistically significant findings.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/3,000 Knh
PI: J. P. Draayer
Louisiana State University
Physics & Engineering

INNOVATIVE *AB INITIO* SYMMETRY-ADAPTED NO-CORE SHELL MODEL FOR ADVANCING FUNDAMENTAL PHYSICS AND ASTROPHYSICS

Research Challenge

One of the quintessential open problems in contemporary physics is to design a comprehensive many-body theory for modeling and predicting nuclear structure and reactions. As short-lived nuclei, currently inaccessible to experiment, are often found to be key to understanding processes in extreme environments from stellar explosions to the interior of nuclear reactors, first-principle nuclear models that hold predictive capabilities have had and will have a tremendous impact on advancing our knowledge at the frontiers of multiple branches of physics. This project uses *ab initio* ("from first principles") theory to find a solution to this problem.

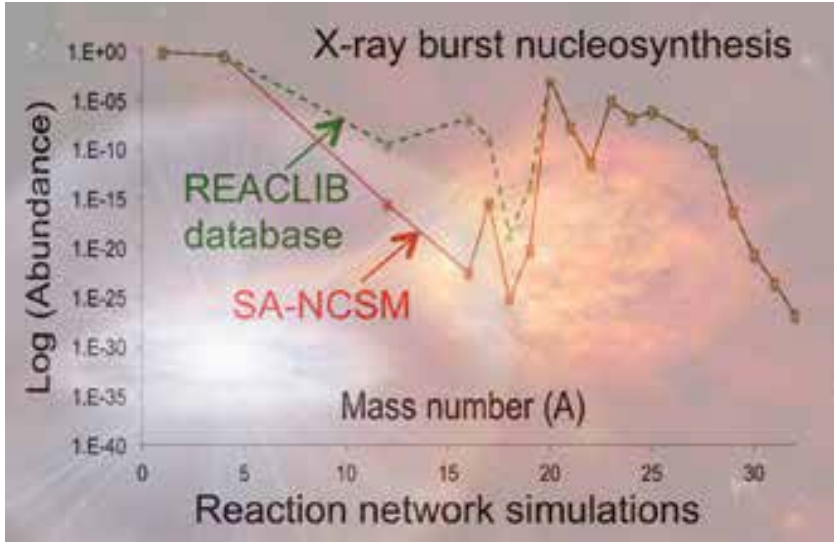

Effect on the abundance pattern from X-ray burst (XRB) nucleosynthesis simulations (based on Hix's Xnet) when reaction rates from the BW-enabled first-principle SA-NCSM simulations of ^{20}Ne are used (compared to current database, for fixed astrophysical conditions).

Methods & Codes

The team has developed an innovative *ab initio* nuclear structure approach, dubbed the symmetry-adapted no-core shell model (SA-NCSM), simulated using the computer code "LSU3shell", which embraces the first-principles concept and capitalizes on a new symmetry of the nucleus. The *ab initio* SA-NCSM solves the time-independent Schrödinger equation as a Hamiltonian matrix eigenvalue problem. These theoretical advances, coupled with the computational power of the Blue Waters system, allow them to reach medium-mass nuclei that are inaccessible experimentally and to other *ab initio* methods.

Results & Impact

- Provided the first *ab initio* description of the open-shell ^{19}Ne, ^{20}Ne, ^{24}Ne and ^{24}Si nuclei, with impact on stellar mergers simulations
- Studied emergent phenomena from first principle in Mg isotopes and their mirror nuclei
- Studied ^{12}C, including the most challenging Hoyle state
- Performed first-principle simulations of ^{48}Ca and ^{48}Ti, with impact on studies of neutrinos and physics beyond the standard model
- Made dramatic improvements to the LSU3shell code

[Ne – Neon, Si – Silicon, Mg – Magnesium, C – Carbon, Ca – Calcium, Ti – Titanium]

Why Blue Waters

SA-NCSM investigations of the intermediate-mass region are beyond the scale of available academic HPC systems. The SA-NCSM drastically reduces the size of the problem and the associated memory requirement down to hundreds of terabytes and petabytes, but this comes at the cost of a major increase in computation. Currently, only the BW system provides resources required for the *ab initio* SA-NCSM studies of medium-mass isotopes. The largest production runs efficiently 717,600 threads running on 22,425 Cray XE6 nodes to solve the nuclear eigenvalue problem with Hamiltonian matrices that occupy up to 400 TB of memory.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/1,100 Knh

PI: Lian Duan
Missouri University of Science & Technology
Physics & Engineering

DNS OF PRESSURE FLUCTUATIONS INDUCED BY SUPERSONIC TURBULENT BOUNDARY LAYERS

Research Challenge

To advance fundamental understanding of the generic statistical and spectral features of boundary-layer-induced pressure fluctuations in wall-bounded turbulence, including the freestream acoustic radiation at supersonic speeds and their dependence on boundary-layer parameters such as the Reynolds number

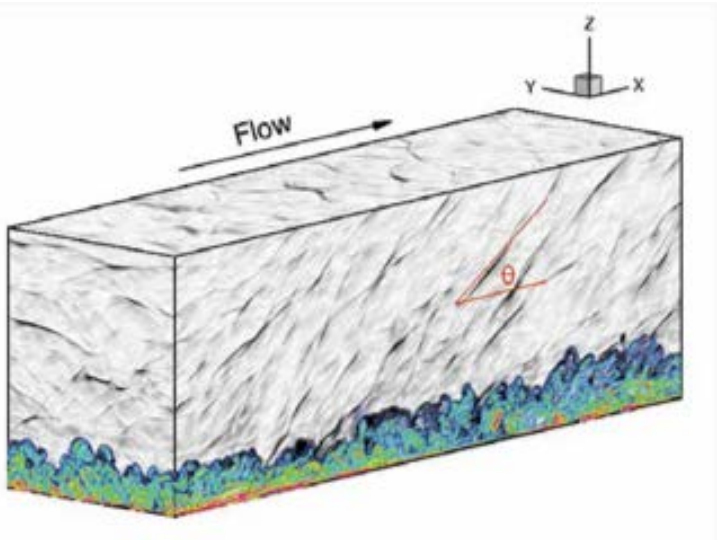

Instantaneous flow visualization for a Mach 6 cold-wall turbulent boundary layer at $Re_\tau \approx 450$. The freestream acoustic noise is visualized using numerical schlieren and the boundary layer is colored by the vorticity magnitude. The angle θ illustrates the preferred direction of radiated acoustic wavefront.

Methods & Codes

Direct Numerical Simulations that solve the compressible Navier-Stokes equations are conducted using an in-house, high-order finite-difference called called HyperWENO (Weighted, Essentially Non-Oscillatory)

Results & Impact

A better understanding of pressure fluctuations is an important aspect of vibrational loading and potentially damaging effects of fatigue and flutter in aircraft

From a theoretical point of view, pressure is of fundamental importance to understanding the turbulent vorticity dynamics and to modeling the pressure-strain terms in the Reynolds stress closure.

Why Blue Waters

Computational demands increase with Reynolds Number, and the proposed research wishes to extend understanding into the regime of higher Reynolds Numbers, using extremely fine meshes and a large number of time steps

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: Exploratory/62.5 Knh
PI: Levent Gurel
University of Illinois at Urbana–Champaign
Computer Science & Engineering

PARALLELIZATION OF THE MULTILEVEL FAST MULTIPOLE ALGORITHM (MLFMA) ON HETEROGENEOUS CPU-GPU ARCHITECTURES

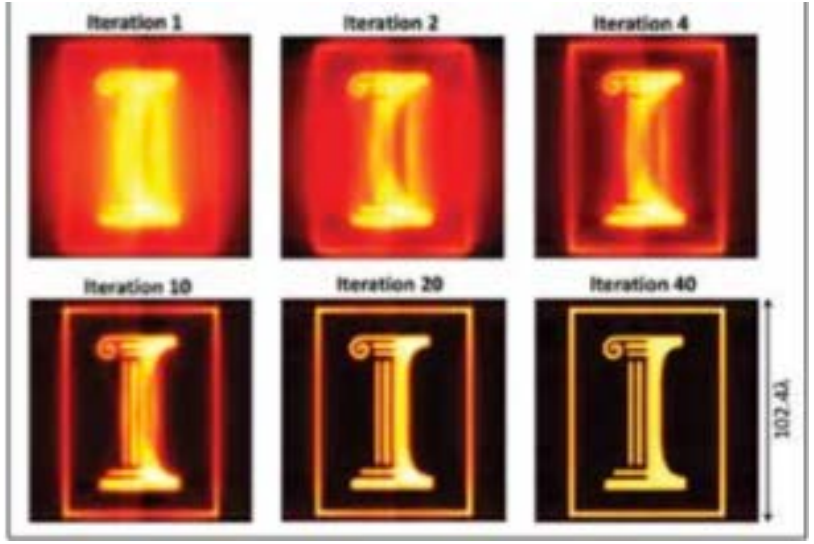

Convergence of the iterative inverse-scattering algorithm. The image details can be well seen after 40 iterations. It takes about two minutes to solve 38,400 forward-scattering problems with MLFMA on 256 XK nodes, and this is 4.34 times faster than employing the same of XE nodes.

Research Challenge

The goal of the experiment was to increase the speed of the Multilevel Fast Multipole Algorithm (MLFMA) across multiple Blue Water nodes, and tom utilize GPUs.

The nature of the algorithm was not easily spread across multiple CPU/GPU to increase speed for computations.

The complexity of calculations could not be easily run on GPUs

Methods & Codes

The available memory on Blue Waters enabled much larger simulations to be performed.

A different compute mechanism was deployed to better utilize Blue Water resources to lower idle time.

Code was developed (Fortran, C++, and CUDA) specifically for this purpose.

Results & Impact

By using GPUs effectively, researchers were able to obtain a 3.97 times speed up compared to using only CPUs, while also achieving a 96% parallelization efficiency.

Using parallelization, the compute time dropped from 11.5 hours (sequential execution) to 7 seconds using 128 GPU nodes.

This allowed near-real-time image viewing.

Why Blue Waters

The available amount of RAM and GPUs on Blue Waters allowed for research to be performed at a much larger scale than had previously been achieved.

Blue Waters staff also provided expertise that enabled researchers to solve many of the issues developing the code base.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: BW Professor/30 Knh
PI: Kathryn Huff
University of Illinois at Urbana-Champaign
Physics & Engineering

COUPLED MULTI-PHYSICS OF ADVANCED MOLTEN SALT NUCLEAR REACTORS

Research Challenge

- The current state-of-the-art advanced nuclear reactor simulation (e.g., the CASL DOE innovation hub) is focused primarily on traditional light-water reactor design types.
- This work extends state of the art nuclear reactor simulation for more advanced reactor designs that have the potential to improve the safety and sustainability of nuclear power.
- High-fidelity simulation of dynamic reactor performance of these designs requires development of models and tools for representing unique materials, geometries, and physical phenomena.

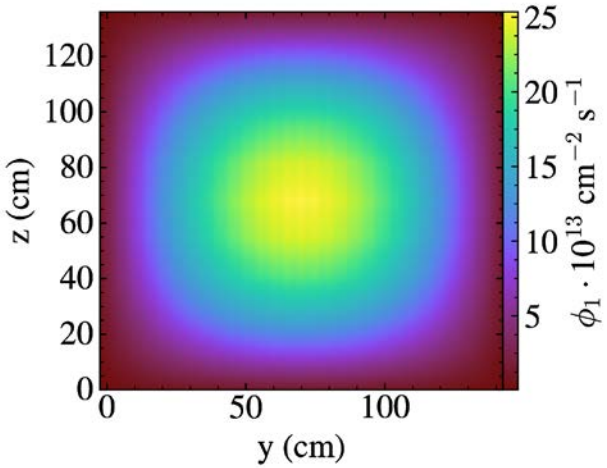

The neutron flux in a 2D cylindrical axisymmetric model of a molten-salt reactor. This flux has the anticipated magnitude and canonical cosine shape (r = 0 is center of core) and is undergoing validation against experimental results from the Molten-Salt Reactor Experiment

Methods & Codes

- Moltres is a collection of physics kernels material definitions, to extend the ecosystem of applications built on the highly scalable, fully implicit, Multiphysics Object-Oriented Simulation Environment (MOOSE) framework from Idaho National Laboratory.
- MOOSE and LibMesh translate data from Moltres into inputs for solution routines.

Results & Impact

- A first-of-its-kind, scalable, finite-element model of the transient neutronics and thermal hydraulics in a liquid-fueled molten salt reactor design.
- Future Moltres work includes generating a high-fidelity, 3D model as well as investigating various transient accident scenarios, additional reactor configurations, and numerous design concepts.

Why Blue Waters

There is a need for many CPU cores to process the many two-dimensional and three-dimensional finite-elements needed to assess nuclear reactor performance under a various conditions and dynamic transients. These simulations often occupy tens of thousands of CPU cores at a time and vary in completion time. Transient and multiscale simulations, which require greater capability per simulation, are on the horizon for our work. These may occupy up to 100,000 CPU cores at a time.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/9,000 Knh
PI: Sohrab Ismail-Beigi
Yale University
Physics & Engineering

UNDERSTANDING HYDROGEN STORAGE IN METAL ORGANIC FRAMEWORKS USING MASSIVELY-PARALLEL ELECTRONIC STRUCTURE CALCULATIONS

Research Challenge

- Efficient storage materials are required to make hydrogen practical for green energy storage.
- Metal-organic frameworks (MOFs) may provide improved hydrogen storage.
- Quantum mechanical simulations based on density functional theory can model MOFs, but is computationally expensive.

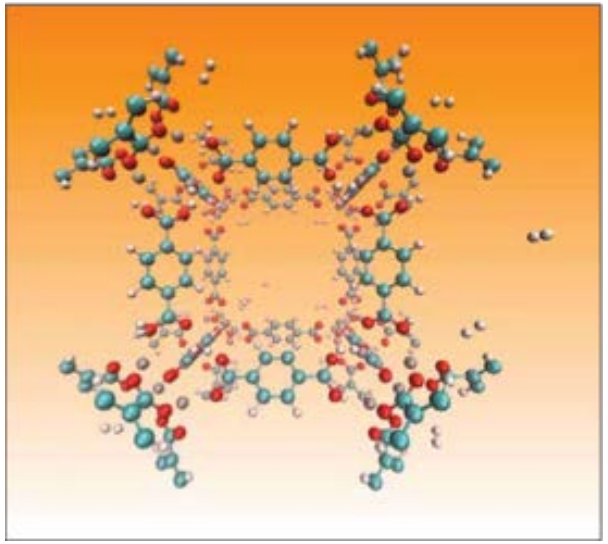

The model Zn-based MOF being studied that includes 43 hydrogen molecules (a total of 510 atoms per unit cell). Only 1/8 of the unit cell is shown.

Methods & Codes

- The Path Integral Car–Parrinello Molecular Dynamics simulation technique (PI-CPAIMD) implements this model.
- PI-CPAIMD Requires a highly-scalable implementation and petascale computational resources.

Results & Impacts

- Initial studies have examined a model MOF system.
- A smaller version of the MOF was studied with CPAIMD long simulations to obtain preliminary results.
- Good agreement with prior results increases confidence that benchmarking has produced physically accurate results.

Why Blue Waters

- Model requires tightly coupled computing nodes; electron waves are delocalized, so all parts of the system end up interacting.
- CPAIMD simulations of the MOF simulation already require massively parallel calculation with hundreds of nodes.
- Including nuclear quantum effects means that only a petascale computer like Blue Waters can deliver results on a reasonable (one- to two-year) timescale.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: Exploratory/30 Knh
PI: Prashant K. Jain
University of Illinois at Urbana–Champaign
Physics & Engineering

ATOMISTIC MODELING OF TRANSFORMATIONS IN NANOCRYSTALS

Research Challenge

Engineered nanocrystals have utility as new materials for battery electrodes, electronics, and optical devices such as infrared scanners. The team uses unconventional methods to engineer the chemical composition and crystal structure of nanocrystals to create new materials for potential use in energy storage and device efficiency. One discovery was a novel crystal form of mercury–cadmium selenide. Blue Waters was used to learn more about the properties of this new material.

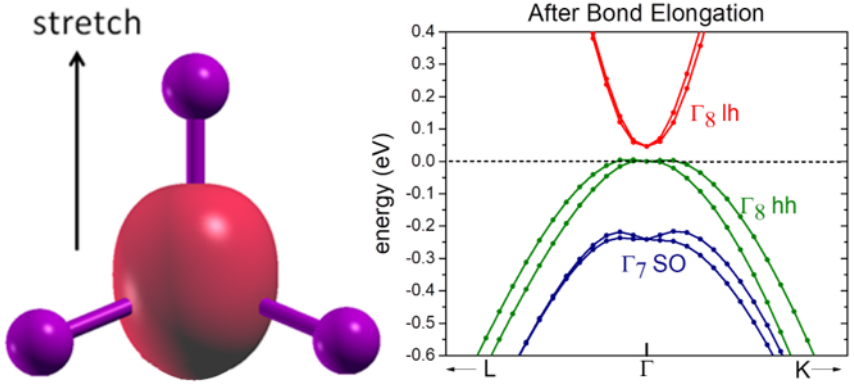

In the novel wurtzite form of HgSe, uniaxial elongation of Hg–Se bonds results in opening a band gap between the highest energy valence (green curve) and lowest energy conduction (red curve) bands. Hg atoms are shown in purple and the electron density is shown in red. Figure from *Chem. Mater.* **2017**, 29, 6356.

Methods & Codes

- Used the open-source Quantum Espresso software suite
- Numerous calculations were run
- Each calculation is distinct from the others in terms of the crystal geometry, chemical formula, or both, of the material.

Results & Impact

The team found that a band gap is created by bond elongation in a novel crystal polymorph of mercury selenide (HgSe) and mercury cadmium selenide. Combined with inverted bands in HgSe, the band gap qualifies this new material as a potential 3D topological insulator (TI). 3D TIs could one day lead to energy-efficient chips that run cooler, permitting the design of more powerful supercomputers by overcoming the problems caused by overheating.

Why Blue Waters

Systematic study of chemical trends requires the creation of a data library from numerous calculations of structures of varying elemental compositions. The computational expense would be prohibitive if not for a Blue Waters allocation. Also, Blue Waters' interconnect/communications hardware lets the Quantum Espresso code run more efficiently as the code's parallelization schemes involve sizable and frequent communication among CPUs. Another plus was access to NCSA staff for help with code preparation and troubleshooting.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: GLCPC/880 Knh
PI: Eric Johnsen
University of Michigan
Physics & Engineering

NUMERICAL SIMULATIONS OF COLLAPSING CAVITATION BUBBLES ON BLUE WATERS

Research Challenge

Cavitation occurs when local pressure reductions of a liquid lead to the formation of vapor bubbles that can rapidly expand and undergo violent collapse, resulting in shock waves with high pressure temperature regions that can damage neighboring solid objects such as propellers or soft tissues. High-fidelity numerical simulations of the collapse of individual and multiple vapor bubbles can provide a clearer image of the detailed nonspherical bubble dynamics, pressure, and temperature fields, and stresses/deformations of the neighboring solid.

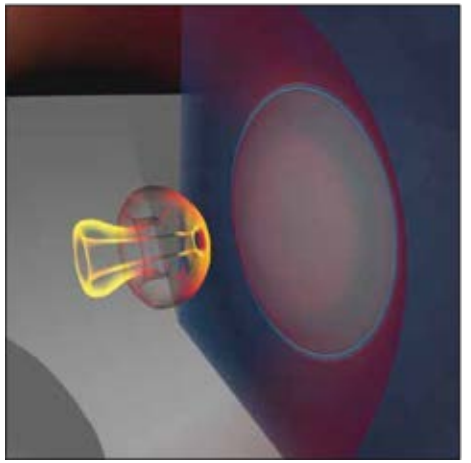

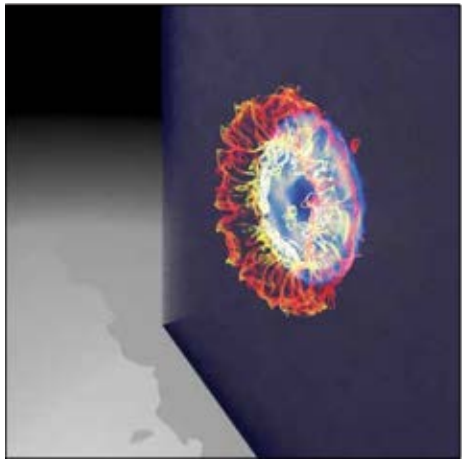

Left: Cavitation bubble secondary collapse due to reflected shock wave.

Right: Afterward, the bubble takes the form of a convoluted vortex ring, which is convected toward the adjacent wall.

Methods & Codes

Developed a computational framework for massively parallel simulations of the three-dimensional compressible Navier-Stokes equations for gas-liquid flows, extended to include viscoelastic deformations of neighboring solids.

The algorithm differentiates between smooth and discontinuous regions of the simulation, applying nondissipative methods where the solution is smooth, and computationally expensive dissipative schemes in discontinuous regions.

Results & Impact

Provides a clearer image of the detailed nonspherical bubble dynamics, pressure, and temperature fields, and stresses/deformations of the neighboring solid

Enables establishment a numerical cloud model that includes multiple bubbles and bubble–bubble interactions in bubble-clouds.
Algorithm has been shown to be over 83% efficient

Why Blue Waters

Blue Waters offers the high degree of parallelism essential to model the high spatial resolution required for more accurate simulations
Blue Waters enables scaling from single-bubble simulations to establishing new models and simulations of multiple interacting cavitation bubbles to more accurately model real-world scenarios

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/1,239 Knh
PI: Gerhard Klimeck
Purdue University
Physics & Engineering

LEADING FUTURE ELECTRONICS INTO THE NANO REGIME USING QUANTUM ATOMISTIC SIMULATIONS IN NEMO5

Research Challenge

The semiconductor industry influences approximately 10% of the world GDP, either directly or through electronic systems. Continuous speed and power consumption improvements must be made to the transistor technology for the stability and growth of the industry. Further improvements in shrinking dimensions will come only through the detailed study of device designs, materials, and of quantum effects such as tunneling, state quantization, and atomistic disorder. This project addresses these issues and others on a variety of semiconductor devices through simulation.

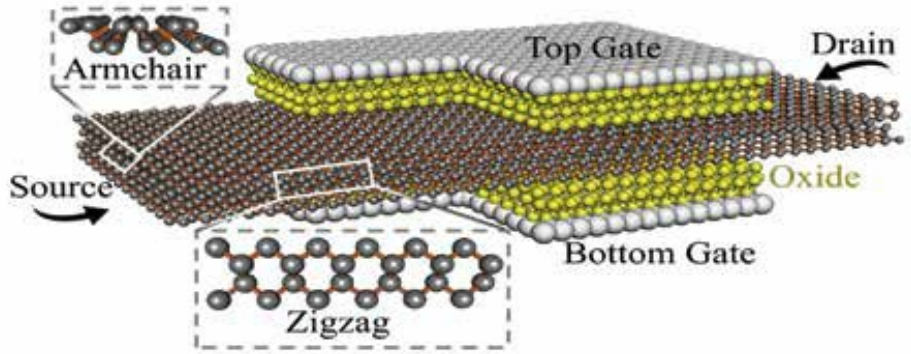

A 3D visualization of a novel sub-10 nm transistor design made from bilayer phosphorene that is being investigated using NEMO5. (Credit: Tarek Ameen)

Methods & Codes

The team uses NEMO5, available at nanoHUB.org, to perform multi-physics atomistic quantum simulations. Topics of study include quantum transport simulations for transistors, the current mechanism of tunneling field-effect transistors, tunneling transistors, quantum computing devices such as quantum dots, multi-scale quantum transport simulations for certain LEDs, and others.

Results & Impact

- Nitride device simulations have been used to suggest improvements to the multi-quantum-well nitride-based LED.
- Simulations have led to the proposal of a new alloy engineered nitride tunneling field-effect transistor as a novel low-power transistor design.
- A new computational method for modeling incoherent scattering phenomena in electron transport, which provides shorter times-to-solution and smaller memory footprints, has been added to NEMO5.

Why Blue Waters

Quantum transport fundamentally deals with non-equilibrium phenomena that involve multi-particle interactions. Such simulations are conceptually and computationally more demanding that traditional ab-initio materials simulations. Blue Waters was used for running such simulations on up to 16,384 cores per simulation. In many cases the work could not be accomplished in a reasonable amount of time without Blue Waters, and for the larger simulations the work could not be accomplished on other available systems.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: Illinois/52 Knh
PI: Deborah A. Levin
University of Illinois at Urbana-Champaign
Physics & Engineering

MODELING PLASMA FLOWS WITH KINETIC APPROACHES USING HYBRID CPU-GPU COMPUTING

Research Challenge

With the space environment becoming a home to constellations of small satellites and cube improved predictability of key surfaces of solar cell arrays and spacecraft charging in the backflow environment of chemical and electric-propulsion (EP) thrusters is crucial. Indirect environmental exposure of spacecraft material can cause appreciable sputtering and erosion, which is hard to quantify and predict over the lifetime of the mission. These processes need to be modeled in spacecraft environment effect models.

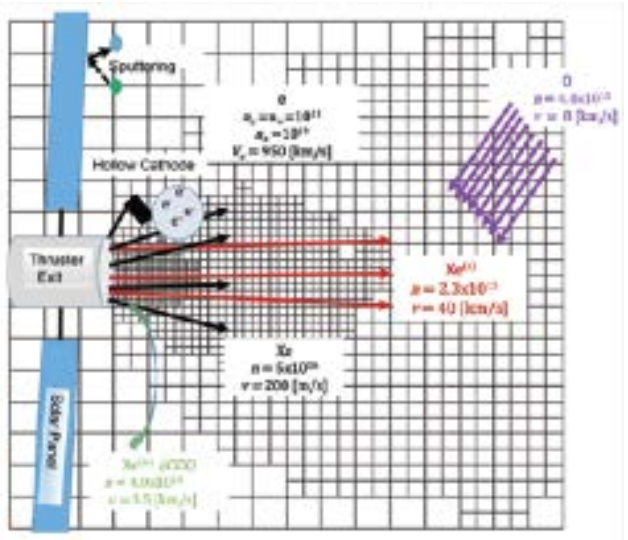

Key species, velocities, and concentrations (m^3) in the backflow and beam region for a spacecraft in a low- to mid-earth orbit space environment. Approximate scale is 1 m.

Methods & Codes

We use our new plasma modeling, which is an outgrowth of our DSMC code, CHAOS (Cuda-based Hybrid Approach for Octree Simulations) to include both neutral and ion species in an external electric field. Implementation of this on an AMR/octree grid is nontrivial and has been accomplished through the use of local (on a single processor) and local–global (across processors) stages.

Results & Impact

From the simulations of mesothermal, collisionless plasmas for shifted electron and ion sources, we discovered very interesting unsteady plume dynamics. The time-varying behavior is very different for the ions versus the electrons. The electron dynamics were found to be complex, with the continual exchange between electron kinetic and potential energy resulting in a meandering/bouncing movement.

Why Blue Waters

Compared to the present state-of-the-art plasma simulations, a uniform grid in 3D would require about a factor of seven more cells than our use of AMR/octree. The use of a single K20x GPU decreases the runtime by a factor of five compared to a single Interlagos processor. Very conservatively, we estimate that the octrees in combination with GPUs, decrease the total runtime by at least a factor of 10, compared to uniform grid solvers on multi-core CPUs.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/21,780 Knh
PI: Paul Mackenzie
Fermi National Accelerator Laboratory
Physics & Engineering

HIGH ENERGY PHYSICS ON BLUE WATERS

Research Challenge

The goal of this project is to carry out groundbreaking studies of the Standard Model of high-energy physics. These calculations are a part of the theoretical side of high energy physics. They address fundamental questions in high-energy physics and support and complement the very large, worldwide experimental studies in this field. This project is a joint effort of the Fermilab Lattice, MILC, and RBC Collaborations. This research aims at deeper understanding of the standard model of elementary particles and at searching for physical phenomena that go beyond it.

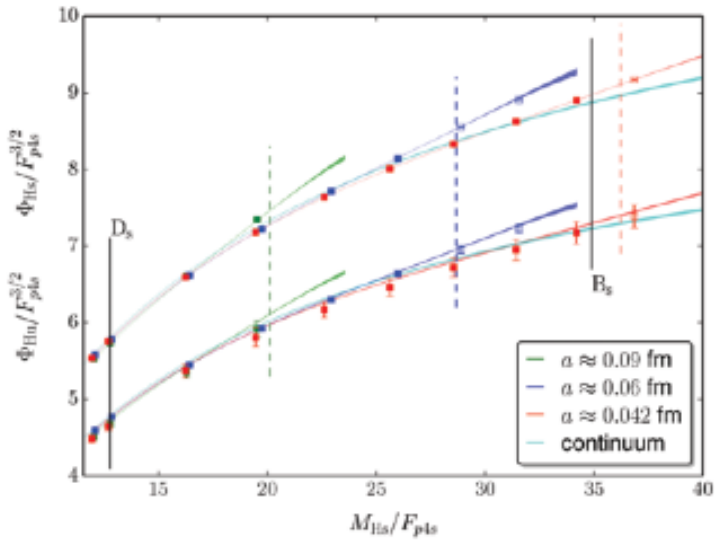

Lattice data (squares) with a chiral-continuum-HQET fitting function. The solid curves are fit results for single lattice spacings and for the continuum limit. The quantity plotted on the x-axis is proportional to the strange quark mass, and those on the y-axis are proportional to the leptonic decay constants for the B and Bs mesons.

Methods & Codes

Lattice quantum chromodynamics (QCD) calculations proceed in two steps. In the first, one uses importance-sampling techniques to generate gauge configurations, which are representative samples from the Feynman path integrals that define QCD. These configurations are saved; in the second step they are analyzed to calculate a wide variety of physical quantities such as mass and width and other quantities which are then compared to experiment.

Results & Impact

Members of our three groups have played a leading role in this transformation through the development of algorithms and community codes, and by carrying out petascale simulations. One particularly striking result is the determination of the leptonic decay constant of the B meson to a precision of 0.4%, a factor-of-five reduction of uncertainty from the world average.

This calculation is illustrated in the figure above.

Why Blue Waters

The advent of petascale computers such as Blue Waters has had a transformational impact on our field. Generating gauge configurations is the rate-limiting step and requires the most capable supercomputers available. Blue Waters has allowed these studies to get much closer to the continuum limit than has been possible before. Blue Waters has enabled these simulations to be carried out with realistic(ally small) quark masses for the very first time.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/9,200 Knh
PI: Burkhard Militzer
University of California, Berkeley
Earth and Planetary Science

FIRST-PRINCIPLES COMPUTER SIMULATIONS OF HYDROCARBONS UNDER FUSION CONDITIONS

Research Challenge

The development of a first-principles methodology for warm dense matter (WDM) applications that treat temperature effects consistently is a key component in plasma science. The progress in high-energy-density physics (HEDP) applications relies on simulations for input and guidance. Hydrocarbons are the primary materials used for the ablator in inertial confinement fusion target capsules. The determination of the correct equation of state (EOS) of hydrocarbon ablators is important to optimize experimental designs to achieve desired density and temperature conditions.

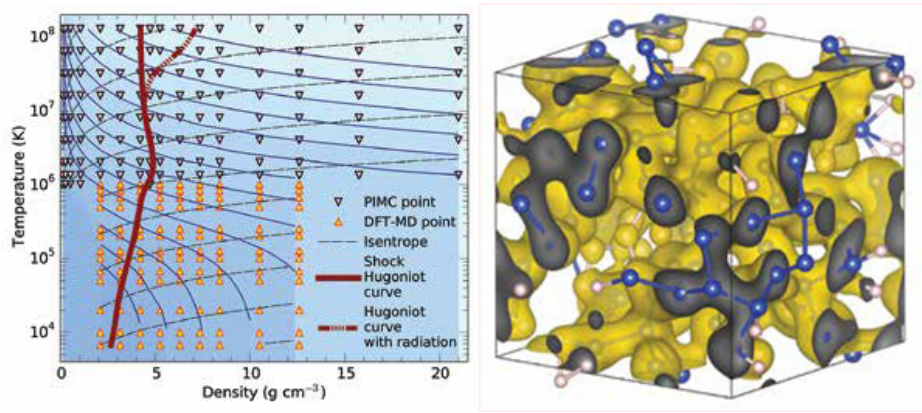

Temperature–density conditions of the path integral Monte Carlo and density functional molecular dynamics simulations of CH plastic materials (left). CH plastic at extreme pressure–temperature conditions studied with atomistic simulations (right).

Methods & Codes

For the low-temperature part of the WDM regime, density functional molecular dynamics is an accurate and efficient first-principles simulation method for these conditions. To study the high-temperature regime, we focus on the development of the path integral Monte Carlo (PIMC) method, which naturally incorporates finite temperature quantum effects by working within the many-body thermal density matrix formalism. All PIMC simulations were performed with the team's own code, CUPID.

Results & Impact

Performed an entirely first-principles determination of hydrocarbon mixtures in the WDM regime by including all nonideal effects. Based on PIMC and DFT-MD, obtained coherent sets of EOS over a wide range of density and temperature conditions and derived the shock Hugoniot curves of a series of hydrocarbon materials. While, in the past, the characterization of one material has taken a Ph.D. thesis, on HPC systems like Blue Waters many materials can now be characterized in parallel within a single year.

Why Blue Waters

The Blue Waters allocation is one order of magnitude larger than any other allocation that the team have obtained elsewhere. This enabled them to perform simulations with more accurate fermion nodes that are more realistic but also much more expensive. Besides the hydrocarbons, they performed simulations for a variety of other materials, including sodium, lithium fluoride, and silicon. Also the team could meet the composition challenge rigorously. In addition to a typical 2D parameter, they performed 3D parameter scans in density–temperature– composition space by performing simulations of C, C_2H, CH, CH_2, CH_3, CH_4 as well as H.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/3,870 Knh

PI: Sarma L. Rani
University of Alabama in Huntsville
Physics and Engineering

DIRECT NUMERICAL SIMULATIONS OF THE RELATIVE MOTION OF HIGH-INERTIA PARTICLES IN ISOTROPIC TURBULENCE

Research Challenge

To gain a deeper understanding of the role of turbulence in driving the relative velocities and positions of high-inertia particles in isotropic turbulence.

The viscous relaxation times of these particles are significantly large, with estimated Stokes number based on the Kolmogorov length scale of approximately 10-100, requiring significant computational resources.

Methods & Codes

Both direct numerical simulations (DNS) and pseudo-spectral method codes, relying on Fourier Transform techniques, are used.

The P3DFFT library is used to carry out the transforms between physical and spectral spaces.

Langevin simulations based on a probability density function kinetic model were used for pair relative motion.

Results & Impacts

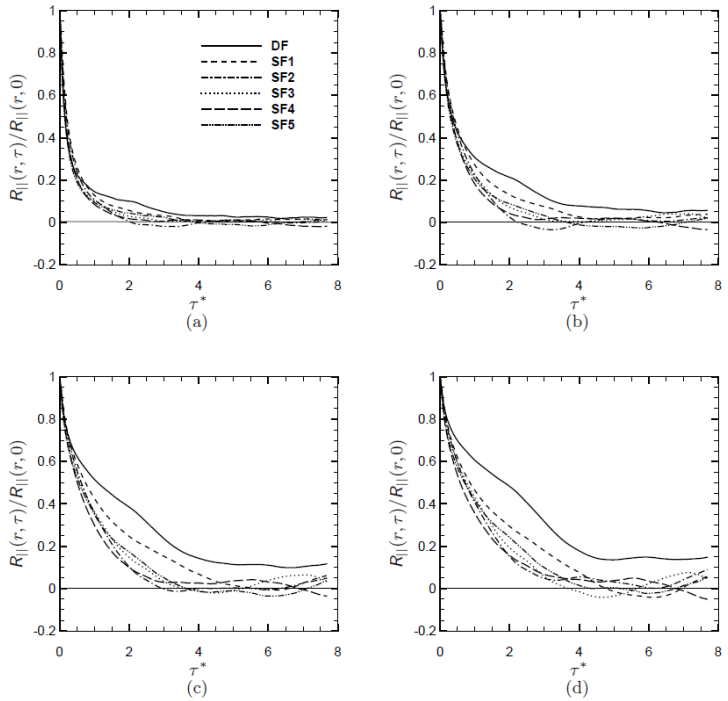

Longitudinal correlation $R_{||}(r,\tau) = \langle \Delta \boldsymbol{u}(\boldsymbol{r},\boldsymbol{x},0)\, \Delta \boldsymbol{u}(\boldsymbol{r},\boldsymbol{x},\tau) \rangle_{||}$ as a function of dimensionless time separation $\tau^* = \tau u_{rms}/L$ for Re_λ = 210 at four dimensionless separations (a) *r/L* = 0.56, (b) *r/L* = 1.12, (c) *r/L* = 2.24, and (d) *r/L* = 3.36. DF stands for deterministic forcing, and SF for stochastic forcing.

The processes being investigated here are an important part of several aspects of astrophysics, including:

- Understanding the interstellar medium
- Protoplanetary disks
- The atmospheres of planets and dwarf stars

Why Blue Waters

DNS of turbulence places a high demand on computational resources, due to the wide range of length and temporal scales involved.

The proposed simulations will run on tens of thousands of cores, and will potentially generate several terabytes of data

Due to these central processing unit time and storage requirements, the Blue Waters supercomputer is the ideal platform to achieve these objectives

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/9.44 Mnh
PI: Caroline Riedl
University of Illinois at Urbana-Champaign
Physics & Engineering

MAPPING PROTON QUARK STRUCTURE USING PETABYTES OF COMPASS DATA

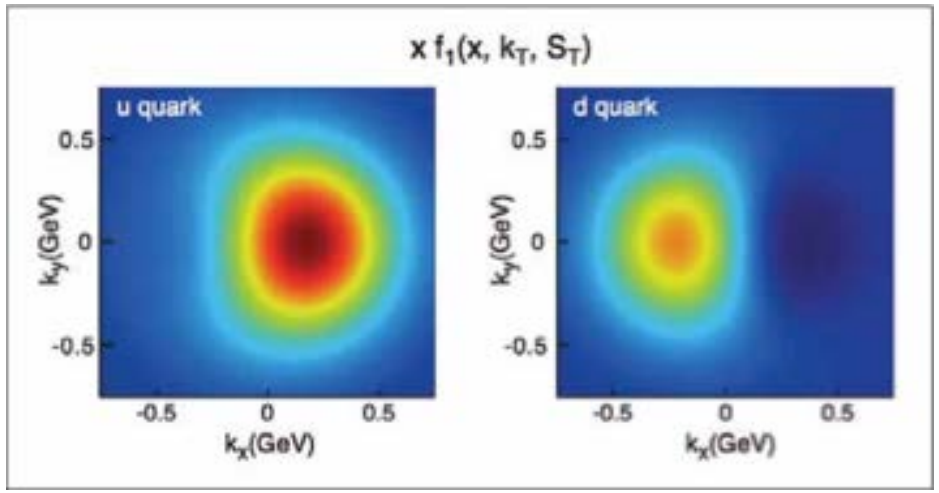

Quark densities in the transverse momentum plane for a transversely polarized proton (y-direction). Deep red (blue) regions indicate large negative (positive) values.

Research Challenge

The COMPASS particle-physics experiment at CERN probes proton substructure by scattering high-energy pion and muon-beams off of nuclear targets to measure the momentum and coordinate phase space of quarks inside the proton. Observing correlations between proton spin and intrinsic transverse momentum of quarks will shed light on quark dynamics inside the proton and will provide a critical test of fundamental predictions derived from quantum chromodynamics, the quantum field theory describing the strong nuclear force.

Methods & Codes

The COMPASS Reconstruction Analysis Library (CORAL) software performs the transition from raw data to physical quantities. The reconstructed information is stored in the form of Data Summary Trees, which are read and analyzed using the COMPASS PHysics Analysis Software Tools (PHAST). The production of Monte Carlo data is performed in two steps: (1) The generation of signal and background events is carried out with event-generator packages. (2) a GEANT4 toolkit is used based on the description of the COMPASS apparatus.

Results & Impact

Simulations of the detectors play a central role in understanding subtle detector effects and removing background events from the data sample.

This Blue Waters project involves students and young postdocs and it will in the future attract more young physicists. It thus offers outstanding educational potential for a significant number of students and postdocs towards building a community capable of using petascale computing.

Why Blue Waters

The measurements will produce 10 petabytes of experimental and simulated data. Blue Waters' balance of processing capabilities with data storage and handling is well suited for the analysis of the large COMPASS data samples as these require significant algorithmic processing per pion/muon- proton-scattering event. In addition to raw data processing and physics-level analysis, Blue Waters allows for detailed simulation of COMPASS detector properties and environmental effects.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: BW Professor/30 Knh; Illinois/550 Knh
PI: Andre Schleife
University of Illinois at Urbana-Champaign
Physics & Engineering

OPTICAL DETERMINATION OF CRYSTAL PHASE IN SEMICONDUCTOR NANOCRYSTALS

Research Challenge

The fundamental connection between crystal structure and optical properties is not well understood, which hampers the use of nanocrystals in light-emitting diodes, solar cells, bioimaging, and consumer electronics. This research will establish a connection among the electronic and optical properties of semiconductor nanocrystals. This will allow tuning the size, shape, composition, and internal structure and accelerate the development of these materials.

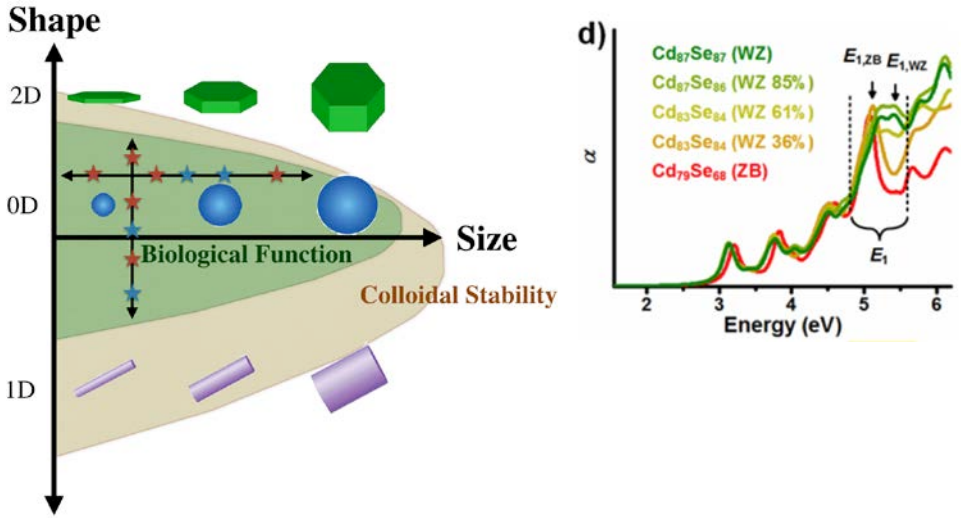

(Left) Shape and structure design space of semiconductor nanocrystals; (Right) Influence of crystal structure on optical properties.

Methods & Codes

Cadmium selenide is an inorganic compound with the formula CdSe. The team uses a combination of cutting-edge experimentation and theoretical spectroscopy to clarify important structure-property relationships for CdSe nanocrystals in two different crystal structures - wurtzite and zincblende. Simulations are accomplished with the Vienna Ab Initio Simulation Package (VASP: https://www.vasp.at) and solving the Bethe-Salpeter equation (https://en.wikipedia.org/wiki/Bethe–Salpeter_equation).

Results & Impact

CdSe-containing quantum-dot light-emitting diodes (QLEDs) are 5.4 times more efficient than previously reported similar devices.

It is possible to distinguish between the two crystal structures using optical absorption spectra and the Blue Waters calculations explain the band structure origin of the spectral features underlying this effect. The results the way for applying accelerated experimentation to nanocrystals in solution.

Why Blue Waters

The analyses are computationally extensive with large matrices that require well parallelized codes on fast supercomputers with fast inter-connects. These calculations are not possible on high throughput systems. The calculations utilize fast CPUs, large amounts of memory per compute core, and fast interconnects. Further, the Blue Waters SEAS staff helped identify and fix performance bottlenecks associated with writing and reading large wave function files.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: Exploratory/50 Knh
PI: Ahmed Taha
University of Illinois at Urbana–Champaign
Physics & Engineering

ADVANCED DIGITAL TECHNOLOGY FOR MATERIALS AND MANUFACTURING

****The allocation covered multiple graduate research projects and a summary for the activities of two of these projects is presented here and notes as P-1 & P-2*

Research Challenge

*P-1: Prediction of age- and disease-related fractures in oral bones and to design of improved dental implant systems

*P-2: Using traditional materials testing approaches to measure the mechanical properties of bone, such as compression, tension, and three-point and four-point bending, are ex vivo and destructive

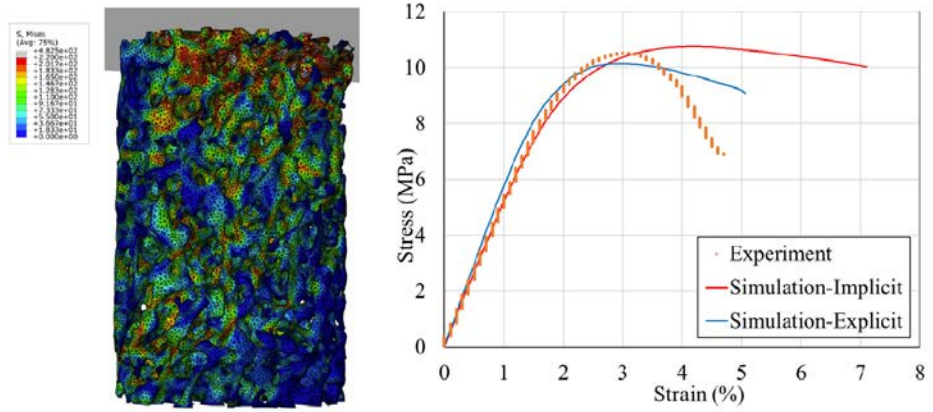

(a) Example of von Mises stress distribution in samples under uniaxial compression, (b) implicit and explicit solver results versus experimental stress–strain curves.

Methods & Codes

*P-1: Use the implicit and explicit solvers of the commercial Finite Element Analysis code “Abaqus” to analyze Micro-computed tomography finite element models of trabecular bone and compared the performance of the two solvers.

*P-2: Use Abaqus to analyze the bone structure, bone strength and fracture resistance using two different experimental methods, BioDent and Osteoprobe RPI, and comparing their efficiency and effectiveness.

Results & Impact

*P-1: There is a good match between micro computed tomography finite element model results using implicit and explicit solvers as show in the above figure

*P-2: The relationship of BioDent RPI instrument to the bone material properties has been established using the finite element method. Also simulating bone fracture using the extended finite element method on a single-osteon cortical bone sample has been evaluated and published. However, completing a study to relate Osteoprobe RPI output to bone material properties and its fracture resistance is currently being developed.

Why Blue Waters

*P-1: Each model has many millions degrees of freedom and nonlinearities making such 3D models impossible to solve without the use of Blue Waters capacity.

*P-2: The problem is highly nonlinear, and multiple iterations are needed to be able to relate Osteoprobe device output to different bone mechanical properties. Each of the Osteoprobe RPI simulation iteration demands high computational power and time.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: Exploratory/50 Knh
PI: Ange-Therese Akono
University of Illinois at Urbana-Champaign
Physics & Engineering

MULTI-SCALE AND MULTI-PHYSICS MODELING OF THE STRENGTH OF GEOPOLYMER COMPOSITES

Research Challenge

To employ finite element analysis and molecular dynamics at extreme scales to investigate the processing-microstructure-properties relationships in inorganic polysialates or geopolymer cements from the nanometer length scale up to the macroscopic length-scale.

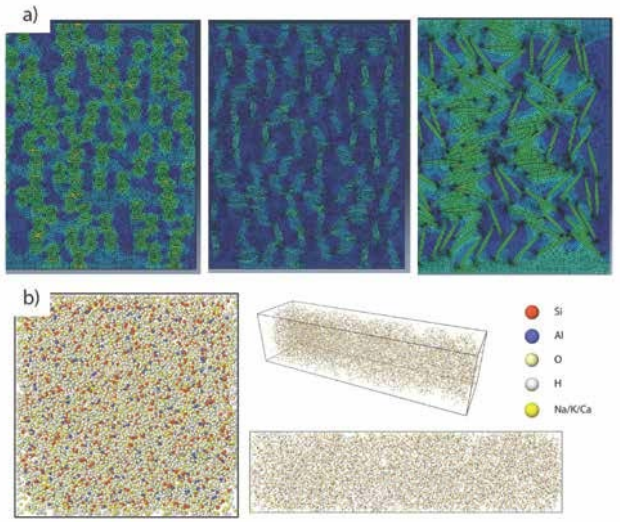

a) FEM model at mesoscale showing the stress distribution as a function of the filler aspect ratio. b) Molecular dynamics modelling of inorganic polysialate disiloxo.

Methods & Codes

- At the nanoscale, atomic simulations were performed in LAMMPS.
- Various molecular structures were recreated using the Avogadro and Packmol software.
- The software OVITO was utilized for visualization.
- At the mesoscale, the finite element package Abaqus was used.
- The microstructure was generated from OOF2D and MATLAB.

Results & Impact

Our theoretical and computational framework was validated on 31 different geopolymer-based systems, based on experiments carried out by various researchers over seven years. To our knowledge, it is the first time that a theoretical model has been proposed to upscale the constitutive behavior of geopolymers.

Why Blue Waters

For molecular dynamics, it was crucial to be able to simulate large systems in a timely fashion.

In the case of finite element simulations, the nonlinearity of the equations due to friction, contact, and plastic flow made it impossible to obtain results using a desktop workstation.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: Illinois/100 Knh
PI: Brian G. Thomas
Co-PIs include: SM Cho, SP Vanka
University of Illinois at Urbana-Champaign
Physics & Engineering

TRANSIENT MULTIPHASE FLOW PHENOMENA AND DEFECT FORMATION IN STEEL CONTINUOUS CASTING

Research Challenge

- Transient multiphase flow phenomena should be better understood in order to reduce the problems that introduce defects during continuous casting.
- The modeling results should be validated with plant measurements and applied to find optimal process conditions, including nozzle port angle, nozzle submergence depth and EMBr field strength.

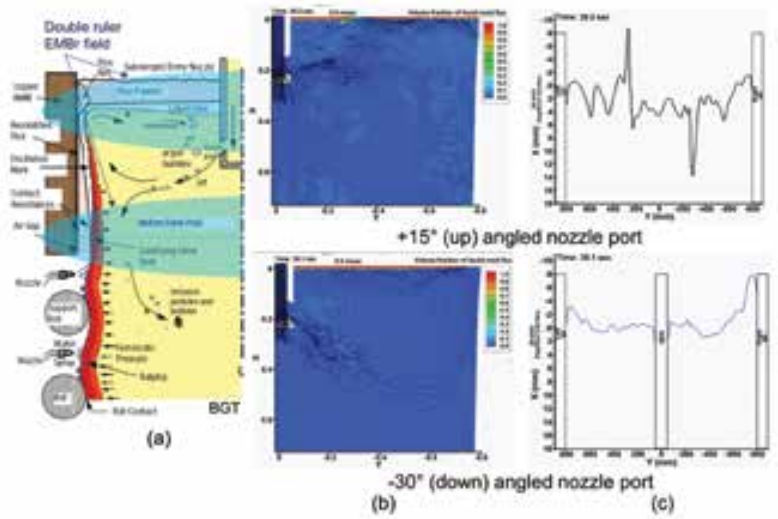

(a) Schematic of multiphase flow phenomena in the steel continuous-casting mold and Effect of nozzle port angle on (b) mold flow patterns, including motion of the top liquid mold flux layer, and (c) level variations at the liquid mold flux/molten steel interface in the meniscus region.

Methods & Codes

- ANSYS Fluent is used to implement Large-Eddy Simulations (LES) coupled with Volume of Fluid (VOF) models.
- Use of GPU-based in-house code, CUFLOW, to create LES coupled with Lagrangian particle capture and MagnetoHydroDynamics models.

Results & Impact

- These studies have enabled better understanding of the complex multiphysics phenomena related to defect formation.
- This better understanding has led to suggestions of nozzle geometry/casting condition combinations that lead to fewer defects, and, consequently, to significant savings to the steel plants.

Acknowledgements

- Continuous Casting Consortium and NSF (Grant CMMI 15-63553) support

Why Blue Waters

- Blue Waters enabled high-resolution multiphase flow simulations of the continuous caster needed for accurate predictions.
- Blue Waters resources (both ANSYS Fluent HPC on XE nodes and our in-house multi-GPU code CUFLOW on XK nodes) showed speed-up breakthroughs (e.g., over 3000x with ANSYS Fluent HPC on BW) needed to provide this modeling capability for the steel continuous casting process.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: Illinois/200 Knh
PI: Rafael Tinoco Lopez
University of Illinois at Urbana–Champaign
Physics & Engineering

HIGH RESOLUTION NUMERICAL SIMULATION OF OSCILLATORY FLOW AND SEDIMENT TRANSPORT THROUGH AQUATIC VEGETATION

Research Challenge

Aquatic vegetation provides a wide range of services to the ecosystem: improving water quality through nutrient uptake and oxygen production, providing flood buffering and coastal protection services, and regulating erosion and deposition patterns, thus playing a paramount role in habitat creation and promotion of biodiversity. While vegetation–flow interactions have been studied for unidirectional flows, less is known about oscillatory conditions. The study is geared toward increasing our understanding of the interactions among vegetation, flow, and sediment under oscillatory flows.

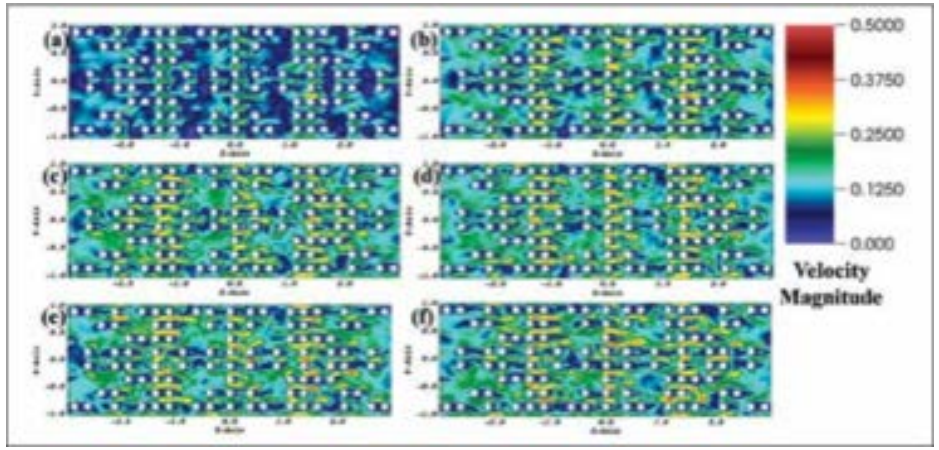

Results from 3D simulation of turbulent flow through a random arrangement of cylinders. About 296 million computational points are used. Instantaneous velocity magnitude at different elevations. Planes at (a) 0.5 %, (b) 1 %, (c) 10 %, (d) 50 %, (e) 75 %, and (f) 95 % water depth.

Methods & Codes

High-resolution Direct Numerical Simulations (DNS) and Large-Eddy Simulations (LES) of the flow at different configurations of the idealized vegetation were conducted using the open-source, spectral element-based higher-order incompressible Navier–Stokes solver *Nek5000.* The Spectral Element Method combines the accuracy of spectral methods and the flexibility of Finite Elements Method.

Results & Impact

A quarter of the full domain was simulated of turbulent flow through a random arrangement of cylinders, with ~296 million computational points. Fields of instantaneous velocity magnitude at different elevations show that the length of the vortex being shed increases with distance from the bottom of the cylinder. This simulation, one of the largest high-resolution eddy-resolved hydrodynamic simulations in this field, will provide as yet unseen details of the physical processes involved.

Why Blue Waters

The study pushes the limit of the scale at which high-resolution simulations are used to study complex multi-phase flow in environmental fluid mechanics . Simulations have been conducted for up to 296 million computational points, scaling to 32,768 MPI ranks. Without access to Blue Waters, completing the study within a realistic timeframe would be impossible. We will work with Blue Waters staff to create animations of the phenomenon using data from the simulations.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: Illinois/202.8 Knh
PI: Lucas K. Wagner
University of Illinois at Urbana-Champaign
Physics & Engineering

QUANTUM MONTE CARLO SIMULATIONS OF MAGNETISM AND MODELS IN CONDENSED MATTER

Research Challenge

There are special materials in which quantum mechanics are noticeable even at the scale of the entire material. Materials such as unconventional high-temperature superconductors, magnetic materials, topological materials, and many others have quantum effects that are very challenging to describe under the normal band structure framework. Current tools also struggle to accurately predict the properties of simple materials.

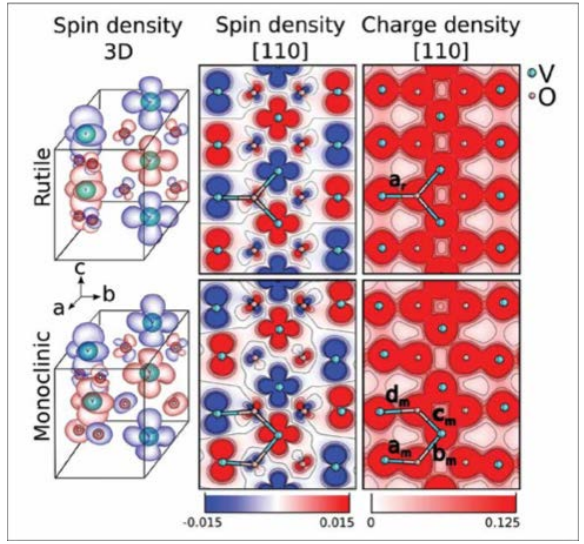

Understanding of a transition of vanadium dioxide from metal to insulator from quantum Monte Carlo (QMC). This material can switch very quickly between these two states, which could enable optical devices and control transparent windows. The QMC calculations were the first to successfully describe this transition from the atoms up, without any empirical parameters.

Methods & Codes

The Qwalk code is developed at University of Illinois. Monte Carlo calculations are used to sample electron position. Correlation effects can be found, based on the magnitude of the sample positions.

Results & Impact

Work performed on BW on vanadium dioxide, which transitions from a metal to an insulator at around 340 °K (150 °F). For many years, it has been a mystery why this happens, mainly because electron correlation is important in the effect, which is difficult to access with inaccurate quantum calculations. This new predictive capability opens the door to computer-based design of these extraordinarily sensitive materials that can be used in new devices.

Why Blue Waters

Solving the Schrödinger equation for complex materials has been out of reach because of the computational cost.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/9,547.165 Knh
PI: P.K. Yeung
Georgia Institute of Technology
Physics & Engineering

INTERMITTENCY, RESOLUTION EFFECTS AND HIGH SCHMIDT NUMBER MIXING IN TURBULENCE

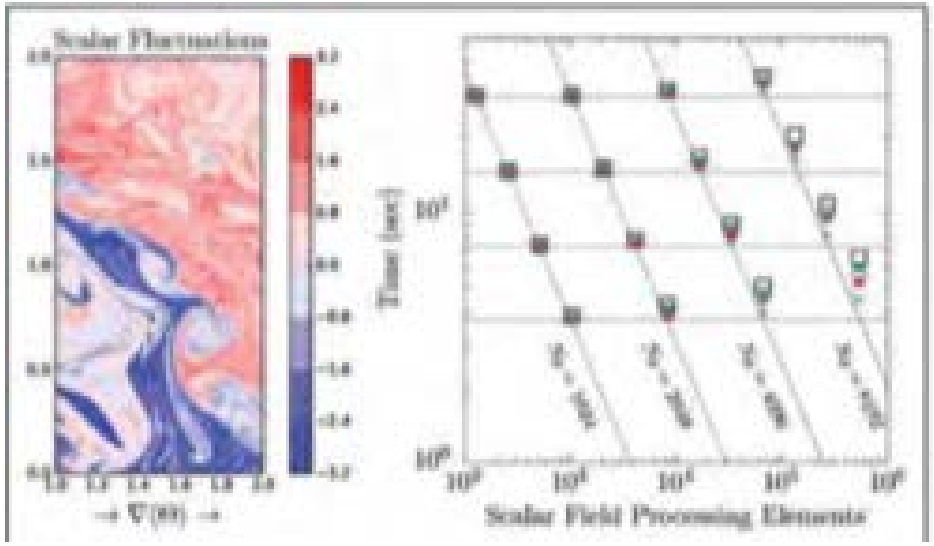

Left—Isocontours of normalized scalar fluctuations in a local region within a two-dimensional plane. Right—Scalability of scalar field computation for different grid resolutions using different versions of the CCD routines, from the baseline single-threaded nonoverlapped version (open squares) to the best version with dedicated communication threads (stars).

Research Challenge

- To use direct numerical simulation (DNS) to understand the nature of intermittency in terms of fluctuations of energy dissipation rate and enstrophy for turbulent flows with disorderly fluctuations, and to critically examine the accuracy and reliability of the simulation results.
- Develop and apply a new parallel algorithm uniquely suited to the study of turbulent mixing at high Schmidt number, where low molecular diffusivity leads to fluctuations at scales smaller than those in the velocity field

Methods & Codes

- For direct numerical simulation (DNS), the code PSDNS uses Fourier pseudo-spectral methods for the velocity field, while performing local averaging both along a line (one dimension) and over a cube (three dimensions).
- For high Schmidt number turbulence mixing, the new algorithm uses the pseudo-spectral method for the velocity field on a coarser grid and combined compact finite differences (CCD) on a finer grid for resolution

Results & Impacts

- Using DNS and three dimensional averaging are helpful in allowing intermittency theories to be evaluated more definitively than in the past.
- Combining the pseudo-spectral methods with CCD, and performing careful optimization, have made the code highly scalable, and enabled studies of turbulent mixing in previously unreachable Schmidt number regimes.

Why Blue Waters

The $8{,}192^3$ grid resolution requires access to a world-class machine such as Blue Waters, and the memory capacity of Blue Waters has made the $16{,}384^3$ resolution, the highest known, possible. Other important machine characteristics include support for Co-Array Fortran, topologically aware scheduling, and a large number of cores per node.

Computer Science and Engineering

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: Illinois/300 Knh
PI: Donna Cox
National Center for Supercomputing Applications
Computer Science & Engineering

CADENS NSF PROJECT: DIGITAL LITERACY, DATA VISUALIZATION, AND THE CINEMATIC PRESENTATION OF SCIENCE

A visualization of the universe 400 million years after the Big Bang.

Research Challenge

- Visualize supercomputer datasets in a cinematic way while maintaining scientific accuracy
- Few existing tools handle both cinematic animation and scientific visualization

Methods & Codes

- yt, a scientific analysis and visualization package for Python
- Houdini, a visual effects tool
- Created middleware and a pipeline to connect yt and Houdini named Ytini.
- Created a Python tool named Blurend to prepare Houdini scene files for rendering on Blue Waters.

Results & Impact

- Rendered ”First Light in the Renaissance Simulations” in 4k stereoscopic and planetarium dome formats.
- Rendered “El Reno 2011 Tornado” in 4k planetarium dome format.
- Processed data on Blue Waters for “Cosmic Bubble Bath” for the documentary “Seeing the Beginning of Time” (available on Amazon Prime)
- Processed data on Blue Waters for “Formation of the Moon” for a planetarium dome show.

Why Blue Waters

- Blue Waters reduced data preparation time from 4 days to 4 hours, speeding development of visualizations significantly.
- Blurend software, created for Blue Waters, facilitated rendering of “Seeing the Beginning of Time” and a tornado simulation in 4k planetarium dome format.
- Collaborators already used Blue Waters which made data transfer unnecessary, making things easier and saving time for scientists and visualization team.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: BW Professor/80 Knh
PI: William Gropp
University of Illinois at Urbana–Champaign
Computer Science & Engineering

ALGORITHMS FOR EXTREME-SCALE SYSTEMS

Research Challenge

At an extreme scale, even small inefficiencies can cascade to limit the overall efficiency of an application. In order to achieve the performance required by applications that apply highly challenging scientific computations, new algorithms and programming approaches are needed. Developing more scalable versions of methods, such as Krylov methods, used in many science applications would immediately benefit performance. Also important is the development of techniques to improve communication between nodes by improving the methods used to exploit internode and intranode communication.

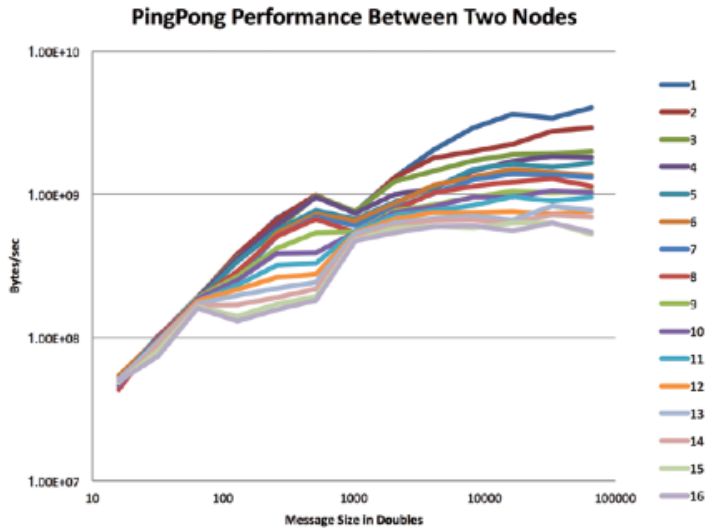

Communication performance between two nodes with 1–16 processes communicating at the same time. Note that as the number of communicating processes increases, the achieved performance per process drops, with nearly an order of magnitude loss between one and 16 communicating processes at large messages.

Methods & Codes

To address the challenges of parallelism and scale, the team developed several codes to benchmark these operations, to gather detailed timing results, and to perform experiments with different approaches. Additionally, researchers developed benchmark codes that better measure the achievable performance of the communication patterns commonly used in applications.

Results & Impacts

Early results using alternative Krylov formulations showed several performance effects that can provide a 200% or more improvement in performance at scale. Initial experiments using smaller core counts, in preparation for a later study using larger runs, will study the impact of network noise on highly parallel computers, and how parallel numeric algorithms can be developed that will perform well despite noise.

Why Blue Waters

Blue Waters provides one of the few environments available for large-scale experiments. Additionally, it provides a highly capable I/O system, which we plan to use in developing improved approaches to extreme scale I/O.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: Illinois/210 Knh
PI: Rakesh Nagi
University of Illinois at Urbana-Champaign
Computer Science & Engineering

PARALLEL ALGORITHMS FOR SOLVING LARGE ASSIGNMENT PROBLEMS

Research Challenge

- The Quadratic Assignment Problem (QAP) was introduced as a mathematical model to locate indivisible economical activities (such as facilities) on a set of locations so as to minimize a quadratic cost function. QAP is a NP-Hard problem.
- A fast and scalable QAP solver can be a powerful tool for researchers working on such NP-hard problems
- A sequential QAP solver can become computationally intensive and therefore, the algorithm can benefit from parallelization on an appropriate parallel architecture, such as Blue Waters
- Multi-Dimensional Assignment (MAP) and Linear Assignment Problem (LAP) can also benefit from GPU acceleration

Methods & Codes

- Parallelized the Lagrangian dual ascent algorithm then had to solve $O(n^4)$ Linear Assignment Problems (LAPs) and adjust $O(n^6)$ Lagrange multipliers to obtain a strong lower bound on the QAP.
- The LAPs are split across these GPUs and solved using a GPU-accelerated Hungarian algorithm , while the $O(n^6)$ Lagrange multipliers are updated by multiple CUDA threads in parallel.

Results & Impact

Able to obtain strong lower bounds on problem instances and optimally solve quadratic assignment problems using only a modest number of GPUs. With some additional work in memory management and CPU + GPU collaboration, some proposed algorithms may be used effectively to solve truly large QAPs with over 30 facilities, which has been impossible so far.

Why Blue Waters

Needed a large number of processors that can explore the solution space in parallel. Additionally, the GPU-accelerated dual ascent procedure benefits from the large number of powerful GPU-enabled processors available at the Blue Waters facility. Also, Blue waters provided a significant cost reduction for this resource intensive task compared to other solution such as AWS.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: Illinois/50 Knh
PI: Luke Olson
University of Illinois at Urbana-Champaign
Computer Science & Engineering

LOCALIZING COMMUNICATION IN SPARSE MATRIX OPERATIONS

Research Challenge

High communication costs associated with sparse matrix operations limit performance for a range of applications on Blue Waters. The key challenge of this project is to limit costly, internode communication through localization.

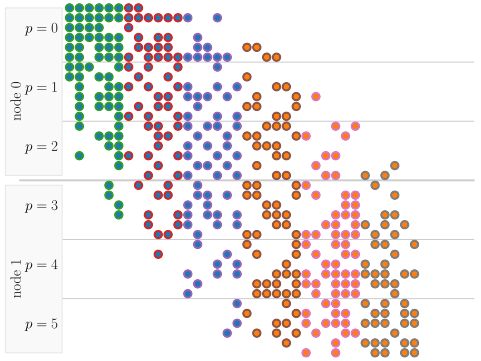

Distribution of data among nodes and processors.

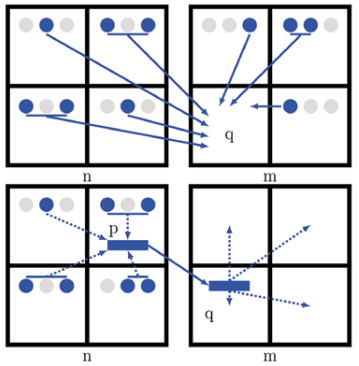

Communication patterns in the standard and topology-aware algorithms.

Methods & Codes

- Method developed reroutes communication inside of sparse matrix operations so that MPI messages are aggregated on a node before executing internode communication on the network
- Method considers the topology for communication.

Results & Impact

- Significant reduction in communication costs in sparse matrix-vector and matrix-matrix operations
- Reduced time-to-solution for simulations that rely on sparse linear systems
- Scalable and efficient multigrid performance

Why Blue Waters

Parallel sparse iterative methods, such as algebraic multigrid (AMG), solve a variety of linear systems in virtually every field of science and engineering. Supercomputers such as Blue Waters provide sufficient memory and bandwidth to solve extremely large systems, enabling the simulation of more complex problems. Blue Waters is an ideal platform for testing scalable algorithms for future machines..

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: Illinois/50 Knh
PI: Edgar Solomonik
University of Illinois at Urbana-Champaign
Computer Science & Engineering

PERFORMANCE EVALUATION OF NEW ALGEBRAIC ALGORITHMS AND LIBRARIES

Research Challenge

Design software and algorithms for applications running on Blue Waters addressing one of the biggest challenges facing parallel scalability of methods in computational science - the overhead of moving data between processors.

- Goal: Develop algorithms and libraries that minimize communication in the number of messages as well as in the amount of data moved.

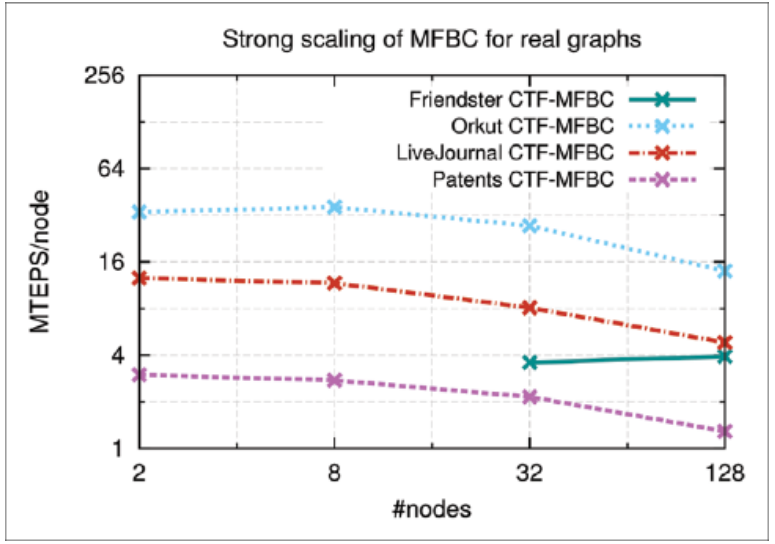

Parallel scalability of betweenness centrality using sparse matrix multiplication for massive graphs.

Methods & Codes

- Various libraries available to application developers on BW.
- Focus Library: Cylops Tensor Framework (CTF) – provides distributed-memory support for sparse and dense tensors. Uses performance models to make runtime mapping decisions using autotuning to train the model parameters.

Results & Impact

- Impacts multiple domains from quantum chemistry to graph analysis
- Scalability of key kernels in quantum chemistry applications such as atomic-to-molecular orbital transformations
- Development of a new, scalable QR factorization

Why Blue Waters

Blue Waters is ***essential for testing and evaluation***. While all of the codes developed are designed to be portable, demonstrating performance on Blue Waters helps foster local collaborations and deployment of parallel numerical library software.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: Illinois/125 Knh
PI: Tandy Warnow
University of Illinois at Urbana-Champaign
Computer Science & Engineering

PARALLEL ALGORITHMS FOR BIG DATA PHYLOGENOMICS, PROTEOMICS, AND METAGENOMICS

Research Challenge

This project aims to develop a better method for analyzing large size genomic data sets, where standard methods either do not run or provide poor accuracy. Data sets are for genome-scale phylogeny estimation, multiple sequence alignments and metagenomics analysis.

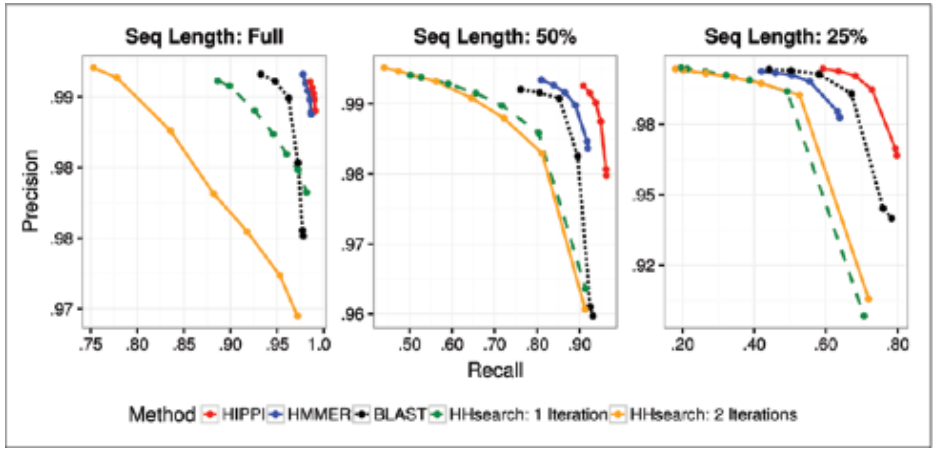

Precision-recall curves for methods for protein family classification, evaluated on one cross-fold subset of the PFAM seed sequence data. Our new method— HIPPI (shown in red)—strictly dominates all the other methods. (Figure reproduced from *BMC Genomics* 17(Suppl 10); 765 (2016), under the Creative Commons Attribution License.)

Methods & Codes

- Applied the method to small subsets of a large data set followed by combining results from the small data sets in parallel.
- Developed a novel machine-learning technique to detect membership in existing protein families.

Results & Impact

- The main outcome of this project is HIPPI (Hierarchical Profile Hidden Markov Models for Protein family Identification), a analysis method for protein family classification.
- This study showed that our HIPPI method outperformed all other existing sequence analysis methods in terms of both precision and recall, especially when analyzing short sequences (see figure).

Why Blue Waters

- This development requires extensive testing, which is not feasible on other platforms.
- The analysis of large biological data sets (and even of moderate-sized data sets) often requires years of CPU time and Blue Waters makes this feasible and enables biological discovery.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

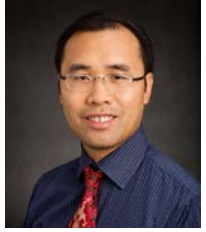

Allocation: Illinois/50 Knh
PI: Tao Xie
University of Illinois at Urbana-Champaign
Computer Science & Engineering

HARDWARE ACCELERATION OF DEEP LEARNING

Research Challenge

Implicit data synchronization across different nodes severely limits the training process. This research pursues a data manager that overlaps data transfer with computation. The team evaluated the performance of three deep neural networks on the Blue Waters GPU nodes. The scalability differences of the neural networks relates to the computation per byte in each layer. We explored the design space of mapping different kinds of neural networks to the GPUs.

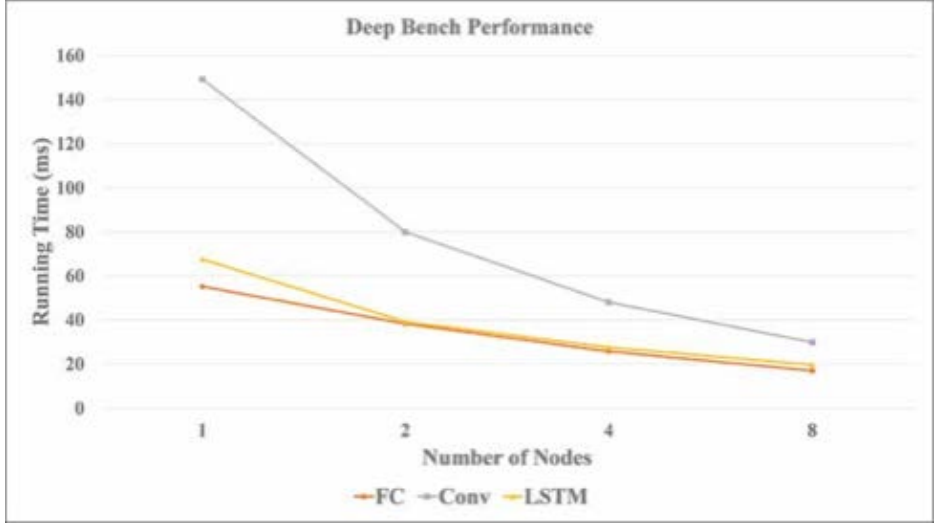

Performance Evaluation of Different Layer Types on Different Numbers of Nodes.

Methods & Codes

To evaluate the performance of different types of neural networks, the research team used the AlexNet neural network for reference of the convolutional layer and fully connected layer topologies. They implemented the neural network layers based on DeepBench, which is a performance benchmark for deep learning hardware accelerators. They modified DeepBench to accommodate the Blue Waters GPU enabled nodes.

Results & Impact

By explicitly overlapping data transfer overhead with computation the team achieved a speedup of 1.6X over the implicit data transfer implementation. We can not achieve the best performance or system efficiency by using one single resource allocation scheme for different types of neural networks. The team will design a new resource allocation and algorithm mapping technique to achieve better system performance.

Why Blue Waters

Blue Waters provides support for research on the optimization of deep learning on computational clusters using GPUs. GPUs are more suitable than CPUs for convolution and matrix multiplications, which are the major computations in deep learning. State-of-the-art deep learning facilities widely employ GPUs as their hardware accelerators.

Biology, Chemistry, and Health

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: Illinois/986 Knh
PI: Aleksei Aksimentiev
University of Illinois at Urbana-Champaign
Biology, Chemistry & Health

MOLECULAR MECHANISM OF NUCLEAR TRANSPORT

Research Challenge

The nuclear pore complex (NPC) regulates the transport of all ribonucleic acid (RNA) and proteins across the nuclear envelope of eukaryotic cells, in which the well-defined chromosomes (bodies containing the hereditary material) are located. Simulations will be conducted to gain new insight into the physical mechanism of nuclear transport, with important implications for several human diseases and the development of novel gene therapies.

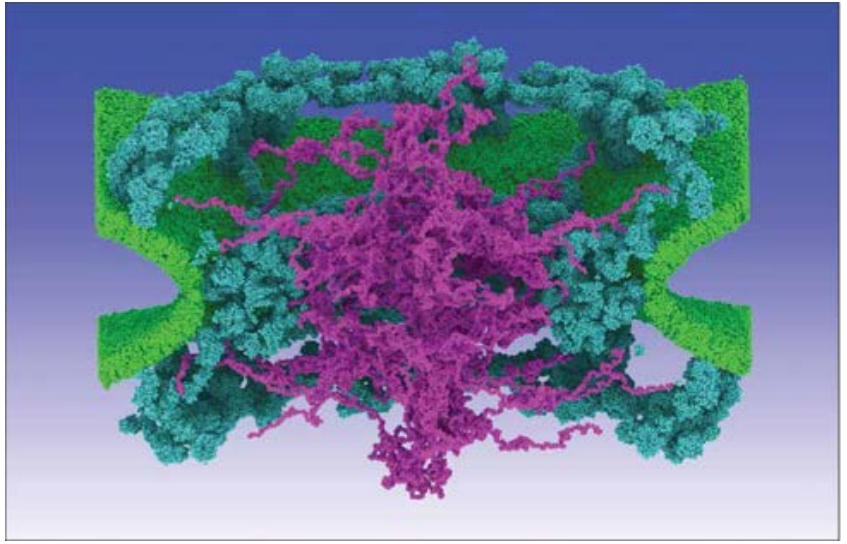

An atomic model of the entire nuclear pore complex, 140 M atoms in total (water omitted for clarity). Colors highlight the outer scaffolding (in cyan), the disordered central channel of nups (in purple), and the nuclear envelope (in green).

Methods & Codes

We performed explicit-solvent all-atom MD simulations with the latest version of NAMD2 of a solution of nucleoporin (nup) fragments and varied the protein volume fraction within a confined volume. Simulations with an applied electric field were performed to calculate ionic conductivity as a function of nup density. We used atomic-resolution Brownian dynamics (ARBD) and traditional all-atom MD to model the entire NPC

Results & Impact

We characterized the structural fluctuations and electrical properties of a solution of nups varying in protein density. We measured the ionic current of each system in the presence of an applied electric field and observed a transition from conducting to not-conducting ions at a critical nup density, matching a key experimental result. This project may offer new insights into the molecular origin of diseases (e.g., cancer, viral infections, and neurodegenerative diseases) with implications for developing gene therapy treatments.

Why Blue Waters

Explicit-solvent all-atom MD simulations are needed to characterize the structural fluctuations and electrical properties of the disordered central channel of the NPC. MD simulations the size of the entire NPC—140 Million atoms—are only possible with the computational power of Blue Waters. The large number of XK nodes with GPU accelerators and fast Gemini interconnect makes it one of the best publicly available systems for performing simulations of the entire NPC.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: Innovation & Exploration /500 Knh
PI: Greg Bauer
National Center for Supercomputing Applications
Biology, Chemistry & Health

IMPROVING NWCHEM SCALABILITY USING THE DATASPACES FRAMEWORK

Research Challenge

- Limited pool of experimental data and the extremely labor-intensive nature of parameter optimization represent the major limiting factors in improving the quality of classical force field parameters, in which these parameters are essential to improve the predictive ability of Molecular Dynamics (MD) simulations.
- The challenge with lack of experimental data in parameter optimizations can be addressed through reliable electronic structure computations at the CCSD(T) level, however a faster and more computationally efficient implementation is needed to optimize the high computational cost of CCSD(T) method.
- Labor-intensive nature of classical force field determination, and the manually driven optimization drives the need to introduce a quality control into the optimization process, in order to be able to reproduce the results, revise and improve the protocol.

Methods & Codes

- Performance optimization of the CCSD(T) method is being developed through offloading large memory arrays from compute nodes to a dedicated data storage employing the DataSpaces data management framework.
- A methodology based on data extracted from a previously optimized version of CCSD(T), used together with scalable tool to optimize Lennard-Jones parameters in the AMBER force field has been developed to generate significant sets of quality parameters for use in subsequent MD simulations.

Results & Impact

- A new approach to parameter optimization for classical force fields that combines a high-level electronic structure calculation method to extract additional previously inaccessible information from experimental data.
- Reengineered optimization procedure that maximizes HPC resources.
- 10x shorter time needed for parameter optimization.
- Transformation of force field optimization from an empirical to a well-structured discipline enables reproducibility of results by the community instead of limiting the accessibility to only a few highly capable teams.

Why Blue Waters

- The fast interconnect and large memory per core of Blue Waters enables CCSD(T) computations of molecular systems encountering a thousand basis functions, which is vital for the success of the developed parameter optimization procedure.
- Availability of large numbers of nodes is essential to address the extremely resource demanding needs of exhaustive exploration of parameter space.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: Illinois/350 Knh
PI: Gustavo Caetano-Anolles
University of Illinois at Urbana-Champaign
Biology, Chemistry & Health

HOW FUNCTION SHAPES DYNAMICS IN PROTEIN EVOLUTION

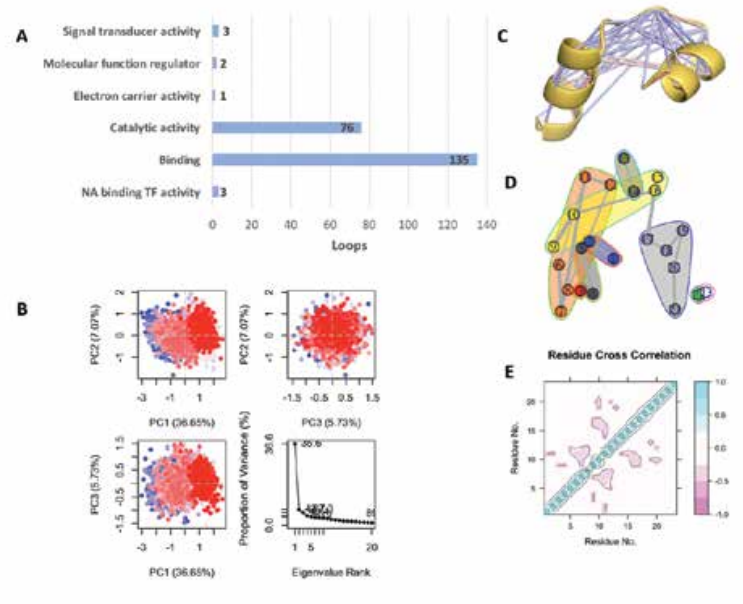

Research Challenge

The main objective of the studies is to bridge disparate disciplines of biology and physics with Molecular Dynamics (MD) simulations performed at nanosecond (ns) timescales to capture evolutionary dynamics on a scale of billions of years. The team explored biophysical variables of the MD simulations by studying community structures of protein loop residues that describe the molecular trajectories of the loop regions.

Methods & Codes

The team constructed a dynamics space, a modified version of the dynamosome, by calculating the eigenvalues of the top five principal components from principal component analysis and centrality metrics from a network based on the dynamic cross-correlation matrix of the motions of protein residues.

Results & Impact

The aim of the investigation is to detect the presence or absence of patterns of motion in molecules. The team focused on an analysis of dynamic network topologies defining a three-dimensional morphospace delimited by the conceptual axes of modularity, scale- freeness, and randomness.

Why Blue Waters

NAMD scales well on the Blue Waters architecture, especially when combined with GPU (graphics processing unit) nodes. This provides a significant boost in acceleration. Apart from system specifications that are well-suited to our project, the domain experts/scientists in the Blue Waters support team have helped smooth out any technical issues that have arisen during the current and previous allocations.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/6,510 Knh
PI: Vincenzo Carnevale
Temple University
Biology, Chemistry & Health

MECHANISM OF TEMPERATURE SENSITIVITY IN TRPV1 CHANNEL

Research Challenge

The team studies the complex molecular mechanism responsible for burning heat sensation in peripheral nerves. The project goal is to develop a molecular model for the opening of TRPV1, an ion channel playing a major role in nociception.

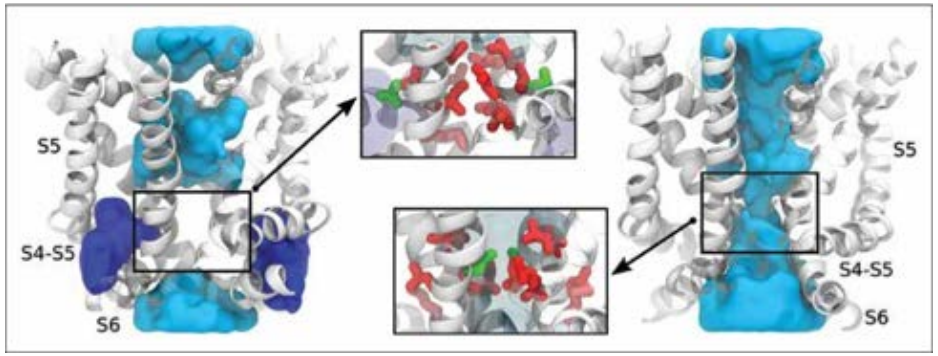

TRPV1 pore domain with hydrated (left) and empty (right) protein cavities (PCs). Water in the PCs and in the central pore is shown in blue and cyan, respectively. In the state with hydrated PCs, there are two interruptions of water density, while with empty PCs the water density is continuous. Residue N676 is shown in green.

Methods & Codes

TRPV1 is an ion channel crucially responsible for transduction of nociceptive stimuli into pain signals. To identify molecules that modulate the channel without interfering with its other useful functions, a microscopic understanding of TRPV1 is needed. This project conducts long time-scale molecular dynamics (MD) simulations using the NAMD package and CHARMM36 force field together with enhanced sampling techniques to estimate free energies.

Results & Impacts

The team found that the lower gate of the ion channel is open or closed depending on the conformation of amino acid N676. Hydration state of the adjacent protein cavity (PC) guides N676 in and out the central pore. N676 upon rotation changes the hydrophobic character of the molecular surface lining the central pore. This microscopic picture provides the basis for rational drug design.

Why Blue Waters

The massively parallel architecture of Blue Waters and large memory capacity enabled MD simulation of a system of approximately 400,000 atoms at the microsecond time scale.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/12,000 Knh
PI: Thomas E. Cheatham III
University of Utah

Biology, Chemistry & Health

EXPLORING THE STRUCTURE AND DYNAMICS OF CONVERGED ENSEMBLES OF DNA AND RNA THROUGH MOLECULAR DYNAMICS SIMULATIONS

Research Challenge

The team works on improving Molecular Dynamics (MD) simulation methods, which have been one of the most important tools in computational chemistry. The presence of very large number of degrees of freedom in the simulated system makes obtaining converged results challenging. The project includes method development and the study of molecular mechanism of DNA-drug interaction.

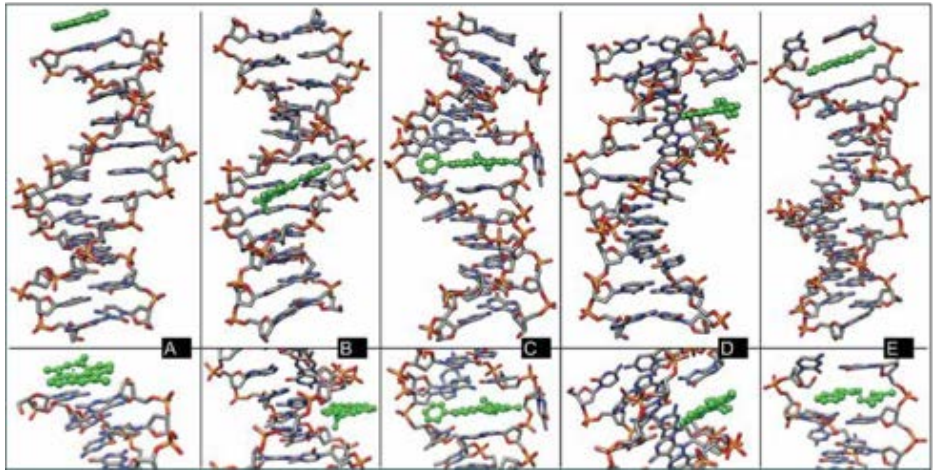

Predicted modes of DNA – drug interaction (drug in green): a) stacking at the end of the DNA; b) and d) minor groove interaction; c) base-pair eversion mechanism; and e) intercalation at the terminal base pairs.

Methods & Codes

The team uses the AMBER package and its GPU-accelerated version, and develops specialized software and methodologies to analyze the vast amount of simulation results to further the understanding of biological DNA, RNA, and protein structures.

Results & Impacts

The team has created new improved simulation techniques. Using these, the team has achieved a converged simulation of several small RNA systems that generated insight into successes and failures of simulation methods. The conducted study of the process of DNA-ligand binding revealed five principal binding mechanisms of a family of planar copper complexes representing medicinal value.

Why Blue Waters

The GPU acceleration and massive CPU parallel architecture of Blue Waters allowed the team to independently run hundreds of biomolecular simulations of myriad systems of DNA and RNA. In addition to that, the highly trained Blue Waters staff have been in close collaboration to develop custom and specialized code capable of performing the analysis of terabytes of simulation data.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/2,900 Knh
PI: Ken A. Dill
Stony Brook University
Biology, Chemistry & Health

PREDICTING PROTEIN STRUCTURES WITH PHYSICAL PETASCALE MOLECULAR SIMULATIONS

Research Challenge

The project addresses a central challenge of biology: the prediction of protein structure from sequence. Accurate determination of protein structure is essential to understanding biological mechanism and designing new drugs. The purpose of this work is to use the petascale capabilities of Blue Waters to overcome limitations of current experimental and computational tools of structure determination.

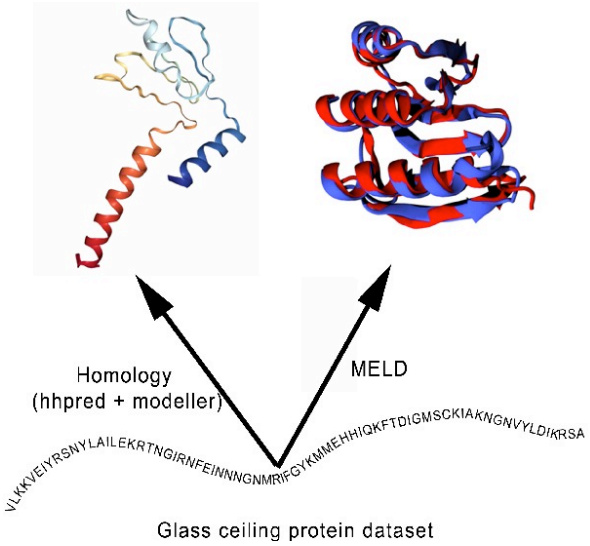

Glass ceiling protein dataset

When predicting the structure based on sequence using state-of-the-art homology methods (left) the predicted structure is wrong. Using MELD and physics, we get the right answer (right: experiment in red, our prediction in blue).

Methods & Codes

The team accelerates molecular dynamics (MD) simulations, and rationally incorporates sparse, noisy, and ambiguous experimental or heuristic information using the MELD (Modeling Employing Limited Data) tool developed by the team. The MD engine relies on the OpenMM program and AMBER package for force fields and setup.

Results & Impacts

The team has participated in CASP, the blind protein structure prediction competition. During the three months of competition, hundreds of targets are released and predictions have to be submitted in a timely manner. Before MELD, atomistic physics-based simulations had not been possible this quickly. Some of the predictions made by the team were the best of the entire competition. The team predicted binding poses and relative binding free energies for the complex P53–MDM2 involved in the development of cancer.

Why Blue Waters

Blue Waters is the only HPC machine in the United States that provides sufficient throughput GPU usage for the time-sensitive CASP competition. Every MELD calculation requires at least 30 GPUs. On Blue Waters, the team could run up to 30 calculations. Without Blue Waters, the CASP challenge workload would have taken about 16 times longer to simulate. The staff is helpful and quick to respond to solve issues, and to help to maximize the Blue Waters usage for efficiency.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: Exploratory/50 Knh
PI: Ahmed Elbanna
University of Illinois at Urbana–Champaign

Biology, Chemistry & Health

MULTISCALE MODELING OF BIOFILM DYNAMICS IN DRINKING WATER DISTRIBUTION SYSTEMS: TOWARD PREDICTIVE MODELING OF PATHOGEN OUTBREAKS

Research Challenge

The quality of drinking water is a critical public health issue. Thin films made of cells and extracellular polymeric substances that are found ubiquitously on the surfaces of drinking water distribution systems are known to attract, harbor, and hide pathogens from disinfectants and have been linked to pathogenic outbreaks. To fight these outbreaks, we need to improve our understanding of the response of these biological thin films to different types of disinfectants and their susceptibility to fracture under different conditions.

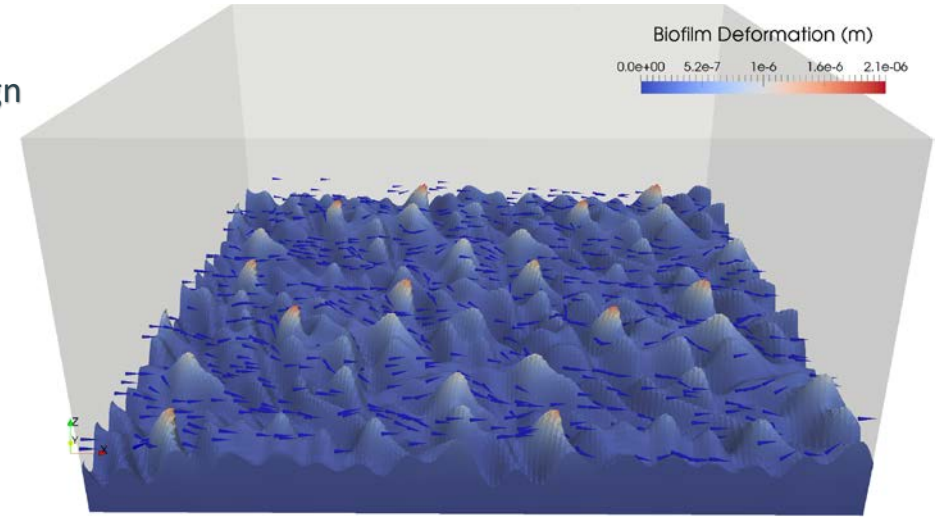

Mechanical deformation of a three-dimensional biofilm (in plane dimensions 1mm x 1mm) with multiscale surface roughness subjected to a fluid flow rate of 1m/s. Displacements are highly localized in the biofilm surface peaks suggesting that these areas are most susceptible to fracture.

Methods & Codes

We developed a continuum nonequilibrium statistical thermodynamics framework for modelling the nonlinear elasto-plastic response of soft amorphous materials. We implemented this model in a nonlinear finite-deformation framework within the MOOSE platform. To describe biofilm–fluid interactions we use a finite-element code developed by Dr. JaeHyuk Kwack, a member of the Blue Waters team.

Results & Impact

Our results have been pushing the limits of the state-of-the-art in modeling biofilm mechanics. To the best of our knowledge, this is the first three-dimensional model of biofilm—fluid system with complex surface geometry that has ever been created. It is capable of modeling 3D turbulent structures near surface irregularities as well as complex stress patterns that are not apparent from 2D simulations.

Why Blue Waters

The progress accomplished to date would not have been possible without Blue Waters. Each simulation of 3D biofilm—fluid interaction generates tens to hundreds of gigabytes of data and requires 10,000 or more core-hours of runtime.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: GLCPC/310 Knh
PI: Peter Freddolino
University of Michigan Medical School
Biology, Chemistry & Health

COMPREHENSIVE *IN SILICO* MAPPING OF DNA-BINDING PROTEIN AFFINITY LANDSCAPES

Research Challenge

Binding of proteins to DNA is an important process that governs gene expression patterns and in many ways dictates the behavior of human cells. Experimental methods used to study affinity landscapes of DNA and proteins carry high costs in terms of both time and money. This project applies methods based on molecular dynamics simulations to develop a cheap and reliable protocol for such studies *in silico*.

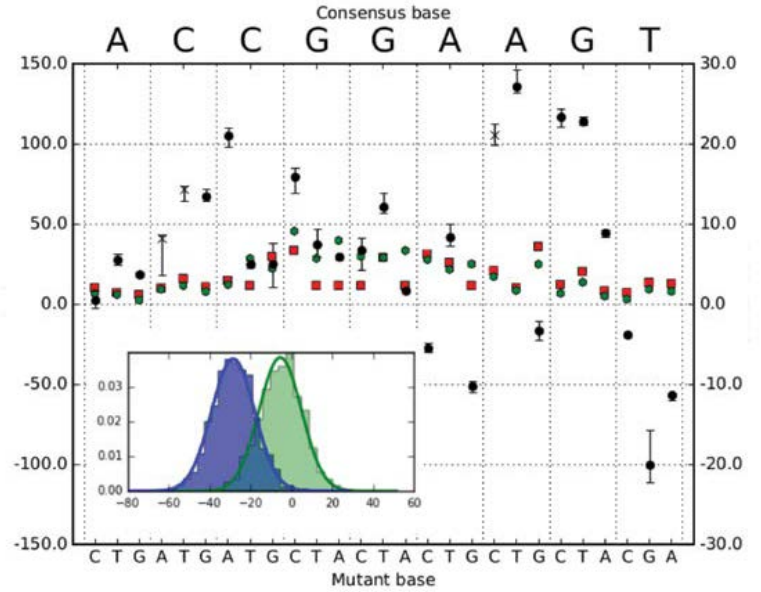

Changes in binding free energy of ELK1 protein to DNA upon single nucleotide mutation. ELK1 plays important roles in long-term memory formation, drug addiction, Alzheimer's disease, Down syndrome, breast cancer, and depression.

Methods & Codes

The Crooks-Gaussian Intersection (CGI) method for calculating free-energy changes due to DNA base pair substitution. This method requires very long equilibrium molecular dynamics simulations of the protein-DNA complex and DNA alone.

Results & Impact

We have shown that equilibrium CGI simulations have poor agreement with experimental observations and produce unrealistic magnitudes of binding free energy changes. Our findings hint at the importance of non-equilibrium interactions between proteins and DNA. A reliable computational method for studying DNA—protein interactions may help better understand the true pathways of drug addiction, Alzheimer's diseases, Down syndrome, breast cancer, and possibly many other pathologies.

Why Blue Waters

This project requires enormous computing power that can handle dozens of simulations of DNA—protein complexes using GPU-accelerated molecular dynamics software, followed by hundreds of state-of-the-art free-energy calculations using CPU-only code. The hybrid architecture of Blue Waters has been ideal for these applications and provided this project with the most efficient environment for every step of the project.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: BW Professor/240 KNH
PI: Sharon Hammes–Schiffer
University of Illinois at Urbana–Champaign
Biology, Chemistry & Health

NON-BORN–OPPENHEIMER EFFECTS BETWEEN ELECTRONS AND PROTONS

Research Challenge

In certain chemical and biological processes, such at the absorption of light, the transfer of a proton is key to understanding the pathway. In such cases in order to accurately model the process the quantum mechanical nature of the proton must be included. Including nuclear quantum effects in traditional quantum chemistry methods is computationally very challenging. This work develops new methods for including nuclear quantum effects and applies those methods to the light absorption cycle for the Flavin adenine dinucleotide (BLUF) photoreceptor protein.

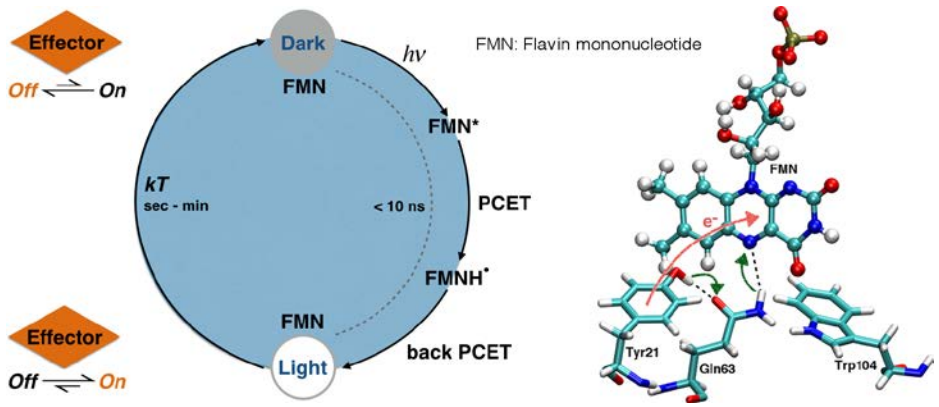

(Left) A schematic diagram of the photocycle in BLUF photoreceptor proteins. The dark state to light state conversion is believed to involve photoinduced PCET and shifts the equilibrium between the “off” and “on” states of the effector domain. (Right) Our studies indicate that the proton relay conformation of the active site Tyr and Gln residues is important for electronic charge transfer from Tyr to the flavin.

Methods & Codes

Using the Nuclear-Electronic Orbital (NEO) method developed by the Hammes–Schiffer group in combination with umbrella sampling techniques implemented in CHARMM, a comprehensive search of the conformation space was carried out. This benefitted from a new integral code that provided a multiple order of magnitude speedup over the previous code.

Results & Impact

The results of the this work clearly demonstrate the effectiveness and efficiency of the NEO-RXCHF method. NEO-RXCHF has also been demonstrated to by much more efficient and computationally scalable than the previous state of the art method.

The NEO-RXCHF method was applied to understanding the signaling mechanism of the BLUF photoreceptor including the geometries of the active sites enabling electron transfer as part of that signaling process.

Why Blue Waters

The simulation is computationally intensive requiring both a large processor count and high-quality interconnect.

Blue Waters staff assistance was greatly beneficial to building and testing the software.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: BW Professor/200 KNH
PI: So Hirata
University of Illinois at Urbana–Champaign
Biology, Chemistry & Health

BRUECKNER–GOLDSTONE QUANTUM MONTE CARLO

Research Challenge

Past computational chemistry simulations to accurately predict molecular and materials properties scale very poorly in terms of chemical system size as well as core count. This project seeks to develop a new stochastic method that is highly scalable in terms of both system size and computer size, yet still provides high accuracy.

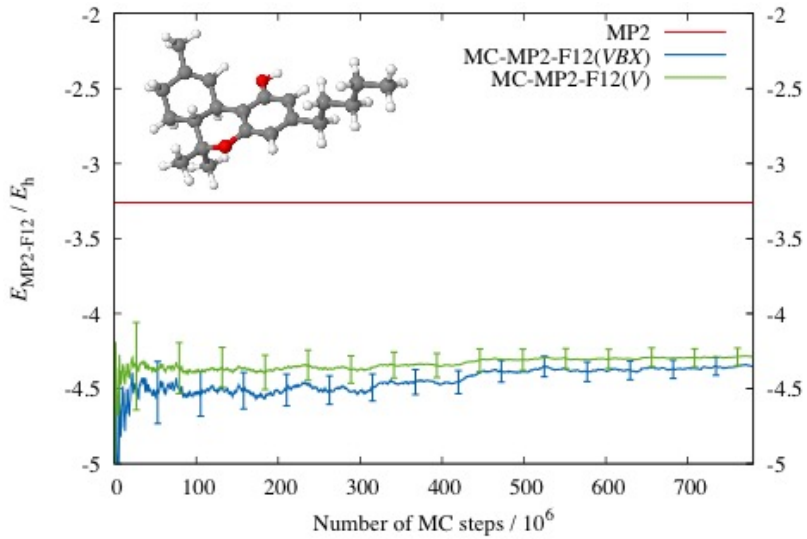

Monte Carlo explicitly correlated MP2 energies of tetrahydrocannabinol

Methods & Codes

The new MC-MP2-F12 method combines stochastic Monte Carlo methods with a transformation of the quantum mechanical Møller-Plesset second order perturbation (MP2) method with a correction to reach the complete basis set limit. The method can compute energy differences on hundreds of GPUs or thousands of CPUs.

Results & Impact

The new methods were applied to demonstrate accurate energetics for tetrahydrocannabinol shown above and C_{60}, C_{70} which play important roles in solar cells.

The methods have demonstrated excellent scalability through 256 GPUs and 4096 CPUs.

Why Blue Waters

Blue Waters provides a stable, easy to use platform allowing the targeting of both CPUs and GPUs and rapid turnaround for testing at large node counts.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: GLCPC/640 Knh
PI: Peter Kasson
University of Virginia
Biology, Chemistry & Health

HOW MEMBRANE ORGANIZATION CONTROLS INFLUENZA INFECTION

Research Challenge

We seek to understand how influenza viral infection responds to membrane perturbation. Initial data suggest that lipid-modifying drugs may have potential as antivirals, but we do not understand their mechanism. Even the fundamental lipid biophysics of viral membrane organization is unknown, and the link from atoms to patients is a grand challenge in understanding and treating viral infection.

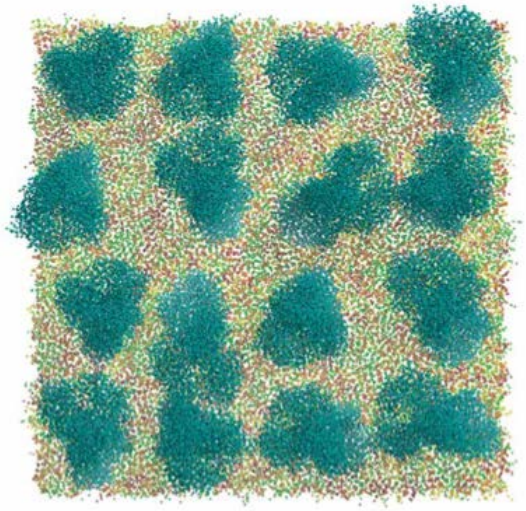

This image shows a model of an influenza membrane patch. Components of the membrane are varied to match experimental perturbations and compared to experimental measurements.

Methods & Codes

We use structural data from X-ray crystallography, cryo-electron microscopy, and mass spectrometry to construct models of influenza viral and target membranes. These membranes are simulated using the molecular simulation software, Gromacs, and are measured under perturbations that match experiments. In-house machine-learning methods analyze the simulations to relate to experiments. This yields a model of membrane response and its infectious consequences at a high level of molecular detail with experimental validation.

Results & Impact

The results have yielded a new membrane-organizing principle that explains the experimental data on influenza virus binding and cell entry. This molecular explanation helps to understand how the influenza virus responds to membrane changes and is likely applicable to cellular signaling as well. Such detailed molecular models of mesoscale phenomena would not be possible without this combination of precise biophysical experiments and large-scale computing.

Why Blue Waters

Blue Waters offers a large number of GPUs that are tightly coupled by a fast Cray interconnect, permitting many simultaneous simulations that each use a number of GPUs to accelerate the calculation.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: Illinois/169 Knh
PI: Zaida Luthey-Schulten
University of Illinois at Urbana-Champaign
Biology, Chemistry & Health

A HYBRID STOCHASTIC-DETERMINISTIC SIMULATION METHOD ENABLES FAST SIMULATION OF CELLULAR PROCESSES IN EUKARYOTES

Research Challenge

Many processes in living cells, especially gene expression, are characterized by low particle numbers, low reaction rates, spatial heterogeneity, and a high degree of randomness.

Simulations of such processes must track the random diffusion and reactions of each particle.

Such simulations become prohibitively expensive in systems with many frequently-reacting particles.

A hybrid method is required to combine both fast and slow processes in a single model.

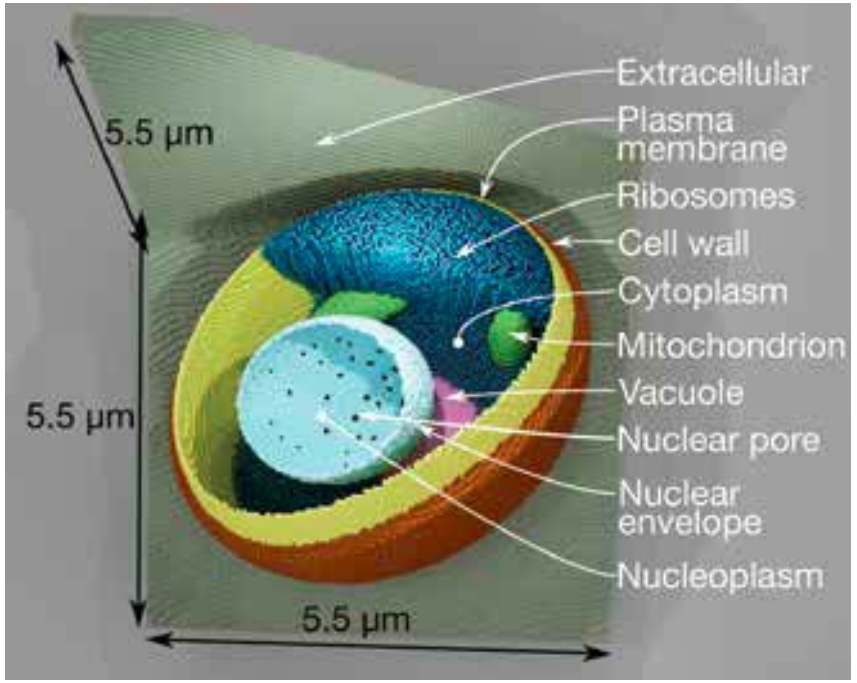

Yeast cell compartmentalization used in the simulations.

Methods & Codes

The Reaction-Diffusion Master Equation is solved by the Stochastic Simulation Algorithm to explicitly model low-concentration chemical species.

A deterministic ordinary differential equation model of the high-concentration chemical species is linked to the stochastic explicit model.

The Lattice Microbes software suite and its pyLM problem-solving environment provide a convenient way to set up simulations of complex biological systems.

Results & Impact

Hybrid simulations are a factor of 50-100 faster than pure stochastic methods in simulating a well-studied galactose switch in S. cerevisiae (yeast).

Researchers can now simulate larger cells, in greater detail, with a large number of species types, many cellular components, and high concentrations of metabolites (sugars, etc.) inside and outside the cell.

Why Blue Waters

Blue Waters was essential to generate over 1,000 replicate hybrid simulations over a simulated time of 750 minutes and a range of concentrations. This was needed to provide sufficient data to make the results statistically reliable and to determine the optimal coupling interval between the stochastic and deterministic models.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: Illinois/200 Knh
PI: Arif Masud
University of Illinois at Urbana–Champaign
Biology, Chemistry & Health

PATIENT-SPECIFIC HPC MODELS AND SIMULATION-BASED IMAGING FOR CARDIOVASCULAR SURGICAL PLANNING

Research Challenge

- To computationally identify regions of atherosclerotic plaque, interior narrowing, or loss of wall elasticity.
- To reconstruct artery bifurcations to add valve replacement via 3D printing.
- Use 3D models as constituent of “Virtual Patient” alternative to angiography.

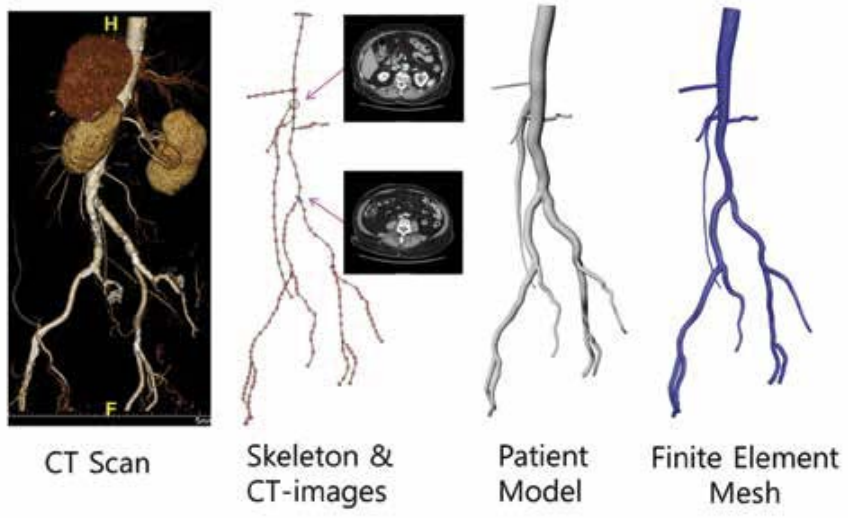

CT image of a branch structure created using the program developed by students on the project team, the surface model of a patient, and the computational grid.

Methods & Codes

We developed a software framework to create high-fidelity patient models from CT scans, using a multiscale, non-Newtonian fluid constitutive model, and a coupling scheme for blood–artery interaction. We also developed a point cloud based method for geometrically reconstructing patient-specific arterial tree models and a graphic visualization package based on open source software from Sandia National Lab.

Results & Impact

Recent advances in computation and imaging allow for determination of flow and pressure from CT scans, but with a massive increase in data. Simulation-based images are essential for clinical diagnostics to aid in analysis. Software developed by this project bridges the gap in bringing HPC resources to clinicians and surgeons.

Why Blue Waters

Novel coupled hierarchical multiscale methods exploited the large local resident memory which was critical for performing large-scale biomechanics simulations.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: Innovation & Exploration/20 Knh
PI: James W. Mazzuca
Alma College
Biology, Chemistry & Health

QUANTUM EFFECTS OF PROTON TRANSFER IN BIOLOGICAL SYSTEMS

Research Challenge

In chemical dynamics calculations, all nuclei are classical particles and all electrons are quantum particles represented by molecular orbitals. There are cases where the quantum mechanical effects of nuclei become important, such as in condensed-phase chemistry in which a hydrogen/proton transfer is the rate-limiting step. We examine a model system to determine how these quantum effects influence the mechanism of charge transfer. [Quantum mechanics is the science of the behavior of matter and its interactions with energy on the scale of atoms and subatomic particles.]

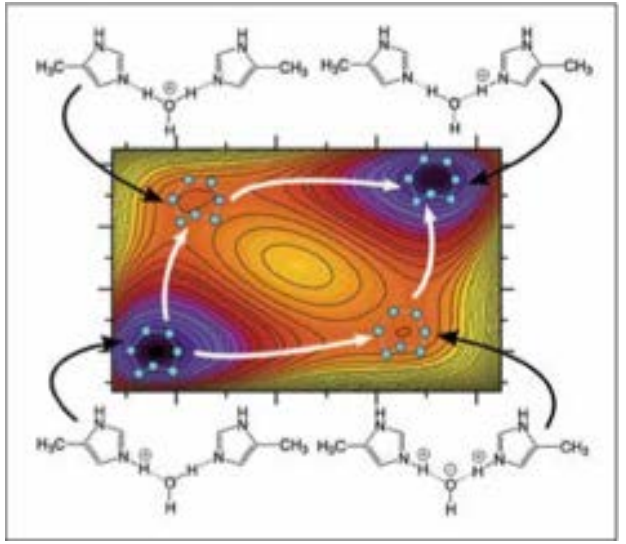

This image shows how potential energy changes as a proton is passed from a donor to acceptor molecule, with the help of an intervening water molecule. Image from *J. Phys. Chem. A*, 121, pp. 819–826, 2017.

Methods & Codes

We use the ring polymer molecular dynamics (RPMD) method to include quantum effects of proton motion in our system. This method is accurate for a wide range of applications, it scales linearly with respect to the system size, and the calculations can be done in parallel across many cores. Our group developed a Fortran-based code, which was highly scalable and well-suited for supercomputing environments such as Blue Waters.

Results & Impact

We find that quantum tunneling of nuclei plays an important role in the charge transfer mechanism, and a physically correct picture of this process depends on the inclusion of these quantum effects.

These calculations have provide us with a very close and detailed view of a physical phenomenon that plays an important role in many biological systems and would not have been possible without highly parallel supercomputing resources.

Why Blue Waters

The RPMD method we have implemented scales extremely well with the highly parallel framework of Blue Waters, which means that even our largest jobs can be completed in under an hour. Additionally, the Blue Waters Student Internship Program provided a perfect gateway into using these resources while funding an undergraduate student's research experience on Blue Waters.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/2,300 Knh
PI: Vijay Pande
Stanford University
Biology, Chemistry & Health

MACHINE LEARNING REVEALS LIGAND-DIRECTED CONFORMATIONAL CHANGE OF μ OPIOID RECEPTOR

Research Challenge

An efficient opioid (medicinally perfect) would be a potent pain reliever without side effects, would show sustained efficacy in chronic treatments, and would not be addictive. The μ Opioid Receptor (μOR) is a G-Protein Coupled Receptor (GPCR) that mediates pain and is a key target for clinically administered analgesics (i.e., pain medicines). Using Blue Waters, the team combined long timescale molecular dynamics (MD) simulations with machine learning to better understand how μOR works by classifying its physiologically-significant states.

μOR samples from at least two deactivating pathways that are gated by two distinct configurations of the DFWY motif

Methods & Codes

The team performed multiple rounds of MD simulations on Blue Waters starting from the active with ligand, active APO, and inactive crystal structures. They used MDTraj to convert and assemble the trajectories. Then, they used the Conformation software package written for and applied to the featurization of this large GPCR MD dataset.

Results & Impact

Using Blue Waters, the team mapped the complete free energy landscape of μOR and discovered novel and significant conformation states not tractable by experiments. The important signal relaying residue switches discovered shed new physical insight into the deactivation mechanism and pathway of μOR. The newly discovered protein structures will be publically available for drug discovery projects.

Why Blue Waters

Blue Waters is an extremely powerful and versatile computational resource. In addition to powerful CPU and GPU hardware, the fast interconnect allows us to do types of calculations (rapid adaptive sampling, Markov State Model construction, force field optimization, etc.) that we could not do on other platforms such as distributed resources (e.g., Folding@home). Also, the availability of the NAMD (nanoscale molecular dynamics) simulation package on Blue Waters has particular advantages for adaptive sampling and restrained equilibrations.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/3,710 Knh

PI: Benoit Roux
University of Chicago

Biology, Chemistry & Health

ELUCIDATING THE MOLECULAR MECHANISM OF C-TYPE INACTIVATION IN POTASSIUM CHANNELS

Research Challenge

The team studies the molecular mechanism of potassium-ion channels, which are of great physiological significance due to their involvement in the function of the central nervous system and heart.

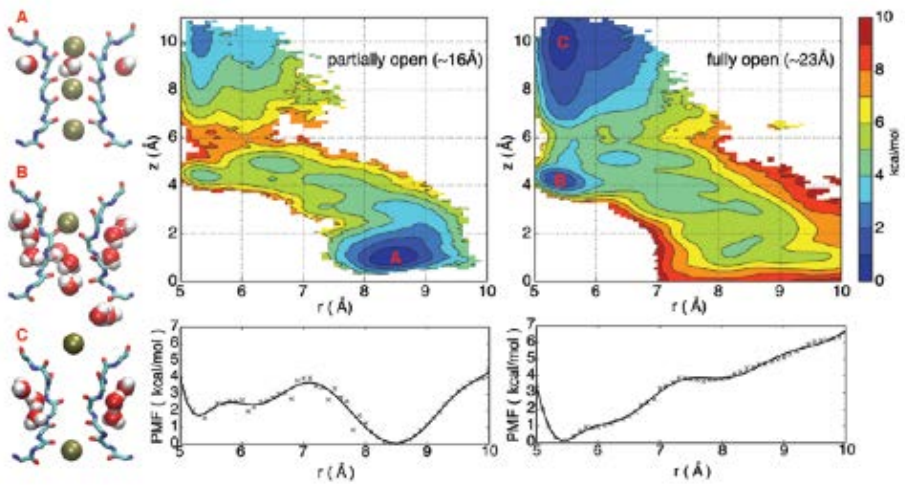

2D map of potential of mean force (PMF) for the conductive-to-constricted transitions for the selectivity filter with different opening degrees of inner gate. The horizontal reaction coordinate r describes the width of the selectivity filter. The vertical reaction coordinate z indicates the position of the external K+ ion along the Z axis relative to the center of the selectivity filter. The lower panel is the one dimension PMF along the reaction coordinate r.

Methods & Codes

The team carries out molecular dynamics (MD) simulations based on an atomic model of the ion channel with the program NAMD using the CHARMM force field PARAM36. The MD and free energy calculations simulate the molecular mechanism of inactivation of the bacterial KcsA channel.

Results & Impacts

Our 2D free energy calculation quantitatively characterized the conformational preferences of the selectivity filter in which a fully open gate (~23 angstroms) highly favors a constricted filter, whereas a partially open gate (~16 angstroms) notably prefers to maintain a conductive filter. The computational results support the notion that the constricted conformation of the selectivity filter correspond to the functional C-type inactivated state of the KcsA channel.

Why Blue Waters

Blue Waters offers the ability to carry out extensive free energy simulations with replica-exchange molecular dynamics using multiple copies of the simulated system.

Jing Li, Jared Ostmeyer, Luis G. Cuello, Eduardo Perozo1, and Benoît Roux (Journal of General Physiology, 2018, in press)

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/707.6 Knh
PI: Ivan Soltesz
Stanford University
Biology, Chemistry & Health

DATA-DRIVEN, BIOLOGICALLY CONSTRAINED COMPUTATIONAL MODEL OF THE HIPPOCAMPAL NETWORK AT FULL SCALE

Research Challenge

The main goal of the project is to study the mechanisms that govern the emergence of characteristic oscillatory behavior in the hippocampal network and its implications for information processing using 1:1 scale biophysical computational models that are closely based on electrophysiological, morphological, and imaging data. The hippocampal circuits that store and recall information are comprised of diverse cell types, each exhibiting distinct dynamics and complex patterns of synaptic connectivity, and the complex interactions between these cell classes can only be resolved via computational modeling.

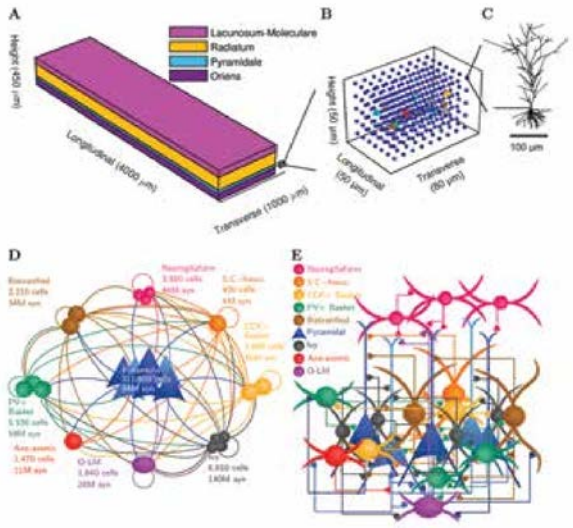

Structure of the CA1 network model. (A) The model network has dimensions identical to the rodent hippocampus, (B) model cell bodies laid out in their respective layers, (C) detailed morphology and synapse placement for each model neuron, (D) illustration of the connectivity types, (E) the characteristics of each connection are constrained by experimental data.

Methods & Codes

The principal simulation environment was NEURON 7.4. NEURON is designed to simulate neuronal models that are described in terms of the membrane properties and geometric structure of neurons, and supports computationally efficient representation of connections among neurons in a network.

Results & Impact

The team extensively validated a first-of-its-kind 1:1 scale, strictly biological data-driven computational network model of the CA1 region of the rodent hippocampus. The model spontaneously generates theta oscillations, which occur in the brain during locomotion and spatial navigation.

Why Blue Waters

A typical behavioral experiment with animals running on a linear track has a duration of tens of seconds. The simulations must therefore have a minimum duration of 10 seconds in order to be comparable to behavioral experiments. Simulations of the CA1 model of 10 seconds of physical time took 14 hours to run on 1,024 Blue Waters nodes. It would not be practical to run such large simulations on any other system. Larger simulations of a model of the entire hippocampus are expected in the future.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: GLCPC/350 Knh
PI: Marcos Sotomayor
The Ohio State University
Biology, Chemistry & Health

STRETCHING THE CADHERIN MOLECULAR VELCRO® OF CELL–CELL JUNCTIONS

Research Challenge

Selective and robust adhesion between cells is essential for multicellular life, and is a basic molecular mechanism in the development, function, and repair of tissues in the human body.

Cadherin molecules protrude from cell surfaces to engage with cadherin molecules on adjacent cells.

Previous work has described detailed interactions between only a single pair of molecules .

Strong adhesion requires a strong and robust lattice of cadherin molecules as studied here.

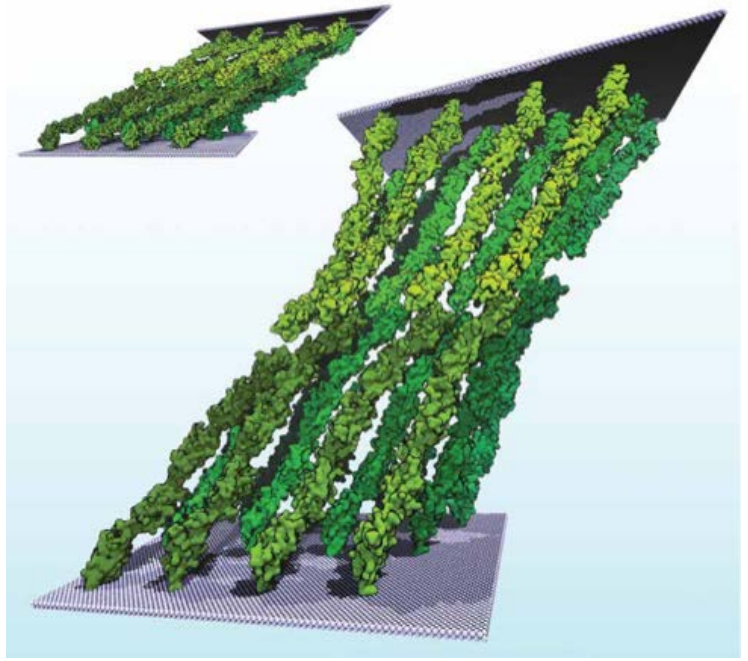

Atomic-resolution simulation of cell–cell adhesion by a lattice of cadherin molecules, shown without tension (top left) and after stretching (center).

Methods & Codes

The mechanical properties of cadherin proteins are best studied in simulations where all atoms, including water and ions, are explicitly modeled.

The programs NAMD and VMD were used to build and simulate the large models of up to 3.7 million atoms required for this work.

The steered molecular dynamics (SMD) technique was used to test the mechanical response of the cadherin complexes to inter-cellular tension.

Results & Impact

Cadherins at cell–cell junctions act as molecular shock absorbers, straightening at low force without breaking their links to the neighboring cell.

At higher tension, links separate without unfolding the cadherin proteins.

Cadherin chains are disordered and floppy when calcium ions are absent.

These simulations provide an unprecedented atomistic view of the mechanisms of cellular adhesion.

Why Blue Waters

Molecular dynamics simulations of large atomistic systems are computationally demanding and cannot be divided into smaller independent simulations to be distributed among poorly networked computational resources. Only a fast networked and massively parallel system like Blue Waters can be used.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/150 Knh
PI: Ashok Srinivasan
Florida State University
Biology, Chemistry & Health

SIMULATION OF VIRAL INFECTION PROPAGATION THROUGH AIRTRAVEL

Research Challenge

The team's goal is to develop models and a novel methodology that can provide insight to decision-makers on policies and procedures that will reduce the likelihood of infection spread during air travel.

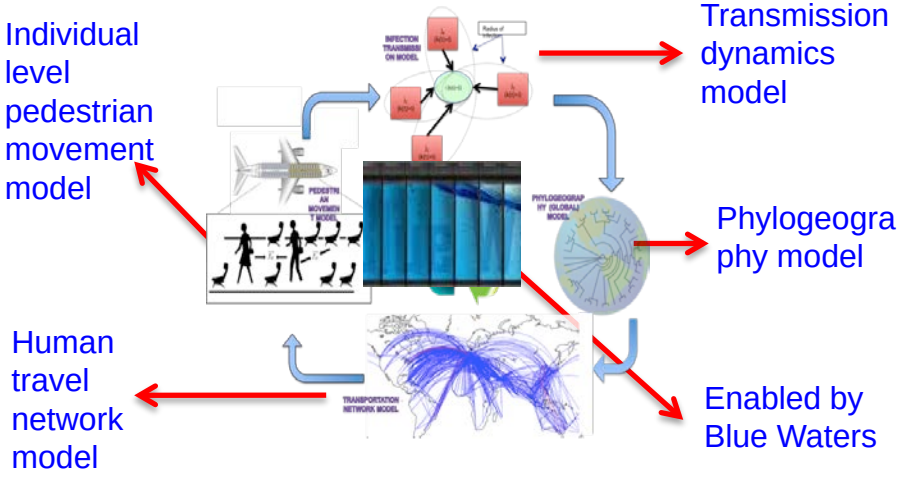

Schematic for combining fine-scale model at airport level with global level phylogeography model to evaluate impact of airport-level public health policies on global disease spread.

Methods & Codes

The team modeled pedestrian movement during air travel as particles based on a force-field approach. The pedestrian trajectory information was then integrated with a model for infection transmission. The results of this model were aggregated at the global level to model the spread of epidemics across large geographic scales.

Results & Impact

The team's results show promise for substantial impact from choices involving procedures for boarding, plane size, and seat assignment on infection spread. Altering pedestrian movement strategies could be used to reduce the spread of disease, especially during epidemics. These results were reported in over 75 news outlets around the world.

FOX NEWS
The Economist
CBS5AZ.com

Why Blue Waters

There was initially a large parameter space of uncertainties, requiring a large computational effort. Subsequent fine-tuning of the models also required low turnaround time and parallelism.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: Illinois/100 Knh
PI: Brad Sutton
University of Illinois at Urbana-Champaign
Biology, Chemistry & Health

HIGH-RESOLUTION MAGNETIC RESONANCE IMAGING OF MECHANICAL PROPERTIES OF THE BRAIN

Research Challenge

A newly developed method, magnetic resonance elastography (MRE) allows us to see the mechanical properties (for example: stiffness and damping ratio) of the brain. We have seen that the mechanical properties are very sensitive to changes that reflect health and performance of the brain. However, advanced image formation approaches and a nonlinear inversion (NLI) are required to get the property maps from the imaging data. These computationally intensive steps result in the need for significant resources for analysis of a single brain.

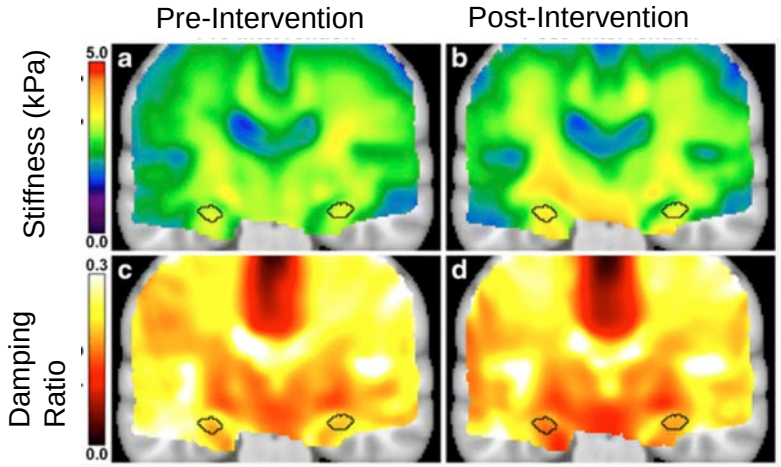

Hippocampal viscoelasticity is improved by exercise training in persons with multiple sclerosis. Exercise therapy improves brain health and cognition, and the MRE measure is used to evaluate the recovery of hippocampal integrity following intervention. It can be used as a biomarker to design and monitor future therapy.

Methods & Codes

MRE uses an external mechanical stimulator to apply microscopic shear waves to the brain. The full 3D waves are captured using synchronized motion sensitized imaging, from which material properties are estimated using NLI. NLI divides the brain into smaller subproblems for optimization of properties at a local level, then reconstituting as a global solution.

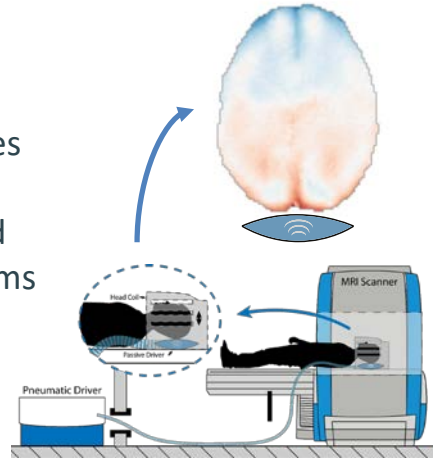

Results & Impact

Two recent studies have shown correlations between memory performance and hippocampal mechanical properties. Two ongoing studies look to characterize the structural–functional properties of the normal aging human brain. The team investigated the relationships among hippocampal viscoelasticity, memory performance, and fitness by using MRE in a study of exercise training in persons with multiple sclerosis (MS) and observed that this intervention significantly improved their hippocampal viscoelasticity and memory performance.

Why Blue Waters

The Blue Waters system provides a unique resource for MRE because it enables the application of highly compute intensive algorithms across the significant numbers of subjects necessary for meaningful neuroscience research. In addition, the availability of this resource, combined with the training and support Blue Waters provides, are transforming how our group approaches neuroimaging by embracing applications were previously computationally intractable.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/300 Knh
PI: Ilias Tagkopoulos
University of California, Davis
Biology, Chemistry & Health

A CRYSTAL BALL OF BACTERIAL BEHAVIOR: FROM DATA TO PREDICTION USING GENOME-SCALE MODELS

Research Challenge

To develop an accurate model, spanning several scales to predict cellular state and traits in novel environments.

There are many obstacles to achieving this vision, including the large complexity of the cellular organization and machinery, the lack of cohesive datasets, and models that are capable of integrating them into one system that is more than the sum of its parts.

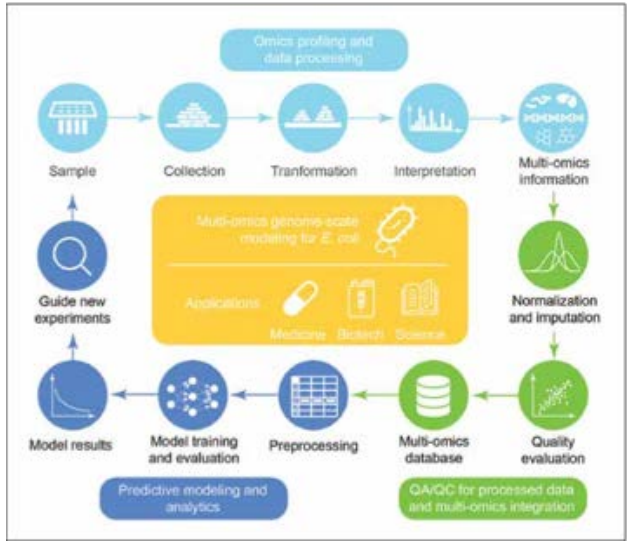

Overview of the processing pipeline and genome-scale modeling for *E. coli*. There are three major steps: 1) omics profiling and data processing, 2) QA/QC for processed data and multi-omics integration, and 3) predictive modeling and analysis

Methods & Codes

The DeepPep code, written in torch/Python, is based on convolutional neural networks and deep-learning techniques. Results show that it can predict proteins from peptides more accurately than any technique until now.

The genome-scale model employs a recurrent neural network and a constrained regression altogether to predict genome-wide responses layer by layer.

Results & Impacts

This work led to the most accurate predictor of bacterial behavior. Applications in clinical microbiology (antibiotics resistance) biotechnology (production of drugs and chemicals).

Why Blue Waters

Large-scale simulations, on the order of thirty million data points, as well as the effectiveness of GPU technology for training deep neural network models, necessitate the use of high-performance computing.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/13,940 Knh

PI: Klaus Schulten (deceased), Emad Tajkhorshid (successor)
University of Illinois at Urbana-Champaign

Biology, Chemistry & Health

STUDYING CELLULAR PROCESSES THROUGH THE COMPUTATIONAL MICROSCOPE

Research Challenge

Infectious viral pathogens such as the human immunodeficiency virus type 1 (HIV-1) and the hepatitis B virus (HBV) are major risks to public health, and millions of people die annually due to a lack of effective anti-viral treatments. Developing novel drug compounds that can target viruses depends heavily on characterizing components of virus structure and the roles these components play in facilitating infection. One such component key to virus function is the capsid, a protein shell that packages the viral genome and regulates its delivery to the host cell nucleus. Virus capsids are currently of great pharmacological interest as drug targets.

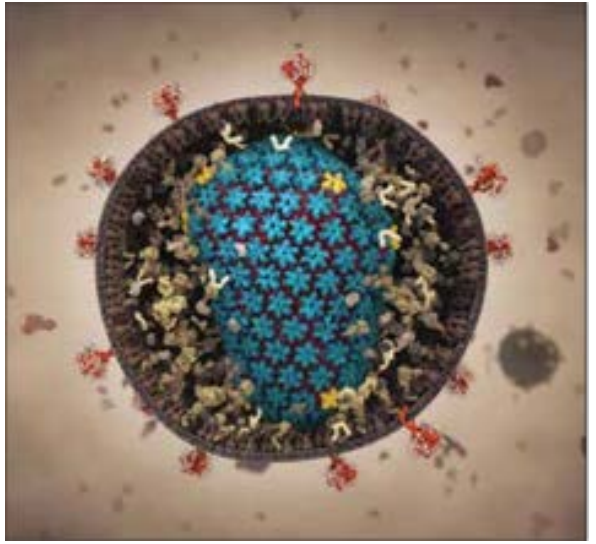

At the core of the HIV-1 virion lies its capsid. Composed of repeating copies of proteins, the capsid protects the viral genome and performs key functions throughout the viral life cycle that are essential to successful infection.

Methods & Codes

Molecular Dynamics (MD) simulation provides a powerful technique to investigate the structure and properties of virus capsids. All-atom simulations are capable of capturing subtle effects on capsid structure and dynamics induced by bound drug molecules. While simulation of capsids comes at great computational expense, access to the NAMD simulation software on Blue Waters has revealed new insights into capsids, as well as to suggest mechanisms by which drug molecules can disrupt them.

Results & Impact

A simulation characterizing the dynamical behavior of the HIV-1 capsid system (64 million atoms) over the timescale of 1 μs, indicate new avenues for the development of drugs that seek to disrupt the capsid by altering its complex biophysical properties. Extensive unbiased simulations have characterized the dynamical behavior of the HBV capsid system (6 million atoms) in the presence and absence of three distinct drug compounds over timescales of 1 μs.

Why Blue Waters

Due to their formidable computational expense, these simulations of virus capsids are only possible on a machine like Blue Waters. Further, analysis of the data sets generated by these simulations is feasible only through access to the massively parallel computing power and high-performance Lustre filesystem provided by a resource like Blue Waters. The exciting discoveries revealed by this project underscore the essential role for Blue Waters in the development of anti-viral treatments, and demonstrate that access to leadership-class computing facilities holds the potential for significant impact on overall public health.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: NSF PRAC/1,200 Knh
PI: Gregory A. Voth
University of Chicago
Biology, Chemistry & Health

LARGE-SCALE COARSE-GRAINED MOLECULAR SIMULATIONS OF THE VIRAL LIFECYCLE OF HIV-1

Research Challenge

- HIV-1 Gag polyproteins self-assemble into the “immature protein lattice” at an infected cell’s membrane as part of the process of releasing viral particles to spread the infection.
- Disruption of this self-assembly is a potential therapeutic target, but difficult to study in conventional experiments.
- Understanding aggregation and interactions of large numbers of biomolecules at cell membranes can advance fundamental biophysical knowledge.

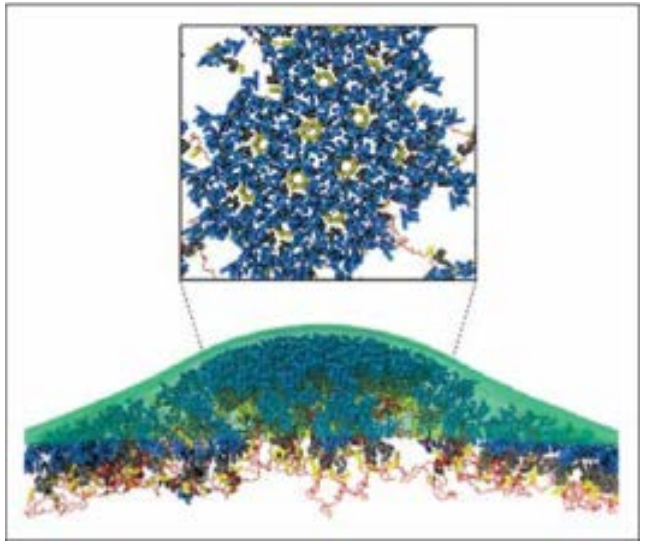

Molecular snapshot depicting early assembly of Gag polyprotein(blue, yellow, and grey tubes) as catalyzed by RNA (red chain) at the site of a deforming membrane puncta (translucent green sheet).

Methods & Codes

- We created large-scale coarse-grained (CG) molecular models of relevant viral components and the cell membrane to investigate immature protein lattice self-assembly.
- CG models are computationally efficient, enabling otherwise-impossible simulated time- and length-scales.

Results & Impacts

- By combining computer simulations with fluorescence localization experiments, we elucidate the interactions that regulate HIV-1 viral assembly dynamics.
- Aspects of the interactions, difficult to control experimentally, are tractable with our CG models.
- Insights gained through CG analysis can assist in designing new therapeutic approaches.

Why Blue Waters

- These simulations required simulating large numbers of individual molecules over long time scales, under a wide range of biologically-relevant conditions.
- The tightly coupled parallel nature of our simulations made the leadership-class computing capabilities and cutting-edge network hardware of Blue Waters crucial in successfully investigating this and other molecular systems of significant biomedical interest.

Social Sciences, Economics and Humanities

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: Illinois/400 Knh
PI: Wendy K. Tam Cho
University of Illinois at Urbana-Champaign
Social Science, Economics, & Humanities

ENABLING REDISTRICTING REFORM: A COMPUTATIONAL STUDY OF ZONING OPTIMIZATION

Research Challenge

Political redistricting is intended to provide fair representation in Congress to all communities and interest groups. Gerrymandering occurs when districts are drawn in a manner that discriminates against a partisan or racial group. Due to a lack of sufficient tools to analyze and synthesize redistricting data, the Supreme Court has found it difficult to identify a workable standard by which we might regulate gerrymandering. This team has developed a computational redistricting tool utilizing massively parallel high-performance computing for redistricting optimization and analysis.

An example of a Congressional electoral map generated by PEAR for the state of North Carolina.

Methods & Codes

The algorithm, PEAR, or Parallel Evolutionary Algorithm for Redistricting, uses MPI nonblocking functions for asynchronous migration, and the C SPRNG 2.0 library to provide a unique random number sequence for each MPI process. They designed spatial evolutionary operators that incorporate spatial characteristics to effectively search the solution space. The parallelization of the algorithm maximizes the overlapping of computing and communication at runtime.

Results & Impact

By incorporating spatial evolutionary operators to handle spatial characteristics and the associated computational challenges, and harnessing massive computing power, PEAR provides a powerful and novel computationally scalable redistricting tool. The discrete optimization framework identifies large sets of quality electoral maps. The project has been the subject of numerous amicus briefs and been discussed in oral arguments before the Supreme Court.

Why Blue Waters

The PEAR library is designed for extreme-scale redistricting applications. From the beginning, it was intended to scale to all of the processor cores on Blue Waters through nonblocking MPI communication calls. The computational approach implemented in the solution requires generating a very large number of electoral maps for quantitative study of redistricting phenomena. Identifying quality electoral maps requires significant computing in the combinatorial optimization process. Generating a large number of statistically independent maps is only feasible on a supercomputer at the scale of Blue Waters.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: GLCPC/250 Knh
PI: Lars Hansen (1), Yongyang Cai (2)
(1) University of Chicago, (2) The Ohio State University
Social Science, Economics, & Humanities

POLICY RESPONSES TO CLIMATE CHANGE IN A DYNAMIC STOCHASTIC ECONOMY

Research Challenge

There are uncertainties in climate and economic systems. Integrated Assessment Models (IAMs) of climate and the economy are designed to analyze the impact of policy responses to climate change. We develop and solve new computational IAMs that merge spatial temperature systems, climate tipping points, economic risks, clean energy usage, regional economic activities, etc. to do economic analysis about the optimal climate policy under uncertainty and risks, and how such a policy will impact economic activities.

Methods & Codes

We developed DSICE (Dynamic Stochastic Integration of Climate and the Economy) as a computational framework for this research. There are 3 parallel packages for stochastic dynamic programming, nonlinear certainty equivalent approximation, and supergames with high parallel efficiency. They are designed to handle large-scale problems and perform well on Blue Waters.

Results & Impact

Economics integrated with climate change introduces problems with the size and complexity that justify using massively parallel systems. The Blue Waters work allows economists to solve computationally intensive problems. We have produced three papers; the first showed an 8-fold increase in social costs of carbon; the second led to a method to solve huge-dimensional stochastic problems efficiently; and the third demonstrated a new parallel algorithm to model strategic interactions among multiple constituents.

Why Blue Waters

Our parallel algorithms include frequent communications, in our master-worker structure, of small or moderate sizes. These algorithms are limited by the high latency inherent in commodity clusters to solve large problems in a reasonable amount of time. The largest problem we have solved on Blue Waters had 372 billion tasks, and Blue Waters has allowed us to solve this and similarly large calculations efficiently.

Graduate Fellows

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: 2016-2017 BW Grad Fellow
Elizabeth Agee
University of Michigan
Bioscience

RESOLVING PLANT FUNCTIONAL BIODIVERSITY TO QUANTIFY FOREST DROUGHT RESISTANCE IN THE AMAZON

Research Challenge

The increased frequency and severity of droughts and their regional consequences have highlighted the potential vulnerability of the Amazon Basin region to heat- and drought-induced stress. This study seeks to understand what controls the response of biodiverse tropical rainforest regions to heat and drought-induced water limitation.

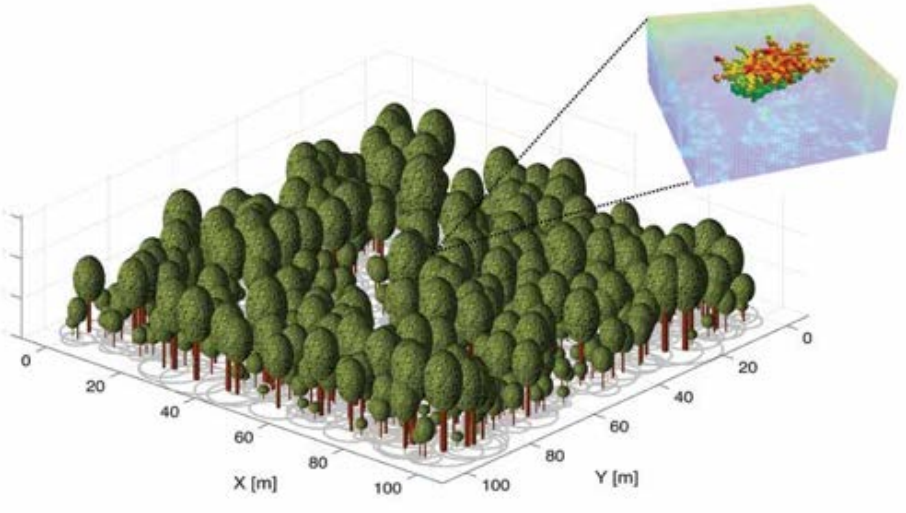

Root water uptake was modeled for individuals within the canopy. This representation shows the spatial distribution of trees and their relative heights. Grey circles indicate the lateral spread of root systems. Root systems overlap (inset), pulling water from shared soil water reserves.

Methods & Codes

Root water uptake has been coupled with the Department of Energy's massively parallel flow and transport model (PFLOTRAN) to simulate water uptake for individual trees.

Using these tools, we are exploring how tree roots contribute to forest drought resilience in areas of the Amazon rainforest during the recent 2015–2016 El Niño drought event.

Results & Impacts

Model simulations are able to compute water uptake for upwards of 1500 trees per hectare—comprising the largest simulations of three-dimensional root water uptake ever attempted. Blue Waters is being used to run a large number of simulations, testing the contributions of individual traits to individual and community drought resilience. These results will help us understand how to improve the next generation of earth system models and the increase our understanding of below-ground processes.

Why Blue Waters

Blue Waters is critical to the ongoing success of this project. Simulations of this complexity and scale require the computational power of this system to make meaningful analyses. Not only are the simulation domains complex, multiple simulations are needed to account for system uncertainty. Blue Waters will allow this analysis to be done with a much larger domain and much higher diversity of species than has been previously possible.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Allocation: 2015-2016 BW Grad Fellow
Maureen T. Brooks
University of Maryland
Biological Oceanography

MODELING NONLINEAR PHYSICAL-BIOLOGICAL INTERACTIONS: EDDIES AND *SARGASSUM* IN THE NORTH ATLANTIC

Research Challenge

- To highlight interactions among the macroalgae *Sargassum* and mesoscale eddies and fronts
- To better predict *Sargassum* dispersal and growth

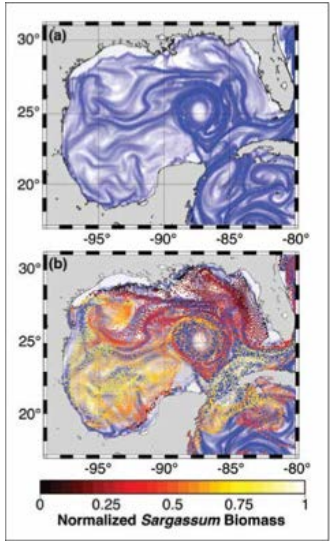

(upper) Attracting Lagrangian coherent structure field for the Gulf of Mexico for November 6, 1993, as calculated by finite-time Lyapunov exponent from a 28-day backwards-time particle integration. (lower) Normalized biomass of *Sargassum* particles, overlaid on the LCS field for the same date. Note higher biomass in the western GoM.

Methods & Codes

A system of four coupled models was used to perform these simulations:

- A modified version of the Hybrid Coordinate Ocean Model (HYCOM) code
- A biogeochemical model including nutrient chemistry and plankton growth
- A Lagrangian particle model to simulate *Sargassum* raft movement
- A *Sargassum* growth model including reproduction and mortality

Results & Impacts

This study found the Gulf of Mexico and Western Tropical Atlantic are the two key regions that influence the *Sargassum* distribution throughout its range.

Eddy activity may help maintain the populations in these key regions.

This is important for tracking the sources of the *Sargassum* blooms that lead to costly wash-up events that have inundated Caribbean and Gulf beaches in recent years.

Why Blue Waters

The HYCOM model has a high computational cost; the research proposed here also requires additional ocean biogeochemistry models.

The professionalism of the NCSA staff has also been key to the success of this project. Their responsiveness and expertise made implementing and running this code on Blue Waters as straightforward as possible.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

2016-2017 BW Grad Fellow
Iryna Butsky
University of Washington
Astronomy & Astrophysics

THE ROLE OF COSMIC RAYS IN ISOLATED DISK GALAXIES

Research Challenge

- To simulate isolated disk galaxies for which cosmic rays are dynamically important in order to study how cosmic ray-driven outflows shape the structure and kinematics of the circumgalactic medium and to compare these with observations
- To implement a new cosmic ray fluid that is compatible with the Riemann solvers in Enzo

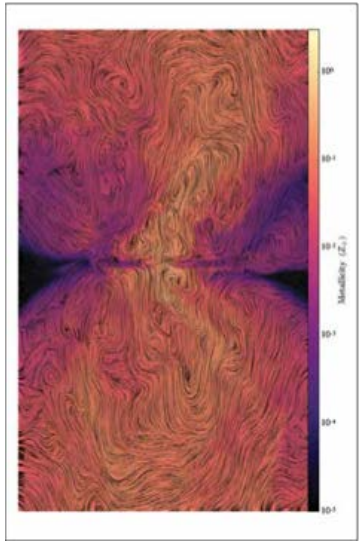

An edge-on view of outflows from an isolated disk galaxy after 1.5 Gyr. The color depicts the metallicity (metal enrichment relative to solar abundances) of the gas, and the streamlines follow the topology of the magnetic field. The dimensions of the image are 90 kpc x 150 kpc.

Methods & Codes

- Enzo code
- Finite volume method for hydrodynamics
- Adaptive mesh refinement
- Gravity
- Cosmic ray (particle) propagation including anisotropic diffusion, streaming, and gas heating

Results & Impact

- Results show evidence of strong, mass-loaded outflows that enrich the circumgalactic medium when cosmic rays are present
- Directly comparing the simulations will place better constraints on theories pertaining to the circumgalactic medium and make predictions for the structure and metallicity distribution of the circumgalactic medium for future observations.

Why Blue Waters

- Galaxy simulations need to resolve a large dynamic range in physical and temporal scales from sub-parsec to kiloparsec (kpc) scale
- Galaxy simulations require the use of massively parallel, high-performance supercomputers such as Blue Waters
- This Graduate Fellow has benefited greatly from the support of Blue Waters staff, who are admirably dedicated to resolving issues in a timely manner.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

2015-2016 BW Grad Fellow
Justin Drake
University of Texas Medical Branch
Biophysics

TOWARD DEVELOPING A THERMODYNAMIC MODEL OF BINDING-INDUCED CONFORMATIONAL TRANSITIONS IN SHORT, DISORDERED PROTEIN REGIONS

Research Challenge

The project studies intrinsically disordered regions (IDRs) in proteins, which are highly dynamic and often undergo conformational transitions upon binding with a protein target. Understanding of the thermodynamics underlying these transitions, particularly conformational entropy and its contribution to protein binding, is necessary to develop new drugs that bind to IDRs.

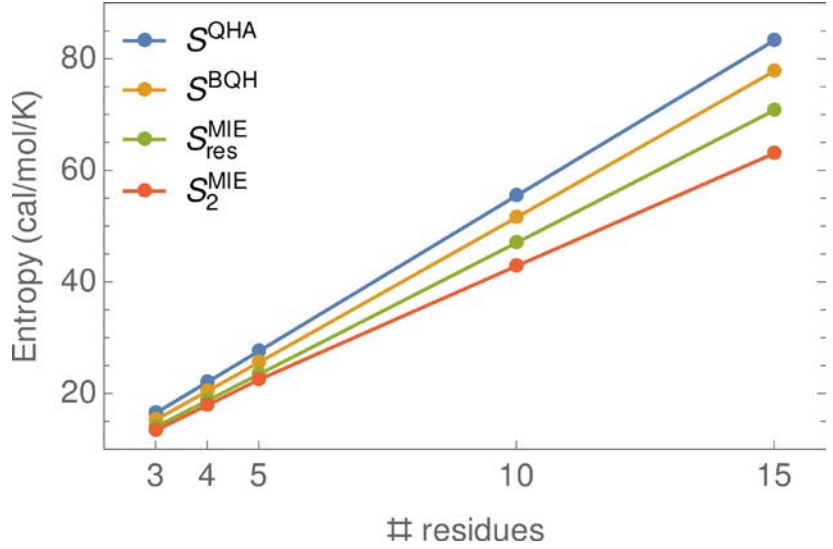

Conformational entropy calculated from the dihedral angles of successively longer glycine polypeptides using the QHA, BQH, and two variants of the MIE method.

Methods & Codes

The project involves all-atom molecular dynamics (MD) simulations to sample the conformational states of successively longer glycine polypeptides (GlyN where N = 3, 4, 5, 10, and 15 residues), and uses a number of methods to calculate the dihedral angle contribution to the conformational entropy. Simulations rely on the NAMD package and Amber ff12SB force field.

Results & Impacts

The obtained data show that conformational entropy scales remarkably linearly with chain length. At a temperature of 300K, the conformational entropy of the short IDR models (e.g., ~21 kcal/mol for Gly15) provides a significant free energy reservoir that proteins may tap through order-disorder transitions to modulate or regulate protein binding.

Why Blue Waters

Access to GPU-optimized MD packages such as NAMD for use with Blue Waters XK nodes, assistance from the project staff to develop an efficient workflow, and the high-throughput capacity of Blue Waters were instrumental for success of the project. Additionally, the Blue Waters Graduate Fellowship provides the opportunity to successfully propose, manage, and partially fund the doctoral research.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

2016 - 2017 BW Grad Fellow
Paul Hime
University of Kentucky
Bioscience

GENOMIC PERSPECTIVES ON THE AMPHIBIAN TREE OF LIFE

Research Challenge

All organisms trace their ancestry back to a single common ancestor nearly four billion years ago. Yet today, life has diversified into tens of millions of species. Reconstructing these evolutionary relationships is a key aim of the field of phylogenetics, and such insights may inform nearly all aspects of modern biology. Yet, it is becoming increasingly clear that different regions of the genome can support conflicting phylogenetic hypotheses, and reconciling these discordant gene genealogies is a key problem facing evolutionary biologists today.

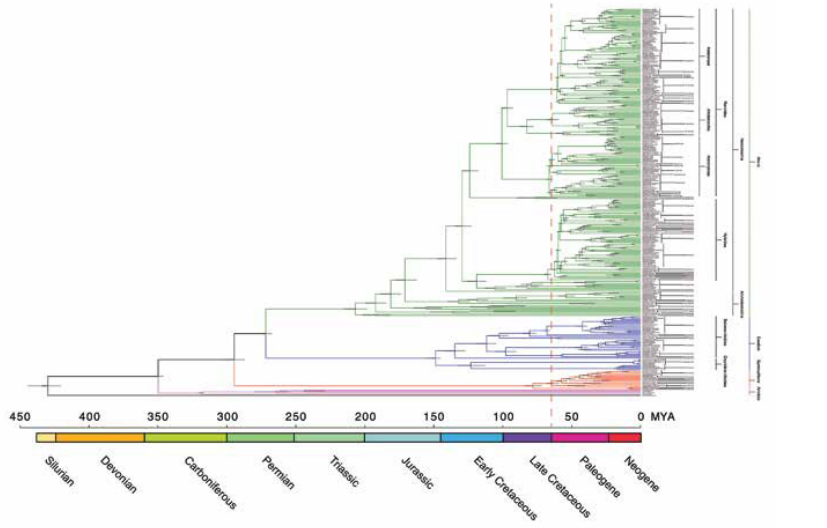

Estimated divergence times across Amphibia—Amniotes, caecilians, salamanders, and frogs are depicted by magenta, red, blue, and green branches, respectively. The dashed red line denotes the Cretaceous–Tertiary boundary.

Methods & Codes

Three possible topologies exist for relationships among frogs, caecilians, and salamanders. Gene trees were estimated in RAxML. The Akaike information criterion (AIC) was used to quantify the direction and magnitude of support across genes. Gene trees were then reconciled into an estimate of the species tree using Astral. This species tree topology was then used in concert with a set of 25 fossil calibrations to estimate divergence times across Amphibia in MCMCTree.

Results & Impacts

This project has demonstrated that different genes in the amphibian genome support conflicting topologies for the relationships among the three amphibian orders. Analyzing hundreds of genes for hundreds of species has traditionally been computationally intractable for empirical data sets, and this study is one of the first to use an information-theoretic framework to address not only the direction of support for phylogenetic hypotheses across the genome but also the magnitude of that support.

Why Blue Waters

The Blue Waters fellowship has provided essential support to enable Paul to make progress on this project and to successfully write and defend his Ph.D. dissertation. Paul was able to leverage the massively parallel resources on Blue Waters to perform topological testing in Neobatrachian frogs at a level of gene- and species-sampling that has not been attempted in any other empirical data set.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

2016 - 2017 BW Grad Fellow
Michael P. Howard
Princeton University
Fluid Systems

MULTISCALE SIMULATIONS OF COMPLEX FLUID RHEOLOGY

Research Challenge

To correctly model the *dynamics of complex fluids* like oil in water mixtures, it is essential to resolve both the direct interactions between solute molecules (oil) as well as the *solvent-mediated interactions* (water). The latter *dominate the computational cost* for a molecular model of such a mixture; however, a molecular-level description of the solvent itself is often not of interest. *A multiscale approach that simplifies the solvent model while preserving its most important interactions interactions is required* to study complex fluids at relevant length and time scales.

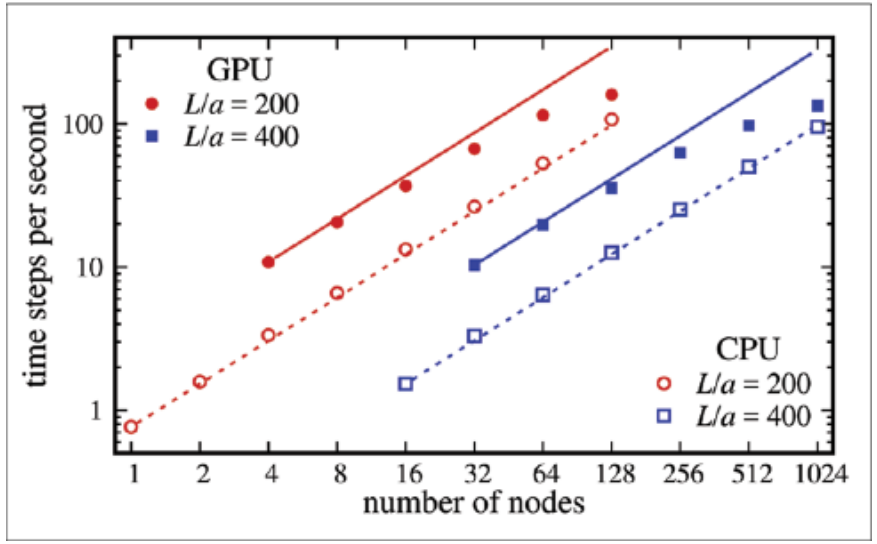

Strong-scaling benchmarks of the MPCD software on Blue Waters for CPU-only (XE nodes) and GPU (XK nodes) implementations.

Methods & Codes

The *multiparticle collision dynamics (MPCD) algorithm* was implemented as part of the open-source simulation package *HOOMD-blue* and optimized for NVIDIA GPUs. All MPCD data were stored independently from HOOMD-blue's molecular dynamics (MD) data to ensure high performance. Nearly all computations are performed exclusively on the GPU, which minimizes latency associated with data transfers between the host and the device.

Results & Impacts

According to strong-scaling benchmarks of the developed MPCD codes on Blue Waters XE and XK nodes, *GPU acceleration on the XK nodes gave a roughly 3x speedup compared to the XE nodes*. The complex fluid and soft matter research communities will significantly benefit from the developed MPCD software, especially the GPU implementation, which will permit studying processes at scales that would otherwise be inaccessible.

Why Blue Waters

Blue Waters is the only system available to them that delivers both the CPU and GPU resources necessary to develop and optimize their software at scale. *The large number of GPUs available in the XK nodes also significantly increases their overall scientific productivity* and allows them to study process dynamics that would be challenging or impossible to obtain with fewer resources.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

2014 - 2015 BW Grad Fellow
Alexandra L. Jones
Princeton University
Molecular Bioscience

MULTISCALE SIMULATIONS OF COMPLEX FLUID RHEOLOGY

Research Challenge

A cloud model that treats broadband integration and 3D radiative transfer in a highly accurate and unbiased way is needed to quantify the bias in the simpler models ubiquitously used. Such a model was not publicly available prior to this project.

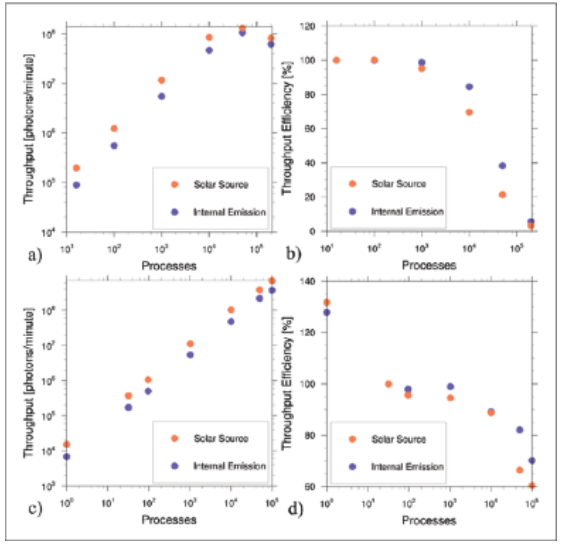

Performance of the "IMC+emission" model runs in both internal emission mode and solar source mode for strong scaling experiments (a-b) and weak scaling experiments (c-d). Throughput (a, c), defined as the number of photons simulated per minute of wall-clock time, and throughput efficiency [%](b, d) are shown as a function of number of processes. Note that logarithmic spacing is used on the horizontal axis in all panels and on the vertical axis for panels a and c.

Methods & Codes

The Monte Carlo approach was chosen in order to exploit the parallelism inherent in the method. Mixed with the radiative transfer equation both weak and strong scaling experiments yield good results and demonstrate parallel efficiency.

Results & Impacts

The goal of this project is to make publicly available to the radiative transfer community the models, tools, data, and products developed to aid in faster and more robust progress in addressing scientific questions about the interactions of clouds and realistic radiative transfer.

Why Blue Waters

Access to profiling and debugging tools such as CrayPat and the Allinea Debugger allowed streamlining of the development process. Having access to a point of contact on the SEAS staff helped me think through issues and find tailored solutions for problems that would have otherwise delayed progress for weeks.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

2016 - 2017 BW Grad Fellow
Andrew Kirby
University of Wyoming
Engineering

HIGH-FIDELITY BLADE-RESOLVED WIND FARM SIMULATIONS

Research Challenge

High-fidelity numerical simulation of wind energy applications is becoming a precedent for future technologies, not only for the wind energy sector but also for lower-fidelity modeling. This is especially important in the study of downstream wake effects on wind turbines that are the primary cause of decreased wind plant power production efficiency.

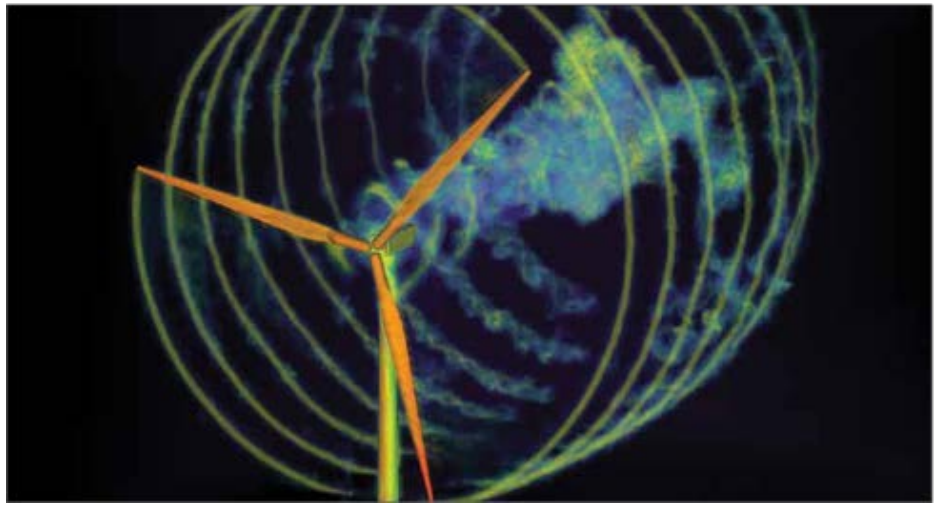

Volume rendering of the vertical wake structures generated downstream of a Siemens SWT-2.3-93 wind turbine with uniform inflow.

Methods & Codes

The numerical methods developed in this research utilize an overset grid paradigm where multiple meshes and multiple flow solvers are used in a coupled manner to enable efficient simulation of this truly multiscale problem. The overset solver is TIOGA and the flow solvers, both developed at the University of Wyoming, are NSU3D and dg4est.

Results & Impacts

This research enabled the largest complete wind farm simulations to date using a high-fidelity blade-resolved turbine model.

Why Blue Waters

Blue Waters offers a unique environment not only as a computational resource but also for its expert project staff. The design of Blue Waters makes it an excellent machine geared toward scientific output rather than just its flop rate.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

2014 - 2015 BW Grad Fellow
Sara Kokkila Schumacher
Stanford University
Chemistry

REDUCING THE COMPUTATIONAL COST OF COUPLED CLUSTERY THEORY

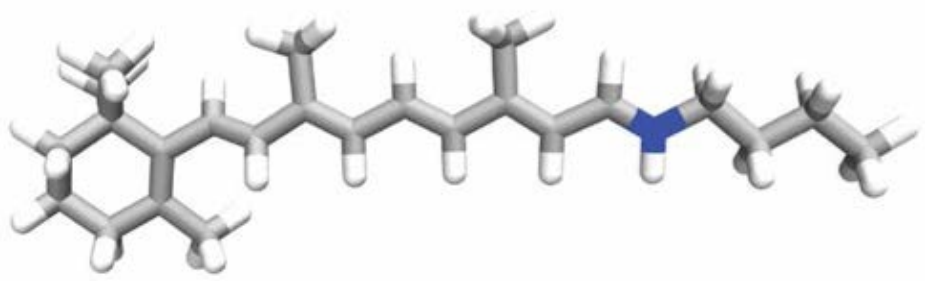

Retinal model used in this study.

Research Challenge

The project aims to develop a parallel implementation of the tensor hypercontraction equation-of-motion second-order approximate coupled cluster singles and doubles (THC- EOM-CC2) electronic structure method to accurately and efficiently simulate the absorption spectra of retinal models.

Methods & Codes

The work involves redesigning the THC-EOM-CC2 algorithm developed in TeraChem to take advantage of parallelism. MPI parallelization takes places at a higher level to take advantage of multiple compute nodes, whereas GPU acceleration happens at a finer level. The grid-based variant of the THC approximation allows the new algorithm to block over grid point indices. Laplace transformation techniques are used to express certain terms by numerical quadrature.

Results & Impacts

Application of the THC approximation to EOM-CC2 (THC-EOM-CC2) lowered the cost of the electronic structure calculation from an order of O(N5) to O(N4). The developed THC-EOM-CC2 performs better than other available excited state methods that are capable of calculation of absorption spectra of retinal models.

Why Blue Waters

The Blue Waters code development environment and tools facilitated debugging different implementations more rapidly. The Blue Waters project staff offered valuable insight into code development for high-performance computing systems. The massively parallel architecture of Blue Waters enabled conducting computationally intensive electronic structure calculations of the ground and multiple excited states of many different configurations of the retinal model.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

2015 - 2016 BW Grad Fellow
Larissa Reames
The University of Oklahoma
Meteorology

SIMULATED EFFECTS OF URBAN ENVIRONMENTS ON THE DYNAMICS OF A SUPERCELL THUNDERSTORM

Research Challenge

Earth's population is increasingly concentrated in urban areas, with nearly two-thirds of the world's population expected to live in urban areas by 2050. However, interactions between urban areas and synoptically-active convection, such as supercells, remain relatively unexamined.

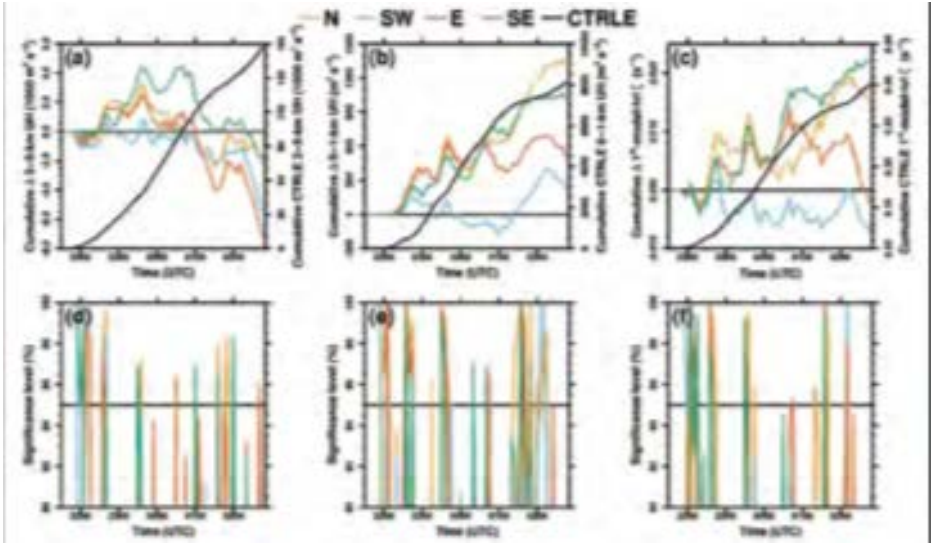

Cumulative CTRLE group average (black) and accumulated difference from CTRLE of N (yellow), SW (blue), E (red), and SE (green) groups of (a) mid-level (2- to 5-km) storm strength, (b) low-level (0 to 1-km) storm strength, and (c) near-surface circulation strength as a function of time.

Methods & Codes

Using the Weather Research and Forecasting (WRF) model—a community mesoscale numerical weather prediction model, the study investigated a total of 134 simulations of a supercell thunderstorm to quantify the impacts of a large Great Plains urban area on the evolution and strength of a supercell thunderstorm, with careful consideration on grid sizes and boundary layer requirements.

Results & Impacts

- Compared full-physics urban simulations to CTRLE; Investigated the effects of the storm having various city-relative path and the storm lifecycle stage during urban interactions; Undertook a factor separation approach to determine the relative importance of the roughness and thermal characteristics of urban areas on storm modification.
- The findings argue for a greater focus on the implications of urban surface roughness.

Why Blue Waters

The large computational and storage requirements of the analysis has made Blue Waters vital to this work. While each simulation was relatively small, the large quantity of of simulations needed to produce significant results required the large computational and data storage capacities of Blue Waters.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

2016 - 2017 BW Grad Fellow
Sherwood Richers
California Institute of Technology
Astronomy & Astrophysics

MONTE CARLO NEUTRINO CLOSURES IN 3D GRMHD SIMULATIONS OF CORE-COLLAPSE SUPERNOVAE AND NEUTRON STAR MERGERS

Research Challenge

- Simulating neutrino radiation requires modeling the six-dimensional integro-differential Boltzmann equation
- Currently impossible for the duration of the supernova explosion
- Understand not just what causes core-collapse supernovae, but also how they happen
- Team makes the first jump to verification of Boltzmann transport in multiple spatial dimensions

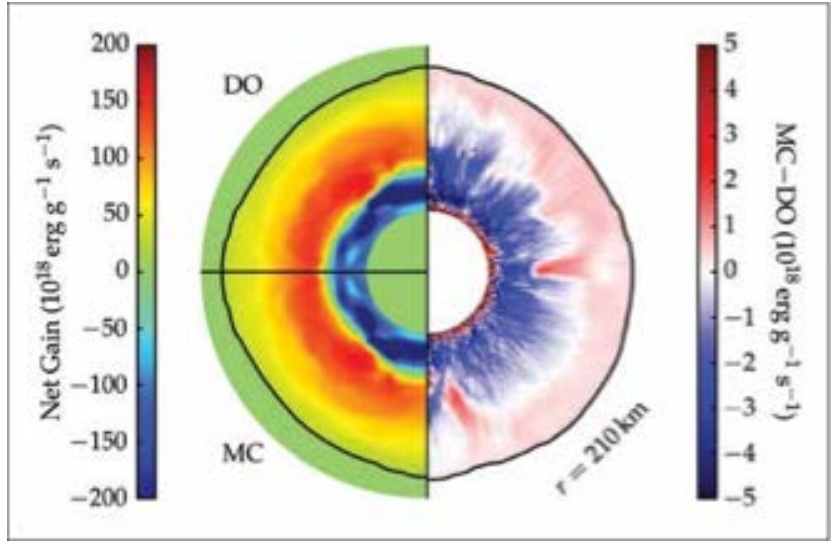

The heating rate in a supernova snapshot computed by the Monte Carlo code Sedonu (bottom left) and a discrete ordinates code of (top left).They agree very well, but the small differences between them (right) could make the difference between an explosion and a dud in a dynamical simulation.

Methods & Codes

- General Relativistic Magnetohydrodynamics (GRMHD)
- Sedonu Monte Carlo neutrino transport code
- Monte Carlo neutrino radiation transport
- GR1D supernova code, relatvistic hydrodynamics, gravity

Results & Impacts

- Optimized Sedonu code
- First multi-dimensional comparison of full Boltzmann neutrino transport methods
- Parameter study for gravitational wave signals from rotating core collapse events using 1,800 simulations; found that current gravitational detectors cannot distinguish between different current descriptions of nuclear matter

Why Blue Waters

- The domain-replication parallelism used requires both a large amount of compute time and a large amount of memory on each node, which Blue Waters provides
- The support staff with both domain and system expertise significantly simplified the task of implementing and optimizing our algorithms on Blue Waters

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

2016 - 2017 BW Grad Fellow
Ronald Stenz
University of North Dakota
Atmospheric & Climate Science

THE IMPACTS OF HYDROMETEOR CENTRIFUGING ON TORNADO DYNAMICS

Research Challenge

To quantify, for the first time, the impacts of centrifuging precipitation on the tornado vorticity budget. Preliminary findings have removed an unrealistic build-up of precipitation in the vortex center (widely seen in tornado simulations) for both idealized vortices and simulations of an entire storm and the tornado it produces

This knowledge will advance our understanding of tornadoes by making the simulations used to study these destructive and dangerous storms more physically realistic

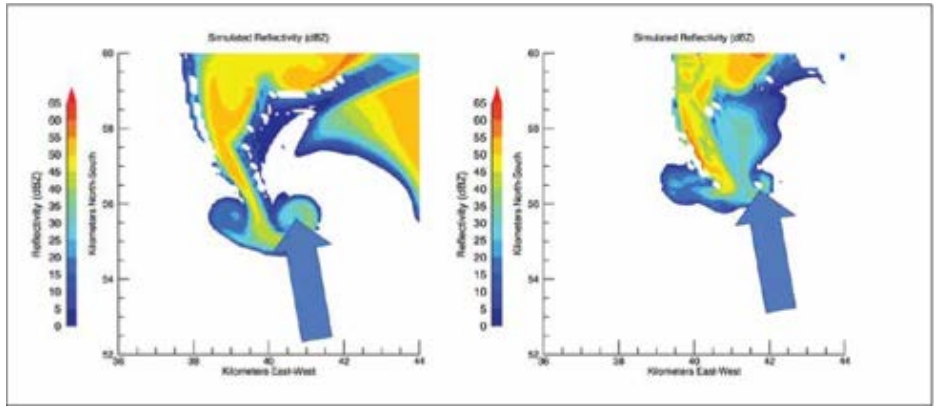

Rain is beginning to accumulate in the simulated tornado center (left). With centrifuging turned on (right), there is an eye-like feature in the simulated tornado center consistent with radar observations of actual tornadoes.

Methods & Codes

Simulations use the Cloud Model 1 (CM1) code

Results & Impacts

A better understanding of tornado dynamics can help to minimize the loss of life caused by storms

Simulations in the proposed research will be more consistent with what is observed in nature, particularly the un-realistic buildup of precipitation in the vortex center

Why Blue Waters

Tornado simulations require thousands of computing cores and produce large amounts of d

ata that must be stored and analyzed

Additionally, the technical and visualization support available with Blue Waters greatly facilitated accomplishment of these research goals.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

2016 - 2017 BW Grad Fellow
Erin Teich
University of Michigan
Applied Physics

GLASSY DYNAMICS AND IDENTITY CRISES IN HARD-PARTICLE SYSTEMS

Research Challenge

A clear picture of the thermodynamics and structural change undergirding glass formation remains unknown. This is due in large part to the significant slowing down of any system as it approaches the glass transition, requiring investigations of glass formation to resolve, dauntingly, system dynamics on time scales that vary by orders of magnitude.

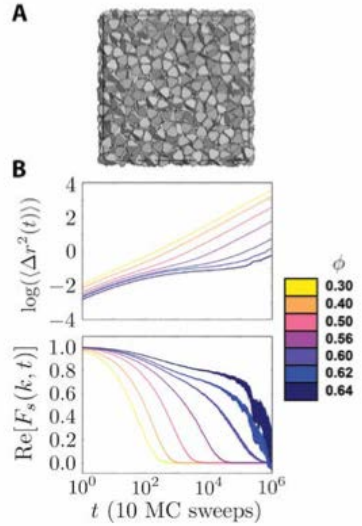

A glass-forming system of hard polyhedral particles. (A) A system snapshot at a volume density of 0.62. (B) Both the mean-squared displacement and the self-intermediate scattering function of all particles in the system over time, for a variety of densities. Plateaus in both parameters are indicative of caging.

Methods & Codes

- Chose non-assembling monatomic systems of hard polyhedral particles, with no interactions aside from those of excluded volume.
- Chose particle shapes consisting of tetrahedra with varying degrees of edge and vertex truncation.
- Used the hard particle Monte Carlo (HPMC) simulation method included in HOOMD-blue tool kit to simulate these systems.

Results & Impacts

- Simulated systems of 4096 particles on a single GPU at various densities, and subsequently measured structural and dynamical information.
- Wrote system trajectories to disk very frequently and collected data for the far longer period of 100 million Monte Carlo (MC) sweeps.
- Produced trajectories as large as 3.25 terabyte (TB) per simulation, each containing about 10 million simulation frames.

Why Blue Waters

HOOMD-blue HPMC's parallel nature on both CPUs and GPUs has made taking advantage of the computational resources on Blue Waters--in particular, its GPU resources--possible . Blue Waters' storage capability has also allowed the generation and storage of tens of TB of raw data on the system for further analysis.

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

2015 - 2016 BW Grad Fellow
Samuel Totorica
Stanford University
Physics

MAGNETIC RECONNECTION IN LASER-DRIVEN PLASMAS: FROM ASTROPHYSICS TO THE LABORATORY *IN SILICO*

Research Challenge

Magnetic reconnection is a fundamental plasma process that converts magnetic field energy into plasma kinetic energy though the breaking and rearrangement of magnetic field lines. The goal of this project is to use simulations to study particles from reconnection in varied plasma conditions, and in particular to investigate whether laser-driven plasma experiments could be used to study the particle acceleration properties of reconnection in the laboratory.

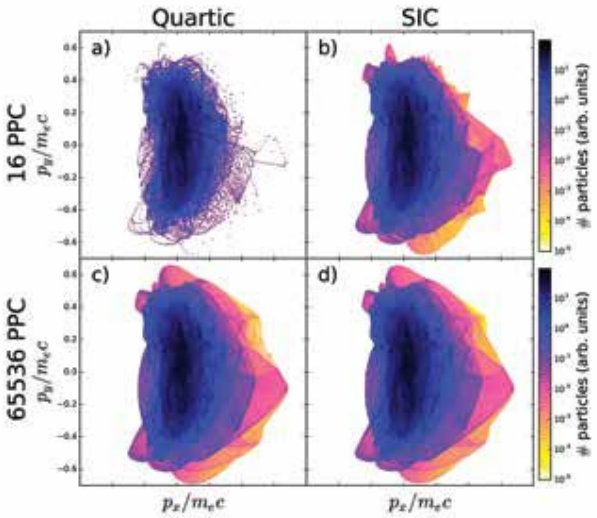

Momentum space for simulations of the Weibel instability with 16 particles-per-cell (PPC) (top) and 65,536 PPC (bottom). The SIC (simplex in cell) deposit (right) captures the tails and shape of the distribution at low PPC much better than a high-order (quartic) particle deposit (left).

Methods & Codes

One of the most powerful tools for *ab initio* (“from first principles”) plasma simulation is the particle-in-cell (PIC) method, which treats the plasma as a collection of discrete simulation particles that interact via electromagnetic forces. The simulations for this project were run using the massively parallel, fully relativistic PIC code OSIRIS, and match the experimental conditions produced by the most energetic laser systems in the world, such as the National Ignition Facility.

Results & Impacts

- The team demonstrated that electrons can be accelerated by reconnection with sufficient quantity and energy to be detected in the laboratory.
- Simulation results are being used to guide several experimental programs in the United States.
- The team has been developing simplex-in-cell (SIC), a novel method for plasma simulation that, under the right conditions, has been shown to reduce the number of required simulation particles by a factor of 1,000.

Why Blue Waters

This project required the use of large-scale 2D and 3D simulations with sufficient size and resolution to bridge the multiscale physics, from fluid dynamics to the kinetic microscopic processes. These computationally demanding simulations can require billions of simulation particles, and demand the cores, memory, and communication performance available on Blue Waters.