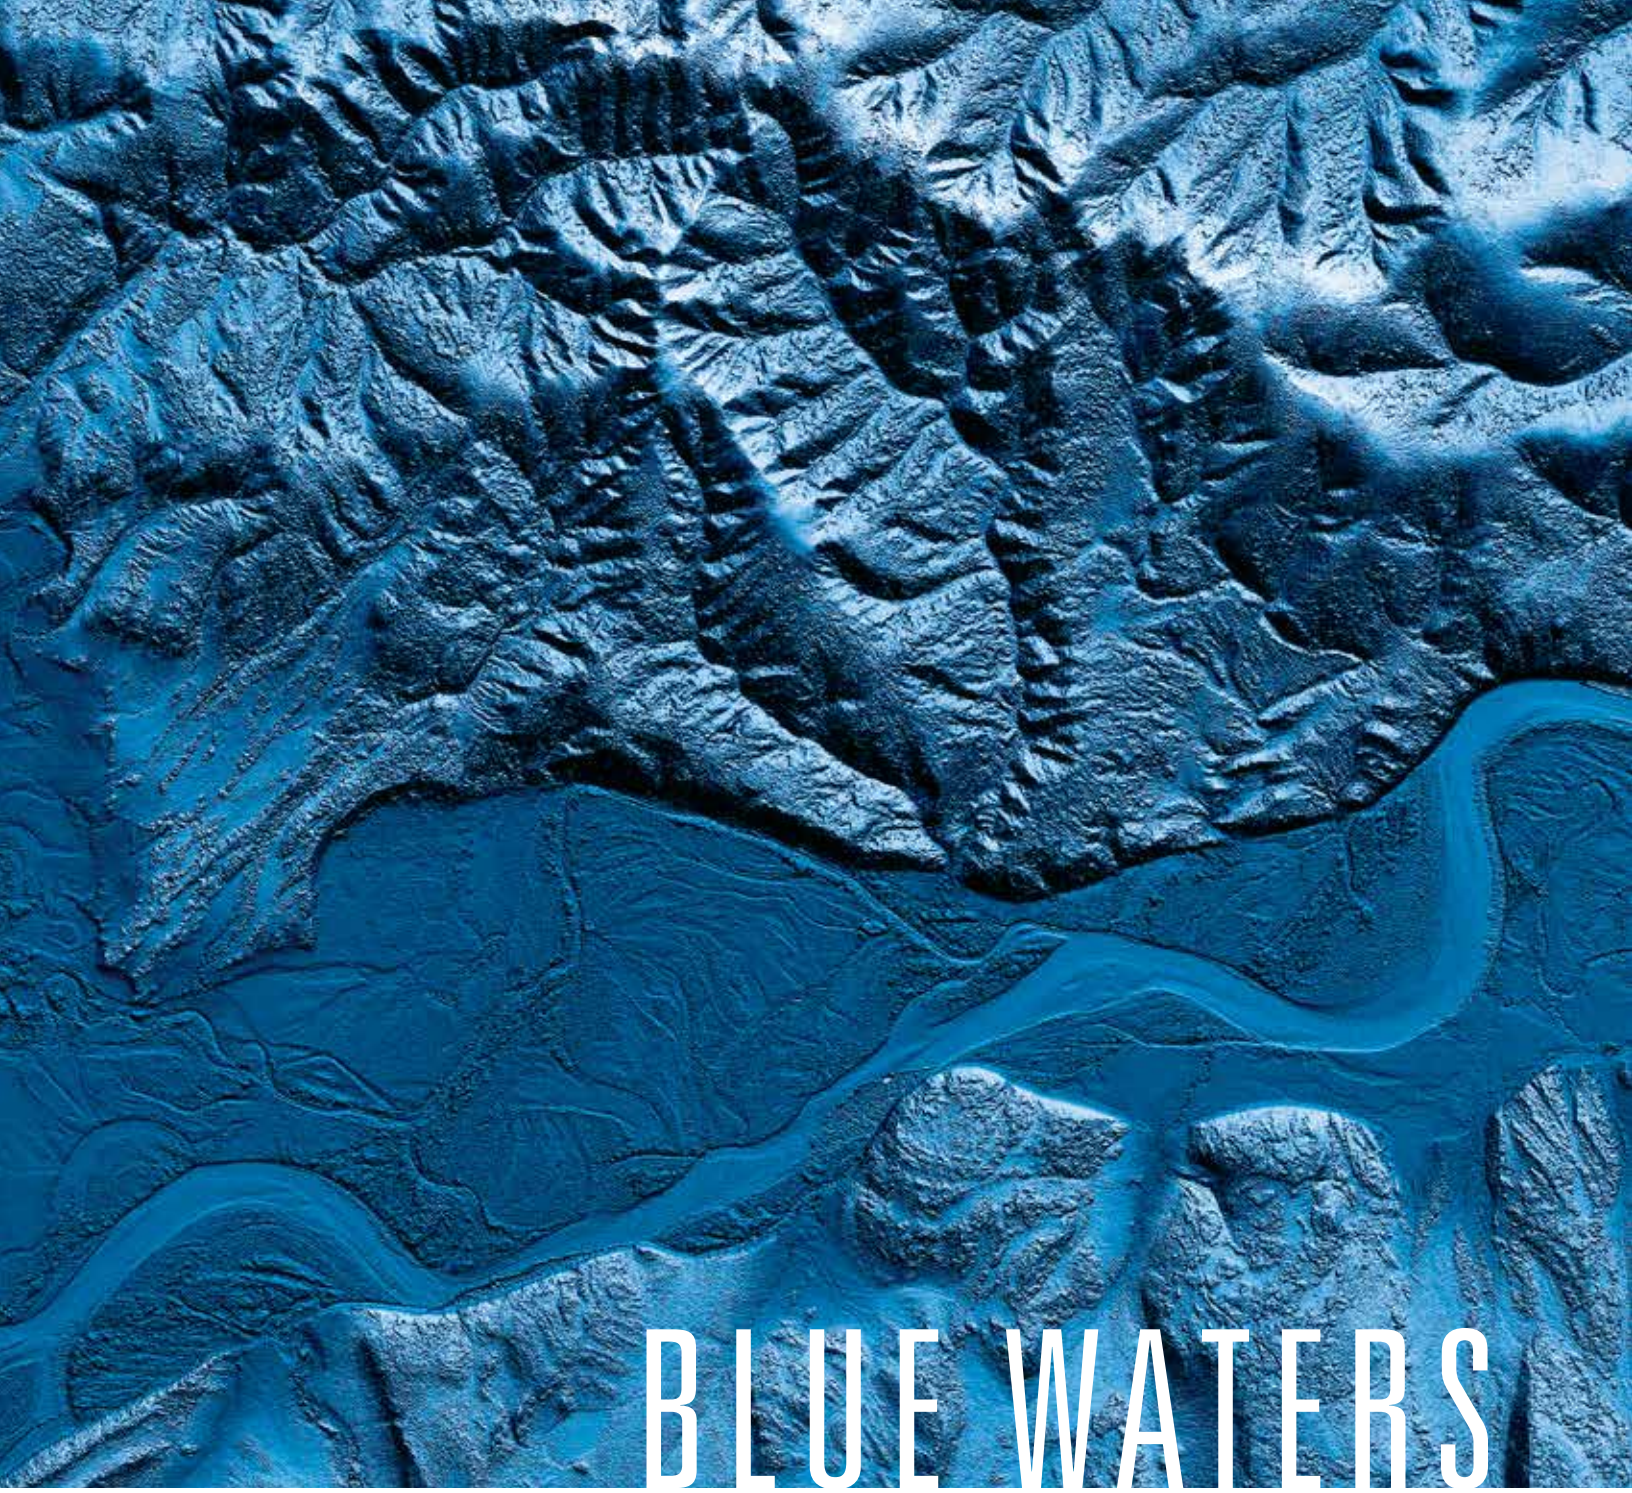

BLUE WATERS

SUSTAINED PETASCALE IN ACTION:
ENABLING TRANSFORMATIVE RESEARCH

2016 ANNUAL REPORT

BLUE WATERS

SUSTAINED PETASCALE IN ACTION:
ENABLING TRANSFORMATIVE RESEARCH

2016 ANNUAL REPORT

Executive Editor
Barbara Jewett

Editor
Lois Yoksoulian

Creative Director
Steve Duensing

Graphic Designer
Christie Klinger

Proofreaders
Hannah Remmert
Susan Szuch

Project Director
William Kramer

The research highlighted in this book is part of the Blue Waters sustained-petascale computing project, which is supported by the National Science Foundation (awards OCI-0725070 and ACI-1238993) and the state of Illinois. Blue Waters is a joint effort of the University of Illinois at Urbana-Champaign and its National Center for Supercomputing Applications.

Visit https://bluewaters.ncsa.illinois.edu/science-teams for the latest on Blue Waters-enabled science and to watch the 2016 Blue Waters Symposium presentations.

TABLE OF CONTENTS

KB = kilobytes
TB = terabytes
PB = petabytes
I/O = input/output
Knh = thousand node hours
Mnh = million node hours

Allocations denoted as type/size in extended abstracts.

A MESSAGE FROM THE NCSA DIRECTORS

We hope you will be as excited as we are by the research vignettes across all areas of science and engineering compiled in this document. By providing computational and analytical researchers with Blue Waters' advanced computing and data capabilities, the National Science Foundation, the University of Illinois, and NCSA are catalyzing discoveries that will have immense impact in a diverse range of fields, sparking new understanding of our world and opening new avenues for future research.

Many of the frontier projects using Blue Waters are working at unprecedented scale, combining data from simulation and observation or experiment, and cutting across disciplinary boundaries. These projects exemplify the trend toward *transdisciplinary digital convergence*, which is at the heart of NCSA's vision for the future of research, education, and innovation. The challenges and opportunities facing science and society demand this new paradigm, which draws together multiple disciplines, research techniques, and powerful cyberinfrastructure such as Blue Waters and the ever more powerful systems that must follow it, to create a groundbreaking new mode of investigation.

For example, consider how Blue Waters and NCSA are helping to change the way genetic medicine is researched and practiced in Africa. Much of what is known about the genetics of diseases is based on people with European ancestry, but the Consortium for Human Heredity and Health in Africa (H3Africa) aims to change that by promoting research that takes into account the genetic diversity of African populations. An international partnership of institutions and researchers collaborated to discover genomic variants in over 300 deeply sequenced human samples to help construct a genotyping chip specific for African populations. Blue Waters was an essential component of this work.

Look at how the Polar Geospatial Center and its partners are addressing the dearth of high-resolution, high-quality elevation data for the Arctic region. The team, led by the University of Minnesota's Paul Morin, uses data from satellite imagery, photogrammetric algorithms, and the power of Blue Waters to construct very high resolution digital elevation models that will be used to support activities including transportation and infrastructure, natural resource and land management, and scientific research. This work will transform our knowledge of the Arctic, turning it from the worst to one of the best mapped regions on Earth.

Our goal is to make NCSA **the** home for this type of convergent digital research. Advanced digital capabilities like Blue Waters are and will continue to be a critical engine for this mode of investigation. Deeper integration of advanced computing systems, data analysis and visualization tools and services, instruments, and technical expertise will be required to fulfill the vision and possibility of convergence; NCSA is pioneering this necessary digital fabric.

Convergence also requires strong transdisciplinary teams that unite a wide range of diverse collaborators. NCSA is working to develop and nurture teams in six key thematic areas: astronomy and physics; bioinformatics and health science; computational and data-enabled science; culture and society; earth and environment; and materials and manufacturing. NCSA strives to offer the fertile creative environment where faculty, staff, students, industry partners, and collaborators from around the globe can come together to take on challenges and advance knowledge in ways that are not possible elsewhere.

In the past few years, Blue Waters and the NCSA team have made possible remarkable work in biology, chemistry, physics, geosciences, cosmology and astrophysics, atmospheric science, and many other fields such as economics and social sciences. Even more remarkable breakthroughs will be possible as creative people and powerful digital tools converge at NCSA to transform the conduct and culture of research. We are enormously excited to be part of this transformation, and to see what comes next.

Ed Seidel
Director, NCSA

William Gropp
Acting Director, NCSA

SHARING THE WONDER OF SCIENCE SUCCESSES

In this, our third annual report of how the Blue Waters Project is enabling the most challenging and complex open science investigations, I am pleased to point out Blue Waters continues to expand the breadth and depth of our impact to all fields of science and continues advancing the development of our petascale workforce.

This annual report attempts to share some of that wonder. Each researcher or research team during the 2015-2016 reporting period was invited to briefly present highlights from their research that has leveraged Blue Waters—the National Science Foundation's most powerful system for computation and data analysis.

You will notice this report contains 30 percent more high-impact result summaries than the 2015 report, which itself was an increase over the 2014 report. The fact that in just three short years of providing one of most effective leadership systems in the world is making such a remarkable difference to fields like physics, bio-medicine, geo and space science, biology, economics, social and political science, big data analysis is a tribute to the vision of our stakeholder, NSF, to our exceptional science and research teams, and to the dedicated Blue Waters staff.

In the following pages, you can read about ground breaking investigations into exploding supernovae and a dwarf "dark galaxy"; understanding, at unprecedented levels, how sunlight is transformed into chemical energy; investigations for new insights into how influenza, ebola, and other virus strains infect people; how the entire polar regions are changing at resolutions and time to insight that is millions of times more than could be done just two years ago; how fluid flows in applications from steel casting to blood flowing in our bodies and what happens when particles and ice are mixed with the flow; how earthquakes, plate tectonics and supervolcanoes involve and influence people; and how computer solutions can assist our political districts for better fairness and effectiveness. Blue Waters is accelerating research across a wide range of science and engineering disciplines.

The Blue Waters mission is to enable the *"best of breed"* science and research investigations that are not feasible or tractable on less capable resources. We then help extend the impact of the first best of breed calculations so that the same or other teams can apply them to equally challenging problems. Hence, we already have best of breed methods and codes introduced in the first year of the project used to bring insight and understanding to problems that are now two to 10 times larger and more complex.

Every day, the Blue Waters Project brings previously impossible or intractable investigations and insights within the reach of researchers across the United States. Every day, the combination of massive computing power and the intellectual might of pioneering scientists and engineers create opportunities for us to better understand and shape our world. While we are now three years into full-service operations for this supercomputer and its associated support, training, and education efforts, time hasn't dulled the sense of wonder and delight I feel when I consider the doors that are being opened to discovery.

The Blue Waters project has always been about much more than just providing computing cycles, memory capacity (Blue Waters still has more memory than any other system in the open research space) and storage bytes. While such high performance resources are necessary for success, they are not sufficient for effectiveness. In other sections of this report, you will hear about our expanded Petascale Application Improvement Discovery (PAID) program, where the project is providing millions of dollars to science teams and computational and data experts to improve the performance of applications (in a measurable manner). The interim report documents the fact that we have been able to help some teams achieve orders of magnitude improvement in productivity and time to solution, otherwise known as performance improvements.

Likewise, the project is proud to play a role educating and developing the next generation extreme scale workforce, through our workshops, symposium, graduate fellowship, undergraduate internships, the original and evolving Virtual School for Computational Science and Engineering, our funding for the HPC University and our training workshops and allocations.

While the scientific insights gained with support from Blue Waters are perhaps the most exciting way to judge the project's impact, there are ways to consider its titanic role in the open-science community. The project has provided over 15 billion core-hour equivalents in its first three years. The geographic reach of the project is also extensive, with 1,980 science and engineering partners from 238 institutions and organizations in 43 states and Puerto Rico, along with 16 countries. Adding our educational and workforce development efforts, Blue Waters has reached investigators and students in 48 of the states plus Puerto Rico and countries around the world.

Indeed, I am grateful that in collaboration with our science team leaders, our Science and Engineering Team Advisory Committee (SETAC) and many others, and others, we are seeing many in Washington and around the country acknowledge the critical role Blue Waters (and other @scale cyber-infrastructure) play in accelerating discovery and competitiveness across the US and that NSF appears to be planning next generation programs for Track-1 level systems for the Open Science Community. This is an important step forward for the entire science and research community.

In support of this effort, I am happy to report to you that the Blue Waters project has proposed and NSF has agreed, to **extend the full service operations of Blue Waters for at least 1 more year, through March 2019**, at no additional cost to our stakeholders and funders.

The extension of our Full Service came about through effective and deliberate project management to strive to meet and exceed all the explicit and implicit expectations we have, all in cost effective manners. As part of this approach, we believe to manage well and improve, we have to objectively measure what we are doing. You will see three sections in this report that discuss our external and internal assessments of the quality of services and programs we provide. The first is a report from our external iSTEM evaluation team for our educational and workforce development programs. The second is a short summary of our control and operational metrics reports that look at the overall set of system services and resources. The third section is the annual report from our SETAC. I am happy to report these and other assessments show that, while the systems, services and programs we provide are already world class, the Blue Waters project is learning and improving to have even more positive impact.

It is also important that I point out the tremendous support and guidance we receive from the entire National Science Foundation, in particular Irene Qualters, Edward Walker and James Kurose. Likewise, the project benefits from the support and guidance from the State of Illinois, the University of Illinois and the National Center for Supercomputing Applications. I particular, I would like to thank Peter Schiffer, Melanie Loots, Edward Seidel, William Gropp, Wen-mei Hwu and Marc Snir, Brett Bode, Gregory Bauer, Jeremy Enos, Cristina Beldica, and Kjellrun Olson for their support and expertise in making the Blue Waters Project successful.

Finally, I take great pride in the Blue Waters project and in the outstanding teams that makes Blue Waters an outstanding resource. Talented, dedicated individuals at the National Center for Supercomputing Applications and our partner institutions strive night and day to maintain the system, support our science and engineering partners in achieving breakthrough research, improve tools and processes, and train the next generation of computational researchers. This publication celebrates their accomplishments as well as those of the researchers, educators, and students they support.

With great respect and thanks
to be working with everyone,

Dr. William T.C. Kramer
Blue Waters Project Director and Principal Investigator
Research Professor of Computing Science

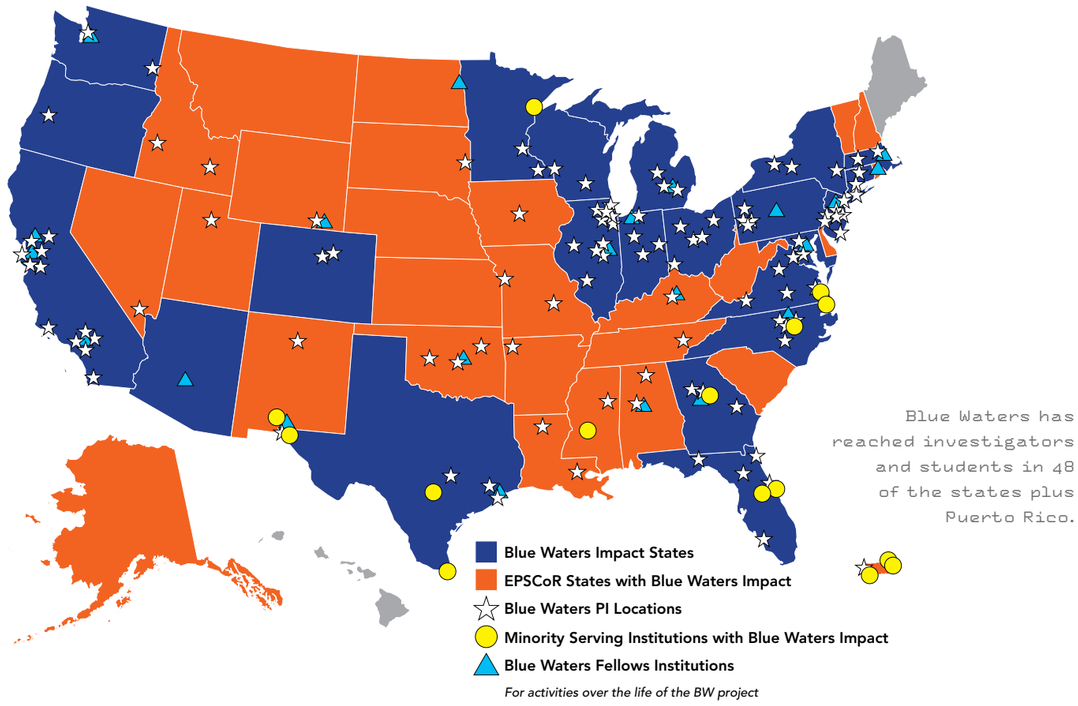

Blue Waters has reached investigators and students in 48 of the states plus Puerto Rico.

PETASCALE APPLICATION IMPROVEMENT DISCOVERY PROGRAM

Advanced hardware is just one step in realizing breakthrough science and engineering research. Applications also must be adjusted to take full advantage of enormous computing systems such as Blue Waters. That's why NCSA and the University of Illinois have undertaken a comprehensive program to provide additional services and support to assist science and engineering teams in enhancing the functionality and performance of applications.

The Petascale Application Improvement Discovery (PAID) program aims to facilitate the creation of new methods and approaches that will dramatically improve the ability to achieve sustained science on petascale systems. It also targets the introduction of new, fundamental application approaches in addition to optimization of current applications for Blue Waters, to increase the knowledge of and use of "best practices" for highly scalable computing and data analysis. The program provides funds for Improvement Method Enabler (IME) teams to work with science and engineering research teams to help create and implement technologies that improve application performance, and provides science and engineering teams the funds to work with the IMEs.

The following brief reports describe the progress made over the past year by the IMEs and the research teams with which they collaborate closely.

TOPOLOGY AWARE / COMMUNICATION PATTERN DETECTION

IME Team leader: Laxmikant (Sanjay) Kale, University of Illinois at Urbana-Champaign

Research team PIs: Scott Nobel (astrophysics), Klaus Schulten (biophysics); Paul Woodward (stellar hydrodynamics), P.K. Yeung (turbulence)

The main goal of this IME effort is to improve performance by optimizing the use of system resources, specifically Blue Waters' compute nodes (CPU and GPU) and interconnect. Through the PAID program, we have collaborated with the developers of scientific applications, learned about the research team's practices, and transmitted knowledge and awareness of new approaches to improve the efficiency and scalability of their applications on Blue Waters.

The two main approaches being used are: (a) job and task placement in Blue Waters to optimize communication; and (b) dynamic load balancing within the application to efficiently use compute resources. Both approaches are aimed at efficient use of hardware resources during job execution, with the result that applications run faster and have better scaling performance. As part of the PAID program, software and technologies are developed that facilitate the adoption of these techniques.

Due to the topology of the interconnect, optimizing job and task placement is a major factor affecting communication performance of applications on Blue Waters. We have developed tools that automatically profile and classify applications based on their communication pattern, to suggest optimal parameters for job placement and for task placement within a job. Using these, we have seen improvements in several scientific applications, each exhibiting different communication patterns. We applied these tools to production-ready applications like the quantum chromodynamics code MILC, achieving a **speedup of more than 2x**, and the molecular dynamics application Qbox, with which we have seen up to **57% improvement** in execution time compared to not using the tools. We are collaborating with another Blue Waters research team to improve the performance of their turbulence code and have improved performance up to 12% in small-sized jobs, with further improvement expected for larger jobs. We plan to package these tools for Blue Waters users and to use them in more applications, such as the CESM climate model.

Load imbalance in parallel applications, whereby the work is not uniformly distributed among compute resources, leads to underutilized resources and increasing time to solution. Dynamic load balancing is an effective technique to address this issue. Charm++ is a parallel programming framework developed over many years, with dynamic load balancing capabilities. Production-ready applications based on Charm++ on Blue Waters include OpenAtom, NAMD, and ChaNGa. Recently, we increased the scalability limit of NAMD by 2x in core count in a collective variable scenario using intra-node load balancing based on Charm++ and OpenMP interoperation, leading to an overall **2x speedup**.

In addition, we have developed a general-purpose library to assist in load balancing decisions for developers of MPI applications. Lastly, we provide support to run MPI applications on AMPI (MPI implementation on top of Charm++), thus obtaining the benefit of Charm++ features, including load balancing. Other important activities include research and development on heterogeneous systems (nodes with CPU and GPU) and techniques to balance load between GPUs and CPUs. These are being incorporated into Charm++ for NAMD, ChaNGa, and other Charm++ applications, and will also be part of the general-purpose load balancing library.

SPIRAL FFT

IME Team Lead: Franz Franchetti, Carnegie Mellon University

Research team PIs: P.K. Yeung (turbulence), Gerhard Klimeck (nanomaterials)

Spiral is an automatic program generation system that generates platform-tuned libraries of signal processing transforms, such as the discrete Fourier transform, discrete cosine transform, and many others. SPIRAL addresses one of the key problems in numerical software and hardware development: how to achieve close to optimal performance with reasonable coding effort. The Spiral PAID project focuses on developing Spiral for several research teams, most notably the group led by Georgia Tech's P.K. Yeung.

The Yeung group studies the complexities of turbulent fluid flow using pseudo-spectral direct numerical simulation (PSDNS). Using the computed fluid velocity fields from PSDNS, they study the advection of passive particles through the flow using a compact finite difference (CCD3D) scheme. This combined PSDNS and CCD3D solution has bearing on physical phenomena such as climatology and chemical mixing.

PSDNS relies on 3D fast Fourier transforms (FFTs). The most computationally expensive portion of the 3D FFT is the network communication required to perform the global transpose necessary

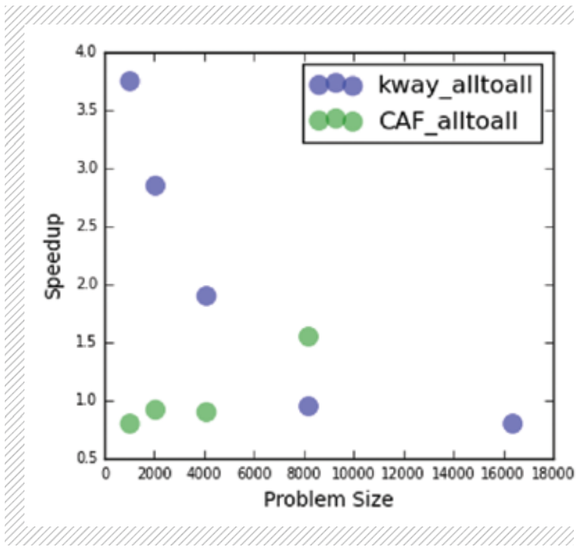

FIGURE 1: Comparison of *kway_alltoall* single-core speedup versus *MPI_alltoall* (Speedup=1) and *CAF_alltoall*. The *kway_alltoall* method uses a k-way recursive algorithm where the optimal k-way value is selected empirically. The kway_alltoall method provides significant performance improvement for the range of problem sizes of interest to the Yeung group's PSDNS and CCD3D routines.

when using the pencil decomposition method. This global transpose is typically accomplished using the vendor-supplied *MPI_alltoall* function.

A new library, ACCFFT, implements a different *kway_alltoall* function which is based on a k-way recursive Bruck algorithm. The value of k-way, and the k-way transpose method, must be chosen carefully based on the platform's network topology, scheduler, and other factors. Spiral helped to optimize and port the ACCFFT library for Blue Waters. An interface was created so Fortran applications (like Yeung's PSDNS and CCD3D) could use the library. Furthermore, the improved *kway_alltoall* function was exposed directly for use in the Yeung PSDNS and CCD3D routines. Optimized values of the k-way parameter are determined initially at program startup. Figure 1 shows the performance improvement (speedup, as compared to *MPI_alltoall*) of the *kway_alltoall* and *CAF_alltoall*, a co-array Fortran (CAF) optimized version provided by Cray that works well at the largest problem sizes of interest.

The Spiral team has turned attention to the CCD3D, focusing first on the finite difference stencil used in CCD3D and building on our previous work on stencils.89 Stencil computations involve arithmetic operations on physically contiguous data elements, e.g., c0*(A[i-1]+A[i]+A[i+1]). The methods for optimization rely on loop transformations to enhance data locality and enable load-balanced parallelism, combined with data layout transformations. A demonstration of the performance improvement of these techniques is shown in Figure 2 for a single core on a Blue Waters XE node (AMD 6276 Interlagos processor). Larger performance improvements are being planned by applying the same stencil optimization strategies to the rest of the CCD3D routine.

MODEL-BASED CODE REFACTORING AND AUTOTUNING

IME Team leader: Mary Hall, University of Utah

Research team PIs: Thomas Cheatham (biophysics)

The Composable High-Level Loop Optimization (CHiLL) framework is an autotuning compiler that transforms loop nest computations to produce performance-portable optimized code; a thin layer that generates CUDA code is called CUDA-CHiLL. CHiLL exploits autotuning to explore different implementations of a computation in a systematic way and identify the best-performing implementation. The optimization strategies can be exposed to the user through transformation recipes, which permits expert developers to guide optimization but still retain high-level code that is portable across optimizations. CHiLL has been used to optimize critical computations in a variety of applications.

CHiLL and CUDA-CHiLL were installed on Blue Waters in Fall 2015. In June 2016, a new version of CHiLL was installed on Blue Waters, which separates the abstract syntax tree (AST) from a generic CHiLL AST. This new version also includes new domain-specific optimizations for sparse matrices, stencils and geometric multigrid.

An important value of the CHiLL PAID project is to assist science teams that use codes that have not been ported to GPUs or that have significant node-level performance issues on Blue Waters.

Since Fall 2015, we have been collaborating with Thomas Cheatham and Daniel Roe, University of Utah, on the CPPTRAJ analysis code for the AMBER molecular dynamics package. CPPTRAJ is MPI-parallelized to handle ensembles of simulations and makes use of OpenMP parallelization for certain computationally demanding tasks. Using OpenMP, scaling is limited to small numbers of threads (e.g., four). A desirable mode of AMBER operation would use shorter runs with immediate analysis to decide where to steer the calculations. Therefore, it would be advantageous for CPPTRAJ to take advantage of the GPU nodes already in use for the molecular dynamics simulations.

We applied CUDA-CHiLL to generate CUDA code for the "Action" kernels of CPPTRAJ, which required some manual changes to the code. The main "Action Closest" kernel was distilled into six core kernels by removing embedded divergent code. We measured the optimized CUDA kernels' performance on Blue Waters; we also integrated them back into the CPPTRAJ code on an input of 15,022 to 4,143 molecules across 2,000 frames. While the individual kernels demonstrated varying speedups over sequential versions ranging from less than **3x up to 278x**, not including copying overheads the overall CPPTRAJ speedup was **6.7x for an XK node** compared to a fully loaded single XE node.

The PAID program and our work with the AMBER team offer the opportunity to better understand the needs of application scientists and identify limitations in our tool. Constructs in CPPTRAJ, including C++ iterators, vector libraries, and structures-of-arrays, were not supported by CUDA-CHiLL, and we manually modified the code to eliminate these. As an outgrowth of this program, we have identified future research issues for CUDA-CHiLL.

PARALLEL I/O PERFORMANCE

IME Team Lead: William Gropp, University of Illinois at Urbana-Champaign

Research team PIs: Tiziana DiMatteo (astrophysics), Lijun Liu (geodynamics), Robert Sugar (high-energy physics), Daniel Bodony (engineering)

File I/O is often a performance problem for applications, and one that is often overlooked. This project takes advantage of system-level tools to identify potential performance problems, combined with deeper investigation into specific applications to address I/O performance problems. We have worked with two teams to develop I/O test cases, established benchmarks for achievable performance, and developed an easy-to-use I/O library for regular meshes that has significantly improved the I/O performance of one application.

We are testing the lattice quantum chromodynamics code used by the MIMD Lattice Computation (MILC) collaboration on a small data set in order to understand how it works. We are able to build a working version from the GitHub source provided by the MILC team. We have determined that one routine appears to be responsible for slow I/O performance. This code shuffles data between nodes to reshape the data before writing to disk and does not exploit code in MILC that is used to reduce the number of processes accessing the file when the number of processes greatly exceeds the stripe count (and hence the parallelism) of the file. The MeshIO library that we have developed and tested on another application may be effective here, and we will be testing that approach.

We have worked with the BlueTides astrophysics team to understand their I/O needs, including the workflow into which the application fits. The current system uses an interface called “bigfile.” A simulator for this has been developed by the science team and released to our project; the simulator is available on GitHub. We are using this code to better understand the research team’s needs.

We are at the beginning of our work with the research team led by Lijun Liu on the CitcomS finite element code, which is designed to solve compressible thermochemical convection problems. Using the Darshan logs from the Darshan I/O characterization tool collected by Blue Waters, we have performed an initial assessment and will determine next steps in collaboration with the research team.

In addition, the code PlasComCM, used by the University of Illinois engineering research team led by Daniel Bodony, was observed to have an I/O performance problem. The original code required about 90 minutes on Blue Waters to sequentially read a file of around 1GB (the performance was better on the system on which this code was often used, possibly due to a different parallel file system configuration). We developed a new I/O library designed for regular meshes, Mesh_IO, and integrated it into the PlasComCM code. The time to read the same file was reduced to 240 seconds using a grouped read strategy and further reduced using MPI-IO to 3 seconds removing file IO as a performance bottleneck. The PlasComCM team also is testing the output routines in our new library and plans to adopt these new routines. In addition, this project provided a good opportunity to test the use of collective MPI-IO, and we expect to apply these routines to some of the codes used by PRAC teams, including the MILC code.

BEST PRACTICE IDENTIFICATION, DISSEMINATION, AND IMPLEMENTATION

IME Team Leaders: William Tang, Princeton University

Research team PIs: Nikolai Pogorelov (heliophysics), Jerry Draayer (astrophysics)

The purpose of this PAID project is to provide relevant and useful information to the Blue Waters science application community. A considerable numbers of key experiences and “lessons learned” gained from the successful development of GTC-P, a highly scalable particle-in-cell code for kinetic plasma simulations on a variety of modern supercomputing platforms, are leveraged for this activity. In particular, expertise is provided for developing particle-based application codes (or codes that share similar features) on large modern GPU/CPU hybrid systems such as Blue Waters by using both CUDA and compiler directives such as OpenACC. In addition, we examine and balance the trade-off of portability vs. speedup in choosing between CUDA and OpenACC in GPU deployment.

During the first year of the PAID program, our team worked with two research teams: the heliophysics group led by Nikolai Pogorelov from the University of Alabama-Huntsville, which uses

MS-FLUKSS, and the astrophysics group led by Louisiana State University's Jerry Draayer, which uses LSU3Shell.

MS-FLUKSS is a fluid and kinetic hybrid code, where the fluid portion is solved by the Chombo package using finite volume methods and the kinetic part is solved by an "in-house" Fortran code using Monte Carlos methods. Our team first identified and separated the kinetic code to be optimized on multi-core CPUs and ported it to many-core GPUs. The initial code investigation also identified the code's performance challenges, including fine-grained data hazards, poor data locality, low arithmetic intensity, and load imbalance. In order to address these challenges, a new multi-threading implementation based on loop-level OpenMP was developed. Compared with the original implementation, a **speedup of up to 2.89x** has been obtained in scaling up to eight threads in a single XK node on Blue Waters. Another key accomplishment includes developing the **first** OpenACC version of the MS-FLUKSS code for deployment on GPUs.

Advances on the LSU3Shell project included delivery of a new matrix construction algorithm and the redesign and implementation of associated data structures. In particular, the new implementation leverages data locality by a smart reuse of computationally expensive data that results in a dramatic **speedup of up to 10,000x** compared to the execution time of the legacy code on the same hardware for the matrix multiple kernel. The code was then further developed for multi-core CPUs using OpenMP directives and many-core GPUs using CUDA.

EFFECTIVE USE OF ACCELERATORS/ HIGHLY PARALLEL HETEROGENEOUS UNITS

IME Team leader: Wen-mei Hwu, University of Illinois at Urbana-Champaign

Research team PIs: Thomas Jordan (seismology), Thomas Quinn (astrophysics), Paul Woodward (stellar hydrodynamics)

The goal of this project is to assist science teams in effectively exploiting Blue Waters' GPU accelerators to reduce time-to-science. In our initial engagement with a science team, the existing application code is reviewed to identify opportunities for acceleration. Then the IME and the research team collaboratively develop a holistic plan for algorithm transformation, GPU kernel development and optimization, and application of performance tools to enable the application to solve larger problems faster. This year we have worked directly with two science teams and collaborated with another IME to work with another science team.

Thomas Jordan and Yifeng Cui at the University of Southern California and the Southern California Earthquake Center have developed AWP-ODC, a wave propagation simulation code that uses GPUs for several time-consuming tasks. This year, the scientific capabilities of the code were broadened to include new frequency-dependent attenuation features and a new numerical algorithm for kernel plasticity, which resulted in **more accurate prediction** of long-period ground motions in the Los Angeles basin during earthquakes.

We optimized three existing GPU kernels by employing texture memory, shared-memory tiling, and register queuing to better use the GPU memory hierarchy, alongside kernel launch and compiler directives to improve GPU utilization. CUDA streams were used to exploit communication/computation overlap. These modifications yield an application-level speedup of **greater than 1.3x** for small cases, which increases as problem size grows. Weak scaling was evaluated on Blue Waters up to the petaflop scale, with parallel efficiency ranging from 80% to 100%.

Thomas Quinn at the University of Washington has developed ChaNGa (Charm N-body GrAvity solver) to perform collisionless N-body simulations. At our initial engagement with the Quinn group, two GPU kernels already existed for time-consuming tasks. These kernels were optimized for the Fermi GPU architecture, and profiling revealed that they underused Blue Waters' newer Kepler GPU architecture.

We worked to re-optimize those kernels for the Kepler GPU architecture by reducing shared-memory use and the associated thread synchronization and using compiler directives to control the resulting increased register usage. The kernels were also modified to use CUDA streams and to take advantage of other newer CUDA features. The two kernels were optimized to produce final **speedups of 2.1x and 1.6x** respectively. Additionally, the application was modified to execute several computational kernels simultaneously, achieving 88% of the peak GPU instruction throughput and a total **application speedup of 2.5x**.

GLOBUS ONLINE

IME Team leader: Ian Foster, Argonne National Laboratory

Research teams: All

A growing number of research problems require computation over vast amounts of data. Blue Waters provides the compute capability but a challenge remains in moving data to and from the supercomputer, and making that data accessible to collaborators at multiple institutions. Combining this computational power with the reliable, secure file transfer and sharing capabilities of Globus makes Blue Water a unique system for enabling largescale science. The integrated solution obviates the need for researchers to build their own data management infrastructure, allowing them instead to focus on core research and easily leverage the advanced capabilities that Blue Waters provides.

The goal for the Globus PAID project is to provide simplified data management capabilities to all of Blue Water's scientific researchers.

Globus is cloudbased research data management software as a service (SaaS) and platform as a service (PaaS) that is well integrated into Blue Water's data management platform. Globus provides high-performance data transfer and management on Blue Water's Lustre (NCSA#BlueWater) and HPSS nearline (NCSA#Nearline) storage systems. From July 2015 through June 2016, 394 users representing more than 100 institutions performed file transfers on NCSA#BlueWaters, transferring more than 18.5 petabytes. From July 2015 through June 2016, 180 users representing more than 50 institutions performed file transfers on NCSA#Nearline, transferring more than 13 petabytes.

In the past year, improvements were released to the Globus service based on feedback from the Blue Waters staff. In particular, enhancements to the management console interface to support local usernames and pause/resume of tasks were added. This allows Blue Water administrators to manage their endpoint using the familiar local usernames instead of Globus usernames and to manage use of their endpoints.

FIGURE 1: Amount of data transferred per month for NCSA#BlueWaters and NCSA#Nearline.

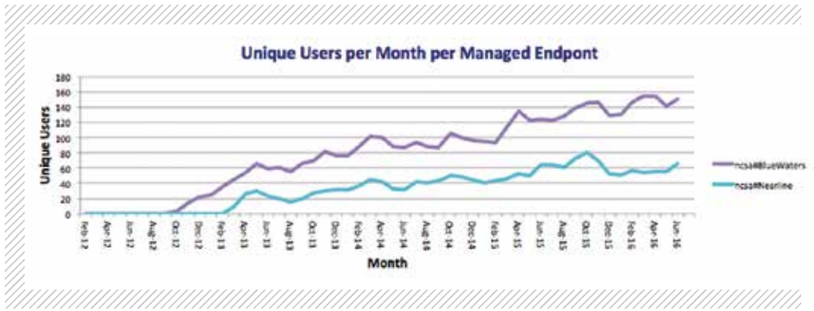

FIGURE 2: The number of unique users per month for NCSA#BlueWaters and NCSA#Nearline.

HDF5

IME Team leader: Gerd Heberk, The HDF Group

Research team PIs: Robert Rauber (climate science), Lijun Liu (geodynamics), P.K. Yeung (turbulence), Leigh Orf (atmospheric science)

In the words of Ken Batcher, "A supercomputer is a device for turning compute-bound problems into I/O-bound problems." If we accept this "definition" then The HDF Group's main task in this PAID project is to assist users whose applications undergo the inevitable transition from being compute-bound to being I/O-bound to: 1) make effective use of the HDF5 library and format; and 2) maximize utilization of the I/O subsystem on Blue Waters.

In practice, our supporting activities include:

- Analyzing and developing I/O strategy
- Designing HDF5 profiles (the data organization and layout in one or across multiple HDF5 files)
- Creating baseline and I/O kernels, performance tuning, and benchmarking
- Performance debugging of high-level (on top of HDF5) I/O libraries
- Conducting tutorials at Blue Waters events and presenting information on new HDF features and releases

The HDF team initiated collaboration efforts with Jim Edwards from the climate modeling team that developed and maintains PIO, an I/O interface library used by CESM and ACME models leveraging NetCDF4 or PNetCDF. NetCDF4, in turn, uses the HDF5 library for I/O, whereas PNetCDF uses the NetCDF3 file format.

The HDF team set up PIO and replicated preliminary tests that showed some of the performance issues encountered by the PIO team. The team did a detailed analysis to establish the causes of the poor NetCDF4 write performance in PIO with the subset method, and reported its findings and recommendations.

Leigh Orf is co-PI of a science team studying the inner workings of supercell thunderstorms that produce long-track tornadoes. The team generates large amounts of simulation data in HDF5 files and uses several advanced features of the HDF5 library, including the ability to perform HDF5 I/O in memory. The HDF Group helped the Orf's team identify the cause of I/O load imbalances that lead to significant performance degradation in large-scale runs. The analysis suggested that a new buffer pre-allocation function would benefit this team and many other users who have adopted a similar I/O strategy. This API is being implemented and will be included in the next HDF5 maintenance release.

On March 31, 2016, The HDF Group reached a milestone with the release of HDF5 1.10.0. After describing some of the new features at the Blue Waters Advanced Users Workshop, we were approached by the science team lead by P.K. Yeung, who studies the complexities of high-Reynolds-number turbulence. A new HDF5 1.10 feature, the so-called Virtual Datasets (VDS), caught the team's attention because it promised to substantially simplify the handling of multiple HDF5 files in their I/O code. Since the new release was just hot-off-the-press, Cray Inc. did not yet provide pre-compiled HDF5 1.10 binaries on Blue Waters. The HDF Group built, tested, and provided instructions on how to deploy the software on Blue Waters, and assisted the team with integrating VDS into their code.

It is not unusual that science teams focus initially on their computational challenges, and, after having overcome those challenges, discover that, as Batcher predicted, their "pain point" has shifted into the realm of I/O performance, just as they are trying to scale their applications to large portions of systems such as Blue Waters. This is an uncomfortable position because high-performance I/O is not a mere afterthought, but the result of a deliberate effort. Education and outreach, early intervention, and smarter tools might help to change attitudes in the long run.

The participation in the PAID program is precious to The HDF Group. It is a unique opportunity to grow our use base, to see up close how science teams use our software, and to identify opportunities for improving the software and its documentation that benefit the Blue Waters research teams, but also the entire high performance computational community.

EVALUATING THE EFFECTIVENESS AND IMPACTS OF BLUE WATERS' EDUCATION, OUTREACH AND WORKFORCE DEVELOPMENT ACTIVITIES

Education, outreach, training and workforce development activities are integral to the Blue Waters project. In order to effectively shape these activities and to assess their impact, Blue Waters employs a robust evaluation process of internal and external assessments. The ultimate goal of this effort is to validate and document the effectiveness of the Blue Waters project's community and user engagement programs and to disseminate findings and best practices through publication and presentation. The external evaluation is led by Professor Lizanne DeStefano, the executive director of the Center for Education Integrating Science, Mathematics, and Computing (CEISMC) as well as an associate dean and professor of psychology at Georgia Tech.

The evaluation team uses a widely accepted Values-Engaged Educative Approach (VEE). Developed with NSF support, the VEE approach defines high-quality STEM educational programming as that which effectively incorporates cutting-edge scientific content, strong instructional pedagogy and sensitivity to diversity and equity issues. In the VEE approach, evaluators work closely with program implementers to promote their understanding of program theory, implementation, and impact.

The evaluations are designed to answer four questions:

1. **Implementation:** Are Blue Waters community and user engagement programs being implemented on schedule and as planned?
2. **Effectiveness:** Are key components of the Blue Waters project (e.g. the Blue Waters Graduate Fellows program, Blue Waters Undergraduate Internship program and the Petascale Institute for undergraduate interns, Blue Waters Symposium, webinars, workshops, Virtual School of Computational Science and Engineering courses, user support services, etc.) operating effectively? How might they be improved?
3. **Impact:** What outcomes (e.g. increased productivity, gains in scientific knowledge, improved technical skills, employment outcomes, etc.) are associated with participation in the Blue Waters community and user engagement programs? How does impact vary across groups? What are the longitudinal trends (e.g. student progress through degrees and professional employment, start-up accounts becoming larger allocations, researchers becoming PIs, etc.)?
4. **Sustainability:** How and to what extent are elements of the Blue Waters community and user engagement programs becoming institutionalized

(such as by attracting long-term support from government agencies, educational institutions, or industry) to ensure sustainability of program components? What opportunities and barriers exist?

The external evaluation team analyzes the annual Blue Waters Symposium, advanced user workshops, the graduate fellowship program, the undergraduate internship program, and other project activities. Their findings are regularly reported back to the Blue Waters project team to enable the project to improve program implementation and outcomes. Typically the results are shared soon after each event completes or quarterly for the programs that span a year or more. The Blue Waters Project has made a number of improvements in our programs based on this feedback.

BLUE WATERS SYMPOSIUM

The annual Blue Waters Symposium brings together leaders in Petascale computational science and engineering to enable these researchers and technologists to share successes and challenges in large-scale heterogeneous computing. The symposium includes presentations from the Blue Waters science teams and keynote addresses from innovative thinkers in science and technologies. The evaluators assess the annual symposium by making observations during the event, conducting a post-event survey of attendees, and interviewing some attendees.

Findings from both 2015 and 2016 indicate that participants were highly satisfied with the overall experience at the Symposium. When asked about the event's greatest strength, respondents mentioned the diversity of topics, methods, and interests (30%), the strong content of the presentations (21%), and the opportunities for networking and discussion (20%).

ADVANCED USER WORKSHOP

Blue Waters provides a variety of training and support opportunities to its user partners. Among these was an Advanced User Workshop held in March 2016. This three-day hands-on workshop encouraged participants to "bring their own code" and included representatives from Cray, Allinea, PGI, and NVIDIA. A post-workshop survey administered by the evaluation team showed that most participants rated the workshop's activities as "valuable" or "very valuable." All respondents also "agreed" or "strongly agreed" with the following statements:

- My primary goal for attending this workshop was achieved.
- I am confident that I can incorporate this information into my work.
- My overall experience met my expectations.
- This workshop was well organized.
- I enjoyed the format of the activities.
- The sessions provided the appropriate amount of depth.
- Overall, I would rate my experience as successful.

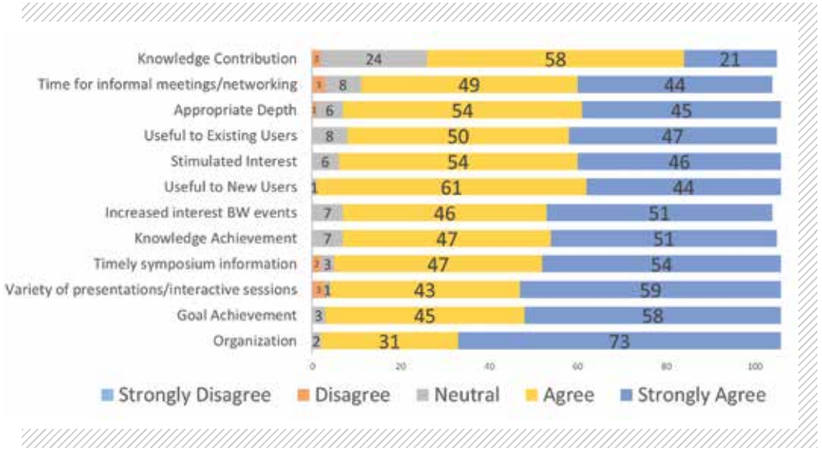

FIGURE 1: After the 2016 Blue Waters Symposium, participants reported a high level of satisfaction with the event. While the 2015 Symposium received a strong satisfaction score from attendees, the 2016 Symposium was rated even more highly.

GRADUATE FELLOWSHIP & UNDERGRADUATE INTERNSHIP PROGRAMS

The competitive Blue Waters Graduate Fellowship program awards elite PhD students who have decided to use high performance computing and analysis in their research. The Fellows are from institutions across the country and from a diverse range of science and engineering domains with a year of substantial financial support and an significant allocation on Blue Waters, in order to advance their computational research. Blue Waters allocations can continue for Fellows even after their fellowship ends as long as they are continuing their education and making progress. Each graduate fellow is paired with a member of the Blue Waters Science and Engineering Application Support (SEAS) experts; these "point of contact" personnel work closely with the fellows to understand their research goals and needs and to help them port and tune their codes as well as with any standard support requests the fellows may have. The program also includes monthly virtual group meetings woth all the Fellows, their POCs , their fellowship coordinator and other Blue Waters staff.

The evaluators solicit feedback from both the fellows themselves and from their research advisors, conducting focus groups, surveys, and interviews to assess the fellowship program.

The undergraduate internship program engages over 20 undergraduate students from across the country in petascale computing research and development projects; the internship includes an intensive two-week workshop covering the fundamentals of high-performance computing, followed by a year-long project mentored by an established faculty researcher and project coordinators.

The evaluation team uses focus groups, surveys, and monthly reports to assess the internship program.

For 2015, the evaluators found that both the fellowship and internship programs positively impacted students' research progress and career plans, and that faculty advisors believed that the programs were effective in providing the students with a positive research experience and opportunities for interactions. In fact, the evaluators comments that the Fellowship program was among the best they are aware of because of the in depth support and working assistance from the Blue Waters staff.

Among the comments from graduate fellows:

"I don't think I would've been able to work in power systems which was really interesting, and this allowed me to not only work with my advisor but another professor in this area. So for me, this has made a huge change in my graduate path."

"Up until this point, I had really only worked with computing on more like the computer cluster scale, and I didn't really realize some of the limitations of that and some of the opportunities and challenges associated with high-performance computing and how that can really make an impact in a study."

"I think that will really help me, not only set up different types of simulations, but also develop the better tools…I think that has helped me gain a broader base. So that'll be nice for further accentuating with a post doc."

Among the comments from undergraduate interns:

"I really hadn't considered a field that I would want to go into until I was a part of this internship, and now I'm declaring a double major in that field, so that's been real life-changing in a way."

"I am so thankful that I got to participate with this. It was a wonderful learning experience, and I enjoyed getting to work with a mentor and getting to use Blue Waters for my work."

"The Blue Waters Internship Program was a life-changing event. The program has given me the experience necessary to work in HPC after graduation. In addition, I have made extremely valuable connections in the supercomputing and HPC world that I know will help me in the future."

SPACE SCIENCE

ASTRONOMY

ASTROPHYSICS

COSMOLOGY

HELIOPHYSICS

VISUALIZATION FOR PUBLIC OUTREACH

FIGURE 1:
Cosmological simulation from John Wise, soon after the first generation of stars appeared. Filaments of dense cool gas are blue-white; gas heated and ionized by ultraviolet light from new stars is orange-to-white; clouds of heavier elements synthesized in early stars and released when they explode as supernovae are paler yellow; clusters of young stars are white dots.

Allocation: Illinois/400 Knh
PI: Donna Cox[1]
Collaborators: John Wise[2], Brian O'Shea[3], Michael Norman[4], Leigh Orf[5]

[1] University of Illinois at Urbana-Champaign
[2] Georgia Institute of Technology
[3] Michigan State University
[4] University of California at San Diego
[5] Central Michigan University

EXECUTIVE SUMMARY

CADENS (The Centrality of Advanced Digitally ENabled Science) is a National Science Foundation-supported project to increase digital literacy and inform the public about computational and data-enabled scientific discovery. CADENS invites scientists in diverse fields to submit their data to the artists and technologists of the Advanced Visualization Laboratory, who then create data-driven visualizations of scientific phenomena for dome productions and documentaries that are seen by hundreds of thousands of people. In 2015-16 the CADENS team used Blue Waters in processing data and creating animations for two major productions. "Solar Superstorms" is a full-dome planetarium documentary featuring simulations on multiple scales of solar magnetism and its influence on Earth and of early-universe star formation. The flat-screen production "Super Tornado: Anatomy of a Mega-Disaster" looked at two storms that created multiple strong tornadoes.

Work is ongoing on a third production, aimed at understanding what the Dark Energy Survey and Large Synoptic Survey Telescope will reveal, with Blue Waters being used to analyze two early-universe cosmological simulations from Michael Norman and Brian O'Shea.

INTRODUCTION

Digital tools and methods are increasingly central to science, engineering, and scholarly research, but many people still envision research as exclusively using analog tools such as microscopes and telescopes. The NSF-supported CADENS project aims to increase public understanding of the crucial role digital tools and techniques play in research across a wide range of fields, while also sharing the latest research results with wide audiences.

METHODS & RESULTS

In order to bring cutting-edge scientific research to life for large audiences of all ages, we create clear and visually appealing animations from large physical simulation data provided by our collaborators. The CADENS team visualized the first generations of stars ionizing the early universe (simulations conducted on Blue Waters by John Wise and by Michael Norman and Brian O'Shea), a solar storm impacting the Earth's magnetosphere (simulation by Homa Karimabadi et al.), and a system of tornadoes that struck El Reno, Oklahoma, in 2011 (simulation conducted on Blue Waters by Leigh Orf).

For the cosmological simulations, we used the astronomical analysis/visualization package yt [1], which supports both OpenMP and MPI parallelism. We relied on a then-experimental version of yt's volume-rendering feature, which was being reinvented at the time, and are grateful to Matthew

Turk, Sam Skillman, and other yt developers for their critical assistance. We created full-dome imagery of several quantities (gas density, temperature, metallicity, stars) and combined them.

With Karimabadi's magnetosphere simulation, we again could use yt, this time to trace magnetic field lines perturbed by the incoming solar coronal mass ejection.

We have also experimented with running the commercial animation software Houdini on Blue Waters. Houdini is a backbone of our image production system, with a high-quality renderer, a rich set of data-flow-based processing operators configured via a graphical interface, and our own group's scientific data reader plugin. We currently use Houdini on a dedicated visualization cluster, but it would be valuable to allow Houdini rendering to operate with Blue Waters data directly. Initial results were promising, and we hope to further develop this capability.

WHY BLUE WATERS

We work with simulation data from collaborators who compute on Blue Waters. The resultant datasets are often so large that they can't readily be moved elsewhere.

Blue Waters staff were very helpful when technical problems arose, whether with negotiating the system-installed Python, getting usable performance reading large memory-mapped files, or acquiring statistical data on jobs we had run.

NEXT GENERATION WORK

As CADENS and future public outreach projects continue, we anticipate working with science collaborators who will use larger computing environments to solve larger problems. There will be a need to visualize their results, and we expect to need a future Track-1 system for at least some of that work.

PUBLICATIONS AND DATA SETS

Super Tornado: Anatomy of a MegaDisaster, available via Amazon Prime, Hulu, IMDB, and YouTube.

Solar Superstorms, available in fulldome format for museums and science centers via Spitz Inc.

Solar Superstorms, alternate expanded flat-screen version broadcast by TVF International.

Solar Superstorms, available via Amazon.com, Amazon Prime, DirectTV, Hulu, and YouTube.

Rivers of Energy Inside the Sun 4K, YouTube.

When A Solar Storm Engulfs Earth, Hulu and YouTube.

FIGURE 2:
Visualization of Leigh Orf's simulation of the tornado that struck El Reno, Oklahoma, on May 24, 2011, showing cloud water, cloud ice, and cloud rain. This was an attempt at a realistic treatment of the data, since these variables are what would be visible to the naked eye when watching the tornado. To indicate the storm's scale, the ground floor is a grid of 1km squares.

HIGH-RESOLUTION SIMULATIONS OF MAGNETOHYDRODYNAMIC TURBULENCE IN ACCRETION DISKS

Allocation: Illinois/440 Knh
PI: Charles F. Gammie[1]
Co-PI: Ben Ryan[1], Sebastien Fromang[2], Pierre Kestener[3]

[1]University of Illinois at Urbana-Champaign
[2]Service d'Astrophysique
[3]CEA/IRFU

EXECUTIVE SUMMARY

Accretion disks are rings of gas that orbit black holes, white dwarfs, neutron stars, young stars, and planets. Disks lie at the heart of many important astrophysical problems, including planet formation and the central engine of quasars. Disk evolution is likely driven by turbulent diffusion of angular momentum. We studied turbulence in differentially rotating, magnetized disks at **unprecedented** resolution using a graphics processing unit (GPU)-accelerated code. Our goal was to decide between earlier, contradictory claims about convergence of disk turbulence. We found, surprisingly, that the intensity of turbulence does not converge.

INTRODUCTION

Accretion disks are rings of gas that orbit black holes, white dwarfs, neutron stars, and young stars and planets. The release of gravitational energy in disks around supermassive black holes powers the quasars, luminous objects in galactic nuclei that outshine all the stars in their parent galaxies. Disks are also the sites of planet and moon formation. Disk evolution is likely driven by turbulent diffusion of angular momentum within the disk. The leading candidate for producing turbulence in disks is the magnetorotational instability (MRI) [1]; our goal in this project was to investigate saturation of MRI-driven turbulence and, with groundbreaking resolution, decide between claims that MRI-driven turbulence in stratified disks is converged [2] and not converged [3] in implicit large eddy simulation (ILES) studies.

METHODS & RESULTS

We studied a model disk problem known as the *stratified shearing box* using the hybrid (MPI+CUDA) ramses-GPU ideal magnetohydrodynamics (MHD) code developed by our collaborators [4]. Using Blue Waters, we were able to integrate at an **unprecedented** resolution of 768 x 512 x 3072 for the hundreds of dynamical times necessary to obtain a statistically significant measurement of turbulent intensity. Our linear resolution **surpassed** all earlier models by nearly a factor of 2. Figure 1 shows the effect of increasing resolution on the density and magnetic field structure in our main series of models.

The result of our **record-breaking** integration was surprising. Although some measures of turbulent intensity did converge, others did *not*. In particular, we found that the characteristic size of turbulent eddies (the correlation length) scales with the number of zones per disk scale height N as $N^{-1/2}$. Figure 2 shows the run of correlation length with a vertical position in the disk over a factor of 16 in resolution, illustrating the weak nonconvergence. This result suggests that many numerical models of MHD turbulence in disks that use ILES, rather than explicit dissipation models, may not be convergent. Our results differ from earlier results showing nonconvergence of ILES in certain ("zero net field") unstratified shearing boxes [5], which show correlation length scaling directly with the grid scale, i.e. as N^{-1}. Confirmation of our results will be required using alternate numerical schemes (and therefore alternate ILES closures), such as athena. Our results also motivate further high-resolution investigation of models with explicit dissipation.

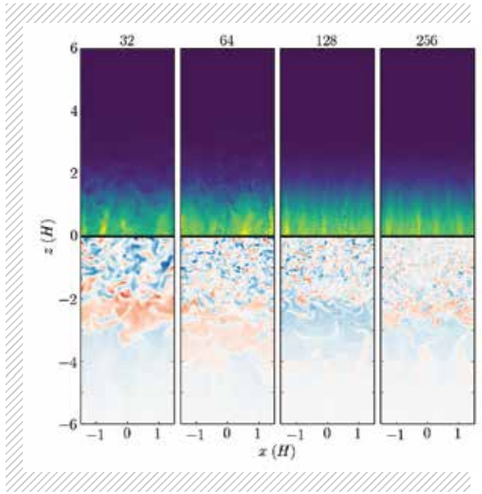

FIGURE 1: Slices through a 3D simulation of the fluid density and azimuthal magnetic field strength over the range of resolutions surveyed in this work. Increasing resolution dramatically decreases the length scale of the magnetic field structure, with a corresponding decrease in turbulent intensity.

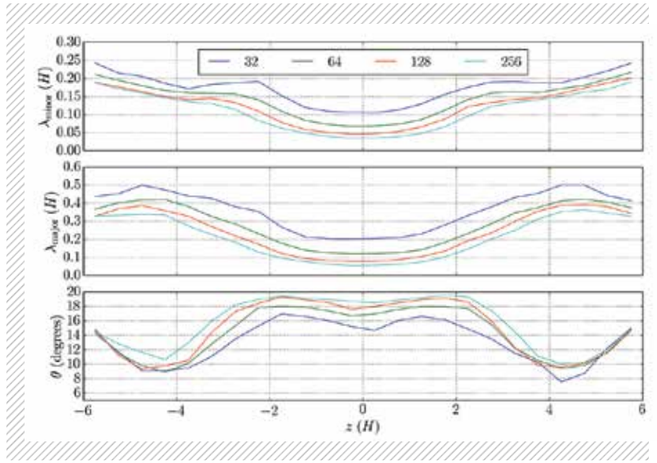

FIGURE 2: Magnetic field correlation function for disk turbulence, from 32 zones per disk scale height H to a record-breaking 256 zones/H. The correlation function in the plane of the disk is characterized by a minor axis correlation length (top panel), major axis correlation length (middle panel) and the angle between the semimajor axis and the orbital velocity vector (bottom panel). While the change in angle with resolution is consistent with convergence, the change in correlation lengths near the midplane of the disk is not.

WHY BLUE WATERS

Blue Waters is the only available facility capable of running this problem. We were able to exploit the large number of GPUs available on Blue Waters, which run our GPU-based code at a factor of about three faster per node over a CPU-based version of the code.

NEXT GENERATION WORK

Supercomputers like Blue Waters now commonly integrate models of hydrodynamic turbulence with explicit dissipation (DNS model) with a several decade inertial range between the outer scale and inner (dissipation) scale. MHD turbulence, by contrast, has both a resistive and a viscous dissipation scale. The ratio of these scales is an important dimensionless parameter of the problem that varies over many orders of magnitude in astrophysical settings. Studies of MHD turbulence in DNS models with a clear separation of viscous and resistive length scales will require next-generation facilities.

COMPARING CAF AND MPI-3 AND SIMULATING MOLECULAR CLOUD TURBULENCE WITH TWO-FLUID MHD

Allocation: GLCPC/552 Knh
PI: Dinshaw S. Balsara[1]
Collaborators: Sudip Garain[1], Jinho Kim[1], Blakesley Burkhart[2], Alexandre Lazarian[3], and Dan Nagle[1]

[1]University of Notre Dame
[2]Harvard University
[3]University of Wisconsin-Madison

EXECUTIVE SUMMARY

This project has three foci. First, we plan to demonstrate that Coarray Fortran standard offers petascale performance that is comparable to or better than the MPI-3 standard has to offer. Both CAF and MPI-3 have also been shown to outperform the MPI-2 standard by a substantial margin. In doing this work, we also found efficient implementation strategies that work well for any one-sided messaging paradigm. Second, we wish to study two-fluid turbulence in molecular clouds. The turbulent plasma that makes up a molecular cloud is predominantly made up of neutral molecules that are threaded by ions that gyrate around a strong magnetic field. The ions and neutral fluids are coupled, but not perfectly, which results in a modified turbulence. We have also extended our two-fluid studies to relativistic two-fluid plasmas that are made of electron-positron fluids.

INTRODUCTION

Observations show that star formation takes place in a partially ionized plasma. Furthermore, Balsara [1] showed that it is critical to represent the plasma as two fluids, one made up of neutrals and the other made up of ions that are threaded by a magnetic field. Previous studies were analytical or restricted to low resolution. Thanks to the availability of Blue Waters, we have been able to carry out large-scale studies of two-fluid turbulence computationally. These have been documented in two recent papers [2,3] and have also contributed to the graduate theses of both Chad Meyer and Blakesley Burkhart, both of which are currently employed based on their computational skills.

METHODS & RESULTS

The first task that we undertook in the 2014-2015 timeframe was to generate a very extensive comparison of Cray's CAF standard with MPI-3 for a range of partial differential equation (PDE) applications. Both CAF and MPI-3 are novel programming paradigms, and it would greatly help the community if their capabilities were documented and published. We are also required to carry out this task as part of our NSF funding (NSF-ACI-1307369; NSF-DMS-1361197). Both CAF and MPI-3 provide for one-sided, non-blocking messaging, which should make them especially well-adapted to exploit Cray's SHMEM library middleware. We have carried out such a weak scalability study and report our results in Garain, Balsara & Reid [4].

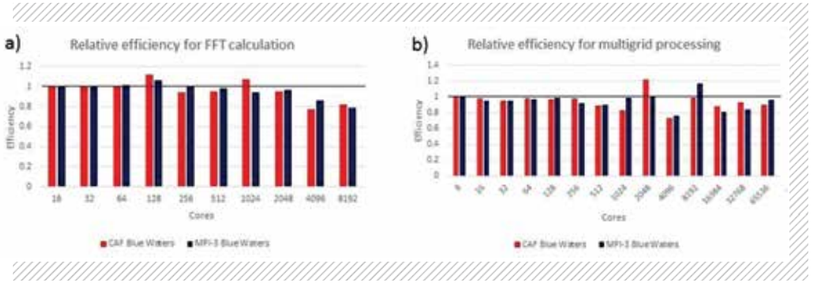

FIGURE 1A, B: From Garain, Balsara & Reid (2015), show a weak scaling study for FFT and Multigrid-based PDE solvers respectively. The relative efficiency for CAF and MPI on Blue Waters is shown in red and blue colors respectively.

WHY BLUE WATERS

Fig. 1 shows the results of weak scaling on Blue Waters for Fast Fourier Transform (FFT) and multigrid applications ranging from 8 to 65,536 cores. Blue Waters is the only available machine on which we can demonstrate this level of extremely scalable performance. We observe that CAF and MPI-3 can keep pace with each other across the entire range of processors. We further find that both CAF and MPI-3 are more than twice as fast as MPI-2 when one approaches large numbers of processors. This observation shows the immense value of these novel programming paradigms when computing at scale. Furthermore, this study would not have been possible without Blue Waters. In Garain, Balsara & Reid [4] we also document best practices for using CAF and MPI-3 for the community. We also show that CAF code is much easier to write and maintain, and the simpler syntax makes the parallelism easier to understand. Educational lectures on CAF were also developed as part of this work and are freely available from our website (http://www.nd.edu/~dbalsara/Numerical-PDE-Course).

NEXT GENERATION WORK

Electron-positron plasmas are very common in strongly relativistic environments, like the atmospheres of pulsars or black hole accretion disks. Thus developing a highly parallel capability to simulate such plasmas is crucial, and was achieved in the work of Balsara, Amano, Garain and Kim [5]. It is thought that the gamma-ray flares from the Crab pulsar are due to rapid magnetic reconnection events. Fig. 2 shows the results of such a simulation from Blue Waters. These are early results stemming from a very productive new direction.

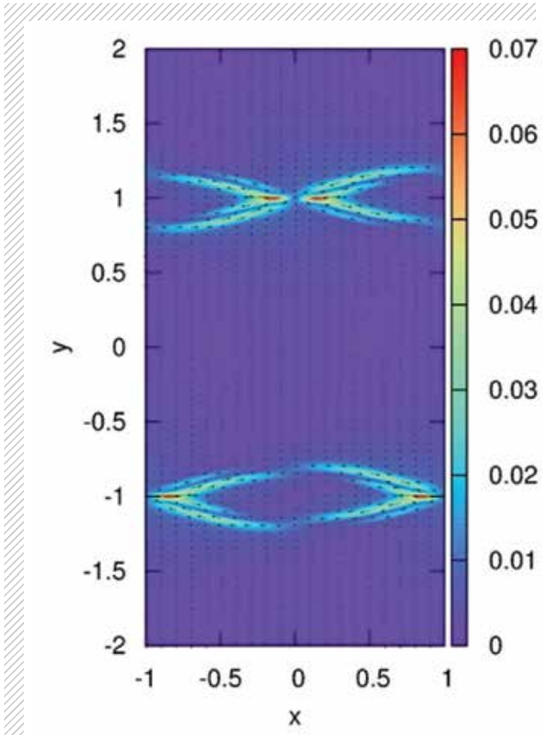

FIGURE 2: Showing the velocity field for reconnection in an electron-positron plasma from Balsara et al. (2015).

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

Allocation: NSF PRAC/5.00 Mnh
PI: Michael L. Norman[1]
Co-PI: James Bordner[1]

[1]University of California, San Diego

EXECUTIVE SUMMARY

We have performed state-of-the-art simulations of the intergalactic medium (IGM) during the epoch of helium reionization. It is believed that the UV radiation emitted from quasars ionized the intergalactic helium over a period of several billion years beginning at a redshift of 4. We 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. For **the first time** we have performed multigroup radiative transfer self-consistently coupled to the cosmological hydrodynamics of the IGM at sufficient resolution and domain size to examine the photoionization and photoheating processes in detail. This is being done to determine whether this is the "missing physics" that will improve the agreement between models and high-precision observations of the IGM, specifically the Lyman-alpha forest observations of distant quasars.

INTRODUCTION

In the past decade, new, more precise observations of the intergalactic medium (IGM)—the hydrogen and helium gas between the galaxies produced in the Big Bang—have revealed a discrepancy with the well-established predictions of our computational models. In particular, precision observations of the IGM using the Keck telescopes in Hawaii show that the temperature and ionization state of the IGM is not what our standard cosmological simulations predict: The IGM is either somewhat hotter than ultraviolet radiation from stars in galaxies can make it, or the IGM is distributed differently in space than the simulations predict, or both. There could be missing sources of heat in our models, such as energy injection by decaying dark matter particles. The discrepancy is perplexing since the standard model predicts the galaxy distribution exceedingly well. The discrepancy suggests that the standard model lacks some essential ingredient, which we refer to simply as "missing physics." The significance of this project to the nation is that it promotes the progress of science in the fundamental field of cosmology, in which the United States is a world leader. The project is addressing the issue of whether we are overlooking a key component of the mass-energy content of the universe. Precise answers require powerful tools, and the Blue Waters supercomputer is the tool for the job.

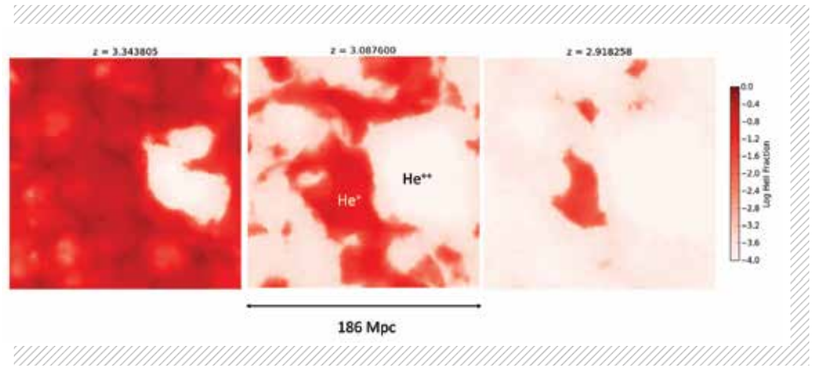

FIGURE 1 : Photoionization of helium in the intergalactic medium by UV radiation from time-dependent quasars. Shown is the projection of the He^+ fraction, which steadily diminishes as He^+ is converted to He^{++}. Simulation is carried out in a 185 Mpc volume using the Enzo hydrodynamic cosmology code coupled to a 5-group radiation diffusion solver. Credit: M. Norman, UCSD.

METHODS & RESULTS

In the first year of this project, we have examined the possibility that late photoheating of the IGM by quasars is the missing physics. In the standard model of the Lyman-alpha forest (which subtly disagrees with observations), quasar ionization is modeled as a homogeneous but time-varying radiation background. Photoheating is treated in the optically thin limit, which underestimates the heating behind optically thick ionization fronts. In reality, quasars are radiating point sources that ionize the helium in the IGM in their vicinity; it is the growth and eventual overlap of these growing spheres of ionization, each centered on a luminous quasar, that ionize the IGM. This is the situation we have simulated on Blue Waters (see Figure 1).

We have carried out a suite of the **first** fully coupled radiation hydrodynamic cosmological simulations that treat the quasars as a time-varying statistical population of point sources. For the **first** time, we perform multigroup flux limited diffusion radiative transfer self-consistently coupled to the cosmological hydrodynamics of the IGM at sufficient resolution and domain size to examine the photoionization and photoheating processes in detail. This is being done to determine whether this is the "missing physics" that will improve the agreement between models and high-precision observations of the IGM, specifically the Lyman-alpha forest observations of distant quasars. We now are analyzing the spectroscopic properties of our simulated Lyman-alpha forest, including these effects.

WHY BLUE WATERS

Blue Waters is required because the simulation is extremely computationally intensive. Large spatial volumes need to be simulated at high resolution, to simultaneously resolve the Lyman-alpha forest absorbers, and at the same time encompass a representative sample of quasars. The ratio of the outer and inner scales is about 2,000, meaning that simulations require $2{,}000^3$ grids. Additionally, a multifrequency treatment of radiative transfer is essential to accurately model the transport of the hard UV spectrum of quasars, including the phenomena of ionization front pre-ionization and spectral hardening of the radiation field.

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

NEXT GENERATION WORK

In the 2019-2020 timeframe the post-petascale version of the Enzo adaptive mesh refinement (AMR) simulation code will be fully operational. Called Enzo-P, it will permit AMR radiation hydrodynamic cosmological simulations to be performed of a size and dynamic range equaling or surpassing the largest pure dark matter N-body simulations carried out today. Enzo-P is an application built on top of the extreme scale AMR framework called Cello, developed at the University of California, San Deigo. Cello, in turn, is based on the Charm++ parallel objects framework that underpins the highly scalable molecular dynamics code NAMD. We will use Enzo-P to simulate the combined effects of galaxies and quasars on the intergalactic medium, including radiative, chemical, and kinetic feedback. We will use the improved knowledge of high redshift galaxies obtained by the James Webb Space Telescope, to be launched in 2018, to create the most physically detailed model of the co-evolving population of galaxies and active galactic nuclei that reionize the universe at z=6-7. We will explore the post reionization evolution of the IGM in the redshift interval 4-6, which is relatively poorly understood at this time.

PREDICTING THE TRANSIENT SIGNALS FROM GALACTIC CENTERS

PI: Scott C. Noble[1]
Co-PIs: Manuela Campanelli[2] and Julian Krolik[3]
Collaborators: Dennis Bowen[2], Roseanne Cheng[3], and Hotaka Shiokawa[4]

[1]The University of Tulsa
[2]Rochester Institute of Technology
[3]Johns Hopkins University
[4]Harvard University

EXECUTIVE SUMMARY

Active supermassive black holes (SMBHs) at the centers of galaxies are home to the most energetic phenomena in the universe and are systems of great interest for high-cadence sky surveys such as the Large Synoptic Survey Telescope (LSST). Insight into their transient nature is contingent on the community developing detailed simulations of electromagnetic emission from SMBHs and SMBH binaries (SMBHBs), which is our goal. With Blue Waters, we have learned how the electromagnetic fluctuations from SMBHBs depend on the mass ratio of the black holes and the disk's magnetic field. We have explored the stability properties of "mini-disks" about each black hole. Finally, we have begun exploring tidal disruptions of stars using our new "multi-patch" system for simulating dynamics at different temporal and spatial scales. Each set of runs advances the field's existing resolution and sophistication limits, thanks to the unique resources of Blue Waters.

INTRODUCTION

Black holes are exotic astrophysical objects that power some of the most energetic phenomena in the universe. The first direct detection of gravitational waves this past year has ushered us into an era of gravitational wave astronomy, and implies that the future space-based observatories will be rich in sources. Our work with Blue Waters will provide the community with the most advanced simulations of two kinds of systems of focus in transient surveys: SMBHB accretion disks and tidal disruption events (TDEs) of stars by SMBHs. Both projects offer the opportunity to test the strong-field limit of gravity, which benefits society fundamentally and provides new results on a topic of significant public interest: black holes and how they influence their environment.

METHODS & RESULTS

Our simulations on Blue Waters rely on the conservative, high-resolution shock-capturing code called HARM3d to solve the equations of general relativistic magnetohydrodynamics (GRMHD). Its coordinate-independent approach allows us to resolve the black holes in SMBHB simulations using a "dual fisheye" system that smoothly transitions from a local Cartesian-like high-resolution grid near each black hole to a spherical coordinate system further away where the circumbinary disk resides, all via a continuous coordinate transformation [1]. Also, HARM3d now accommodates a multi-program/multi-data mode, with which we may solve the equations independently, only interacting via boundary conditions. We call this approach "multi-patch" and it is being used for the **first time** in TDE simulations on Blue Waters, and will be used extensively for SMBHB simulations in the near future. It offers the ability to evolve the patches with different time steps, grid spacings, physics modules (e.g., gravity models), reference frames, and coordinate topologies—all to maximize the simulation's runtime performance and scalability.

We have expanded our earlier work on circumbinary disks [2,3], to consider SMBHBs of different mass ratio and magnetic field distributions. We find that once the mass ratio drops below 1/5, the binary's tidal torque diminishes to a degree that the disk's activity resembles that about a single black hole. Taking advantage of Blue Waters' large resources, we quadrupled the size of the initial reservoir of mass and magnetic field. This simulation confirms that a sustained source of magnetic field can mitigate the development of the circumbinary disk's electromagnetic signal. This work implies that these strong periodic signals may only arise from near-equal mass binaries in weakly magnetized environments.

Our new "mini-disk" initial data prescription, which is being demonstrated for the **first time** in 3D on Blue Waters, will allow us to bypass a large portion of the relaxation phase of the evolution and simulate the approximate steady-state of the circumbinary flow for longer. With Blue Waters, we have measured the tidal truncation radius of the mini-disks in the relativistic regime for the first time. We find that Newtonian predictions [4-6] remain valid down to binary separations of approximately 50M, but relativistic corrections are required for smaller separations.

The full TDE, from disruption to accretion, has not been simulated before because of the prohibitive expense of numerical methods which treat the problem in one computational domain. Our initial 3D hydrodynamic calculations are made affordable by using the multi-patch scheme, which is used for the first time for this project (Fig. 1). In addition to validating the multi-patch infrastructure, it paves the way for the **first ever** MHD TDE evolutions we hope to perform on Blue Waters soon.

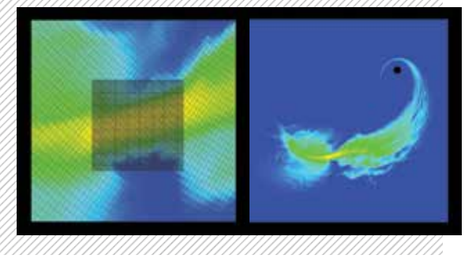

FIGURE 1: Logarithm of the density of debris from a star tidally disrupted by a black hole. Image credit: R. Cheng (JHU).

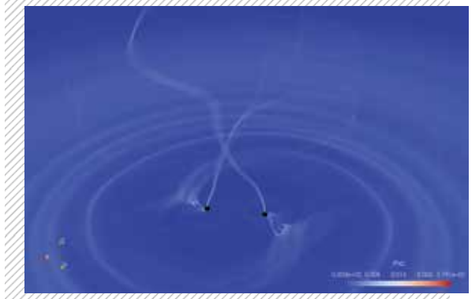

FIGURE 2: Magnetic field lines (white curves) emanating from two black holes. Equatorial slice of density of the accreting gas (background). Image credit: M. Vanmoer (NCSA).

WHY BLUE WATERS

Blue Waters provides a unique resource with which we may run at scales and core counts not possible on other systems. Using multiple patches multiplies the required core-count by the number of patches used. Further, our circumbinary disk calculations require tens of millions of time steps to cover the large range in time scales inherent to the problem.

An incredible resource made possible by our Blue Waters allocation is access to NCSA's visualization team. We have worked closely with Mark Vanmoer and Roberto Sisneros on how to best visualize magnetic field lines and volumetric ray-casts of density from our circumbinary accretion disk simulations (Fig. 2).

NEXT GENERATION WORK

To properly consider the effects of opacity and radiation pressure, the radiation's field equations must be evolved in tandem, adding extra dimensions to the problem and amplifying the workload dramatically. We look forward to the future generation of supercomputers for including radiation in our GRMHD simulations. In particular, we anticipate relying on graphics processing unit (GPU)/integer-core computing to relieve some of the extra computational burden, which will either require us to develop integer-core abilities into HARM3d or have the abilities automatically provided by the next generation of load balancing and distribution environments (e.g., Charm++ [7]).

Further, we expect to use the multi-patch scheme in a variety of new ways in the future. Since it uses the multi-program/multi-data program model, we may run microphysics codes (e.g., for local turbulence) alongside a global calculation to provide it with sub-cell-level input. Such calculations will again multiply the required resources and core-counts, and provide demand for the nation's continuing support of **cutting-edge** computing resources.

CONNECTING GALAXIES IN THE EARLY UNIVERSE TO THE MILKY WAY

PI: Brian O'Shea[1]
Co-PIs: David Collins[2] and John Wise[3]
Collaborators: Devin Silvia[1], Cameron Hummels[4], and Britton Smith[5]

[1]Michigan State University
[2]Florida State University
[3]Georgia Institute of Technology
[4]California Institute of Technology

EXECUTIVE SUMMARY

In this project, we are investigating the growth and evolution of galaxies like the Milky Way and its satellites over the entire age of the Universe, and we are addressing observational questions raised by a variety of astronomical observations. In particular, we are examining the life cycle of gas in galaxies, the evolution of the Milky Way's earliest progenitor galaxies, and the ways in which ionizing radiation and magnetic fields are created and escape from galaxies. All of these problems require simulations with an extremely high dynamic range in space and time, complex physics (including radiation transport, magnetohydrodynamics, and complex gas chemistry), and large simulation volumes. We use the Enzo code, which has been modified to scale to large core counts on Blue Waters, the only machine available where the heavy data and communication needs can be satisfied.

INTRODUCTION

Despite the wealth of astronomical observations examining the properties of galaxies, a detailed and self-consistent theoretical model of galaxy formation does not yet exist. The evolution of galaxies has been examined in great detail and in many wavelength bands from approximately 600 million years after the Big Bang to the present day. These observations show that the galaxies that we can see have undergone radical changes in size, appearance, and content over the last 13 billion years and that the properties of galaxies are intimately tied to their environment and formation history [1].

Complementary observations of our own Milky Way have provided a rich dataset of the kinematics and elemental abundances of local stellar populations, including large numbers of metal-poor stars in the halo of our galaxy and local dwarf galaxies [2]. In principle, this "galactic fossil record" can probe the entire merger and star formation history of the Milky Way and its satellites, and is an extremely useful counterpoint to direct observations of distant galaxies. Similarly, observations of diffuse clouds of gas in the halo of the Milky Way provide important constraints on how gas reaches the disks of galaxies, thus feeding star formation [3]. This complements observations of the same process in distant galaxies from surveys like the Hubble Space Telescope's COS-Halos project [4].

The volume of observational data on galaxy formation, which is already staggering, will increase exponentially over the next decade as observatories such as the Large Synoptic Survey Telescope and James Webb Space Telescope come online. At present, however, we lack the theoretical models necessary to interpret such observational datasets adequately. This lack of theoretical models is a result of the inherent challenges in the modeling of galaxies, which includes the necessity for both high dynamic range and complex, interconnected physics. As such, it is necessary to resort to large-scale numerical simulations to accurately model these physical phenomena.

METHODS & RESULTS

This project was carried out using the Enzo code [5], which is used for the simulation of cosmological and astrophysical phenomena. Enzo uses an adaptive mesh refinement scheme to achieve high spatial and temporal resolution and provides a huge range of physics modules including dark matter, hydrodynamics, magnetohydrodynamics, radiation transport, non-equilibrium many-species gas chemistry, and prescriptions for star and black hole formation and feedback. Enzo is publicly available and used to model many types of astrophysical phenomena.

Thus far, we have primarily focused on the analysis of the results of our previous Blue Waters-related project and have run large simulations of the formation of Milky Way-type galaxies. Our most interesting results have been the characterization of the properties of galaxies in the early universe – namely, the galaxies that form the seeds that will one day become the Milky Way and its satellites. We have found that these galaxies have very intermittent star formation, form stars very inefficiently at lower masses (as compared to higher-mass galaxies), and that the properties of individual galaxies' stellar and gas content, including metallicity and formation history, vary significantly from galaxy to galaxy.

WHY BLUE WATERS

The simulations required to properly model galaxies over the age of the universe require extremely high spatial and temporal dynamic range, and also require complex physics – most importantly for our problem –radiation transport, magnetohydrodynamics, and non-equilibrium gas chemistry. Furthermore, large simulation volumes (and thus many resolution elements) are needed to model enough galaxies to be able to adequately compare theory to observations in a statistically meaningful way. Taken together, this requires the use of a supercomputer with large memory and disk space (to accommodate the tremendous dataset sizes), 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 system available to the academic community that fits all of these requirements.

NEXT GENERATION WORK

A future-generation system will allow us to include even more physics – simultaneously including much more sophisticated models for radiation transport that resolve line transport as well as continuum transport, for example – and increased spatial and temporal dynamical range. The goal for Milky Way-type galaxies is to be able to simulate the entire system from the "virial radius" (the scale defining the "edge" of the Milky Way's neighborhood – approximately one million light years from the center of the galaxy) down to the size of individual star-forming clouds (a few light years), all within a cosmological context. This capability would allow us to perform incredibly realistic simulations of cosmological structure including all possible important physical phenomena.

FIGURE 1 (LEFT): This image shows several galaxies in a small region of the universe approximately 500 million years after the Big Bang. These galaxies will eventually become part of a galaxy like our own Milky Way. The blue filamentary structure shows neutral hydrogen (i.e., the "cosmic web"), which feeds cold gas into galaxies to form stars. The red and yellow clouds show gas that has been ionized by massive stars in the individual galaxies, and the white glowing regions show radiation from populations of massive stars that are currently ionizing the universe. Image credit: Bob Patterson (NCSA Advanced Visualization Laboratory).

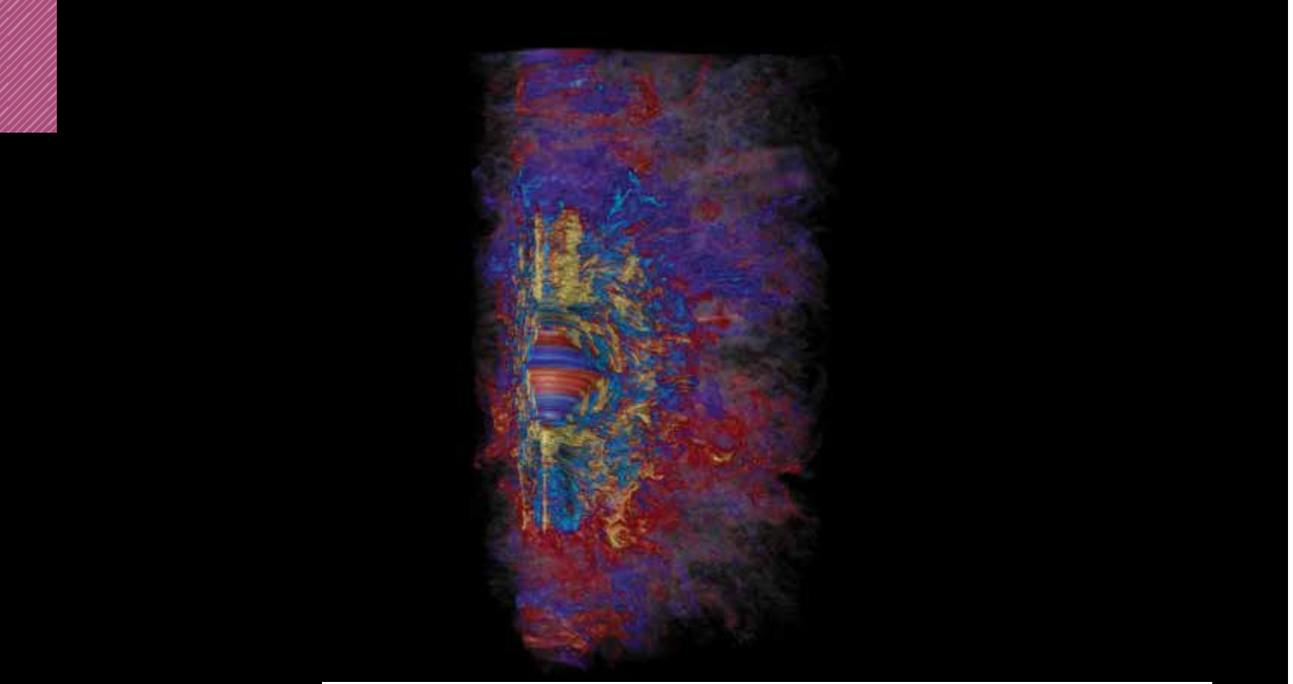

MAGNETARS, BLACK HOLE COLLISIONS FOR LIGO, AND A NEXT GENERATION NUMERICAL RELATIVITY CODE

FIGURE 1: Volume rendering of the toroidal magnetic field strength in a rapidly rotating core-collapse supernovae. The simulation shows a forming magnetar as the turbulent engine of the explosion. Colors represent magnetic field strength. Yellow and light blue indicate magnetar-strength positive/negative field and red and blue weaker field.

Image credit: Robert R. Sisneros (NCSA) and Philipp Mösta (UC Berkeley).

PI: Christian D. Ott[1]
Co-PIs: Mark Scheel[1], Peter Diener[2]
Collaborators: Philipp Mösta[3], Luke Roberts[1], David Radice[1], Sherwood Richers[1], Roland Haas[5], Erik Schnetter[2], Lawrence Kidder[4], Saul Teukolsky[4], Scott Field[4], and Francois Foucart[3]

[1]California Institute of Technology
[2]Lousinana State University
[3]University of California, Berkeley
[4]Cornell University
[5]Albert Einstein Institute

EXECUTIVE SUMMARY

Neutron stars with an extremely strong magnetic field are called magnetars. An open question is how they obtain their magnetic field. We have carried out 3D simulations of a magnetized differentially rotating proto-neutron star. These simulations show that the magnetorotational instability is active in such proto-neutron stars. It efficiently amplifies magnetic field locally and a dynamo process rearranges and orders the field, creating a large-scale global magnetic field as needed for driving supernova explosions.

The National Science Foundation's (NSF) Laser Interferometer Gravitational Wave Observatory (LIGO) has made its first detection of gravitational waves. We responded with target-of-opportunity numerical relativity simulations on Blue Waters that are now helping LIGO to extract information more reliably from the waves.

To more efficiently use large core counts on Blue Waters and in preparation for the next-generation high-performance computing architectures, we have designed, implemented, and tested the new general-relativistic magnetohydrodynamics (MHD) code, SpECTRE. It employs the Discontinuous Galerkin method; is based on the Charm++ parallel programming system, and it exhibits near perfect strong scaling to 128,000 Blue Waters cores.

INTRODUCTION

Numerical general relativity simulations discretize Einstein's equations and are essential to study phenomena involving strong gravity, such as the collision of black holes or the formation of neutron stars in core-collapse supernova explosions.

Magnetars, neutron stars with extremely strong magnetic fields ($B \geq 10^{15}$ Gauss), are believed to power energetic supernova explosions (so-called "hypernovae") and cosmic gamma-ray bursts. MHD turbulence has been suggested as a mechanism for building up the extremely strong magnetar magnetic field. This MHD turbulence is driven by the magnetorotational instability (MRI), whose fastest growing mode has a very short wavelength, making it difficult to resolve in simulations. An additional complication is that the MRI builds up magnetic field strength locally, but what is needed for a magnetar is a large-scale (global) magnetic field that can force the high-density stellar plasma into coherent motion and thus drive a supernova explosion.

Gravitational Waves (GWs) from a coalescing pair of black holes were observed by LIGO on September 14, 2015. This first direct detection of GWs opened up the new field of GW Astronomy. Numerical relativity simulations of coalescing binary systems (and other GW sources) provide essential GW signal predictions for LIGO – these predictions are needed for comparison with GW observations to extract the astrophysical parameters (e.g., mass and spin of black holes) of the sources and for testing General Relativity.

Parallel scaling and efficiency on heterogeneous architectures are key requirements for petascale (and future exascale) simulation codes. Much current computational astrophysics and numerical relativity codes use traditional data-parallel approaches and sequential execution models that suffer from load balancing issues, allow no or only limited execution concurrency and make it difficult to benefit from accelerators. High-order finite-difference and finite-volume methods require the communication of three or more subdomain boundary points in each direction, creating large communication overheads. The Discontinuous Galerkin (DG) method can drastically reduce the number of points that must be communicated. DG, in combination with a task-based parallelization strategy (which allows task concurrency), provides dynamical load balancing, and can hide latencies, may be the route forward for simulation codes.

METHODS & RESULTS

Making Magnetars

We employed our 3D general-relativistic MHD code Zelmani for our magnetar MHD turbulence simulations. Zelmani is open-source, is based on the Cactus framework, and uses the Einstein Toolkit. We mapped our initial conditions from a full 3D adaptive mesh refinement simulation of rapidly rotating stellar collapse with an initial magnetic field of 10^{10} Gauss. We then carried out simulations at four different resolutions (500 m, 200 m, 100 m, 50 m), where the two highest resolutions are sufficient to resolve the key driver of turbulence in this system, the MRI. The 50 m simulation was run on 132,768 Blue Waters cores, used 7.8 TB of runtime memory, and created 500 TB of output data.

Previous local simulations have shown that the instability creates small patches of the magnetar-strength magnetic field, but have not been able to address whether these patches can be turned into a global dynamically relevant toroidal magnetic field structure needed to power an explosion.

Our simulations show that not only does the MRI produce local magnetar-strength magnetic field, but that a subsequent dynamo process connects this local field into a global field that is able to support an explosion. This dynamo process, which is akin to local storm cells merging to form large-scale storm patterns like a hurricane, works extremely efficiently and we predict the formation of a large-scale toroidal field independent of the initial magnetization of the progenitor star. Figure 1 shows an example magnetic field configuration obtained in our highest-resolution simulation.

Our simulations demonstrate that rapidly rotating stellar collapse is a viable formation channel for magnetars. Magnetars are abundantly observed in the universe and have recently gained a lot of attention as possible engines driving hypernovae and superluminous supernovae, some of which are also connected to long gamma-ray bursts.

Binary Black Hole Waveforms for Gravitational Wave Astronomy with LIGO

We simulate binary black hole (BBH) coalescence using the Spectral Einstein Code (SpEC). SpEC uses multi-domain pseudospectral methods to discretize Einstein's equations. Pseudospectral methods provide exponential convergence and are optimal for simulating gravitational fields, which are guaranteed to be smooth. The black hole interiors are excised from the computational grid to avoid singularities.

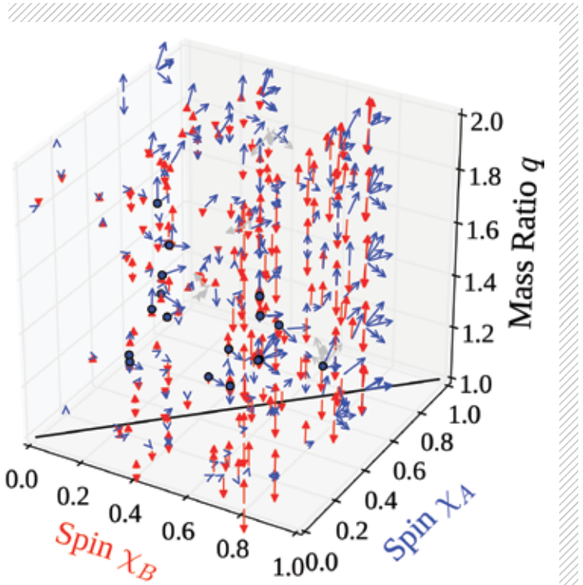

FIGURE 2: Parameters of the 285 binary black hole systems simulated for LIGO. The location of each circle indicates the mass ratio and dimensionless spin magnitudes for a system. The arrows indicate the direction of spin of each black hole (blue arrows for hole A, the larger black hole; red arrows for hole B) and are scaled by the magnitude of the spin.

An astrophysical binary black hole (BBH) system has seven parameters: the mass ratio M_A/M_B and two spin vectors with three components each. The total mass of the system is a scale factor that can be removed from theoretical predictions, but can be inferred from the frequency of the observed waves. Using LIGO data and approximate analytic models for BBH gravitational waveforms, some of the BBH parameters have been measured for the BBH system that LIGO observed. To improve these measurements, numerical relativity waveforms can be used. However, Bayesian parameter estimation algorithms need millions of waveforms as they sample the parameter space near the detected signal. This is problematic because each numerical relativity simulation takes weeks to months to complete.

To address this problem, we have developed a technique to interpolate between numerical relativity waveforms that we *did* simulate to construct waveforms at parameters that we *did not* simulate. The interpolant waveform can be computed in a fraction of a second (compared to the tens to hundreds of thousand central processing unit hours required to construct such a waveform using numerical relativity). Furthermore, the interpolant can be made indistinguishable from the results of numerical relativity to within a specifiable error tolerance. It can thus serve as a *surrogate model* that replaces additional numerical relativity simulations. This error tolerance dictates how many numerical relativity simulations are needed – typically many orders of magnitude fewer than the millions that are needed for parameter estimation – and a greedy algorithm determines where in parameter space these simulations must be performed.

Motivated by LIGO's detection, we have carried out target-of-opportunity SpEC BBH simulations using Blue Waters essentially as a high-throughput capacity computing facility to run hundreds of BBH simulation, each using only two or three nodes. Figure 2 shows the simulated points of the parameter space. In total, we simulated 286 BBH systems, and each was simulated with three different resolutions to estimate the truncation error. Using these waveforms, we built a surrogate model that outperforms previous analytic waveform models and enables more accurate parameter estimation.

A next-generation simulation code: SpECTRE

SpECTRE is a new numerical relativity and general-relativistic hydrodynamics code that uses the DG method with a task-based parallelization strategy facilitated by the Charm++ parallel programming system. DG and task-based parallelization make a promising combination that will allow multi-physics applications to be treated both accurately and efficiently. On Blue Waters, SpECTRE now scales to more than 128,000 cores on a relativistic MHD test case simulating the interactions of shocks (Fig. 3). The Charm++ task-based parallelism library includes a dynamic runtime system to assign task ordering as well as communication optimizations tailored to Cray's Gemini network.

Building on this success, we are now adding the remaining infrastructure and physics to simulate realistic astrophysical systems. To tackle challenging multi-scale problems which naturally arise in such systems, we are adding local time stepping techniques and adaptive mesh refinement strategies to either split the elements into smaller elements or increase the number of basis functions in an element.

WHY BLUE WATERS

Our MHD turbulence simulations that resolve the MRI would have been **completely impossible without Blue Waters'** capability. Blue Waters was essential for these simulations and directly facilitated the breakthrough in our understanding how magnetar-strength magnetic field can be generated. Blue Waters' capacity is crucial for a quick turnaround in our target-of-opportunity numerical relativity BBH simulations. No other machine allows us to run hundreds of simulations concurrently to generate the theoretical predictions needed for LIGO. Finally, Blue Waters is a highly-valuable development and testing resource for our next-generation code SpECTRE. We couldn't develop this code without being able to test it at scale on Blue Waters.

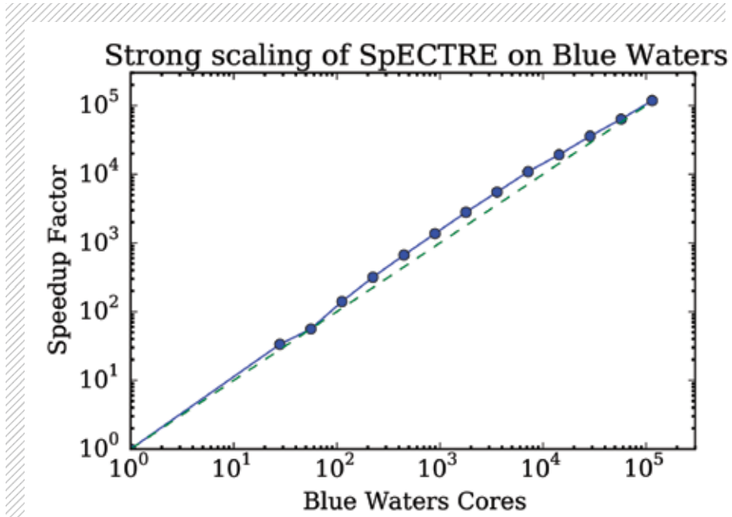

FIGURE 3: Strong scaling of the next-generation general-relativistic MHD code SpECTRE for an MHD test problem on Blue Waters.

NEXT GENERATION WORK

We are working with high intensity on our new code SpECTRE, which we plan to deploy on a next-generation Track-1 system for more detailed, higher-resolution, and longer (in physical time) simulations of core-collapse supernovae, magnetars, and binary systems of black holes and neutron stars. Present simulations, even with Blue Waters, are limited in accuracy and, importantly, in the physical time they can track. Our magnetar MHD simulation, for example, could cover only tens millisecond of the magnetar's life while we would **need to simulate at least a hundred times longer** to fully study its magnetic field evolution and its impact on the supernova explosion.

PUBLICATIONS AND DATA SETS

Mösta P., et al., A large-scale dynamo and magnetoturbulence in rapidly rotating core-collapse supernovae, *Nature*, 528:376 (2015), doi: 10.1038/nature15755

Radice D., et al., Neutrino-driven Convection in Core-collapse Supernovae: High-resolution Simulations, *Astrophys. J.*, 820:76 (2016), doi: 10.3847/0004-637X/820/1/76

Radice D., et al., The One-Armed Spiral Instability in Neutron Star Mergers and its Detectability in Gravitational Waves, *Phys. Rev. D.*, arXiv:1603.05726 (2016).

Radice D., et al., Dynamical Mass Ejection from Binary Neutron Star Mergers. *Mon. Not. Roy. Ast. Soc.*, arXiv:1601.02426 (2016).

Bernuzzi S., et al., How Loud Are Neutron Star Mergers? *Mon. Not. Roy. Ast. Soc.*, arXiv:1512.06397 (2016).

MODELING HELIOSPHERIC PHENOMENA WITH THE MULTI-SCALE FLUID-KINETIC SIMULATIONS SUITE: FROM THE SOLAR SURFACE TO THE LOCAL INTERSTELLAR MEDIUM

Allocation: NSF PRAC/1.21 mnh
PI: Nikolai Pogorelov[1]
Co-PIs: Sergey Borovikov[1] and Jacob Heerikhuisen[1]
Collaborators: Matthew Bedford[1], Raymond Fermo[1], Tae Kim[1], and Mehmet Sarp Yalim[1]

[1]University of Alabama in Huntsville

EXECUTIVE SUMMARY

Blue Waters was used to investigate physical phenomena occurring when the solar wind (SW) interacts with the local interstellar medium (LISM): (1) the origin of the SW on the solar surface and its further acceleration to supersonic velocities; (2) the effect of transient phenomena on space weather on Earth; (3) the SW propagation throughout the heliosphere towards the heliopause and perturbations it creates in the LISM; (4) mixing of the SW and LISM plasma at the heliospheric interface owing to the heliopause (HP) instability and magnetic reconnection; (5) the influence of the heliosphere on the observed anisotropy of TeV galactic cosmic rays; and (6) the application of observations from multiple spacecraft to reconstruct otherwise missing properties of the SW and LISM. Our simulations are important for the explanation of IBEX (Interstellar Boundary Explorer), New Horizons, Ulysses, and Voyager measurements, as well as multi-TeV cosmic ray observations.

FIGURE 1: Magnetic field line distributions around the Sun are shown together with the radial magnetic field contours on the photosphere.

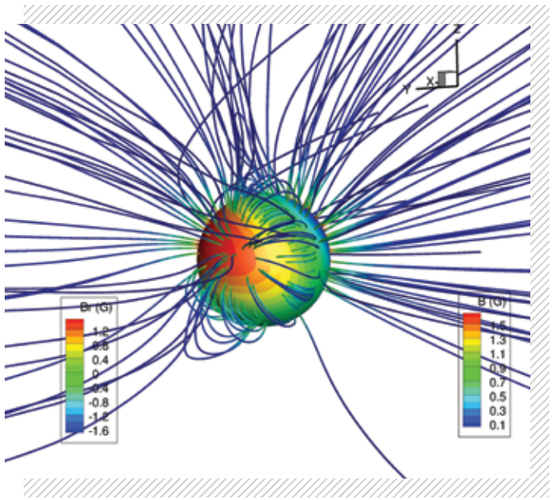

INTRODUCTION

Voyager 1 and 2 (V1 and V2) spacecraft crossed the heliospheric termination shock (TS) in December 2004 and in August 2007, respectively, and after 38 years of historic discoveries, V1 started sampling the LISM [1] while V2 is approaching the heliopause—a tangential discontinuity separating the solar wind from the LISM. The Voyagers acquire in situ information about the local properties of the Solar Wind plasma, energetic particles, and magnetic field at the heliospheric boundary [2]. V1 data related to the LISM properties gives the heliospheric studies community a unique opportunity to constrain models with observational data. Another constraint on the LISM properties is derived from the presence of pickup ions (PUIs) that are born when the thermal solar wind ions exchange charge with the LISM neutral atoms. On the other hand, the IBEX is measuring line-of-sight integrated fluxes of energetic neutral atoms (ENAs) in different energy bands [3]. Since most ENAs are created during charge exchange between hot PUIs and LISM neutral atoms, they bear the plasma properties of the region they are created. Data-driven simulations of the SW allow us to specify the inner boundary conditions. The combination of IBEX, New Horizons, SOHO, Ulysses, Voyager, and multi-TeV air shower data makes it possible to investigate physical phenomena at the heliospheric interface.

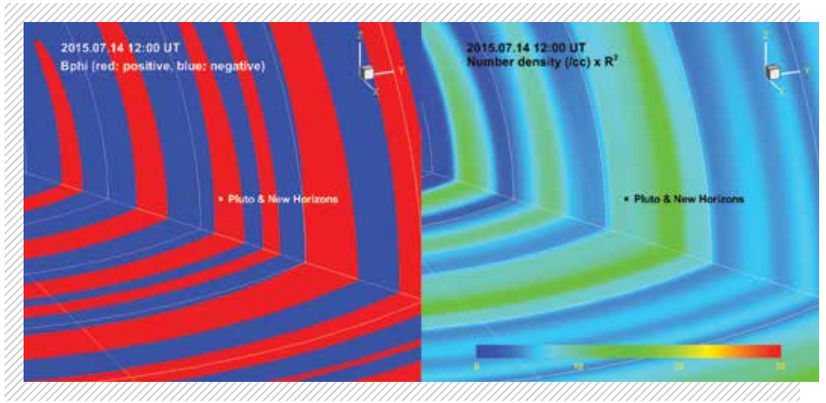

FIGURE 2: The plots magnetic field, B, polarity (left panel) and solar wind proton number density (right panel) on 14 July 2015 12:00 UT. Red and blue colors indicate positive and negative signs of the toroidal component of B, respectively. Proton density is shown scaled to 1 AU values.

METHODS & RESULTS

We solve 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 describing different atom populations dependent on the domains of their origin. We have developed both fluid dynamics and kinetic models for PUIs and turbulence generated by kinetic instabilities of their distribution function. All these are components of a Multi-Scale Fluid-Kinetic Simulation Suite (MS-FLUKSS) – an adaptive mesh refinement code we have built on the Chombo framework.

The major results of the work are (1) we have performed data-driven solar wind flow simulations starting from the solar surface (Fig. 1); (2) we have calculated solar wind propagation from Earth orbit to Pluto along the New Horizons spacecraft trajectory (Fig. 2) and further to the heliopause and demonstrated good agreement with observational data; (3) we have explained the existence of extended regions of the solar wind sunward flow near the heliopause and spontaneous transition to turbulence; (4) we have performed high-resolution simulations of the heliopause instability and identified the areas of possible magnetic reconnection near the heliopause crossed by Voyager trajectories (Fig. 3), which allowed us to put forward a possible explanation of V1 observations that showed a few consecutive increases and decreases in the galactic and anomalous cosmic ray flux intensities; (5) we have analyzed the heliotail flow and quantity distribution in the heliospheric bow wave for different LISM conditions and revealed [4] that the observed multi-TeV cosmic ray anisotropy may be explained by the LISM magnetic field distortion by the heliosphere; and (6) we have investigated the effect of non-thermal ions on time-dependent plasma distributions in the solar wind and LISM.

By addressing the basic physical phenomena occurring at the interface of the heliosphere and LISM, our research is of importance for the solar and heliospheric physics, physics of the interstellar medium, and plasma physics in general. Our collaboration with the Blue Waters team further promotes the application of adaptive technologies to contemporary plasma physics problems through the development of publicly available packages suitable for multiple applications.

WHY BLUE WATERS

(1) Our simulations are computationally intensive. Neutral atoms are modeled kinetically and we need of the order of 1012 particles in our Monte Carlo simulations. These simulations require particle splitting, multiple grids, and careful coupling with the MHD module. They also produce multiple data sets sometimes exceeding 1 TB, which requires hybrid parallelization. (2) Computational region sizes are very large, as in the case of long-tail simulations to about 10,000 AU. (3) Very deep adaptive mesh refinement is necessary near magnetic reconnection sites. With the support of the Blue Waters PAID program we are able to upgrade MS-FLUKSS to use GPUs. We will use our new allocation (PRAC project ACI-1615206) to perform simulations of the turbulent SW-LISM interaction.

FIGURE 3: Distributions of the magnetic field strength (right panel) and plasma density (left panel) in the plane defined by the current directions of the Voyager 1 and Voyager 2 trajectories demonstrate the heliopause instability at V1 and signatures of magnetic reconnection at V2.

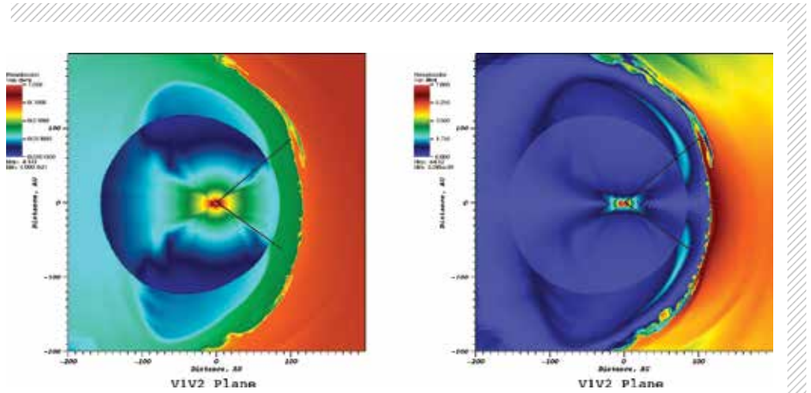

NEXT GENERATION WORK

On a next-generation Track-1 system in the 2019-2020 timeframe, we plan to extend MS-FLUKSS by adding the capability to perform local hybrid plasma simulations (particle ions and electron fluid) in the framework of the global MHD model. This will allow us to investigate micro-scale phenomena, e.g. particle acceleration and kinetic instabilities, in the realistic environment.

PUBLICATIONS

Heerikhuisen, J., E.J. Zirnstein, and N. Pogorelov, Kappa-distributed protons in the solar wind and their charge-exchange coupling to energetic hydrogen. J. Geophys. Res. Space Phys., 120:3 (2015), pp. 1516-1525.

Fermo, R. L., N.V. Pogorelov, and L.F. Burlaga, Transient shocks beyond the heliopause. J. Phys. Conf. Ser., 642 (2015), 012008.

Luo, X., et al., A Numerical Simulation of Cosmic-Ray Modulation near the Heliopause. Astrophys. J., 808:1 (2015), p. 802.

Manoharan, P., et al., Modeling solar wind with boundary conditions from interplanetary scintillations. J. Phys. Conf. Ser., 642 (2015), 012016.

Pogorelov, N. V., S.N. Borovikov, H. Heerikhuisen, and M. Zhang, The heliotail, Astrophys. J. Lett., 812:1 (2015), L6.

Pogorelov, N. V. and S.N. Borovikov, Mixing of the Interstellar and Solar Plasmas at the Heliospheric Interface. In Numerical Modeling of Space Plasma Flows: ASTRONUM-2014, 498 (2015), pp. 160-167.

Pogorelov, N. V., The Heliotail: Theory and modeling. J. Phys. Conf. Ser., 719 (2016), 012013.

Zhang, M., X. Luo, and N. Pogorelov, Where is the cosmic-ray modulation boundary of the heliosphere? Phys. Plasmas, 22:9 (2015), 091501.

Zirnstein, E. J., et al., Local Interstellar Magnetic Field Determined from the Interstellar Boundary Explorer Ribbon. Astrophys. J., 818:1 (2016), L18.

DEPLOYMENT OF THE DARK ENERGY SURVEY WORKFLOWS

PI: Donald Petravick[1]
Collaborators: Gregory Daues[1], Robert Gruendl[1], and Felipe Menanteau[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

The Dark Energy Survey (DES) is performing a 5,000 square-degree wide field survey in five optical bands of the southern sky and a 30 square-degree deep supernova survey with the aim of understanding the nature of dark energy and the accelerating universe. DES uses the new three square-degree charged couple device (CCD) camera, DECam, installed at the prime focus on the Blanco four meter telescope to record the positions and shapes of 300 million galaxies up to redshift 1.4. During a normal night of observations, DES produces about 1 TB of raw data, including science and calibration images, which are transported automatically from Chile to NCSA to be archived and reduced. The DES data management system (DESDM) is used for the processing, calibration and archiving of this data, which has been developed by collaborating DES institutions led by NCSA. The DESDM team at NCSA has successfully deployed the DES Production Pipeline on Blue Waters. Over the course of a year of investigations, several software and network improvements were made by the Blue Waters team to accommodate our workflows, and we were able to commission our production framework successfully in the fall of 2015. Moreover, we were able to make use of the remainder of our initial allocation to process 10,814 DECam exposures on the XE Compute Nodes. This corresponds to 15% of the total data volume (over 70k exposures) that DESDM processed for the Y3A1 release.

INTRODUCTION

The goal of the DES is to understand the origin of cosmic acceleration and the nature of dark energy using four complementary methods: weak gravitational lensing, galaxy cluster counts, large-

FIGURE 1: An example of one of the 10,814 DECam exposures processed by Blue Waters in early 2016 using the FINALCUT pipeline during our initial DD time allocation.

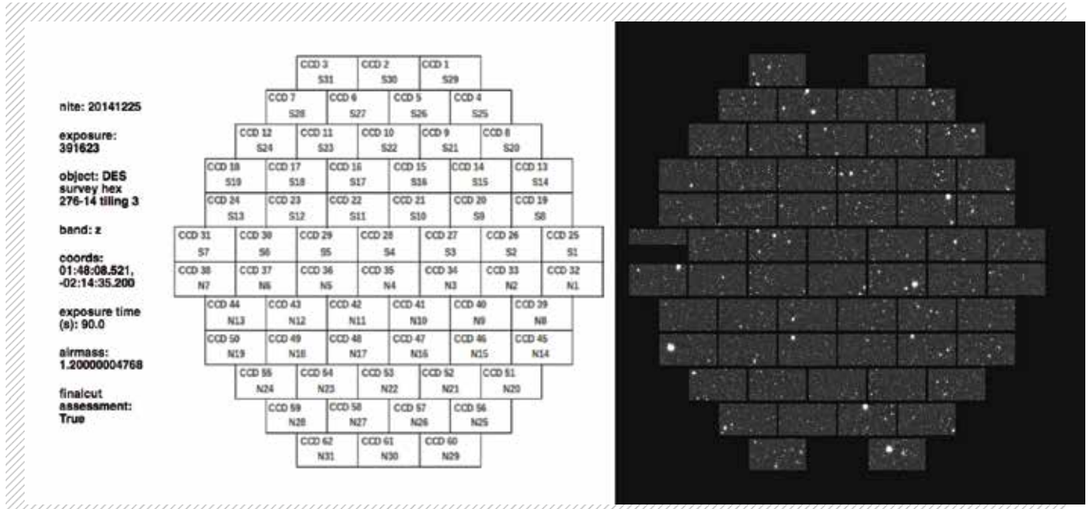

FIGURE 2: An example of a section near the rich galaxy cluster RXJ2248 from a deep stacked co-add tile color image using the *g,r,i* filters and that was generated using the DESDM pipeline.

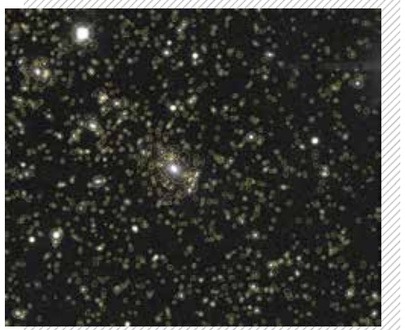

FIGURE 3: The same tile section showing ellipses for the objects detected and catalogued by the DESDM pipeline.

scale galaxy clustering (including baryon acoustic oscillations), and Type Ia supernovae. DES comprises two multi-band imaging surveys; a 5,000 square degree g,r,i,z,Y wide-survey of the southern sky to approximately 24th mag and a deeper time-domain 30 square-degree g, r, i, z deep DES Supernova (SN) Survey with a cadence of approximately five days.

DES uses the state-of-the-art three square-degree Dark Energy Camera (DECam), a 570 megapixel camera installed at the prime focus of the Blanco four meter telescope at the Cerro Tololo Interamerican Observatory (CTIO) in Northern Chile. DECam consists of 62 fully depleted, 250 micron thick 2048×4096 CCDs combined with four 2048×2048 guider and eight 2048x2048 autofocus CCDs.

For 525 nights, between 2013-2018, DES is scanning the sky to perform a 5,000 square-degree wide field survey. Over five observing seasons, DES will measure shapes, positions, fluxes, and colors for approximately 300 million galaxies and will discover and measure light curves for 3,500 supernovae and use these measurements to deliver powerful, new constraints on cosmic acceleration and dark energy. Each image arrives from CTIO in Chile to NCSA within minutes of being observed and it is usually processed by the nightly processing pipeline within 24 hours.

METHODS & RESULTS

One of main goals of the DESDM Project within the DES collaboration is the operation of the data reduction pipelines using high-performance computing (HPC) facilities to generate the survey data products. The DESDM data reduction process is composed of several pipelines or workflows that, starting from raw images, remove known instrumental signatures and sky background, masks defects, and detects and measures properties of objects to produce catalogs and calibrated images.

In 2014, we requested a Blue Waters allocation to explore the feasibility of running DESDM software and workflows on a shared Track-1 system. After nearly a year of investigations, where several software and network improvements were made by the Blue Waters team to accommodate our workflows, we were able to successfully deploy our production framework in the fall of 2015.

After each observing DES season, all DECam exposures meeting the survey data-quality criteria are systematically reprocessed using the FINALCUT and COADD pipelines. During the first trimester of 2016, we needed more computing resources to process the single-epoch FINALCUT campaign for the year three annual release (Y3A1) and used the remainder of our allocation to process 10,814 DECam exposures on the XE Compute Nodes. This corresponds to 15% of the total data volume (over 70k exposures) that DESDM processed for the Y3A1 data release. In Figure 1, we show an example of a DECam exposure processed using Blue Waters. In Figures 2 and 3, we show a section of a deep stacked co-add DES tile near the center of the cluster RXJ2248.

The DESDM system relies on the HTCondor software (CHTC UW-Madison) to manage jobs within a directed acyclic graph-based workflow. Porting this to Blue Waters represents a new and novel use of this system. Lessons from our implementation could serve as a model for other astronomy projects and pipeline that are not currently using Blue Waters.

WHY BLUE WATERS

The yearly re-processing campaigns of DESDM impose large seasonal variations in the demand for computing resources. The petascale size of Blue Waters, coupled with the perfectly parallel nature of our pipelines, allows DESDM to elastically expand its working pool of compute nodes to accommodate the burst in demand arising from the year annual processing for the collaboration.

DESDM is led by NCSA, where all images are archived and served in the community. Data processing on Blue Waters at NCSA is more robust and preformat than distributing workloads to remote sites. Moreover, the proximity between DESDM scientists and Blue Waters staff enables rapid feedback and clear communication, which are important for the success of complex implementations.

NEXT GENERATION WORK

We would like to continue production of our yearly annual releases on Blue Waters, which will continue to increase in size with each observing season. Our last data release is scheduled for 2021.

We want to improve the level of automation of HTCondor-based workflow on Blue Waters for the subsequent production campaigns.

We are also interested in the use of User Defined Images (UDI) for standard deployment of our software stacks to HPC systems.

AB INITIO MODELS OF SOLAR ACTIVITY

Allocation: NSF PRAC/9.01 Mnh
PI: Robert Stein[1]
Co-PIs: Mats Carlsson[2], Aake Nordlund[3], William Abbett[4], and Bart De Pontieu[5]
Collaborators: Viggo Hansteen[2]

[1]Michigan State University
[2]Oslo University
[3]Copenhagen University
[4]University of California Berkley
[5]Lockheed-Martin Space Sciences Lab

EXECUTIVE SUMMARY

The goal of this project is to understand how solar magneto-convection forms and controls solar active regions–how magnetic flux emerges through the solar surface and how that contributes to the heating of the chromosphere and corona and the acceleration of charged particles into the interplanetary medium.

INTRODUCTION

This project is motivated by society's vulnerability to harmful space weather. X-ray bursts and the high-energy particles associated with intense solar activity can harm astronauts, disable satellites, and hamper terrestrial systems for communication, guidance, and power distribution. Earth's heliospheric environment is controlled by magnetic fields produced by a subsurface convective dynamo. Some of the fields produced emerge through the solar surface into its atmosphere. Convective motions move these field foot points around in the photosphere. This tangles them and causes reconnection higher in the sun's atmosphere, heating the chromosphere and corona. Reconnection of active region fields produce X-rays and accelerate charged particles to high energies and drives them from the Sun into the heliosphere.

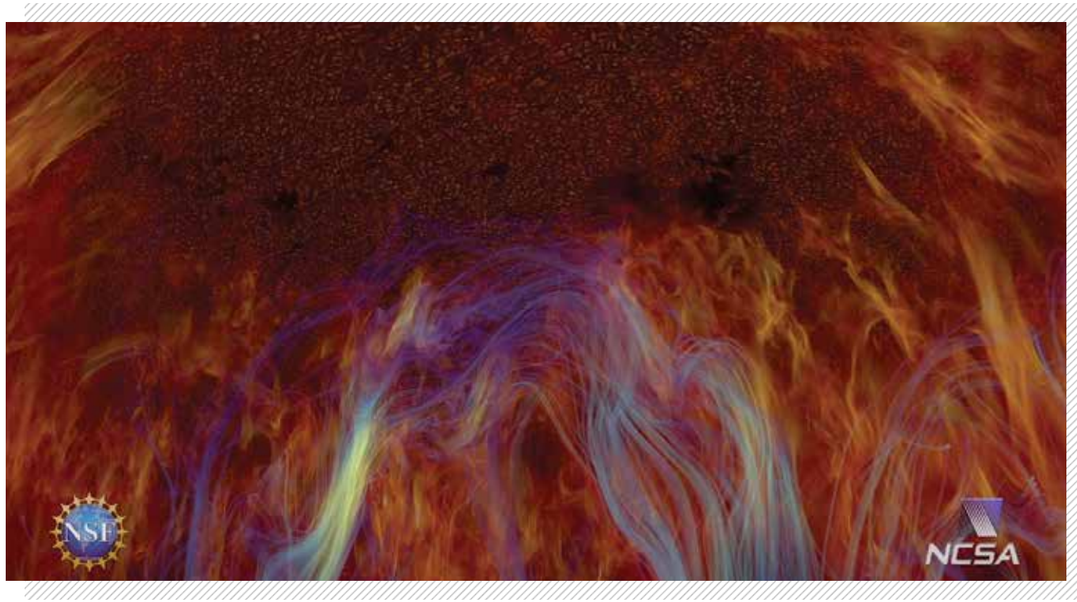

FIGURE 1: Image based on data from STAGGER calculations, from the planetarium show "Solar Superstorms" produced at NCSA showing the emergence of magnetic field lines through the solar surface.

METHODS & RESULTS

A finite-difference code is used to numerically solve the conservation equations for mass, momentum, internal energy. The induction equation for the magnetic field and the equation for non-gray radiation transport are used to model the magneto-convection and the behavior of the overlying chromosphere and corona.

The bulk of this project year was spent finding and removing an incompatibility introduced between the interior calculation scheme and the bottom boundary conditions. It was finally possible to start the relaxation of a 192 Megameter (Mm)—one million meter—by 20 Mm deep model of the outer solar convection zone using a grid of 4,032 by 4,032 by 500. This large size make it possible to use data from interior dynamo simulations to control the bottom boundary conditions for the study of active region emergence.

WHY BLUE WATERS

Blue Waters was particularly well suited for this task. Although it is up to four times slower on a per core basis than the newer Pleiades supercomputer, it was possible to run efficiently on 2,016 nodes or 64,512 integer cores on Blue Waters, whereas only 4,032 cores can be used on Pleiades because of excessive queue wait times for larger jobs.

NEXT GENERATION WORK

A new code, DISPATCH, is being developed which uses task based scheduling of small Cartesian patches that evolve with local time stepping. This type of code is ideal for obtaining **unprecedented** performance and scaling the current petascale and the near-future exascale computing eras, and with its large number of nodes and emphasis on a small number of large-scale projects, the Blue Waters system is particularly well suited for DISPATCH.

EVOLUTION OF THE SMALL GALAXY POPULATION FROM HIGH REDSHIFT TO THE PRESENT

Allocation: NSF PRAC/11.9 mnh
PI: Thomas Quinn[1]
Co-PI: Fabio Governato[1]
Collaborators: Lauren Anderson[1], Michael Tremmel[1], Sanjay Kale[2], and Harshitha Menon[2]

[1]University of Washington
[2]University of Illinois at Urbana-Champaign

FIGURE 1: A large galaxy from our cosmological simulations is modeled with different active galactic nucleus (AGN) feedback and dynamics. The morphology of the galaxy is significantly different depending on whether we use (from left to right) our new AGN model, no AGN feedback, poor black hole dynamics or a less sophisticated accretion model.

EXECUTIVE SUMMARY

Creating robust models of the formation and the evolution of galaxies requires the simulation of a cosmologically significant volume with sufficient resolution and subgrid physics to model individual star-forming regions within galaxies. This project aims to perform such modeling with the specific goal of interpreting Hubble Space Telescope observations of high redshift galaxies. We are using the highly scalable N-body/Smooth Particle Hydrodynamics code, ChaNGa, based on the Charm++ runtime system on Blue Waters to perform a simulation of a 25 Mpc cubed volume of the universe with a physically motivated star formation/supernovae feedback model. This past year's accomplishments include incorporating realistic black hole formation, dynamics, and feedback into our model which allows us to create and advance more realistic modeling of the knee in the galaxy luminosity function at high redshift. Comparisons with Hubble data show that we can reproduce the high redshift galaxy population with models that also reproduce the morphologies of present-day galaxies.

INTRODUCTION

The cold dark matter (CDM) paradigm for structure formation has had many successes over a large range of scales, from cosmic microwave background fluctuations on the scale of the horizon to the formation and clustering of individual galaxies. However, at the low end of the galaxy luminosity function, the CDM theory and observations are somewhat at odds. In particular, the existence of bulgeless, cored small galaxies is not a natural prediction of CDM. However, these are the scales where the baryonic physics of gas cooling, star formation, and feedback can significantly impact the overall mass of the galaxy. Furthermore, accurately modeling the star formation process requires a spatial resolution of order 100 parsecs to resolve the molecular star-forming regions of the interstellar medium. On the other hand, survey volumes addressing small galaxies, including recently approved Hubble Space Telescope (HST) programs, are over 10,000 cubic Mpc. Only with large simulations can we perform proper comparisons with these programs to address the following basic issues of the CDM model:

- Does the standard ΛCDM model produce the correct number densities of galaxies as a function of mass or luminosity?
- What role to these galaxies play in the evolution of the baryons in the Universe?
- How do these galaxies relate to the galaxies we can study in detail in the local universe?

METHODS & RESULTS

We used the highly scalable N-body/smooth particle hydrodynamics code ChaNGa to simulate the volumes surveyed by HST with sufficient resolution to make robust predictions of the luminosity function, star formation rate, and morphologies appropriate for these surveys. The results of the simulations were processed by our parallel data reduction pipeline that creates simulated observations. These results can be directly compared with results from observational programs.

Results from our simulations indicate that we can reproduce the high-redshift galaxy luminosity functions observed by Hubble. The simulations also predicted the numbers of galaxies fainter than those observed so far. Depending on how much UV radiation can escape from these low-mass galaxies, these galaxies can produce enough stellar radiation to completely reionize the intergalactic medium. This reionization occurs on a timescale that is consistent with observations of absorption in high redshift quasars and polarization of the cosmic microwave background.

We also made significant improvements in the algorithms used for the treatment of Black Hole formation, dynamics, and feedback. With this new treatment, we can reproduce observed star formation and black hole accretion histories, star formation histories and supermassive black hole masses in Milky Way-sized galaxies over the age of the universe.

WHY BLUE WATERS

We used the same mass and resolution requirements from our previous resolution tests to reliably model galaxy form and structure. Therefore, the size of the simulations we performed are set by the sub-volume of the universe we wish to model. HST surveys of high-redshift, star-forming galaxies cover a volume comparable to a cube 25 Mpc. This volume not only allowed us to make direct comparisons with HST surveys but also to enhance their value by providing a better understanding as to how these galaxies evolve to the present.

Proper treatment of black hole dynamics also requires high dark matter mass resolution. This treatment is needed so that the sinking of black holes to the centers of galaxies is not interrupted by two-body scattering on the dark matter. Combining these two resolution requirements with the complete simulation requires an order of magnitude more computationally intensive than what we could previously use. Therefore, a sustained petascale facility like Blue Waters was necessary.

NEXT GENERATION WORK

Nevertheless, this simulation was still a compromise. For example, if we wish to understand how high-redshift galaxies influence surrounding intergalactic gas, a much larger volume of gas will be needed in our simulations. Intergalactic gas is studied from absorption observations of background quasars using the HST Cosmic Origins Spectrograph. Statistical samples of this gas require a volume of order 60 Mpc; over an order of magnitude larger than our current simulation. Simulations made with samples of this magnitude is the only way we will be able to understand the extent to which star formation, supernovae, and active galactic nuclei in individual galaxies influence the surrounding gas, and to conduct a proper census of the majority of the baryonic matter in the universe. The next generation of Track-1 computational resources will be required to run simulations using higher volumes of intergalactic gas.

PUBLICATIONS AND DATA SETS

Menon, H., et al., Adaptive techniques for clustered N-body cosmological simulations. *CompAC*, 2 (2015), 1.

Tremmel, M., Governato, F., Volonteri, M., Quinn, T. Off the beaten path: A new approach to realistically model the orbital decay of supermassive black holes in galaxy formation simulations. Mon. Not. R. Astron. Soc., 451:2 (2015), pp. 1868-1874.

Loebman, S., et al., Big-data management use-case: A cloud service for creating and analyzing galactic merger trees. *DanaC: Workshop on Data analytics in the Cloud*. Snowbird, UT. June 22 - 27, 2014.

FIGURE 2:
Distribution of gas (left) and stars (right) in a group of galaxies at z = 0.25. Stars are colored by their formation time such that young stars are white. Comparing left and right images, one can find a number of gas-free galaxies with older stellar populations.

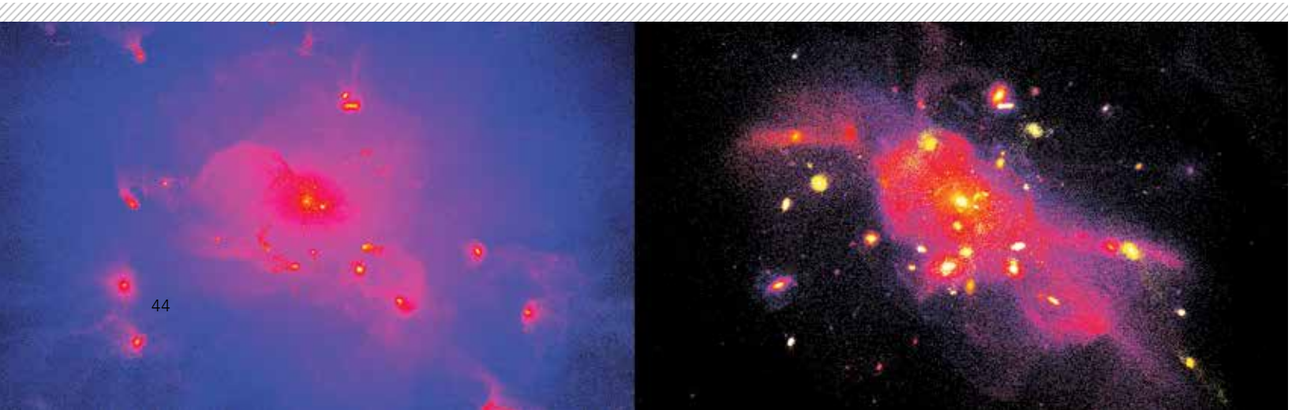

BINARY NEUTRON STARS MERGERS: A JET ENGINE FOR SHORT GAMMA-RAY BURSTS

Allocation: Illinois/958 Knh
PI: Stuart L. Shapiro[1]
Co-PIs: Ryan Lang and Milton Ruiz[1]
Collaborator: Vasileios Paschalidis[2]

[1]University of Illinois at Urbana-Champaign
[2]Princeton University

EXECUTIVE SUMMARY

The LIGO-Virgo Collaborations recently reported the first direct detection of a gravitational-wave (GW) signal produced by the inspiral and coalescence of a binary black hole (BHBH) system. This breakthrough marks the beginning of the era of gravitational-wave astrophysics. GWs are expected to be generated not only by BHBH binaries but also by neutron star-neutron star (NSNS) and black hole-neutron star (BHNS) binaries. Merging NSNSs and BHNSs are not only important sources of GWs but also the two most popular candidates to produce relativistic jets and serve as the engines which power short-hard gamma-ray bursts (sGRBs). Simultaneous observation of GWs and gamma rays from these systems is the holy grail of multimessenger astronomy. We have performed ideal magnetohydrodynamic simulations of NSNS systems in full general relativity and have shown unambiguously that they can indeed launch incipient jets even when the initial B field is confined to the interior of the stars.

INTRODUCTION

A century after the General Relativity (GR) was published, the LIGO-VIRGO collaboration reported, for the first time, the direct detection of the GWs (event GW150914) [1]. This detection provided a spectacular confirmation of GR theory as the fundamental theory of gravitation and confirmed the existence of BHs and BHBHs. Most importantly, this breakthrough opens a new window to our universe, as GWs can provide us with information that cannot be obtained from the typical electromagnetic spectrum. This includes the observation of BHBH and measurement of their properties, the study of the early universe before the recombination era, as well as the nature of matter above nuclear density. Also, GW signals are expected to be generated not only by BH binaries but also from NSNSs and BHNSs, among other compact objects. Many of these sources are likely to also generate electromagnetic (EM) radiation counterparts to the GWs. Detecting both GW and EM radiation from the same cosmic source will constitute a major advance in multimessenger astronomy.

GRBs were first discovered in 1967, and theorists have been working to explain them ever since. The mergers of a NS with a companion NS or BH are the two most popular candidate progenitors of short gamma ray bursts, those with a duration less than two seconds. These systems are thus excellent candidates for multimessenger detection. To verify the binary-short gamma ray bursts association and properly interpret the GW and EM signals we will receive, we need to model these systems and simulate their evolution in full general relativity with magnetohydrodynamics (GRMHD). Our work to date has focused on studying the merger of magnetized BHNS and NSNS systems.

METHODS & RESULTS

GRMHD numerical simulations require the solution of the GR equations to determine the gravitational field, the relativistic MHD equations to determine the flow of matter, and Maxwell's equations to determine the electromagnetic fields. Together the equations constitute a large system of highly nonlinear, multidimensional, partial differential equations in space and time.

Recently, we demonstrated that mergers of magnetized BHNS systems can launch jets and be the engines that power short gamma ray bursts [2]. The key ingredient for generating a jet was found

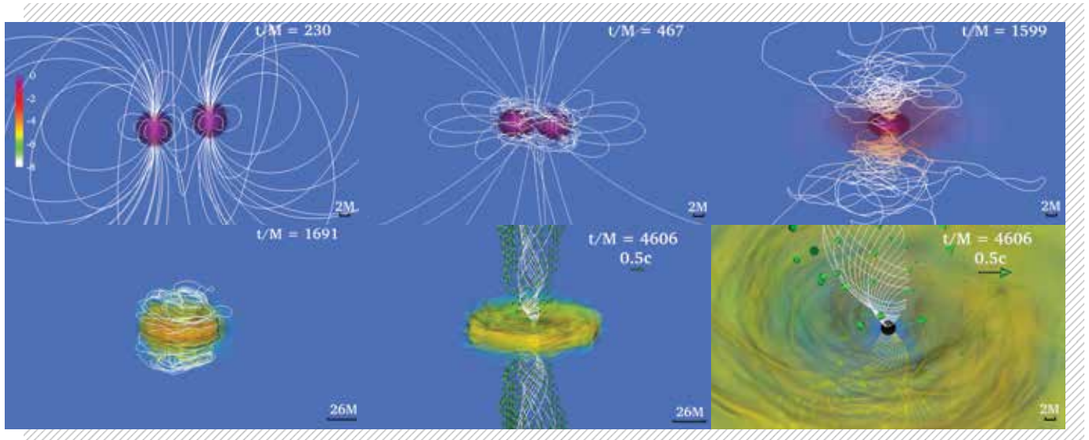

FIGURE 1: Snapshots of the rest-mass density, normalized to its initial maximum value (log scale), at selected times for the Pulsar-like case. Arrows indicate plasma velocities and white lines show the magnetic field lines. Bottom panels highlight the system after an incipient jet is launched.

to be the initial endowment of the NS with a dipole B field that extends into the NS exterior, as in a pulsar magnetosphere. By contrast, if the initial B field is confined to the interior of the NS, no jet is observed [3]. These results prompt the query of whether NSNS mergers produce jets in the same way as BHNS systems, or does of the mechanism requires an initial BH.

Previous ideal GRMHD simulations by Rezolla et al. suggest that NSNS mergers may launch a relativistic jet [4], while those by Kiuchi et al., which focus on different initial configurations, show otherwise [5]. Both of these studies have considered only scenarios where the B-field is initially confined to the interior of the two NSs.

Using our latest adaptive-mesh refinement GRMHD code we performed simulations of equal mass NSNS binaries initially on a quasicircular orbit that undergo merger [6]. The initial stars are irrotational, n=1 polytropes and are magnetized. We explore two types of B field geometries: one where each star is endowed with a dipolar B field extending from the interior into the exterior (Fig. 1, upper left), as in a pulsar, and the other where the dipolar field is initially confined to the interior. In both cases, the adopted B fields are dynamically unimportant initially. To ensure reliable evolution of the exterior field and properly mimic the conditions that likely characterize the exterior magnetosphere, in the pulsar-like case, we initially imposed a low-density atmosphere such that the NS exterior was described by a constant plasma parameter, defined as the ratio of the gas pressure to the magnetic pressure. We then continued the evolution through the merger of the two NS and subsequent formation of a hypermassive NS (Fig. 1, upper center and right) that undergoes delayed collapse to a BH immersed in a magnetized accretion disk (Fig. 1, lower left and center).

Unlike in the BHNS case, the B field does not grow following BH formation: the existence of the HMNS phase instead allows the B field to build to saturation levels prior to BH formation. We do observe a gradual growth in the ratio B2/2ρ above the BH pole due to the emptying of the funnel as matter accretes onto the BH. At ~66 ms, following the merger of the two NSs, incipient jets are launched even when the initial B field is confined to the interior of the stars (Fig. 1, lower right). The duration of the accretion and the lifetime of the jet is ~0.1 seconds, which is consistent with short-duration short gamma ray bursts. The luminosity is 1051 erg/s, also consistent with observed short gamma ray bursts values. Our simulations provide theoretical corroboration that mergers of NSNS systems can indeed launch jets and be the central engines that power short gamma ray bursts.

WHY BLUE WATERS

By adding OpenMP support to our message interface passing (MPI)-based code, scalability on multi-core machines has improved greatly. With the Blue Waters next-generation interconnect and processors, our hybrid OpenMP/MPI code exhibits greater scalability and performance than on any other supercomputer we have used. Recently, we were able to build our code with the Blue Waters Intel compilers. This resulted in a 30% boost of our code's performance, making Blue Waters unique for tackling the astrophysical problems we want to address.

Blue Waters is also used by our undergraduate research team to make visualizations (e.g., Fig. 1) and movies of our simulations with the VisIT software.

PUBLICATIONS

Gold, R., V. Paschalidis, Z. B. Etienne, S. L. Shapiro, and H. P. Pfeiffer, Accretion disks around binary black holes of unequal mass: GRMHD simulations near decoupling. *Phys. Rev. D*, 89 (2014), 064060.

Ruiz, M., V. Paschalidis, and S. L. Shapiro, Pulsar spin-down luminosity: Simulations in general relativity. *Phys. Rev. D*, 89 (2014), 084045.

Etienne, Z. B., et al, Improved moving puncture gauge conditions for compact binary evolutions. *Phys. Rev. D*, 90 (2014), 064032.

Gold, R., V., et al., Accretion disks around binary black holes of unequal mass: GRMHD simulations of postdecoupling and merger. *Phys. Rev. D*, 90 (2014), 104030.

Paschalidis, V., M. Ruiz, and Shapiro S. L., Relativistic simulations of black hole-neutron star coalescence: the jet emerges. *ApJ*, 806 (2015), L14.

M. Ruiz, R. N. Lang, V. Paschalidis, and S.L. Shapiro, Binary neutron star mergers: a jet engine for short gamma-ray burst. *Astrophys. J. Lett.*, 824.1 (2016), L6.

THE MOST MASSIVE GALAXIES AND BLACK HOLES AT THE COSMIC DAWN OF THE UNIVERSE

Allocation: NSF PRAC/2.89 Mnh
PI: Tiziana Di Matteo[1]
Co-PIs: Yu Feng[2] and Rupert Croft[1]
Collaborators: Dacen Waters[1], Steven Wilkins[3]

[1]Carnegie Mellon University
[2]University of California Berkeley
[3]University of Sussex

EXECUTIVE SUMMARY

Our team has led the development of cosmological codes adapted to petascale supercomputers and used Blue Waters to understand how the first massive black holes and galaxies were formed, from the smallest to the rarest and most luminous. Using nearly one trillion particles, we have carried out the BlueTides simulation on Blue Waters. BlueTides can answer questions in cosmology which require simulations of the entire visible universe at high mass and spatial resolution. We can directly predict what should be seen in future observations that will probe the cosmic dawn of the universe and the formation of the first galaxies and black holes.

INTRODUCTION

Survey astronomy has enabled the study of galaxy and large-scale structure formation at low redshifts to mature into a precise science. Current galaxy surveys at high redshift, however, have covered very small volumes of space in the early universe during the epoch of formation of the first galaxies and quasars. The search for the earliest objects in the universe is extremely challenging. From the observational point of view, the field will be transformed by the next generation telescopes (JWST, WFIRST, etc.). In the coming decade, a new generation of astronomical instruments will observe the universe at the time of the formation of the first stars and quasars, opening up the "last frontier" in astronomy and cosmology.

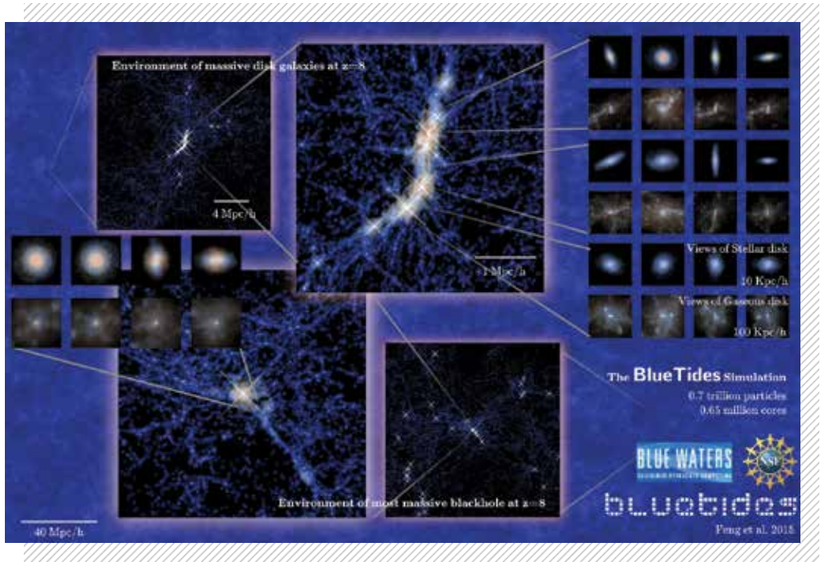

FIGURE 1 : Image of the BlueTides simulation. The large scale density field of the early universe is shown in the background. The insets show zooms of two regions of the large scale density where the first disk galaxies form (galaxies shown in the small left images) and the most massive black hole (left side).

Numerical simulations of galaxy formation have been limited by the volume at which they can evolve. From the simulation-theoretical perspective, large-scale uniform volume hydrodynamic simulations of the high redshift universe are a problem ideally suited to modern petascale facilities like Blue Waters. It is now feasible to run memory-limited computations with the resources that computer time panels can allocate, and therefore reach **unprecedented** volumes and resolutions in the early universe. We have carried out a program to simulate high redshift quasars and first galaxies using Blue Waters and the MP-Gadget cosmological hydrodynamic simulation code. Radical updates to the code efficiency, and hydrodynamic formulation and star formation modeling allow us to meet the challenge of simulating the next generation fields and effectively utilize Blue Waters. The BlueTides simulation is unique in that it probes directly the range of scales (masses and epochs) of galaxies and quasars that will be discovered shortly.

METHODS & RESULTS

Massive disk galaxies in the early universe

Our current run of BlueTides allowed us to apply observational selection algorithms (SourceExtractor) to the simulated sky maps and build catalogs of millions of galaxies. The high resolution of BlueTides made it possible to produce detailed images of individual galaxies and uncover a striking and unexpected population of large Milky Way-sized disk galaxies (Fig. 2) present when the universe was 5% of its present age. Both of these achievements will be of great benefit to the burgeoning frontier fields that will utilize JWST and WFIRST.

The origin of the first massive black holes

Quasars, powered by supermassive black holes, represent a fascinating and unique population of objects at the intersection of cosmology and astrophysics. The growth of the most massive black holes in the early universe, consistent with the detection of highly luminous quasars before the universe is a billion years old, implies that a sustained accretion of material is required to grow and power them. Given a black hole seed scenario, the question remains as to what the fundamental condition in the early universe was to allow the fastest black hole growth. BlueTides allows us to explore the conditions conducive to the growth of the earliest supermassive black holes. The most massive black holes approach 100 million solar masses at z=8 and are found in extremely compact spheroid-dominated host galaxies. More importantly, the role of the initial tidal field sets the condition for early black hole growth. In regions of low tidal fields, gas accretes

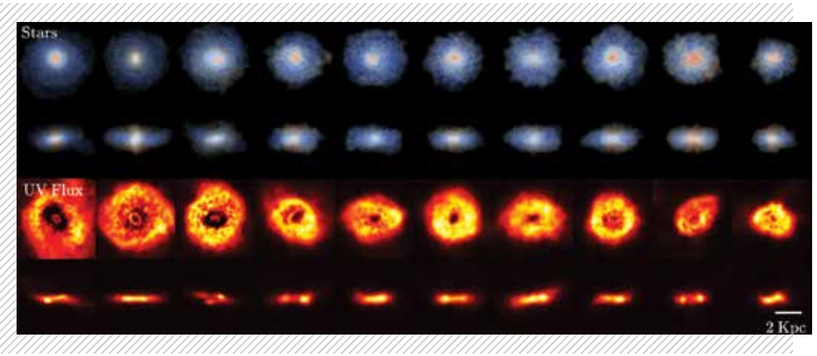

FIGURE 2: A sample of disk galaxies selected from the BlueTides simulation at redshift z=8. We show galaxies identified kinematically to be rotating disks. For each galaxy we show a face-on and side-on view. The colors of the top two rows represent the stellar surface density coded by the stellar age, face-on and side-on. The colors of the bottom two rows represent the star-formation surface density. The face-on images have included the effect of dust-extinction.

"cold" onto the black hole and falls along thin, radial filaments straight into the center, forming the most compact galaxies and most massive black holes at earliest times (Fig. 2). In regions of high tidal fields, larger, more coherent angular momenta influence the formation of the first population of massive compact disks (Fig. 1).

WHY BLUE WATERS

A complete simulation of the universe at the epochs we are studying requires a small enough particle mass to model the dwarf galaxies which significantly contribute to the summed ionizing photon output of all sources. It also calls for an enormous volume: 1 cubic Gigaparsec (3x1019 cubic light years) to capture the rarest and brightest objects—the first quasars. The first requirement is, therefore, equivalent to a high particle density and the second to a large particle volume.

Blue Waters made this qualitative advance possible, producing arguably the **first** complete simulation (at least regarding hydrodynamics and gravitational physics) of the creation of the earliest galaxies and large-scale structures in the universe. Our application runs required almost full system capacity: we used 20,250 nodes (648,000 core equivalents—the new version of the code can scale higher, but we left a safety margin) using 57 GB/node (89%). This application thus uses 1.15 PB of memory—something only Blue Waters can provide, and which is **90% of the available memory**.

Our project establishes a theoretical framework for understanding the cosmic dawn of the universe, the results of which will benefit a broad spectrum of scientific communities, including cosmology, high-energy astrophysics, and anybody studying galaxy evolution. This study will provide answers to broad questions about the history and future evolution of the universe and the formation of structure.

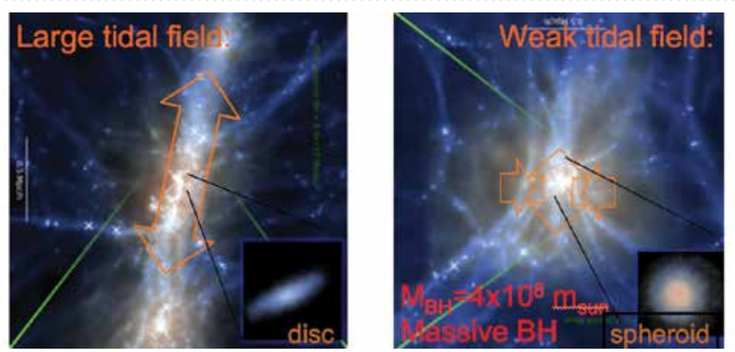

FIGURE 3: The most massive disk galaxy environment is compared to the most massive black hole and host galaxy. Tidal fields are strong and form large scale filaments and the formation of first disk galaxies (top), low tidal fields allow the formation of the most massive black holes in the center (bottom panels).

NEXT GENERATION WORK

BlueTides is a **path-finder** for developing methods and calculations for future cosmological hydrodynamic simulations of galaxy formation with volumes and resolutions suitable for creating models for next-generation surveys. Our simulations are important as they blaze a trail for future calculations in the future Track 1 systems. Currently, we are evolving models forward in time for only one billion years, rather than the 14 billion years necessary to cover the history of the universe to the present day that next generation systems, assuming they have enough memory capacity, would enable.

PUBLICATIONS AND DATA SETS

Croft, R., T. Di Matteo, N. Khandai, and Y. Feng, Petascale cosmology: simulation of structure formation. *Comput. Sci. Eng.*, 17:2 (2015), DOI:10.1109/MCSE.2015.5

Y. Feng, M. Straka, R. Croft, T. Di Matteo, Sorting At Scale on BlueWaters 2015, Proceedings of Cray User Group Conference 2015.

Y. Feng, et al., The formation of Milky Way-mass disk galaxies in the first 500 million years of a cold dark matter universe. *APJ Lett.*, 808:1 (2015), L17.

Y. Feng, et al., BlueTides: First galaxies and reionization. *Mon. Not. R. Astron. Soc.*, 455:3 (2016), pp. 2778-2791.

Waters, D., et al., Monsters in the Dark: Predictions for Luminous Galaxies in the Early Universe from the BlueTides Simulation. *Mon. Not. R. Astron. Soc.: Lett.*, (2016), slw100.

Wilkins, S.M., et al., The Lyman-continuum photon production efficiency in the high-redshift Universe. *Mon. Not. R. Astron. Soc.: Lett.*, 458:1 (2016), L6-L9.

Waters, D., et al., Forecasts for the WFIRST High Latitude Survey using the BlueTides Simulation. *arXiv preprint arXiv*:1605.05670 (2016).

Wilkins, S.M, et al., The Photometric Properties of Galaxies in the Early Universe. *Mon. Not. R. Astron. Soc.*, (2016): stw1154.

CORE-COLLAPSE SUPERNOVAE THROUGH COSMIC TIME

Allocation: NSF PRAC/8.94 Mnh
PI: Eric J. Lentz[1]
Co-PIs: Stephen W. Bruenn[2], W. Raphael Hix[3], O. E. Bronson Messer[3], and Anthony Mezzacappa[1]
Collaborator: J. Austin Harris[4]

[1]University of Tennessee
[2]Florida Atlantic University
[3]Oak Ridge National Laboratory
[4]Lawrence Berkeley National Laboratory

EXECUTIVE SUMMARY

Explosions of massive stars (core-collapse supernovae; CCSNe) have a significant impact on the development of galaxies and their heavy element content. We compute 3D simulations of CCSNe that include the complex range of CCSN physics, across the range of input conditions representing the history of massive stars in the universe to obtain the variety of outcomes seen in nature. Our computations account for the appropriate nuclear processes needed to generate and eject heavy elements that are needed to form planets. We have discovered a **previously unseen** nuclear burning behavior that can only be observed through the completeness of our simulations and the included nuclear physics.

INTRODUCTION

Massive stars (mass greater than eight solar masses) are relatively rare, yet they play a significant role in the evolution of galaxies, particularly through their explosive finales as CCSNe. Energy from CCSNe triggers new star formation and elements synthesized in massive stars and CCSNe are the ingredients for terrestrial planets in those star systems. The conversion of gravitational potential energy from the collapse of the stellar core into an expulsion of the stellar envelope is a complex physical process. This physical process combines gravitation, nuclear physics, neutrino physics (neutrinos transport the needed energy to drive the explosion from the collapsed core), and turbulent fluid dynamics with a rich phenomenology.

METHODS & RESULTS

To compute CCSN models in the necessary 3D, we have developed *Chimera*, a program that accounts for neutrino transport and opacities, nuclear equations of state and reaction networks, compressible fluid dynamics, and self-gravity [1, 2]. This project addresses the wide range of pre-supernova stellar configurations from the range of initial masses and the build-up of heavy elements (mostly from previous CCSNe) through cosmic history by sampling in both dimensions. From these models we will address the nature of the CCSN mechanism and the production of elements in the explosions.

In the low-mass, primordial composition simulation in our grid, we have identified a previously unseen burning mode in stellar collapse. During collapse, compression of the silicon shell intensifies the burning of the remaining oxygen at the bottom of the layer until it triggers a silicon flash. The silicon flash burns much of the silicon shell to iron-peak elements and some of the overlying oxygen-neon shell to silicon. The deposited energy alters the collapse dynamics and helps intensify the explosion driven by neutrino heating from the collapsed iron core interior. Silicon flashes have been observed in pre-supernova stellar evolution models for similar progenitors [3], but were not seen in previous work with the same progenitor [4] as nuclear burning was not adequately included.

WHY BLUE WATERS

Computing stellar explosions in 3D requires large and long computations. Blue Waters provides the capacity needed to accommodate our simulation requirements.

NEXT GENERATION WORK

More powerful and capable machines will permit improvements in the computationally expensive portions of CCSN simulations (neutrino transport, nuclear networks, resolution). This will help to better realize the nature of the explosions and increased simulation counts to better account for the variety of inputs and outcomes.

SIMULATING PLASMA TURBULENCE FROM DRIVING SCALES TO DISSIPATIVE SCALES

Allocation: GLCPC/560 Knh
PI: Kirit Makwana[1]
Co-PIs: Fausto Cattaneo[2], Fan Guo[3], Hui Li[3], and Xiaocan Li[4]
Collaborators: Vladimir Zhdankin[5] and William Daughton[5]

[1]Center for mathematical Plasma Astrophysics
[2]University of Chicago
[3]Los Alamos National Laboratory
[4]University of Alabama-Huntsville
[5]University of Colorado-Boulder

EXECUTIVE SUMMARY

Magnetized plasmas are ubiquitous in space, astrophysical, and laboratory environments. Energy is injected into them at large scales by external forcing or some intrinsic instability of the system. This energy gets transferred to smaller scales by the nonlinear interaction of the system and is ultimately converted into heat. We simulate this entire process using both magnetohydrodynamic (MHD) and kinetic (particle-in-cell) codes for different plasma parameters. We observe striking similarities between the two codes in their energy dynamics and energy spectrum. Thin, current-sheet-like, dissipative structures are formed in both codes. A statistical analysis of their morphological characteristics is performed, to reveal that their length scales with the driving scale of the turbulence. The kinetic simulation reveals that their thickness is the skin-depth scale. The dissipation is concentrated in current sheets with magnetic-field-parallel dissipation dominant at lower plasma beta. This leads to non-thermal particle energization at low plasma beta.

INTRODUCTION

Plasma turbulence is present in a variety of laboratory, space, and astrophysical plasmas. For example, the solar corona is interspersed by magnetic field lines that are constantly moving, led by their footpoint motion [1]. This motion launches Alfven waves in the corona which then reflect, interact, and produce turbulence. Turbulence cascades energy from larger to smaller scales, where it is converted into heat. This process has the potential to explain the observed heating of solar corona and solar wind [2]. Similar processes occur or are thought to occur in laboratory tokamaks, planetary magnetospheres, accretion disks, and nearly all turbulent plasmas. Understanding the turbulent cascade process and its consequential heating is thus a very important problem with numerous applications to various systems.

Past simulations have mostly utilized the MHD framework, which is applicable at macro-length scales larger than the ion gyro-radius scale. However, in typical collisionless plasmas, the conversion of turbulent energy into heat occurs due to wave-particle interactions below this scale. To understand these interactions, kinetic simulations are required. We simulated the turbulent cascade of energy using both MHD and kinetic codes. This simulation tells us whether MHD is the correct limit of plasma turbulence at macro scales, and also where MHD breaks down and kinetic physics becomes important. We were able to analyze the dissipative, current sheet structures that form in both simulations. This project helps us in characterizing turbulent energy dissipation in the solar wind and other similar plasmas, which can be further used in interpreting their observations and making predictions.

METHODS & RESULTS

We used the MHD code PLUTO [3] and the particle-in-cell code VPIC [4]. We simulated decaying plasma turbulence by specifying an initial condition and allowing the system to decay. The initial condition was an ensemble of superimposed shear Alfven waves, same for both the codes. As these waves interact, they generate turbulence, which cascades energy to smaller scales where it is converted into heat.

We found remarkably similar decay rates in both the simulations, despite the fact that the small-scale dissipation physics in the two codes is very different. Our MHD simulations relied purely on artificial numerical dissipation, whereas the PIC simulations included all kinetic processes involved in dissipation. Still, the dissipation rate was same, indicating that the dissipation scales adjust themselves to balance the rate of energy transfer arriving from larger scales. The energy spectra also matched between the two codes. This match showed that MHD simulations of turbulence produce reliable energy dynamics.

At smaller scales, important differences appear between the two descriptions. We observed the formation of current sheets in both the codes. The thickness of current sheets in MHD depends directly upon the grid-size, which is unphysical. However, VPIC correctly identifies the current sheet thickness as the skin-depth, which is a kinetic scale. We also find the generation of a significant non-thermal tail in the particle energy distribution function in VPIC at low plasma beta (ratio of thermal to magnetic energy density), as shown in Figure 1. It is observed that particle heating mainly takes place via the work done by the parallel electric field on the current density. However, this parallel electric field is not necessarily associated with reconnection.

This project shows that we can achieve a direct comparison between the MHD and fully kinetic description of plasma turbulence. It gives us the confidence that MHD simulations of turbulence produce reliable energy dynamics. Conversely, it also shows that Blue Waters and high-performance computing are now capable of reproducing MHD results from first principle, particle-in-cell simulations. This result opens up the small-scale physics of energy dissipation to the investigation by first principle codes. We can expect to unravel the mystery of collisionless dissipation in hot plasmas, which will inform us about plasma heating and particle energization in the solar corona, solar wind and in planetary magnetospheres. This discovery will help in understanding space weather and its implications for our technology which is hugely affected by space weather.

WHY BLUE WATERS

The petascale computing ability of Blue Waters was essential for carrying out this project. We used the particle-in-cell code VPIC for carrying out the kinetic simulations. VPIC is a part of the NCSA Blue Waters Sustained Petascale Performance (SPP) suite. NCSA and Cray worked to improve compiler optimization of loops not already using optimized vector compiler intrinsic functions, optimizations to eliminate extra data copies, added FMA4 compiler intrinsic functions to improve compute performance and used Cray I/O buffering functionality. Our simulations utilized a 1176x1176x1176 cell domain with 5.0E+12 particles. Blue Waters is the only system where such a large simulation can be run. The Blue Waters staff was also helpful in the data analysis and visualization. They helped us in using Paraview to visualize the data in 3D, handling and storing our data, and troubleshooting routine technical issues that came up.

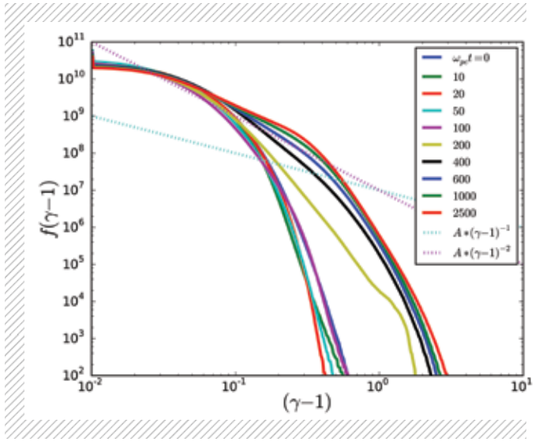

FIGURE 1: This image shows the time evolution of the particle distribution as a function of kinetic energy in a low plasma beta simulation. It begins with a Maxwellian shape initially, and particles are energized to higher energies as the turbulence dissipates with time. It develops a flattening between energies of 0.05 and 0.3, indicating a non-thermal feature.

NEXT GENERATION WORK

The next Track-1 system can help us in simulating plasmas with even lower beta, thereby giving significant non-thermal particle energization. The present system only allows simulations with an unrealistic ion to electron mass ratio of unity. The next Track-1 system can help us simulate realistic mass-ratio plasmas. These improvements will help us in studying relativistic astrophysical plasmas.

PUBLICATIONS AND DATA SETS

Makwana, K., et al., Energy dynamics and current sheet structure in fluid and kinetic simulations of decaying magnetohydrodynamic turbulence. *Physics of Plasmas*, 22 (2015), 042902.

3-D SIMULATIONS OF I-PROCESS NUCLEOSYNTHESIS IN THE EARLY UNIVERSE

Allocation: NSF PRAC/3Mnh
PI: Paul R. Woodward[1]
Collaborators: Falk Herwig[2], Chris Fryer[3], William Dai[3], Michael Knox[4], Pei-Hung Lin[5], Ted Wetherbee[6]

[1]Minnesota Institute for Astrophysics, University of Minnesota
[2]University of Victoria, B.C., Canada
[3]Los Alamos National Laboratory
[4]University of Minnesota, Laboratory for Computational Scence & Engineering
[5]Lawrence Livermore National Laboratory
[6]Fond du Lac Tribal and Community College

EXECUTIVE SUMMARY

We are exploiting the scale and speed of Blue Waters to enable 3-D simulations of brief events in the advanced evolution of stars that can have profound impacts upon their production of heavy elements. We are focusing on hydrogen ingestion flash events, because these have been identified as potential sites for the origin of observed, strongly anomalous abundance signatures in stars that formed in the early universe. In this second year of our three-year project, we have simulated H-ingestion in a very low metallicity giant star formed in the early universe. Our results reveal an unstable, global oscillatory burning of the ingested hydrogen which gives rise to quite different behavior than was anticipated from the earlier 1-D modeling of this hydrogen-ingestion flash phenomenon. We are also building simplified models of this behavior and are restructuring our PPMstar simulation code to accurately handle more challenging cases such as nuclear burning shell mergers in massive stars.

INTRODUCTION

We are interested in understanding the origin of the elements in the developing universe. The elements heavier than hydrogen and helium were manufactured within stars and later expelled into the interstellar gas to become incorporated in later generations of stars and planets. The first generations of stars played a particularly important role. The late stages of evolution of these stars can be strongly affected by hydrogen ingestion events. These events occur, for example, when a convection zone above a helium burning shell in the asymptotic giant branch (AGB) stage of evolution of such stars reaches up to unprocessed hydrogen-helium gas above it. This gas is lighter than the helium and carbon mixture of such a helium shell flash convection zone, and therefore is more buoyant. It resists being entrained into the gas of the convection zone, yet can nevertheless be dragged downward as the result of a process we call convective boundary mixing. In order to understand the H-ingestion flashes, as well as the evolution of many other types of stars, such as the pre-supernova evolution of stars that eventually explode, it is critically important to be able to quantitatively simulate this convective boundary mixing. Neutrons produced in H-ingestion flash events are captured by trace concentrations of heavy elements in the gas to build up an entire series of progressively heavier nuclei. This material can ultimately be expelled from the star along with its outer envelope as it forms a planetary nebula.

METHODS & RESULTS

Despite the relative brevity of these H-ingestion flash events in the lives of stars, simulating this entire process, even on Blue Waters, is a challenge. In the top two images in the figure, we see results for a very low metallicity AGB star of the early universe simulated on a grid of 3.6 billion cells for over 9 million time steps on 0.44 million processor cores of Blue Waters running at about 0.41 Pflop/s. The simulated time interval is nearly 2 days in the life of the star, but in order to follow this star through the flash we increased the driving helium burning luminosity by a factor of about 30 to more quickly traverse a long, slow initial ingestion period leading up to the flash itself. In the top-left image, rapid burning of locally ingested hydrogen-rich gas has

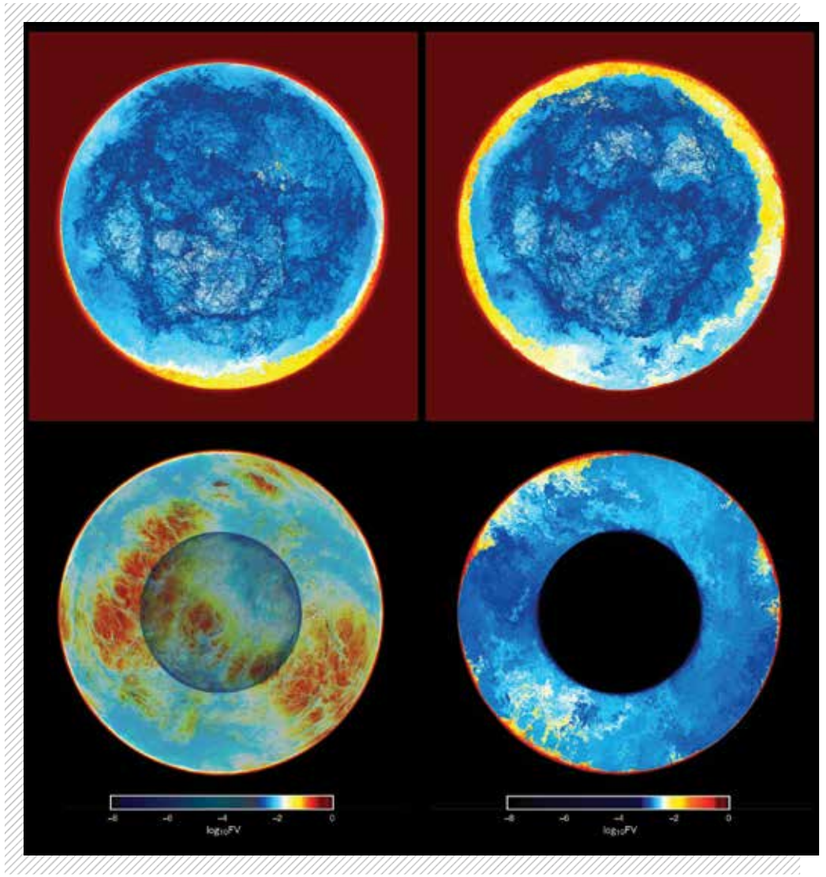

FIGURE 1: In the top 2 images the process of entrainment of hydrogen-rich gas into the helium shell flash convection zone of a very low metallicity star and the development of the Global Oscillation of Shell Hydrogen ingestion (GOSH) is shown. The front half of the star has been cut away, and the central degenerate carbon-oxygen core, which will ultimately become a white dwarf star, has been made transparent. Only mixtures of the helium and carbon gas of the convection zone with entrained hydrogen-helium gas from above it are made visible in this volume rendering. Concentrations of entrained gas from large to small range in color from red (1.6x10-2) to yellow (1x10-3), white (1.6x10-4), aqua (2.5x10-5), and finally dark blue (3x10-6). In the bottom pair of images, we see a similar display, but with a different mapping of concentration to color that is shown, for a study of entrainment of gas from above the convection zone caused by oxygen burning in a massive star (see Jones et al. 2015).

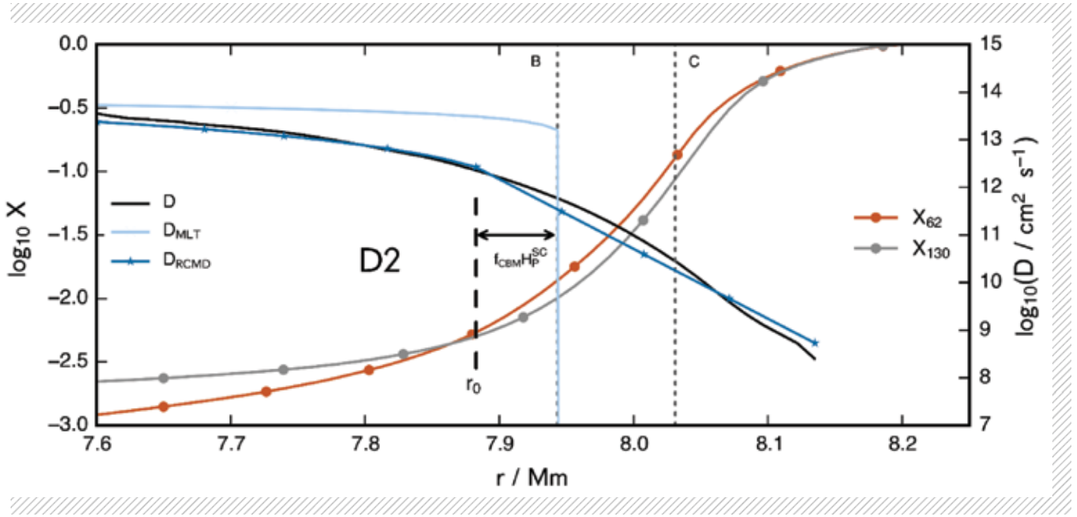

FIGURE 2: An analysis of this oxygen shell entrainment simulation in terms of 1-D stellar evolution code requirements for approximating the process in 1D. The solid black line shows an effective diffusion coefficient that produces the spherically averaged result of the 3-D simulation, the light blue line gives the diffusion coefficient as derived from standard mixing length theory, and the darker blue line gives the result for a revised 1-D model that our work recommends.

created a high pressure at the bottom right of this hemisphere of the star that is driving pressure waves outward around the star, peeling off a layer of newly entrained hydrogen-rich material as they go. About 6 minutes later, in the top-right image, these waves have reached all the way around the star, where burning of the newly entrained material will produce a new pressure wave racing back in the other direction, setting off a growing global oscillation of shell hydrogen ingestion, (GOSH). Working with a case that evolves more rapidly, a study of the oxygen burning shell in a massive star, shown in the middle of the figure, we find that a simplified model of convective mixing previously used in 1-D stellar evolution calculations by our collaborator, Falk Herwig, produces a good representation of the 1-D averaged behavior if we use the correct coefficients for the model, as shown at the bottom of the figure. In the coming year we will experiment with a 1D-3D simulation approach in which we use 3-D simulations performed at fairly widely spaced time intervals to determine the proper model coefficients that enable a 1-D stellar evolution approach to carry the simulation to the threshold of truly 3-D behavior, where we must proceed without the help of the 1-D model.

Our team is now working intensely to restructure our code completely, so that it can serve us well for the next decade. The design of the new code explicitly counts upon the very large numbers of simultaneous threads that must execute on the nodes of modern machines. These are exploited in a new approach to dynamic load balancing both on the node and over the machine interconnect. Our stellar hydrodynamics problems do not require many grid refinement levels, and therefore we are able to introduce simplifications into our design for just three such levels that will enable our runs to scale with high efficiency to over 10,000 nodes, even on machines whose nodes support many more simultaneous threads than the present Blue Waters.

WHY BLUE WATERS

3-D simulations like the ones shown here, made possible by the sustained petaflops computing capability of Blue Waters, allow us to explore and understand phases of stellar evolution that are very long in duration compared to a stellar explosion, but are nevertheless very short compared to those behaviors which are well approximated by 1-D stellar evolution codes. With the insights provided by these simulations, we are building new 1-D models and 1D-3D computational techniques that will allow us to follow stars through these stages of their lives with greater confidence.

NEXT GENERATION WORK

One challenging goal of our present code developments is to enable the accurate simulation at affordable cost of the merger of two nuclear burning shells initially separated by a convection zone above which is a thin stable layer before the next burning shell begins. On a future Track-1 system, we would like to follow a massive star through multiple hydrogen ingestion and shell merger events to the point just before its core collapses and a supernova explosion results.

PUBLICATIONS AND DATA SETS

S. Jones, R. Andrassy, S. Sandalski, A. Davis, P. Woodward, F. Herwig; Idealized hydrodynamic simulations of turbulent oxygen-burning shell convection in 4π geometry. Monthly Notices of the Royal Astronomical Society 2016; 465 (3): 2991-3010. doi: 10.1093/mnras/stw2783

Woodward, P. R., Herwig, F., and Lin, P.-H. 2015. Hydrodynamic Simulations of H Entrainment at the Top of He-Shell Flash Convection. *Astrophysical Journal* 798, 1-26.

Herwig, F., P. R. Woodward, P.-H. Lin, Mike Knox, and C. L. Fryer 2014. Global Non-Spherical Oscillations in 3 D 4π Simulations of the H-Ingestion Flash. *Astrophysical Journal Letters* 792, L3.

Woodward, P. R., J. Jayaraj, P.-H. Lin, M. Knox, J. Greensky, S. D. Hammond, S. E. Anderson, "Scaling the Multifluid PPM Code on Blue Waters and Intel MIC," Proc. Extreme Scaling 2013 Workshop, Boulder CO, August, 2013. https://www.xsede.org/documents/271087/586927/Woodward.pdf.

Woodward, P. R., J. Jayaraj, P.-H. Lin, M. Knox, D. H. Porter, C. L. Fryer, G. Dimonte, C. C. Joggerst, G. M. Rockefeller, W. W. Dai, R. J. Kares, and V. A. Thomas, "Simulating Turbulent Mixing from Richtmyer-Meshkov and Rayleigh-Taylor Instabilities in Converging Geometries using Moving Cartesian Grids," Proc. NECDC2012, Oct., 2012, Livermore, Ca., LA-UR-13-20949; also available at www.lcse.umn.edu/NECDC2012.

Lin, P.-H., and P. R. Woodward 2016, "Transforming the Multifluid PPM Algorithm to Run on GPUs," *Journal of Parallel and Distributed Computing* 93, 56-65.

Wetherbee, T., E. Jones, M. Knox, S. Sandalski, and P. Woodward 2015. In-Core Volume Rendering for Cartesian Grid Fluid Dynamics Simulations. *Proceedings of the 2015 XSEDE Conference: Scientific Advancements Enabled by Enhanced Cyberinfrastructure*, Article No. 35.

GEOSCIENCE

WEATHER

CLIMATE

GEOLOGY

ENVIRONMENT

GEOPHYSICS

ECOHYDROLOGY

SPACE WEATHER

THE TERRA DATA FUSION PROJECT

Allocation: Blue Waters Professor/240 Knh
PI: Larry Di Girolamo[1]
Collaborators: Guangyu Zhao[1], John Towns[1], Shaowen Wang[1], Yan Liu[1], and Kent Yang[2]

[1]University of Illinois at Urbana-Champaign
[2]The HDF Group

EXECUTIVE SUMMARY

The Terra satellite is the flagship of NASA's Earth Observing System. Data from its five instruments reside at different data centers in different file formats and projections and total ~1 petabyte. This non-uniformity and geographic distribution makes it impractical to address many Earth Science questions. Here, we initiated the Terra Data Fusion Project through collaborative efforts between NASA, Hierarchical Data Format (HDF) Group, and NCSA. Large Terra data transfers involving the Blue Waters team and NASA resulted in improvements to NASA's cyberinfrastructure. The data was used to (1) characterize drop size distribution of liquid water clouds over the oceans, revealing enormous positive biases (2 to 11 μm in zonal means) in drop sizes relative to previous satellite-based knowledge, and (2) examine decadal trends in Earth's color and texture, producing the **first** climatological color and texture maps of the Earth, and showing that the Earth has been getting bluer and smoother over the Terra record.

FIGURE 1: RGB composite image generated using the radiance averaged between May 2000 and Dec. 2014 for the MISR (a) AN, (b) DA and (c) DF cameras, along with trends in color shift for the (d) AN, (e) DA and (f) DF cameras with the global mean trends removed. Only trends with p-value ≤ 0.05 are displayed. Negative trends represent redder over time, while positive trends represent bluer.

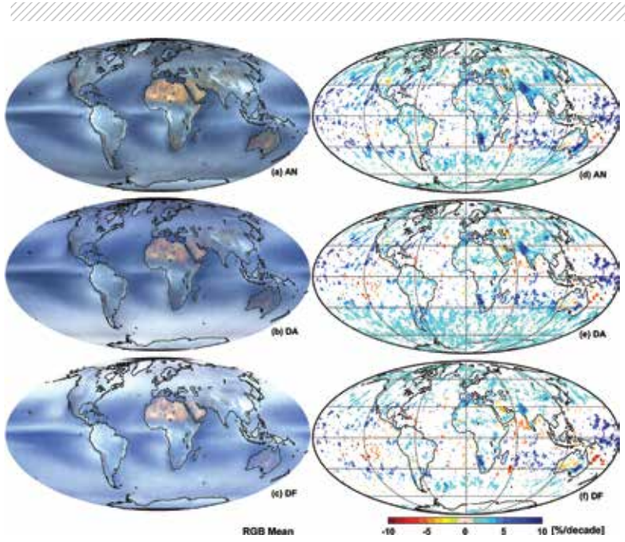

INTRODUCTION

Terra was launched in 1999 and continues to collect data for Earth sciences using five instruments: the Moderate-resolution Imaging Spectroradiometer (MODIS), the Multi-angle Imaging SpectroRadiometer (MISR), the Advanced Spaceborne Thermal Emission and Reflection Radiometer (ASTER), the Clouds and Earth's Radiant Energy System (CERES), and the Measurements of Pollution in the Troposphere (MOPITT). Terra data is amongst the most popular NASA datasets, serving not only the scientific community, but also governmental, commercial, and educational communities.

The need for data fusion and the ability for scientists to perform large-scale analytics with long records have never been greater [1]. The challenge is particularly acute for Terra, given its growing data volume (>1 petabyte), the storage of different instrument data at different NASA centers, the different data file formats and projections, and inadequate NASA cyberinfrastructure [2]. We recently initiated the Terra Data Fusion Project to tackle two long-standing problems: 1) How to efficiently generate and deliver Terra data fusion products; 2) How to facilitate the use of Terra data fusion products by the community in generating new products and knowledge through national computing facilities, and disseminate these new products and knowledge through national data sharing services.

The solutions to these questions will (1) facilitate greater ease in creating new geophysical retrieval algorithms that provide greater accuracy than the current single instrument algorithms, (2) provide an easy mechanism for other researchers to access and process the entire Terra record, (3) greatly reduce error and redundancy among scientists who are using multiple instrument datasets, (4) provide greater insight into geophysical processes through synergistic use of fusion products, and (5) provide a framework for fusion that could extend to other NASA missions and constellations. The result will facilitate discovery and accelerate progress in Earth Science research.

METHODS & RESULTS

Key steps in the Terra Data Fusion Project include: (1) transferring the entire Terra record (Level 1B radiance; >1 petabyte) to Blue Waters from NASA centers, (2) building software optimized for whole mission processing on Blue Waters to create basic fusion products; (3) optimizing data granularity and HDF application programming interface settings that best support parallel I/O on Blue Waters; (4) archiving and distribution of Terra fusion products through National Data Services and searchable through existing NASA services.

Thus far, we successfully transferred 480 terabytes of MISR and MODIS radiance data. Blue Waters staff worked with the NASA networking and Langley Atmospheric Science Data Center (ASDC) teams to diagnose and remedy performance obstacles for this data flow. This effort led to a 250% increase in data transfer rates out of the ASDC to the broader community.

Initial science investigation using this dataset was carried out in two studies. Fusing MISR and MODIS products, we characterized effective radius of the liquid cloud drop size distribution over oceans. Details were given in Blue Waters Annual Report 2015 and in [3]. Our results paint a radically different picture on the distributions of cloud drop sizes in our atmosphere compared to what was previously determined from the original MODIS data; e.g., 2 to 11 μm differences in the zonal means.

We also processed the 15-year MISR data to examine the decadal trends in the Earth's color and texture [4]. Globally, we show that the Earth has been appearing relatively bluer (up to 1.6% per decade from both nadir and oblique views) and smoother (up to 1.5% per decade only from oblique views) over the past 15 years. Regional shifts in color [Figure 1(d)(e)(f)] and texture [Figure 2(d)(e)(f)], which are significantly larger than global means, are observed, particularly over polar-regions, along the boundaries of the subtropical highs, the tropical western Pacific, Southwestern Asia, and Australia. In [4], we demonstrated that the large regional trends cannot be explained either by uncertainties in radiometric calibration or variability in total or spectral solar irradiance; hence they reflect changes internal to Earth's climate system. The 15-year-mean true color composites [Figure 1(a)(b)(c)] and texture [Figure 2(a)(b)(c)] images of Earth at both nadir and oblique views are a first—providing us a climatological view of what the Earth looks like at ~10:30 AM, the time at which the MISR has been designated to see the daylight side of the Earth in its sun-synchronous orbit.

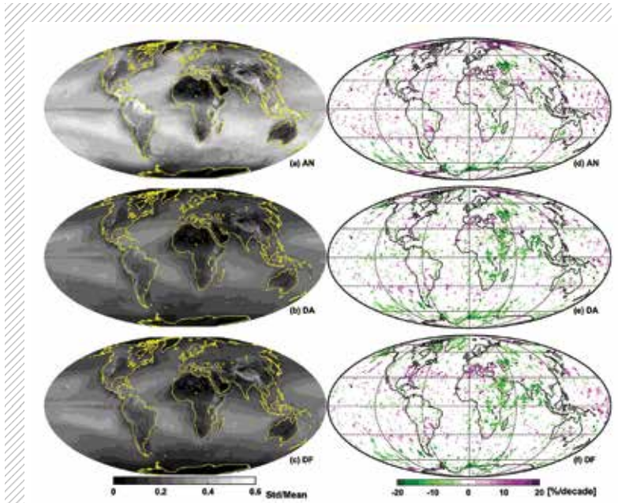

FIGURE 2: Global distribution of mean texture index averaged between May 2000 and Dec. 2014 for the (a) AN, (b) DA and DF cameras along with trends in texture index for the (d) AN, (e) DA and (f) DF cameras. Only trends with p-value ≤ 0.05 are displayed. Negative trends represent smoother over time, while positive trends represent rougher.

WHY BLUE WATERS

Key advantages of using Blue Waters for access, usage, and distribution of Terra fusion products are that the Terra data and processing are local, with access and sharing that are **global**. It has been demonstrated that having the Terra data local, with processing tuned to a massively parallel system with excellent sharing services, in one of the largest storage and bandwidth computing facilities in the country, provides an excellent framework for large-scale processing, analytics, and mining of the entire Terra record. In addition, the project staff provides expertise critically needed to maximize workflows.

NEXT GENERATION WORK

Terra is just one Earth Science data set. Fusion with other instrument records and meteorological reanalysis data (all of which are growing exponentially) for advancing Earth Sciences will require a Track-1 system that is accessible by the community.

PUBLICATIONS AND DATA SETS

Liang, L., L. Di Girolamo, and W. Sun, Bias in MODIS cloud drop effective radius for oceanic water clouds as deduced from optical thickness variability across scattering angles, *J. Geophys. Res. Atmos.*, 120, (2015), doi:10.1002/2015JD023256

The main data set applied to this study is the fusion product generated by fusing the MISR Level1B radiance product (Version F03_0024) and the MODIS level2 cloud product (Collection 6).

Zhao, G., et al., Regional changes in Earth's color and texture as observed from space over a 15-year period (2016), *IEEE Trans. Geosci. Remote Sens.* doi:10.1109/TGRS.2016.2538723

The main dataset applied to this study is the MISR level1B radiance product (Version F03_0024).

PETASCALE MODELING OF CONVECTIVE STORMS UNDER CLIMATE CHANGE AND VARIABILITY

Allocation: Illinois/240 Knh

PI: Robert J. Trapp[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

This research seeks to answer the question of how current-day extreme storm events might be realized under human-induced climate change. The "pseudo-global warming" (PGW) methodology has been adapted for this purpose. Modified atmospheric states drawn from global climate model (GCM) output were used to constrain Weather Research and Forecasting Mode (WRF) model simulations of the extreme events at high resolution. Comparison of an ensemble of these simulations with control simulations facilitated assessment of PGW effects.

A major conclusion thus far is that the pseudo-global warming modifications do not induce a change in the convective morphology of the events considered. In other words, when a current-day tornado-bearing supercell is placed in a future climate, it becomes a more intense supercell rather than a benign thunderstorm. This successful application of the PGW methodology has motivated our ongoing work with land falling hurricanes.

INTRODUCTION

A persistent uncertainty in climate change assessments regards how the frequency and intensity of local, high-impact thunderstorms, and even large thunderstorm systems including hurricanes, might be affected by human-enhanced greenhouse gas concentrations. 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. Modeling approaches such as dynamical downscaling have addressed this resolution issue, but their applications thus far have generally been unconcerned with historical events, and therefore about how these events might be projected in the future.

METHODS & RESULTS

The pseudo-global warming methodology was adapted for the purpose of investigating the impact of human-induced climate change on three high-end tornado events. Modified atmospheric states drawn from global climate model output were used to constrain WRF model simulations of these events at high resolution. Comparison of an ensemble of these simulations with control simulations facilitated assessment of PGW effects.

The conclusion that the PGW modifications do not induce a change in the convective morphology of the three events considered is due, in part, to the PGW-enhanced convective available potential energy (CAPE), and in spite of the PGW-reduced environmental wind shear. Other PGW modifications precluded storms from forming: the combined effects of increased convective inhibition and decreased forcing led to a failure of convection initiation in many of the experiments.

Our application of the PGW methodology has provided some additional insight into, and perhaps some alternative interpretations of, the results from prior studies that used environmental proxies. We found that the updrafts simulated under PGW were more relatively intense, but not in proportion to the projected higher levels of CAPE. As estimated by the amount of hail in updraft cores, we found that the effects of precipitation loading and its associated reduction of updraft buoyancy were higher in PGW updrafts. The conclusion is that projected extreme values of CAPE have the potential to lead to convective updrafts that are strong, but not necessarily extremely strong.

Our new work is employing finer-resolution grids to determine how these PGW effects will impact tornado intensity. We are also applying the PGW methodology to land falling hurricanes. A specific question here is whether a prolific tornado-generator, like Hurricane Ivan (2004), will become more hazardous in a future climate. Finally, we are developing a PGW-based framework for climate-change attribution studies of current-day extreme storms.

WHY BLUE WATERS

The episodic nature and relatively small size of thunderstorms and tornadoes necessitate a research approach that can account for temporal scales that range from decades to minutes, and spatial scales ranging from thousands of kilometers to hundreds of meters. Blue Waters is providing us with the resources needed to achieve this **unprecedented** level of climate simulation.

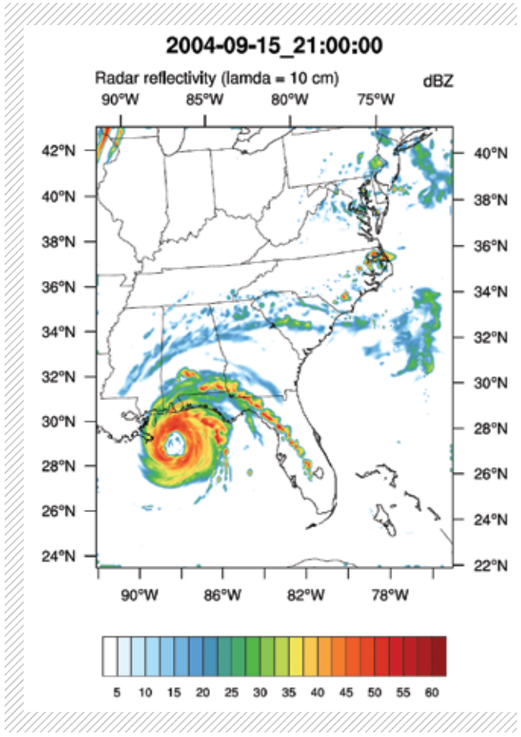

NEXT GENERATION WORK

The next-generation work will involve the use of the Model for Prediction Across Scales (MPAS) to move toward understanding convective storms under climate change and variability. MPAS is an emerging global atmospheric model with variable-resolution grids. A variable-resolution model simultaneously allows a high concentration of grid points (and hence high resolution) over one region of the global domain such as the contiguous U.S., and coarsely spaced grid points elsewhere. Thus, processes from the high-resolution region are allowed to feedback naturally to the remainder of the global atmosphere.

PUBLICATIONS

Trapp, R. J., and K. A. Hoogewind, 2016: The realization of extreme tornadic storm events under future anthropogenic climate change. *J. Climate*, DOI: http://dx.doi.org/10.1175/JCLI-D-15-0623.1

FIGURE 1: Simulated radar reflectivity from a high-resolution WRF model simulation of Hurricane Ivan. The radar reflectivity portrays the structure of Ivan at 2100 UTC on 15 September 2004, as this hurricane was making landfall along the Gulf Coast. The rainband extending through Florida is of particular interest, because it is associated with numerous tornadoes.

EXTREME-SCALE MODELING—UNDERSTANDING ECOHYDROLOGIC DYNAMICS UNDER CLIMATE CHANGE

Allocation: Illinois/50.0 Knh
PI: Praveen Kumar[1]
Co-PIs: Phong V. V. Le[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

The elevation of atmospheric CO_2 increases the ratio of carbon fixation to water loss of plants or water-use efficiency. This shift in ecosystem functioning is central to understanding the cycles of water, energy and carbon under climate change. However, the magnitude of the effects on ecohydrologic dynamics, such as soil moisture content, surface runoff, and ponding dynamics controlled by microtopographic variability on the land surface, remains unclear. The goal of this project is to develop efficient and robust predictive models on a hybrid central processing unit-graphics processing unit (CPU-GPU) architecture for capturing the mechanism of vegetation acclimation and its link to hydrologic processes. This utilization of Blue Waters marks the **first** investigation on the impacts of vegetation acclimation under climate change on ecohydrologic dynamics at light detection and ranging (LIDAR)-resolution over large scales.

INTRODUCTION

Amongst the most important potential environmental changes is the elevation of atmospheric carbon dioxide concentrations\(CO_2), predicted to increase to 550 ppm by 2050 and probably exceed 700 ppm by the end of the 21st century [1]. One of the primary concerns with rising CO_2 under climate change is its potential to alter the hydrologic cycle through vegetation acclimation and modifications in evapotranspiration. Several models are being used to investigate the heterogeneity and process complexity of the soil-vegetation-atmosphere interactions in ecosystems under global warming. The common approach among these models is the coupling of a land surface model (LSM) with a distributed, either physically-based or simplified, hydrological model to capture the feedback cycles between the biosphere and atmosphere. However, these models have not incorporated capabilities to capture processes dominated by micro-topographic features on the land surface.

This work directly targets the development of a predictive capability for investigating micro-topographic controls on ecohydrologic dynamics under climate change over large scales in a massively parallel architecture that allows us to explicitly incorporate emerging high-resolution LIDAR measurements.

METHODS & RESULTS

We coupled a multi-layer canopy model [2,3] (MLCan) with a conjunctive surface-subsurface flow model [4] (GCSFlow) to capture the acclimatory responses of vegetation to global warming, then predict how these changes affect ecohydrologic dynamics on landscapes at LIDAR-scale resolution. While MLCan is a biophysical model that simulates the eco-physiological acclimations of vegetation under climate change, GCS-flow is an integrated surface–sub-surface flow model utilizing LIDAR-resolution topographic data to capture the micro-topographic controls on hydrologic processes. The MLCan-Flow3D model is implemented on a hybrid CPU-GPU parallel computing environment to overcome challenges associated with the high density of computational grid and nonlinear solvers. Specifically, MLCan is implemented at all nodes covered by vegetation using Message Passing Interface (MPI), and GCS-flow is implemented using the CUDA platform.

We performed simulations using LIDAR topographic data in the Goose Creek watershed of the Upper Sangamon River Basin (USRB) in central Illinois. This watershed is intensively managed for agriculture and is part of the Critical Zone Observatory for Intensively Managed Landscapes (IML- CZO). Maize and soybean are two major crops planted in rotation every year in the study area. The model runs are performed with three different scenarios contrasted to evaluate the impacts of vegetation acclimation on ecohydrologic dynamics under climate change.

Changes in ecohydrological dynamics dominated by micro-topographic features in response to elevated CO_2 and air temperature increase are presented in Figure 1. We showed that rising CO_2 is likely to decrease evapotranspiration, thus increase soil moisture and surface water and ponding dynamics. However, as higher temperature is also considered, there is a net increase in evapotranspiration, leading to a reduction in soil moisture storage and ponding persistence. Thus far, using Blue Waters we have demonstrated that hybrid computing is feasible for detailed, extreme-scale ecohydrologic modeling, which has been previously assumed to be an intractable computational problem.

WHY BLUE WATERS

Blue Waters has been critical for this project. To our knowledge, all existing ecohydrologic models have not incorporated capabilities to simultaneously capture vegetation acclimation under climate change and processes dominated by micro-topographic features. The unique computational power of Blue Waters associated with MPI-CUDA-aware capability is allowing us to conduct detailed ecohydrologic simulations at **emerging** LIDAR resolution over large domains for the **first** time and to rigorously study the impacts of micro-topographic variability on ecohydrologic dynamics.

NEXT GENERATION WORK

We aim to develop additional components based on the current model that provides powerful tools to conduct simulations and explore various scientific questions in ecohydrology and biogeochemistry.

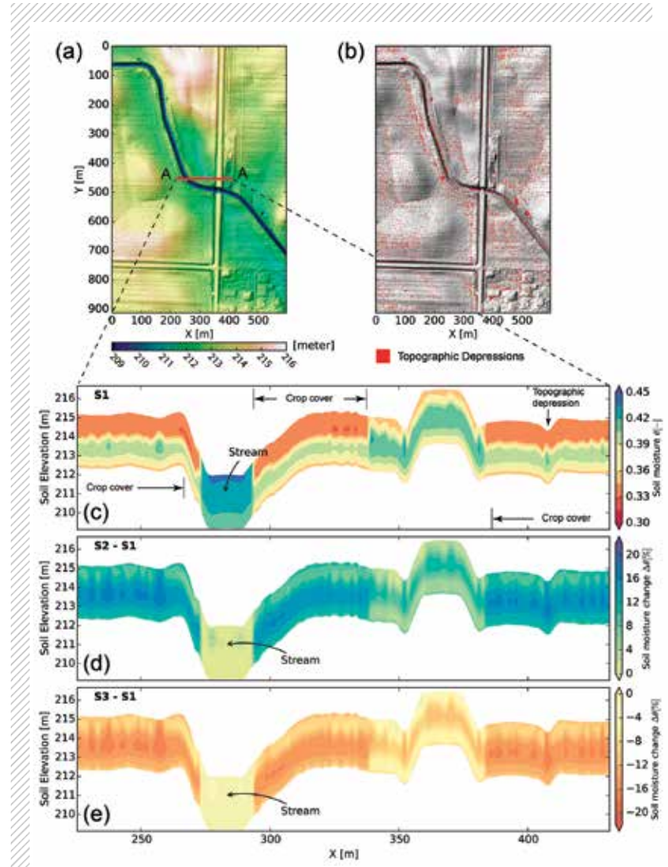

FIGURE 1: Comparison of soil moisture profile under present (S1) and elevated (S2, S3) CO_2 conditions. (a) LIDAR data of a zoomed area within the simulation domain. (b) Map of topographic depressions identified using LIDAR data in Goose Creek watershed. (c) Snapshot of soil moisture profile over depth in cross-section A-A at DOY 230 in 2005 growing season. (d) Difference of soil moisture profile over depth between S1 and S2 scenarios. (e) Difference of soil moisture profile over depth between S1 and S3 scenarios.

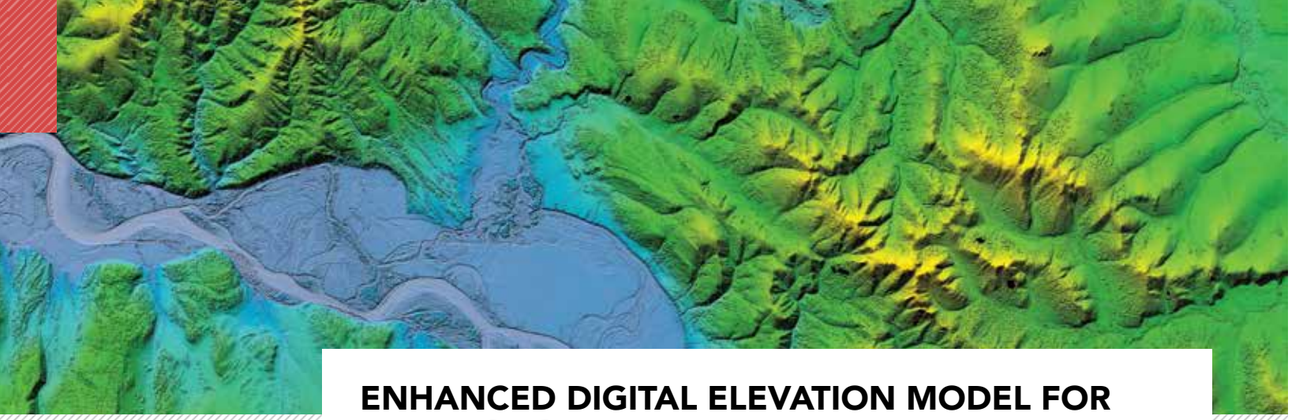

FIGURE 1: Meandering through Brooks 2m sm cc. This is a river meandering through the Brooks Range on its way to the North Slope. Note that the texture next to the rivers and along some of the mountain slopes is vegetation. In one image you can see the interplay between tectonics, vegetation, erosion and sediment transfer in a remote and unmodified high-relief environment.

ENHANCED DIGITAL ELEVATION MODEL FOR THE ARCTIC

Allocation: NSF PRAC/19.5 Mnh; Innovation 3.5 Mnh
PI: Paul Morin[1]
Collaborators: Charles Nguyen[1], Ian Howat[2], MJ Noh[2], Michael Willis[3], Brian Bates[4], and Cathleen Williamson[4]

[1]University of Minnesota
[2]The Ohio State University
[3]Cornell University
[4]National Geospatial-Intelligence Agency

EXECUTIVE SUMMARY

The earth's land surface topography is arguably the most fundamental single data set in the geosciences, geographical sciences, and civil engineering. It is essential to disciplines ranging from the location of rivers and the extent of watersheds in hydrology, to permafrost collapse in built-up areas, to the change in the shape of volcanoes in volcanology. The Polar Geospatial Center (PGC) and its partners at The Ohio State University (OSU) and Cornell University are adapting PGC's digital elevation model (DEM) production capabilities from a small area, on-demand production to systematically process and create high resolution stereo elevation mosaics of the entire Arctic sub-meter commercial imagery archive. Such a DEM not only catapults the Arctic from the worst to among the best mapped regions on Earth, it also allows precise detection of change over time, creating a new quality basis for measuring this rapidly evolving landscape.

Using Blue Waters, the National Geospatial-Intelligence Agency (NGA), Digital Globe (DG) and PGC have built a nearly seamless archive of polar sub-meter stereo imagery that consists of almost 50,000 stereo pair strips from the Worldview 1, 2, and 3 satellites. Using photogrammetric algorithms to construct DEMs from the stereo satellite image pairs allows mapping surface features at the same scale as airborne LIDAR without the cost or logistics constraints. This imagery collection increases at a rate to more than the equivalent in area of California every day while there is sufficient Arctic sun. The data will be used by the Arctic research and government communities to support activities including transportation, national defense, land management, sustainable development, and scientific studies. Further, repeating DEMs with frequencies of months or even days, can be used for change detection, with applications ranging from studies of land use such as deforestation, to land and water resource management, to environmental change.

INTRODUCTION

There is little high-resolution, consistent, high-quality elevation data in the Arctic. In 2000, the Shuttle Radar Topography Mission (SRTM) began by acquiring Synthetic Aperture data for the earth that was processed into an elevation model with a 30 meter resolution. This mission had limited geographic coverage between the latitudes of 60N and 56S because of the shuttle's orbital inclination.

METHODS & RESULTS

At their most basic level, all stereo-photogrammetric DEM extraction algorithms use a statistical procedure, usually cross-correlation, to locate conjugate pixels in one image to those of the same point on the ground in a second image acquired from a different camera position. The difference in position of conjugate pixels between nearly simultaneous images provides the parallax from which the relative height of the surface can be derived using known ground-to-image geometry. In the case of the WorldView satellites, ground-to-imagery geometry is estimated from a set of rational polynomial coefficients (RPCs), determined from the satellite orbital position, which constrain a cubic model of the spatial relationship between image and ground. Locations of conjugate points between images are typically determined starting from an initial estimate provided by the RPC positioning, followed by determining the precise location from statistical image matching.

Our team has spent three years developing an efficient algorithm for constructing photogrammetric DEMs from satellite imagery to create a fully automated system capable of handling large amounts of data. Development of the Surface Extraction from TIN-based Search-space Minimization (SETSM) algorithm began to facilitate an automated processing pipeline for the PGC operations. SETSM DEMs have been extensively validated [1], are node parallelized using the OpenMP application programming interface, and have been applied to processing large area DEM mosaics in proof-of-concept studies. Uniquely, SETSM's structure eliminates the need for an existing (i.e. "seed") elevation measurement for a priori constraint or any data specific, user-defined search parameters, making it a truly automated algorithm. SETSM only requires the two stereo images and the RPC file.

The DEM extraction workflow starts with a preprocessing step that corrects the source imagery for sensor-specific detector alignment artifacts and outputs a GeoTIFF-formatted set of source rasters. Once the source imagery is corrected, SETSM takes the two source images and derives increasingly detailed elevation models using its pyramid-based approach. The initial step halts the process at 8 meter resolution. This is done for two reasons: 1) any poor-quality source imagery will fail to reach this step, allowing us to identify problematic jobs early on, and 2) the 8 meter product is extremely useful to the scientific community because of the reduced file size. We convert the TIN-based (vector) elevation model to a raster and use the success rate to assess the feasibility of continuing to process to 2 meter. This quality-assessment filter reduces the amount of computer time spent on data unlikely to result in high-quality output. For image pairs that pass this filter, we submit another set of jobs to refine the 8 meter elevation model to 2 meter. The final step in the workflow is to convert the 2 meter TIN-based model to a raster.

WHY BLUE WATERS

We will generate a publicly available, 2 meter posting elevation model of the entire Arctic landmass poleward of 60N and extended south to include all of Alaska, Greenland, and Kamchatka. This ArcticDEM will be constructed from stereo-photogrammetric DEMs extracted from pairs of sub-meter resolution WorldView satellite imagery and vertically registered using ground control from GPS-surveyed points, coordinated airborne LIDAR surveys by NASA's Operation IceBridge, and historical elevations from the NASA IceSAT mission. This project is to be performed during the U.S.'s tenure as Arctic Council chair (2015–2017) [2]. Due to the large footprint of the project (18M km^2) and the limited 2 year timeframe for production, it will require over 19 million node hours (600 million core hours equivalents) to complete.

To simply explain the impact of Blue Waters, we have 12.5^2 improvements in resolution, 58,500 times improvement in time to solution compared to a single workstation, and 220 times improvement in cost. This equates to a **9 billion time** productivity improvement.

NEXT GENERATION WORK

The next generation of this work will involve accelerated DEM projections, expanding the range of application of ground control from a wide variety of sources, and the mosaicing of 20 trillion 2m x 2m elevation grid cells.

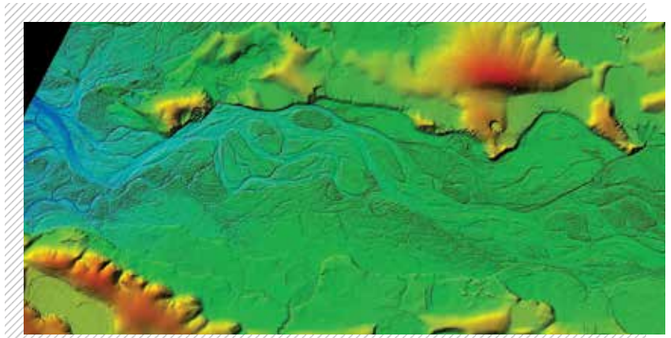

FIGURE 2: Western Alaska Braided Stream. Braided streams are found in rivers with a high sediment load and little vegetation. This example shows the ever changing bars being corralled by topography above and vegetation (the textured area) to the south.

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

Allocation: NSF PRAC/2.30 Mnh
PI: Jamesina J. Simpson[1]
Collaborators: Bach Nguyen[1], Santosh Pokhrel[1], Alireza Samimi[1], Miguel Rodriguez[1]

[1]University of Utah

EXECUTIVE SUMMARY

The largest documented geomagnetic storm occurred in 1859. This storm caused telegraph operators communicating over 100-km-long wire lines to experience electric shocks, some nearly fatal. The historical record suggests that extreme space weather is likely to impact the Earth again in the future. However, modern electro-technologies will be affected by space weather to a much larger degree than in the past. We are using global Maxwell's equations models of the Earth-ionosphere waveguide to calculate location-specific space weather hazards to electric power grids in order to prevent blackouts. Blue Waters is permitting us to account, at high resolutions, for the Earth's topography, oceans, variable composition of the lithosphere, as well as the variable ionospheric composition, and source conditions according to time, altitude, and position around the globe. Previous analytical and computational approaches were localized in nature, assumed highly simplified geometries, and could not model arbitrary (realistic) source waveforms in time or space.

INTRODUCTION

The historical record indicates the possibility of extremely intense space weather events directed toward the Earth from the sun. The largest documented geomagnetic storm occurred in 1859 [1] and caused telegraph operators communicating over 100-km-long wire lines to experience electric shocks, some nearly fatal [2]. Further, business transactions requiring telegraphic exchanges were completely shut down in the world's major capitals [2].

A 2008 National Academies report [3] indicates that extreme space weather events, "though rare, are likely to occur again some time in the future." However, a reoccurrence of an 1859-magnitude (coronal mass ejection-driven geomagnetic) storm could disrupt today's society to a much greater degree due to the proliferation of vital but vulnerable electro-technologies. Interruptions to radio communications, commercial airline flight plans, satellite operations, transportation, banking, financial systems, home and industrial computer electronics, and power grids are just some examples. The National Academies report estimates the overall economic cost of one such extreme event as ranging from millions to trillions of dollars, with a recovery time of four to 10 years [3].

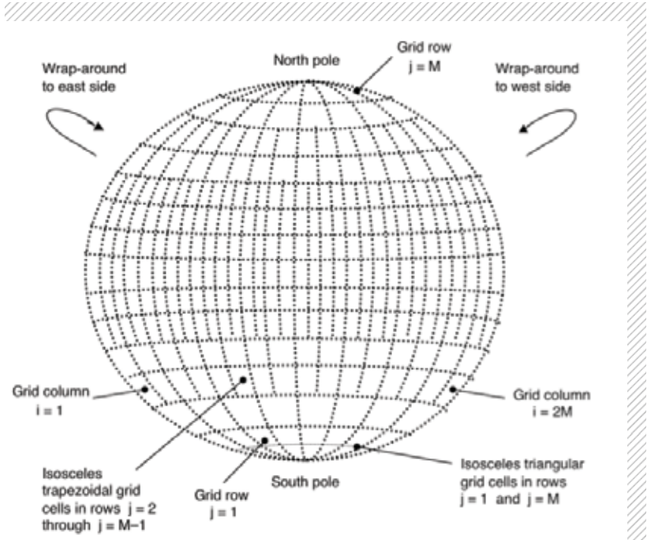

FIGURE 1: Layout of the 3D FDTD grid as seen from a constant radial coordinate.

METHODS & RESULTS

Our goal is to greatly improve our ability to understand and predict space weather hazards in the near-Earth environment, especially on the operation of electric power grids. To achieve this goal, we are advancing and applying detailed, high-resolution Maxwell's equations models of the Earth-ionosphere waveguide developed by Professor Simpson over the past 15 years (e.g. [4, 5]). These models are based on the finite-difference time-domain (FDTD) method. FDTD allows for the modeling of realistic time-waveforms of disturbed ionospheric currents and resulting electromagnetic fields at the Earth's surface induced by space weather. FDTD is also a grid-based approach that permits modeling of intricate details on a global scale, such as the Earth's complete topography, oceans, variable composition of the lithosphere, as well as the variable ionospheric composition and disturbed ionospheric current systems according to altitude and position around the globe.

Using the global FDTD models, we are generating location-specific ground-level electromagnetic field data to help predict the induced voltages on electric power grids during space weather events. Fig. 1 illustrates a planar cut of the three-dimensional FDTD grid as seen from constant radial coordinate. The top half of Fig. 2 illustrates an example snapshot of the disturbed ionospheric electric fields during the October 2003 "Halloween" geomagnetic storm. These disturbed ionospheric fields are used as sources to the FDTD grid at ~100 km altitude, and then the ground-level electromagnetic fields are calculated. The ground-level electromagnetic fields corresponding to the time of the source currents in the top half of Fig. 2 are shown in the lower half of Fig. 2. Individual power grid operators may use the FDTD-computed results to design and implement effective mitigation strategies to protect the grid from voltages induced by geomagnetic storms.

WHY BLUE WATERS

FDTD can account for highly detailed three-dimensional geometries and material compositions. However, it is computationally expensive, especially when modeling the entire world. Blue Waters has helped us improve the parallelization of our global model, so that we can now model at resolutions on the order of 1 km x 1 km x 1 km and higher. Before the start of our Blue Waters project, our highest grid resolution was 40 km x 40 km x 5 km, so this is an **8,000 times improvement**. Achieving these high resolutions has been challenging because, as shown in Fig. 1, dividing the grid into equal sections for each processing core is challenging due to the merging of grid cells in the polar regions.

Furthermore, Blue Waters is allowing us to model **more realistic** ionospheric source time-waveforms and spatial distribution than previously possible. Hazards to electric power grids critically depend on the complex distribution of storm-driven ionospheric sources overhead, the grid's vicinity to ocean-continent boundaries, and the underlying rock structure. The FDTD-calculated results may be instrumental in helping protect individual power grids substations.

Blue Waters project **staff have been critical** to our success. They improved the efficiency of our model by 4% by helping us incorporate non-blocking message-passing interface (MPI) to send and receive commands into a section of our code. Also, the staff significantly aided our productivity by rapidly addressing issues and questions.

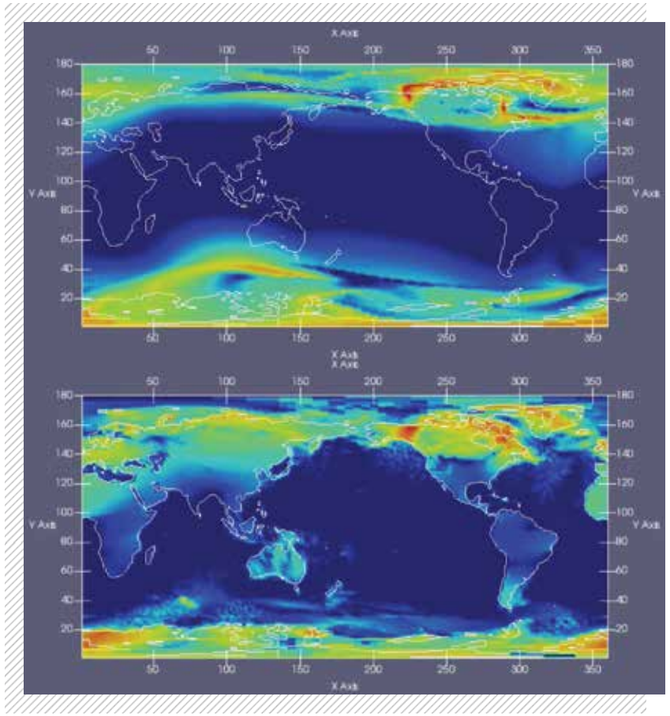

FIGURE 2: Snapshot of the electric field source amplitude versus position during the October 2003 Halloween geomagnetic storms as calculated by BATS-R-US model developed at the University of Michigan (top), and the resulting surface-level electric field values calculated by the global FDTD model (bottom).

NEXT GENERATION WORK

The next Track-1 system would allow us to extend our models higher into the ionosphere so that instead of projecting currents down to ~100 km in altitude, we could model the actual propagation and resulting physics of those currents. The FDTD models could be coupled to models of the magnetosphere to provide a more complete physics analysis of the effect of space weather on the earth.

PUBLICATIONS AND DATA SETS

Samimi, A., and J. J. Simpson, Parallelization of 3-D global FDTD Earth-ionosphere waveguide models at resolutions on the order of ~1 km and higher, *IEEE Antennas and Wireless Propagation Letters* (in press).

Samimi, A., M. Rodriguez, N. Dupree, R. Moore, and J. J. Simpson, The application of global 3-D FDTD Earth-ionosphere models to VLF propagation: Comparison with LWPC, *IEEE AP-S International Symposium and USNC/URSI National Radio Science Meeting*, Puerto Rico, June (2016).

PHYSICS-BASED STRONG GROUND MOTION SIMULATIONS

Allocation: NSF PRAC/6.60 Mnh
PI: Thomas H. Jordan[1]
Co-PIs: Jacobo Bielak[2], Kim Olsen[3], and Yifeng Cui[4]
Collaborators: S. Azizzadeh-Roodpish[5], S. Callaghan[6], D. Gill[6] , C. A. Goulet[6], R. W. Graves[6], M. M. Huda[5], W. Hwu[7], N. Khoshnevis[5], P. J. Maechling[6], K. R. Milner[6], C. Pearson[7], D. Restrepo[8], D. Roten[3], W. H. Savran[3], R. Taborda[5], and K. B. Withers[3]

[1]University of Southern California
[2]Carnegie Mellon University
[3]San Diego State University
[4]San Diego Supercomputer Center
[5]University of Memphis
[6]Southern California Earthquake Center
[7]University of Illinois at Urbana-Champaign
[8]EAFIT University

EXECUTIVE SUMMARY

Earth scientists, engineers, and computer scientists working with the Southern California Earthquake Center (SCEC) use physics-based numerical simulations and high-performance computing (HPC) to improve the understanding of seismic hazards, and earthquake processes and their effects. This past year, SCEC teams used Blue Waters to perform deterministic earthquake ground motion simulations with frequencies up to 8 Hz, while introducing new physics required for more realistic ground motion simulations, including rough-fault geometrical complexity, frequency-dependent attenuation, material plasticity, small-scale material heterogeneities, and surface topography. Earthquake simulations using our improved numerical models were validated against records from past earthquakes. We also increased the computational performance of our research software through graphics processing unit (GPU) code and parallel I/O improvements, and through workflow optimizations.

INTRODUCTION

The SCEC performs fundamental research in earthquake system science and develops predictive models of earthquake processes. SCEC scientists develop and apply the best available geoscientific understanding of faulting and wave propagation processes, together with state-of-the-art computation techniques, to produce the next generation of physics-based seismic hazard models. SCEC's research program is a collaboration among several user communities with shared interests in reducing seismic risk and enhancing seismic resilience. SCEC's computational research activities help to educate a diverse STEM workforce from the undergraduate to the early-career levels, and cross-train scientists and engineers in challenging HPC environments.

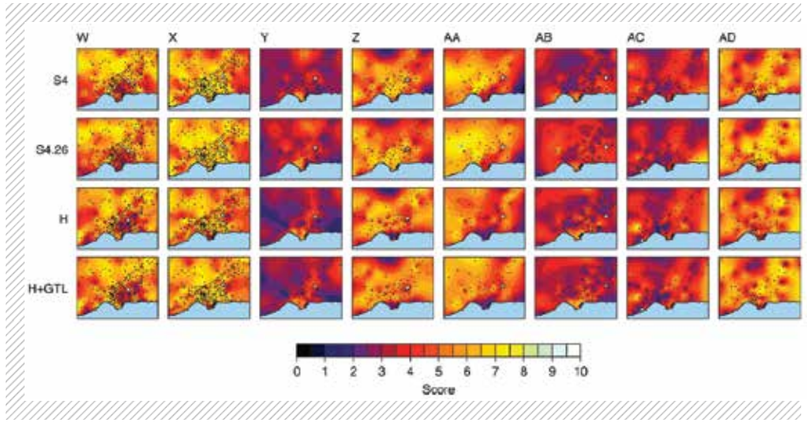

FIGURE 1: Goodness-of-fit (GOF) maps for all events with 53 or more stations used for validation. Contours indicate the score obtained by averaging the GOF values for all three components of motion (EW, NS and UD). Dots correspond to the location of stations and stars indicate the epicenters for each event. Event labels at the top of each set of four maps correspond to the results obtained using alternate velocity models (CVM-S4, CVM-S4.26, CVM-H and CVM-H+GTL), as indicated with labels on the left margin.

METHODS & RESULTS

SCEC researchers used Blue Waters to perform simulations of earthquake faulting and wave propagation at frequencies of interest to civil engineers. A significant focus of our Blue Waters computational research this year involved validating simulations against data, with much of this effort led by engineering seismologists and engineers, who recognize the potential of SCEC's efforts in physics-based ground-motion prediction.

SCEC approaches seismic hazard analysis as an earthquake system science problem that requires integration of several interrelated computational models, including accurate 3D earth structural models, friction-based fault rupture models, and anelastic wave propagation (AWP) models. SCEC's approach iteratively improves these models, re-validates them against observed ground motions, and then re-combines the models, producing an improvement in broad-impact seismic hazard computational methods.

A team led by Ricardo Taborda at University of Memphis used Blue Waters to evaluate four existing southern California velocity models by assessing how well each predicted ground motion in the greater Los Angeles region when used as inputs to deterministic wave propagation simulations. These evaluations were performed by running multiple earthquake simulations and then using quantitative comparisons between simulated motions and a collection event data. The team used Blue Waters to simulate earthquakes within a domain with a surface area of 180 km x 135 km. Each earthquake was modeled as a point source with rupture parameters scaled according to magnitude. Hercules—finite-element software developed by SCEC-affiliated scientists—was used to simulate the ground motions for each earthquake and velocity model combination. Hercules has shown to be a reliable tool for 3D earthquake ground motion simulation [1,2]. The group simulated 30 moderate-magnitude earthquakes (3.5 to 5.5) and compared synthetics with data recorded by seismic networks on over 800 stations. Each of the 120 simulations (30 earthquakes, four velocity models) was run with a maximum frequency of 1 Hz and a minimum shear wave velocity of 200 m/s. The comparisons between data and synthetics were ranked quantitatively using standard seismological goodness-of-fit (GOF) criteria. The regional distribution of the GOF results for all events and models were analyzed and ranked according to the performance of each velocity model (Fig.1). The group identified one of the southern

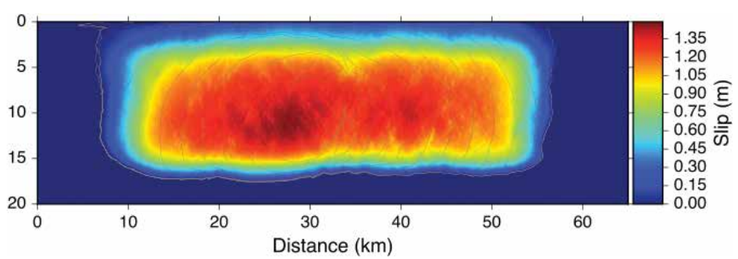

FIGURE 2: Sample final slip on fault for a moment magnitude 6.9 earthquake showing 1s (1 Hz) rupture contours in gray. The variation in the final slip is correlated with the topographic complexities on the fault surface.

California velocity models that yields consistently better results.

A team lead by Kim Olsen at San Diego State University used Blue Waters to generate a database of dynamic rupture sources on topographically complex faults (Fig. 2). Each dynamic source computation used 8,192 CPU cores for approximately 5.5 hours using the SORD code [3,4]. The ruptures were verified against ground motion prediction equations (GMPEs) used in engineering practice by running wave propagation simulations for a region approximately 40 km from all sides of the fault. The wave propagation simulations were completed using AWP-ODC—finite-difference software developed by SCEC-affiliated scientists—on 22,272 cores for approximately two hours. Each simulation produced 52.5 seconds of 0–7.5 Hz wave propagation. This group found that the dynamic rupture-based source models produce realistic ground motions at high frequencies, indicating that their database of ruptures contains suitable high-frequency source models for statistical analysis.

The SCEC PAID Project team, led by Yifeng Cui at the San Diego Supercomputer Center, worked with the GPU IME team led by Wen-mei Hwu to optimize the nonlinear AWP-ODC GPU code for scalability and efficiency on Blue Waters XK7 nodes. The team integrated several improvements into the AWP numerical models, including support for plasticity yielding and frequency-dependent attenuation. They also optimized code performance through yield factor interpolation, memory tuning to increase occupancy, communication overlap using multi-streaming, and parallel I/O to support concurrent source inputs. Simulation of non-linear material behavior requires the addition of 17 new variables compared to linear computations, posing challenges regarding solution time and memory management. The improved nonlinear code has proven to be highly scalable and efficient and has achieved better performance than the linear code despite the additional variables and processing. The team then used the improved software on Blue Waters to perform 0–4 Hz nonlinear ShakeOut scenario earthquake simulations (Fig. 3). These results represent an advance in our ability to perform earthquake simulations at frequencies up to 4 Hz because our codes now include the advanced physics, including the small-scale complexity of the source, nonlinear effects, and frequency-dependent inelastic attenuation that are needed to accurately simulate these higher frequency ground motions.

WHY BLUE WATERS

Earthquake simulations at the spatiotemporal scales required for probabilistic seismic hazard analysis present demanding computational challenges. SCEC researchers use Blue Waters to addressing scientific problems that limit the accuracy and scale in current numerical representations of earthquake processes. Blue Waters provides the computational scale, data management capabilities, and variety of computing capabilities needed to perform our research.

NEXT GENERATION WORK

SCEC's earthquake system science computational requirements will continue to increase as new physical properties and principles are included in the simulations. Higher frequency calculations are important to engineering end-users, and they are much more computationally intensive. Once individual earthquake simulations are validated, ensembles of simulations are needed to make probabilistic calculations. Our computational requirements will grow dramatically as computational-intensive seismic hazard techniques are applied to more regions.

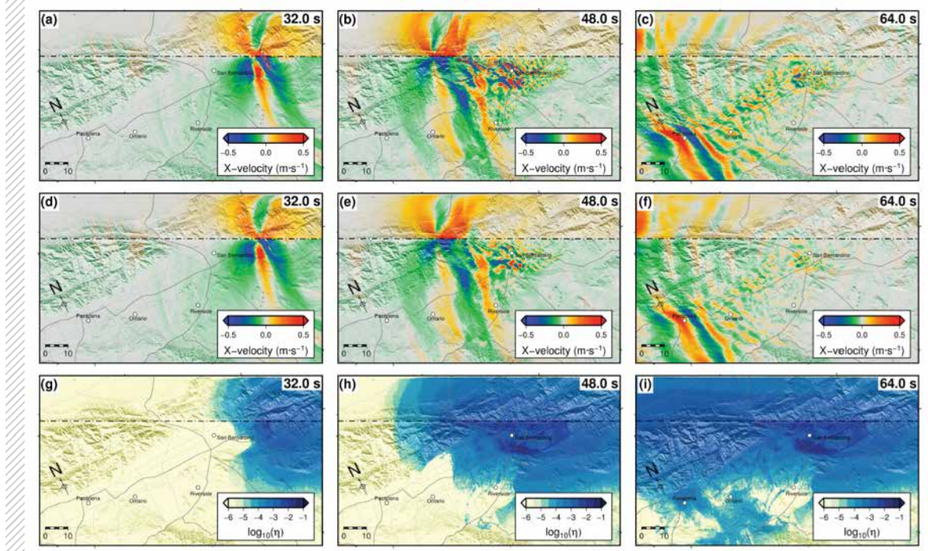

FIGURE 3: Snapshots from a 4 Hz San Andreas simulation inside the blue rectangle and the dashed line shows the fault trace. Maps a, b, and, c show fault-parallel velocity for the linear cases, and Maps d, e, and f for the nonlinear cases. Maps g, h, and i depict the evolution of permanent plastic strain at the surface obtained from the nonlinear simulation.

PUBLICATIONS AND DATA SETS

Taborda, R., S. Azizzadeh-Roodpish, N. Khoshnevis, and K. Cheng, Evaluation of the southern California seismic velocity models through simulation of recorded events, *Geophys. J. Int.*, 205:3 (2016), pp. 1342–1364, doi: 10.1093/gji/ggw085

Taborda, R., N. Khoshnevis, S. Nzizzadeh-Roodpish, and M.M. Huda, Influence of the source, seismic velocity, and attenuation models on the validation of ground motion simulations, in *Proceedings of the World Conference on Earthquake Engineering*, Santiago de Chile, Chile, January 2017, no. 4574.

Savran, W. H., and K. B. Olsen, Model for small-scale crustal heterogeneity in Los Angeles basin based on inversion of sonic log data, *Geophys. J. Int.*, 205:2 (2016), pp. 856–863, doi:10.1093/gji/ggw050

ANALYZING TROPICAL CYCLONE-CLIMATE INTERACTIONS USING THE COMMUNITY EARTH SYSTEM MODEL (CESM)

Allocation: Illinois/660 Knh
PI: Ryan L. Sriver[1]
Co-PI: Hui Li[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Our research seeks to advance understanding of the relationship between tropical cyclones (TCs) and Earth's climate system using a high-resolution state-of-the-art Earth system model (Community Earth System Model—CESM). Here, we highlight results from a series of ocean-only simulations, in which we analyze the effect of tropical cyclone wind forcing on the global ocean using three different horizontal ocean grid spacing (3˚, 1˚, and 0.1˚). Findings indicate that TCs significantly contribute to global ocean heat and energy budgets, pointing to important connections between TCs and ocean dynamics, which can influence seasonal to inter-annual climate variability (such as El Nino) and large-scale circulation patterns in the atmosphere and ocean. Furthermore, TCs' contribution to the climate system may have important implications for anthropogenic global warming through feedbacks in the coupled system, which is a scientific questions we seek to answer with our work.

INTRODUCTION

Tropical cyclones are among the world's most dangerous and destructive natural hazards. TC strength and frequency largely depends on the ambient environmental conditions and vary with changes in the earth's climate. TCs can also play an active role in the earth's climate system through complex ocean-atmosphere interactions. Extreme TC winds cause vigorous vertical mixing in the upper ocean that disturbs stratification, resulting in anomalous heat gain in the ocean interior, with the potential to alter ocean temperature distributions and influence large-scale circulations in the atmosphere and ocean. Understanding the connection between tropical cyclones and climate may be fundamentally important to advancing our understanding of climate variability and constrain uncertainties for future climate change projections.

Tropical cyclones' potential influences on climate are mostly missing from today's generation of climate models, in part due to the extreme computational costs associated with simulating the relevant small-scale physical processes and complex ocean-atmosphere interactions under TC conditions. In particular, ocean model resolution largely affects the representation of small-scale oceanic features such as mesoscale ocean eddies, which are significant contributors to modulating upper ocean stratification and global ocean heat transport. Mesoscale eddies are explicitly resolved in the 0.1˚ ocean model, whereas they are partially (fully) parameterized in the 1˚ (3˚) model. The current work aims to investigate TCs' cumulative impacts on global ocean using the ocean-sea ice configuration of the Community Earth System Model (CESM) [1], and examine the sensitivity of these TC-related impacts to model horizontal grid spacing.

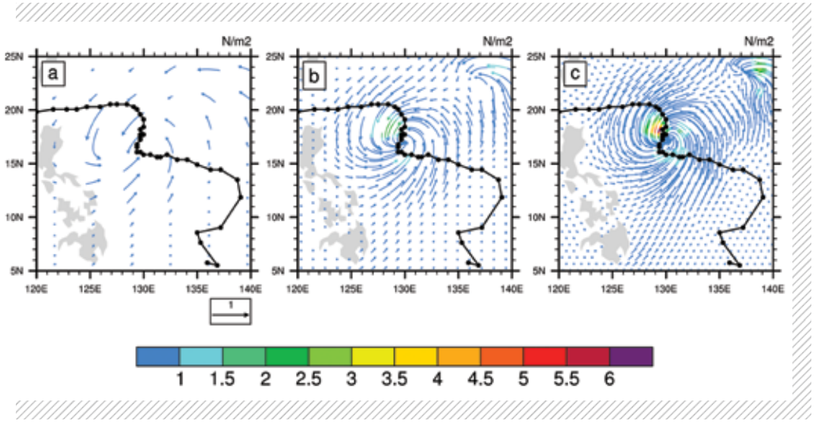

FIGURE 1: Tropical cyclone wind drag on the ocean surface as received in (a) 3˚, (b) 1˚, and (c) 0.1˚ ocean models.

METHODS & RESULTS

We conduct a suite of ocean-only simulations with varying ocean horizontal grid spacing (3˚, 1˚, and 0.1˚) using the Community Earth System Model (CESM). The ocean models are forced with identical atmospheric boundary conditions, in which we incorporate TC wind fields obtained from a high-resolution fully coupled CESM simulation [2] where TCs' geographical distribution, seasonal variation, and storm intensity are generally in agreement with the observational records. We also perform corresponding control simulations where the atmosphere contains no TC feature, to diagnose TCs' impact on the ocean properties.

Figure 1 shows an example of a TC event and the corresponding ocean surface wind drag perceived by the 3˚, 1˚ and 0.1˚ ocean models. Even though the models are forced with identical TC winds, the TC strength and wind patterns are different. The 0.1˚ model can preserve most of the TC characteristics, while its lower resolution counterparts underestimate the wind strength, which could lead to underestimated ocean vertical mixing and the subsequent heat uptake. Analysis reveals that the model simulates key characteristics of transient ocean responses to TCs, including mixing-induced sea surface cooling and sub-surface warming—responses that are sensitive to ocean grid resolution. Moreover, the 0.1˚ model better captures important dynamic processes such as ocean mixed layer deepening, eddy-induced vertical heat advection, and the zonal heat transport by equatorial waves.

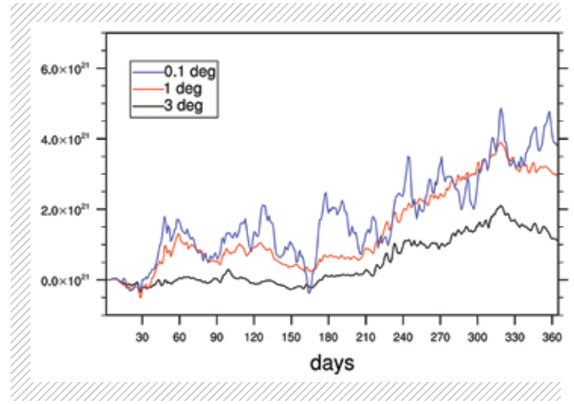

FIGURE 2: Time series of TC-induced ocean heat uptake in the 3˚ (black), 1˚ (red), and 0.1˚ (blue) ocean models. (Unit: J)

We estimate annually accumulated ocean heat gain injected irreversibly by TCs in the model cases. Figure 2 shows the time series of TC-induced ocean heat uptake over a one-year period. The value at the end of the year indicates the annually accumulated effect of TCs on the ocean heat anomalies. Results show that TCs could significantly contribute to global ocean heat and energy budgets, the magnitude of which depends on ocean grid resolution. The 0.1˚ model produces the most significant annual ocean heat uptake, amounting to 4e+21 Joules, or 0.13 petawatts in terms of annual heating rate. Differences between modeling scenarios suggest that enhanced ocean model resolution with resolved mesoscale eddies can greatly improve the representations of TCs' heating effect on global ocean.

WHY BLUE WATERS

A major challenge of high-resolution climate modeling studies is the heavy demand on computational resources. Our work involves a fully coupled CESM simulation with grid resolutions fine enough to capture realistic tropical cyclone circulations and ocean dynamics within a global framework. This effort represents the **cutting edge** of Earth system modeling, and it requires petascale computational power and a highly scalable platform. Blue Waters provides a unique opportunity to analyze tropical cyclone-climate connections using a state-of-the-art, comprehensive Earth system model. Results from this study will enable us to answer key science questions about tropical cyclones and climate that can fundamentally improve our understanding of climate variability and future changes.

DEPENDENCE OF THE DIRECTIONAL INTENSITY AND POLARIZATION OF LIGHT SCATTERED BY SMALL ICE CRYSTALS ON THEIR SHAPE AND SIZE: APPLICATIONS FOR AIRBORNE CLOUD PROBES

Allocation: Illinois/750Knh
PI: Junshik Um[1]
Collaborator: Greg M. McFarquhar[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Current probes installed on aircraft for measuring the sizes of ice crystals smaller than 50 micrometers assume the intensity of light scattered by particles passing through the sample volume is a function of particle size. The relationship between particle size and scattered light used in current probes is based on Mie theory, which assumes the refractive index of a particle is known, and that all particles are spherical. Not only are small crystals not spherical, but they also can have a wide variety of shapes. To improve airborne measurements of small crystals, precise relationships between the light scattered by nonspherical particles and their size and shape are required and are based on accurate calculations of crystal scattering properties. However, such calculations use numerically exact methods and hence require a large amount of computing time and memory, which increase with crystal size. Thus, we use Blue Waters for these calculations.

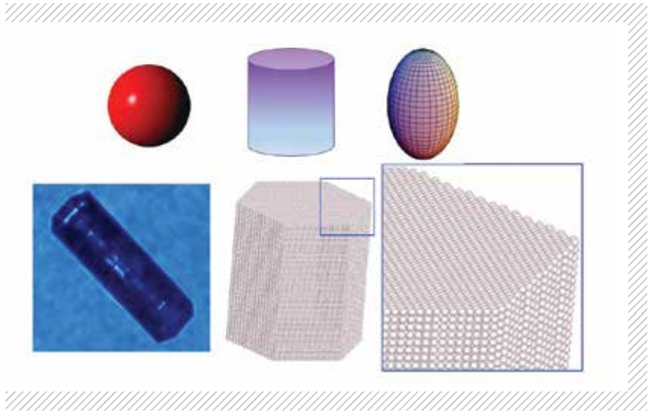

FIGURE 1: Top: A sphere, cylinder, and spheroid represent shapes of ice crystals in previous studies. Bottom: Real image (left) of ice crystals and its representation in DDA calculations.

INTRODUCTION

Cloud particle size distributions measured by cloud probes installed on aircraft are used to both develop parameterizations for and evaluate the results of satellite retrieval algorithms, and numerical models. Some current cloud probes (e.g., forward scattering spectrometer probes) measure the sizes of ice crystals smaller than 50 micrometers assuming that the intensity of light scattered by particles in the forward, and sometimes backward direction is a function of particle size. The retrieval of ice crystal sizes from satellites and application of parameterization schemes for numerical models also rely on relationships between light scattering and particle size.

The relationship between particle size and scattered light used in current forward scattering probes is based on Mie theory, which assumes that the refractive index of a particle is known and that all particles are spherical. Small crystals are not spherical, and a wide variety of shapes have been used to represent them. Although it is known that the scattering properties of non-spherical ice crystals differ from those of spheres, the impacts of non-sphericity on retrieved sizes for airborne probes, satellite retrievals, and model parameterizations are unknown. Further, more advanced forward scattering probes and satellite sensors that use the polarization and intensity of light also require knowledge of the scattering properties of non-spherical ice crystals.

Although a few studies have calculated the scattering properties of nonspherical particles assuming cylindrical and spheroidal shapes (Fig. 1), the shapes used are far different from realistic shapes of small ice crystals. Thus, to improve airborne measurements of small crystals, satellite retrieval algorithms, and numerical models, precise relationships between the intensity and polarization of light scattered in multiple directions with the size and shape of nonspherical crystals are required. However, accurate calculations of the scattering properties of crystals using numerically exact methods (i.e., solving Maxwell's equations) require large amounts of computing time and memory, which increase with crystal size. Thus, Blue Waters is an important resource for these calculations.

METHODS & RESULTS

The scattering properties (i.e., angular dependence of scattering, intensity of scattering, and polarization) of ice crystals with a maximum dimension (D_{max}) smaller than 50 micrometers are calculated at a non-absorbing wavelength (λ=0.55 μm) using a numerically exact method (i.e., the discrete dipole approximation, DDA) [1]. For these calculations, hexagonal ice crystals with varying aspect ratios (AR=length/width, AR=0.1, 0.25, 0.5, 1.0, 2.0, and 4.0) are used to represent the shapes of natural small ice crystals (Fig. 1). In DDA calculations, a hexagonal ice crystal is divided into elementary polarization units called dipoles (Fig. 1) with each or a group of dipoles assigned to a single core of Blue Waters for the calculations.

To quantify errors in the sizing of ice crystals used in current forward scattering probes, Figure 2 shows differences in the calculated D_{max} assuming spherical (i.e., Mie calculations) and hexagonal (i.e., DDA calculations) shapes of ice crystals. For a given amount of scattered light measured by a forward scattering probe, multiple selections of particle size are possible because of variations in crystal shapes. In Figure 2, blue dots represent the average differences in calculated diameters, whereas red dots are the maximum and minimum differences in diameters. These differences represent the errors in sizing nonspherical ice crystals using current forward scattering probes. The errors depend on the sizes and ARs of ice crystals (i.e., each panel of Fig. 2). It is shown that differences in sizing particles are smallest for D_{max}=~10 μm, while larger for D_{max}<10 μm (due to light interference structures) and for D_{max}>10 μm (due to nonsphericity) regardless of the AR. The differences are larger when the shapes are more nonspherical (i.e., AR=0.1 and 4.0). The average differences are up to 41% with a prolate shape (i.e., AR=4.0). These differences, representing errors in the retrieval of particle sizes from current forward scattering probes, determined using realistic shapes (i.e., hexagonal shapes) of ice crystals are larger than those determined assuming either cylinder or spheroid shapes in previous studies. Thus, to improve the retrieval of particle sizes from current forward scattering probes, accurate calculations of the scattering properties of realistic shapes of ice crystals are required.

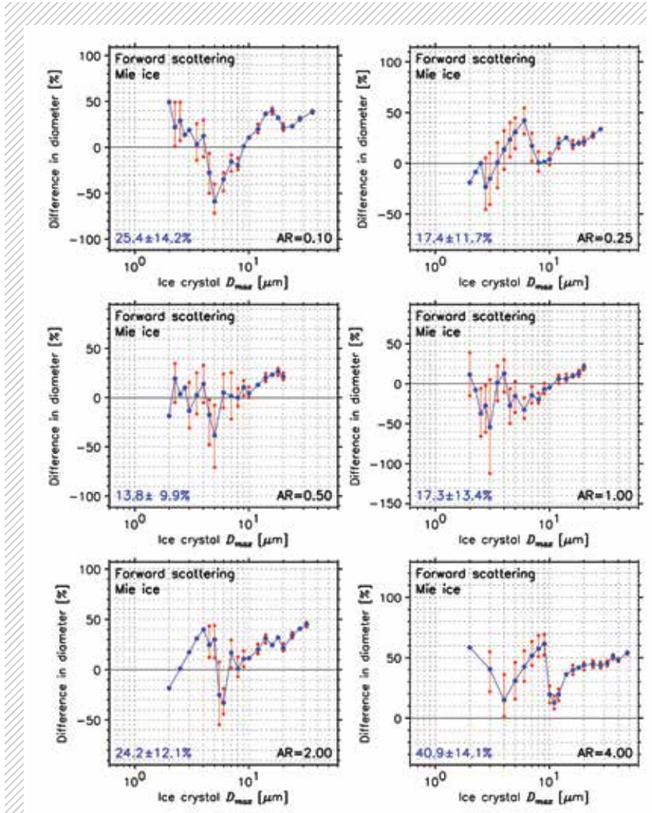

FIGURE 2: Differences (100% × (D_{max} – Mie D_{max})/D_{max}) between actual sizes of ice crystals and those determined based on Mie calculations in forward directions as a function of AR of hexagonal ice crystals (each panel). Mean and standard deviation of absolute values of blue circles are embedded in each panel.

WHY BLUE WATERS

Much larger computing resources are required to compute accurate scattering properties of realistic shapes (e.g., hexagonal columns and plates) of ice crystals compared to previous studies that used approximate shapes (e.g., spheres, cylinders and spheroids). The exact discrete dipole method was

needed for particles with the small sizes used here. Although the exact methods provide more accurate results, they require more computing time and memory that rapidly increases with particle size. The accuracy of measurements of particle sizes with airborne probes, radiative transfer models, and satellite retrieval algorithms depends heavily on accurate calculations of single-scattering properties of ice crystals. The petascale Blue Waters is an important resource for completing these calculations.

NEXT GENERATION WORK

Natural ice crystals larger than 50 micrometers have various shapes with much more complex internal (e.g., inclusions) and external features (e.g., surface roughness) compared to small ice crystals. Thus, calculations of scattering properties of larger ice crystals will require much larger computing resources than available on Blue Waters. We want to perform such calculations using a next-generation Track-1 system.

PUBLICATIONS AND DATA SETS

Um, J., and G. M. McFarquhar, Formation of atmospheric halos and applicability of geometric optics for calculating single-scattering properties of hexagonal ice crystals: Impacts of aspect ratio and ice crystal size. *J. Quant. Spectrosc. Radiat. Transfer*, 165 (2015), pp. 134-152, doi:10.1016/j.jqsrt.2015.07.001

Um, J., and G. M. McFarquhar, Light scattering by atmospheric ice crystals: Application to forward scattering probes. *International Symposium on Radiation*, Auckland, New Zealand, April 16-22, 2016.

McFarquhar, G. M., and J. Um, Light scattering by atmospheric hexagonal ice crystals for determination of applicability of geometric optics and formation of atmospheric circumscribed halos, *International Symposium on Radiation*, Auckland, New Zealand, April 16-22, 2016.

BUILDING A DATA ASSIMILATION FRAMEWORK FOR FORECASTING VOLCANIC ACTIVITY DURING PERIODS OF UNREST

Allocation: Illinois/25.0 Knh
PI: Patricia M. Gregg[1]
Collaborators: Yan Zhan[1] and J. Cory Pettijohn[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

A primary goal of the University of Illinois Urbana-Champaign Volcano Lab is to develop innovative strategies for combining volcano-monitoring datasets with thermomechanical models to better understand the dynamics of triggering a volcanic eruption. We use Blue Waters to develop and test a framework for **multi-data stream data assimilation** by conducting a series of eruption "hind casts" for recent eruptions at Sinabung Volcano in Indonesia. Blue Waters is uniquely capable of handling the computational expense for our ensemble-based modeling approach, which has compute times and storage requirements considerably outside the capabilities of traditional high-performance computing (HPC) resources. Ultimately, this work will provide a critical foundation for future interdisciplinary efforts to model volcano evolution and mitigate volcano disasters for vulnerable populations worldwide.

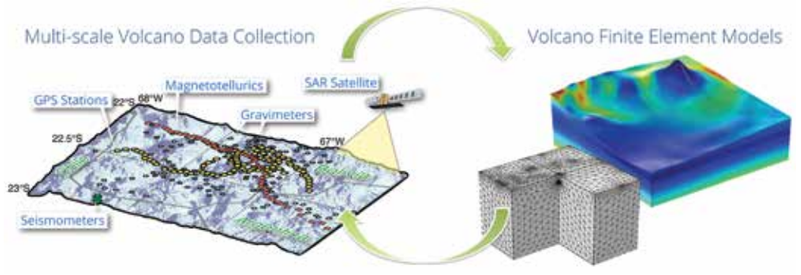

FIGURE 1: Statistical data assimilation techniques combine disparate observations with models to produce forecasts of an evolving system. The data inform the trajectory of the models and the models in turn inform future data targets and collection strategies. Left: A map depicting the NSF PLUTONS Project at Uturuncu Volcano, Boliva [7]. Right: FEM mesh and modeled stress for Sinabung Volcano, Indonesia.

INTRODUCTION

A primary motivation for investigating volcanic systems is to develop the ability to predict eruptions and mitigate disaster for vulnerable populations. In recent decades, the evaluation of volcanic activity has been greatly enhanced by coupling remote (e.g., satellite and global seismic arrays) and local observations (e.g., GPS, gas emissions, and seismometers) to provide early warning of an imminent eruption or information on the evolution of a magma system during a volcano crisis. Concurrently, thermomechanical models of magma reservoirs have significantly advanced our understanding of eruption-triggering mechanisms beyond the temporal and spatial limitations of our observations [1,2].

Volcano monitoring datasets are commonly analyzed using analytical inversions techniques or by optimizing finite element models [3,4]. While these approaches work well for combining models with one or two data streams, they are static assessments of the system state that do not provide updates or forecasts and are limited in their scope. Alternatively, statistical data assimilation methods were developed to systematically link data with models to provide model updates. Significant advancements in data assimilation have been made in many fields, including engineering, hydrology, physical oceanography, and climate modeling, to incorporate disparate datasets into dynamic, nonlinear models and provide model forecasts [5]. Sequential model-data fusion methods provide a framework for integrating large, disparate data sets into time forward, forecast models (Fig. 1). Data are used to nudge the model trajectory and provide updates of the system's evolution, and models inform future data targets.

Our current efforts on Blue Waters are focused on developing strategies for rapid assimilation of monitoring datasets into evolving geodynamic models to provide near-real-time forecasts and assessment of volcanic unrest. To that end, we are adapting data assimilation methods developed in other fields to combine observations from volcanoes experiencing unrest with thermomechanical finite element models (FEMs) to calculate volcano evolution [6].

METHODS & RESULTS

We have adapted the Ensemble Kalman Filter (EnKF, [5]) sequential data assimilation method to assimilate volcano-monitoring data from satellite and ground-based observations into geodynamic models. The EnKF utilizes a Markov Chain Monte Carlo (MCMC) approach to create suites, or ensembles, of models that are updated sequentially as new observations become available.

Preliminary results indicate that the EnKF is a powerful tool well suited for the problem of forecasting volcanoes experiencing unrest. Our Blue Waters Exploratory Allocation is allowing us to test the feasibility of a large-scale data assimilation approach for volcano monitoring. This is the **first** effort of its kind and has great potential for significantly advancing the field of volcano hazards.

WHY BLUE WATERS

The EnKF is an ensemble based sequential data assimilation method that requires calculating hundreds to thousands of finite element models at each time step. While the EnKF analysis step has been optimized to run very swiftly, the computational expense of running and storing hundred to thousands of finite element models for each time step in the EnKF analysis is cost prohibitive for traditional HPC resources. Blue Waters is uniquely positioned to handle our computational needs and has allowed us to make rapid progress and develop **ambitious** approaches without being hampered by computational limitations.

FIGURE 1: Vorticity field of the 30 meter simulation shortly following tornadogenesis. Red shading indicates cyclonic (counterclockwise) vorticity, which is most intense within the tornado (dark red vertical tube), but is abundant behind and along the leading edge of the thunderstorm's cold pool to the right of the tornado vortex.

SIMULATING THE MOST VIOLENT THUNDERSTORMS

Allocation: Illinois/510 Knh
PI: Robert Wilhelmson[1]
Co-PI: Leigh Orf[2]
Collaborators: Bruce Lee[3], Catherine Finley[4], and Adam Houston[5]

[1]University of Illinois at Urbana-Champaign
[2]University of Wisconsin-Madison
[3]High Impact Weather Research & Consulting, LLC
[4]St. Louis University
[5]University of Nebraska-Lincoln

EXECUTIVE SUMMARY

Tornadoes remain the subject of intense research due to the danger they present to society. The least common but most devastating tornadoes occur within supercell thunderstorms. This research focuses on supercells that produce strong, long-lived tornadoes that would be ranked EF4 or EF5 (the most intense categories of the Enhanced Fujita scale) in nature. Thus far, we have simulated devastating tornadoes spawned by simulated supercells initialized within the 24 May 2011 central Oklahoma environments, where a supercell producing a long-lived EF5 tornado occurred that day. Simulations run at ultra-high resolution and reveal flow features that contribute towards the genesis and maintenance of a long-lived devastating tornado. These features include a streamwise vorticity current (SVC, helically rotating air moving parallel to its rotational axis) formed by the consolidation of horizontal vorticity originating within the storm's cold pool. The SVC is tilted vertically by the storm's intense updraft, helping to strengthen the storm.

INTRODUCTION

To accurately predict the formation and behavior of tornado-producing thunderstorms, one must first understand how such thunderstorms work, and that is the focus of our research. We are currently unable to predict whether a supercell (either forecast or already formed) will produce a tornado at all, much less a long-lived extremely violent tornado. Our research is particularly important to people who live in tornado-prone regions, such as "Tornado Alley" in the Great Plains of the United States.

METHODS & RESULTS

The actively-developed CM1 cloud model, which was designed to run efficiently on computers such as Blue Waters, was chosen as the primary application. Substantial effort was made to rewrite CM1's I/O layer to best exploit the parallel Lustre filesystem of Blue Waters and to simplify post-processing and visualization. Following this, various supercell environments were explored until a simulation produced a long-lived devastating tornado. This simulation was visualized and analyzed, and an overview of the simulation was published in the Bulletin of the American Meteorological Society, a high-impact journal.

WHY BLUE WATERS

Our group has been utilizing Blue Waters from the start. Blue Waters has been essential for this research because of the scale of the problem we are trying to solve, which requires tens to hundreds of thousands

FIGURE 2: Cloud and rain fields of the 20 meter simulation during the most intense phase of the simulation. Multiple vortices are apparent within the cloud field, which is a feature of devastating tornadoes often observed in the field.

of computing cores. The large amount of scratch and nearline storage are essential to the success of this project. We have been very pleased with the Blue Waters support staff who have responded to issues in a timely and professional manner.

NEXT GENERATION WORK

Recently, a simulation of a supercell initialized in the 24 May 2011 environment at a resolution of 20 meters was completed. Similar to the 30-meter simulation, a long-lived EF5 tornado occurs, but with significantly more physical realism based upon volume rendering of the cloud, rain, and vorticity fields. We will continue to examine this simulation as well as explore the sensitivity of the simulation to parameters such as surface treatment and microphysics. We also propose to execute simulations in other environments where devastating tornado-producing thunderstorms occurred.

PUBLICATIONS AND DATA SETS

Orf, L., R. Wilhelmson, and L. Wicker, Visualization of a simulated long-track EF5 tornado embedded within a supercell thunderstorm. *Parallel Comput.*, 55 (2016), p. 28–34, doi:10.1016/j.parco.2015.10.014

Orf, L., R. Wilhelmson, B. Lee, C. Finley, and A. Houston, 2016: Evolution of a Long-Track Violent Tornado within a Simulated Supercell. *Bull. Am. Meteorol. Soc.*, in press, doi:10.1175/BAMS-D-15-00073.1

FIGURE 3: Vorticity field of the 20 meter simulation towards the end of the tornado's lifecycle. At this stage, the tornado has widened and taken on a multiple vortex structure.

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

Allocation: Illinois/300 Knh
PI: Matthew West[1]
Co-PI: Nicole Riemer[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

This research addresses key uncertainties associated with aerosol-climate impacts. Aerosol particles influence the large-scale dynamics of the atmosphere and climate because they directly interact with solar radiation by scattering and absorbing light, and indirectly by acting as cloud condensation nuclei. Their sizes range from nanometers to micrometers, and a major source of difficulty in understanding the climate impact of aerosols is attributed to scale interactions. There is a high computational cost required for modeling those interactions. The particle-resolved 3D model WRF-PartMC-MOSAIC, which has the unique ability to track size and composition information on a per-particle level, was developed to address this problem while remaining computationally feasible. In combination with efficient algorithms, a resource with the capabilities of Blue Waters is essential to be able to address questions regarding the importance of aerosol composition for prediction of weather and climate at the regional scale.

INTRODUCTION

Many of the greatest challenges in atmospheric modeling and simulation involve the treatment of aerosol particles, ranging from the prediction of local effects on human health [1] to the understanding the global radiation budget via the aerosol indirect and direct effects [2]. Models provide important insights in the study of aerosols but experience a trade-off between the representation of physical detail and spatial resolution. Aerosol modeling has proven difficult because of the complex microscale physics of individual particles, which are not individually resolved in models largely due to computational constraints. Current methods of representing the high-dimensional and multi-scale nature of aerosol populations still make large simplifications. While this makes computation much cheaper, it introduces unknown errors into model calculations. This has far-reaching consequences for the estimation of climate-relevant aerosol quantities, such as the aerosols' ability to scatter and absorb sunlight as well as their ability to form clouds.

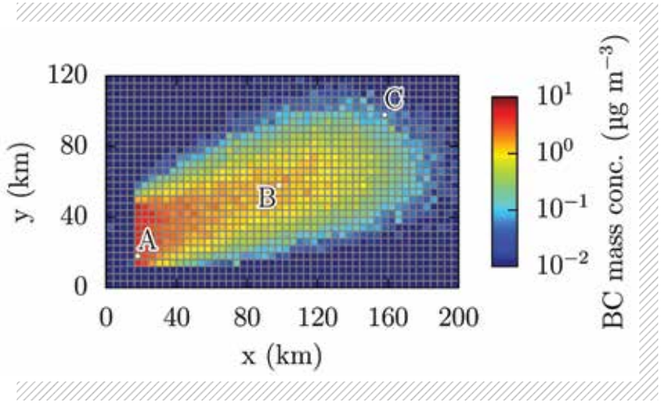

FIGURE 1: Black carbon mass concentration for an urban plume scenario where black carbon particles are emitted and transported by simulated wind fields. Location A is located near the emission source while Locations B and C are located downwind from the emission source.

METHODS & RESULTS

To overcome the current limitations in representing aerosols and associated uncertainties, the particle-resolved model PartMC-MOSAIC [3] was coupled to the state-of-the-art 3D Weather Research and Forecast (WRF) model. The two models complement each other. PartMC-MOSAIC is a highly detailed aerosol model that tracks the size and complex composition of individual particles in the atmosphere but is unable to resolve spatial heterogeneities of aerosol populations. The 3D, regional WRF model is an advanced numerical weather model that captures the transport of chemical species in the atmosphere but assumes a simplified aerosol representation. The combined model has the advantages of both.

Figure 1 shows the spatial distribution of black carbon mass concentrations near the surface for a scenario simulated on Blue Waters. While this is a bulk quantity common to any chemical transport model, the particle-resolved aerosol representation provides greater detail of particle composition. Figure 2 shows the aerosol mixing state of black carbon at three highlighted locations in the domain shown in Figure 1. At each of these locations, a complex continuum of mixing states is present where particles of similar diameters can have very different chemical composition. Particles are emitted with particular fractions of black carbon determined by their source characteristics, and composition evolves due to coagulation and condensation of secondary gas species. Location A shows the presence of freshly emitted black carbon-containing particles, which consists of high black carbon mass fractions. Location B, located downwind from the emission source, has particles that have undergone aging, which decreased the black carbon mass fraction by adding mass from condensation of secondary gasses. Meanwhile, very little black carbon emissions have been transported to Location C. By using the physical process-level particle-resolved modeling, this phenomenon is accurately captured, allowing for the **most accurate** simulations aerosol impacts on climate to date.

WHY BLUE WATERS

Blue Waters allows for a **cutting-edge** model that pushes 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 challenges due to computationally intensive equations per particle and memory requirements to track complex particle composition. Blue Waters provides the tens of thousands of cores, fast interconnections between those cores, and sufficient memory per process required for the simulation of aerosols at high spatial and compositional resolution.

NEXT GENERATION WORK

The next-generation Track-1 systems will allow future particle-resolved simulations at higher spatial resolution, for larger regions, and with a greater number of tracked particles to provide greater accuracy. These simulations will target regions where single particle measurements are made available from field campaigns such as the Carbonaceous Aerosols and Radiative Effects Study (CARES) [4] and future upcoming aerosol studies.

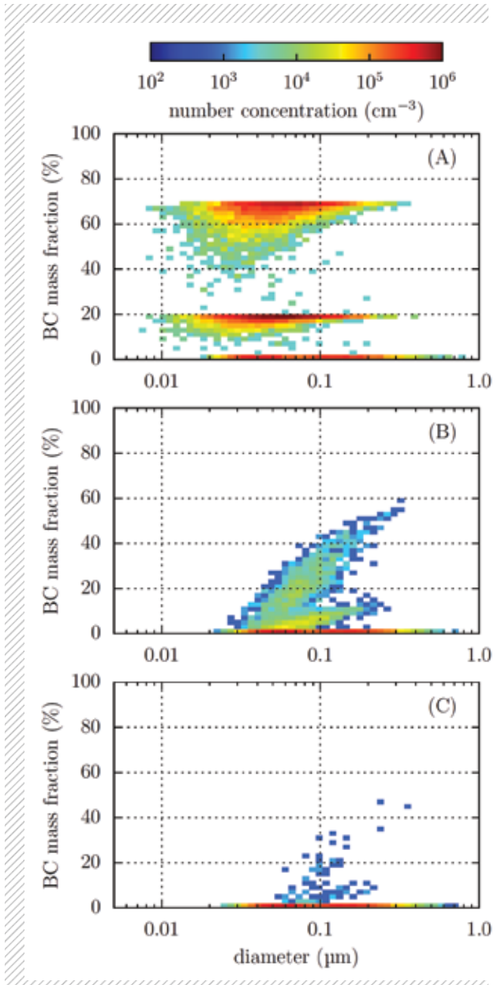

FIGURE 2: Two-dimensional number distributions as a function of particle dry diameter and black carbon mass fraction for locations shown in Figure 1. The number distribution indicates the amount of particles within a range of diameters and a range of fractions of black carbon mass.

LARGE EDDY SIMULATION OF SEDIMENT TRANSPORT AND HYDRODYNAMICS AT RIVER BIFURCATIONS

Allocation: Illinois/250 Knh
PI: Marcelo H. Garcia[1]
Co-PI: Paul Fischer[1]
Collaborator: Som Dutta[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Bifurcations are fundamental features of all river systems. Our current study focuses on a specific class of bifurcations called diversions. Experiments have shown that the distribution of near-bed sediment between the downstream channels at a diversion is not proportional to the water flow distribution, with a disproportionately higher amount of sediment going into the lateral channel. A better understanding of this phenomenon will help in efficient design of river diversions, which are used for navigation and flood-mitigation and have been proposed as a solution for reclaiming deltas that are slipping under the sea due to rising sea levels. The current study employs large eddy simulations (LES) and direct numerical simulations (DNS) of the flow in idealized diversions of different configurations, with the sediment modeled as Lagrangian particles. The simulations have been conducted using a highly scalable spectral element based incompressible Navier-Stokes solver, Nek5000. The simulation conditions are comparable to laboratory experiments, which make these simulations some of the **largest** and **most complex** to date in the area of river mechanics. Blue Waters provides the computing resources to undertake these highly-resolved simulations. Currently, simulations have been conducted for up to 243.648 million computational points, with strong scaling being shown up to 32,768 MPI ranks.

INTRODUCTION

Bifurcations are fundamental features of all river systems. The current study focuses on a specific class of bifurcations called diversions, where one source channel splits into a main channel and a lateral branch. Laboratory experiments conducted by Bulle in 1926 [1] and later by other investigators have shown that the distribution of near-bed sediment between the downstream channels is not proportional to the flow distribution, with a disproportionately higher amount of sediment going into the lateral channel. This non-linear phenomenon is often referred to as the Bulle Effect. The current study investigates the mechanisms behind this phenomenon through high-resolution numerical simulations of the flow and sediment transport simulations at the scale of, and for configurations similar to, Bulle's experiments.

A better understanding of this phenomenon will help in efficient design of river diversions, which are used for navigation and flood-mitigation. Diversions have been proposed as a way to reclaim deltas that are sinking under the sea due to sea levels that are rising as a consequence of climate change [2]. A prime example is the Mississippi Delta; different diversion designs currently are being explored for diverting water and sediment from the Mississippi River [3]. Our study will also help more accurately predict the short and long-term geomorphological evolution of river bifurcations, thus furthering the state of the art in the field of river mechanics. Our research will also provide insights that will help improve the numerical models used for simulating field-scale bifurcations. A better understanding of the fundamental mechanism behind the Bulle Effect will also help shed light on the related phenomenon of vorticity-driven sediment transport, which affects both natural and manmade systems [4]. Finally, bifurcations are not only found in rivers but also in other places, such as the carotid bifurcation in the human body [5]. Thus, the current study will contribute to our general understanding of dynamics and transport at bifurcations.

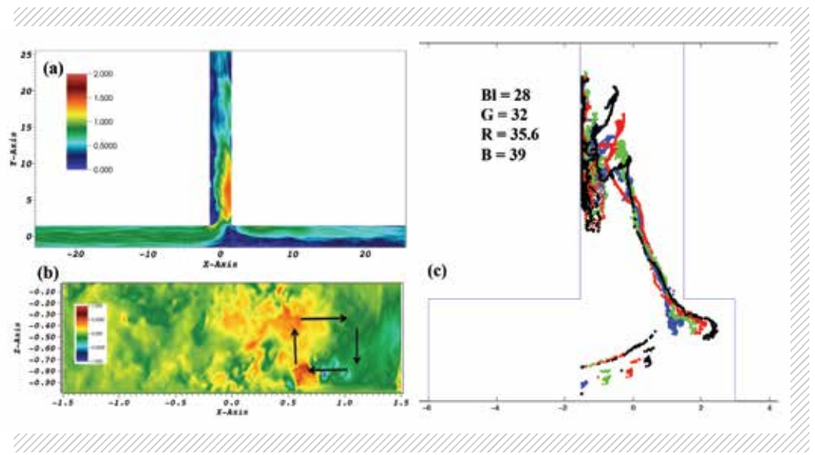

FIGURE 1: (a) Time-averaged velocity magnitude at 5 percent of the height of the channel. (b) Instantaneous velocity in the vertical direction, showing a clear signal of clockwise rotating vortex. (c) Evolution of horizontal positions of the sand particles in time, with Bl being Blue and B black. (c) Clearly shows that most of the sediment enters the lateral channel, which can be attributed to the majority of the flow near the bottom entering the lateral channel [see (a)] even when the total flow split is 50:50.

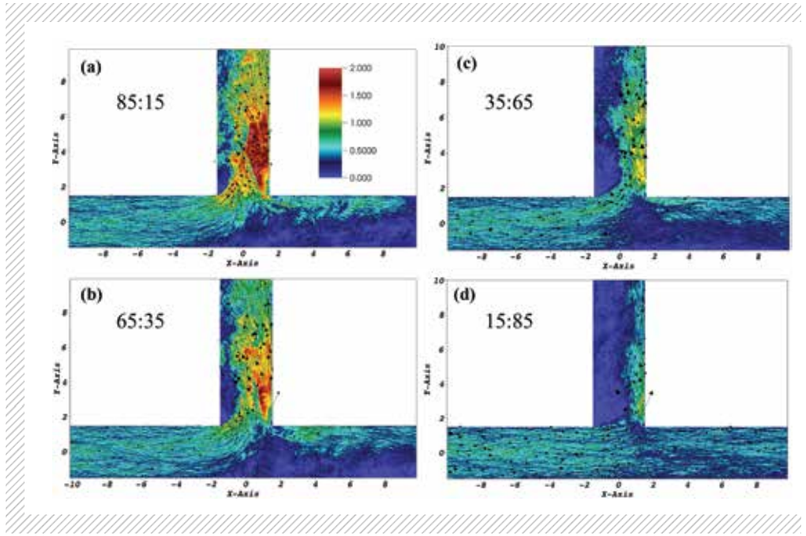

FIGURE 2: The figure shows the instantaneous velocity magnitude of the flow near the bottom of the channel, for a flow of bulk Reynolds number = 25000. Four cases with different % of the total flow going into the lateral channel have been presented; e.g. in case (a) 85 % of the flow is going into the lateral channel. The figures show that even for cases in which, a smaller % of the total flow moves into the lateral channel [case (c)], majority of the flow near the bottom continues to move into the lateral channel.

METHODS & RESULTS

High-resolution large eddy simulations (LES) and direct numerical simulations (DNS) have been conducted of the flow in idealized diversions of different configurations, with sediment being modeled as Lagrangian particles. Use of LES or DNS depends on the Reynolds number of the flow, with simulations conducted for Reynolds number values from 10 to 25,000. For the cases with Reynolds number of 25,000, the conditions are comparable to the laboratory experiments, which make these some of the **largest** and **most complex** simulations in the area of river mechanics. The simulations have been conducted using the open-source, spectral element based higher-order incompressible Navier-Stokes solver Nek5000 [6]. The spectral element method (SEM) combines the accuracy of spectral methods and the flexibility of numerical methods based on local approaches, like the finite elements method [7]. Sediment transport in the flow has been modeled using Lagrangian particle tracking. For simulating transport of poly-disperse sediment particles efficiently, a semi-implicit Lagrangian particle algorithm has been developed for the current study [1].

The first case presented here is of a 90-degree diversion at Reynolds number 20,000, with 50% of the total flow going into the lateral channel [2]. The time-averaged velocity magnitude of the flow near the bottom has been plotted (Fig. 1a) along with the position of the sediment particles (Fig. 1c). The majority of the sediment moves into the lateral channel, which is in agreement with the laboratory experiments. This highly non-linear behavior can be attributed to the fact that most of the flow near the bottom moves into the lateral channel, even though only 50% of the total flow moves into the lateral channel. Another important flow feature that influences the dynamics of the sediment at the diversion is the clockwise rotating vortex near the right-wall of the lateral channel (Fig. 1c). Instantaneous velocity magnitude near the bottom of the channel for different flow splits has also been presented for the 90-degree diversion at Reynolds number of 25,000 (Fig. 2). It shows that even when only 35% of the total flow moves into the lateral channel, most of the flow near the bottom moves into the lateral-channel (Fig. 2c). These results hint toward the underlying mechanism behind the non-linear Bulle Effect.

We are conducting more simulations with poly-disperse sediment for different flow-splits, Reynolds numbers and diversion-angles. The completed study will not only help to fully understand the underlying mechanism behind the Bulle Effect, it will also help in developing a reduced-order model for the phenomenon. This study provides new insights into the hydrodynamics and sediment transport at bifurcations, and it also shows that high-resolution LES can be used to study complex river-mechanics problems.

WHY BLUE WATERS

The current study **pushes the limit** of the scale at which high-resolution large-eddy simulations have been used to study complex multi-phase river mechanics problems, warranting the use of Blue Waters, which can provide sustained computing power at an **unprecedented scale**. For the current study, simulations have been conducted for up to 243.648 million computational points, with the code scaling strongly up to 32,768 MPI ranks. Without access to a petascale high-performance computing system like Blue Waters, completing the study in any realistic time-frame would be impossible. One of the most useful ways to understand a phenomenon is through visualization, thus we are working with Blue Waters staff to create an animation of the flow and sediment transport for one of the simulated cases.

NEXT GENERATION WORK

Access to the next generation of Track-1 HPC system will allow us to step up the scale at which we work, thus allowing us to conduct high-resolution LES of environmental flows at scales and with complexity similar to that of nature. This will allow us to fathom the underlying mechanisms of different environmental phenomena, thus aiding in improved predictions of different natural processes.

HIGH-RESOLUTION EARTH SYSTEM MODELING FOR INTERNATIONAL CLIMATE ASSESSMENT

Allocation: NSF PRAC/31.5 mnh
PI: Robert Rauber[1]
Co-PIs: Warren Washington[2], Gerald Meehl[2], Joe Tribbia[2], and Justin Small[2]
Collaborators: Donald Wuebbles[1], Zachary Zobel[1], Xin-Zhong Liang[3] and Chao Sun[3]

[1]University of Illinois at Urbana-Champaign
[2]National Center for Atmospheric Research
[3]University of Maryland

EXECUTIVE SUMMARY

This is a collaborative project to investigate climate change and regional climate processes using higher model resolutions that would not be possible without resources such as Blue Waters. These simulations allow us to investigate climate change, specifically regarding extremes; regional model performance through dynamic downscaling; and to examine regional model uncertainties about atmospheric processes. These studies are on the **leading edge** of high resolution global and regional climate modeling and provide the pathway for the next generation climate models and assessments. Our studies also allow better understanding of model uncertainties. No previous global modeling study has provided such high-resolution information at the regional scale to fully analyze the potential impacts of climate change on human society across many different sectors (e.g., health, food, water, energy, transportation) and on ecosystems.

INTRODUCTION

This project has several objectives. One objective is to quantify changes in future climate extremes using two high-resolution versions of the Community Earth System Model (CESM): one with a high-resolution atmosphere (0.25° atmosphere, 1.0° ocean) and the other with high-resolution atmosphere and ocean (0.25° atmosphere, 0.1° ocean) where ocean eddies are derived internally instead of being parameterized. Another objective is to evaluate the effects of dynamical downscaling (DD). DD inputs time varying boundary conditions from a lower-resolution, limited-area, global climate model (GCM) into a regional climate model (RCM) [1,2], which is the Weather Research and Forecast (WRF) model. Finally, work on climate parameterization uncertainty continues via a multiple physics ensemble (MPE) analysis.

METHODS & RESULTS

CESM: The CESM pre-industrial (PI) control simulation with a fully coupled 0.25° atmosphere and 1.0° ocean has completed 100 years of simulation. In extended control simulations, it is desirable to have the radiative balance at the top of the atmosphere below |0.1| W/m^2 so that model drift is small and follow-on simulations, such as twentieth century historical or associated climate sensitivity simulations, can be branched with suitably balanced and realistic initial conditions. The PI control run on Blue Waters reached 0.03 W/m^2.

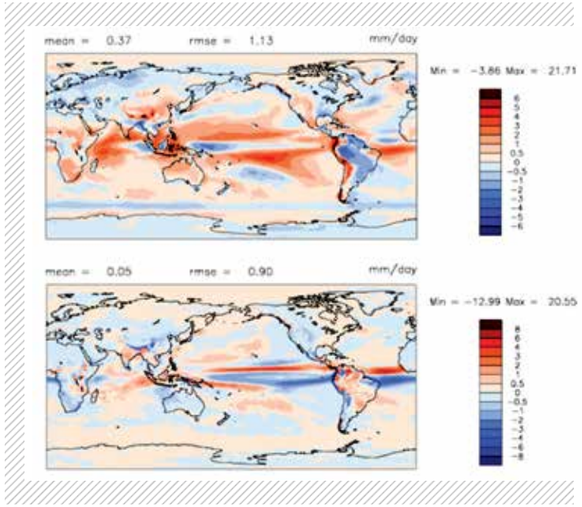

FIGURE 1: Annual mean total precipitation rate difference between a baseline low resolution (1.0°) PI control and GPCP observations (top panel) and between the high resolution (0.25°) PI control conducted on Blue Waters and the baseline low resolution PI control (bottom panel). Wherever the color is opposite in these two panels are regions in which high resolution represents precipitation rate better than low resolution.

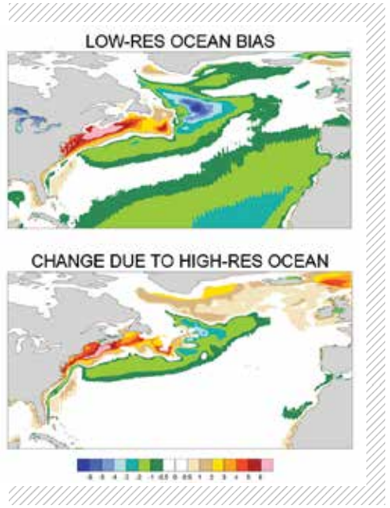

FIGURE 2: Sea surface temperature difference (°C) between the high resolution model (0.25° atmosphere / 1° ocean) and (top panel) observations [5] and (bottom panel) the high resolution ocean model (0.25° atmosphere / 0.1° ocean). Wherever the color is the same in these two panels are regions in which high resolution represents sea surface temperature better than low resolution.

The expectation in conducting climate simulations at higher resolution is not necessarily to improve the modeled mean climate state but to capture many aspects of extreme events and diurnal cycle patterns that cannot be resolved at a lower resolution. However, we have observed improvements in the mean state, notably in precipitation, in the PI control and anticipate that this will translate to a better representation of precipitation extremes. Figure 1 compares precipitation rate from the low and high-resolution models to observations. Three regions of improvement stand out: the Andes and interior Amazon, the Indian Monsoon region, and North America.

We are also studying Earth's climate at an **unprecedented** resolution (0.25° atmosphere and 0.1° ocean) with a CESM configuration validated recently for present-day climate [3]. It produces a smaller climatological sea surface temperature (SST) bias in the topics than standard resolution runs, and also a more realistic interannual variability of Eastern Pacific SSTs. At this very high resolution, increased confidence is given to potential future regional climate changes arising from natural and anthropogenic forcing mechanisms. For example, by resolving ocean mesoscale features, we get a better representation of the SSTs off the US East coast, due to a better Gulf Stream path (Fig. 2), which in turn modifies the atmospheric precipitation and storm tracks.

Dynamic Downscaling: The WRF model of the RCM was used to evaluate the performance of six dynamically downscaled decadal historical simulations with a 12 km resolution for a large domain (7200 km × 6180 km) covering most of North America. This study is based on high-spatial-resolution (12 km) RCM simulations using one numerical model to evaluate the model performance and the uncertainties coming from different boundary conditions and different model setups. The initial and boundary conditions are taken from three separate GCMs and one from reanalysis data. Figure 3 shows the initial findings for changes in seasonal and annual precipitation using CCSM4 boundary conditions with the RCM 8.5 scenario for 2085-2094 [4]. The hatched areas on the figures are statistically significant changes in seasonal precipitation for this decade. While this figure focuses on seasonal averages, further analysis will be needed to help identify regional and local climate extremes for using different GCMs as boundary conditions.

CWRF Uncertainty Analyses: We have continued using the regional CWRF model [5] in an MPE configuration at 30 km grid spacing over the United States to examine uncertainties in the treatment of cloud, aerosol, and radiative processes. When driven by the ECMWF-Interim Reanalysis (ERI), the CWRF ensemble mean exhibits higher prediction skills than the ERI. The variabilities of almost all ensemble members are closer to the normalized standard deviation of the observations than that of the ERI data. This is physically meaningful because the CWRF has a higher resolution and can simulate the spatial variability more realistically.

WHY BLUE WATERS

In climate modeling it is necessary to (a) quantify model characteristics/sensitivity, (b) produce a sufficiently long control, and (c) simulate historical and future mitigation scenarios. Based on the number of simulated model years at high resolution, very large allocations are required that are not attainable without systems like Blue Waters. For example, the computing requirements for simulations using the high-resolution ocean model (0.1°) have increased four-fold over the coupled model with a 1.0° ocean.

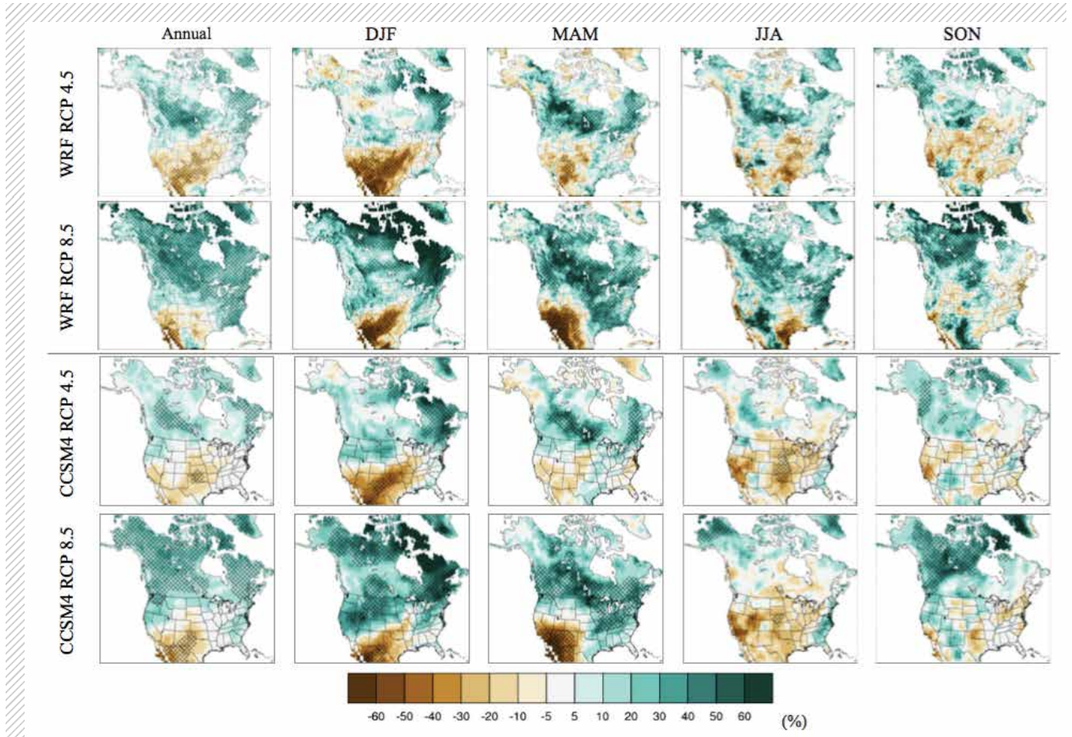

NEXT GENERATION WORK

The next generation systems will be required to achieve even higher resolution for the atmosphere model, eventually to represent small-scale processes without parameterizations, nested grids to study regional impacts, and more complex biogeochemistry. Striving for a 6 km horizontal resolution should enhance understanding of tropical cyclones and more accurately account for sophisticated land-ice processes (i.e., melting ice contribution to sea level rise). The 6km horizontal resolution translates to approximately 100 times more computation.

FIGURE 3: Difference in yearly and seasonal rainfall between 2085-2094 and 1995-2004 [4]. WRF simulations were driven by CCSM4 boundary conditions using bias correction for both decades. Hatched areas represent locations that experience statistically significant changes.

PUBLICATIONS AND DATA SETS

Zarzycki, C. M., et al., Impact of surface coupling grids on tropical cyclone extremes in high-resolution atmospheric simulations. *Geosci. Mod. Dev.*, 9:2 (2016), pp. 779-788.

HIGH-RESOLUTION SIMULATIONS OF CUMULUS ENTRAINMENT

Allocation: Blue Waters Professor/240 Knh
PI: Sonia Lasher-Trapp[1]
Collaborators: Daniel Moser[1] and Dave Leon[2]

[1]University of Illinois at Urbana-Champaign
[2]University of Wyoming

EXECUTIVE SUMMARY

Understanding and predicting the rate at which liquid and ice are depleted in cumulus clouds by the entrainment of dry air, which affects their vertical development, longevity, and ability to precipitate, has been elusive. Observations alone are not sufficiently comprehensive to determine the key scales and motions involved. Simulations performed on Blue Waters are enabling us to address this problem numerically to determine (a) physically the most important scales of turbulence and (b) the minimum model resolution required for predicting the bulk cloud properties resulting from entrainment. A new approach for direct calculation of entrainment in a single simulated cloud has been applied for this purpose. The simulation results are consistent with new observational quantification of eddy sizes at the cloud edge, as detected by airborne cloud LIDAR.

INTRODUCTION

Deep convective clouds produce the majority of the earth's precipitation, and yet it is difficult to predict if developing cumulus clouds will attain the depth and longevity required to produce heavy rainfall, and/or become severe thunderstorms. Entrainment is the term for the process by which clouds bring dry air from outside the cloud inward. In time, entrainment not only reduces the cloud buoyancy, limiting its vertical development, but also depletes its liquid water by evaporation, limiting precipitation formation. A long-standing problem in meteorological models has been to reproduce how quickly entrainment accomplishes these results. This problem affects a broad range of atmospheric science problems, ranging from short-term weather forecasts in numerical weather prediction models to climate forecasts in regional and global climate models.

METHODS & RESULTS

Using a thermodynamic atmospheric profile collected from an observational study of real cumuli as input to the CM1 cloud-resolving model, we continually decreased the model grid spacing from 200 m down to 20 m to represent increasingly smaller "eddies" in a single cumulus cloud (Fig. 1). We applied a recently developed numerical technique [1] to find the edge of the high liquid water core within the cloud, and then at this edge directly calculated the air inflow and outflow from the core surface to quantify the net entrainment as a function of space, and of time. We then created one-minute averaged vertical profiles of the entrainment to compare among different simulations, and to compare to the amount of liquid water in the core of the cloud versus time.

The results show that as the model grid spacing is decreased, more entrainment occurs in the simulated clouds. The entrainment at 20 m grid spacing is nearly double that at 100 m, and about 1 ½ times that at 50 m grid spacing. The increased entrainment depletes the cloud core more quickly (Fig. 2), in closer agreement to airborne observations of multiple clouds on this day. It is not yet clear that the model results have converged at 20 m grid spacing, and so runs with smaller grid spacing are yet being conducted to resolve even smaller eddies. New wavelet analysis of airborne cloud LIDAR data suggests that eddies at the cloud tops on the order of 30-50 m wide are prevalent (Figure 3), and thus will likely need to be represented in our simulations to attain the most realistic computational results.

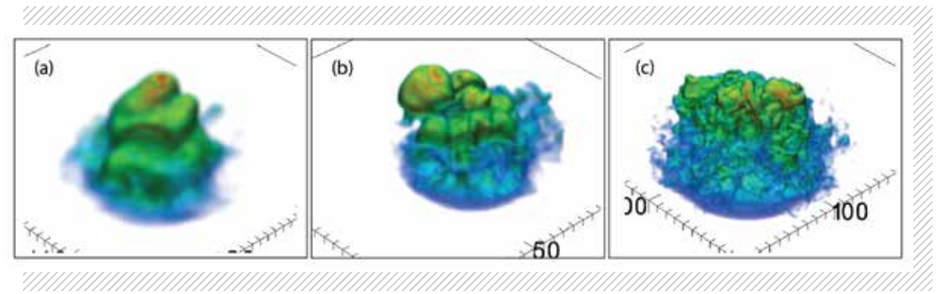

FIGURE 1: Cloud water shown at 1500 s for simulations run with (a) 100 m, (b) 50 m, and (c) 20 m grid spacing. The representation of increasingly smaller entraining eddies results from smaller grid spacing, and leads to fewer and weaker areas of high values of cloud water (red).

WHY BLUE WATERS

Our Blue Waters allocation is essential for testing the resolution-dependency of the entrainment process, in particular for determining the sizes of the eddies that are most critical to represent in these and future simulations of cumulus entrainment. Blue Waters allows us to push the spatial scale limit much farther than in the past, with its huge number of nodes, its high speed, and its large storage capability for high-resolution model output and analysis. The hardware needed to run these kinds of simulations quickly supersedes the limits of most computers.

NEXT GENERATION WORK

Our work over the past couple of years has focused on understanding what spatial and temporal scales must be considered in high-resolution simulations of cumulus clouds. This helps to accurately capture the most important cloud motions for evolution of the bulk properties (like total cloud water mass) in space and time. As we complete that investigation over the next year, including running ensembles to assess the generality of the results heretofore performed with single cloud simulations, we will also begin pursuing the question of how an accurate representation of cumulus entrainment in model simulations can affect the development, longevity, and severity of thunderstorms and the rain and/or hail that they produce. Such simulations will need to be run over a larger domain and for longer periods of time, and require detailed information on the cloud particles (water droplets, ice, raindrops, hail). By the time the Track-1 system is in place in 2019/2020, we will have some initial simulations performed, but the new system will allow us to run ensembles of simulations to assess the generality of the results for better understanding and prediction of flash floods and hail storms.

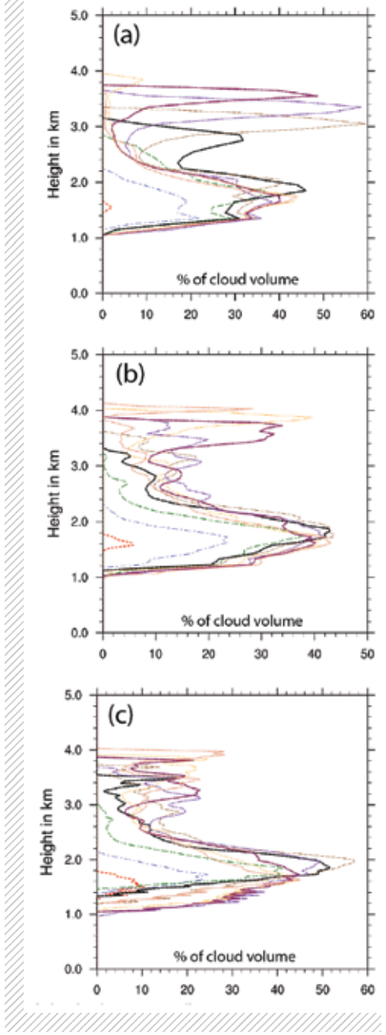

FIGURE 2: Vertical profiles of the percentage of the cloud volume at the given altitude that contains high values of cloud water (i.e. exceeding 80% of the maximum theoretical upper limit). The results are most important above 2 km height, where local areas of high liquid water can promote precipitation. As more entrainment is represented at smaller grid spacing (a) 100 m, (b) 50 m, (c) 20 m, less of the cloud volume contains regions of high cloud water.

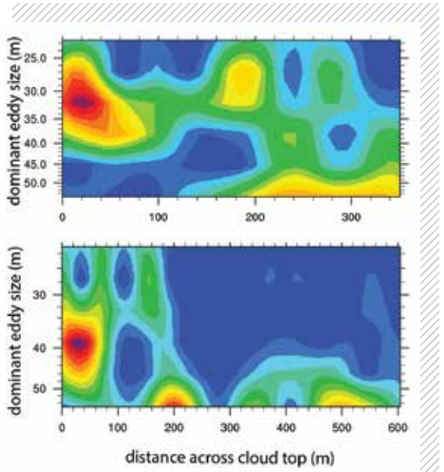

PUBLICATIONS AND DATA SETS

Lasher-Trapp, S., D. H. Moser, D. C. Leon, J. French, and A. M. Blyth, High Resolution Simulations of Cumulus Entrainment. *17th International Conference on Clouds and Precipitation*, Manchester, England, July 25-29, 2016.

FIGURE 3: Examples of wavelet analysis performed from scans across the tops of two real cumulus clouds with an airborne cloud lidar. "Hot spots" (red) show areas where cloud top eddies of the size shown on the vertical axis are prevalent. This observational evidence thus warrants additional simulations at higher resolution to investigate if these smaller eddies contribute significantly to entrainment.

FORECASTING GLOBAL CROP PRODUCTIVITY USING NOVEL SATELLITE DATA AND PROCESS-BASED MODELS

Allocation: Blue Waters Professor/240 Knh
PI: Kaiyu Guan[1]
Co-PI: Bin Peng[1]
Collaborator: Min Chen[2,3]

[1]University of Illinois at Urbana-Champaign
[2]Carnegie Institution for Science
[3]Stanford University

EXECUTIVE SUMMARY

The ultimate goal of this project is to improve predictability for global crop yield by integrating site measurements, advanced remote sensing observation, and process-based modeling. We took our first step forward by focusing on the high-temperature impacts on corn/soybean yield in the U.S. Corn Belt. Different pathways of high-temperature impacts on crop yield are considered in our newly developed CLM-APSIM modeling framework, which combines the strengths of earth system model and agronomy crop models. We are conducting parameter sensitivity analysis and optimization as well as a set of historical simulation experiments aimed at disentangling the contribution of different mechanisms to high-temperature impacts on crop yield. Projection runs will also be conducted shortly to explore the impact of high temperatures on crop yield under various climate change scenarios.

INTRODUCTION

Temperatures over the terrestrial areas of the planet are expected to increase at least twice the global average, with specific regions experiencing increases of between 3-4°C by the mid-2000s. With this expected climate change scenario, high temperature can severely affect crop productivity through different pathways. Individual processes have been well studied, but mostly in isolation, prompting a need for a holistic approach to quantify different processes at one time and in actual environments. Furthermore, when projecting future high-temperature impact on crop productivity under climate change, instead of lumping all effects into a single regression coefficient as done in most previous statistical studies, we should explicitly simulate different processes in a process-based manner which makes designing detailed climate change adaptation strategies for agriculture systems achievable.

Corn and soybean are the two most widely planted crops in the U.S., with planted acreage of 92.9 and 74.8 million acres in 2012. The U.S. produces about 40% of the global corn and soybean crops. Understanding the mechanisms of corn/soybean response to abiotic stresses and providing possible breeding priority for improved varieties has critical significance for the U.S. agricultural economy and also for individual farmers.

METHODS & RESULTS

In this study, we aim to quantify various processes of high-temperature impacts on corn/soybean yield at various scales. Specifically, we consider the following major processes: (1) direct temperature effects on photosynthesis and respiration; (2) sped-up growth rate and the shortening of growing season; (3) heat stress during reproductive stage (flowering and grain-filling); (4) high-temperature-induced increase of atmospheric water demands. By combing the strength of the Community Land Model (CLM) in modeling surface hydrology and photosynthesis and that of the Agricultural Production Systems sIMulator (APSIM) in modeling crop phenology, the newly developed CLM-APSIM crop modeling framework enables us diagnose the effect of high temperature stress through different processes on various crop phonology stages.

First, ground measurements obtained from the advanced SoyFACE facility at the University of Illinois at Urbana-Champaign are used to validate, tune, and improve the CLM-APSIM modeling framework at the site level. Second, parameter screening and sensitivity analysis (SA) experiments are conducted at the local site scale. We use the Sobol sensitivity analysis, which is a global SA method that decomposes the variance of the model output into contributions from each parameter. It's interactions with other parameters through tens of thousands of ensemble simulations of the CLM-APSIM model following the improved Monte Carlo scheme. Third, we calibrate the CLM-APSIM model at grid scale across U.S. through model-data fusion or data assimilation techniques. Novel satellite data, such as fluorescence from OCO-2, soil moisture from SMAP, backscattering information from RapidSCAT, will be used to constrain the CLM-APSIM model at grid scale. We finally use the CLM-APSIM modeling framework to project crop yield for the whole U.S. Corn Belt under different baseline climate scenarios (efficient versus business-as-usual ones).

WHY BLUE WATERS

The computational demand of our satellite data interpretation and modeling effort is huge (~5PB storage for satellite/model inputs and outputs) and the Blue Waters facility offers us the best solution for large processing element demands, high-frequency I/O, and output post-processing and visualization. Interactions with the system staff and those in NCSA lead to a strong technical support for our project.

NEXT GENERATION WORK

As described above, the ultimate goal of this project is to improve the forecasting skill of global crop yield under climate change, though we just focused on U.S. Corn Belt at present. We plan to extend our modeling work to the whole globe on a possible next-generation Track-1 system in the 2019-2020 timeframe. The global scale research is essential as climate change-induced temperature increasing has been diagnosed everywhere across the global terrestrial areas. Understanding the high-temperature impacts on crop yield at the global scale will be helpful to the world's food security and agriculture economy.

PHYSICS & ENGINEERING

MATERIALS

QUANTUM

FLUIDS

NANOTECHNOLOGY

NUCLEAR PHYSICS

BIOMECHANICS

PLASMA PHYSICS

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

Allocation: Illinois/50.0 Knh
PI: Jean Paul Allain[1]
Collaborators: Michael Lively[1], Brandon J. Holybee[1], Michael Toriyama[1], Scott Norris[2], and Stefan Facsko[3]

[1]University of Illinois at Urbana-Champaign
[2]Southern Methodist University
[3]HZDR (Helmholtz-Zentrum Dresden-Rossendorf), Institute of Ion Beam Physics and Materials Research

FIGURE 1: Upper Left: Side-view cutaway snapshot of the initial (i.e. early-stage irradiated by 500 eV Kr ions) surface used for the simulations. Lower Left: Side-view cutaway snapshot of the "control" surface after a ~4 ns simulated relaxation period. Right: Time-evolution of the Ga-Ga and Sb-Sb pair distribution functions, showing that Ga remains in an amorphous state while Sb self-arranges into nanocrystalline lattices.

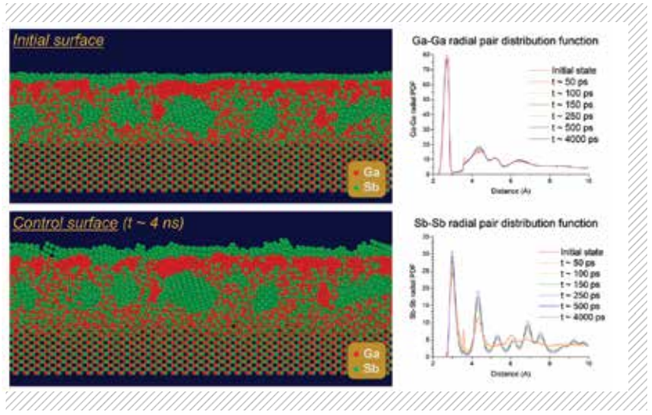

EXECUTIVE SUMMARY

Directed irradiation synthesis of ordered nanodots and other nanostructures on III-V semiconductors has tremendous potential as a method for rational design of advanced nanomaterials for various industrial applications. Massive-scale atomistic simulations used while developing this technique are necessary to address knowledge gaps in existing theories of nanopattern formation and growth. Molecular dynamics studies have been carried out on Blue Waters to simulate the ion bombardment up to 8.4×10^{13} cm^{-2} fluence of a 100×100 nm^2 gallium antimonide (GaSb) surface of altered compositional state resembling experimental observations at the onset of pattern formation. The dominant observation is the formation of Sb nanocrystals in Sb-enriched GaSb regions leading to disparate mobilities of Ga versus Sb. The Sb crystallization is not directly ion-driven during the simulated time but is an intrinsic material response to the existing ion-induced compositional instability. The direct ion-induced mechanisms driving this instability remain the subject of additional, detailed study.

INTRODUCTION

The use of plasma- and ion-based synthesis has gained attention as scalable techniques to create a variety of surfaces using a single process step have been developed. The interfaces between plasma and material become an open thermodynamic system driven from equilibrium by physical mechanisms including high-energy kinetic disordering, compositional phase dynamics, and the emergence of metastable material states. The instabilities that arise lead to the evolution of well-ordered nanostructures, the compositional and morphological characteristics of which dictate the material properties.

Of particular interest is the formation of hexagonally-ordered quantum dots at the surfaces of III-V semiconductors such as GaSb exposed to energetic ion beams [1]. However, little progress has been made in modeling the evolution of III-V surfaces under ion irradiation, and thus the driving mechanisms behind pattern formation are poorly understood—in particular how compositional gradients near the surface of multi-component materials drive nanopatterning. Consequently, the dependence of the nanopattern on experimental parameters such as the incident ion energy or mass is understood in a broad, qualitative sense, making it difficult to design surfaces in regimes which have not been explored. Lacking is a comprehensive, atomistic computational modeling paradigm which addresses these questions of structural and compositional dynamics, coupled with directed in-operando experimental validation tools and techniques.

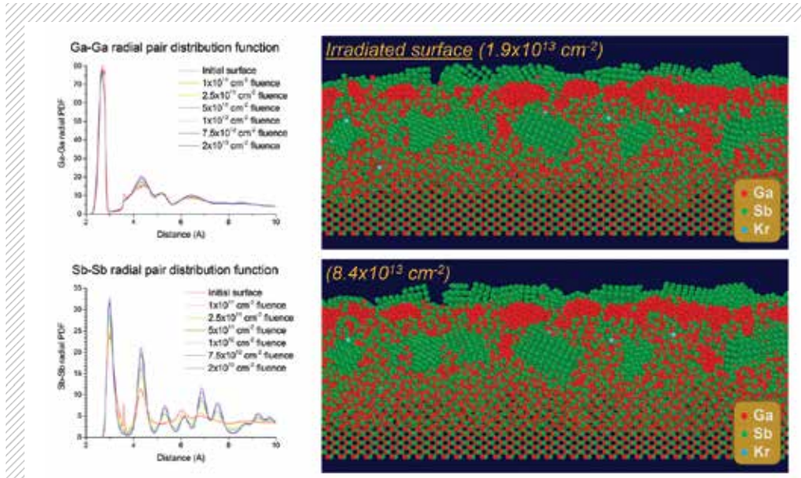

FIGURE 2: Left: Time-evolution of the Ga-Ga and Sb-Sb pair distribution functions for the irradiated surface, most importantly indicating that the self-ordering of Sb atoms is not directly influenced by the incident ions, by comparison to the same data in Figure 1. Right: Side-view cutaway snapshots of the surface after simulated ion bombardment to fluences of 1.9×10^{13} cm^{-2} and 8.4×10^{13} cm^{-2}.

Another challenge is identifying physical scales where experiments and models can be compared. Ion-induced mechanisms are thought to be ultrafast events that could be coupled to atomistic modeling at simulated time scales of the order of a few nanoseconds. Additionally, the ability to fabricate III-V nanodot arrays rapidly and reliably is critical for industrial applications. The understanding gained from this work lays a foundation for the development of advanced nanomaterials and surfaces to be used for devices in fields as diverse as semiconductors, biomaterials and biosensors, or next-generation energy technology.

METHODS & RESULTS

Simulations are started from an initial surface created with a compositional profile mimicking the early-stage ion-irradiated surface at the onset of patterning, which has been measured experimentally (surface representative of 500 eV Kr^+ irradiation to a fluence of nearly 10^{16} cm^{-2}). This technique allows detailed study of the mechanisms active during the actual pattern formation process, elucidating the influence of the surface composition on the ion-induced momentum and energy transfer and the presence of compositionally-driven surface responses which are not directly ion-induced. This initial surface is shown in the upper left of Figure 1.

In addition to the ion-bombardment simulations, a shorter "control" study was performed to identify compositionally-driven mechanisms which were not influenced by the ion impacts. For this study, the surface was allowed to run without ion impacts for ~4 ns. The results are shown in Figure 1. The Sb clusters in Sb-enriched regions self-arrange into nanocrystalline lattices, while the Ga clusters in Ga-enriched regions remain amorphous. While initially surprising, this result could be expected based on previous experimental observations [2,3]. This phenomenon suggests that the amorphous Ga atoms retain a much higher mobility than the Sb atoms which are strongly bound into rigid, crystalline structures. Indeed, this conjecture also agrees with experimental observations that Ga has orders-of-magnitude higher diffusivity than Sb in GaSb [4,5].

Results from the ion-bombardment simulations are shown in Figure 2. Qualitatively, the surface looks very similar to the "control" surface, and the coordination analysis indicates no change in the dominant Sb-crystallization phenomenon. This observation confirms that the formation of Sb nanocrystals in the sub-surface is not directly ion-driven during the simulated time, but is an intrinsic response of GaSb to the already-present compositional instability. The initial ion impacts induce mixing and redistribution of the near-surface

atoms, leading to a metastable compositional depth profile which includes the Sb-enriched region. The atoms in this region phase-separate laterally into pure Sb clusters and 50-50 GaSb, the former of which crystallize as observed here, forming the "seeds" which will drive the formation of nanopatterns on the surface.

Further analysis is needed to investigate the influence of compositional variations on the direct ion-impact energy and momentum transfer to the surface atoms. Such information is critical for determining the mechanisms and parameters relevant to generalized models of surface nanopatterning [6,7], allowing the prediction of pattern parameters and surface properties as functions of experimental parameters such as ion species, incident energy, or choice of surface component materials.

WHY BLUE WATERS

The work completed involved conducting MD simulations of 3 million atoms for nearly 50 million time steps. On a conventional supercomputing cluster this would take multiple years to complete. Large output data files were written at a fairly high frequency, necessitating a high-performance file system to minimize the impact on the overall simulation time. Production runs on Blue Waters were able to simulate up to 10 million time steps on 128 nodes in just 48 hours, allowing the majority of the production runs to be completed in less than two months actual time.

SIMULATING STRONGLY CORRELATED SYSTEMS: FROM FRUSTRATED MAGNETS TO MANY-BODY LOCALIZATION

Allocation: Blue Waters Professor/250 Knh
PI: Bryan Clark[1]
Collaborators: Hitesh Changlani[1], Eduardo Fradkin[1], Krishna Kumar[2], David Luitz[1], and Xiongjie Yu[1]

[1]University of Illinois at Urbana-Champaign
[2]Tufts University

EXECUTIVE SUMMARY

The key problem in condensed matter physics is connecting microscopic details to emergent phenomena. My group has used Blue Waters to develop this connection. We have focused on three main emergent phenomena: many-body localization, spin liquids, and superconducting systems. We describe the former two in this report.

Frustrated magnets arise from geometrically frustrated lattices, such as the Kagome lattice. We find strong numerical evidence for a chiral spin liquid in a natural spin system on the Kagome lattice under a magnetic field. Many-body localization is a phenomenon where statistical mechanics breaks down and quantum mechanics survives at infinite temperature. Using a **newly developed** method we show the saturation of entanglement for many-body localized states.

INTRODUCTION

Frustrated Magnets

In many cases, the electrons in a material are localized in space. The important physics can be captured by viewing the electrons simply as spins which decorate a geometric lattice and want to anti-align. On a square lattice there is a simple pattern which satisfies all electrons; every other electron points up (respectively down). In contrast, when the electrons are on a triangular lattice, this is no longer the case. Instead, while two electrons can be anti-aligned, the third electron must be aligned with one of the other two, causing geometric frustration. One of the most frustrated lattices is the Kagome lattice, which consists of corner sharing triangles. This inherent frustration makes it possible for exotic phases to arise. One such phase is the chiral spin liquid.

Many-Body Localization

Many-body localization is a phenomenon where statistical mechanics [1,2] breaks down and quantum mechanics manifests itself at finite (even infinite) temperature. This is particularly surprising as quantum phenomena are typically only prevalent at zero temperature. From a technological point of view, this phase is important because it may allow for the construction of a quantum computer and quantum memory that is robust to thermal noise.

METHODS & RESULTS

Frustrated Magnetism

Our approach to understanding electrons on the Kagome lattice was to perform an exact quantum simulation (using exact diagonalization) for the largest system size possible given our computational constraints. We simulate a system of spins with the nearest neighbor XY interactions under a magnetic field. Our computational ability allows us to simulate 48 spins requiring finding the lowest eigenvector of hundreds of matrices each of which are of size 377 million by 377 million. In doing so, we find ample numerical evidence that this system forms a chiral spin liquid.

The strongest evidence for the presence of a chiral spin liquid is as follows: A normal state of matter is insensitive to the topology on which the system lives. A magnet on a sphere, a donut, or a cylinder looks roughly the same. Topological states of matter such as spin liquids know the topology of their system and can "feel" the effect of distant boundaries. In our simulations, we slowly twist the boundaries of the system so that the spins interact differently across them. Once we make a full twist, the geometry of the system returns to where it started and amazingly enough, the system does not. Instead, the system picks up an additional quantized phase. The presence of this phase is a key signal of the topological phase of matter and shows that our system is a chiral spin liquid.

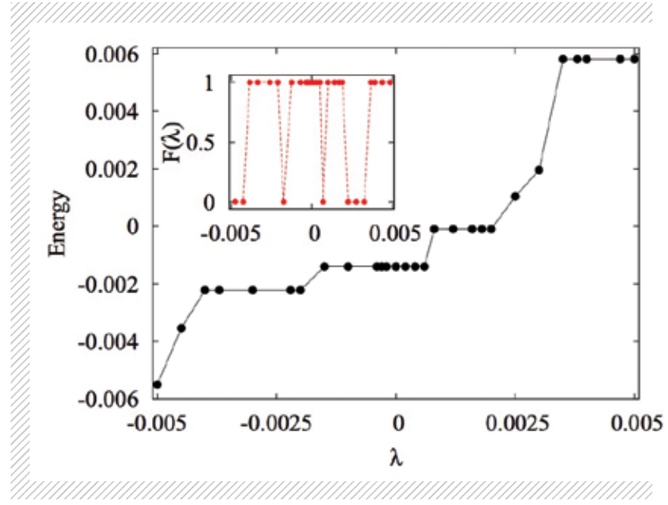

FIGURE 1: Example of sweeping the target energy lambda of a SIMPS run. Each dot corresponds to the eigenstate energy identified. The inset shows the overlap of nearest neighbor eigenstates. The fact that they are always one or zero validates that each state is individually an eigenstate.

Spin liquids are important for two reasons. First, they are the prototype for a phase of matter which stretches beyond the theoretical boundaries of the typical way physicists describe phases—the Landau paradigm. Secondly, they have the potential to be essential for the implementation of quantum computers that are robust to errors.

Many-Body Localization

Unfortunately, in the field of many-body localization, exact methods are limited to 22 sites. More sites are needed to better understand the phase, so we have developed a **novel** density matrix renormalization group (DMRG) algorithm (called the Shift and Invert MPS, or SIMPS). We applied this new algorithm to the disordered Heisenberg model scaling to an order of magnitude more sites than previously possible. With this capability we were able to measure the probability of entanglement, show the saturation of entanglement, and generate thousands of local excitations.

WHY BLUE WATERS

Frustrated Magnetism

Blue Waters was essential to "numerically prove" the existence of our chiral spin-liquid. Because of symmetry constraints, 48 spins was the smallest number of spins for which one could reasonably

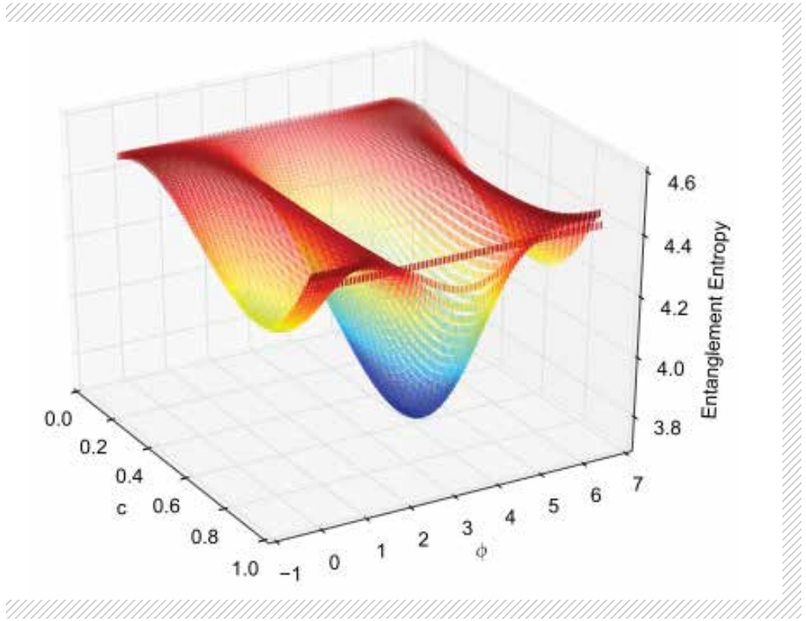

FIGURE 2: Renyi entanglement entropy for the 42b cluster as a function of c, φ obtained from the reduced density matrix along two topologically non-trivial cuts. The linear combinations corresponding to local minima of the entanglement suggesting the existence of a topological phase.

compute all properties of the topological phase, but computing the entire set of twists required over 100,000 node hours.

Many-Body Localization

Blue Waters was essential to this project because of the multiple layers of computation needed. Many-body localized phenomena involve disordered systems. Therefore, any computation requires an average of over thousands of disordered realization. For each of these disordered realizations, a sweep over multiple target energies lambda is required. Finally, each (lambda, disordered configuration) point requires tens of applications of the SIMPS algorithms to converge to the excitation.

NEXT GENERATION WORK

Frustrated Magnetism

We discovered a chiral spin liquid in our system at a single phase point. We would like to discover whole regions that have spin liquid behavior because that will give us the best chance of finding realizations in nature of these exotic phenomena. Unfortunately, each point in this region requires an entire calculation comparable to the one we have done for the chiral spin liquid. It is only with Track-1 systems that this will become possible.

Many-Body Localization

Our SIMPS methodology currently allows us to access states within the many-body localized region, but fails at the critical point, and understanding the nature of this critical point is key to understanding many-body localization. We are developing new tensor network methodologies which may be able to access this point, but the computational complexity of running them will require the next-generation machines to achieve the system sizes necessary.

PUBLICATIONS AND DATA SETS

Chen, X., et al., Many-body localization transition in Rokhsar-Kivelson-type wave functions. *Phys. Rev. B*, 92:21 (2015), p. 214204.

Yu, X., D. Pekker, and B.K. Clark, Finding matrix product state representations of highly-excited eigenstates of many-body localized Hamiltonians. *arXiv* preprint arXiv:1509.01244 (2015).

Yu, X., D.J. Luitz, and B.K. Clark, Bimodal entanglement entropy distribution in the many-body localization transition, *arXiv* preprint arXiv:1606.01260 (2016).

TRANSIENT TWO-PHASE FLOW AND ELECTRO-MAGNETIC FIELD EFFECT IN STEEL CONTINUOUS CASTING

Allocation: Illinois/200 Knh
PI: Brian G. Thomas[1,2]
Co-PIs: Seong-Mook Cho[1], Kai Jin[1], S.P. Vanka[1], Hyunjin Yang[1], Matthew Zappulla[1], Xiaolu Yan[1], Ahmed Taha[1], and Seid Koric[1]

[1]University of Illinois at Urbana-Champaign
[2]Colorado School of Mines

EXECUTIVE SUMMARY

This project aims to advance current state-of-the-art, computationally-intensive models of multiphase phenomena including turbulent fluid flow, particle transport, and MagnetoHydroDynamics (MHD) in the continuous casting of steel. Also, the project applies these current models to gain practical insights into the transient flow phenomena related to defect formations and to improve this important commercial process. In this study, a transient, two-phase Large-Eddy-Simulation model of molten steel-argon gas flow was applied to investigate anisotropic (directionally dependent) turbulent flow in the caster, with and without a double-ruler electro-magnetic field. The model calculations have been validated with plant measurements and applied to understand the mechanism of the flow variations, which is important to surface defect formation in the final product. Also, the effects of the magnetic field on the flow stability in the caster has been quantified with different process parameters.

INTRODUCTION

Continuous casting is used to manufacture more than 95% of steel in the world [1]. Molten steel flows from a tundish container, through a slide gate control valve, and down a vertically bifurcated nozzle into the mold (Fig. 1). Once in the mold, molten steel solidifies against the water-cooled, copper mold walls to form a solid shell. Transient fluid-flow phenomena in the mold are very important to quality and defects in the final product. Abnormal surface flow aggravates level fluctuations, shear instability of the molten slag/steel interface, and vortex formation near the submerged entry nozzle (SEN), which leads to mold slag entrainment. These detrimental flow phenomena in the mold become more complex with argon gas, which is injected into the nozzle to prevent clogging. Furthermore, the argon gas bubbles may become entrapped into the solidifying steel shell, forming other defects. To stabilize and optimize the transient fluid flow in the mold, electro-magnetic fields are often applied, especially at high casting speed [2-4].

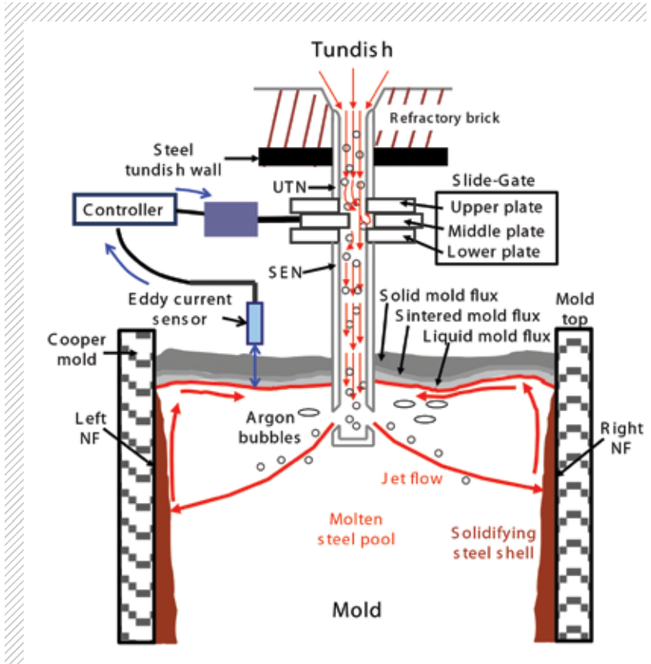

FIGURE 1: Schematic of the continuous steel slab-casting process.

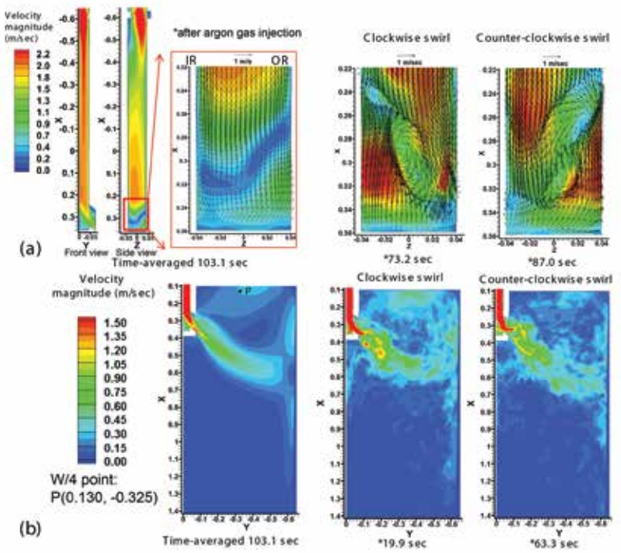

FIGURE 2: Time-averaged and instantaneous flow patterns in (a) nozzle and (b) mold.

In the present work, Large Eddy Simulations (LES) of multiphase flow are performed to quantify the transient behavior of molten steel-argon gas flow in the nozzle and mold during nominally-steady continuous casting of steel slabs. The model has been validated with plant measurements and applied to investigate the mechanism of multiphase flow stability, argon-bubble entrapment into the steel shell and effects of double-ruler Electro-Magnetic Braking (EMBr) on the transient flow variations for different process conditions.

METHODS & RESULTS

To calculate transient flow of molten steel and argon gas in the nozzle and mold, a three-dimensional finite-volume LES model coupled with Discrete Phase Model (DPM) for particle tracking has been developed to consider the interaction between the molten steel flow and argon bubble motion. This model was implemented into a special multi-license version of ANSYS Fluent software on Blue Waters.

Time-averaged and instantaneous flow patterns in the nozzle and mold are shown in Figure 2 [3]. The asymmetrically-positioned slide gate middle plate, with its open area near the Outside Radius (OR) side, produces swirl flows in the nozzle well bottom, alternating chaotically between clockwise and counter-clockwise rotation. Instantaneous mold flow patterns show up-and-down wobbling of the jet in the mold, which produces different impingement points on the narrow face. When the nozzle flow from the slide gate flips sides down the nozzle (Fig. 2a), it produces counter-clockwise swirl exiting the ports, impinges downwards on the narrow faces, and results in slow surface flow in the mold. When the nozzle flow goes straight down to produce clockwise swirl exiting the port, the accompanying faster jet bends upwards, impinges the narrow face horizontally, and produces 50% faster surface flow (Fig. 3a).

Time-averaged and instantaneous profiles of surface velocity magnitude predicted by the LES model are compared with nail board measurements in Figure 3b. Each line shows surface velocity magnitude profiles across the mold width at the center-plane 10mm below the interface between the molten steel and liquid mold flux layers. Symbols with error bars present time averages and standard deviations of 10 nail-dipping tests at each measurement location [4]. The model predictions of the average velocity profile, and its time and spatial variations, all agree with the measurements. This confirms that the LES-DPM model on Blue Waters is an accurate tool to predict complex mold-flow phenomena including multiphase effects.

Related work coupling turbulent flow with DPM to model bubble transport and capture, based on local force balances on argon-gas bubbles at the solidification front, has revealed that biased flow across the mold surface can aggravate bubble capture [5].

The effect of double-ruler EMBr was added to the LES-DPM model by implementing the magnetic induction MHD equations. The EMBr in the current work has one peak centered across the mold above the port and the other centered below the nozzle port. With EMBr, swirl direction flipping occurs more frequently, flow asymmetries are reduced, turbulence from jet wobbling is reduced, and thus, surface velocity fluctuations are smaller [3]. Furthermore, an efficient multi-GPU based code, CUFLOW, was used to perform LES to investigate the effects of several SEN depths and double-ruler EMBr strengths [6], to suggest the best choice of operating conditions to control surface flow.

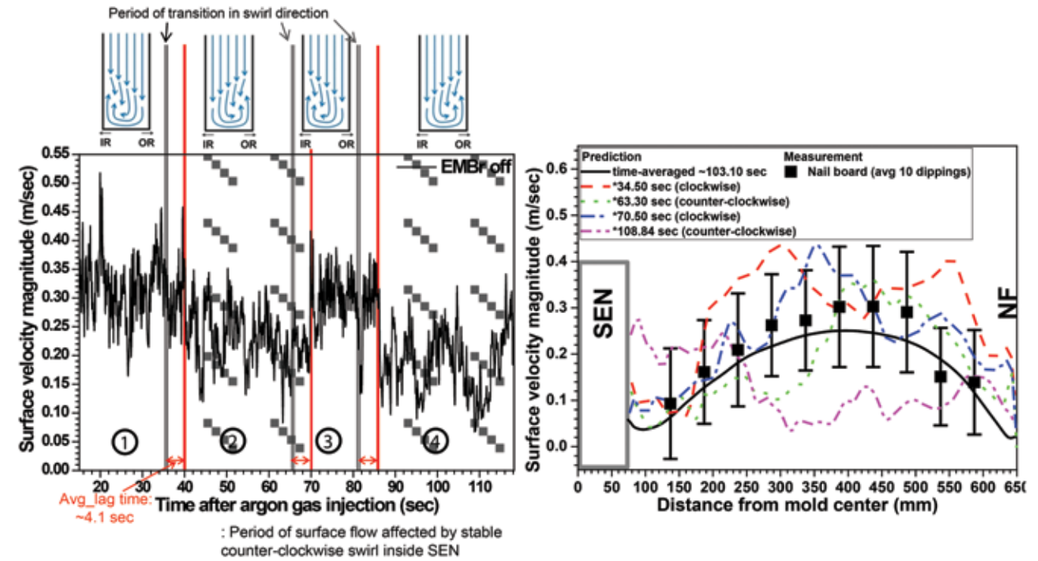

WHY BLUE WATERS

Multiphase flow simulations in the continuous caster are enabled by Blue Waters, including the large domain volume (e.g. 0.2 m X 1.6m X 4.6m) and fine mesh (e.g. ~22 million cells) needed to properly capture and resolve the turbulent flow in the casting process. Specifically, the Fluent-14.5 HPC software on Blue Waters ran 3,357 times faster than on an ordinary work station PC (Dell T7600: Intel® Xeon® CPU E5-2603 @ 1.80GHz, RAM 40.0 GB, using 6 cores) with 1,120 floating cores (70 XE nodes).

NEXT GENERATION WORK

A validated two-dimensional thermal-fluid model of solidification phenomena at the slag/molten steel interface near the meniscus region [7], will be extended into the full 3-dimensional LES mold discussed here, to investigate how anisotropic multiphase flow variations affect initial solidification, with and without magnetic fields.

PUBLICATIONS AND DATA SETS

Jin, K., B.G. Thomas, and X. Ruan, Modeling and Measurements of Multiphase Flow and Bubble Entrapment in Steel Continuous Casting., *Metall. Mater. Trans. B*, 47B:1 (2016), pp. 548-565. DOI: 10.1007/s11663-015-0525-5

Cho, S-M., B G. Thomas, and S-H. Kim, Transient two-phase flow in slide-gate nozzle and mold of continuous steel slab casting with and without double-ruler electro-magnetic braking. *Metall. Mater. Trans. B*, (2016) pp. 1-19.

Jin, K., S. P. Vanka, B.G. Thomas, and X. Ruan, Large eddy simulations of the effects of double-ruler electromagnetic braking and nozzle submergence depth on molten steel flow in a commercial continuous casting mold. *TMS Annual Meeting, CFD Modeling and Simulation in Materials Processing Symposium*, Nashville, TN, Mar. 14-18, 2016, TMS, Warrendale, PA, 2016, pp. 159-166.

FIGURE 3: (a) Instantaneous surface velocity magnitude at W/4 point and (b) comparison of predicted surface velocity profiles with plant nail-board measurements

THE LIQUID-LIQUID TRANSITION IN DENSE HYDROGEN

Allocation: Blue Waters Professor/240 Knh
PI: David Ceperley[1]
Collaborators: Carlo Pierleoni[2], Giovanni Rillo[2], Raymond Clay[1], Miguel Morales[3], and Marcus Holzmann[4].

[1]University of Illinois at Urbana-Champaign
[2]University of L'Aquila
[3]Lawrence Livermore National Laboratory
[4]CNRS Grenoble France

EXECUTIVE SUMMARY

The phase diagram of high-pressure hydrogen is of great interest for fundamental research, planetary physics, and energy applications. A first-order phase transition in the fluid phase between a molecular insulating fluid and a monoatomic metallic fluid has been predicted. The existence and precise location of the transition line are relevant for planetary models. Recent experiments reported contrasting results about the location of the transition. Theoretical results based on density functional theory are also very scattered. We performed highly accurate coupled electron-ion Monte Carlo calculations of this transition, finding results that lie between the two experimental predictions, close to that measured in diamond anvil cell experiments but at 25–30 GPa higher pressure. The transition exhibited a discontinuity in the specific volume, a sudden dissociation of the molecules, a jump in electrical conductivity, and loss of electron localization, all indications of a weak first-order transition.

INTRODUCTION

The properties of hydrogen under extreme conditions of pressure and temperature are of great interest for fundamental research, planetary physics, and energy applications [1]. Hydrogen accounts for much of the visible mass in the universe. The properties of hydrogen and helium are needed for understanding the giant planets (Jupiter and Saturn), but experiments under the relevant conditions are challenging. Even though hydrogen is the first element in the periodic table, calculating its properties is not simple since both the electronic and protonic correlations are quantum and correlated.

There is a long-standing question [1] regarding how hydrogen makes a transition from a molecular insulating state to an atomic metallic state as pressure and temperature are increased. A first-order phase transition in the fluid phase between a molecular insulating fluid and a monoatomic metallic fluid has been predicted [2]. Recent studies [3-4] reported contrasting results about the location of the transition different by a factor of two (in pressure). Theoretical results based on density functional theory are also very scattered and hence not predictive [5]. These findings motivated us to repeat our earlier calculations on Blue Waters, to better control the convergence and utilize recent improvements in methodology.

METHODS & RESULTS

Over the past decade, we have developed new Quantum Monte Carlo simulation methods to treat quantum systems at low temperature. The quantum Monte Carlo method we use (Coupled Electron Ion Monte Carlo) works with the full interaction between the electrons and protons and treats both particles fully quantum mechanically. In contrast to density functional calculations, all effects of electronic correlation are explicitly included. This concept is particularly important in hydrogen, because of possible self-interaction effects, difficulty in treating the hydrogen bond breaking and the large van der Waals interactions. We model hydrogen with about 100 electrons and protons in a periodic cell. Special methods are used to extrapolate accurately to the thermodynamic limit. With our method, we simulated hydrogen for temperatures in the range of 200K to 5000K, and at relevant pressures, 100GPa to 500GPa.

We find pressures for the transition that lie between the two experimental predictions, close to that measured in diamond anvil cell experiments, but at 25–30 GPa higher. This observation is shown for both hydrogen and deuterium (Fig. 1). The transition along a line of constant temperature exhibited a discontinuity in the specific volume, a sudden dissociation of the molecules, a jump in electrical conductivity, and loss of electron localization. For temperatures below 2000K, we observe a first order transition between an insulating molecular liquid and a denser metallic atomic liquid. Our predicted transition pressures are intermediate between the two experimental observations [3,4]. Future work will be to perform further simulations and analysis to understand the divergent results of the experiments and the unusual properties of the molecular and atomic liquid. New experiments using the National Ignition Facility (NIF) at Lawrence Livermore National Laboratory have also been recently performed though the results are proprietary.

Our calculations are needed to validate our computational method and to resolve the different experimental measurements. It is essential for progress in the high-pressure community to determine the difference between the experiments and computation. After validation, the method can be used with more confidence in modeling the wide variety of astrophysical objects observed, composed largely of hydrogen and helium under extreme conditions.

WHY BLUE WATERS

Without access to Blue Waters we would not have been able to perform this study because of the computational demands of the simultaneous treatment of quantum electrons and protons.

NEXT GENERATION WORK

Our method is highly amenable to next generation systems. We would be able to perform more such calculations and systems containing elements heavier than hydrogen and helium, essential not only in astrophysics but for materials as well. We will also be able to compute other properties relevant for modeling planetary interiors.

PUBLICATIONS AND DATA SETS

Pierleoni, C., et al., Liquid-liquid phase transition in hydrogen by coupled electron-ion Monte Carlo simulations, *Proc. Nat. Acad. Science (US)* 113 (2016), p. 4953-4957.

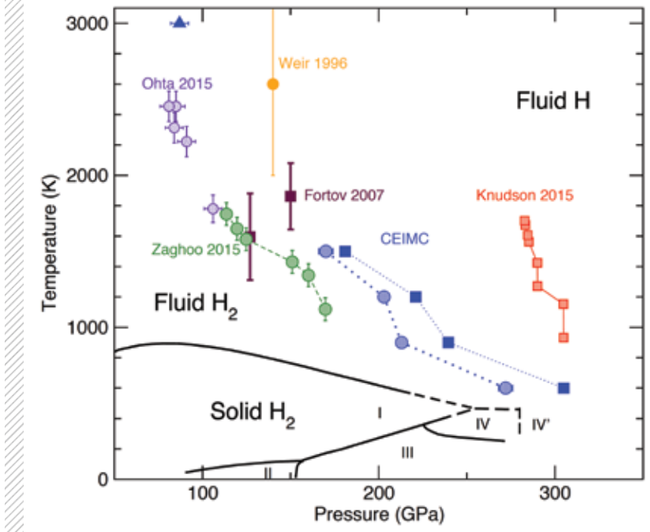

FIGURE 1: The predicted phase diagram (temperature versus pressure) of dense hydrogen. The blue circles (hydrogen) and squares (deuterium) show the results from the Blue Waters calculation [5]. The red squares [3] and green circles [4] show the different experimental measurements.

REDUCING JET AIRCRAFT NOISE BY HARNESSING THE HETEROGENEOUS XK NODES

Allocation: Blue Waters Professor/120 Knh
PI: Daniel J. Bodony[1]
Collaborators: Simon Garcia De Gonzalo[1], Wen-mei Hwu[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Reducing the noise emitted by commercial and military aircraft to alleviate their adverse environmental and health impacts requires computational resources that exceed current petascale computers. Future exascale computers are expected to have architectures where most of the floating-point computing capacity is handled by accelerators; however, no single architecture is expected to be dominant. We focus on the programming models and runtime systems required to support computational science on future exascale computers.

INTRODUCTION

Gas turbine-powered commercial and military aircraft create intense noise at take-off that negatively impacts near-airport communities and creates a health hazard for carrier-borne personnel. Predicting and, ultimately, reducing jet noise requires immense computing resources because of the large range of temporal and spatial scales present in the hot, turbulent exhaust gases. Future exascale computers are expected to open new opportunities in jet noise reduction, but their architectures will differ from current leadership-class computers to improve power-per-FLOP performance. In particular, the bulk of the floating-point operations will be performed on accelerators that may be highly simplified computing elements similar to modern-day GPUs connected to task-managing CPUs, host-less accelerators similar to Intel Xeon Phis, or single-chip heterogeneous system architectures (HSA) in which the GPU and CPU sit on the same die. Without architecture clarity, computational science codes must adapt and be sufficiently flexible to run on all three kinds of systems. We use the XK nodes on Blue Waters as a stand-in example of one type of a future exascale computer to test a new programming model and runtime system capable of using all three types of future computers. We focus on the specific needs of typical computational fluid dynamics algorithms with low FLOP-to-load ratios.

METHODS & RESULTS

For low FLOP-to-load ratio algorithms the biggest challenge to achieving performance is the bandwidth between the memory system and the computing elements. The three classes of proposed exascale architectures have very different memory-compute layouts, so our work has focused on isolating the

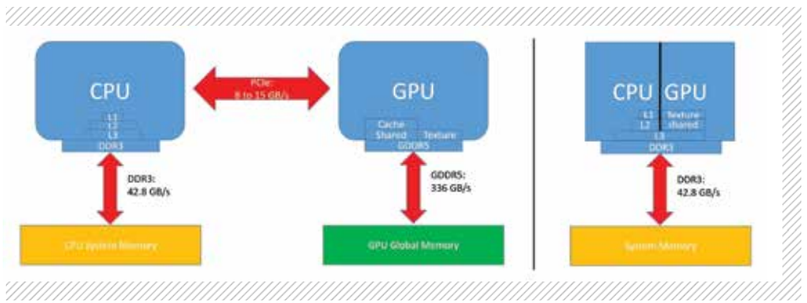

FIGURE 1: Different memory architectures and bandwidths for legacy CPU-accelerator systems (left) and cache-coherent HSA systems (right).

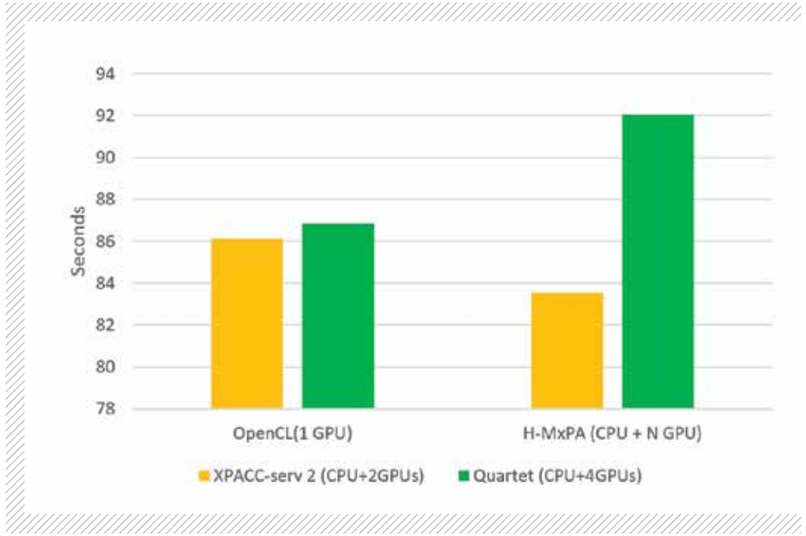

FIGURE 2: Performance of RK4 update step on two legacy CPU-GPU systems (smaller is better). XPACC-serv-2 has one Intel E5-2680V2 connected to two NVIDIA K40m GPUs, while Quartet has one Intel i7-3820 connected to four NVIDIA K20c GPUs.

expression of the algorithm from the underlying hardware description and permitting on-the-fly runtime optimization (see Fig. 1). The *H-MxPA* [1] cross-compilation and runtime tool uses C++AMP [2] to express the algorithm in a platform-independent manner, a compilation back-end supported by LLVM [3], and a dynamic runtime that partitions work between the available CPU (if present) and computing element(s) (if present) by closely monitoring performance. In addition, on modern CPU-GPU nodes like Blue Waters' XK7s, memory coherency must also be ensured between the host and device by the runtime.

Applying *H-MxPA* to the update step in the standard fourth order Runge-Kutta (RK4) algorithm, which contains six memory operations (five loads and one store) and six FLOPs (four additions and two multiplies), to legacy CPU-GPU nodes showed the loss of performance due to the memory bandwidth and cache coherency requirements (Fig. 2). For these runs the work distribution was hand-tuned. These performance data show explicitly the penalties incurred on CPU-GPU systems where host and device memories are connected via a PCI bus and speak to the performance improvement possible when the two elements are co-located on a single die, as on an HSA-type machine. It is noteworthy that the pre-exascale machines to be built by IBM-NVIDIA for the Department of Energy use NV-LINK and hardware-supported cache coherency [4].

WHY BLUE WATERS

The XK nodes of Blue Waters represent an early vision of the future exascale computers that will be realized in the 2020-2025 timeframe. The architecture and software stack of Blue Waters' XK nodes provide relevant representative platforms on which to develop software technologies capable of using heterogeneous computers. The size of the Blue Waters XK partition and its connection to the XE partition provide a unique opportunity to build and test computational science algorithms that can handle intra- and inter-node inhomogeneity at leadership-class scale.

NEXT GENERATION WORK

The cross-compilation and runtime development being performed on Blue Waters will become the foundation on which our next-generation computational fluid dynamics (CFD) code will be built. With the flexibility provided by *H-MxPA* we will use the future Track-1 system to demonstrate the efficiency with which computational science can be performed on the pre-exascale computer as well as advance the science of CFD codes used by the engineering and science communities.

TOWARDS LARGE-SCALE KINETIC SIMULATIONS OF THE PLASMA-MATERIAL INTERFACE

Allocation: Illinois/50.0 Knh
PI: Davide Curreli[1]
Co-PIs: Rinat Khaziev[1]
Collaborators: Ryan Mokos[2] and Gregory H. Bauer[2]

[1]University of Illinois at Urbana-Champaign
[2]National Center for Supercomputing Applications/University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

In the edge region of magnetically-confined plasmas, the interaction of plasmas and material surfaces poses significant challenges to the survivability of plasma-facing components, currently limiting the successful development of commercially-viable nuclear fusion reactors. Taming the plasma-material interface is one of the top priorities of fusion science research in order to achieve a demonstration fusion power plant. When exposed to plasma irradiation, plasma-facing materials exhibit evidence of surface morphology modifications and nanostructuring, with detrimental consequences on the thermo-mechanical integrity of the wall.

We have improved the on-node performance and tested the scalability on Blue Waters of the hPIC code (HPC platform for Plasma-Material Interactions and Nanostructuring), a fully-kinetic platform for kinetic analysis of plasma-material interactions, including both a Boltzmann description of the near-wall plasma and the multi-physics response of the material surface.

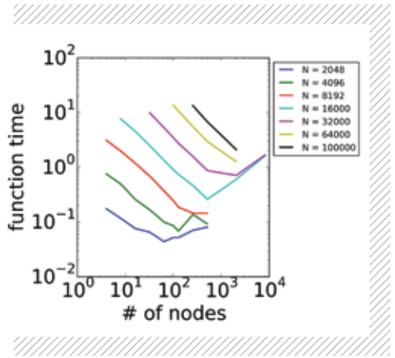

FIGURE 1: Blue Waters scaling of the field module of hPIC with multigrid preconditioner on a square plasma domain of size of NxN. As expected, the scaling deteriorates when small problems are solved at the largest node count.

INTRODUCTION

Understanding the interaction of plasmas with material walls is relevant to the design of the next generation of fusion devices and to industrial processes exploiting plasmas for material manufacturing. In the presence of physical boundaries, plasmas form a near-wall layer, called the plasma sheath, in which the charged particles normally at equilibrium inside the plasma bulk are supersonically accelerated toward the wall. Such acceleration increases the kinetic energy of the particles to levels that might overcome the erosion threshold (sputtering), causing the release of particles from the material wall. The dynamics of this process can be well resolved with electrostatic particle-in-cell (PIC) simulations [1], a well-known numerical method for the solution of the Boltzmann kinetic equation of an ensemble of interacting particles. A typical iteration cycle of a PIC code [2] includes: (1) particle push, during which the particles are moved and their properties (mass, charge) are weighted on the mesh, and (2) the calculation of the electric and magnetic fields from the Maxwell equations and their interpolation at the particles' location. For electrostatic PIC codes it is safe to assume that magnetic field configuration is fixed or quasistatic, and the only equation to solve is the Poisson equation. The major difference between electrostatic PIC (ES-PIC) codes and electromagnetic PIC (EM-PIC) codes is that the latter benefit from the locality of the Maxwell-Faraday and Ampere's equations. Several electromagnetic PIC codes have been developed, showing excellent scalability up to the largest number of nodes on high-performance computing (HPC) systems.

However, existing electrostatic PIC codes are not well supported on HPC systems. Nevertheless, electrostatic PICs are preferable for plasma-material interaction (PMI) problems because they allow less restrictive conditions on the time step and because electromagnetic modes add negligible features in a PMI context. Electrostatic PICs require the solution of an elliptic problem, rather than a set of hyperbolic equations as in electromagnetic PIC codes. In an HPC context, such a difference translates to the need for a more complex communication strategy for ES-PICs. The parallel data partitioning of an ES-PIC code is dictated by the partitioning of the linear solver and by the geometric partitioning of the particle arrays, requiring a different parallelization scheme. The goal of this work was to adapt our electrostatic PIC code hPIC for runs on petascale and exascale supercomputers and to characterize the code's performance at large node numbers.

METHODS & RESULTS

In our implementation of the ES-PIC method we rely on Poisson solvers developed by the FASTMath-SciDAC institute [4], offering a variety of scalable methods for the solution of linear and non-linear algebraic problems. Our code hPIC has been linked to the PETSc library as a back-end. Most of the hPIC development effort has been devoted to improving the on-node performance of each code component and to testing the code on a large number of nodes on Blue Waters. Tests have demonstrated that algebraic multigrid preconditioners with conjugate gradient offer the best solution in terms of parallel performance. Tests were performed on square plasma domains of size NxN (with N being the number of grid nodes along one dimension). For the largest test case evaluated (N=100,000), which corresponds approximately to a plasma domain of $1m^2$ with plasma properties such as those found in the scrape-off layer of fusion devices, the time required to solve the electrostatic problem (Figure 1) was of the order of 2 seconds on 8,192 Blue Waters nodes (262,144 cores). An efficient implementation of the particle parallelization scheme has been adopted, storing the particles along the corresponding field arrays of the Poisson solver, requiring the communication of only those particles no longer associated with the field arrays locally stored at the node level. Tests have shown that this particle communication scheme typically requires less than 1 millisecond on Blue Waters for all cases relevant to our applications, with excellent weak-scaling properties (Figure 2).

WHY BLUE WATERS

Simulations of large plasma volumes at resolutions large enough to resolve the plasma sheath require extensive computational resources, both for node number (more than 1,000 Blue Waters XE nodes) and **in-node memory**. The Blue Waters support staff facilitated fast transition from previous HPC platforms to Blue Waters. Blue Waters offered an excellent environment for code development, both for on-node optimization and testing/profiling, which uniquely contributed to the code optimization at all steps.

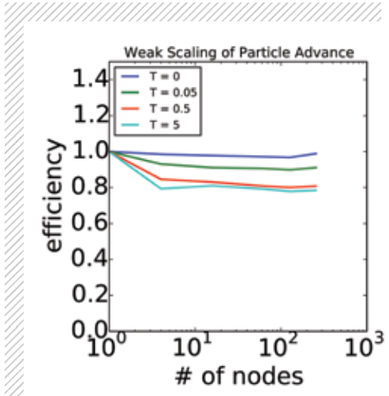

FIGURE 2: Weak scaling test of hPIC on Blue Waters with 32 million particles per node, mesh of 55,000x55,000 on each Blue Waters node (1 million particles per MPI process, 100 particles per cell, 100x100 mesh on each MPI process). The weak scaling holds well with efficiency close to 100% for T=0 (cold plasma), ~91% for T=0.05 (typical ions) and ~80% for T=5 (typical electron plasma).

NEXT GENERATION WORK

The HPC resources available with the next-generation of Track-1 systems will allow us to perform large-scale 2D-3V plasma simulation of the plasma-material interface, potentially simulating a large portion of the divertor private region of a tokamak (a device that uses a magnetic field to confine plasma in a torus), a large portion of the scrape-off layer facing the first wall, or a full-scale industrial plasma device. The code is also ready to explore the first small 3D simulations of the plasma material boundary. Three-dimensional effects are expected to play a vital role at the boundary of stellarator fusion devices, like the Hybrid Illinois Device for Research and Applications (HIDRA) recently acquired by the University of Illinois at Urbana-Champaign.

QMCDB: A LIVING DATABASE TO ACCELERATE WORLDWIDE DEVELOPMENT AND USAGE OF QUANTUM MONTE CARLO METHODS

Allocation: Illinois/50.0 Knh
PI: Elif Ertekin[1]
Co-PI: Lucas Wagner[1] and Ray Plante[2]

[1]University of Illinois at Urbana-Champaign
[2]National Institute of Standards and Testing

EXECUTIVE SUMMARY

Blue Waters has enabled us to carry out automated, high-throughput quantum Monte Carlo calculations of condensed matter systems as part of our development of QMCDB (Quantum Monte Carlo DataBase), a database of materials properties calculated via the highly-accurate quantum Monte Carlo technique. The systems simulated are to be incorporated into our database, which will be made available to the international materials modeling community. This platform for easy, searchable data exchange will accelerate the knowledge base around the use of QMC for materials modeling and enable its evolution from a physics approach to a tool for real engineering materials design. This work would not be possible without Blue Waters, which allows the calculation of a large class of materials ranging from classic to exotic semiconductor materials, photovoltaics, thermoelectrics, and metallic systems; we can take advantage of the near-linear scaling of our methods and code up to several thousand nodes.

INTRODUCTION

Quantum Monte Carlo (QMC) methods are a suite of tools for direct stochastic solution of the many-body interacting Schrödinger equation. Although QMC methods are considered to be one of the highest-accuracy, first-principles materials modeling methods available, and demonstrate a long and distinguished history of benchmark calculations, their usage for materials design and discovery has historically been limited by their large computational cost. With the high-performance computing (HPC) capabilities of Blue Waters, however, it is now possible to extend this method to the realm of high-throughput materials computation and discovery. The goal of our work, therefore, is to develop the first database of materials computations based on quantum Monte Carlo results. We expect that this database will serve as a shared community resource to accelerate the use of this high-accuracy method and advance the community's knowledge of best practices in the application of QMC to real materials design and discovery.

METHODS & RESULTS

The prospect of materials design using HPC is one of the most exciting for future technologies. This has been already achieved to some degree using established methods on model problems. However, the properties of modern materials are rather complex, and current numerical methods can fail to describe them quantitatively. Today, QMC is a state-of-the-art suite of tools for high-accuracy *ab initio* modeling. It shows great promise for high-accuracy materials modeling and is already well established in the physics community (model systems, effective Hamiltonians), but its application to real materials with chemical identity remains fairly young. The reason for this is that historically, quantum Monte Carlo methods were perceived as being too computationally costly for wide-scale adoption. Modern implementations of the algorithms, together with the leadership-class computing facilities offered by Blue Waters open a new opportunity in this field.

Our goal is to accelerate the development of the QMC community's collective knowledge base around the use of this method for predictive modeling of real materials. We have used Blue Waters to carry out high-throughput calculations of the properties of a large class of semiconductors and oxides, the results of which are to be incorporated into a newly developed database. This database will serve as a vehicle to quickly overcome the current expertise hurdle and bring the QMC methodology into the standard computational modeling toolkit.

Our efforts thus far have focused on the calculation of thermodynamic properties (total energies and formation enthalpies) and the band gaps of a wide class of semiconductor materials spanning from conventional (silicon, etc.) to more exotic materials (wide band gap oxides, correlated systems, photovoltaic thin film materials). While today there are several competing materials databases for density functional theory, there is no existing QMC database. Thus, Blue Waters has provided us the opportunity to establish this tool for the worldwide QMC community, **for the first time**.

To date, our work has encompassed:

- The implementation of an automated framework for carrying out QMC calculations of solid materials on Blue Waters. The framework is working and is now being used together with Blue Waters to carry out our automated quantum Monte Carlo calculations of semiconductor materials.
- The establishment of QMCDB: Our MongoDB database "QMCDB" is now active through National Data Service Labs, using cloud capabilities. The database is described in detail on the National Data Service Project Wiki page, at http://wiki.nationaldataservice.org/QMCDBProject. Once we populate the database with our initial set of 20 materials calculations (largest QMC data set to date) with provenance, it will become publicly accessible.

Ultimately, the impact will be to enable quantum Monte Carlo methods to emerge as a standard component of the computational materials modeling toolkit, enabling **unprecedented** high-accuracy simulation of complex materials, correlated systems, high-T_c superconductors, and other historically challenging materials.

WHY BLUE WATERS

Blue Waters is the key to carrying out the comprehensive set of QMC materials calculations that populate our database. QMC methods exhibit near-linear scaling on the entire Blue Waters platform, which has allowed us to calculate properties of an extensive set of materials that would otherwise not be possible. Historically, the wide scale adoption of QMC as a method for materials modeling has been limited by its large computational cost, but Blue Waters allows us a first key opportunity to overcome this barrier.

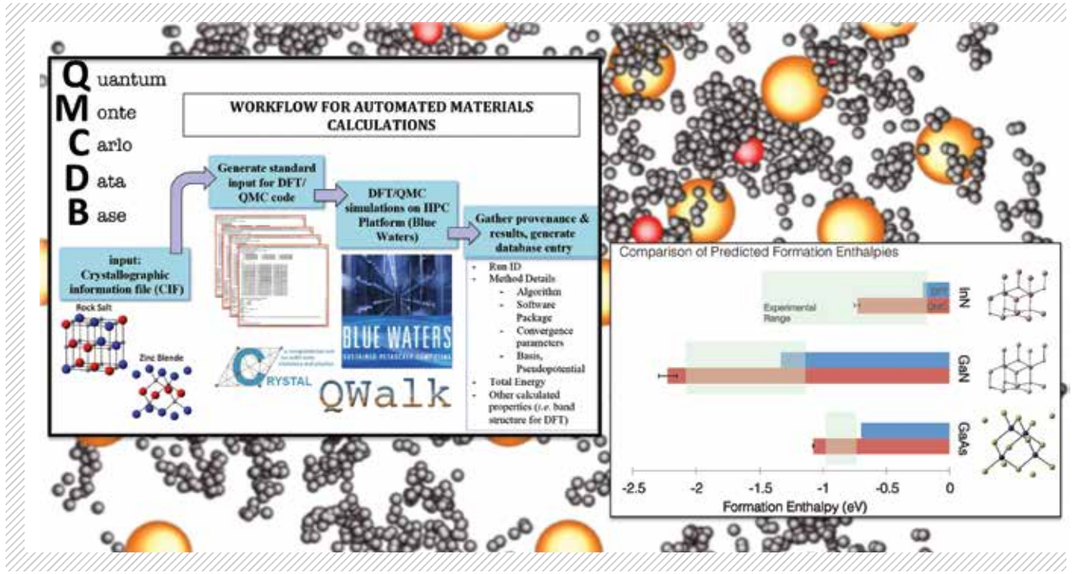

FIGURE 1: A snapshot of a configuration of electrons around ionic nuclei in the magnesium oxide solid. Foreground: workflow for automated quantum Monte Carlo calculations of solids and incorporation into quantum Monte Carlo database, and selected results for calculated formation enthalpies of conventional semiconducting solid materials.

NEXT GENERATION WORK

Next generation Track-1 systems in the coming years will enable us to carry out more systematic analyses on a larger class of materials. Our initial work shows that simulations of larger sized systems are important. Also, more systematic studies will enable us to carry out big data assessments of correlated many-body wave functions, towards the idea of "correlated electron genomics."

PUBLICATIONS AND DATA SETS

Schiller, J.A., L.K. Wagner, and E. Ertekin, Phase stability and properties of manganese oxide polymorphs: Assessment and insights from diffusion Monte Carlo, *Phys. Rev. B*, 92 (2015), 235209 doi: 10.1103/PhysRevB.92.235209

Yu, J., L. K. Wagner, and E. Ertekin, Towards a systematic assessment of errors in diffusion Monte Carlo calculations of semiconductors: case study of zinc selenide and zinc oxide, *J. Chem. Phys.* 143 (2015), 224707, doi:10.1063/1.4937421

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

Allocation: NSF PRAC/3.00 Mnh
PI: J. P. Draayer[1]
CO-PIs: K. D. Launey[1], T. Dytrych[2], J. P. Vary[3], P. Maris[3], U. Catalyurek[4], and M. Sosonkina[5]
Collaborators: R. Baker[1], D. Langr[6], T. Oberhuber[6], E. Saule[7], W. M. Tang[8], and B. Wang[8]

[1]Louisiana State University
[2]Czech Academy of Sciences and Louisiana State University
[3]Iowa State University
[4]The Ohio State University
[5]Old Dominion University
[6]Czech Technical University
[7]University of North Carolina
[8]Princeton University

EXECUTIVE SUMMARY

We use Blue Waters to carry out large-scale modeling of light and medium-mass nuclei, including short-lived nuclei not yet accessible to experiment but critical to understanding astrophysical processes, which are the focus of current and next-generation rare isotope experimental facilities. The scale of computational challenges inherent to modeling such intricate quantum many-body systems makes the utilization of Blue Waters resources essential for addressing long-lasting challenges of importance to nuclear theory and experiment, as well as astrophysics. **Breakthrough** theoretical advances [1] coupled with Blue Waters' **cutting-edge** computational power have opened a new region, the intermediate-mass nuclei from fluorine (F) to calcium (Ca) isotopes, for **first** investigations with *ab initio* methods. This breakthrough fundamentally advances our understanding of nucleosynthesis, as nuclear energy spectra and reaction rates for many short-lived nuclei involved in nucleosynthesis are **not yet accessible by experiment or reliably measured** for the astrophysically relevant energy regime.

INTRODUCTION

One of the quintessential open challenges in contemporary physics is to design a comprehensive many-body theory for modeling and predicting nuclear structure and reactions starting from inter-nucleon forces that are consistent with the underlying theory of Quantum Chromodynamics (QCD). The ultimate goal of *ab initio* theory is to find a solution to this problem, which is a computationally-intensive endeavor due to a dual challenge, namely, the non-perturbative nature of QCD in the low-energy regime and the complexity of many-particle nuclei. Short-lived nuclei, currently **inaccessible to experiment**, are 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 been and will be demonstrating a tremendous impact on advancing our knowledge at the frontiers of astrophysics, neutrino physics, and applied physics.

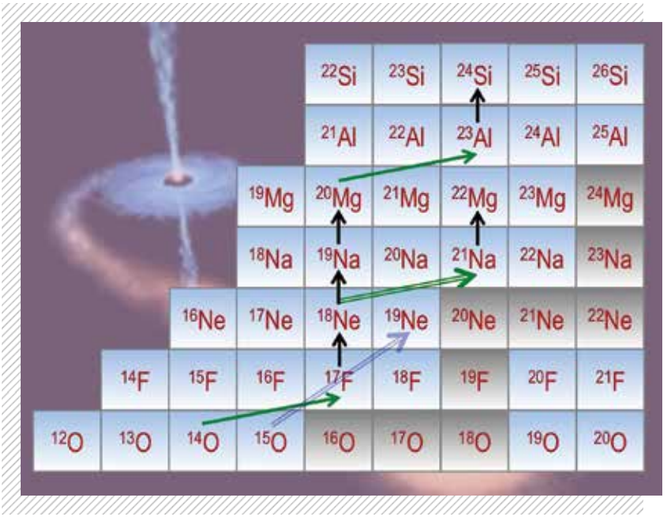

FIGURE 1: Nuclear region of interest to X-ray burst nucleosynthesis—it is now opened up for *ab initio* investigations by the unique capabilities of our SA-NCSM symmetry-guided concept and Blue Waters.

METHODS & RESULTS

We have developed an innovative *ab initio* nuclear structure approach, symmetry-adapted no-core shell model (SA-NCSM) [1], with concomitant computer code dubbed *LSU3shell* [2], that 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. The main computational task is to evaluate a large symmetric Hamiltonian matrix and to obtain the lowest-lying eigenvectors that correspond to the experimental regime. Accuracy is based on the degree of convergence, which is linked to the size of the model space that can be achieved. The SA-NCSM utilizes physically relevant model space of significantly reduced dimensionality compared to ultra-large model spaces encountered by standard *ab initio* approaches. These theoretical advances coupled with the computational power of Blue Waters allow us to reach medium-mass nuclei that are inaccessible experimentally and fall far beyond the limits of other *ab initio* methods. The nuclei of interest represent a considerable challenge requiring computational power of near full-capacity of Blue Waters and its system memory. The following list describes the results and projected studies:

- We have provided the **first** *ab initio* description of the open-shell ^{20}Ne and ^{18}F nuclei. Following this success, we target Ne, Mg, and Si isotopes of importance to nuclear reaction studies. Such reactions in the intermediate-mass region are key to further understanding phenomena like X-ray burst nucleosynthesis and the Ne-Na and Mg-Al cycles.
- Work in progress focuses on one of the most challenging problems in nuclear physics today: achieving an *ab initio* description of the so-called Hoyle state in ^{12}C, the resulting state of the essential stellar triple-alpha process, key to modeling nucleosynthesis and stellar explosions. We have also calculated important low-lying states in ^{12}C, including negative-parity states and giant monopole and quadrupole resonances.
- We have studied ^{10}He (the focus of current experimental proposals), with impact on experimental techniques around the neutron-drip line that involve halo nuclei. Further calculations in even larger model spaces, only feasible with Blue Waters, will be key to resolving the observed inconsistency among several experimental results.

WHY BLUE WATERS

Ab initio nuclear structure studies represent an extremely computing-intensive endeavor. To illustrate the level of complexity, applications to medium-mass nuclei require over hundreds of exabytes of memory to store the Hamiltonian

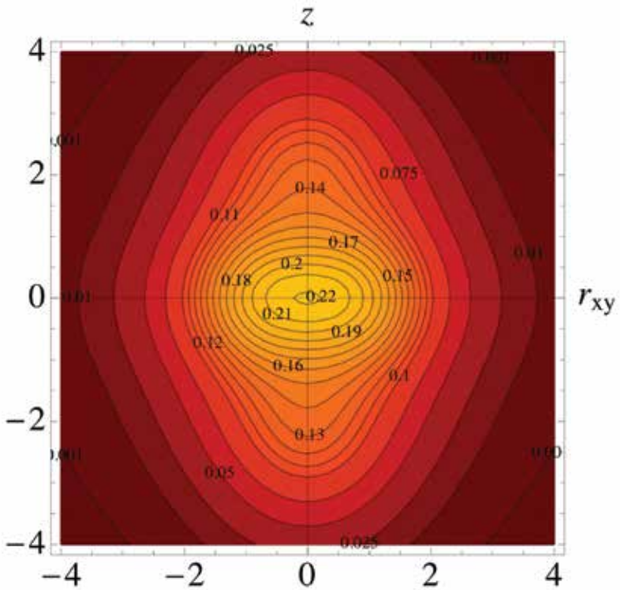

FIGURE 2: Density profile of the challenging ^{10}He ground-state resonance from first principles, revealing a surprising intrinsic structure.

matrix. 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 computing intensity. As a result, SA-NCSM investigations of the intermediate-mass region are beyond the scale of available academic high-performance computing systems. Currently, only Blue Waters provides resources required for the *ab initio* SA-NCSM studies of medium-mass isotopes with cutting-edge accuracy. To capitalize on this opportunity, we drew from the experience and expertise of the Blue Waters staff and managed to improve the scalability of our code. As a result, our largest production runs utilized efficiently 717,600 concurrent threads running on **22,425 Cray XE6 compute nodes** to solve the nuclear eigenvalue problem with Hamiltonian matrices that occupy up to 400 TB of memory. Clearly, Blue Waters represents a unique computational platform that already plays a crucial role in advancing *ab initio* nuclear theory toward new domains.

NEXT GENERATION WORK

The major increase in computational power provided by the second generation of Track-1 system, along with emergent algorithms designed to take advantage of modern massively parallel architectures, will enable *ab initio* theories to start providing information of unprecedented quality for probing fundamental symmetries and physics beyond the standard model. Furthermore, increased computational resources will allow the SA-NCSM framework to address even heavier nuclear systems. For example, addressing neutrinoless double beta decay for ^{48}Ca should become feasible at a level that will reduce large uncertainties in the nuclear structure matrix elements and allows one to determine the neutrino type from planned experiments, which represents one of the most fundamental problems in physics today.

PUBLICATIONS AND DATA SETS

Launey K.D., T. Dytrych, and J.P. Draayer, Symmetry-guided large-scale shell-model theory. *Prog. Part. Nucl. Phys.*, 89 (2016), p. 101, doi:10.1016/j.ppnp.2016.02.001

Dytrych T., et al., Emergence of Simple Patterns in Complex Atomic Nuclei from First Principles. *J. of Phys: Conf. Ser.*, 639 (2015) 012008, doi:10.1088/1742-6596/639/1/012008

Dytrych, T., et al., Electron-scattering form factors for Li-6 in the *ab initio* symmetry-guided framework. *Phys. Rev. C*, 91 (2015), 024326, doi:10.1103/PhysRevC.91.024326

Launey, K. D., et al., Approximate symmetries in atomic nuclei from a large-scale shell- model perspective. *Int. J. Mod. Phys. E*, 24 (2015), 1530005, doi:10.1142/ S0218301315300052

Launey, K. D., et al., Emergent Symmetries in Atomic Nuclei from First Principles. *J. of Phys: Conf. Ser.*, 597 (2015), 012054, doi:10.1088/1742-6596/597/1/012054

Dytrych, T., et al., Emergence of Simple Patterns in Complex Atomic Nuclei from First Principles. *Proc. Int. Conf. Nuclear Theory in the Supercomputing Era 2014*, (2015).

Draayer, J. P., et al., HPC-enabled Nuclear Structure Studies – Description and Applications of the Symmetry-adapted No-Core Shell Model. *J. Phys: Conf. Ser.*, 580 (2015) 012044, doi:10.1088/1742-6596/580/1/012044

Draayer, J. P., et al., Dominant Modes in Light Nuclei – *Ab initio* View of Emergent Symmetries. *J. Phys: Conf. Ser.*, 578 (2015), 012010, doi:10.1088/1742-6596/578/1/012010

Draayer, J. P., K. D. Launey, and T. Dytrych, Fuel of the cosmos. *International Innovation*, 160 (2015) 20.

Langr D., et al., Algorithm 947: Paraperm: Parallel Generation of Random Permutations with MPI. *ACM Transactions on Mathematical Software*, 41 (2014), 5:1, doi:10.1145/2669372

Launey, K. D., et al., Emergence of cluster structures and collectivity within a no-core shell-model framework. *J. of Phys: Conf. Ser.*, 569 (2014), 012061, doi:10.1088/1742-6596/569/1/012061

Dytrych, T., K. D. Launey, and J. P. Draayer, Symmetry-adapted no-core shell model. in *McGraw-Hill Yearbook of Science & Technology*, 2014, doi:10.1036/1097-8542.YB140314

Draayer, J. P., et al., Keynote talk: Unraveling Mysteries of the Strong Interaction – 'Top Down' versus 'Bottom Up' Considerations. *Proc. Int. Conf. Nuclear Theory in the Supercomputing Era 2013*, (2014), pp. 47–61.

Dytrych, T., et al., Utilizing Symmetry Coupling Schemes in *Ab initio* Nuclear Structure Calculations. *Proc. Int. Conf. Nuclear Theory in the Supercomputing Era 2013*, (2014), pp. 62–72.

Launey, K. D., et al., Symmetry-adapted No-core Shell Model for Light Nuclei. *Proc. Fifth Int. Conf. ICFN5 2012*, (2013), pp. 29–38, doi:10.1142/9789814525435_0003

NANOSCALE MECHANICS OF DEFORMATION IN HIGH-CAPACITY LITHIUM-ION BATTERIES

Allocation: Illinois/291 Knh
PI: Huck Beng Chew[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Silicon (Si) is one of the most promising electrode materials for high-performance lithium (Li) ion batteries because it has an order of magnitude higher specific capacity compared to conventional graphite electrodes. However, Si electrodes crack massively during Li insertion since they expand by 300% when fully-lithiated; they also delaminate from the current collector after many charge cycles. We have conducted large-scale parallel molecular dynamics simulations and density functional theory calculations on Blue Waters to uncover the underlying mechanisms for cracking and delamination of the Si electrode during charge cycling; these mechanisms are in excellent agreement with experiments. Blue Waters resources were needed because of the computational scale of the problem and the many computational runs needed to cover the entire parameter space. Our results have provided rich insights into the design and engineering of damage tolerant electrode materials for high-capacity Li-ion batteries.

INTRODUCTION

Lithium ion batteries are high-energy-dense systems that store energy by insertion of Li ions into solid electrodes. Silicon is one of the most promising electrode materials for high-performance Li-ion batteries since it possesses the highest known specific capacity of 4200 mAh/g, which is an order of magnitude greater than conventional graphite electrodes. During lithiation, the Si electrodes form Li_xSi compounds and undergo huge volume expansion of about 300% since one Si atom can theoretically bond with a maximum of x = 3.75 Li atoms. When attached to a metal current collector, such as copper (Cu), the massive and inhomogeneous

FIGURE 1: Sliding and delamination of Li_xSi thin film electrodes from the Cu current collector. (A) Atomic structure of the interdiffused Li-Si-Cu interphase between a Li_xSi electrode and a Cu-current collector. (B) Interface sliding facilitated by the formation of well-delineated and weakly bonded Si-Cu and Li-Cu crystalline atomic layers within this interphase structure. (C) Shear stress versus shear strain response demonstrating distinct regions of stress build-up and release leading to interface sliding, and stress accumulation leading to interface delamination.

volume changes during repeated lithiation and delithiation charge cycles lead to colossal cracking of the Si electrode. Recent studies have shown that Si electrodes of small feature sizes, such as nanowires, nanoparticles, porous structures, and thin films, display significantly higher reversible charge capacities and longer cycle life. In fact, a critical feature size of these nanostructured Si electrodes exists, below which fracture would be completely mitigated. It is believed that the improved fracture resistance originates from the ability of the nanoscale structure to accommodate the lithiation-induced strain by plastic deformation, resulting in lower stresses present during volume changes. However, the delamination of crack-free nanostructured Si electrodes from current collectors after a critical number of charge cycles has been widely reported, resulting in the loss of electrical contact and consequent capacity fade. Even though the cracking of Li_xSi thin films can be mitigated through patterning individual Si islands, the uncracked electrode still delaminates from the current collector after a critical number of charge cycles. To date, much is still unknown about the interface bonding the Si electrode and a metal current collector, such as Cu. Studies have suggested that sliding readily occurs along the Si-Cu interface to accommodate the massive volume changes in lithiated-Si during charge cycling. However, understanding the mechanisms of interface sliding and delamination is complicated by significant intermixing of Cu, Si, and Li atoms at the interface between a lithiated-Si film and the Cu substrate.

METHODS & RESULTS

Using first-principle calculations, we recreate model structures of the interdiffused Li-Si-Cu interphase (Fig. 1A) and show that the interdiffusion among Li, Si, and Cu atoms leads to the formation of well-delineated, crystalline Si–Cu and Li–Cu atomic layers at intermediate Li concentrations (Fig. 1B). These atomic layers are weakly bonded in shear, and readily slide to relieve the interfacial stresses during lithiation processes. Ideally, interface sliding between the Si electrode and the Cu current collector will help limit film stresses introduced by the lithiation process. However, sliding between the Si–Cu and Li–Cu atomic layers cannot occur indefinitely. The formation of pinning defects in the form of $LiSi_3$ compounds along the interface can eventually inhibit sliding (Fig. 1C). The consequential buildup of interfacial stresses leads to delamination failure of the Si electrode from the Cu-current collector. Understanding the atomic-scale mechanisms that promote or impede sliding provides the critical first steps toward designing Si-Cu interface structures to mitigate electrode failure.

WHY BLUE WATERS

To our knowledge, detailed analyses of the sliding and delamination processes of lithiated-Si electrodes from the current collector are virtually non-existent due to the computational complexity of such systems. Firstly, the lithiated-Si structures are amorphous. Recreating these structures requires substantial intermixing between Li and Si atoms through a heating and quenching process in *ab initio* molecular dynamics (MD) simulations and density functional theory (DFT) calculations which are computationally expensive. The computational complexity, coupled with a large number of computational runs to elucidate the interface mechanics as a function of lithium concentration, makes this a process that requires the capacity of Blue Waters. The proposed DFT and *ab initio* MD calculations were performed on Vienna *Ab initio* Simulation Package (VASP).

NEXT GENERATION WORK

Our next focus will be on elucidating the deformation and fracture mechanics of the Solid Electrolyte Interphase, which is critical to the performance of high-capacity lithium ion battery electrodes.

PUBLICATIONS AND DATA SETS

Wang, H., and H.B. Chew, Molecular dynamics simulations of plasticity and cracking in lithiated silicon electrodes. *Extreme Mech. Lett.*, (2016), p. 503–513.

Harpale, A., M. Panesi, and H.B. Chew, Plasma-graphene interaction and its effects on nanoscale patterning. *Phys. Rev. B*, 93 (2016), p. 035416.

Wang, H., X. Wang, S. Xia, and H.B. Chew, Brittle-to-ductile transition of lithiated silicon electrodes: Crazing to stable nanopore growth. *Jour. Chem. Phys.*, 143 (2015), p. 104703.

Li, R., and H.B. Chew, Closed and open-ended stacking fault tetrahedra formation along the interfaces of Cu-Al nanolayered metals. *Philos. Mag.*, 95 (2015), p. 2747-2763.

Wang, H., et al., Atomic-scale mechanisms of sliding along an interdiffused Li-Si-Cu interface. *Nano Lett.*, 15 (2015), 1716-1721.

Harpale, A., M. Panesi, and H.B. Chew, Communication: Surface-to-bulk diffusion of isolated versus interacting C atoms in Ni(111) and Cu(111) substrates: A first principle investigation. *Jour. Chem. Phys.*, 142 (2015), p. 061101.

Li, R., and H.B. Chew, Planar-to-wavy transition of Cu/Ag nanolayered metals: A precursor mechanism to twinning. *Philos. Mag.*,95 (2015), p. 1029-1048.

LATTICE QCD ON BLUE WATERS

PI: Robert Sugar[1]
Collaborators: Alexei Bazavov[2], Kate Clark[3], Carleton DeTar[4] Daping Du[5], Robert Edwards[6], Steven Gottlieb[2], Balling Joo[6], Kostas Orginos[7], Thomas Primer[8], David Richards[6], Doug Toussaint[8], and Frank Winter[6]

[1]University of California Santa Barbara
[2]Indiana University
[3]NVIDIA
[4]University of Utah
[5]Syracuse University
[6]Thomas Jefferson National Accelerator Facility
[7]College of William and Mary
[8]University of Arizona

EXECUTIVE SUMMARY

The goal of this project is to develop highly optimized code for the study of quantum chromodynamics (QCD) on Blue Waters to carry out calculations that will have a major impact on high-energy and nuclear physics. We have optimized and used the Chroma code for the simulation of Clover quarks and the MILC code for the simulation of HISQ quarks. Our long-term objectives with highly improved staggered quark (HISQ) are to generate gauge configurations with physical-mass up, down, strange and charm quarks, to use these configurations to calculate fundamental parameters of the standard model of high energy physics, and to perform precise tests of the standard model. The objective of our Clover quark program is to determine the excited mass spectrum of strongly interacting particles (hadrons) within QCD.

INTRODUCTION

The standard model of high energy physics encompasses our current knowledge of the fundamental interactions of nature. The model has successfully explained a wealth of data from accelerator and cosmic ray experiments over the last 40 years. However, it has been difficult to extract many of the most interesting predictions of quantum chromodynamics (QCD), the component of the standard model that describes the strong interactions. The only way to do so, from first principles and with controlled errors, is through large-scale numerical simulations. These simulations are needed to obtain a quantitative understanding of the physical phenomena controlled by the strong interactions, to determine some of the fundamental parameters of the standard model, and to test them. Despite the successes of the standard model, high-

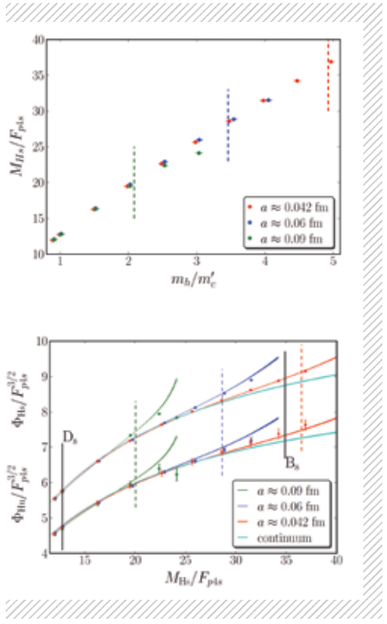

FIGURE 1: Heavy-quark discretization effects in the calculation of the leptonic decay of the B and Bs mesons. The upper panel shows the heavy-strange meson mass M_Hs in units of the quantity Fp4s, as a function of the ratio of the heavy-quark mass, m_h, to the simulation charm-quark mass, m_h/m_c. The lower panel shows the decay constants of a heavy meson made up of a heavy-quark and an up quark, and of a heavy quark and a strange quark, Phi_Hu and Phi_Hs, as a function of M_Hs. We are interested in the extrapolation of M_Hs to M_Bs.

energy and nuclear physicists believe that a more general theory will be required to understand physics at the shortest distances. Thus, QCD simulations play an important role in efforts to obtain a deeper understanding of the fundamental laws of physics.

METHODS & RESULTS

Our objective is to perform calculations of QCD to the precision needed to support large experimental programs in high-energy and nuclear physics. We are using two formulations of lattice quarks. The HISQ formulation is being used to calculate parameters of the standard model, and to perform precise tests on it. In particular, the HISQ formulation is being used to calculate the masses of quarks, which are the fundamental building blocks of strongly interacting matter, and to determine elements of the Cabibo-Kobayashi-Maskawa (CKM) matrix, which are the weak interaction transition couplings between quarks.

Our primary objective with the Clover formulation of lattice quarks is to perform a calculation of the mass spectrum of strongly interacting particles (hadrons). The determination of the excited-state spectrum of hadrons within QCD is a major objective for experiments and is a focus of the $310 million upgrade of Thomas Jefferson National Accelerator Facility. In particular, the GlueX experiment at Jefferson Laboratory will search for "exotic" mesons. These particles are a signature for new states of matter, specifically the presence of gluonic degrees of freedom, predicted by QCD, but thus far not clearly observed. The spectroscopy effort is intended to determine whether the equations of QCD do, in fact, realize the existence of such exotic states of matter. These calculations will be performed before the experiments and, will therefore, provide crucial information about the decay signatures of such exotic states that **will inform and guide the experimental searches**.

Lattice QCD calculations have two steps: First, one generates and saves gauge configurations, which are representative samples of the QCD ground state. Second, the gauge configurations are used to measure a wide range of physical quantities. Generating gauge configurations is the rate-limiting step and requires the most capable computers available. The most computationally expensive component of the second step is to calculate the Green's functions for the propagation of quarks in the gauge configurations. The light quarks used in this calculation also require highly capable computers.

We have made major progress in our efforts to generate gauge configurations and quark propagators using Blue Waters, including the most challenging ensembles undertaken to date. The new HISQ configurations have been used to make the most precise determination to date of the decay properties of some mesons containing strange and charm quarks [1-4]. This work has led to the evaluation of several CKM matrix elements that are important for tests of the standard model; and produced the most precise ratios among up, down, strange, and charm quark masses. The HISQ configurations are also being used in a new, very promising calculation of the decay properties of B and Bs mesons [4]. Important advances have been made in the development of code for the generation of gauge configurations and quark propagators with the Clover formulation [5], and through new approaches for the determination of resonance parameters. **Ground-breaking** results have been obtained for the scattering of K, pi and eta mesons [6-9].

WHY BLUE WATERS

The precision possible with lattice QCD calculations in recent years has reached within a fraction of a percent. Such precision is needed to test the standard model and to obtain a detailed understanding of physical phenomena controlled by the strong interactions. The advent of petascale computers, like Blue Waters, are playing a critical role in these advances.

NEXT GENERATION WORK

The next generation Track-1 systems will allow us to account for the electromagnetic effects and isospin violations in the quark masses and to work at even smaller lattice spacings. These advances will enable us to determine a wide range of quantities at the sub-percent level, making lattice QCD calculations reliable tools in the support of the large experimental programs in high energy and nuclear physics.

PUBLICATIONS AND DATA SETS

Bazavov, A., et al., Determination of |V_{us}| from a lattice-QCD calculation of the K -> pi l nu semileptonic form factor with physical quark masses, *Phys. Rev. Lett.* 112 (2014), 112001 [arXiv:1312.1228].

Bazavov, A., et al., Charmed and light pseudoscalar meson decay constants from four-flavor lattice QCD with physical light quarks, *Phys. Rev.*, D90 (2014) 074509 [arXiv:1407.3772].

Bazavov, A., et al., Gradient Flow Analysis on MILC HISQ Ensembles, *Proceedings of Science (Lattice 2014)* 090 (2014) [arXiv:1411.0068].

Bazavov, A., et al., Charmed and light pseudoscalar meson decay constants from HISQ simulations, *Proceedings of Science (Lattice 2014)* 382 (2014) [arXiv:1411.2667].

Bazavov, A., et al., Symanzik flow on HISQ ensembles, *Proceedings of Science (Lattice 2013)* 269 (2013) [arXiv:1311.1474].

Bazavov, A., et al., Leptonic decay-constant ratio fK/fpi from lattice QCD with physical light quarks. *Phys. Rev. Lett.* 110 (2013),172003 [arXiv:1301.5855].

Bolton, D.R., R. A. Briceno, and D. J. Wilson, Connecting physical resonant amplitudes and lattice QCD, *Phys. Lett. B* 757 (2016), pp. 50-56. [arXiv:1507.07928 [hep-ph]].

Briceno, R.A., et al., The resonant pi gamma -> pi pi amplitude from Quantum Chromodynamics, *Phys. Rev. Lett.* 115 (2015), 24200, [arXiv:1507.06622 [hep-ph]].

Clark, M.A., R.G. Edwards and B. Joo, IPDS '14, Proceedings of the A Framework for Lattice QCD Calculations on GPUs, F.T. Winter, *2014 IEEE 28th International Parallel and Distributed Processing.*

DeTar, C., et al., Charmonium mass splittings at the physical point, *Proceedings of Science (Lattice 2012)* 257 (2012) [arXiv:1211.2253].

Dudek, J.J., R.G. Edwards, C.E. Thomas, and D.J. Wilson. Resonances in coupled pi K - eta K scattering from quantum chromodynamics. *Phys. Rev. Lett.* 113 (2014), 182001, [arXiv:1406.4158 [hep-ph]].

Gamiz, E. et al., K semileptonic form factor with HISQ fermions at the physical point, *Proceedings of Science (Lattice 2013)* 395 (2013) [arXiv:1311.7264].

Gamiz, E. et al., Charmed and strange pseudoscalar meson decay constants from HISQ simulations, *Proceedings of Science (Lattice 2013)* 382 (2013), [arXiv:1312.0149].

Mohler, D., et al., Low lying charmonium states at the physical point, *Proceedings of Science (Lattice 2014)* 085 (2014) [arXiv:1412.1057].

Primer, T., et al., Kaon and D meson semileptonic form factors from lattice QCD, *Proceedings of Science (Lattice 2014)* 374 (2014) [arXiv:1411.1651].

Wilson, D.J., J. J. Dudek, R.G. Edwards, and C.E. Thomas. Resonances in coupled pi K, eta K scattering from lattice QCD, *Phys. Rev.* D91 (2015) 054008, [arXiv:1411.2004 [hep-ph]].

Wilson, D.J., et al., Coupled pi-pi, K K scattering in P-wave and the rho resonance from lattice QCD, *Phys. Rev.* D92 (2015), 094502, [arXiv:1507.02599 [hep-ph]].

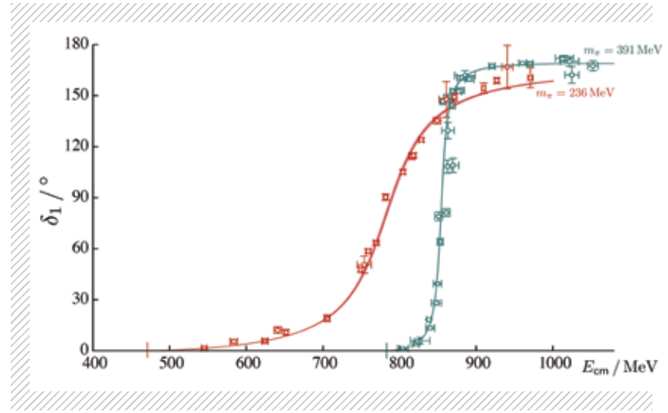

FIGURE 2: Isospin=1, P-wave phase-shift in the pi-pi elastic scattering region (below the Kbar-K threshold). Phase-shift points mapped from finite-volume energies in several volumes and moving frames using the Luscher formalism. Curve shows a Breit-Wigner fit to the energy dependence which shows a clear resonant line-shape. The mass and width of the resonance are in line with expectations for a calculation with m_pi = 391 MeV and m_pi = 236 MeV.

LARGE-SCALE, LONG-TIME MOLECULAR DYNAMICS SIMULATION OF CRYSTAL GROWTH

Allocation: NSF PRAC/1.00 Mnh
PI: Sharon C. Glotzer[1]
Co-PIs: Joshua A. Anderson[1], Michael Engel[1]

[1]University of Michigan

EXECUTIVE SUMMARY

In order to discover the fundamental principles underlying the process of crystal growth we study soft-matter systems that self-assemble into a plethora of different structure types. Our open-source simulation package HOOMD-blue [1], which is designed for GPU architectures, offers molecular dynamics [2] as well as hard-particle Monte Carlo simulation techniques [3] and grants us access to the workings of a large range of systems with various interactions. We are simulating the growth of a variety of crystal structures from the fluid. By comparing the growth of simple and complex or even aperiodic structures, we aim for a universal understanding of the formation of ordered structures from disordered fluids. We anticipate that our studies will impact multiple fields of chemistry, physics, materials science, chemical engineering, geology, pharmaceutical sciences, and any disciplines in which crystal formation, nucleation and growth, and/or the emergence of order from disorder are of interest.

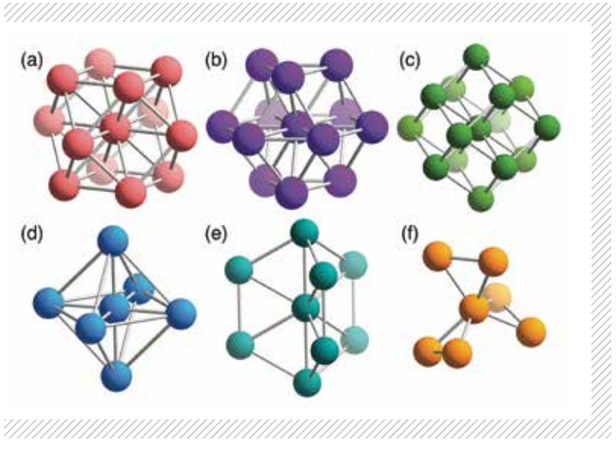

FIGURE 1: Coordination polyhedra of the simple crystal structures grown with attractive isotropic pair potentials. (a) *cF*4-Cu, (b) *hP*2-Mg, (c) *cI*2-W, (d) *cP*1-Po, (e) *hP*2 (compressed), (f) *cP*4-Li.

INTRODUCTION

Crystal growth is still a largely obscure process that researchers have been trying to clarify for decades [4]. According to classical theories, crystals grow from the liquid in an atom-by-atom fashion. This model seems to work well enough for simple crystal structures, which consist of particles or atoms of only one kind that all have the same crystallographic environment in the final crystal. However, if the ensuing crystal structure will contain dozens, hundreds, or thousands of atoms in a unit cell—or if it is aperiodic, as are quasicrystals, meaning that no unit cell exists at all—how does each atom find its position in the crystal?

While the observation of crystal growth in experiments has been, until very recently, an almost impossible undertaking in atomic systems and remains challenging on the soft-matter length scale, molecular dynamics simulation can grant access to this phenomenon on adjustable time and length scales. By simulating crystal structures of varying complexities with HOOMD-blue and observing particle attachment and particle environments prior to and after their attachment to a growing crystal, we can identify characteristic growth modes and compare structural motifs in the fluid and crystalline phases.

METHODS & RESULTS

In previous studies, we identified different crystal structures that occur in systems growing from an oscillating pair potential [5], as well as a Lennard-Jones Gauss potential. We are now investigating crystal growth with atomistic resolution in order to resolve the relationship between the complexity of a crystal structure and its growth mechanism.

We are performing molecular dynamics simulations of the growth of simple crystal structures. We are investigating the densest sphere packings: the cubic close packing (*ccp*, also: *cF*4-Cu structure type) and the hexagonal close packing (*hcp*, also: *hP*2-Mg structure type), as well as the body-centered cubic packing (*cI*2-W structure type), the simple cubic structure (*cP*1-Po), a simple hexagonal structure (hP1 structure type), a simple chiral cubic structure (*cP*4-Li structure type), and a compressed version of the hexagonal close packing (*hP*2) with different resulting particle environments.

For all of those structures we have now determined the respective crystallization temperatures and extracted snapshots of a newly formed crystal nucleus at that temperature. These snapshots will be used for seeded growth simulations conducted at the crystallization temperature, in order to (a) bypass nucleation, which is a rare event and an interesting mechanism that should be studied separately, and (b) avoid multiple nucleation events, which would make the data much harder to interpret. These seeded simulations can then be successively scaled up to systems containing many millions of particles in order to reach a regime where size-effects of the crystallite play a vanishing role.

In the following months, increasingly complex, periodic crystal structures will be simulated, including different Frank-Kasper phases, clathrates, and other, previously unreported structures. Furthermore, the growth of an icosahedral quasicrystal [5] will be simulated and analyzed. Consequently we will be able to better understand the difference between the growth of an aperiodic and a periodic crystal structure, or otherwise describe how both can be regarded as variations of the same principle mechanism.

By tuning the parameters of the isotropic pair potential that we are using, we can access phase regions that exhibit crystal growth from either the fluid or the gas phase. Contrasting both regimes will help us understand how universal the observed motifs of crystal growth truly are.

WHY BLUE WATERS

Blue Waters is the **only** NSF-funded system that offers large-scale access to GPUs. GPUs can accelerate soft matter simulations by more than an order of magnitude vs. CPU-based calculations, and therefore drastically improve the turn-around for research. The performance of the HOOMD-blue code was previously benchmarked with Blue Waters [6], demonstrating its ability to scale on up to 1,024 nodes. Blue Waters allows us to run the large simulations needed to investigate the growth of crystal structures that have large unit cells, such as clathrates, or that are aperiodic, like quasicrystals. The HOOMD-blue molecular dynamics package, which we develop and have heavily optimized for excellent scaling performance, is the fastest code available for this type of molecular dynamics simulation.

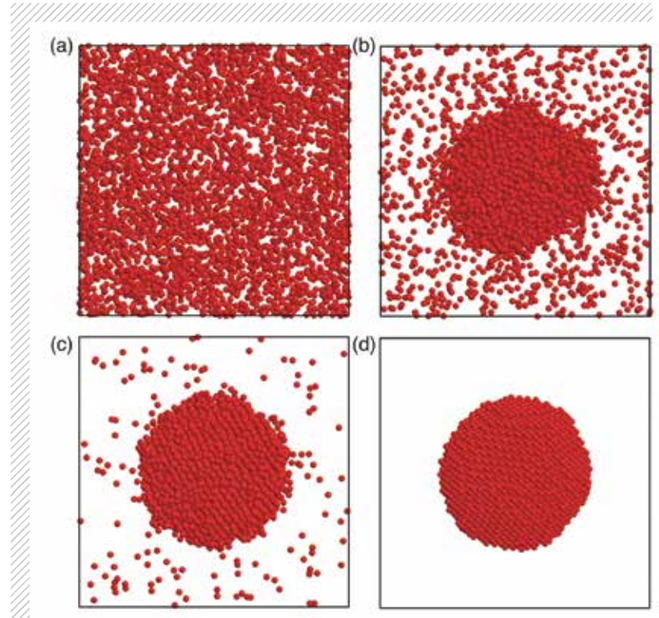

FIGURE 2: Simulation snapshots of (a) the fluid phase, (b) the emerging liquid droplet, (c) the forming crystallite, and (d) the final crystal, at successively decreasing temperatures.

NEXT GENERATION WORK

Within the framework of this project, we are focusing on the exploration of crystal growth. The nucleation of a crystalline seed is a physically distinct process that requires far larger compute resources to explore, as it is a rare event that needs to be probed with statistical methods and hence a large number of very large sample runs. We envision that studying the nucleation of diverse crystal structures from various initial states would become feasible with the next-generation Track-1 computing architecture.

MULTISCALE MODELING OF BONE FRACTURE AND STRENGTH

Allocation: Illinois/50.0 Knh
PI: Iwona Jasiuk[1]
Co-PI: Seid Koric[1]
Collaborator: Fereshteh Sabet[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

In this project, nonlinear finite element models of trabecular bone were used with the aim of building a predictive tool to capture plasticity and strength of this biological material. The complex, random, and irregular features of bone were effectively imaged using micro-computed tomography to build three-dimensional finite element models. The complexity of the structure, material and geometric nonlinearities, along with complex contact conditions, make the numerical analysis very challenging even on the latest high-performance computing (HPC) platforms. Our results closely match experimentally measured stiffness and strength and capture experimentally observed plasticity of the trabecular bone, demonstrating the potential of the model to predict the mechanical behavior of bone. This model provides a crucial step in creating a multiscale model of whole bone based on patient-specific data. Such a model will serve as a novel tool for more accurate prediction of osteoporosis, a bone disease characterized by bone's susceptibility to fracture.

INTRODUCTION

Osteoporosis, which is a growing clinical problem in aging societies, is a bone disease characterized by low bone density and deterioration of bone's structure leading to bone fragility and increased risk of fractures [1]. It is a silent disease with no symptoms prior to fractures and no cure, but treatments can slow its progress. Thus, its early and accurate diagnosis is crucial. Currently, bone quality is assessed clinically by measuring the bone mineral density while other factors such as bone's complex hierarchical structure also contribute to bone's properties. Thus, a new approach is needed for more accurate diagnosis of osteoporosis. Computational mechanics modeling can provide an effective new tool for the clinical assessment of bone.

Bone is a multifunctional biological tissue. One function is to serve as a structural support for soft tissues in the body. As a structural material bone has an ideal combination of properties when healthy: high stiffness, strength, and fracture toughness, and light weight. Bone is made of a compact cortical bone forming an outer core and a spongy trabecular bone filling an inside space and ends of long bones (Fig. 1). Such geometry is optimal as it minimizes weight, reduces bearing stresses at joints, and allows the body to withstand high functional loads. These superior properties are due to the hierarchical structure of bone spanning from molecular to macroscopic levels [2].

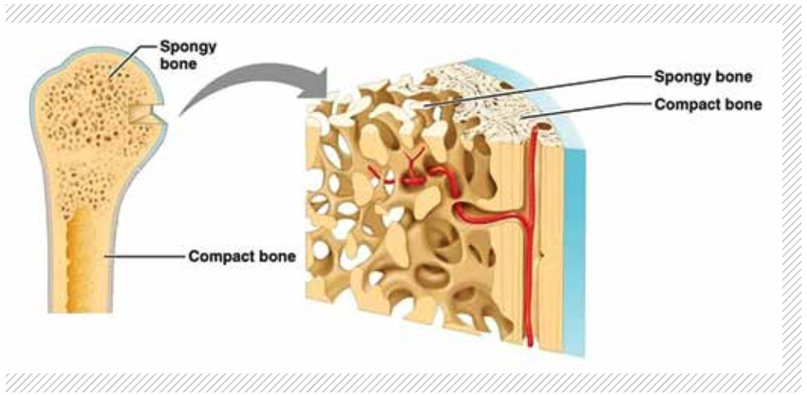

FIGURE 1: Structure of bone showing cortical (compact) and trabecular (spongy) bone.

Dense cortical bone and porous trabecular bone work together to provide excellent load-bearing properties. Osteoporosis-related fractures mostly occur in trabeculae-rich areas of bone. Trabecular bone is also the primary site for insertion of orthopedic implant systems. Thus, mechanical properties of trabecular bones are of high clinical and research interest for prediction of age- and disease-related bone fractures, optimizing treatments to reduce fracture risks, as well as designing improved implant systems [3-4].

In this project, trabecular bone was simulated using non-linear finite element (FE) models to gain a better understanding of its mechanical behavior. Such models have the potential to build a structure-based predictive tool to assess strength and damage locations in patient-specific bones. It can also serve as a guide to engineers for design of **novel** synthetic bio-mimetic and bio-inspired materials for a wide range of engineering applications.

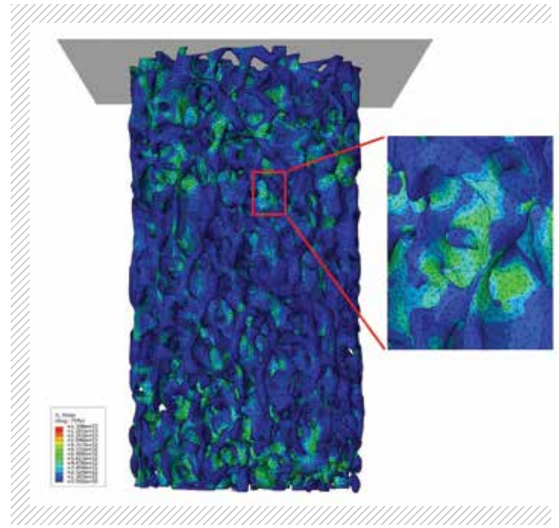

FIGURE 2: Distribution of stress in a trabecular bone sample under compression.

METHODS & RESULTS

In this project, nonlinear micro-computed tomography (micro-CT) FE model of trabecular bone was built as a cost-effective compliment to experiments and conventional methods that are expensive, time-consuming and prone to artifacts. In this model, the complex, random, and irregular features of microstructure were captured with microstructure data from the micro-CT imaging. The model was then meshed and analyzed under loading with numerical simulations on trabecular bone samples with complex random microstructures.

An example of our trabecular bone FE model along with stress distribution is shown in Figure 2. Figure 3 shows a sample stress-strain response of a trabecular bone sample under compression obtained from our model and compared to experiments. Our results indicate a close match with experimentally measured stiffness and strength. Plasticity of trabecular bone is also well captured in our simulations, showing the model's potential to predict the mechanical behavior of bone.

The effectiveness of micro-CT FE models as predictive tools for mechanical properties of trabecular bone is directly dependent on proper calibration and validation. As such, this project investigates the influence of using different constitutive laws and parameters at the tissue level on apparent response through a systematic study to serve as a guide for modeling bone.

WHY BLUE WATERS

Our simulations involved a large number of finite elements, material, and geometric nonlinearities, along with complex contact conditions which made the numerical analysis challenging, even on the latest HPC platforms. The model analysis was performed utilizing up to 256 CPU cores and the large memory capabilities of Blue Waters. It has been recently shown that the similar multifrontal solver had enough scalability and robustness to perform computations on large ill-conditioned FE analysis problems on many thousands of CPU cores [5], thus potentially opening the door for future higher fidelity and complexity simulation studies in biomechanics.

NEXT GENERATION WORK

We plan to extend our modeling to a whole-bone level with the aim of building a patient-specific predictive tool for clinical use capable of assessing bone quality and bone response under loading. Once the model is established for healthy bone, bones

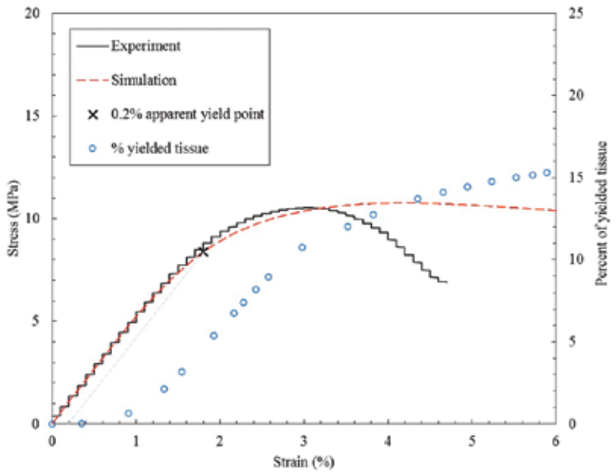

affected by diseases like osteoporosis, or diabetes and also aged bone will be modeled to better understand how the properties of bone are affected by various factors, assess effects of medications and to predict risk of bone fracture for patients.

PUBLICATIONS AND DATA SETS

Koric, S., F.A. Sabet, O. Jin, I. Jasiuk, Direct numerical simulation of bone plasticity and strength. *International Symposium on Plasticity*, Keauhou, Hawaii (2016).

Sabet, F., et al., Hierarchical modeling of plasticity and strength of trabecular bone, *XXIV ICTAM*, 21-26 August (2016) Montreal, Canada.

FIGURE 3: Prediction of stress-strain behavior of trabecular bone under compression along with percent of yielded tissue.

NANOSCALE ELECTRONIC DEVICES WITH NEMO5

Allocation: NSF PRAC/1.24 Mnh
PI: Gerhard Klimeck[1]
Collaborators: James Charles[1], Fan Chen[1], Jim Fonseca[1], Junzhe Geng[1], Xinchen Guo[1], Hesameddin Ilatikhameneh[1], Tillmann Kubis[1], Daniel Mejia[1], Bozidar Novakovic[1], Michael Povolotskyi[1], Rajib Rahman[1], Prasad Sarangapani[1], Archana Tankasala[1], Yu Wang[1], Lloyd C. L. Hollenberg[2], Sven Rogge[3], Joe Salfi[3], and Michelle Simmons[3]

[1]Purdue University
[2]The University of Melbourne
[3]The University of New South Wales

EXECUTIVE SUMMARY

Relentless downscaling of transistor size has continued according to Moore's Law for the past 50 years. According to the International Technology Roadmap for Semiconductors (ITRS), transistor size will continue to decrease in the next 10 years, but foundational issues with currently unknown technology approaches must be pursued. The number of atoms in critical dimensions is now countable. As the materials and designs become more dependent on atomic details, the overall geometry constitutes a new material that cannot be found as such in nature. NEMO5, the software package developed by the Institute for Nanoelectronic Modeling (iNEMO), is designed to comprehend the critical multi-scale, multi-physics phenomena through efficient computational approaches and to quantitatively model new generations of nanoelectronic devices including transistors and quantum dots, as well as to predict novel device architectures and phenomena [1,2].

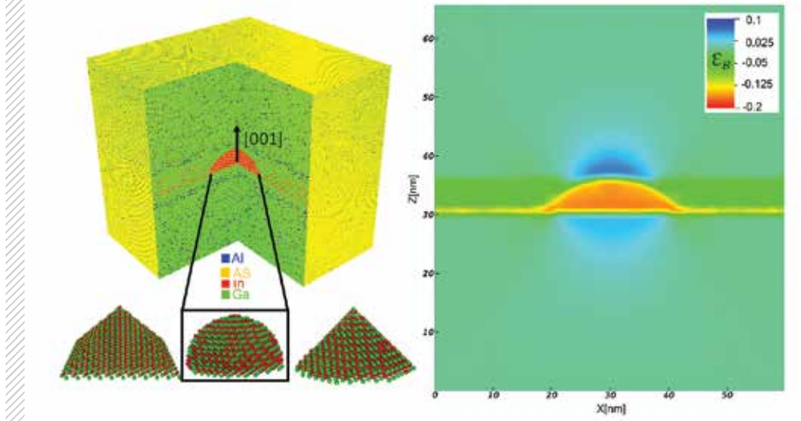

FIGURE 1: NEMO5 can simulate alloys such as this AlGaAs/InGaAs quantum dot system. On the left is the atomistic structure of a simulated 10 million atom system. The right shows the biaxial strain distribution dome-shaped quantum dot of size 20 nm by 5nm. Strain calculations are one of the first steps toward ensuring accurate treatment of the electronic qualities of devices.

INTRODUCTION

The U.S. is a market leader in the semiconductor industry, which produces many high-paying, high-technology jobs. The U.S. semiconductor industry is one of the nation's largest export industries, and the U.S. holds one-third of the global semiconductor device market worth over $300 billion. Simultaneously, the end of Moore's Law scaling as we know it will be reached in 10 years, with devices expected to be about 5 nm long and 1 nm in their critical active region width. 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. Fundamental questions remain about the downscaling of the CMOS (complementary metal–oxide–semiconductor) switch and its eventual replacement. What is the influence of atomistic local disorder from alloy, line-edge roughness, dopant placement, and fringe electric fields? How do lattice distortions due to strain affect carrier transport in nanometer-scale semiconductor devices such as nanowires, finFETs, quantum dots, and impurity arrays? Can power consumption be reduced by inserting new materials and device concepts?

NEMO5 is developed and used by the Institute for NanoElectronic Modelling (iNEMO) at Purdue to address these fundamental questions on a variety of semiconductor devices. Besides enabling basic engineering, physics, and materials science research, NEMO5 is used by leading semiconductor firms to design future devices. The source code, binaries and support for academic use are available through nanoHUB.org.

METHODS & RESULTS

iNEMO's research on Blue Waters encompasses simulations of incoherent scattering effects on tunnel FET devices, time-resolved quantum transport, a compact model for self-assembled quantum dot heterostructures, and a multiscale approach for nitride-based light emitting diodes.

The current mechanism of tunneling field effect transistors (TFETs) is through interband tunneling rather than thermionic emission in typical MOSFETs (metal-oxide semiconductor field-effect transistors). Tunneling from the valence to conduction band has the potential to yield significantly improved subthreshold slopes to allow lower supply voltages and better efficiency than MOSFETs. Electron transport through the overall device, including the source and drain, entails significant amounts of computationally demanding scattering, which cannot be ignored in realistic device performance predictions, and these simulations agree well with experimental resistivity data. In an ultra-scaled silicon TFET, simulation also found that incoherent scattering has a significant impact on performance.

Time-resolved quantum transport data allow more accurate calculation of energy/delay device

characteristics during turn-on for studying novel effects based on wave function phase manipulation, and as an alternative research path to simulating dissipation and nonlocal scattering. Time-resolved quantum transport simulations are performed with the semi-empirical tight binding method. Initial work shows that this approach is valid up to about 1 mV/ps, corresponding to a few GHz in realistic transistors.

Self-assembled quantum dots are highly strained heterostructures with optoelectronic applications, such as infrared photodetectors, intermediate band solar cells, optical amplifiers, and quantum dot lasers. A universal behavior in terms of strain magnitude and profile is observed in atomistic strain simulations of dots with different shapes and materials. Atomistic strain simulations are more accurate but more expensive than analytic continuum solutions. Simulations on Blue Waters showed that both techniques indicate that the strain depends on the aspect ratio of the dot, and not on the individual dimensions. This has allowed for the formulation of a compact model of strain effects on self-assembled quantum dots.

Multi-quantum well LEDs have carrier flow through complex quantum states and the NEGF (non-equilibrium Greens function) approach has been used to model nitride-based diodes that provide blue mid-to-high power light. Strong electron-electron/phonon scattering thermalizes carrier distribution in the wells. A multiscale approach models the barriers in non-equilibrium with the wells in equilibrium. Simulations of 120 nm long devices show IV characteristics matching experimental results when accounting for temperature differences and external resistance. This research has revealed how p-side wells provide higher radiative recombination, supporting experimental evidence, due to more deeply filled electronic states, leading to a stronger overlap in the electron-hole densities.

WHY BLUE WATERS

In many cases the work could not be accomplished in a reasonable amount of time without Blue Waters, and for larger simulations the work could not be accomplished on other available systems. Blue Waters staff provide **exemplary** support and user outreach to guide system usage, help with issues as they arise, and assist with code performance and scaling. In particular, they have provided custom scripts and permissions to facilitate management of an allocation with a large research group.

NEXT GENERATION WORK

A next-generation Track-1 system will enable efficient reliability predictions of modern nanodevices—a crucial milestone in their design. Atomistic deviations due to statistical and fabrication fluctuations can have increasingly large effects on a device's operation as devices are scaled, and it is imperative to have a strong grasp on these phenomena in the coming years.

PUBLICATIONS AND DATA SETS

Novakovic, B., and G. Klimeck, Atomistic quantum transport approach to time-resolved device simulations, *International Conference on Simulation of Semiconductor Processes and Devices (SISPAD)*, Washington, D.C., Sept. 9-11, 2015, doi:10.1109/SISPAD.2015.7292245

Charles, J., P. Sarangapani, R. Golizadeh-Mojarad, R. Andrawis, D. Lemus, X. Guo, D. Mejia, J. E. Fonseca, M. Povolotskyi, T. Kubis, and G. Klimeck, Incoherent Transport in NEMO5: Realistic and Efficient Scattering on Phonons, *Journal of Computational Electronics*, accepted.

Wang, Y., A. Tankasala, L. C. L. Hollenberg, G. Klimeck, M. Y. Simmons, and R. Rahman, Highly tunable exchange in donor qubits in silicon, *Nature Quantum Information*, Article number: 16008 (2016). doi:10.1038/npjqi.2016.8

Tankasala, A., J. Salfi, S. Rogge, G. Klimeck, and R. Rahman, Atomistic simulations of negatively charged donor states probed in STM experiments. *APS Meeting Abstracts*, Baltimore, MD, March 14-18 (2016).

Tankasala, A., Y. Wang, J. Bocquel, B. Voison, M. Usman, S. Rogge, M. Simmons, L. Hollenberg, R. Rahman, and G. Klimeck, Modeling exchange interaction in bulk and sub-surface donor pairs in silicon for two-qubit gates, Poster Session, ARO Quantum Computing Program Review, San Diego, Aug 10-12 (2015).

Ilatikhameneh, H., T. Ameen, G. Klimeck, and R. Rahman, Universal Behavior of Strain in Quantum Dots, *arXiv*:1509.08004 [cond-mat.mes-hall], Sept. 2015.

COMPRESSIBLE LARGE EDDY SIMULATION OF A FILM-COOLED STAGE-ONE NOZZLE AT DIFFERENT FREESTREAM TURBULENCE LEVELS

Allocation: Innovation and Exploration/300 Knh
PI: Gregory M. Laskowski[1]
Co-PI: Gustavo Ledezma[2]
Collaborators: Qiqi Wang[3], ZJ Wang[4], James Kopriva[1], Rathakrishnan Bhaskaran[2], Lawrence Cheung[2], and Umesh Paliath[2]

[1]GE Aviation
[2]GE Global Research
[3]MIT
[4]University of Kansas

EXECUTIVE SUMMARY

Gas turbines are the backbone of aircraft propulsion. Therefore, technologies that improve the fuel efficiency of these turbines can have a significant impact on the U.S. economy and can substantially reduce polluting emissions. The quest for greater turbine efficiency has led to increased firing temperatures, subjecting turbine components to extremely high temperatures. One common technique for abating these temperatures is film cooling, which bleeds cool air from the compressor stage of the engine and discharges it through small holes in the turbine blade walls, providing an insulating layer of cool air. The focus of this project is to investigate the impact of freestream turbulence on film cooling, including mixing, turbulence decay, and boundary layer development.

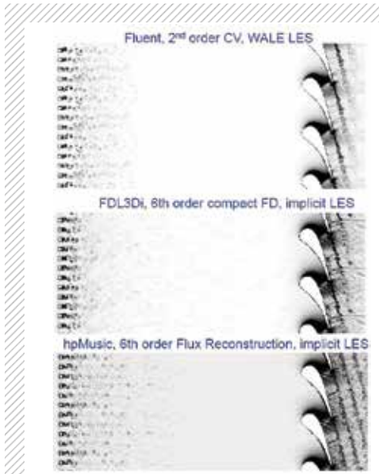

FIGURE 1: Numerical schlieren for an uncooled transonic vane and 5% turbulence intensity (a) unstructured 2nd order solver (b) structured 6th order solver (c) unstructured 6th order solver.

INTRODUCTION

Because of their power density and efficiency, gas turbines are, and will continue to be, the backbone of narrow- and wide-body aircraft propulsion. According to the Federal Aviation Administration, the United States alone consumed approximately 35.6 billion gallons of aviation fuel in 2012. Technologies that can further improve fuel efficiency can have a significant impact on the U.S. economy while reducing emissions.

The modern high-bypass turbofan engine is based on the Brayton cycle. Air is ingested in the engine and passes through the fan. A majority of the airflow bypasses the core of the engine to increase propulsive efficiency. The core flow is compressed, increasing total pressure (Pt) and total temperature (Tt). Fuel is added and burned in the combustor, increasing Tt and slightly reducing Pt, and the flow is expanded through the turbine, reducing Pt and Tt. Work extracted by the turbine drives the compressor. The turbine also drives a fan at the front of the engine in order to increase the mass flow through the engine, thus increasing the propulsive efficiency. In order to improve thermal efficiencies the turbine inlet temperature and compressor pressure ratio have historically increased with time.

The need for higher turbine efficiencies to reduce fuel consumption keeps pushing the firing temperature up. As a result, all components in the high-pressure turbine (HPT) are subject to extremely

FIGURE 2: Numerical schlieren for a cooled transonic vane at 6% and 20% turbulence intensity for unstructured second order solver.

high temperatures that, in many cases, exceed their melting point. One of the most common techniques to protect HPT components against the high thermal loads is film cooling, which bleeds airfoil cooling flow into the hot gas path so the coolant forms an insulating layer between the airfoil and the hot gases. The air used for film cooling is by-pass air from the late stages of the compressor. Therefore, there is a performance penalty due to the extraction of the cooling air that could instead have been used to extract work in the HPT.

The focus of this project is to investigate the impact of freestream turbulence on film cooling, including mixing, turbulence decay, and boundary layer development. In parallel, the investigation is evaluating the importance of higher order numerical schemes in capturing the turbulent mixing.

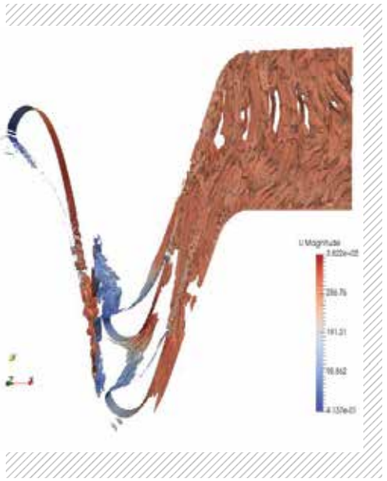

FIGURE 3: Setting the stage for a stage analysis with a high-pressure turbine stationary vane and moving blade.

METHODS & RESULTS

The availability of massively parallel computer environments has increased the popularity of scale-resolving computational fluid dynamics techniques such as direct numerical simulation (DNS), large eddy simulation (LES), and hybrid RANS-LES methods (HLES). Despite being computationally expensive, scale-resolved methods often bring significant improvements in accuracy over RANS (Reynolds-averaged Navier-Stokes) methods and provide an opportunity to understand the fundamental physics. LES resolves a significant part of the energy spectrum and models only the smaller, more dissipative subgrid scales present in the flow, resulting in accurate predictions of turbulent mixing.

For this study, we modified the uncooled turbine vane used by Arts and Rouvroit [1], adding a generic film-cooling hole shape based on Saumweber *et al.* [2] on the suction side of the vane. Figure 1 shows simulation results for uncooled geometry using second order unstructured Ansys Fluent v17, sixth order structured FDL3Di [3], and sixth order unstructured hpMusic [4]. Simulations for high and low levels of turbulence with the film cooling hole are ongoing (Figure 2).

WHY BLUE WATERS

Aero-thermal LES of turbomachinery problems at engine relevant Reynolds numbers is computationally intensive [5]. Physical testing has clearly shown that turbulence is an important contributor to turbulent mixing in aero-thermal applications of turbomachinery, but often testing is not capable of explaining the reason [1,2]. This is due, in part, to the range of length and time scales associated with turbulence and the challenges of measuring

single-point and multi-point correlations in high Reynolds number and high Mach number flows associated with rotating machinery. Simulations on Blue Waters are enabling us to replicate observations in physical testing while providing data that can provide insight into the mechanisms responsible for those observations.

NEXT GENERATION WORK

The current simulations are simplified to provide insight into the complex mixing of a single hole. In reality, the HPT airfoils are cooled by many holes. Furthermore, a high pressure turbine stage consists of a stationary airfoil (stator) and a rotating airfoil (rotor). The former turns the flow for the latter which extracts work to drive the compressor. Executing a cooled stage analysis would be the natural next step, followed by optimization of features. GE is collaborating with MIT to develop optimization capabilities in an LES framework [6], and preliminary simulations of a stator/rotor are ongoing (Figure 3).

PUBLICATIONS AND DATA SETS

Laskowski, G. M., J. Kopriva, V. Michelassi, S. Shankaran, U. Paliath, R. Bhaskaran, Q. Wang, C. Talnikar, Z. J. Wang, and F. Jia Future directions of high-fidelity CFD for aero-thermal turbomachinery research, analysis and design, 46th AIAA Fluid Dynamics Conference, AIAA Aviation (2016).

NUMERICAL STUDY OF THE MANY-BODY LOCALIZATION TRANSITION

Allocation: Illinois/50.0 Knh
PI: David J. Luitz[1]
Collaborators: Xiongjie Yu[1] and Bryan K. Clark[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

We study the details of the many-body localization (MBL) transition and its adjacent phases in one-dimensional quantum spin systems through extensive numerical simulations using an exact diagonalization technique, which allows for the solution of the Schrödinger equation in finite spin chains. This transition is a formidable example of the breakdown of statistical mechanics in strongly interacting disordered systems, which do not thermalize.

Our main finding—using concepts from quantum information—is that in the vicinity of the critical point the system displays a mixture of localized and delocalized states, pointing to a highly nontrivial nature of the critical regime. We also find evidence in the probability distributions of matrix elements of local observables that in the regime of intermediate disorder, where the transport is subdiffusive, the generally assumed ansatz of the eigenstate thermalization hypothesis (ETH) in the scaling of the variance of these distributions as well as their shape has to be modified.

INTRODUCTION

While statistical mechanics usually assumes the presence of a heat bath, in the advent of cold atomic systems the study of completely isolated quantum systems has gained fundamental importance. In the absence of any coupling to the outside world, the dynamics of quantum systems are governed by the unitary time evolution described by the Schrödinger equation, and it is not clear how these systems would reach a thermal state. However, experiments of generic quantum systems point to a rapid thermalization. An important step towards an understanding of this phenomenon was the idea that quantum chaos leads to thermalization via the ETH [2,3,8]. While the ETH can explain thermalization in

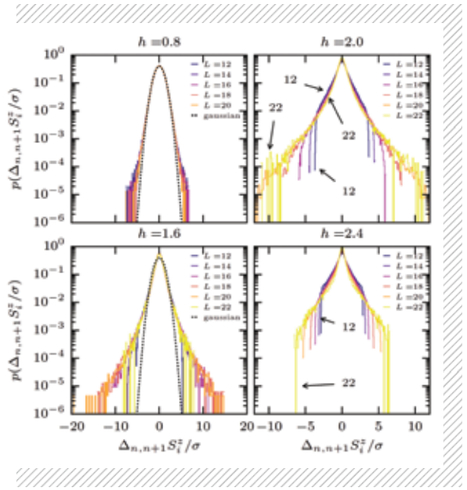

FIGURE 1: Probability distribution of the difference in local magnetization for states from the center of the spectrum of the random Heisenberg chain at disorder strength h. The distributions are normalized by their variance to compare their shapes for different chains of length L and their mean is zero. At weak disorder, the distributions are close to Gaussian and become increasingly non-Gaussian in the regime where transport is subdiffusive. The MBL critical point in this system is around $h_c \sim 3.7(1)$ [5].

many generic systems, noteworthy counterexamples have been discovered, one of which are many-body localized (MBL) systems [1], which violate ETH and do not thermalize.

MBL systems have been mostly explored in one dimension, due to the enormous computational complexity of the problem, caused by the exponential growth of the many-body Hilbert space, and the best-studied systems to date are interacting quantum spin chains with disorder (cf. e.g. [4,5]). In these, a transition between a thermal ETH phase and an MBL phase can be observed and the focus of our work lies in the transition and its vicinity to better understand how the thermalization mechanism breaks down at a critical strength of disorder. We use both concepts from ETH, mostly local operators like the local magnetization, and concepts from quantum information theory, like the entanglement entropy, to study the transition.

METHODS & RESULTS

We use state-of-the-art large-scale exact diagonalization techniques for the sparse Hamiltonian to obtain high energy eigenstates (cf. [5]), which are crucial for directly addressing the ETH. To leading order, local observables in the basis of eigenstates are required by the ETH to be diagonal matrices with Gaussian fluctuations around the mean (and around zero in the off-diagonal entries), whose variance vanishes exponentially with system size. We have studied the behavior of these matrix elements by a systematic study of the probability distributions of the local magnetization for various system sizes and strengths of disorder. This was achieved by calculating histograms of the matrix elements in the eigenbasis at zero energy transfer over a large number of eigenstates and realizations of disorder. Our main finding is that at weak disorder, the ETH ansatz is well-verified and the distributions are very close to a Gaussian distribution, whose variance scales with the exponential law in system size as predicted by ETH. However, at intermediate disorder, this is no longer true and the distributions become strongly non-Gaussian, while the variance of the distribution decreases slower than expected from ETH. In the MBL phase, ETH is violated and the variance of the (non-Gaussian) distributions does not decrease with system size, thus not leading to thermalization.

In ongoing work, we explore whether the deviation from the scaling of the variance from the ETH ansatz as well as the non-Gaussian distributions are linked to the subdiffusive transport, observed in the same parameter regime by previous studies [6, 7].

The MBL transition can also be addressed using concepts from quantum information theory, most prominently by the entanglement entropy (EE). The entanglement entropy can be viewed as the thermodynamic entropy of a subsystem if the rest of the system is considered to be the heat bath. This leads to the requirement from ETH, to the fact that thermodynamic entropies are extensive, and that if the system thermalizes, the EE has to scale by a volume law. Contrarily, in the MBL phase, this is explicitly broken and the EE displays an area law. We consider the transition between these two behaviors in our current work and develop a formalism based on general concepts (strong subadditivity of EE) that allows for a quantification of the EE scaling of single eigenstates in inhomogeneous systems. The erratic behavior of the EE as a function of subsystem size in disordered systems seems at first not to allow for such an analysis but we prove in [9] that in periodic chains, the EE is a concave function of subsystem size, if one averages over all cuts of the same length (cf. Fig. 2). By an extensive analysis of the probability distributions of the slope of this cut averaged EE as a function of subsystem size, we find that in the critical regime, a mixture of volume and area law states appears, which opens new questions about the nature of the MBL transition.

WHY BLUE WATERS

The numerical study of quantum many-body systems is an extremely hard problem that requires massive computational resources. Here, we consider disordered systems, which makes the problem harder by several orders of magnitude as one has to average over many configurations of disorder. The disorder corresponds basically to the effect of "dirt" on the system and is modeled by random potentials.

Solving the Schrödinger equation for one realization of such a dirty system already requires more than one node of Blue Waters for large systems and repeating this calculation for hundreds to thousands of configurations is only feasible on a massively parallel setup like Blue Waters.

PUBLICATIONS AND DATA SETS

Luitz, D. J., Long tail distributions near the many-body localization transition. *Phys. Rev. B* 93 (2016), 134201, DOI: 10.1103/PhysRevB.93.134201

Yu, X., D.J. Luitz, and B.K. Clark, Bimodal entanglement entropy distribution in the many-body localization transition, *arXiv*:1606.01260.

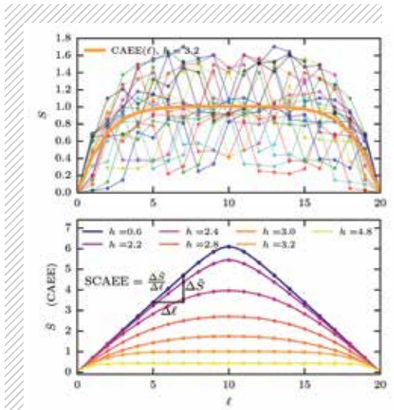

FIGURE 2: Top: Entanglement entropy as a function of subsystem size for one high-energy eigenstate of a random Heisenberg chain of length L=20 and at disorder strength h=3.2 for all possible left cut positions. The erratic behavior is completely removed by the average over all cut positions (CAEE). Bottom: Typical cut averaged EE curves for single eigenstates at various disorder strengths with well-defined slope (SCAEE).

DNS AND STOCHASTIC STUDY OF THE RELATIVE MOTION OF HIGH INERTIA PARTICLES IN ISOTROPIC TURBULENCE

Allocation: NSF PRAC/1.77 mnh
PI: Sarma L. Rani[1]
Collaborator: Rohit Dhariwal[1]

[1]University of Alabama in Huntsville

EXECUTIVE SUMMARY

The objective of our research is to investigate the role of turbulence in driving the relative velocities and positions of inertial particles in isotropic turbulence. First, we studied the effects of turbulence on the relative motion of high-inertia particle pairs in isotropic turbulence. Accordingly, we performed direct numerical simulations (DNS), as well as Langevin simulations (LS) based on a probability density function (PDF) kinetic model for pair relative motion. Recently, we developed a stochastic theory that derives closures in the limit of high Stokes number for the diffusivity tensor in the PDF equation for particle pairs. The diffusivity contained the time integral of the Eulerian two-time correlation of fluid relative velocities seen by pairs that are nearly stationary. The two-time correlation was determined analytically using the

approximation that the temporal change in the fluid relative velocities seen by a pair occurs principally due to the advection of smaller eddies past the pair by large scale eddies. Hence, two diffusivity expressions were obtained based on if the pair center of mass remained fixed during flow time scales, or moved in response to integral-scale eddies.

A quantitative analysis of the aforementioned stochastic theory is performed through a comparison of the particle pair statistics obtained using LS with those from DNS. LS consists of evolving the Langevin equations for pair separation and relative velocity, which is statistically equivalent to solving the Fokker-Planck form of the pair PDF equation. LS of particle pair dispersion were performed using three closure forms of the diffusivity—i.e., one containing the time integral of the Eulerian two-time correlation of the seen fluid relative velocities, and two analytical diffusivity expressions. In the first closure form, the two-time correlation was computed using DNS of forced isotropic turbulence laden with stationary particles. The three diffusivities are extensively analyzed to quantify the effects of the approximations made in deriving them. Pair relative-motion statistics obtained from the three sets of Langevin simulations are compared with the results from the DNS of (moving) particle-laden forced isotropic turbulence for St_η = 10, 20, 40, 80 and Re_λ = 76, 131. Here, St_η is the particle Stokes number based on the Kolmogorov time scale, and Re_λ is the Taylor micro-scale Reynolds number.

INTRODUCTION

The turbulence-driven relative motion of high-inertia particles is relevant in astrophysical scenarios, such as the interstellar medium, protoplanetary disks, and the atmospheres of planets and dwarf stars. Specifically, the "sticking" of dust particles in protoplanetary disks is believed to be the mechanism for planetesimal formation. An intriguing question that astrophysicists are investigating concerns the effects of turbulence on the dispersion, sedimentation, collisional coalescence and fragmentation of dust grains. The viscous relaxation times, τ_v, of these particles are significantly large, with estimated St_η ~ 10-100, where $St_\eta = \tau_v/\tau_\eta$ is the Stokes number based on the Kolmogorov time scale τ_η.

The two principal quantities describing the relative motion of inertial particles in a turbulent flow are the radial distribution function (RDF), which is a measure of the particle spatial clustering, and the PDF of pair relative velocities, which quantifies the particle encounter rate. The RDF and the relative velocity PDF are both key inputs to the particle collision kernel and depend sensitively on the Stokes number. Both statistics can be determined through DNS of particle-laden turbulent flows. However, DNS suffers from the well-known computational limitation on the Reynolds numbers that can be achieved. This drawback of DNS is one of the motivating factors for developing PDF equation-based stochastic models for particle-laden turbulent flows.

We developed a stochastic theory for the relative velocities and positions of high-inertia pairs in forced isotropic turbulence [1]. The theory involved deriving a closure for the diffusivity tensor characterizing the relative-velocity-space diffusion current in the PDF kinetic equation of particle-pair separation and relative velocity. Since we had considered the $St_r >> 1$ limit, the pair PDF equation is of the Fokker-Planck form (St_r is the Stokes number based on the time-scale τ_r of eddies whose size is of the order of pair separation r). Using the diffusivity formulation, one can perform Langevin simulations of pair relative velocities and positions, which is equivalent to simulating the Fokker-Planck equation.

In this context, the current study has two main objectives. First, we perform a quantitative analysis of the three forms of the diffusivity derived in [1]. The insights gained will help us understand the implications of the approximations made in deriving the diffusivities, as well as guide future improvements to the theory. The second objective is to compute the relative motion statistics of particle pairs using both DNS and Langevin simulations (LS), and compare the corresponding results.

METHODS & RESULTS

DNS of forced isotropic turbulence were performed using a discrete Fourier-expansion- based pseudospectral method. Simulations were performed over a cubic domain of length 2π discretized using N^3 grid points, with periodic boundary conditions. The fluid velocity is advanced in time by solving the Navier-Stokes equations in rotational form, as well as the continuity equation for an incompressible fluid. Direct evaluation of the non-linear convective terms in the Navier-Stokes equations is computationally intensive. Hence, a pseudospectral approach is adopted wherein the non-linear terms are first computed in physical space, and then transformed into the spectral space. P3DFFT library [2] was used to carry out the transforms between physical and spectral spaces.

RDF is a well-established measure of particle clustering. In figure 1, the RDF is presented as a function of St_η at four separations r/η = 6, 12, 18, and 24 (η is the Kol-mogorov length scale). LS results are compared with the data from the DNS performed in [3] DNS and [4]. The Février et al. [3] data were for Re_λ = 69, while the current DNS data are for Re_λ = 76. There is excellent agreement between the LS RDF and the two sets of DNS RDFs at all four separations, particularly for St_η > 10. The [4] theory significantly over predicts the RDFs for high Stokes numbers at all separations.

WHY BLUE WATERS

Direct numerical simulations are the most accurate numerical approach to resolve all the temporal and length scales in a turbulent flow. However, DNS of particle-laden turbulent flows are computationally intensive and requires running code using tens of thousands of cores. Also, each DNS run is expected to generate several terabytes of data. Due to these CPU time and storage requirements, Blue Waters is the **ideal platform** to achieve our objective and proved to be an invaluable resource in computing key inputs to our stochastic theory. For instance, the theory requires as an input the two-time Eulerian correlations of fluid relative velocities seen by particle pairs. Evaluation of the two-time correlation for nearly half a trillion pairs is a highly computationally intensive process. We were only able to compute this quantity because of the Blue Waters access, where we ran the code on 20,000 cores.

PUBLICATIONS

Dhariwal, R., Sarma L. Rani, and D. L. Koch, Analysis and comparison with DNS of a stochastic model for the relative motion of high-stokes-number particles in isotropic turbulence. *68th Ann. Mtg. APS DFD*, Boston, MA, 2015.

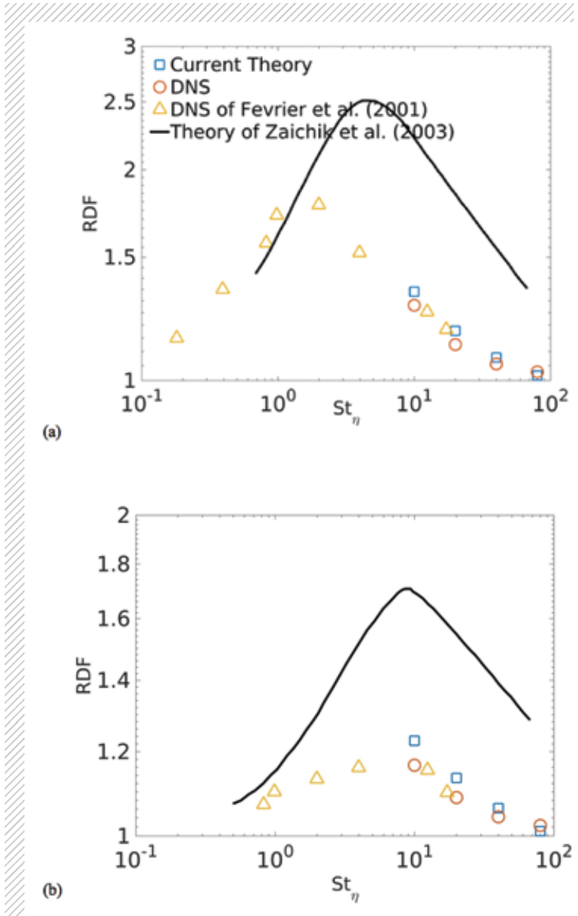

FIGURE 1: Radial distribution function (RDF) versus St_η at specific radial separations: (a) r/η = 6, and (b) r/η = 18. In each plot, squares and circles represent data from CF1 and current DNS at Re_λ = 76; triangles represent DNS data at Re_λ = 69 taken from [3]. Solid line represents data from [4] theory for Re_λ = 69.

THE COMPLEXITIES OF HIGH REYNOLDS NUMBER TURBULENCE

PI: P.K Yeung[1]
Co-PIs: A. Majumdar[2], R.D. Moser[3], D. Pekurovsky[2], and J.J. Riley[4]
Collaborators: B.L. Sawford[5] and K.R. Sreenivasan[6]

[1]Georgia Institute of Technology
[2]San Diego Supercomputer Center
[3]University of Texas at Austin
[4]University of Washington
[5]Monash University
[6]New York University

EXECUTIVE SUMMARY

A half-trillion-grid points simulation of turbulence performed on Blue Waters has provided clear evidence of extreme events where local measures of the deformation and rotation of local fluid elements are much stronger than and have topological properties different from previous understanding. More recently, a substantial algorithmic change in the interpolation of fluid particle velocities on a distributed domain has allowed us to track the motion of some 300 million fluid particles at an affordable cost, both forward and backward in time. Small-scale intermittency is examined as a function of scale separation in space, and through statistics taken along the trajectories of a large collection of fluid particles, taken singly or in pairs. The effects of extreme events are of special interest. New progress is also made for the study of turbulent mixing of passive scalars of very low molecular diffusivity, using a dual-resolution algorithm based on a highly scalable compact finite-difference scheme for the scalar fields.

INTRODUCTION

Fluid motions in many fields of science and engineering are typically turbulent, with disorderly fluctuations over a range of scales. One of our fundamental objectives is to understand the nature of intense fluctuations, which are highly localized in time and space, and to use this understanding to address the effects of fine-scale intermittency [1] in applications. As reported in a recent publication, we have observed the robust existence of "extreme events" where dissipation rates and entropy fluctuations occur at magnitudes as large as $O(10^5)$ times the mean, are nearly coincident in space, and have topological features different from conventional thought. Since these extreme events are localized in space they are important in various flow phenomena dominated by the small-scale motions in turbulence.

Much of the reason for the importance of turbulence itself is enhanced mixing and dispersion of fluid elements with distinct properties or other entities such as heat, chemical species, or passive contaminants carried in the flow [2,3]. For example, examination of fluid particle trajectories going backward in time provides crucial information on how entities originally far apart in space are brought together, which is an important question when disease agents in the environment are involved. The effectiveness of mixing is dependent on a coupling between advective transport by the velocity field and molecular diffusion, which may be strong or weak depending on the nature of the diffusing substance or property. Current work with Blue Waters includes the analysis of recent data and conduct of new simulations, with the latter designed to address both dispersion and mixing in parameter regimes difficult to reach without resources of the magnitude provided on Blue Waters.

METHODS & RESULTS

The only technical approach capable of truly capturing extreme events localized in both time and space is direct numerical simulation, in which we solve exact equations for conservation of mass and momentum. Because of the focus on small scales, we consider isotropic turbulence on a 3D periodic domain. Although Fourier pseudo-spectral methods for this type of geometry are communication-intensive, remote memory addressing and topologically aware scheduling on Blue Waters have allowed us to perform a production $8{,}192^3$ simulation, which surpasses recent work in both the Reynolds number reached and small-scale resolution. The new simulation data have been analyzed from probabilistic, spatial and temporal viewpoints. Extreme events are seen to possess a spatial structure (Fig. 1) differing from conventional notions. Several numerical tests have confirmed these results are robust.

Given the instantaneous velocity field at a mixed set of grid points, the main task of tracking fluid particles is to calculate the particle velocity by interpolation from the velocity values defined at neighboring grid points. We use cubic splines interpolation which is fourth-order accurate and twice differentiable. For each particle, this scheme requires information at 64 grid points that, in general, will reside in different sub-domains and hence require communication. To reduce communication costs, we divide particles among message passing interface tasks dynamically, based on their instantaneous positions instead of the particle IDs. With this algorithm, we have tracked 300 million fluid particles and obtained values of all velocity gradients evaluated at the particle positions as well. Since the governing equations are not reversible in time, we study backward dispersion by post-processing of pairs and tetrads selected based on their final-time positions. Details of this backward-tracking procedure are given in Buaria, Sawford & Yeung (2015) and a Ph.D. thesis (D. Buaria, 2016). A long-standing challenge in the study of dispersion is the pursuit of so-called *Richardson scaling* which occurs at intermediate times long enough for the memory of initial conditions on pairwise separations to fade. This regime is difficult to reach and, especially for backward statistics, requires high Reynolds number—for which our $8{,}192^3$ simulation appears to be just sufficient for the classical scaling to materialize in a convincing manner.

For turbulent mixing, we are working with funding from the Blue Waters PAID subaward to simulate passive scalars at high Schmidt number (i.e. low diffusivity) which requires finer resolution than the velocity field. Previous works were usually limited in either the Schmidt number or the Reynolds number. However, our new collaborator T. Gotoh [4], has developed a new algorithm where the velocity field is computed on a coarser grid using pseudo-spectral methods, but the scalar field is computed on a finer grid using high-order compact finite difference methods. Compact finite difference methods have the advantage of inherently modest communication cost. Test runs to date performed with a new simulation code for this work are very encouraging in both scientific fidelity and performance.

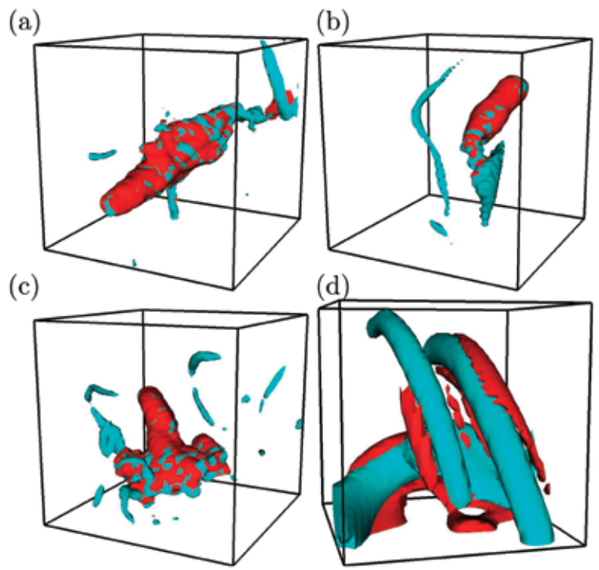

FIGURE 1: Perspective view of 3D color contour surfaces of intense dissipation (red) and entropy (cyan) from three $8{,}192^3$ snapshots (A,B,C), showing a structure distinct from sheets and tubes in image from $1{,}024^3$ at lower Reynolds number (lower right).

WHY BLUE WATERS

In general, an $8{,}192^3$ simulation is almost 16 times as expensive as one at $4{,}096^3$, but it is necessary (as verified by numerical tests) so that high-quality results can be obtained at sufficiently high Reynolds number. A large allocation of time on a multi-petaflops computer such as Blue Waters is thus vital. The outstanding level of support we received through the PRAC program has also been crucial. Generous provision of mass storage resources associated with Blue Waters has allowed us to build a petabyte-sized computational laboratory for answering long-standing questions in the study of turbulence.

NEXT GENERATION WORK

Possible future targets at similar or greater rigor include various turbulent flows subjected to other external influences such as buoyancy, solid-body rotation, or electromagnetic forces.

PUBLICATIONS AND DATA SETS

Buaria D., B.L. Sawford, and P.K. Yeung, Characteristics of two-particle backward dispersion in turbulence at different Reynolds numbers. *Phys. Fluids*, 27 (2015) p. 105101.

Iyer, K.P., K.R. Sreenivasan, and P.K. Yeung, Refined similarity hypothesis using 3D local averages. *Phys. Rev. E*, 92 (2015), p. 063024.

Yeung, P.K, High-Reynolds-number turbulence in a petascale computational laboratory. Invited keynote lecture, *EuHiT Turbulence Conference*, Gottingen, Germany, May 2016.

Yeung, P.K., Extreme events and acceleration statistics at high Reynolds number. Invited speaker, *International Congress on Theoretical and Applied Mechanics*, Montreal, Canada, August 2016.

Yeung, P.K., X.M. Zhai, and K.R. Sreenivasan, Extreme events in computational turbulence. *PNAS*, 112:41 (2015), pp. 12633-12638.

LATTICE SCREENING AND OPTICAL PROPERTIES OF NOVEL PEROVSKITE PHOTOVOLTAIC MATERIALS

Allocation: Illinois/500 knh
PI: André Schleife[1]
Co-PI: Joshua A. Leveillee[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Due to the extremely quick rise of their photoconversion efficiency, hybrid organo-metal-halide perovskites have recently garnered a large amount of attention as potential materials for efficient, cost-effective, and broadly applicable next-generation photovoltaics. However, the influence of excitonic effects on optical absorption and exciton binding energies in these materials is not well understood. In particular, lattice and free-carrier contributions to dielectric screening are currently under investigation. Spin-orbit interaction plays an important role in these materials as well. We use many-body perturbation theory to compute optical properties and an approximate approach to explore the influence of lattice screening. Our results show that when all these effects are taken into account, very good agreement with the experiment is obtained. This work constitutes the **first** step towards a full, first-principles treatment of these effects that will be broadly applicable for material design of novel photovoltaics.

INTRODUCTION

Hybrid organo-metal-halide perovskites of the form ABX_3 have recently garnered a large amount of attention [1]. In this formulation, A is an organic cation, B is a metal cation, and X is the halide anion. Elements such as B={Pb, Sn, Ge} and X={I, Br, Cl} have been investigated, and CH_3NH_3 is of large interest for A. As a result of their exceptional optical absorption, these materials were used as dye sensitizers. However, current interest is attributed to the quick rise of photo-conversion efficiencies to more than 20% within a few years. The high photo-conversion efficiency of prototypical $CH_3NH_3PbI_3$ needs to be understood to successfully design materials with better performance or improved stability.

For successful device engineering, detailed knowledge of the electronic structure and optical properties is needed. In particular, the influence of excitonic effects on optical absorption and exciton binding energies is not well understood. Exciton binding energies are important because

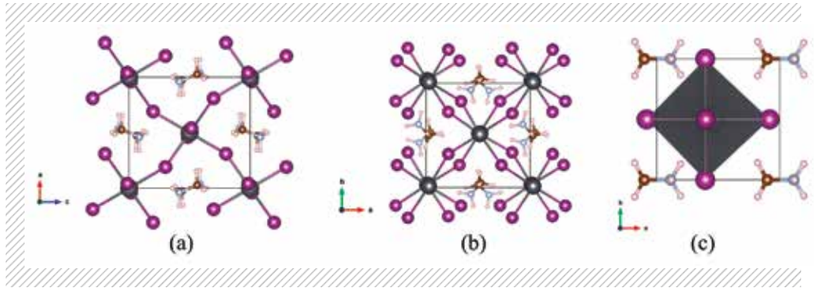

FIGURE 1: Unit cells of the low-temperature orthorhombic phase (a), the room-temperature tetragonal phase used in operational devices (b), and the high-temperature cubic phase (c).

they critically determine how efficiently electrons and holes can be separated in a photovoltaic device. Concurrently, a fundamental understanding is needed because excitonic effects and binding energies sensitively depend on the screening of the electron-hole interaction in the material. Better understanding will influence general first-principles models and could lead to an efficient, cost-effective, and broadly applicable set of materials for next-generation photovoltaics.

METHODS & RESULTS

We study optical properties using first-principles, theoretical spectroscopy based many-body perturbation theory [2]. We compute optical absorption spectra by solving a Bethe-Salpeter equation for the optical polarization function. Single-particle energies in the excitonic Hamiltonian are approximated using a generalized-gradient approximation for exchange and correlation. The spin-orbit interaction is approximately included for orbital energies of collinear spins, but not for wave functions used to compute Coulomb matrix elements. Using this approach on Blue Waters, we can accomplish the numerical challenge of converging the optical absorption spectra on Brillouin zone sampling.

The potentially large influence of lattice screening on the electron-hole interaction makes the description of these materials particularly challenging. We employ an approximate technique by Bechstedt et al. to incorporate the influence of lattice polarizability [3]. Experimental work has determined the static dielectric constant to be in the range of 30 to 32 [4], but even values as large as 1,000 were reported [5]. The approximation used here is merely the first step towards a **more sophisticated**, first-principles approach to clarify the influence of free-carrier and lattice screening.

We compute the optical-absorption spectra of three different polymorphs of $CH_3NH_3PbI_3$: the low-temperature orthorhombic phase, the room-temperature tetragonal phase used in operational devices, and the high-temperature cubic phase (Fig. 1). Our results for the cubic phase in figure 2 indicate a significant influence of the spin-orbit interaction due to heavy lead (Pb) atoms in the material. We also find a significant influence of the lattice screening on the optical-absorption spectra. If the lattice contribution is fully taken into account, excitonic effects are strongly reduced, and the spectrum approaches the density function theory result.

We find very good agreement with an experimental result [6], indicating that not only is lattice screening is important, but also that our approximate

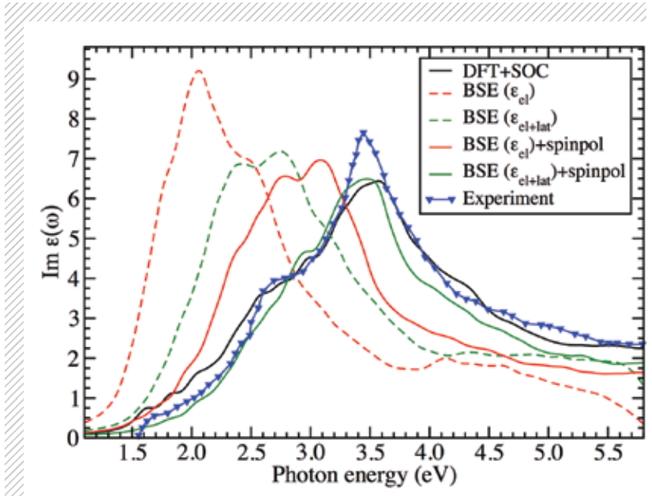

FIGURE 2: Imaginary part of the frequency-dependent complex dielectric function of $CH_3NH_3PbI_3$ computed using different levels of theory is compared to data from experiment.

treatment of the spin-orbit effect is a reasonable approximation. These results are critically important for development of a comprehensive first-principles approach.

WHY BLUE WATERS

The solution of the Bethe-Salpeter equation is computationally challenging, as it requires computing very large exciton Hamiltonian matrices (ranks more than 100,000). We use either an iterative diagonalization scheme to compute their eigenvalues, or we employ a time-propagation approach to compute optical absorption spectra. Each run requires large amounts of memory, disk storage, and fast communication between the two. Many calculations are needed to ensure convergence of spectra and exciton binding energies on Brillouin zone sampling. Blue Waters provides an outstanding computational package that allows us to carry out these simulations for complicated materials such as $CH_3NH_3PbI_3$.

Interactions with the Blue Waters team were extraordinarily helpful. As a result, we are now involved in the Joint Laboratory for Extreme Scale Computing (created as part of the Blue Waters Project) aimed at using the efficient ChASE iterative diagonalization scheme that also runs on graphics processing units (GPUs). While this is work in progress, the Blue Waters project was instrumental in initiating and facilitating this work.

NEXT GENERATION WORK

A next-generation Track-1 system will be instrumental for advanced computational material science research. In our particular field, the most pressing goal is to connect accurate atomistic studies, which fully account for free-carrier screening and electron-phonon effects such as lattice screening, with mesoscale simulations. Furthermore, large-scale materials design requires a large number of such accurate calculations, which, due to their extreme computational cost, can only be achieved on a future Track-1 system. Finally, extending this work towards nanoscale materials such that semiconductor nanocrystals or nanoplatelets will push the computational capabilities of current supercomputers, requires the availability of a future Track-1 system to be successful.

This material is based upon work supported by the National Science Foundation under Grant No. CBET-1437230.

PARTICLE-RESOLVED DIRECT NUMERICAL SIMULATIONS OF FLUID-SOLID HEAT TRANSFER

PI: Shankar Subramaniam[1]
Co-PI: Bo Sun[1]

[1]Iowa State University

EXECUTIVE SUMMARY

Heat transfer between solid particles in a fluid flow occur in multiple engineering applications, such as pneumatic conveying. The purpose of this work is to simulate fluid-solid heat transfer using particle-resolved, direct numerical simulation (PR-DNS). Gas-solid heat transfer has previously been simulated and modeled using our PR-DNS approach in a steady flow through a fixed bed of spherical particles. To extend these models to account for liquid-solid heat transfer, such as in a flow of sand particles in water, we need higher resolution simulations to capture the thermal boundary layers surrounding individual particles. Blue Waters enables the study of this heat transfer problem in liquid-solid flow regimes. We simulate heat transfer in steady flow past a fixed bed of spherical particles with high resolution. The PR-DNS database allows us to extend the models for gas-solid heat transfer to liquid-solid heat transfer.

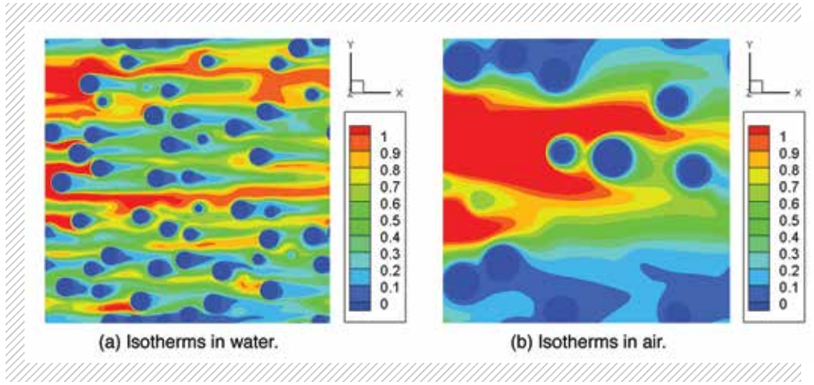

FIGURE 1: A contour plot of the non-dimensional fluid temperature field $\Phi = (T(x,t)-T_s)/(T_{m,in}-T_s)$, where T is the fluid temperature, T_s is the sphere temperature, and $T_{m,in}$ is the bulk fluid temperature, in the cross-sectional plane of a three dimensional periodic computational cubic box at solid volume fraction of 0.1, particle Reynolds number of 20 for (a) Prandtl number of 10 corresponding to water and (b) Prandtl number of 0.7 corresponding to air. The length of the computational cubic box is (a) $15D$ and (b) $7.5D$, where D is the particle diameter. The flow direction is from left to right. The differences in the shapes of the isotherms between liquid and gas are clearly visible and the benefits of high resolution are also apparent.

INTRODUCTION

An improved understanding of fluid-solid heat transfer is crucial for process and component design in multiple engineering applications such as pneumatic conveying systems that transfer powders, granules, and other dry bulk materials through an enclosed pipeline using a combination of pressure differential and the flow of a gas, such as air or nitrogen. The use of computational fluid dynamics (CFD) simulations of multiphase flow are an efficient alternative to experiments for process and design optimization and are becoming more common. Predictive CFD with accurate sub-models has the potential to improve the efficiency of CO_2 capture, as well as clean energy generation technologies. The predictive capability of multiphase CFD simulations depends on models for interphase transfer terms such as the closure model for interphase heat exchange.

Although improved gas-solid heat transfer models for CFD simulations have been proposed [1], they are not verified for liquid-solid heat transfer. Extending these improved models to liquid-solid heat transfer requires high-resolution PR-DNS data that capture the flow and thermal features in the boundary layer surrounding individual particles. Since liquids diffuse momentum faster than heat, the thermal boundary layer in liquid-solid flows is thinner than in gas-solid flows. Therefore, in water-solid flow, higher grid resolution is needed to capture the thermal boundary layer accurately. Resources like Blue Waters are needed to simulate the physics accurately. The outcome of physics-based predictive models of liquid-solid heat transfer will result in the better design of pipelines to transport materials safely and efficiently.

METHODS & RESULTS

To simulate heat transfer in liquid-solid flow accurately, PR-DNS using the Particle-resolved Uncontaminated-fluid Reconcilable Immersed Boundary Method (PUReIBM) [2, 3] approach have been performed with high grid resolution. PUReIBM solves mass and momentum equations, and the convective-diffusive scalar transport equation in the liquid phase by imposing exact no-slip and no-penetration boundary conditions on the surface of each isothermal particle. The solid phase is represented using an immersed boundary forcing in the computational domain.

Figure 1(a) shows contours of non-dimensional temperature in steady flow past a fixed homogeneous bed of 644 monodisperse spheres in the cross-sectional plane of a three dimensional periodic cubic box in a dilute flow (solid volume fraction of 0.1). For this simulation of liquid-solid flow with heat transfer, the grid resolution is $D_m = 80$, where $D_m = D/\Delta x$, D is the sphere diameter and Δx is the grid spacing. The flow direction is from left to right. It is observed that compared with the gas-solid flow in Figure 1(b), the thinner thermal boundary layer forms around each sphere. A thermal wake behind each particle [4] is also seen for the high Prandtl number of 10. Based on these PR-DNS data, extended models for liquid-solid heat transfer are being developed. The extended models for fluid-solid flow will be used in CFD simulations of industrial applications and enable the designer to optimize the design of industrial systems more accurately and efficiently.

WHY BLUE WATERS

Blue Waters has allowed us to perform simulations of heat transfer in liquid-solid flow with high resolution that is necessary to capture the correct physics. Capturing the thermal boundary layer is critical for fluid-solid flow since the quantification of fluid-solid heat transfer depends on local temperature gradient along the particle surface. Simulation of heat transfer in liquid-solid flow is time-consuming and costly. The case, as shown in Figure 1, requires $1{,}200^3$ grid nodes to solve for velocity and temperature fields at a grid resolution of 80. To satisfy the requirements above, Blue Waters is an essential tool for our research to simulate the physics accurately.

NEXT GENERATION WORK

We would like to perform PR-DNS simulation of a fluidized bed with a million particles on Blue Waters (using more than 5,000 central processing units) and expect fast data transfer from Blue Waters to local clusters. The goal is to understand instabilities and clustering formation in gas-solid flows, verifying scale separation in the particle phase and provide the ensemble of realizations for small-scale PR-DNS that is used for parametric studies.

PUBLICATIONS AND DATA SETS

Sun, B., Modeling Heat and Mass Transfer in Reacting Gas-solid Flow Using Particle-resolved Direct Numerical Simulation. Ph.D. thesis, Iowa State University, 2016.

CONFINEMENT AND ORIENTATION EFFECTS IN SEMI-FLEXIBLE POLYMER STRUCTURES

Allocation: GLCPC/448 Knh
PI: David Ackerman[1]

[1]Iowa State University

EXECUTIVE SUMMARY

Polymer chains composed of multiple components have the ability to self-assemble into fascinating microscopic structures. The structures they form are reproducible, highly regular, and can be controlled by manipulating the composition of the polymer chains. The chains themselves can range from highly flexible, to extremely rigid. Between these extremes, the class of semi-flexible diblock co-polymers (block co-polymers with two distinct blocks) contains polymers useful in flexible electronics, biomedical application, and nano-scale templating. Critical to all of these uses is an understanding of how to control the structures on a nanoscale level. Computational simulations are an efficient way to characterize the structures and explore the polymer compositions that lead to them. Until very recently the computational resources to study semi-flexible polymers did not exist. Now, through the power of the Blue Waters system, it is possible to study these polymers and apply the knowledge of how the structures form to materials design.

INTRODUCTION

This research studies the structure of polymers on the microscopic scale. The particular class of polymers are semi-flexible diblock co-polymers which are a mixture of two components, potentially with vastly different properties. There is great interest in these types of polymers because the mixture of properties enables the resulting material to exhibit behavior desirable in industrial applications, including organic electronics [1], lithographic templating [2], self-assembly [3], and many others.

As the long polymer chains interact with each other, they self-assemble into a wide range of microscopic structures. These structures affect the properties of the resulting material. By combining polymer blocks with differing properties, bulk and surface behavior can be tuned by manipulating the way the polymers organize. As an extension of this, the structures themselves can also be used in building microscopic devices or can be used as a template to control placement of particles in a regular pattern.

Being able to predict and control how polymers organize enables construction of devices that meet specified design goals. If, for example, an application requires a microparticle that exhibits a specific property at regularly spaced intervals along its surface, this can be achieved by selecting one of the polymer components to have that property and then designing the diblock polymer to form a structure that places that polymer at the desired spacing along the surface. This work enables control by determining what structures can be formed and understanding the conditions that are necessary to form them.

Past work has focused predominantly on bulk melts or surface coatings. With the strong interest in nanoparticles and nanostructures, this work uses an alternate calculation method to work with arbitrarily shaped particles to study the structures that form from polymers in nanoparticles. In addition to the ability to model arbitrary shapes, a key aspect of this work is that it utilizes a method to study chains of varying flexibility. **Cutting-edge** polymer applications including, for example, flexible organics often use polymer chains described as semi-flexible. The model used in this work is tailored to these types of chains and can accurately capture the physics they exhibit. While the model itself [4] is not new, the computational power required to use it has prevented the application of it for all but the most basic systems [5]. The power of the Blue Waters system allows this work to focus on a class of polymers and applications that have not been well studied in past simulations.

METHODS & RESULTS

This work has been performed by modeling microparticles of polymers and generating structures in them using a self-consistent field approach that searches for the minimum energy structure with a finite element method. A broad sweep of polymer properties has identified several candidate structures and work is underway to refine them to create a phase diagram of the structures that form from a given set of conditions. The value of the phase diagram is in its ability to predict the results arising from a given combination of polymer properties. A scientist wishing to generate a particular structure can utilize a phase diagram to design the polymer chains to obtain the final microstructure. Since it is typically possible to control the chain composition to a high degree, the properties required from the phase diagram can be created in an experimental setting. This allows these phase diagrams to be readily applied in experiments.

WHY BLUE WATERS

This research requires a mixture of numerous smaller simulations requiring only a few dozen computational nodes to develop candidate microstructures and several jobs requiring thousands of nodes to refine them. The Blue Waters system provides the high processing throughput to run the smaller jobs in sufficient number to generate the required candidates. It also provides the computational power needed to efficiently run simulations requiring thousands of nodes. Beyond raw computational power, the balance of high system memory, fast inter-process communication, and **high-performance** disk I/O provided the capability to minimize scaling bottlenecks that often arise with calculations on hundreds of thousands of cores. The Blue Waters staff was responsive to requests. They were able to supply information on future software capabilities that was helpful in the development process.

NEXT GENERATION WORK

On a next-generation Track-1 system, this work could be expanded to enable finer resolution of micro-structures. This could allow identification of smaller scale features. Improvements in the simulation software's design to exploit next-generation computational architecture provides an exciting opportunity to **improve efficiency**.

PUBLICATIONS AND DATA SETS

Ackerman, D., and B. Ganapathysubramanian. Parallel framework for wormlike chains using self consistent field theory. *ASME 2015 Applied Mechanics and Materials Conference*, Seattle, WA. July 1, 2015.

VALIDATION OF MULTIPHYSICS CAPABILITIES OF THE MASSIVELY PARALLEL FINITE ELEMENT CODE ALYA

Allocation: Industry/60.0 Knh
PI: Seid Koric[1]
Co-PIs: Ahmed Taha[1], Mariano Vasquez[2], Guillaume Houzeaux[2], Eva Casoni[2], and Daniel Mira[2]
Collaborators: Paula Cordoba[2], Adriá Quintanas[3], Madhu Vellakal[1], Erman Gulyeruz[1], and Qiyue Lu[1]

[1]University of Illinois at Urbana-Champaign
[2]Barcelona Supercomputing Center
[3]Universitat de Girona

EXECUTIVE SUMMARY

Demonstrating the exceptional level of parallel scalability of the massively parallel multiphysics finite element code Alya from the Barcelona Supercomputing Center (BSC) [1] on the petascale architecture of Blue Waters was one of the major global high-performance computing (HPC) **breakthroughs** in 2014. Making sure that Alya is modelling physics accurately is the next step achieved in this project, which is of high importance for its users. NCSA's Private Sector Program (PSP) team worked jointly with BSC's Alya team comparing and validating the multiphysics capabilities of Alya with the experiments or the corresponding models built in the mature commercial codes that have been validated and in use at NCSA for decades by both academia and industry. The validation examples range from turbulence models coupled with combustion in real industrial geometries to nonlinear deformation in solids. We have found that Alya is capable of producing accurate and credible results consistent with the experiments and the corresponding results from the independent software vendor (ISV) codes.

FIGURE 1: LES simulation of the reacting flow field in a swirling combustor for two regimes: premixed and non-premixed conditions. The plot shows an iso-surface of stoichiometric mixture fraction colored by temperature (top-left), temperature field (bottom-left) and comparison with experiments (right).

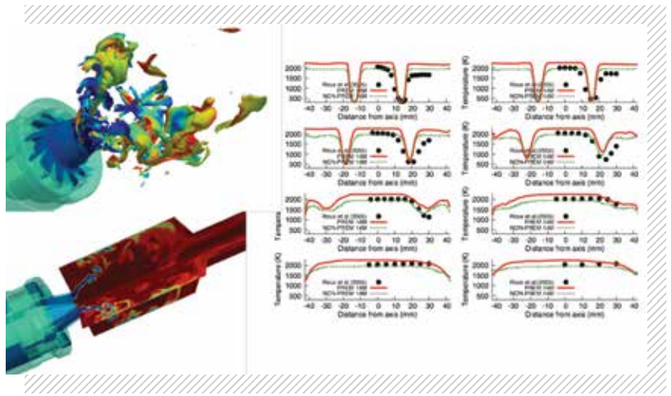

INTRODUCTION

While emerging petascale computing is already a strategic enabler of large-scale simulations in many scientific areas such as astronomy, biology, material science, and chemistry, paradoxically for many engineering researchers, the existing hardware and software often cannot be used to solve their problems. On one hand, current HPC systems in production often lack the computational power, network bandwidth and data storage needed for solving tomorrow's real-world engineering challenges. On the other hand, even the most powerful hardware will fail to deliver on its full potential unless matched with appropriate algorithms designed specifically for such environments.

The common belief that engineering simulation codes do not scale efficiently in large supercomputers was proven false in 2014 using NCSA's Blue Waters. This breakthrough occurred when NCSA worked with BSC to scale Alya to the unprecedented level of 100,000 cores [1], simulating complex multiphysics problems such as airflow in the human body, contraction of the heart, and combustion in a kiln furnace. The previous scaling work coupled with the validation in this work, and accompanied by the growth in petascale computing, will pave the way to higher fidelity and complexity simulation studies in many fields of engineering, achieving drastically faster design cycles, fewer costly physical models, and safer, lighter and more efficient final products.

METHODS & RESULTS

Two fundamental problems of the aerospace sector were solved using the multiphysics code Alya:

Combustion

The design of modern combustion systems addresses challenges in aspects related to the reduction of pollutant emissions, increments in fuel flexibility and avoiding thermo-acoustic instabilities. To achieve these goals, numerical simulations are increasingly important during the design process as they provide detailed insights into the physical processes at a relatively low cost. The combustion process is modeled here using a thermochemical database using premixed flamelets. This approach allows the addition of detailed finite rate chemical kinetics into computational fluid dynamics and to takes into account turbulence/chemistry interactions without accounting for transport equations for all the species involved in the combustion process at all time scales [2]. Therefore, the stiffness of the problem is reduced and permits us to obtain accurate predictions of radicals and products for different combustion regimes. The code Alya has been applied to several industrial configurations, and sample results of a partially premixed swirling combustor located at DLR (German Aerospace Center) along with some validation curves are shown here.

Buckling and post-buckling analysis of a full fuselage barrel

Since the beginning of the aerospace industry, one of the main objectives in the structural analysis field has been to determine a panel's buckling and post-buckling capacity, trying to assure the structural integrity up to ultimate load, which is a critical factor in certification. A highly dynamic and non-linear buckling and post-buckling analysis of a full fuselage barrel loaded in uniaxial compression is conducted. The problem has been validated within the European project SHERLOC (CS2-AIR-GAM-2014-2015-01). A mesh of 4 million 3D solid elements with high-aspect ratio was generated to perform the finite element analysis using Alya [3].

The load-displacement curve is shown in Figure 2. The onset of buckling occurs at around 7mm of longitudinal displacement. After buckling takes place, the slope of the load-displacement changes. As the displacement increases different buckling modes start to appear: at axial displacement of 7mm the radial displacement plot clearly shows a buckled shape consisting of several half waves along the length of the bay. As the load increases the number of half waves increases indicating mode jumps have occurred. During post-buckling, as the local buckling in the skin and stringers occur, the distribution of stresses changes dramatically. The stresses in the skin show concentrations at the intersection with the stringers and stress concentrations around the mouse holes in the frames are still high, as can be seen in the right figures.

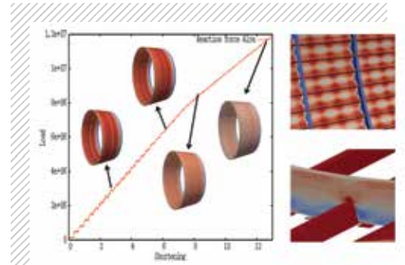

FIGURE 2: Load-shortening curve and Von-Misses stresses in post-buckling analysis.

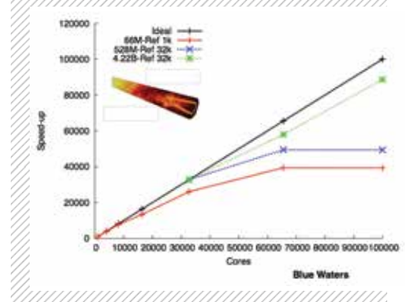

FIGURE 3: Scalability curve for the reacting flow of a rotary kiln using the parallel multiphysics code Alya [1].

WHY BLUE WATERS

Blue Waters is the only system that allows massively parallel multiphysics codes, such as Alya, to be tested by taking full advantage of large amounts of distributed memory, hundreds of thousands of computing cores, and low latencies and increased bandwidth of leading interconnect network technologies. In this context, Blue Waters was used to evaluate the parallel performance of Alya for several multiphysics applications showing good scalability up to 100,000 cores [1]. In particular, a combustion problem was tested on 4.3 billion elements on 100,000 cores (Fig. 3).

Other government-funded HPC resources worldwide are either unreachable for this kind of international collaboration, or their extent and specific hardware design lacks the capability to

support the size and complexity of multiphysics calculations performed in this project.

NEXT GENERATION WORK

With almost 90% parallel efficiency on 100,000 XE6 cores, Alya has a potential to scale to a million cores on the next Track-1 system. Among many exciting possibilities, this would allow direct numerical simulation of turbulence on larger fluid domains, or directl modeling of meso-scale phenomena in many materials, and thus provide great validation tools for today's approximate continuum research methods such as large eddy simulation turbulence or multiscale methods.

PUBLICATIONS

Arnold, G., M. Gajbe, S. Koric, and J. Urbanic. XSEDE OpenACC workshop enables Blue Waters Researchers to Accelerate Key Algorithms, *XSEDE 2014*, Atlanta, GA, July 2014.

Barabash, R.I., et al., Finite element simulation and X-ray microdiffraction study of strain partitioning in a layered nanocomposite, *J. Crystallography*, In Press, 2016.

Giles, M., S. Koric, E. Burness, U. S. Industrial Supercomputing at NCSA, Chapter in: A. Osseyran, M. Giles editors. Industrial Applications of High-Performance Computing, CRC Press, 2015, pp.179-194.

Gimelshein, N. E., A. Gupta, S. C. Rennich, and S. Koric, GPU acceleration of WSMP. *GPU Technology Conference 2015*, San Jose, CA, March 17–20, 2015.

Gupta, A., N. Gimelshein, S. Koric, and S. Rennich, Effective minimally-invasive GPU acceleration of distributed sparse matrix factorization, *Euro-Par 2016*, Grenoble, France, August 24-16, 2016.

Kale, S., S. Koric, and M. Ostoja-Starzewski, Stochastic continuum damage mechanics using spring lattice models. *Ap. Mech. Mat.*, 784(2015), pp. 350-357.

Kale, S., A. Saharan, S. Koric, and M. Ostoja-Starzewski, Scaling and bounds in thermal conductivity of planar Gaussian correlated microstructures. *J. Ap. Phys.* 117 (2015), 104301-1.

Koric, S. and A. Gupta, Sparse matrix factorization in the implicit finite element method on petascale architecture, *Arch. Comput. Method. E.*, 32(2016), pp. 281-292.

Koric, S., F. Sabet, O. Jin, and I. Jasiuk, Direct numerical simulation of bone plasticity and strength, *International Symposium on Plasticity 2016 and Its Current Applications*, Keauhou Bay, HI, January, 2016.

Koric, S., A. Murali, and B. G. Thomas, Elasto visco-plastic model of steel solidification with local damage and failure, *International Symposium on Plasticity 2015 and Its Current Applications*, Montego Bay, Jamaica, January 2015.

Koric, S., and B. G. Thomas, Visco-plasic multi-physics modeling of steel solidification, *Plasticity Symposium 2014*, Freeport, BA, January 3-8, 2014.

Koric, S, Engineering breakthroughs at NCSA, *4th International Industrial Supercomputing Workshop, Amsterdam*, The Netherlands, October 23-24, 2013

Koric, S., B. G. Thomas, and L. C. Hibbeler, Multiphysics modeling of metal solidification processes with abaqus, *Simulia Community Conference Proceedings*, Vienna, Austria, May 22-24, 2013.

Kwack, J., G. H Bauer, and S. Koric, Performance test of parallel linear equation solvers on Blue Waters – Cray XE6/XK system, *Cray User Conference (CUG)*, London, UK, May 8-12, 2016.

Puzyrev, V., S. Koric and S. Wilkin, Evaluation of parallel direct sparse linear solvers in electromagnetic geophysical problems, *Comput. Geosci.*, 89(2016), pp. 79-87.

Puzyrev, V., and S. Koric, 3D multi-source CSEM simulations: feasibility and comparison of parallel direct solvers; Society of Exploration Geophysicists (SEG) Technical Program Expanded Abstracts 2015, New Orleans, LA, October 2015.

Torlak, M., A. Halac, S. Roesler and S. Koric, DES of turbulent flow around a sphere with trip wire, *Proceedings of the Eight International Symposium On Turbulence, Heat and Mass Transfer*, Sarajevo, Bosnia and Herzegovina, October 2015.

Wen, H., S. Koric, Y. Xin, K. Hsia and L. Xiuling, Precision structural engineering of self-rolled-up 3D nanomembranes guided by transient quasi-static FEM modeling. *Nano Lett.*, 14:9 (2014), pp. 6293-6297.

TRANSFORMATIVE PETASCALE PARTICLE-IN-CELL SIMULATIONS OF HIGH ENERGY DENSITY PLASMAS

Allocation: NSF PRAC/3.18 Mnh
PI: Warren B. Mori[1]
Co-PI: Frank Tsung[1]
Collaborator: Luis Silva[2]

[1]University of California at Los Angeles
[2]Instituto Superior Técnico

EXECUTIVE SUMMARY

The goal of this project is to use state-of-the-art simulation tools to unravel the complex physics that is inherent in high-energy density plasma physics. The simulations performed on Blue Waters are helping to answer important questions related to developing new accelerator technology that could be the basis of a next generation linear collider or a highly compact coherent x-ray source, and to successfully achieving inertial confinement fusion in the laboratory.

INTRODUCTION

Our current focus is to unravel key physics within two areas of high energy density plasma physics. Particle accelerators at the energy frontier are the most complex tools built and have been the tool for discovery in high-energy physics for nearly the past 100 years. The newest of these machines is the Large Hadron Collider at CERN, which is 30km in circumference and collides 7 TeV proton beams together that cost billions of dollars. The cost and size of these machines makes it unlikely that a next-generation machine will be built in the foreseeable future.

One option to continue further investigations with existing machines is to use relativistically moving plasma waves as the accelerating structure, or plasma-based acceleration. The space charge force of a relativistic charged particle beam or the radiation pressure of a laser can excite these waves as wakefields behind the beam or laser. These wakefields can support acceleration gradients more than three orders of magnitude larger than current technology. These plasma wave wakefields can also be the basis for building a compact coherent x-ray source that could fit in universities. We are

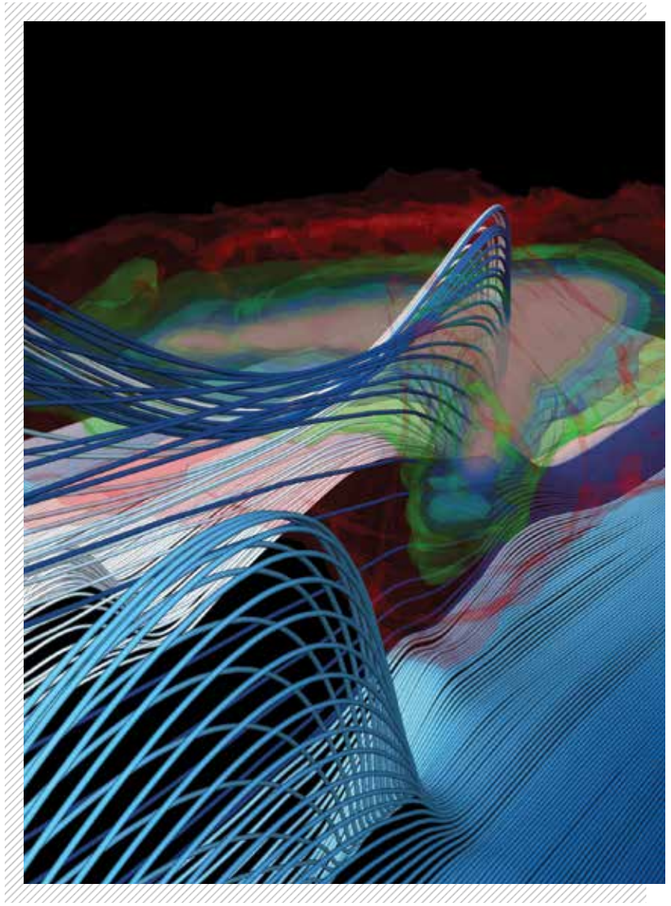

FIGURE 1: QuickPIC simulations of positron accelerations using plasmas. The figure shows isosurfaces of plasma density and positron trajectories (as blue lines).

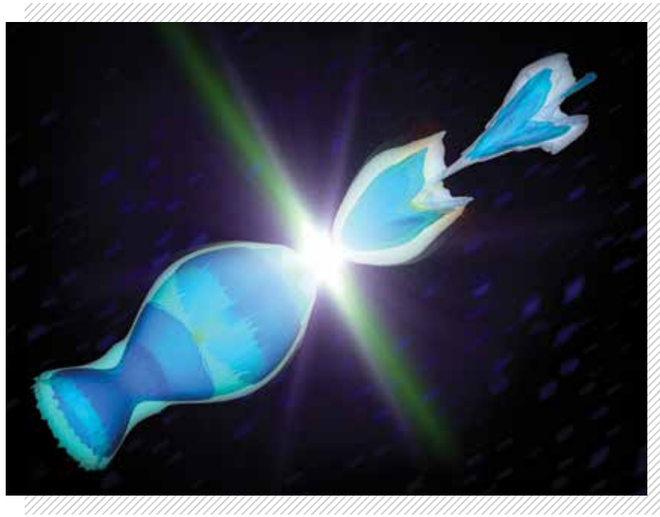

FIGURE 2: QuickPIC simulations of an electron and positron collider using plasmas. In the figure, electrons are moving right left (from the lower left corner of the plot) and positrons are moving left (from the upper right corner). The two bunches meet in the center and collides.

using Blue Waters and highly accurate computer simulation software to study the physics that need to be understood to enable a collider or a compact coherent x-ray source to be designed and built based on plasma-based acceleration.

Controlled fusion energy offers the possibility of an unlimited supply of relatively clean energy. In laser-driven inertial fusion energy, lasers either directly or indirectly drive the spherical implosion of a small amount of deuterium and tritium to densities more than 1000 times solid density and to temperatures exceeding 106 °K. This requires the implosion be very symmetric and that the lasers hit the target accurately. The lasers need to propagate through high energy density plasma where they are susceptible to a host of "instabilities" that can absorb, reflect, and bend them before they hit their target. We use highly accurate models to understand laser-plasma interactions in high energy density plasmas. The goal is to use this understanding to develop ways to control and eliminate the deleterious instabilities.

METHODS & RESULTS

In the past 18 months our major results include:

- QuickPIC simulations of an experiment at SLAC show that a single positron beam can evolve into a self-loaded configuration leading to some positrons forming a monoenergetic tail. This work may point towards developing beam loading configurations for the positron arm of a future linear collider based on plasma wave wakefields [1],
- Highly resolved QuickPIC Simulations that show ion motion within the trailing bunch does not necessarily lead to catastrophic emittance (which measures the angular divergence of the beam) growth for plasma based acceleration based linear collider designs. Simulations using Blue Waters indicate that there are fully self-consistent beam loading scenarios for the electron arm for a future linear collider based on plasma wave wakefields,
- OSIRIS simulation of injection schemes to produce high brightness beams with unique characteristics. Plasma wave wakefields may also be a component in a compact coherent x-ray source. Using Blue Waters we have found that using ionization injection or down ramp transition region in front of an accelerator section permits the controlled injection of kA of higher current of electron beams with very small angular divergences and energy spreads. These beams can have three orders of magnitude higher brightnesses than conventional electron beam sources. Also, the macrobunches can be prebunched on nanometer scales,
- Quasi-3D OSIRIS simulations using Blue Waters allowed us to explore laser wakefield accelerators (LWFA) in the self-guided nonlinear blowout regime for current high power lasers. We found by using a 30 Joule laser it is possible to generate 8 GeV electron beams without the need for plasma channels [2],
- 3D OSIRIS simulations of Lawrence Livermore LWFA experiments in the self-modulated regime [3]. Using Blue Waters we carried out fully resolved 3D simulations of a ~pico-second class laser undergoing self-modulation. These lead to a broad spectrum of hundreds of MeV electrons from a combination of acceleration in the plasma waves and laser fields (including backscattered or reflected laser light). These electrons radiate a broad spectrum of incoherent x-rays. This scheme could lead to the development of a directional, small-divergence, and short-duration picosecond x-ray probe beam with an energy greater than 50 keV, which is desirable for high energy density science experiments.

WHY BLUE WATERS

Blue Waters provides a unique platform that enables us to perform many large scale simulations and support ongoing experiments promptly and produce works which appear in journals such as *Nature* and *Physical Review Letters*.

NEXT GENERATION WORK

Future track-1 supercomputers will most likely build on GPU or Intel Phi based hardwares with many levels of parallelism, and the UCLA simulation group have been very actively porting our codes to these architectures. In the past few years we have developed algorithms which can be optimized on generic many core architectures, including GPU's and Intel Phi's. Adding these algorithms to our production codes will allow us to investigate physics relevant to plasma-based accelerators and laser fusion. With the increased memory and processing speed of future track-1 supercomputers, we plan to investigate:

- Beam loading scenarios for linear collider designs based on plasma wakefields, including Ion effects in plasma wakefield accelerators (PWFA) with very narrow beams,
- LWFA's and PWFA's with high brightness beams,
- The long time behavior of relativistic shocks and their generation of energetic particles,
- Laser plasma interactions relevant to inertial fusion energy in two and three dimensions.

PUBLICATIONS AND DATA SETS

An W., et al., QuickPIC Simulation of Plasma Wake Field Accelerator Experiments at FACET, *2015 International Conference on Numerical Simulation of Plasmas*. August 12-14, 2015, Golden, CO.

An W., et al., Simulations of Plasma Wakefield Acceleration at FACET and Beyond, *2015 Computational Accelerator Physics Conference.* October 12-17, 2015, Shanghai, China.

Corde, S., et al., Multi-gigaelectronvolt acceleration of positrons in a self-loaded plasma wakefield. *Nature*, 524 (2015), p. 442.

Mori, W.B., Synergy Between Experiments and Simulations in Laser and Beam-Driven Plasma Acceleration and Light Sources. *Division of Plasma Physics APS 2015*, November 16-20, 2015 Savannah, GA.

Pollock, B.B. et al., Formation of Ultrarelativistic Electron Rings from a Laser-Wakefield Accelerator, *Phys. Rev. Lett.* , 115 (2015), 055004.

Tsung, F.S., et al, Controlling Laser Plasma Interactions with Temporal Bandwidths, *2016 Anomalous Absorption Conference*, Old Saybrook, CT, May, 2016.

Xu, X., et al., Phase Space Matching of Staging Plasma and Traditional Accelerator Using Longitudinally Tailored Plasma Structure. *6th International Particle Accelerator Conference*, May 4th - 8th, Richmond, VA, USA

Xu, X., et al., Phase Space Dynamics in Plasma Based Wakefield Acceleration, *2015 CPS Fall Meeting, Division of Plasma Physics*, Sep. 11th - 13th, Changchun, Jilin, China

Xu, X., et al.,, et al., Lorentz boosted frame simulation of laser wakefield acceleration in 3D and quasi-3D geometry, *The International Computational Accelerator Physics Conference 2015*, Oct. 12th - 16th, Shanghai, China

Xu, X., et al.,, et al., Exact Phase Space Matching for Plasma and Traditional Acceleration Components using Longitudinally Tailored Plasma Profiles / Prospects for generating high brightness and low energy spread electron beams through self-injection schemes, *FACET Workshop*, Oct. 12th - 16th, San Francisco, CA, USA

Xu, X., et al.,, Intrinsic phase space discretization of charge in laser-triggered ionization injection, *57th Annual Meeting of the APS Division of Plasma Physics*, Oct. 15th - 21th, Savannah, GA.

NON-ADIABATIC ULTRAFAST ELECTRON-ION DYNAMICS NEAR ALUMINUM SURFACES

Allocation: Blue Waters Professor/240 Knh
PI: André Schleife[1]
Co-PI: Alina Kononov

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Computational physics and materials research greatly benefits from high-performance computing, since modern first-principles approaches provide insight with unprecedented precision. Currently, highly accurate simulations of real-time quantum dynamics are coming within reach, allowing researchers to overcome the Born-Oppenheimer approximation systematically. We explore the feasibility of Ehrenfest molecular dynamics, based on real-time time-dependent density functional theory, to investigate secondary electrons near surfaces of aluminum. We study the dependence of electronic stopping and secondary-electron emission on the velocity of fast projectile ions. From non-adiabatic electron-ion dynamics emerging after the impact, we gain detailed insight into the materials physics on an attosecond (10-18 secs) time scale. These simulations allow us to determine precisely the transition from "surface-like" to "bulk-like" systems. While large simulation cells and short simulation time steps are required, Blue Waters enables these simulations and allows us to study complex quantum dynamics quantitatively.

INTRODUCTION

Fast particles entering a material with kinetic energies on the order of keV or MeV produce effects on various length and time scales. At early stages, deposition of energy into the electronic system of the target is the dominating mechanism. It is mediated by the electron-ion interaction and leads to excitations of the target's electronic system. Details of these processes are not well understood: Existing theoretical models require ad-hoc assumptions and rely on separate theories, e.g. for the charge state of the projectile. Early models suffer from large uncertainties and are not suitable for the computational materials design needed for novel technological applications.

Scenarios of practical interest include fusion reactors, where highly energetic particles interact with the reactor wall. At initial stages, the vast majority of energy is deposited into the electronic system of the host. Near the surface, excited electrons can overcome the work function of the material and escape into the reactor. These secondary electrons detrimentally affect the plasma properties. Another example of paramount importance for technological applications is modern helium microscopy [1]; secondary electrons emitted from the sample are collected and create image contrast. Helium microscopy is an appealing technique, especially for insulating uncoated samples. A better understanding of the dependence of secondary-electron emission on the velocity and impact angle of the projectile is necessary. Accurate first-principles techniques that are capable of describing these effects are, thus, desirable for a large community of researchers.

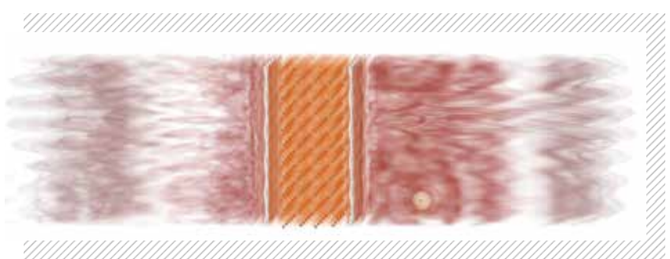

FIGURE 1: Aluminum slab of 200 atoms after being hit by a fast hydrogen projectile. Electrons are localized around the hydrogen projectile and secondary electrons are ejected into the vacuum.

METHODS & RESULTS

In this project we explore Ehrenfest molecular dynamics: we use our new large-scale implementation within the Qbox/Qb@ll code [2], based on real-time time-dependent density functional theory (RT-TDDFT). Nuclei are treated as classical point charges coupled to electrons via Coulomb interaction and electrons are treated quantum-mechanically. The electron-ion interaction is described using non-local pseudo-potentials and the electron-electron interaction is a sum of Hartree and exchange-correlation potential. Kohn-Sham wave functions are expanded into plane waves, and we also use the common adiabatic local-density approximation.

Using this technique we investigated the transition from "surface like" to "bulk like": We varied the number of atoms that constitute an aluminum slab between 100 and 500 (Fig. 1) and computed position-dependent electronic stopping. Figure 2 shows that immediately after entering the material, the projectile (velocity = 1.0 at. u.) experiences larger stopping, which we attribute to the fact that its charge has not equilibrated after entering the slab. Once the projectile traveled about 20 aB, equilibration is observed as the electronic stopping approaches the bulk value [3]. At this point, the projectile attracted some of the host electrons, reducing its initial fully ionized state. Our simulations allow extracting length and time scales of this equilibration, and we can reconstruct the effective charge state of the projectile as it travels through the material.

By fitting the charge-density distribution around the projectile to hydrogen orbitals, we compute the number of electrons around the projectile. For a velocity of v = 1.0 at. u., we found that about 0.3 electrons move with the projectile after it leaves the material. We are analyzing this data for different vacuum lengths, numbers of slab atoms, and projectile velocities to develop a comprehensive picture. We also extract the total number of emitted secondary electrons. These simulations directly provide insight into the underlying physics and allow us to extract quantitative information that can be used to parameterize larger-scale numerical models. Based on our accurate, quantitative data we envision multi-scale computational design of materials under radiation conditions.

WHY BLUE WATERS

First-principles Ehrenfest simulations come with high computational cost and are only possible due to massively parallel implementations of this technique and using high-performance supercomputers with low-latency communication. Thousands of time steps (as short as 0.1 atto-seconds) are needed, and many simulations are required to study vacuum size, numbers of slab atoms, and projectile velocities. Since our code scales very well to hundreds of thousands of cores on Blue Waters, we make efficient use of this machine when running for many hours on hundreds of nodes. Through the Blue Waters team, we have excellent access to visualization and data analysis experts with whom we discussed interpretation and analysis of the vast amount of data (hundreds of megabyte per time step).

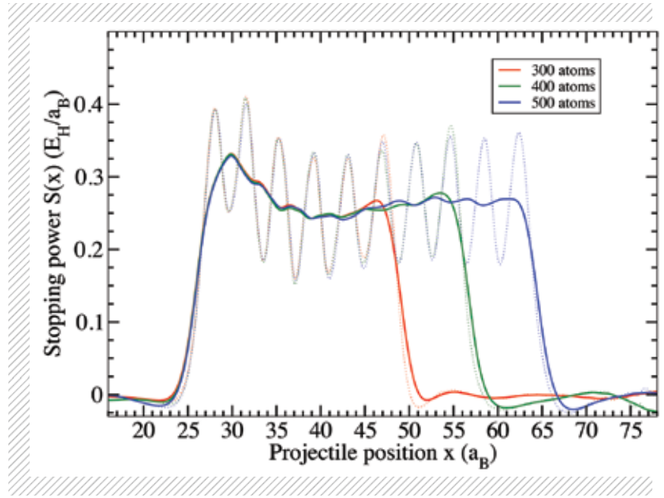

FIGURE 2: Instantaneous electronic stopping of a hydrogen projectile in aluminum slabs of different thickness.

NEXT GENERATION WORK

From our studies, we found that including the effect of excited electrons in first-principles molecular-dynamics simulations is critically important in many cases of practical relevance. Ehrenfest molecular dynamics can be a practical first step towards achieving this. However, the computational cost is high. Using Blue Waters, we can study systems of practical importance. However, we envision that next-generation Track-1 systems will allow us to explore larger length and time scales. They will inevitably be needed to achieve our vision of routinely using quantum-dynamics simulations and of going beyond Ehrenfest molecular dynamics for computational materials design.

KINETIC MODELING AND SIMULATION OF HYPERSONIC, SHOCK-BOUNDARY LAYER INTERACTIONS USING PETASCALE COMPUTING

Allocation: Illinois/70.0 Knh
PI: Deborah A. Levin[1]

[1]University of Illinois at Urbana Champaign

EXECUTIVE SUMMARY

The modeling of shock-boundary layer interactions of three-dimensional hypersonic flows using kinetic, particle methods such as direct simulation Monte Carlo (DSMC), provides the highest fidelity in understanding thermochemical non-equilibrium processes when extremely complex shock interactions are present. We have developed an Octree based (message passing interface) MPI-parallelized code that takes advantage of Adaptive Mesh Refinement (AMR) techniques to maximize the placement of computational particles in flow regions where the collision frequency is highest. Such continuum-like conditions create massive computational loads in simulating several billions of particles. To perform such continuum-like computations, it was important that the code have superior algorithmic efficiency amenable to high scalability. This code required the implementation of Morton-Z space filling curve that facilitates a direct access of leaf cells in an Octree and careful attention to efficient use of the cache.

FIGURE 1: Need for linearization for hypersonic modeling of shock-boundary layer interactions.

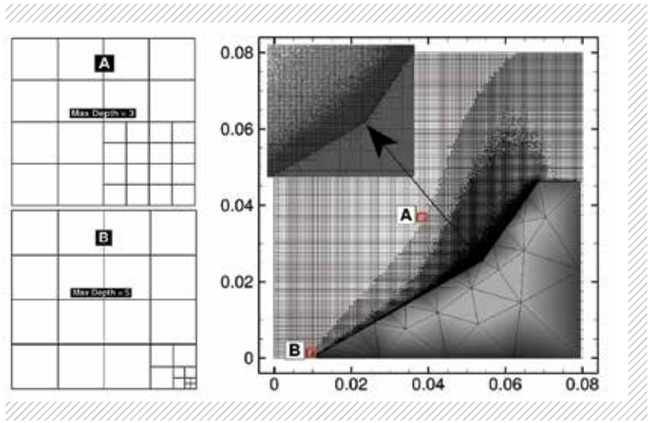

INTRODUCTION

Hypersonic flow over a double wedge configuration at continuum-like free stream conditions has been a challenging problem because of the multiple shock-shock and shock-boundary layer interaction, separated flows near the hinge, sheer layer, and three-dimensional effects. These conditions generate a mesh that is highly non-uniform because of very high levels of refinement near the wedge as shown in Fig. 1. It can be seen that the dense mesh obtained after the shock just near the hinge for the aforementioned double wedge case is highly non-uniform because of the high gradients. The Octree cells lying near the surface are highly refined as compared to those in the free stream and inside the geometry. Therefore, these cells have to perform a lot of work while others wait idly at the end of each time step, which is unfavorable as the full capacity of all the processors is not utilized and we essentially see a high degree of imbalance of the load among the processors. Furthermore, at high computational loads, the communication time starts to increase.

METHODS & RESULTS

The main tasks in any DSMC code are (1) the movement of particles, (2) their mapping to the correct computational cell after movement, (3) sampling or the calculation of the particle macroscopic properties, and (4) performing binary collisions. Note that; the DSMC algorithm is efficient because it is assumed that particle movement and collisions may be decoupled. The earlier version of our code required recursive access of particle data by linked-lists to accommodate the dynamic changes in the number of particles and the number of computational cells. As a result, it took three times longer, per particle, than a Fortran 77 DSMC code. Upon close examination, it was found that both codes took the same amount of time to perform collisions, but the main difference was due to the recursive pointer based access of cells in an AMR-based unstructured code. To avoid this cost but retain the flexibility of the unstructured mesh, a Morton-Z space filling curve that facilitates direct access to leaf cells in an Octree was implemented. In addition, the entire data structure of the SUGAR code was replaced by arrays instead of a pointer-based system, which now ensured that the code gave superior performance, per processor.

After achieving the optimum algorithmic efficiency along with preserving the unstructured nature of the mesh, the scalability of the code was improved. Figure 2 shows the speed up plot for the strong scaling in comparison with the earlier version of the SUGAR code. The earlier version of the code was partitioned using a simple 2-D blocking algorithm where the division of the domain was done at the root (Octree) level. It was further coupled with the state-of-the-art graph partitioning tool Zoltan in an attempt to obtain better load balancing and reduce the communication between the processors. The speed up curve for a case of the flow of argon over a hemisphere using 2-D blocking algorithm and Scotch graph partitioner is shown in the figure. It can be seen that the maximum speed up by a just factor of four was observed with a 16 times increase in the number of ranks, and the use of Scotch made no difference. However, after many improvements, the current version of the SUGAR code shows a near-ideal speed up by a factor of 64. This increase in speed was a result of the following improvements: 1) The domain was portioned now at the leaf level, now possible with the implementation of space filling curves. This also gave better flexibility in balancing the computational loads. 2) Instead of manually putting the weighting criterion based on the number of particles, cells, or geometry panels in a leaf cell, the code was modified to take into account the time taken to perform each of these tasks after the steady state is reached by profiling these times, it comes up with an optimum weighting criterion during the runtime. 3) Attempts to reduce the communication among neighboring processors by creating a list of possible communication pairs resulted in a tremendous improvement in strong scaling. 4) The introduction of particles into the computational domain at the inlet was parallelized, and efforts for further modifications are ongoing. 5) More efficient use of memory. Regarding weak scaling, the SUGAR code can simulate 3 billion particles with 90% parallel efficiency and the tests for a higher number of processors is ongoing.

WHY BLUE WATERS

The main advantage of Blue Waters is its massive number of nodes. A single actual low Knudsen number double-wedge case takes approximately 10,500 node-hours.

NEXT GENERATION WORK

At present, the SUGAR solver can simulate hypersonic neutral and charged particle flows through highly irregular porous geometries and efforts are directed to include chemistry and radiation. We envision new architectures with considerably more memory per CPU to enable us to achieve the levels of parallelization required to solve unsteady, transitional flows.

PUBLICATIONS AND DATA SETS

Sawant, S., R. et al., Study of Shock-Shock interactions Using an Unstructured AMR Octree DSMC Code, AIAA SciTech 2016, January 4th 2016, Paper No. 2016-0501

Jambunathan, R. and D. Levin, Advanced Parallelization Strategies for Modeling Flow Through Ablative Thermal Protection Systems, AIAA Paper No. 2016-3387, (2016), 46th, AIAA Thermophysics Conference, 10.2514/6.2016-3387.

FIGURE 2: Strong scaling performance of the improved SUGAR DSMC code.

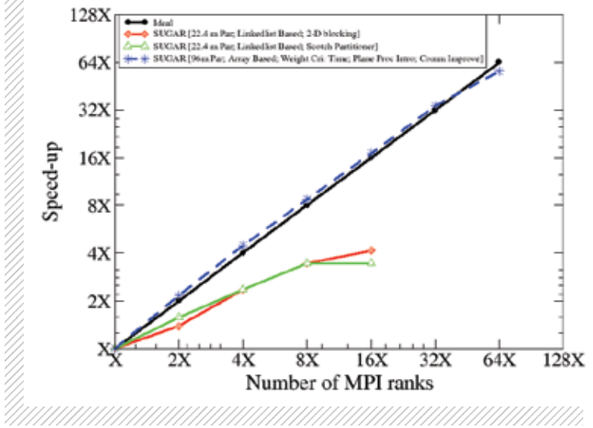

DIRECT NUMERICAL SIMULATION OF FULLY RESOLVED DROPLETS IN A TURBULENT FLOW

Allocation: NSF PRAC/2.35 Mnh
PI: Said Elghobashi[1]
Collaborator: Michele Rosso[1]

[1]University of California, Irvine

EXECUTIVE SUMMARY

The objective of our numerical study is to enhance the understanding of liquid droplet vaporization and mixing processes in a turbulent flow. The study employs direct numerical simulations (DNS) to examine the two-way interactions between freely-moving vaporizing droplets and isotropic turbulence. The droplets are fully resolved in 3D space and time (i.e., not treated as point particles), and all the scales of the turbulent motion are resolved down to the smallest relevant length- and time-scales (the Kolmogorov scales). Emphasis is placed on the two-way exchange of mass, momentum and energy between the vaporizing droplets and the surrounding turbulent gas. We simulated the motion of 1000 liquid droplets in isotropic turbulent flow to study their dispersion and compare it to the dispersion of solid particles. The goal is to understand the effect of the surface tension of the droplets on their dispersion.

INTRODUCTION

All liquid fuel combustion devices, mobile or stationary, use atomizers to produce sprays of fine droplets. The fuel droplets must first vaporize before they mix with the surrounding air. This mixing is followed by a chemical reaction between the fuel vapor and air, which converts the chemical bonding energy into thermal energy resulting in the volumetric expansion of the gas mixture. This volumetric expansion produces the desired mechanical energy for rotary or reciprocating engines.

Understanding the physical details of the droplet vaporization and mixing processes in a turbulent flow is an essential prerequisite to understanding the chemical reaction process and the eventual control/optimization of the energy conversion process. The results of the proposed study will have a significant impact on the efficient utilization of energy. This impact stems from the fact that the vaporization rate is the main controlling mechanism of fuel droplet combustion and that liquid fuels are the primary source of energy for all modes of transportation and will remain as such for the foreseeable future.

METHODS & RESULTS

Our numerical procedure solves the discretized incompressible Navier-Stokes and continuity equations in the liquid and gas phases with appropriate jump conditions at the interface between the two phases. To implicitly capture the interface between the two phases, we employ the accurate conservative level set method. The numerical solution of the discretized equations uses the conjugate gradient method preconditioned by a V-cycle geometric multigrid (GMG) solver. To understand the effects of surface tension on the droplets' motion, we compared the dispersion characteristics of finite size liquid droplets and finite size solid particles in isotropic turbulence at moderate values of Reynolds numbers, in zero gravity. The droplets and particles have equal diameters (larger than the Kolmogorov length scale) and equal ratios of particle (droplet) density to the carrier fluid density. The level set method is used for DNS of the droplets where a variable-density projection method is used to impose the incompressibility constraint. The immersed boundary method is used for DNS of the solid particles. Our results show that in isotropic turbulence, the dispersion of liquid droplets in a given direction is larger than that of solid particles due to the reduced decay rate of turbulence kinetic energy via the four-way coupling effects of the droplets.

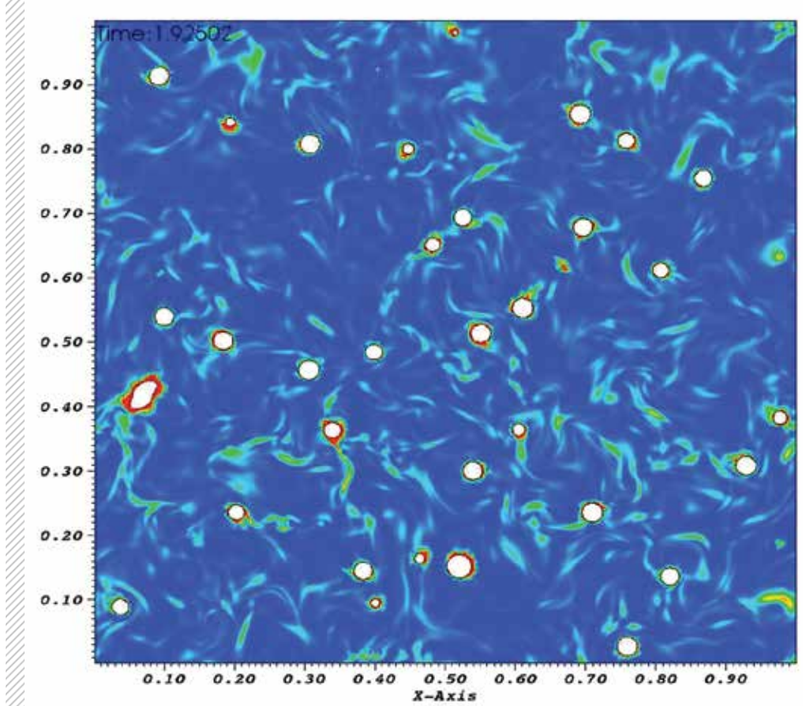

FIGURE 1: Contours of the dissipation rate of turbulence kinetic energy around non-vaporizing droplets in isotropic turbulence. Red and blue contours indicate maximum and minimum dissipation rates, respectively.

WHY BLUE WATERS

Blue Waters is indispensable for the DNS of a droplet-laden turbulent flow since the carrier flow is time-dependent, three-dimensional and contains a wide spectrum of length- and timescales. Furthermore, the interface between the liquid and gas phases must be resolved in time and space as the droplets move, including their shape change and possible droplet-droplet collision and merging. The Blue Waters staff is essential for the success of our study. Their continuous and prompt assistance is greatly appreciated.

NEXT GENERATION WORK

Our next generation work will emphasize the two-way exchange of mass, momentum, and energy between the vaporizing droplets and the surrounding turbulent gas. We plan to study droplet-laden turbulent shear flows after understanding droplet-laden isotropic turbulence.

PUBLICATIONS AND DATA SETS

Rosso, M., H. Wang, and S. Elghobashi, Dispersion of finite size droplets and solid particles in isotropic turbulence. *ICMF-2016 – 9th International Conference on Multiphase Flow*, May 22nd – 27th 2016, Firenze, Italy.

FIRST PRINCIPLE AND MODELING OF TURBULENT TWO-PHASE FLOWS

Allocation: Blue Waters Professor/245 Knh
PI: Vincent Le Chenadec[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Aircraft gas turbines often rely on the combustion of high-energy-density liquid fuel to meet weight and volume restrictions. The efficiency of the conversion and the emission of harmful pollutants depend directly on the mixing of the fuel and oxidizer, which itself results from a cascade of mechanisms initiated by the atomization of a coherent liquid stream. In this context, the challenges met by experimental diagnostics are numerous and include the wide range of scales at play, as well as the lack of optical access. These shortcomings have spurred the development of numerical strategies to simulate and model the atomization process. Computational approaches can potentially provide invaluable insight into the effects of fluid properties and nozzle design, as well as control strategies. There are, however, some obstacles to reaching predictive simulation and modeling capabilities, and petascale systems such as Blue Waters are central to achieving this goal.

INTRODUCTION

The reliable prediction of primary atomization and the resultant spray population statistics are vital for the design of low-emissions and stable gas turbine combustors. Atomization of a liquid jet is controlled by unique small-scale physics that rely on the interaction between vorticity generation at the small scales and the interface motion, which in turn is controlled by surface tension forces. Although a critical sub-component of a liquid fueled combustor, the liquid spray breakup resulting from complex fuel injectors is poorly understood. Yet it has a significant influence on the engine performance in terms of emissions, fuel consumption, thermo-acoustic instabilities and durability. The modeling of liquid fuel injectors represents a formidable challenge. The main quantity of interest is the spray penetration, characterized by the droplet number density function that includes the distribution of droplet sizes and velocities. This information is subsequently used to design the primary reaction zone. For instance, insufficient atomization could result in fuel-rich pockets that promote soot formation. Similarly, the nature of atomization will also control ignition tendencies in high-altitude relight conditions. Hence, reliable computational models that predict the final spray properties are indispensable for the successful design of aircraft combustors.

METHODS & RESULTS

Improving the current understanding and the predictive capabilities of primary atomization requires progress in two distinct areas. The first is Direct Numerical Simulation (DNS), which consists of solving the equations that govern the mixing of the liquid and gas without any simplifying assumptions. Hence, it represents the highest achievable level of fidelity. This level of fidelity comes at a price, however, in that the governing equations are very stiff. Tailored (physics compatible) numerical algorithms are therefore required. Also, in the regime of interest, the governing equations (the two-phase Navier-Stokes equations) are chaotic, and feature a wide range of scales, which require the use of extremely large computational meshes. For the foreseeable future, it is, therefore, undeniable that DNS of full-scale industrial configurations will remain scarce. This introduces the second area that requires advances, namely the modeling of primary atomization. Models based on the Reynolds-Averaged Navier-Stokes or the Large Eddy Simulation formalisms are characterized by their reliance on additional modeling assumptions. They, of course, come with a significant reduction in computational cost, the trade-off being their lower fidelity. An additional advantage is their improved controllability over the local instantaneous formulation used in DNS, which is chaotic and therefore harder to optimize or control.

The first component of the project leverages recent and ongoing developments in the field of multiphase computational fluid dynamics, which are unlocking the predictive capabilities of DNS of turbulent two-phase flows, and hence enabling the exploration of the physics of atomization. Reaching the regimes of interest to practical applications, however, requires harnessing the performances of large-scale systems such as Blue Waters. This means, in particular, designing software that achieves good load balancing where conflicting algorithmic requirements meet, e.g. the local application of computational geometry (highly non-linear) in the interface vicinity vs. the global use of iterative linear solvers in implicit and differential algebraic equations. Emerging heterogeneous architectures, in particular the XK nodes on Blue Waters, therefore represent an unprecedented opportunity.

DNS therefore represents a unique tool to quantify the effects of different atomizer designs, or to perform parametric study over low dimensional spaces (given associated computational cost). But it can also provide valuable insight into the modeling of primary atomization. To date, documented studies in this area have mostly been limited to asymptotic limits, or to restrictive phenomenological or empirical hypothesis. Analysis of the very large data sets generated by the DNS software using statistical tools as well as decomposition techniques has proved useful to model single-phase turbulence and combustion, and is, therefore, the second component of the project. Given the volume of the data sets involved, as well as the specificities of the decomposition techniques, this analysis can only be achieved using the fast storage and large bandwidth achieved by the Blue Waters system.

WHY BLUE WATERS

Blue Waters is critical to not only perform the highly resolved simulations required to explore and model primary atomization in canonical flow configurations, but also to push the **frontiers** of current computational infrastructures by leveraging hybrid (i.e. shared and distributed) and heterogeneous (e.g. CPU and GPU) programming, in order to simulate practical engineering applications, and ultimately develop new nozzle designs.

COMPUTER SCIENCE & ENGINEERING

OPTIMIZATION

SCALABILITY

PARALLEL FUNCTIONALITY

COMPUTER VISION

VISUALIZATION

EFFICIENT, SCALABLE AND FAULT TOLERANT GENOMICS PIPELINE

ALLOCATION: Illinois/50.0 Knh
PI: Ravishankar K. Iyer[1]
Co-PI: Zbigniew Kalbarczyk[1]
Collaborators: Saurabh Jha[1], Subho Banerjee[1], Phuong Cao[1], Valerio Formicola[1], and Hao Jin[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Our overarching investigations use Blue Waters to address two important complex data-driven problems and propose new design and analysis enhancements. First, computational genomics: We have developed a software suite, IGen, for runtime optimizations of genomic workloads, including short-read alignment and human variant discovery, which are significantly faster on a single XE-6 compute node and scale out nearly linearly, leveraging available graphics processing units (GPUs) on XK-7 hybrid nodes (Fig. 1). Second, data-driven resilience at extreme scale: Working with Los Alamos National Laboratory, National Energy Research Scientific Computing Center, and CRAY, the project has developed a new design and assessment of resilience (reliability and security) of extreme-scale systems, using real data from Blue Waters failures and attacks. The tools use large-scale probabilistic graphs to drive machine learning at scale for runtime detection and mitigation of failures and attacks.

FIGURE 1: Comparing IGen's performance with unoptimized tools.

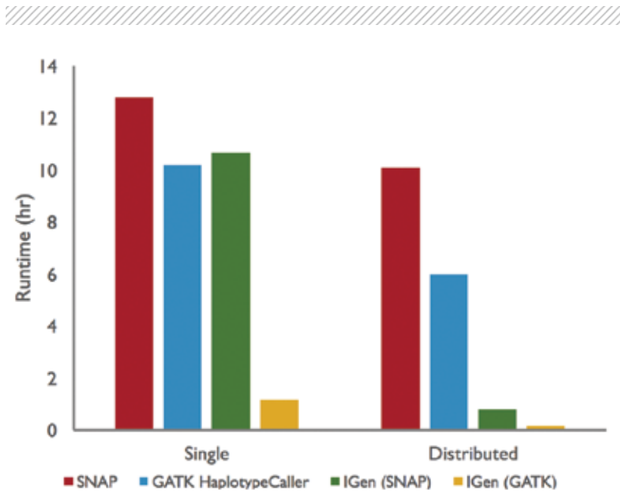

INTRODUCTION

Whole-genome sequencing and analysis is an increasingly important part of the standard of care in many hospitals and will continue to be in the years to come. The standard pipeline used for this purpose involves high computational and storage costs, which is a major hurdle for the routine implementation of genome-based individualized medicine. In this project, we have demonstrated the potential for runtime-level performance optimizations for a large number of computational-genomics workloads and built a prototype application called IGen for short-read alignment and human variant discovery. IGen is significantly (Fig. 1) faster on a single XE-6 compute node, scales out nearly linearly as a message passing interface (MPI) job, and leverages the GPUs available on the XK-7 hybrid nodes. IGen will allow large hospitals to make use of high-performance computing resources such as Blue Waters to perform their analysis in a cost-effective manner. However, even at the scale of a single large hospital, thousands of compute nodes will have to be used to keep up with sequence data being analyzed for every arriving patient. At this scale, fault tolerance becomes an important issue, which must be tackled, keeping in mind the application as well as the system configuration. In this project, we analyze the log information from various sources (e.g., syslogs) to understand the fault-isolation domains in the subsystems and their complex interactions leading to unusual failure modes.

METHODS & RESULTS

Efficient Genomic Pipelines

Common Mathematical Kernels: Static analysis of the algorithms used in computational-genomics tools revealed the existence of common algorithmic kernels. These kernels form the basis of the mathematical models used to analyze genomic data, and a profiling study on Blue Waters revealed that these contribute to the bulk of processing time. We built efficient implementations of these kernels for the alignment and variant-calling steps for the AMD central processing units (CPUs) and the NVIDIA GPUs in Blue Waters. These were compiled into a curated list of "best-performant" kernels as a part of the Illinois Genomics Execution Environment (IGen) library. The IGen library also provides primitives for handling genomic data by utilizing the parallel file access (MPI file I/O) instead of the POSIX-based I/O used in the traditional tools.

Data Flow-Based Runtime System: The static analysis mentioned above also led us to the observation that most computational-genomics applications can be expressed as directed acyclic graphs (DAGs) with kernels as vertices and data dependencies between kernels as edges. We built a runtime system called ExEn (the Execution Engine) to execute computational-genomics applications as DAGs across CPUs and GPUs on a single node and scale to multiple nodes using MPI one-sided communication primitives. At the core of the ExEn runtime is a scheduling algorithm that can perform task placement using three constraints:

- Locality: Minimizing data movement, keeping in mind ccNUMA (Cache Coherent Non-Uniform Memory Access) domains, accelerator-to-host communication, and communication over the Gemini network.
- Processor affinity: Given implementations of some kernel functions on both CPUs and GPUs, deciding on a partition of work between the processors, keeping in mind the structure of the DAG and the historic performance of a kernel on that device for a "median" dataset.
- Shared resource contention: Accounting for microarchitectural resources on processors to ensure that colocated tasks do not significantly interfere with each other's performance.

Extreme-Scale Resilience

To understand the triggers of the recovery mechanisms and its failure/success, we created an augmented LogDiver tool (ALDT). The ALDT creates recovery-sequence clusters by performing temporospatial modeling of error logs, which helps to track failure/recovery-related events, which can affect the currently triggered recovery. Recovery-sequence clusters help to understand the failure propagation path and quantify the impact on the system/applications (Fig. 2). This abstraction was used to understand the reasons for the failure of the recovery mechanisms in the Blue Waters interconnect networks. The following insights on interconnect-related failures and recovery were developed:

- Our analyses identified the following causes for the failure of recovery mechanisms: (a) failures during recovery, (b) lack of hardware support to maintain consistent routing tables across the interconnect network, and (c) bad coordination (handshake and timeout issues) between services and different recoveries during the recovery period. Further, this understanding will allow system designers to create fault-injection test beds to evaluate and check next-generation systems.
- Failures during recovery are assumed to be low-probability events and are ignored in theory and practice. However, we show that this is no longer the case. Specifically, we show that 24% of the link-

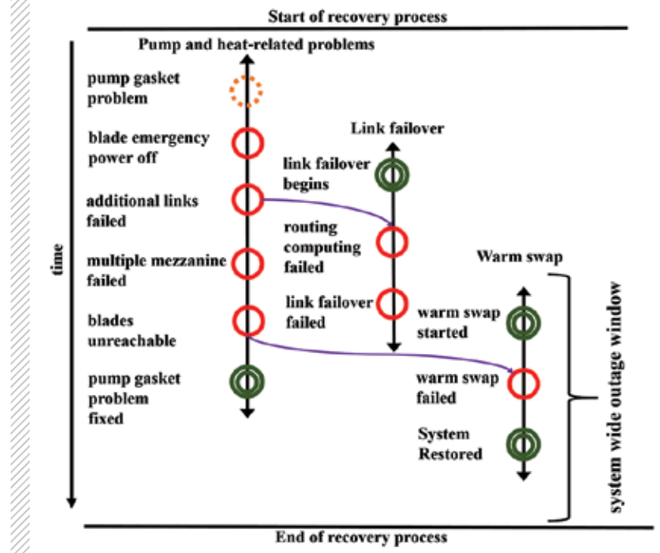

FIGURE 2: Failure propagation path showing sequence of events leading to system-wide outage.

recovery operations and 8% of blade-swapping procedures failed, which led to the failure of 20% of the active applications during these recovery epochs. A successful recovery does not guarantee to protect the application and system but only 0.02% of the active applications failed during the successful recovery period.

- Using real attack data and associated alerts as drivers, we have developed and evaluated *AttackTagger*, an adaptive learning-based IDS. The approach is based on probabilistic graphical models—specifically factor graphs—which integrates security alerts from multiple sources for accurate and preemptive detection. The method was validated using real data from attacks at NCSA.

WHY BLUE WATERS

Blue Waters is one of the few systems that can scale computations to tens or hundreds of thousands of cores on CPUs and GPUs. It also enables the study of failures in production petascale systems with its unique mix of XE6 and XK7 nodes. This capacity allows us to understand the performance–fault-tolerance continuum in HPC systems by enabling the investigation of application-level designs for mixed CPU and GPU node systems, and fault isolation in system components to mitigate failures at the application level. This allows us to design high-performance and resilient genomics pipelines that can make use of HPC systems.

NEXT GENERATION WORK

We ascertained the performance pathologies of several common kernels used in popular computational-genomics tools and built a scheduling algorithm that is able to dynamically decide task placement in a heterogeneous cluster. Also, we analyzed log data produced by Blue Waters to discover and quantify new failure modes that could not be efficiently handled by system-level recovery mechanisms. Our future work will bring together these observations to build a holistic runtime system that will jointly reason about performance and resiliency for coordinated placement and checkpointing decisions.

PUBLICATIONS AND DATA SETS

Banerjee, S. S., et al., Efficient and Scalable Workflows for Genomic Analysis. Proceedings of the *ACM International Workshop on Data-Intensive Distributed Computing (DIDC '16)* (ACM, New York, New York, June 1, 2016), pp. 27–36.

Jha, S., et al., Understanding Gemini Interconnect Failovers on Blue Waters. *Proceedings of the Cray User's Group (CUG)*, London, England, May 8–12, 2016.

Banerjee, S. S., et al., Sequences to Systems. Coordinated Science Laboratory Technical Report UILU-ENG-16-2201.

Di Martino, C., et al., LogDiver: A Tool for Measuring Resilience of Extreme-Scale Systems and Applications. Proceedings of the *5th Workshop on Fault Tolerance for HPC at eXtreme Scale (FTXS '15)* (ACM, New York, New York, June 15–19, 2015), pp. 11–18.

Cao, P., E. C. Badger, Z. T. Kalbarczyk, and R. K. Iyer, A Framework for Generation, Replay, and Analysis of Real-World Attack Variants. Proceedings of the *Symposium and Bootcamp on the Science of Security (HotSoS 2016)* (ACM, New York, New York, April 19–21, 2016), pp. 28–37.

LARGE-SCALE LEARNING FOR VIDEO UNDERSTANDING

Allocation: GLCPC/377 Knh
PI: Jia Deng[1]
Co-PI: Jason Corso[1]

[1]University of Michigan

EXECUTIVE SUMMARY

Video understanding, endowing computers with the ability to interpret videos as humans do, is one of the fundamental challenges of computer vision and artificial intelligence. Video data is ubiquitous, but our ability to perform automated analysis on such data is still primitive. In this project, we use Blue Waters to advance the research in video understanding, including recognizing human actions and activities, extracting high-level semantics from instructional videos, and accelerating deep neural network computation.

INTRODUCTION

Video understanding, endowing computers with the ability to interpret videos as humans do, is one of the fundamental challenges of computer vision and artificial intelligence. Video data is ubiquitous and is projected to account for 79% of all consumer Internet traffic in 2018 [1]. Yet our ability to perform automated analysis on such data is still primitive. We use Blue Waters to advance the research in video understanding.

One important problem is to understand human actions and activities. That is, given a visual input such as a video frame, generate a list of human action categories and their locations, for example, predicting that there is a person riding a horse at the lower left region of the given video frame. Automated recognition of human actions is key to the success of many important applications, such as human-computer interaction, robotics, and smart healthcare systems.

Another problem is to understand the high-level semantics of instructional videos with a focus on cooking activities. For a given cooking video, such as making a peanut butter and jelly sandwich, we seek to learn a cross-modal model of the temporal structure and constraints of the cooking process from both visual and audio content. The learned model will be able to generate a visual-textual summary of the process and will serve as a natural index for search and query across many such processes.

A third problem is in the area of robotic localization and mapping, which have recently become important in the context of autonomous driving. With the availability of more computing power, techniques like deep reinforcement learning [2] have become viable in many practical problems, including autonomous driving. An interesting question is how to improve deep reinforcement learning algorithms using games and simulations. We are interested in evaluating this technique on real-world problems such as learning the concept of objects and learning simple physics rules.

A fourth problem is how to make video understanding algorithms efficient. One aspect is how to accelerate the computation of deep neural networks (DNN) [3, 4], which are widely used for many subtasks for video understanding.

METHODS & RESULTS

Our main approach is machine learning. For the problem of understanding human actions and activities, we investigated deep neural networks (DNN) [3, 4]. The recent development of DNNs has led to large improvements for object recognition [5]. But compared to objects, human actions are far more complex. We have found that naively applying DNN-based object recognition algorithms for human actions does not perform well. We thus investigated novel DNN architectures for recognizing human actions. We have developed a novel multi-stream architecture that can integrate cues from humans, objects, and scene context. Our approach has achieved **state-of-the-art performance** on a large-scale action recognition benchmark [6].

For the problem of extracting semantics from instructional videos, we collected a large cooking video dataset, YouCook2, by querying YouTube. The videos were then annotated with sentence descriptions, as well as timestamps denoting the particular cooking styles, such as grilling and frying. As a first step, we generated an intermediate video representation by grouping similar pixels in space and time in a video using our streaming hierarchical video segmentation method [7]. We further extracted the semantic entities, such as objects and actions, from videos using random field modeling [8]. Our methods achieve much better results than the previous state of the art in the field.

For the problem of deep reinforcement learning, we have set up a simulation in the robot simulator Gazebo with wooden blocks and a robotic arm with two degrees of freedom. For the problem of accelerating DNN computation, we investigated custom data representations and ran simulations on how they can be used to speed up DNN computation through new hardware design.

WHY BLUE WATERS

Running DNNs and processing video require intensive parallel computation on both CPUs and GPUs. In order to perform thorough experiments, we need access to a large number of CPUs and GPUs. This makes Blue Waters **essential** for our research. At the peak of our usage, we were able to use a large number of CPUs and GPUs concurrently, which allowed us to process a large amount of video and explore a large design space of models. In addition, we received timely help from the Blue Waters team, which was also **indispensable** for our project.

NEXT GENERATION WORK

We hope to use Blue Waters to continue advancing video understanding, including more accurate understanding of human actions, deeper semantics from instructional videos, and faster DNN computation.

ALGORITHMS FOR EXTREME SCALE SYSTEMS

Allocation: Illinois/100 Knh
PI: William Gropp[1]
Collaborator: Luke Olson[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Continued performance enhancement of large-scale computer systems will come from greater parallelism at all levels. At the node level, this is seen in the increasing number of cores per processor and the use of large numbers of simpler computing elements in general purpose graphics processing units (GPGPUs). The largest systems must network tens of thousands of nodes together to achieve the performance required for the most challenging computations. Successful use of these systems requires new algorithms. Over the last year, we have shown the benefit of lightweight intranode balancing on scalability and performance. We continue to explore alternative formulations of conjugate gradient that eliminate some of the strict barrier synchronization and better use memory hierarchy, ways to reduce the impact of communication on the scalability of algebraic multigrid, and algorithmic approaches to resilience that exploit the multilevel representation in multigrid methodology.

INTRODUCTION

At extreme scale, even small inefficiencies can cascade to limit the overall efficiency of an application. New algorithms and programming approaches are needed to address barriers to performance. This work directly targets current barriers for effective use of extreme scale systems by applications. For example, Krylov methods such as conjugate gradient are used in many applications currently being run on Blue Waters (MILC code is one well-known example). Developing and demonstrating a more scalable version of this algorithm would immediately benefit those applications. In the longer term, the techniques that are developed will provide guidance for the development of highly scalable applications.

METHODS & RESULTS

Early results with alternative Krylov formulations have revealed several performance effects that can provide a factor of two or greater improvement in performance at scale. Current work has been limited by the fact that the non-blocking MPI_Allreduce on Blue Waters is functional but does not provide the expected (or perhaps hoped for) performance, particularly regarding the ability to overlap the Allreduce operation with other communication and computation. However, even with this limitation, we have seen a benefit in using non-blocking collective operations regarding a reduction in the sensitivity of the application to performance jitter and other irregularities.

WHY BLUE WATERS

Scalability research relies on the ability to run experiments at large scale, requiring tens of thousands of nodes and hundreds of thousands of processes and cores. Blue Waters provides one of the few available environments where such large-scale experiments can be run. In addition, only Blue Waters provides a highly capable I/O system, which we will use in developing improved approaches to extreme-scale I/O.

NEXT GENERATION WORK

We expect the next generation systems to rely on many more cores per node and to use different network topologies compared to Blue Waters. Additionally, there are likely opportunities to exploit new network capabilities and provide algorithms and programming systems adapted to the new memory, processor, and interconnect architectures.

PUBLICATIONS AND DATA SETS

Kale, V., Low-overhead scheduling for improving performance of scientific applications, Ph.D. Thesis, Spring 2015.

Eller, P., and W. Gropp, Exploiting nonblocking collective operations in conjugate gradient, *Copper Mountain Meeting,* 2015.

Bienz, A., Analyzing the performance of a sparse matrix vector multiply for extreme scale computers, *SC15*, Austin, TX, Nov. 2015.

Eller, P., Non-blocking preconditioned conjugate gradient methods for extreme-scale computing, *SC15*, Austin, TX, Nov. 2015.

PARALLEL ALGORITHMS FOR SOLVING LARGE ASSIGNMENT PROBLEMS

PI: Rakesh Nagi[1]
Co-PI: Ketan Date[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

The goal of our project is to develop fast and scalable algorithms for solving large instances of Linear Assignment Problem (LAP) and Quadratic Assignment Problem (QAP) using Graphics Processing Units (GPUs). LAP is polynomial-time solvable with cubic worst-case complexity, while the QAP is strongly non-deterministic polynomial-time hard (NP-Hard). To solve a linearized model of the QAP using branch-and-bound, lower bounds must be calculated using the Lagrangian dual technique, in which a large number of LAPs must be solved efficiently. Additionally, in a branch-and-bound scheme, a large number of nodes must be explored to find a provable optimal solution. To this end, we have used Blue Waters to develop: (1) A GPU-accelerated Hungarian algorithm for solving large LAPs in an efficient manner; (2) A GPU-accelerated Lagrangian dual ascent heuristic for obtaining lower bounds on the QAP. These algorithms will be used in a parallel branch-and-bound scheme to solve large QAPs to optimality.

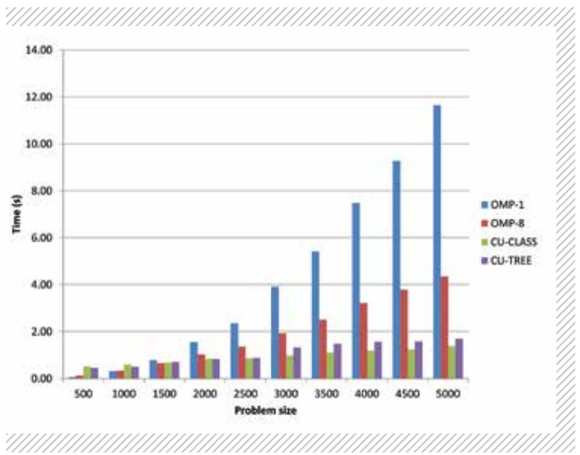

FIGURE 1: Comparison of execution times for CPU-based (OMP-1 and OMP-8) and GPU-based (CU-CLASS and CU-TREE) Hungarian algorithms.

INTRODUCTION

Assignment Problems are fundamental to the discovery in diverse branches of science and engineering. Some of their applications include information fusion, protein-protein interaction analysis, facilities design, vehicle routing, and resource scheduling. To gain meaningful insights, many applications demand quick solutions to large instances of Assignment Problems containing hundreds of thousands of vertices. This makes it incredibly challenging for the sequential algorithms designed for a single processor. Therefore, designing fast and scalable algorithms suitable for the state-of-the-art parallel programming architectures is essential. In this research, we intend to propose novel parallel algorithms for the Compute Unified Device Architecture (CUDA) enabled NVIDIA GPUs to solve the following two Assignment Problems.

Linear Assignment Problem (LAP): The objective of the LAP is to assign "n" resources to "n" tasks such that the total cost of the assignment is minimized. LAP can be solved in polynomial time using one of the many sequential/parallel algorithms that have been proposed in the literature. We chose to parallelize the famous Hungarian algorithm [1] on a GPU, whose theoretical complexity is O(n3).

Quadratic Assignment Problem (QAP): QAP was introduced by [2] 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. One of the ways of solving the QAP is to convert it into a Mixed Integer Linear Program (MILP) by introducing additional variables and constraints and solve it using branch-and-bound with the help of a strong lower bounding technique. We chose to parallelize the Lagrangian dual ascent algorithm for Level-2 Refactorization-Linearization Technique (RLT2) proposed by [3], in which we need to solve O(n4) LAPs and adjust O(n6) Lagrange multipliers to obtain a strong lower bound on the QAP.

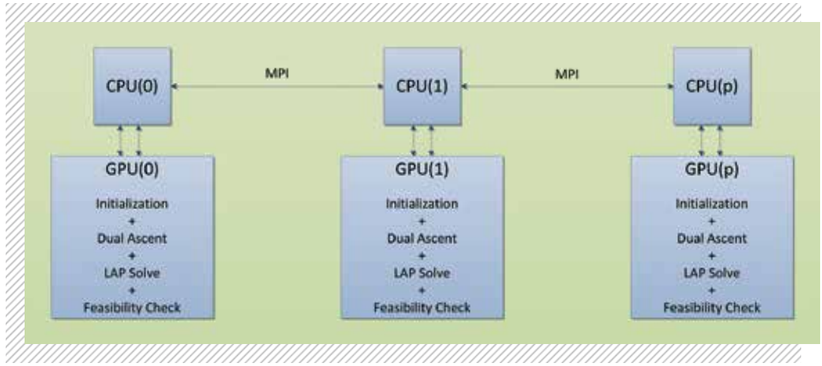

FIGURE 2: MPI+CUDA hybrid parallelization scheme for Lagrangian dual ascent.

METHODS & RESULTS

GPU-accelerated LAP solver: We have proposed two GPU-accelerated variants of the Hungarian algorithm. Our main contribution is an efficient GPU-based parallel algorithm for the augmenting path search, which is the most time intensive step. We show that our algorithm(s) can find multiple vertex-disjoint augmenting, which drastically reduces the execution time. Extensive numerical tests reveal that for problems with n > 1000, our algorithm(s) (CU-CLASS and CU-TREE) are substantially faster than the sequential and OpenMP implementations (OMP-1 and OMP-8) solved on a multi-core CPU (Fig. 1).

GPU-accelerated RLT2 solve: We designed a parallel Lagrangian dual ascent heuristic for solving RLT2 using hybrid MPI+CUDA architecture (Fig. 2), for which we used multiple GPUs. The O(n4) LAPs are split across these GPUs and solved using our GPU-accelerated Hungarian algorithm, while the O(n6) Lagrange multipliers are updated by multiple CUDA threads in parallel. This algorithm is scalable and provides good parallel speedup (Fig. 3) for problem instances Nug18, Nug20, Nug22, and Nug25 from the QAPLIB [4].

Parallel branch-and-bound solver for QAP: In this work, we used GPU-accelerated RLT2 solver in a branch-and-bound scheme to solve QAP instances to optimality. For a node in the search tree, we fix a facility to a location and solve the corresponding RLT2 sub-problem, whose objective value provides a lower bound on the QAP. If this value is greater than the incumbent solution then the node is fathomed, otherwise, it is branched further. Each node is processed using a bank of GPUs. By using multiple such banks, we can process multiple nodes in parallel. Table 1 shows the results for problem instances Nug18, Nug20, and Nug22 from the QAPLIB.

Problem	No. of banks	GPUs per bank	Nodes explored	Time (min)	Avg. Bank utilization
Nug18	10	9	514	16.75	0.916
Nug20	10	10	669	38.31	0.859
Nug22	12	12	969	56.63	0.931

TABLE 1: Branch-and-bound results for medium-sized QAPs.

WHY BLUE WATERS

In a typical branch-and-bound tree, we need to explore a large number of nodes in order to find an optimal solution. Also, as the problem size grows, the number of nodes that need to be explored grows exponentially. Therefore, we need a large number of processors which 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. Coupling the parallel branch-and-bound with the fast GPU-based lower bounding techniques will enable us to solve large-sized problems from the QAPLIB, which still remain unsolved.

NEXT GENERATION WORK

Our ultimate objective through this research is to provide efficient solution methods for a class of Assignment Problems. These problems arise in diverse branches of science and engineering and

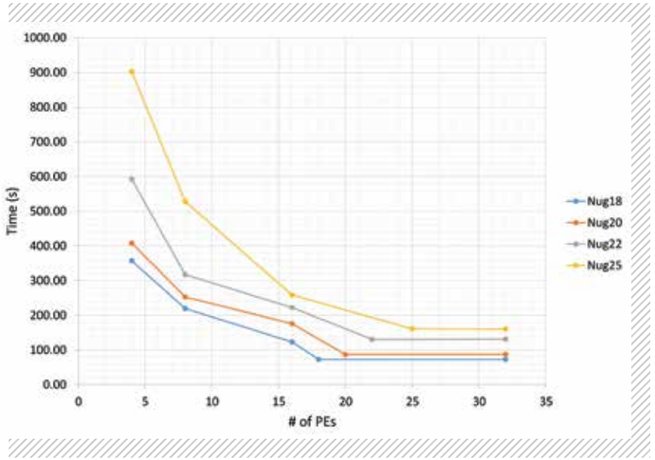

FIGURE 3: Scalability results for 200 iterations of parallel Lagrangian dual ascent on instances from QAPLIB.

they are part of many cutting-edge projects with high socio-economic impact. Using our methods, scientists and engineers will be able to solve Assignment Problems containing hundreds of thousands of vertices within a matter of minutes, leading to **transformative** discoveries. The parallel programming library developed during this research will provide a platform for solving large-scale Assignment Problems and comparing the performance of different algorithms, thus ensuring continued advancement of science and engineering.

PUBLICATIONS AND DATA SETS

Date, K., R. Nagi, GPU-accelerated Hungarian algorithms for the Linear Assignment Problem. *Parallel Comput.*, (2016), http://dx.doi.org/10.1016/j.parco.2016.05.012

THE NEXT GENERATION OF LARGE-SCALE SPARSE MATRIX COMPUTATIONS

FIGURE 1: Communication in sparse matrix-vector multiplication

PI: Luke Olson[1]
Collaborators: Amanda Bienz[1], Bill Gropp[1], and Andrew Reisner[1]

[1]University of Illinois at Urbana-Champaign

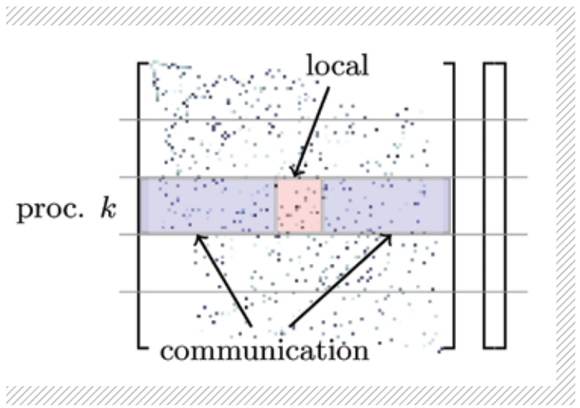

EXECUTIVE SUMMARY

Sparse matrix computations are a common element in a range of high-performance computing applications and often dominate computation, particularly at scale. The goal of this project is to develop numerical schemes that limit parallel communication in sparse matrix methods, leading to efficient and scalable sparse operations without loss of accuracy.

Iterative methods for approximating the solution to sparse linear systems rely on sparse matrix-vector multiplication as a key computational kernel. Yet communication costs drive the complexity of these operations. The focus of this project is on algorithm design to alleviate communication costs at scale.

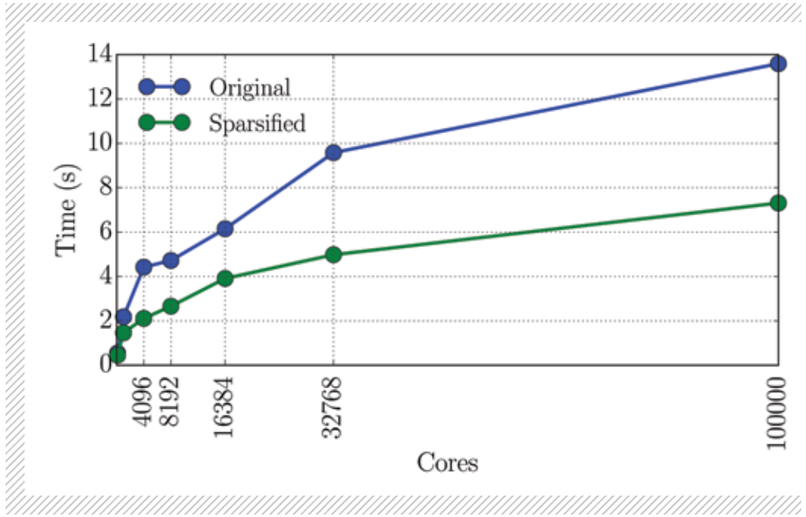

FIGURE 2: Effect of matrix sparsification on solver time.

This project has contributed two key developments: the ability to *sparsify* the matrix operations to limit communication and the method of redundant computation to localize computation, thus limiting communication distance and time.

INTRODUCTION

Sparse matrix problems of the form $Ax=b$ are common in many large-scale simulations. In this project, multigrid methods are considered as a method for the iterative approximation for this problem. Multigrid methods construct a series or hierarchy of smaller sparse matrices that are used to form an iterative refinement of the solution. This leads to fewer floating-point operations in the solution to $Ax=b$, but potentially very high communication costs and memory movement relative to the amount data in the computation.

A key kernel in this process is that of sparse matrix-vector multiplication $w \leftarrow A\ v$. Here, the sparsity pattern—or the location of non-zero entries in the matrix—governs the computational expense of the operation. A common bottleneck in the computation is the scenario where the sparse matrix resides on a high number of cores with a low number of matrix elements per core (Fig. 1).

Achieving efficiency in these computations is critical as many applications rely on these sparse operations throughout the simulation, leading to significant resource usage.

METHODS & RESULTS

Two methods were explored to address the communication demands in the problem. First, a method of sparsification whereby specific entries in the matrix are eliminated from the computation is explored. In multigrid, the matrices in the hierarchy are successively smaller, but relatively more dense, leading to wider communication patterns. A method is designed in [1] that avoids communicating work associated with *weak* or low-influence matrix entries in the problem. This leads to a notable speed-up of 2x, particularly at large core counts (Fig. 2). In addition, the *topology* of the system is also incorporated in the communication of sparse entries, preferring local on-node communication in the algorithm. This leads to a significant reduction in total message traffic, thereby reducing system contention and total simulation time (Fig. 3).

Ultimately, sparse matrices in the multigrid hierarchy are too small to reside on the full set of processing elements. Therefore a redistribution of data is needed to map matrix data to a subset of processors. This comes at a cost, which was explored in this project. By performing a redistribution of data based on a performance model, nearly scalable results can be achieved. As an example, for p processors,

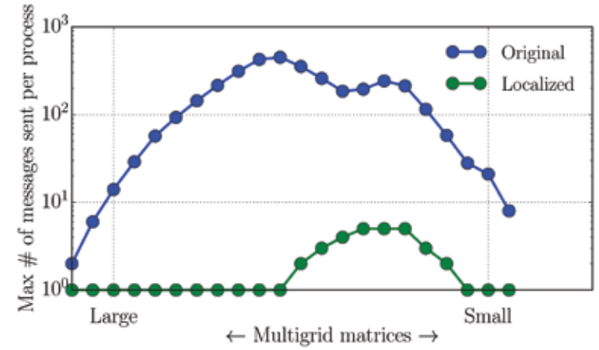

FIGURE 3: Effect of localized message communication on MPI messages.

it was observed that the standard approach to redistribution of data yielded a growth of $p^{1.1}$ in simulation time, while a redistribution designed to limit communication yields around $p^{0.17}$. This is a signficant saving, particularly as p approaches one million cores.

WHY BLUE WATERS

Blue Waters was key to this work as it provided access to the fast Cray Gemini interconnect as well as to the large core count needed for accurate scalability studies. Sparse matrix algorithms need to take advantage of multiple aspects of the compute architecture, and Blue Waters allowed for the development of both on-node features and long-distance communication decisions in the multigrid algorithm. As multigrid methods are expected to be a key solver technology on future systems, access to Blue Waters has been instrumental in advancing these methods.

NEXT GENERATION WORK

Increased concurrency will continue to challenge the scalability of sparse matrix computations. One goal is to capture the topology features of these systems more directly in the algorithms. This project provided several steps toward that goal. A future element of this work would be to add heterogeneous compute elements to the multigrid algorithm and to enhance the underlying performance models to effectively capture their use.

EVALUATING DATA-DRIVEN DETECTORS OF ELECTRICITY THEFT IN SMART GRIDS

Allocation: Illinois/50.0 knh
PI: William H. Sanders[1]
Co-PI: Varun Badrinath Krishna[1]
Collaborator: Juran Kirihara[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Electricity theft is a billion-dollar problem faced by utilities around the world and current measures are ineffective against sophisticated theft attacks that compromise the integrity of smart meter communications. We are devising algorithms that detect such theft attacks, and that are based on mathematical techniques in statistics and machine learning. The goal is to detect and mitigate theft by identifying anomalies in consumption patterns of electricity consumers. We used Blue Waters to evaluate the effectiveness of the autoregressive integrated moving average (ARIMA)-based approach to detect simulated anomalies in smart meter data. The best parameters for these algorithms need to be found using scanning techniques, and they need to account for a wide range of attack parameters that produce anomalies. Our evaluation is based on a large dataset obtained from a real smart meter deployment.

INTRODUCTION

Bloomberg News reported that electricity theft in India contributes to blackouts and costs $17 billion in lost revenue annually. According to the World Bank, electricity theft contributes to a loss in electricity delivery of over 25% of generated supply in India, 16% in Brazil, 6% in China and the U.S., and 5% in Australia. Theft in these countries is almost always achieved by tapping into electric distribution lines. To detect these thefts, utility companies such as BC Hydro have been convincing consumers to install smart meters. However, there has been some push-back as consumers have begun to realize that smart meters are vulnerable to cyber intrusions. In 2010, the Cyber Intelligence Section of the FBI reported that smart meter consumptions were being under-reported in Puerto Rico, leading to annual losses for the utility estimated at $400 million. In 2014, BBC News reported that smart meters in Spain were hacked to cut power bills. Given that smart meters can be compromised, the roll-out efforts of utilities such as BC Hydro may only increase the attack surface for cyber intrusion-based theft methods.

We identified seven classes of electricity attacks, some of which distribute the monetary loss to consumers, at no loss to the utility. Therefore, this problem is not only important to utilities, but also consumers around the world.

METHODS & RESULTS

The methods in this project are detailed in our earlier work [2], where we simulated electricity theft attacks on 500 consumers and tested our detector's false positive and false negative rates on them. The detector fitted an ARIMA model to the consumption data time series and then flagged outliers using a confidence interval created from the model.

We used methods in [3] to fit the ARIMA time series but learned from the larger simulation on Blue Waters that these methods do not scale well and are very sensitive to outliers. Also, we found unexpected results that used Python packages built by third parties. Upon further investigation, we identified errors in the algorithms coded in those Python packages (specifically the statsmodels.tsa.arima_model.ARIMA package). The algorithms simply do not implement ARIMA models correctly and use the differencing order term in the ARIMA model in a manner that is inconsistent with the theory.

WHY BLUE WATERS

Our earlier work published in [2] was performed at a much smaller scale (500 consumers) on a regular server rack, consuming inordinate amounts of processing time. We wanted to perform evaluation studies of our detector at a larger scale (2900 consumers). Also, we wanted to try out many parameters for our detector at that scale. Without Blue Waters, it would have taken years to complete these tasks.

NEXT GENERATION WORK

Our experience with Blue Waters helped us identify problems with third party software packages. We are redeveloping those packages ourselves and hope to complete them in time for the September 2016 call for proposals. With the correct algorithms, we hope to discover insights on our electricity theft detector from our use of Blue Waters.

IMPACT DRIVEN RESEARCH: EVALUATING AND IMPROVING THE DE FACTO STANDARD FOR PARALLEL PARTICLE ADVECTION

Allocation: Innovation and Exploration/100 Knh
PI: Robert Sisneros[1]
Co-PI: David Pugmire[2]

[1]University of Illinois at Urbana-Champaign
[2]Oak Ridge National Laboratory

EXECUTIVE SUMMARY

Foundational visualization algorithms are central to the production of visualization tools running at computing centers around the world but consume tremendous amounts of limited resources. We believe that understanding the performance characteristics of these algorithms is critical in being good stewards of those unique resources. In this project, we investigate the foundational algorithm for parallel particle advection by focusing on a parameter sweep of the de facto standard algorithm "Parallelize Over Data" (POD). Our work has led to several discoveries: a glaring and resolvable issue in poor default values, identification of the best choices of parameters for future tests, and the existence of viable testing alternatives that save time and resources. Our work sets forth a framework for applying solutions to a practical and often overlooked area in computer visualization.

INTRODUCTION

The understanding of vector fields resulting from large scientific simulations is an important and often difficult task. The predominant analytical techniques are derived by calculating the paths that a set of weightless particles travel along when released in a flow. There are a rich set of direct and derived observable quantities that have resulted in several such methods: streamlines, pathlines, stream surfaces, and finite time Lyapunov exponents. However, there are significant computational challenges of these particle-based visualization approaches when they are applied to very large vector field data. Furthermore, the degree to which this negatively impacts performance is data-dependent. Finally, this dependency does not apply solely to the flow field, as many of the above mentioned techniques are highly sensitive to initial seed placement. Therefore, analysis typically includes a trial-and-error exploration where a single flow field for each instance of the calculation may have drastically different computational profiles.

As a result of these challenges, particle advection-based algorithms are susceptible to overload/starvation. Efficient computation requires careful balancing of computational demands placed on I/O, memory, communication, and processors. Algorithms for load balanced particle advection in small parallel settings and have been implemented for several of these components into a test branch of the source code of VisIt, an open source analysis tool. VisIt's particle advection framework provides access to current data decomposition and domain distribution specifics, as well as the capability for the on-demand loading of a domain to any processor.

METHODS & RESULTS

We performed an extensive parameter sweep of the POD algorithm using the implementation found in VisIt along with an updated version that performs communication asynchronously. We developed a framework that manages generating and executing arbitrary test configurations of parameter settings and performed tests on multiple datasets (see Figure). To aid in our analysis, we developed an imbalance metric to help unify comparability and created new visual metaphors for viewing the multiple axes of the study.

We have made significant findings upon analysis of our parameter sweeps to date. First, we find that optimal configurations are not the widely used default, and, in fact, we find that the default configurations to be quite poor. Second, we find that it is not only possible to predict proper configurations, but that the same configuration is likely applicable in both the asynchronous and synchronous versions of the algorithm. On any given architecture we expect that the default configuration may not be optimal, but through testing, this is correctable and an updated system-specific default will save computing resources. Furthermore, resources may be saved by testing, given that asynchronous tests are the only ones needed. Additional analyses made possible through the imbalance metric has allowed us to further pinpoint additional parameters to ignore in future testing. Moreover, we have also identified areas for directed future analysis.

NEXT GENERATION WORK

We want to bridge the gap between the state-of-the-art and the state-of-the-practice for all foundational visualization algorithms. Furthermore, we hope to promote the efficient use of world-class computing facilities through evangelizing such a practical, yet impactful research paradigm.

PUBLICATIONS AND DATA SETS

Sisneros, R., and D. Pugmire, Tuned to terrible: A study of parallel particle advection state of the practice. *High Performance Data Analysis and Visualization (HPDAV), IPDPS 2016*, Chicago, IL. May 2016.

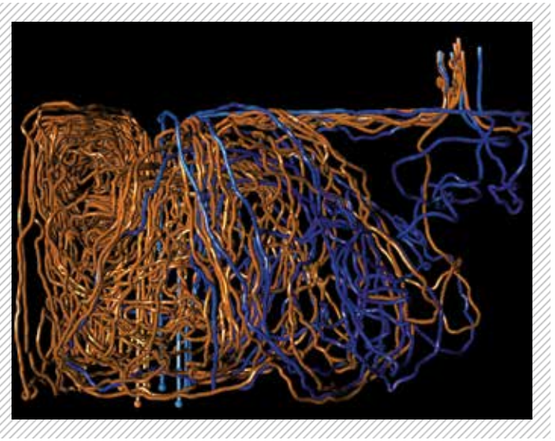

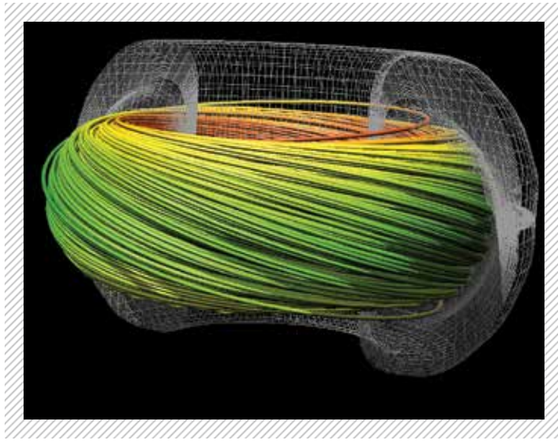

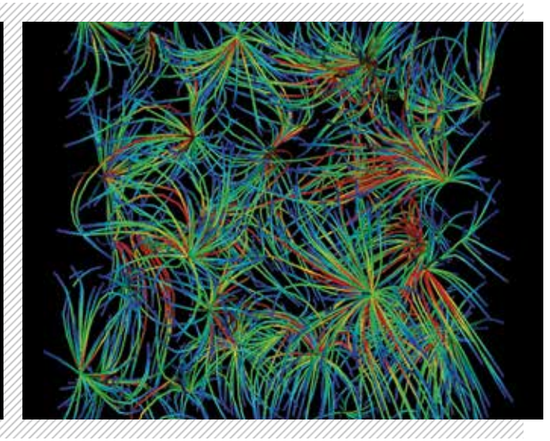

FIGURE 1: Four datasets were used for the parameter sweep in this work. Typical flow visualizations from each are, left to right: astrophysics dataset, thermal hydraulics dataset, fusion dataset, and synthetic test dataset.

SOFTWARE EXTREME BENCHMARKING

Allocation: Industry/595 Knh
PI: Ahmed Taha[1]

[1]University of Illinois Urbana-Champaign

EXECUTIVE SUMMARY

The current allocation is a continuation for running extreme benchmarking for key computational fluid dynamics (CFD) commercial codes on the Blue Waters. The allocation helped us reaching new "global" records in scaling two codes: Ansys-Fluent to 114,000 cores and Starccm+ of CD-Adapco to 102,000 cores.

INTRODUCTION

In service of the computational fluid dynamics (CFD) simulation community and in particular NCSA's Private Sector Program (PSP) industrial partners, PSP provides "extreme" benchmarking and scalability studies on leading commercially oriented CFD codes, provided by Independent Software Vendors (ISVs), which are used by and are of high interest to many industrial partners. The objective of the project is to continue exploring the scaling limits of these codes to inform users of the scaling capabilities for better production and workflow. Two world-class CFD codes are included in the study: Fluent from Ansys,inc. and Starccm+ of CD-Adapco.

The project was carried out in close collaboration between the projects primary investigator, Ahmed Taha, independent software vendors (ISV) Ansys and CD-Adpco, and the Blue Waters vendor (Cray, Inc.).

METHODS & RESULTS

In order for this project to serve the real-life industrial needs, four main parameters were essential in running this study: (1) a real-life application, (2) big mesh size, (3) complicated geometry, and (4) sophisticated physics. Close collaborations with the ISVs and Cray were essential in adjusting environments.

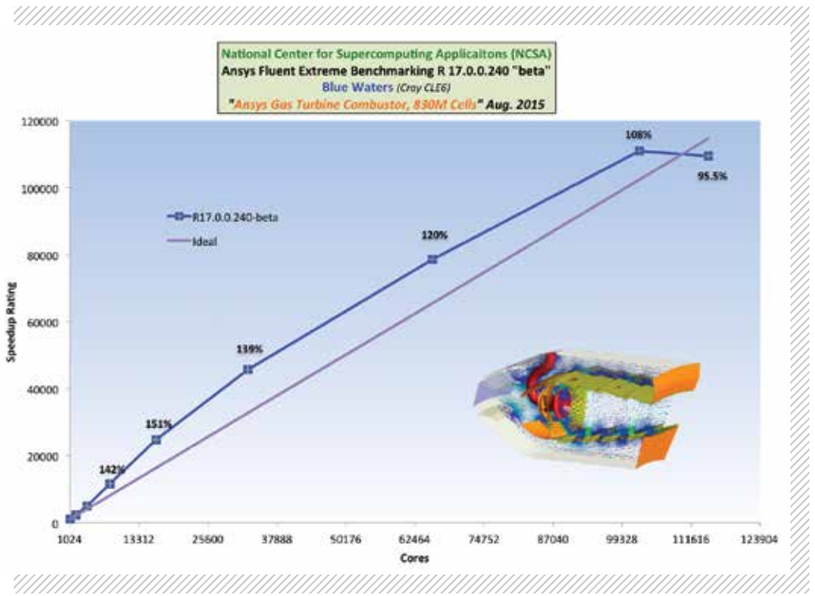

FIGURE 1: Ansys-Fluent Speedup Scaling Curve, 830M Cells, and 114,000 Cores.

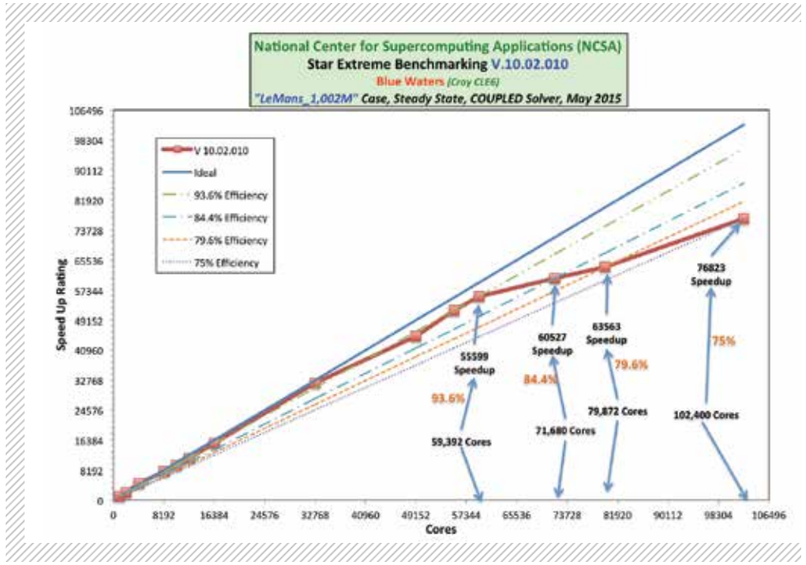

FIGURE 2: Starccm+ Speedup Scaling Curve, 1B Cells, and 102,000 Cores.

Case Study 1: Ansys-Fluent scaled to 10,000 cores using Ansys Benchmark Truck model of 111 million cells. Using successive upgrades of Fluent, the same model was scaled to 20,000 cores, followed by scaling to 36,000 cores using a bigger mesh of 830 million cells for the Ansys generic gas turbine combustor model in mid- 2015 with 80% scaling efficiency. The project concluded its 2015 BW allocation by achieving the highest scaling record of 114,000 cores with 95.5% scaling efficiency using the "beta" version of Ansys-Fluent 17.0 (Fig.1).

Case Study 2: In early 2015, Starccm+ scaled to 13,000 cores using the LeMans racing car model of 100 million cells using Starccm+ V.9.06.011. In March, a bigger mesh of 514M cells for the same model scaled to 20,000 cells. CD-Adapco provided a bigger mesh for the same LeMans model with 1 billion cells, which we scaled to 59,000 cores with 93.6% scaling efficiency then to 102,000 cores with 75% scaling efficiency using Starcccm+ V.10.02.010 (Fig.2). Per this achievement, NCSA currently holds the global record for scaling Starccm+ to 102,000 cores.

Among the many benefits of this extreme scaling study are:

- It provided users with information about the robustness of the flow solvers of the tested codes, which are presented in the high fidelity simulations running at such massive core counts.
- It drove the mesh count, allowing more details on the complicated physics involved for more thorough design studies.
- It showed how time to solution is reduced so customers that have many cores can take advantage of this. However, this is not just for large core counts, even on smaller core counts, problems can be simulated with faster turn around because of the good scaling.
- It showed that extreme scaling it is not only important for mega-jobs of huge mesh sizes. Success on extreme scaling also means that you can run smaller problems on more cores than before and see great scaling.

The study results added to the depth of the PSP strategic relationship with the ISV community.

WHY BLUE WATERS

Blue Waters has been instrumental in allowing this work to reach high levels of success. The massive computational resources and expert staff of NCSA and Blue Waters make it the only place capable of handling a project of this magnitude.

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

Allocation: Illinois/50.0 Knh
PI: Levent Gürel[1]
Co-PI: Prof. Wen-mei Hwu[1]
Collaborators: Weng Cho Chew[1] and Narayan Aluru[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Our main research goal is to solve extremely large scientific problems involving big-data issues. We have been working on scalable parallel implementations of the multilevel fast multipole algorithm (MLFMA) on massively-parallel supercomputers. We derive our immediate applications from computational electromagnetics, acoustics, optics, elastics, but the methods we develop have applications in a plethora of disciplines, such as quantum mechanics, fluid dynamics, astrophysics, molecular dynamics, structural mechanics, heat equation, potential theory, to name a few. Thanks to the massive parallelization offered by Blue Waters, we have been able to solve larger problems in shorter times. Most recently, we have achieved solutions of electromagnetics scattering problems involving 1.5 billion unknowns. Every phase of such a solution requires the handling of 1,500,000,000 x 1,500,000,000 dense matrices and hence incorporates many aspects of big-data computing. These cutting-edge solutions provide us with new insights into how to improve our solver algorithms and parallelization strategies. We plan to tackle increasingly larger problems and also to benefit from the graphics processing unit resources available on Blue Waters by honing our heterogeneous computing skills.

INTRODUCTION

We have been working on solutions to extremely large problems that are derived mainly from electromagnetics, acoustics, and optics, to name a few. An electromagnetic scattering problem may become very large as either the frequency increases or the target size gets larger. Simulating the scattering of high-frequency radar waves from a large aircraft, for instance, may require the solution of millions or even billions of unknowns. Similarly, medical imaging with microwaves requires solutions of thousands of large problems. For such large-scale simulations, we use our in-house parallel implementation of the multilevel fast multipole algorithm (MLFMA) [1-4].

MLFMA is a reduced-complexity solver, and as such, it enables the solution of unprecedentedly large problems not only in computational electromagnetics but in many other disciplines as well. Parallelization of MLFMA on homogeneous central processing unit (CPU) architectures is not trivial and even more difficult on heterogeneous CPU-graphics processing unit (GPU) architectures [5]. We have been addressing several challenges to obtain a scalable parallel implementation of MLFMA on the CPU nodes of Blue Waters. We plan to overcome many more challenges to achieve a scalable parallel implementation of MLFMA on the heterogeneous CPU-GPU architecture of Blue Waters. This will allow us to address previously unsolvable scientific problems, as an outcome of this exploratory project.

METHODS & RESULTS

Our research revolves around our in-house parallel implementation of the MLFMA to compute solutions to various problems, mostly derived from computational electromagnetics (CEM). Prior to Blue Waters, our computing facilities were limited to 16 computing nodes and up to 256 processes. Using those facilities, the largest problems we were able to solve involved 212 million unknowns, mainly limited by the available memory. We achieved out-of-core solutions by using solid-state drives (SSDs) local to each node in order to overcome the memory limitations. This way we were able to extend our solution capability up to 1.1 billion unknowns, at the expense of runtime.

On Blue Waters, we first had to adapt our code to use more computing nodes and more processes. MLFMA uses an "internode communication map" to determine which node will communicate with whom. These maps had to be updated and generated for the larger number of nodes available. As a result of all of our adaptations, we have so far been able to run simulations on up to 512 nodes and up to 1024 processes.

We have increased the number of unknowns in the problems we can solve, by achieving the solution of an electromagnetic scattering problem with 1.5 billion unknowns. Figure 1 demonstrates the accuracy of the solution, where the target object is a simple but extremely large conducting sphere of 1200 wavelengths in diameter. This canonical object is chosen because of the availability of the semi-analytical Mie-series solutions for spheres so that computational results can be compared and validated. A surface integral equation is used to formulate the physical scattering problem in mathematical terms, and a method-of-moments discretization is used to transform the continuum formulation to a matrix equation, whose solution is well suited to digital computers. The ensuing 1,500,000,000 x 1,500,000,000 dense matrix equation is solved with MLFMA and parallelization on Blue Waters.

We will continue to scale up the problem size that can be efficiently solved by our parallel MLFMA solvers. We are currently investigating the bottlenecks that degrade the efficiency of scaling beyond 512 nodes. Our goal is to scale to at least 4096 nodes and achieve the solution of an electromagnetic scattering problem with 5 billion unknowns.

WHY BLUE WATERS

As the problem size (i.e., the number of unknowns) grows, memory requirement increases to the point where both the total memory needed to solve the problem and also the memory available at a single computing node may become insufficient. Prior to Blue Waters, we had reached this limit and we had been experimenting with out-of-core methods to use the disk storage as additional memory, even though we had to endure a huge penalty in the time required to solve problems. The huge number of computing nodes makes available a much larger total memory, thus enabling the solution of such enormous problems that were **impossible for us to solve before** Blue Waters. Additionally, improving parallelization and using more nodes lead to much shorter solution times.

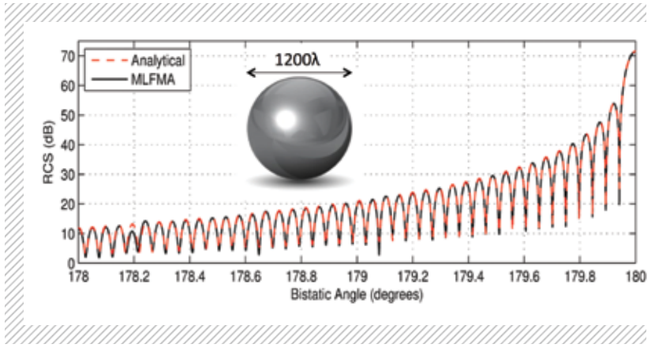

FIGURE 1: Solution of an extremely large electromagnetic scattering problem involving 1.5 billion unknowns. Scattering from a conducting sphere of 1200 wavelengths in diameter requires the solution of a 1,500,000,000 x 1,500,000,000 dense matrix equation. The computed bistatic scattering values are compared to semi-analytical Mie-series solution to demonstrate the high accuracy of the solution.

NEXT GENERATION WORK

We plan to develop scientific solvers for exascale computing. Specifically, we intend to develop parallel implementations of fast-multipole and Fourier-transform software. We consider the software aspect of parallel and heterogeneous computing an essential component for extracting exascale performance from immediately-next-generation supercomputers that will be available within the next 5-10 years. In other words, brute-force use of raw hardware power is neither sufficient nor rational; it will be the intelligent use of software that will unleash the exascale potential of the next-generation supercomputers.

PUBLICATIONS AND DATA SETS

Hidayetoğlu, M., et al., Parallel Solutions of Inverse Multiple Scattering Problems with Born-Type Fast Solvers, *Progress in Electromagnetics Research Symposium (PEIRS2016)*, Shanghai, China, Aug. 2016.

El Hajj, I., KLAP: Kernel Launch Aggregation and Promotion for Optimizing Dynamic Parallelism, *IEEE/ACM International Symposium on Microarchitecture (MICRO'16)*, Taipei, Taiwan, Oct. 2016.

HIGH-SPEED LINK SIMULATION USING X PARAMETERS, VOLTERRA SERIES AND THE LATENCY INSERTION METHOD (LIM)

Allocation: Illinois/50 Knh
PI: Jose Schutt-Aine[1]
Co-PI: Raj Mittra[2]

[1]University of Illinois
[2]The Pennsylvania State University

EXECUTIVE SUMMARY

In today's high-speed communication era, humongous data throughput is being routed through communication channels. Circuit design and analysis techniques have become more challenging and will remain a critical component in ensuring the success and viability in the design of short-range high-speed input/output (I/O) links. Because of the large amount of data, the simulation of such systems is a daunting task. Present systems require bit error rates (BER) in the order of 10-15. A deterministic simulation of such a large number of bits would require several thousand years using a modern-day computer. We used the Blue Waters computational platform to explore the feasibility of simulating high-speed serial links. By properly combining the computational powers of the Blue Waters system with new algorithms developed in our research group, we reduced the computational time by several orders of magnitude and made the design cycle for high-speed links more manageable.

INTRODUCTION

In today's high-speed communication era, bandwidth is increasingly becoming the major limiting factor in channel performance. In order to keep up with the humongous data throughput needed, high edge rates are routed through communication channels consisting of complex arrangements of high-density interconnects, packages, and integrated circuits. As short-range data rates keep increasing to multi gigabits per second (Gbps), interactions between channel links, integrated circuits, packages, and power distribution networks take place and become more difficult to predict and manage. Noise sources such as crosstalk and jitter limit the performance of these systems and force the data rates to be well below the Shannon limit of channel capacity. Circuit design and analysis techniques have become more challenging and will remain a critical component in ensuring the success and viability in the design of short-range high-speed input/output (I/O) links.

To circumvent the performance limitation and voltage scaling problem posed by the conventional parallel buses, the industry has migrated from the multi-drop parallel buses to point-to-point serial buses. In a serial bus, a device called SerDes (Serializer/Deserializer) is used to transmit and receive data over the serial link. The SerDes can be either a stand-alone device or, in most cases, an intellectual property (IP) core integrated into a serial bus controller or an ASIC. The timing skew problem encountered in a parallel bus is solved by embedding the clock signal into the data stream. Since there is no separate clock signal in a serial bus, timing skew between clock and data (which, together with the minimum setup and hold time, determines the maximum data transfer rate) no longer exists. Consequently, a serial bus operates at a much higher data rate than its counterpart in a parallel configuration.

The design of a multi-gigabit SerDes can be challenging because of the high-speed, mixed-signal circuitry involved, as well as stringent electrical specifications. Robust design of these systems requires access to reliable modeling and simulation tools that can help predict the performance and optimize the design of these systems. Simulation techniques for the prediction of signal propagation in these environments have become a subject of increased interest over the past few years. Behavioral and macro-modeling techniques have been among the most popular methods used to predict the performance of these systems.

METHODS & RESULTS

The standard tool for circuit simulation is SPICE. However, for very large networks, the computational time becomes prohibitively large. Our research group introduced the latency insertion method (LIM) for time-domain simulation of large networks [1]. The method is based on a finite difference time domain (FDTD) formulation of the branch and node equations and obtains the solution in a leap-frogging method similar to the Yee algorithm [2] used in electromagnetics. No matrix inversion is performed and the method has linear numerical complexity.

The advantage of the LIM method rests in the fact that it is significantly faster than SPICE. Moreover, this speedup advantage increases as the number of nodes in the network increases. Thus LIM is a time-domain formulation that leads to the generation of update algorithms for the simulation of networks. This algorithm exhibits linear computational complexity; in addition, because of the time domain nature of the formulation, it can handle nonlinearities.

The latency insertion method has proven to be a viable circuit simulation approach that is a good alternative to SPICE. However, transistor models for a LIM-based simulator have yet to be developed. Thus far, the models have neglected the nonlinear charge storage effects which may be critical in short-channel technologies. Given that the LIM algorithm has been found to be efficient in simulating large transistor circuits, it is important that the basic model used be very accurate. In our work, a dynamic model for short-channel MOSFET transistors was developed for LIM-based simulation. The model incorporates the nonlinear charge storage effects. A modular representation is adopted for the MOSFET, which leads to general branch and node update equations. A dynamic LIM-based model was developed for an NMOS transistor which led to the associated update equations. A simulation program was implemented. Results are examined and comparisons are made with existing techniques.

An approach for the transient simulation of circuits through the latency insertion model using advanced models for MOS transistors was implemented. By taking into account the dynamic charge storage effects in short-channel devices a more accurate simulation of high-speed digital and analog circuits via the latency insertion method was performed. The approach makes use of the SPICE LEVEL 3 transistor model for MOSFETs. The use of the latency insertion method allows better convergence and higher computational speed for the simulation. Computer simulations showed improvement in accuracy by using the high-level models. However, when the number of transistors becomes large, the computational simulation becomes very slow, which warrants the use of a platform such as Blue Waters.

In recent years, macromodeling has received increasing attention due to the ever-growing complexity of networks used for computer and communications applications. In particular, blackbox macromodeling has become an attractive approach since it allows us to bypass the complex topologies of networks through the use of behavioral descriptions at the ports and terminals of these systems. S-parameter blackbox macromodeling has become a convenient vehicle to exchange behavioral data pertaining to packages, connectors, channels used in high-speed links and other types of communication networks. As an example, in signal integrity applications, it is often required to perform the simulation of systems for which equivalent circuits are not available. This occurs in instances where the data from the system are obtained from measurements or from a field solver. In these cases, only the network port parameters are known as a function of frequency and are used to represent the network as a macromodel. The final objective is to use the exchanged behavioral model in conjunction with a circuit simulator such as SPICE. In order for such a simulation to become possible, two key steps must be performed.

First, a model-order reduction process must be carried out to extract the multiport poles and residues of the network. The poles and residue contain the "character" of the system and may describe its behavior over a wide range of frequencies. Today, the most commonly used pole/residue extraction method is the vector fitting method which has proven to be robust and reliable. Once the poles and residues of a system are obtained, a simulation of the network can be performed using time-domain simulation techniques such as recursive convolution. In many cases however, a circuit description is often desired in order to be combined with external circuitry that may contain nonlinear elements. The synthesis of such circuits has been explored in the literature using various techniques. In particular, the pole-residue representation can be used to perform a circuit synthesis of the macromodel in which a SPICE-compatible netlist is generated and can be

used for simulation. In our effort, we developed a synthesis technique for the implementation of equivalent circuits based on scattering parameter representation of linear broadband networks. Several circuit topologies for complex pole residue pairs were defined and tested and found to accurately describe the behavior of the blackbox networks. When the number of ports becomes large, the computational efficiency of the method is very limited, which warrants the use of a platform such as Blue Waters.

WHY BLUE WATERS

We plan to use the Blue Waters computational platform to explore the feasibility in simulating high-speed serial links that would necessitate several thousand years using standard algorithms. By properly combining the computational powers of the Blue Waters system with new simulation algorithms developed in our research group, we plan to reduce the computational time by several orders of magnitude. More specifically, our research will include the following tasks: use LIM on Blue Waters to make transistor-level simulations involving a large numbers of devices (this would not be feasible without Blue Waters); develop the circuit synthesis of blackbox macromodels with large numbers of ports; develop techniques to parse, analyze, and process the big data generated by measurement or calculation of X parameters; and process the generated data to develop deterministic and stochastic simulation methods for analog blocks using Volterra series and the latency insertion method (LIM) on the Blue Waters platform.

NEXT GENERATION WORK

In the 2019-2020 time frame, we plan to demonstrate the simulation of integrated circuits with large numbers of transistors as well as high-speed links.

HARDWARE ACCELERATION OF DEEP LEARNING

Allocation: Illinois/50.0 Knh
PI: Tao Xie[1]
Co-PI: Yuan Xie[2]

[1]University of Illinois at Urbana-Champaign
[2]University of California at Santa Barbara

EXECUTIVE SUMMARY

Our project aims to use Blue Waters for hardware acceleration of deep learning for big data image analytics. To achieve near real-time learning, efforts are required for hardware scaling out (increasing the number of compute nodes in a cluster) and scaling up (improving the throughput of a single node by inserting hardware accelerators). We evaluated the performance of scaling up using the graphics processing unit (GPU) enabled node XK7 for training convolutional neural networks. Our key observation thus far is that implicit data synchronization across different nodes severely slows down the training process. We propose a data manager that explicitly covers the data transfer overhead with computation. Our first step was to test the proposed strategy on a single XK7 node of Blue Waters, and results show a speedup of 1.6 times that of the implicit data transfer implementation.

INTRODUCTION

Deep learning has been used in applications such as image classification, speech processing, and object recognition. A very large amount of training data is necessary for deep neural networks, therefore this work requires computing power capable of matching the advanced state-of-the-art accuracy of these tasks. Mainstream deep learning facilities are central processing unit (CPU)-based clusters, which usually consist of thousands of compute nodes. As the major computation of deep learning is convolution and matrix multiplication, which are suitable for GPUs to process, most modern deep learning facilities are equipped with GPUs as hardware accelerators. However, straightforward implementation of deep neural networks on GPU-enabled compute nodes will lead to under-utilization of computing resources. In this work, we evaluated the performance of a convolution neural network on a GPU-enabled cluster (XK7 nodes on Blue Water). We observed that the performance bottleneck of the convolutional neural network training is the implicit data synchronization across different nodes, which is used to confirm the correctness of the output of each layer in the neural network that is distributed in different nodes. To alleviate the performance degradation caused by the synchronization, we propose a method to hide the data transfer explicitly by computation. This method dramatically increases the utilization of the compute units on each GPU chip and thus improves the overall training performance.

METHODS & RESULTS

Traditionally distributed convolutional neural networks allocate one compute node for each replica as part of the training model. The training process of each layer in each replica consists of three phases: 1) receiving output data from the previous layer of other replicas and feeding them into the device memory; 2) executing the device kernel to compute the output data of current layer; and 3) sending the output data of the current layer to other replicas.

The performance bottleneck of this straightforward implementation is that steps 1 and 3 introduce very long latency, which is determined by the longest unpredictable network latency between different compute nodes. The reason for the performance loss is under-utilization of hardware resources of the GPU accelerators. In Steps 1 and 3, compute units are completely idle while waiting for data synchronization. In Step 2, the direct memory access units are inactive while the compute units are executing the kernel functions. Therefore, overlapping the data synchronization and the kernel execution will reduce the total running time of the training process.

To combat performance loss due to the under-utilization of hardware resources, we break down the replicas into finer grained replicas with the same number of compute nodes used. This way, the kernel execution of one replica can run simultaneously with data synchronization of other replicas. Ideally, if the replica break-down does not introduce any communication overhead, there will be more replicas in each node and better performance will be achieved. However, the communication overhead is not negligible, so we have to find the best number of replicas per node. We find the best ratio is two replicas per node, and the speedup of this configuration is 1.6 times over that of straightforward implementation.

WHY BLUE WATERS

Blue Waters offers XK7 nodes which consist of one AMD 8 floating point core CPU and one NVIDIA K20X GPU. As GPUs are more suitable than CPUs for convolution and matrix multiplications, state-of-the-art deep learning facilities widely employ GPUs as their hardware accelerators. Blue Waters offers an opportunity to perform research on equipment optimized for deep learning cluster with GPUs. Also, Blue Waters provides a CUDA API programming environment, which enables us to customize the specific functions to be executed on GPUs.

NEXT GENERATION WORK

A field-programmable gate array (FPGA) is an alternative hardware accelerator to GPUs. It can be quickly used to prototype new hardware designs with very high power efficiency. In the next generation work plan could employ a FPGA to accelerate deep learning algorithms.

BIOLOGY, CHEMISTRY & HEALTH

MOLECULAR

CELLULAR

MEDICINE

BIOPHYSICS

GENOMICS

BIOINFORMATICS

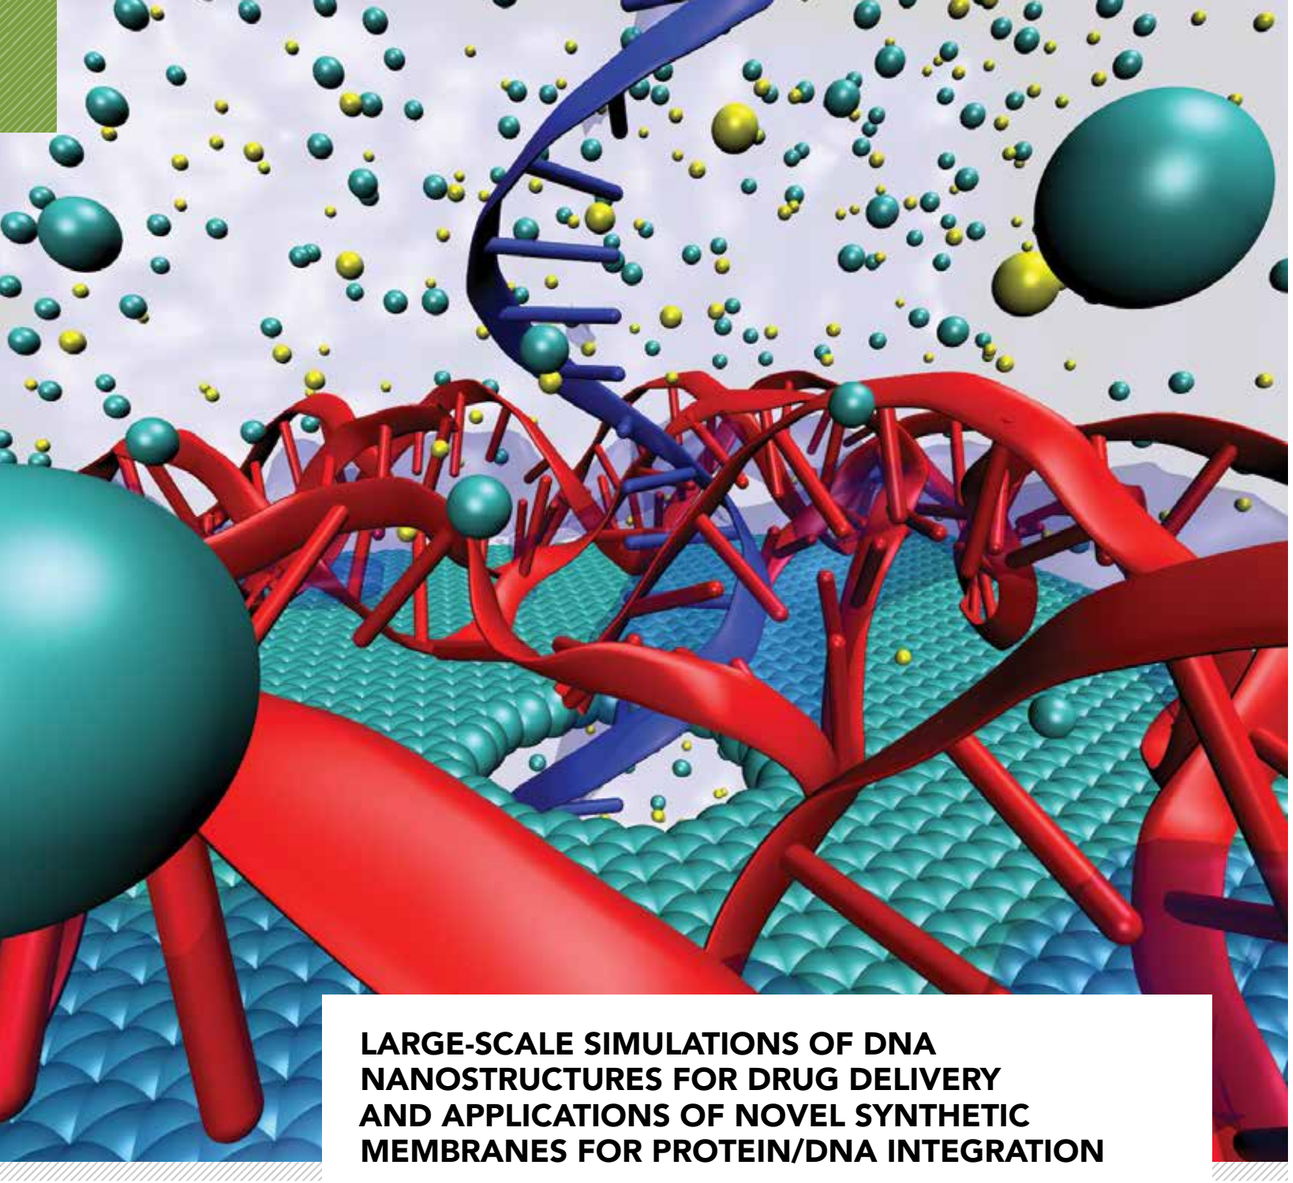

FIGURE 1: Single strand of DNA (blue) translocating through a single layer DNA origami (red) mounted on a graphene sheet (cyan).

LARGE-SCALE SIMULATIONS OF DNA NANOSTRUCTURES FOR DRUG DELIVERY AND APPLICATIONS OF NOVEL SYNTHETIC MEMBRANES FOR PROTEIN/DNA INTEGRATION

Allocation: Illinois/700 Knh
PI: Narayana R. Aluru[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

DNA nanostructures have the advantage of being biocompatible and programmable. The self-assembly features of DNA can be used to design nanopores within the DNA structure for sensing and detecting biological molecules. DNA base detection, which is important in mapping human health, faces a key challenge of high translocation speed in nanopores. Here, we propose a DNA origami-graphene nanopore for DNA sequencing. The DNA origami sits on top of a graphene sheet acting as a substrate. We show that the translocation speed of DNA strands inside the DNA origami nanopore is reduced due to the attractive interactions of nucleotides of the pore and the DNA strands. We also find that the proposed nanopore leads to distinct signals for different types of DNA bases making it a **promising** hybrid material (DNA origami-graphene) for DNA detection.

METHODS AND RESULTS

DNA sequencing using nanopore technology has evolved significantly during the last few years. Oxford nanopore technology is currently fabricating a USB-stick size device that can sequence the DNA rapidly—in a couple of hours. In recent years, both biological and synthetic nanopores were used for "label-free," high-resolution DNA sequencing. The challenges posed to DNA sequencing using nanopore technology are the low signal to noise ratio, pore degradation due to multiple uses, the identification of single base in real time and the high speed of translocation [1,2]. Engineering the translocation of DNA through biological/synthetic nanopore has been defined as one of the challenging problems of biotechnology. Although a number of studies have been performed on different types of pores regarding translocation speed and ionic current blockade, the fundamental understanding of the pore architecture, material and size of the pore on the quality of DNA sequencing and speed of translocation is still lacking. In this study, we will examine DNA origami-graphene hybrid nanopores for DNA sequencing by using extensive molecular dynamics (MD) simulations.

The self-assembly properties of atoms and molecules have frequently been used to create arbitrary 3D nanostructures. The DNA origami technique is a **novel** method, which takes advantage of self-organization properties of DNA molecules and allows folding DNA single strands to construct a complex shape at nanoscale. DNA origami sheets can be created via self-assembly process. Here, we introduce a nanopore in a DNA origami nano sheet mounted on top of a graphene sheet acting as a substrate (Fig. 1). Using extensive MD simulations, this nanopore is shown to be able to detect and sense DNA bases. The DNA origami nanopore is designed in such a way that the edge of the pore carries a specific nucleotide type. We show that the nanopore in DNA origami-graphene results in distinct dwell times for the four DNA base types, while a bare graphene nanopore gives rise to almost identical dwell times for the four DNA base types. The difference in dwell times is due to the strength of attractive interaction between the nanopore and DNA translocating strand. Depending on the nanopore edge functionalization, the base pairing between the bases of the pore and strand is different for different translocating bases. In addition to these distinguishable dwell times, this base pairing results in high residence time of the DNA strand inside the pore or lower speed of translocation.

WHY BLUE WATERS

We performed extensive molecular dynamics simulations which involve up to 120,000 atoms. These expensive computations are not possible to perform without a petascale supercomputer. Also, the MD package (NAMD) we used scales almost linearly with the number of cores up to 1,000 in our test on Blue Waters.

VIRTUAL FLU, A DIFFERENT KIND OF COMPUTER VIRUS

FIGURE 1: The influenza surface coat. The hemagglutinin and neuraminidase glycoproteins are represented in lavender and magenta, respectively. The enveloping lipid bilayer is represented in glossy pink. Credit: Jacob D. Durrant.

Allocation: NSF PRAC/6.00 Mnh
PI: Rommie E. Amaro[1]
Collaborator: Alasdair Steven[2]

[1]University of California, San Diego
[2]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Influenza infection is routinely responsible for hundreds of thousands of deaths annually, punctuated by catastrophic pandemics roughly every 25 years. Some experimental techniques used to study the virus suffer from serious drawbacks. For example, structural-biology methods like electron microscopy, X-ray crystallography, and nuclear magnetic resonance spectroscopy (NMR) typically cannot provide either the atomic resolution, size scaling, or molecular-dynamics data required to answer a number of pharmacologically important questions.

To address these concerns, we have constructed an atomic-resolution model of the **entire influenza viral coat**, containing 160 million atoms. We recently simulated this large-scale system for 121 nanoseconds in order to gain a more complete understanding of the influenza infection process. This research now allows us to explore novel opportunities for drug and vaccine development *in silico*.

INTRODUCTION

We focused our research on the viral surface coat because of the important role it plays in both the initial and final stages of the influenza infection process. The viral coat is composed of a lipid bilayer from which two spike-like glycoproteins, neuraminidase and hemagglutinin, protrude (Fig. 1). When a viral particle first approaches a human host cell, the hemagglutinin proteins latch onto sialic acid molecules attached to the cell surface, ultimately leading to the molecules being transported into the cell. A second surface-coat protein, the M2 proton channel, acidifies the virus interior once the virion enters the endosome pocket within the cell, triggering conformational changes that facilitate replication of the virus [1]. Following replication, the viral progeny bud from the host cell but remain attached to its external surface by the same sialic acid connections. Influenza's neuraminidase glycoprotein is responsible for severing those tethers, allowing the newly formed viruses to depart and infect the next cell [2].

The viral coat has been extensively studied precisely because it is so critical for infection. By analyzing the molecular structures of the various components of the surface coat in isolation, researchers have produced a number of anti-flu drugs currently used clinically (e.g., Tamiflu). Unfortunately, the flu is highly adaptable, and resistance to these medicines has already been documented [1, 3-5]. There is an urgent need for novel therapeutics. Studying the structures and motions of the various surface-coat components when assembled into their natural multi-component environment, rather than in isolation, provides pharmacologically relevant insights that will help us combat future pandemics. No current experimental technique is capable of providing an accurate spatial and temporal model of the entire *dynamic* surface coat at the resolution needed for drug discovery; fortunately, modeling and simulation can serve as a "computational microscope" that provide the needed information.

METHODS & RESULTS

Guided by experimental data, we constructed a surface-coat model of the 2009 H1N1 pandemic flu strain. When immersed in a bath of virtual water with the appropriate electrolytes, each system contains over 160 million atoms (Fig. 2). Our collaborator Alasdair Steven at the University of Illinois at Urbana-Champaign used electron microscopy to identify the general shape of the influenza virus and the approximate locations of the glycoprotein spikes [6]. We then used computational methods developed in our lab to wrap this virus volume in a virtual lipid bilayer [7] and to position atomic-resolution models of the glycoproteins at the appropriate locations. An appropriate number of M2 channels were also added to the structure. We were thus able to transform the low-resolution microscopic data into a high-resolution, atomistic model suitable for molecular dynamics simulations and, ultimately, drug discovery. We finished simulations of this virtual virus several months ago on Blue Waters. Ongoing and preliminary analyses of the simulations have suggested several interesting findings. First, we analyzed the apo neuraminidase molecules scattered across the virion surface. The R371 and R292 residues, which form electrostatic interactions with most endogenous and pharmaceutical ligands, sample many apo conformations other than that of the crystal structure, suggesting that binding occurs via a conformational-shift mechanism. The

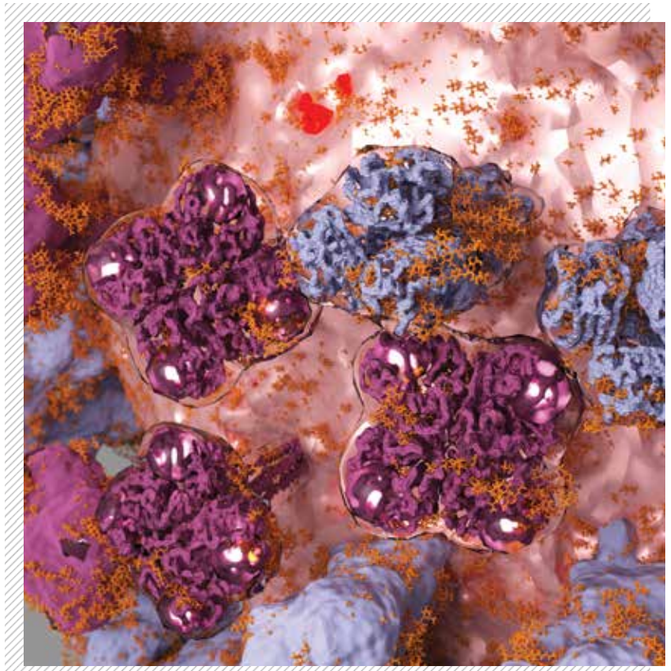

FIGURE 2: The viral surface coat. Selected neuraminidase and hemagglutinin glycoproteins are rendered transparent, revealing the underlying protein backbone. An M2 proton channel, in red, is also visible. This image highlights the fact that these renderings were generated from high-resolution (ultimately atomic-resolution) data. Credit: Jacob D. Durrant.

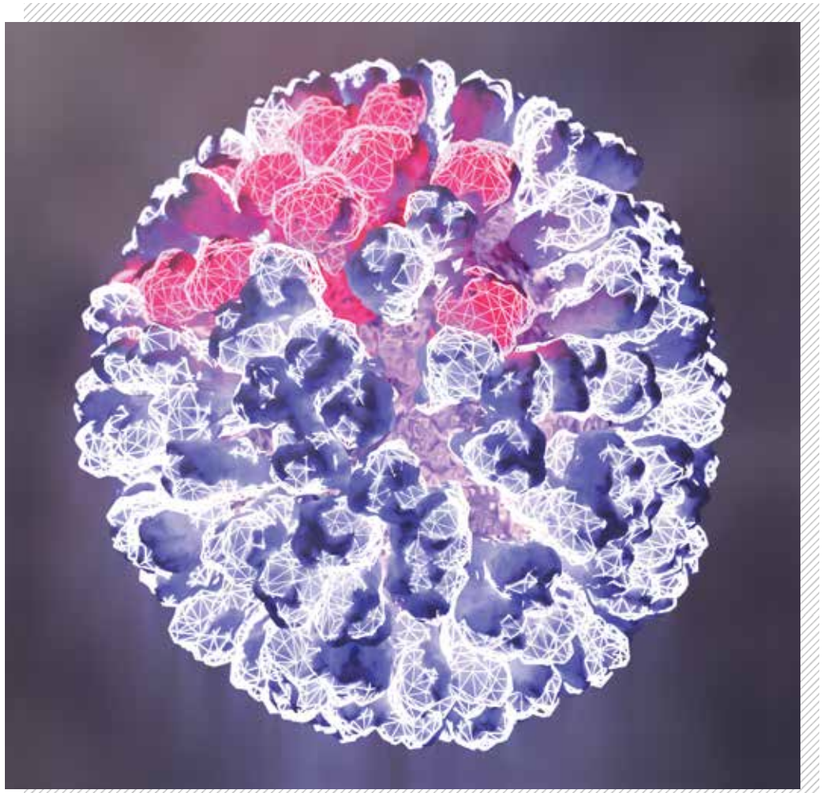

FIGURE 3: Surface-coat dynamics: the first and last frames of the whole-virion simulation. Neuraminidase and hemagglutinin are shown in red and blue, respectively. The glycoprotein conformations of the first and last frames are shown in solid and glowing-mesh representations, respectively. As expected for a simulation of this duration, glycoprotein diffusion through the lipid bilayer was limited. However, the atomic-resolution motions of individual glycoproteins were substantial because they were sampled over both simulated time and space, thanks to the multiple copies of each glycoprotein scattered across the model surface. Credit: Jacob D. Durrant.

motions of these residues permit conformations that have additional druggable hotspots beyond those of the sialic acid-binding and 150-loop regions. Furthermore, we are currently building Markov state models to explore the pharmacologically relevant kinetics of 150-loop opening and closing, and the neuraminidase conformations sampled by this large-scale simulation may prove useful for future virtual-screening efforts as well.

We are similarly analyzing the many hemagglutinin molecules included in the whole-virion-coat simulations. We are hopeful that Brownian dynamics simulations in the context of the whole viral particle will provide useful insights into the mechanism of broadly neutralizing antibodies.

The M2 channels of our model also sample many conformations, ranging from open to closed. We have characterized the volume distributions of these channels and hope to build a Markov state model to describe the opening/closing kinetics of this crucial surface-coat component. The many M2-channel conformations sampled may also prove useful in future small-molecule virtual screens.

WHY BLUE WATERS

Blue Waters has been critical for this project. To our knowledge, a molecular dynamics simulation on so grand a scale has **never before** been attempted. Very few supercomputers are capable of the petascale performance required. **Without Blue Waters, the current work would be impossible.**

These simulations are providing important information about the surface motions and electric fields that surround the viral particle. These "dynamics" and "electrostatics" govern not only the infection process, but also drug and vaccine/antibody binding.

MOLECULAR DYNAMICS OF SELF-ASSEMBLED DNA SYSTEMS

Allocation: Blue Waters Professor/240 Knh
PI: Aleksei Aksimentiev[1]
Collaborators: Chen-Yu Li[1], Jejoong Yoo[1], Christopher Maffeo[1], Scott Michael Slone[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

DNA nanotechnology utilizes self-assembly for the high-throughput construction of sub-micron-size objects with nanometer precision. In comparison to conventional nanofabrication approaches, the DNA origami method is relatively low cost, easy to use, and has an infinite number of possible applications. Using Blue Waters, we have explored the ability of DNA nanostructures to function as membrane channels, carried out a **landmark simulation** of a DNA origami sculpture, and characterized the mechanical properties and ionic conductivity of DNA brick structures. The results of our simulations have contributed to the development of a web server for prediction of DNA origami structures and a web tool for designing nanostructures using the DNA brick methods.

INTRODUCTION

DNA origami is an experimental technique that allows folding of a long DNA molecule into an arbitrary three-dimensional shape [1]. Over the past ten years, the DNA origami method has **advanced** to encompass self-assembly of complex 3D objects with sub-nanometer precision including static structures [2], as well as objects that perform active functions [3].

Predictive computational modeling of DNA origami is an attractive alternative to experimental characterization of such self-assembled objects. Currently, the most accurate computational method is all-atom molecular dynamics (MD). In 2013, we reported the **first** MD study of several model DNA origami systems [4]. Last year, we explored the possibility of using DNA origami in nanopore sensing applications [5]. Our most recent work includes a study of a DNA-based channel embedded in a lipid bilayer, a **landmark** simulation of a DNA sculpture that demonstrates the predictive power of the MD method, and a detailed comparison between DNA origami and an alternative self-assembly approach known as DNA bricks.

METHODS & RESULTS

In living cells, membrane protein channels control the transport of molecules across the cell membrane. Recently, several experimental groups demonstrated assembly and insertion of DNA channels into lipid bilayer membranes [6]. A typical DNA channel is made by arranging several parallel DNA double helices to form a polygon. The central cavity of the polygon is the transmembrane pore. Using

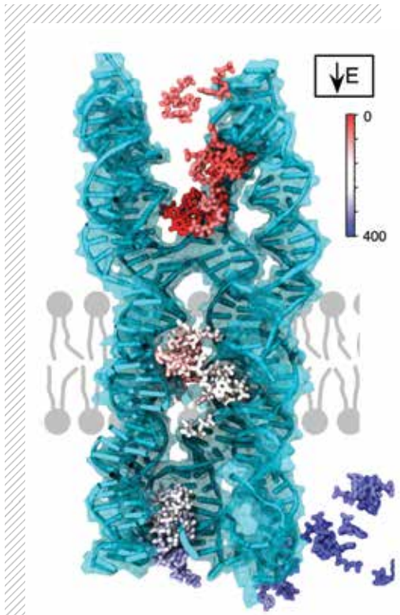

FIGURE 1: A time-lapse illustration of the MD trajectory showing an ATP molecule passing from one side of the DNA channel to the other. The color indicates the progress of the simulation using the red (beginning)–white–blue (end) scheme. Because of the electro-osmotic flow, the ATP molecule moves in the direction of the electric field, opposite to the direction prescribed by its negative charge.

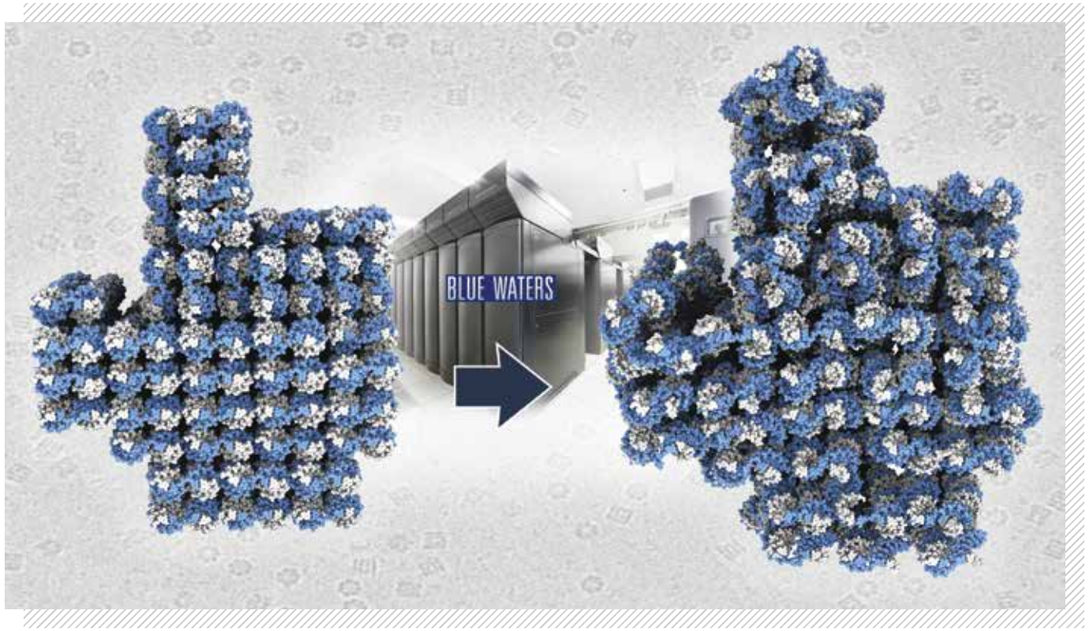

FIGURE 2: Molecular dynamics simulation predicts the solution structure of a 5 megadalton DNA origami object. The background image was adapted from "CryoEM GroEL" by Vossman under a Creative Commons license.

Blue Waters, we carried out the **first** all-atom MD study of the DNA membrane channels. We showed that, while overall remaining stable, the local structure of the channels undergoes considerable fluctuations, departing from the idealized design. The transmembrane ionic current flows both through the central pore of the channel as well as along the DNA walls and through the gaps in the DNA structure. Surprisingly, we found the conductance of DNA channels to depend on the membrane tension, making them potentially suitable for force-sensing applications. Finally, we showed that electro-osmosis governs the transport of drug-like molecules through the DNA channels.

To test the accuracy of our all-atom MD method, we carried out a **landmark** simulation of a DNA origami sculpture whose 3D structure was determined through cryo-electron microscopy [7]. Over the course of the all-atom explicit solvent simulation, the structure was observed to depart from its ideal design, changing its shape toward the experimentally determined structure. Next, we showed that elastic network-guided simulations performed without solvent could yield similarly accurate structural models at a **fraction** of the computational cost, making them suitable for **prototyping** and validation of self-assembled DNA nanostructures. A web-server implementation of our elastic network-guided methods has made it available to a **broader** scientific community.

In contrast to DNA origami, the DNA bricks [8] method produces custom three-dimensional objects using only short DNA fragments. As a result of their design, the assembled DNA brick structures have fewer inter-helical connections than equivalent DNA origami structures. Using the MD method, we directly compared the structure, mechanical properties and ionic conductivity of DNA brick and DNA origami structures. In comparison to equivalent DNA origami structures, the DNA brick structures were found to be less rigid and less dense. Subject to an external electric field, a DNA brick plate was found to be more permeable to ions than an equivalent DNA origami plate because of its lower density and larger cross-section area. Based on the results of this study, we have developed a web tool that considerably **simplifies** the design of DNA brick structures.

WHY BLUE WATERS

Explicit solvent all-atom MD simulation is the only computational method that can treat DNA origami objects enhanced by non-standard functional groups and characterize their transport properties. Because of the size of the DNA origami structures, such MD simulations are computationally demanding. The large number of XK nodes on Blue Waters with graphics processing unit accelerators connected by the fast Gemini interconnect makes it one of the best publicly available systems for performing DNA origami simulations.

NEXT GENERATION WORK

Using the next-generation Track-1 system, we hope to apply our all-atom MD method to dynamic and externally actuated DNA nanostructures, exploring their applications in nanoscale engineering.

PUBLICATIONS AND DATA SETS

Yoo, J., and A. Aksimentiev, Molecular dynamics of membrane-spanning DNA channels: conductance mechanism, electro-osmotic transport, and mechanical gating. *J. Phys. Chem. Lett.*, 6:23 (2015), pp. 4680—4687.

Maffeo, C., J. Yoo, and A. Aksimentiev, De Novo Reconstruction of DNA Origami Structures through atomistic molecular dynamics simulation. *Nucleic Acids Res.*, 44:7 (2016), pp. 3013—3019.

Slone, S., C. Li, J. Yoo, and A. Aksimentiev, Molecular mechanics of DNA bricks: In situ structure, mechanical properties and ionic conductivity. *New J. Phys.*, 18: 055012 (2016).

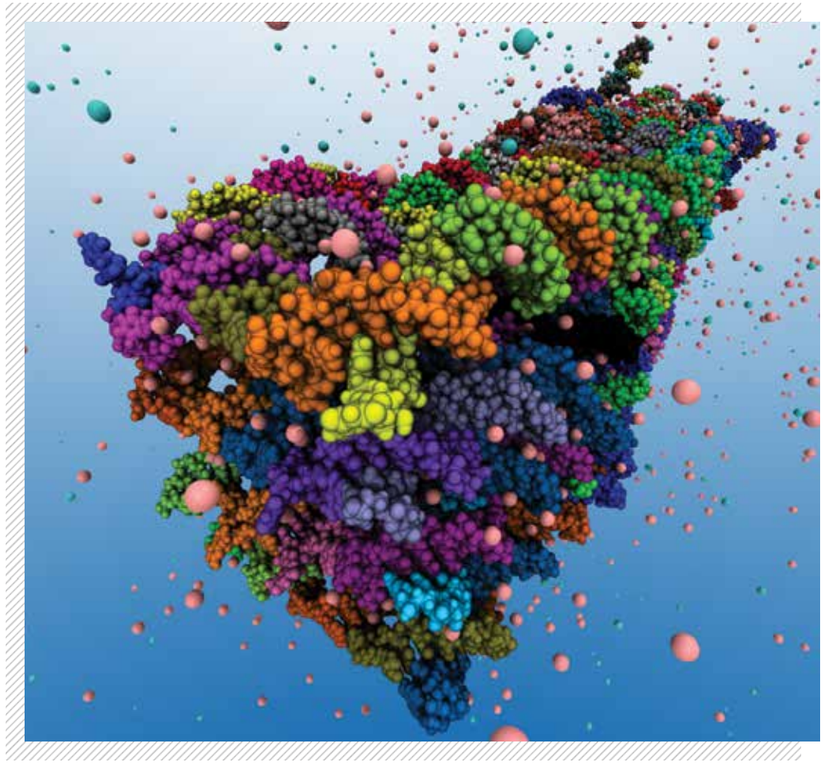

FIGURE 3: The microscopic conformation of a DNA brick structure at the end of a molecular dynamics simulation. Individual DNA strands are shown in distinct colors; magnesium and chloride ions are shown as green and pink spheres, respectively, water is not shown.

MOLECULAR MECHANISM OF SEQUENCE-DEPENDENT DNA LOOPING

Allocation: Illinois/1.20 Mnh
PI: Aleksei Aksimentiev[1]
Collaborators: Jejoong Yoo[1], Hajin Kim[2], Thuy Ngo[1], James Wilson[1], Taekjip Ha[3]

[1]University of Illinois at Urbana-Champaign
[2]Ulsan National Institute of Science and Technology, Ulsan, Korea
[3]Johns Hopkins University

EXECUTIVE SUMMARY

The molecular mechanisms underlying regulation of gene expression are central to cell biology but are still not fully understood. Our group uses the Blue Waters supercomputer to determine how the intrinsic properties of DNA affect gene regulation. One mechanism of gene regulation involves the bending of rigid DNA molecules around histone proteins, restricting access to the encoded information. Using molecular dynamics simulations, we elucidated the effect of epigenetic modifications on the bending propensity of DNA molecules. Another mechanism of gene regulation involves restructuring entire collections of genes. Using enhanced sampling methods, we have shown that AT-rich and methylated DNA experience enhanced attraction at physiological conditions, capable of causing large-scale chromosome compaction. Both mechanisms support the notion of "hidden" codes in DNA that contributes to the regulation of gene expression.

INTRODUCTION

Eukaryotic cells store their genetic information in long (millions of base pairs) molecules of DNA. Inside the nucleus, the DNA is wrapped around histone proteins into bundles called nucleosomes, each consisting of approximately 150 base pairs of DNA forming a double-loop around a histone core. At a given stage of a cell's life cycle, only a fraction of a cell's genes are active; the collection of active genes determines the function of a cell. One method of turning genes on and off is through compaction of the DNA, which makes its genetic code inaccessible to the transcription machinery [1,2,3]. Here, we describe two mechanisms that can realize DNA compaction: one that occurs inside a single nucleosome, and another that occurs among multiple nucleosomes. Within a nucleosome, decreasing DNA flexibility increases the likelihood of DNA unwrapping from the nucleosome, exposing the DNA's genetic code. In the case of multiple nucleosomes, attractive interactions between the nucleosomes can modulate the accessibility of multiple DNA fragments. In collaboration with the Ha group (Johns Hopkins University), we demonstrate that flexibility of individual DNA fragments and mutual attraction between several DNA fragments depend on the DNA sequence and can be modulated by chemical modifications of the DNA in the absence of any auxiliary proteins.

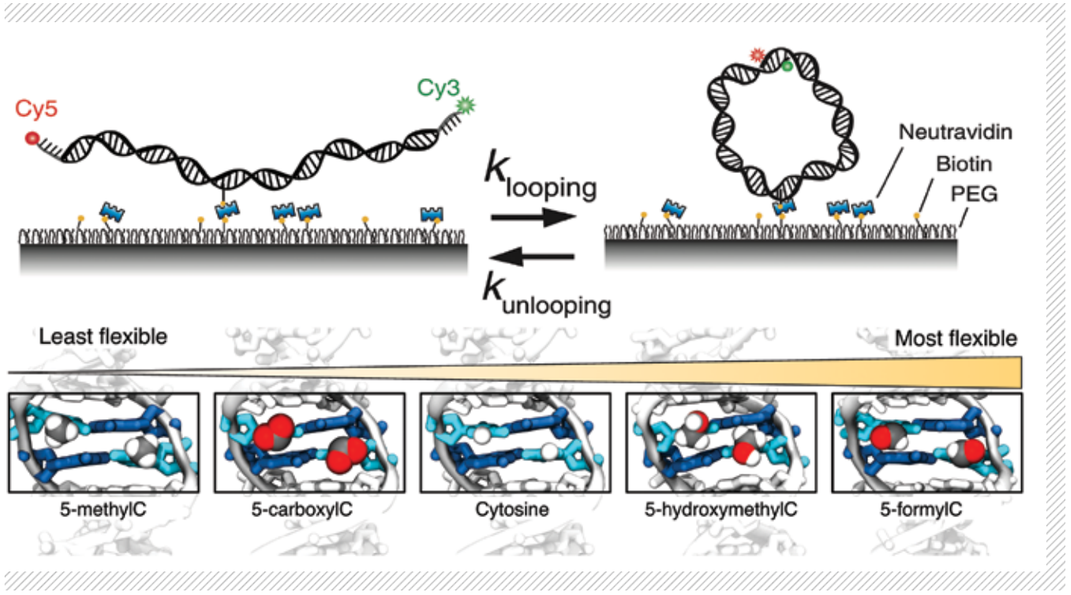

FIGURE 1: Schematic representation of single-molecule DNA cyclization assay. Population of looped DNA is monitored over time using the Fluorescence Resonance Energy Transfer technique. The modified CpG sites are ranked according to their simulated flexibility: 5-formyl C > 5-hydroxymethyl C ≥ 5-carboxyl C ≥ unmodified C > 5-methyl C.

METHODS & RESULTS

In collaboration with the Ha group, we characterized the effect of chemical modifications of cytosine (C) on DNA flexibility. We tested the following four biologically significant modifications: 5-methylcytosine (5-mC), 5-hydroxymethylcytosine (5-hmC), 5-formylcytosine (5-fC), and 5-carboxylcytosine (5-caC). In both simulation and experiment, the 5-fC modification was found to increase DNA flexibility considerably, 5-hmC to enhance flexibility but to a lesser degree than 5-fC, 5-mC to reduce flexibility, and 5-caC to have no measurable effect. Analysis of the molecular dynamics (MD) trajectories showed that chemical modifications affect DNA flexibility via steric, hydrophobic, and electrostatic mechanisms. The modulation of DNA flexibility increased the mechanical stability of the nucleosome and vice versa, suggesting a gene regulation mechanism where cytosine modifications change the accessibility of nucleosomal DNA through their effects on DNA flexibility. In a separate study performed in collaboration with the Meni Wanunu group (Northeastern University), we have shown that oxidized products of thymine (T) can also enhance the flexibility and hydrophilicity of double-stranded DNA.

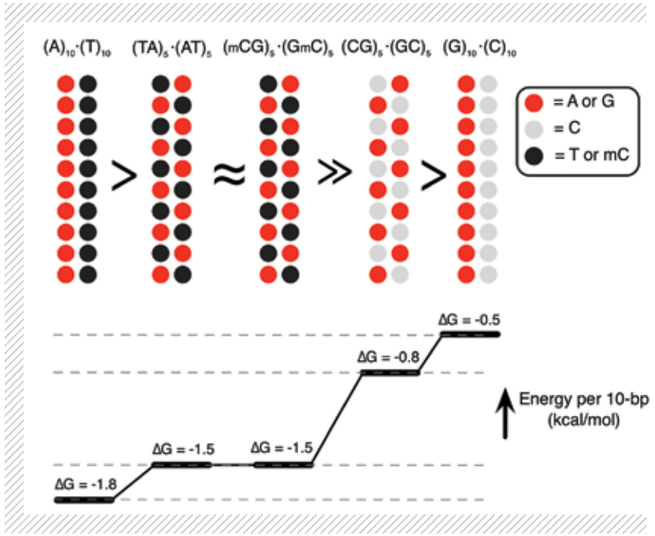

FIGURE 2: The number and spatial arrangement of nucleotides carrying a methyl group (T or mC) determine the interaction free energy of two DNA molecules.

By combining single molecule experiments and MD simulations, we demonstrated a **novel** mechanism of DNA compaction controlled by the methylation pattern of DNA. Prior experiments established a correlation between DNA compaction (high-order chromatin folding) and the nucleotide content: DNA segments containing a higher than average percentage of AT base pairs, a larger than average number of methylated cytosine nucleotides, or both, were found to form more compact folds. One popular explanation of this observation posits the existence of yet unidentified proteins that recognize DNA sequences and their modifications and bring the recognized DNA domains together [4]. Alternatively, the DNA compaction can be induced by direct physical interactions between DNA segments mediated by polyamines, small charged molecules abundant in living cells. To test the latter possibility, we performed a series of free energy calculations at a physiological concentration of spermine, a model polyamine. The computed free energies clearly showed that AT-rich DNA domains attract each other more strongly than GC-rich domains do. Furthermore, methylation of cytosines was found to increase the attraction between two GC-rich domains to the level of two AT-rich domains. Subsequent single molecule experiments confirmed the predictions of our MD simulations.

Our findings suggest a **tantalizing** possibility that polyamine-mediated inter-DNA attraction can play a major role in high-level chromatin folding.

Our inquiry into the molecular mechanism of DNA—DNA interactions found that the standard parameterization of nonbonded interactions in popular MD force fields, CHARMM and AMBER, are not accurate enough to characterize the dependence of the DNA—DNA interactions on the DNA sequence and its modifications. For example, the effective force between two DNA molecules in a 100-mM di-lysine (Lys2+) solution is attractive in both standard CHARMM and AMBER models, whereas the DNA molecules are experimentally known to repel one another at identical conditions. To improve the accuracy of the MD force fields, we reparameterized the strength of amine–phosphate and amine–carboxylate interactions against independent sets of osmotic pressure data. Our extensive validation simulations performed on Blue Waters have shown that our improved parameter set can **significantly enhance** the realism of MD simulations for a broad class of biomolecular systems, including protein folding, protein–DNA interactions, and DNA condensation.

FIGURE 3: Refinement of non-bonded interactions for accurate simulations of inter-molecular forces. The interactions indicated by arrows were reparameterized to reproduce the experimental osmotic pressure data.

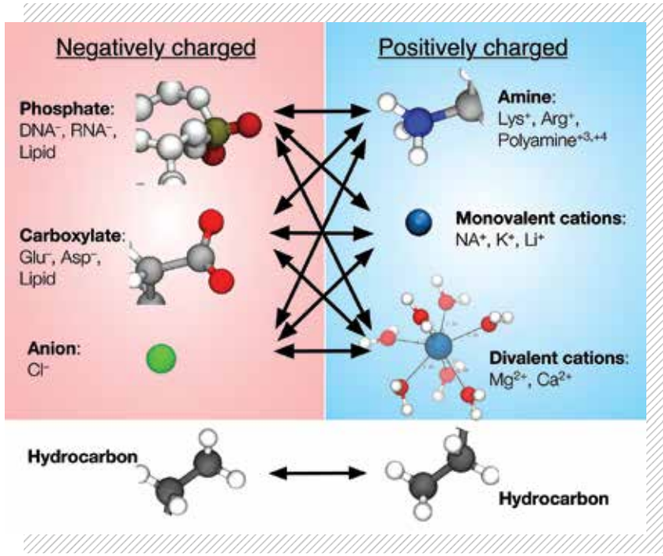

WHY BLUE WATERS

Extensive sampling of biomolecular conformations was essential for characterization of DNA flexibility and calculation of inter-DNA forces. Using Blue Waters was essential to achieve the unprecedented accuracy of our simulations that matched and sometimes exceeded state-of-the-art experimental techniques.

NEXT GENERATION WORK

Using Blue Waters, we plan to create a genome-wide map of DNA flexibility that will elucidate the effect of DNA sequence on gene regulation.

PUBLICATIONS AND DATA SETS

Yoo, J., H. Kim, A. Aksimentiev, and T. Ha. Direct evidence for sequence-dependent attraction between double-stranded DNA controlled by methylation, *Nat. Commun.*, 7:11045, 2016. DOI:10.1038/ncomms11045

Ngo T., et al., Effect of cytosine modifications on DNA flexibility and nucleosome mechanical stability, *Nat. Commun.*, 7:10813, 2016. DOI:10.1038/ncomms10813

Yoo, J. and A. Aksimentiev. The structure and intermolecular forces of DNA condensate, *Nucleic Acids Res.*, 44:2036–2046, 2016. DOI:10.1093/nar/gkw081

Yoo, J. and A., Aksimentiev. Improved parameterization of amine–carboxylate and amine–phosphate interactions for molecular dynamics simulations using the CHARMM and AMBER force fields, *J. Chem. Theory Comput.*, 12:430–443, 2016. DOI:10.1021/acs.jctc.5b00967

Yoo, J., J. Wilson, and A. Aksimentiev. Improved model of hydrated calcium ion for molecular dynamics simulations using classical biomolecular force fields, *Biopolymers*, DOI:10.1002/bip.22868

Carson, S., et al., Hydroxymethyluracil Modifications Enhance the Flexibility and Hydrophilicity of Double-Stranded DNA. Nucleic Acids Res., 44: 2085–2092 (2016). DOI: 10.1093/nar/gkv1199

SEQUENCE SIMILARITY NETWORKS FOR THE PROTEIN "UNIVERSE"

Allocation: Illinois/744 Knh
PI: John A. Gerlt[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

We are devising strategies and tools to facilitate prediction of the *in vitro* activities and *in vivo* metabolic functions of uncharacterized enzymes discovered in genome projects. We used Blue Waters to establish a protocol for generating a library of sequence similarity networks (SSNs) for all Pfam protein families in the UniProt protein sequence databases for dissemination to the scientific community. We have calculated 1) all-by-all Basic Local Alignment Search Tool (BLAST) sequence relationships, 2) statistical analyses of the BLAST results; and 3) merged sets of input sequences based on sequence identity. Based on our experiences, we have defined protocols for the regular (every eight weeks) generation of the library of sequence similarity networks.

INTRODUCTION

The current UniProtKB database contains more than 60M nonredundant sequences. The functions for less than 1% of the entries have been manually curated; the functional annotations for the remaining entries are assigned by automated procedures. As a result, the conservative estimate is that the annotations for at least 50% of the entries are uncertain or incorrect. The majority of the entries are obtained from genome sequencing projects, the rationale being that knowledge of the complete complement of proteins and enzymes encoded by an organism will allow its biological and physiological capabilities to be understood. However, if at least 50% of the proteins and enzymes have uncertain or unknown functions, the considerable investments in genome projects cannot be realized. Because of the very large number of proteins and enzymes for which sequences have or will become available, strategies for predicting their functions must be high throughput and large scale, i.e., computation based.

METHODS & RESULTS

During the past year, we have continued to develop strategies to maximize the usage of RAM to enable the all-by-all sequence comparison using BLAST for the largest Pfam protein families. This has been problematic due to 1) wall time restrictions of 24 hours that recently were increased to 48 hours, and 2) the limited amount of RAM that is available for the input and output (64 GB/node) that makes node usage "inefficient."

We now are exploring the use of DIAMOND, a recently developed alternative to BLAST, for the all-by-all sequence comparisons. We have observed that DIAMOND provides a greater than or equal to 10-fold increase in the rate of all-by-all sequence comparisons relative to BLAST for almost all Pfam protein families. This significant decrease in time "solves" the wall time problem for all Pfam protein families and allows a greater number of nodes to be assigned to the largest Pfam protein families, thereby allowing the all-by-all sequence comparisons to be accomplished. We expect to be able to begin the production phase of this project in which the library of SSNs for all 16,295 Pfam protein families can be updated every eight weeks (with each update of the InterPro protein sequence analysis and classification database).

WHY BLUE WATERS

The project uses an embarrassingly parallel computing model to perform the all-by-all sequence comparison and, in principle, could be run on any cluster of sufficient size. However, because of 1) the scale of the computation (number and sizes of Pfam protein families) and 2) the time sensitivity of the production of the output relative to IntroPro database updates, only a resource at the scale of Blue Waters can perform the job in a reasonable time frame.

NEXT GENERATION WORK

We hope to more efficiently (with respect to RAM and node usage) perform the all-by-all sequence comparisons for all Pfam protein families, and eventually larger clans so that these can be disseminated to the community with each update of the InterPro database.

MAKING ANCESTRAL TREES USING BAYESIAN INFERENCE TO IDENTIFY DISEASE-CAUSING GENETIC VARIANTS

Allocation: Illinois/50.0 Knh
PI: Don Armstrong[1]
Collaborators: Derek Wildman[1] and Monica Uddin[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

A complex network of different interacting genetic variants contributes to many common human diseases. Discovering effective treatments for these diseases requires distinguishing the genetic variants that cause diseases from those variants that are just associated with the disease. Ancestral trees will enable us to identify variants which are associated merely due to ancestry, and decrease the rates of false positives in case-control studies due to admixture. The size of the human genomes (3.2 Gbp x 2000 ≈ 6 TiB) coupled with the gigantic number of possible trees (2000!! or $\approx 4\times10^{2868}$) requires supercomputing on the scale of Blue Waters. In a preliminary allocation, we were able to demonstrate the feasibility of utilizing Blue Waters to generate ancestral trees on large datasets (6 TB) using maximum likelihood and were able to identify bottlenecks in the Bayesian approach.

FIGURE 1: Linear scaling of time in seconds to calculate 10 trees after the 10th, 20th, etc. tree with number of Blue Waters XE nodes starting with 1 million variants and 1000 individuals with the same run parameters.

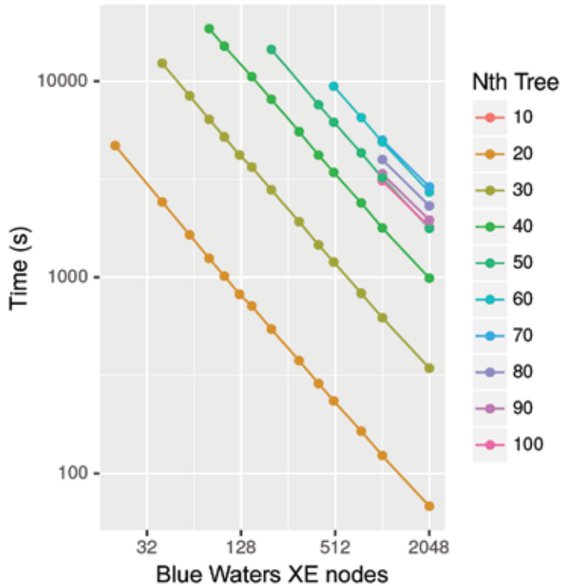

INTRODUCTION

Many human diseases such as diabetes, cancer, cardiovascular disease, and mental illness are caused in part by a complex network of genetic variants which interact with each other and environmental factors. Ancestral trees, which depict the descent from ancestors of a set of genomic regions, enable us to: 1) distinguish between common and rare variants; 2) identify variants which are associated with disorders; 3) identify sets of cases and controls which are ethnically matched for a particular genomic region. However, the generation of trees is nondeterministic polynomial time (NP) complex; the number of rooted binary trees for 2000 individuals is 2000!!, or $\approx 4\times10^{2868}$. When coupled with the size of the human genome (3.2 Gbp) by 2000 individuals (6 TB of information, uncompressed), and the number of variables (mutation rate, selective forces) which can vary at each position and over ancestral time, the computational problem becomes enormous. Making accurate estimates of ancestral trees requires computational resources on the order of Blue Waters.

METHODS & RESULTS

We used two existing programs which use Bayesian inference (MrBayes) and maximum likelihood (ExaML[1]) to estimate ancestral trees on smaller regions of genomes in 1000 individual genomes from the 1000 Genomes Project [2]. Initial experiments identified scaling issues with MrBayes beyond four nodes which we are currently working on resolving using an XSEDE allocation. ExaML was able to scale one million variants linearly to 2048 nodes (Fig.1), and additional scaling to larger numbers of nodes and variants should be possible once we resolve issues with the checkpoint code and initial tree generation. The trees generated from this analysis are congruent with the ethnicity of the individuals (Fig. 2) and we expect to find similar patterns when we can calculate the trees across the entire genome using a full allocation. During this preliminary allocation, we utilized 44,000 node hours.

WHY BLUE WATERS

Blue Waters is one of the only systems that has the computational and I/O resources at the scale necessary to calculate enough trees to examine the tree space in sufficient detail necessary to obtain trees with global maximum likelihood as opposed to a tree with local maximum likelihood. Even our preliminary analyses are beyond the scale of other existing computational resources.

NEXT GENERATION WORK

We hope to extend the initial tree that we generated with additional whole genome sequences as projects like H3 Africa expand our knowledge of the complete variation and population migrations represented through human history.

FIGURE 2: Preliminary ancestral tree from 1 million variants which was calculated from the longest run (8 hours) on 2048 Blue Waters XE nodes. Colors represent reported ethnicities from a subset of the 1000 Genomes Project and are relatively congruent with similar ethnicities and known human migration patterns.

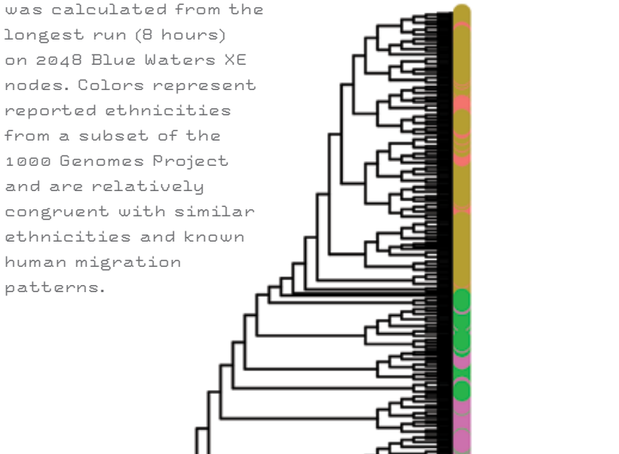

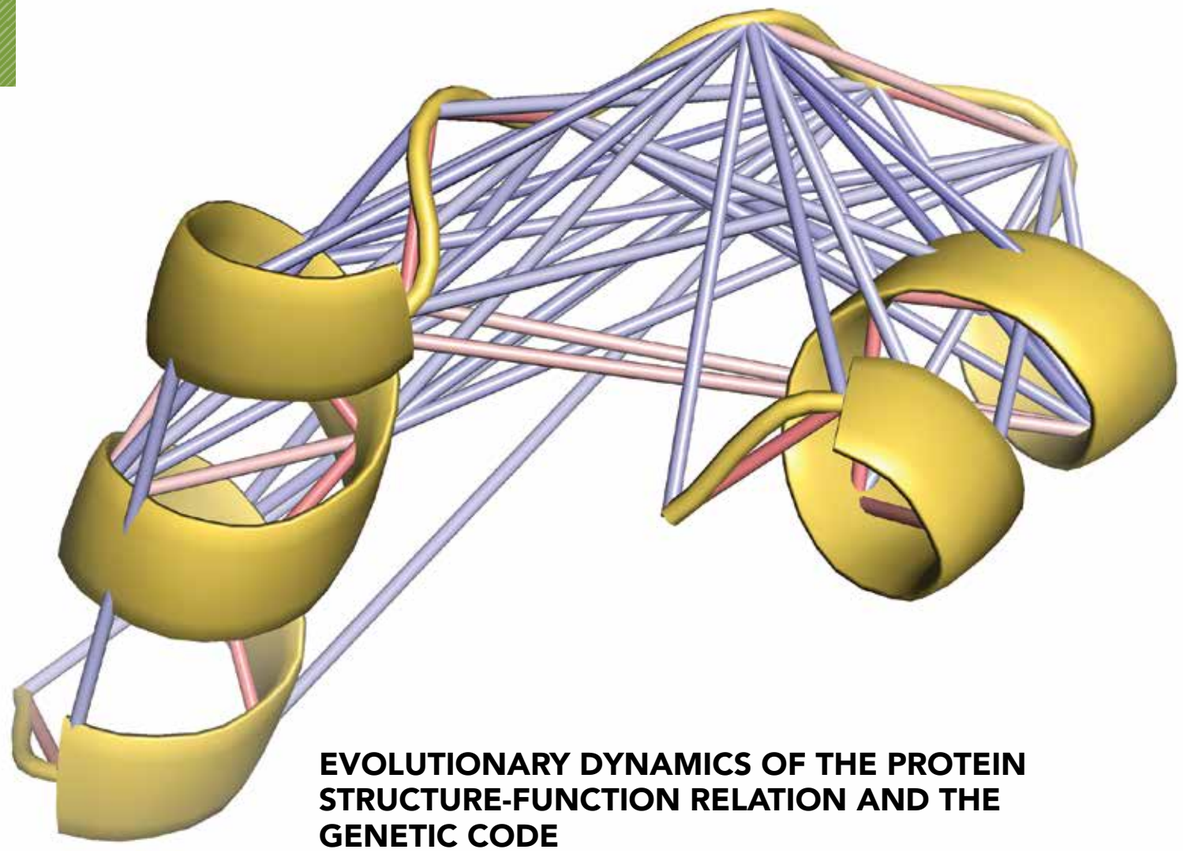

EVOLUTIONARY DYNAMICS OF THE PROTEIN STRUCTURE-FUNCTION RELATION AND THE GENETIC CODE

FIGURE 1: Protein loop 1B7Y_B_408 associated with the a.6.1.1 SCOP domain with residues connected to each other based on positive (red) and negative (blue) correlations of motions during the MD simulation.

Allocation: Illinois/350 Knh
PI: Gustavo Caetano-Anolles[1]
Co-PI: Frauke Gräter[2]
Collaborator: Fizza Mughal[1]

[1]University of Illinois at Urbana-Champaign
[2]Heidelberg Institute for Theoretical Studies

EXECUTIVE SUMMARY

The molecular functions of a protein are determined by the dynamics of its flexible structural components, also known as protein loops. These loop regions exhibit distinct sets of molecular motions, which associate with specific functions. Here we study the biophysics of loop motions in protein structural domains with ages spanning the entire timeline of protein evolution. We are currently working on the molecular dynamic simulations of an initial set of 87 loops embedded within the aminoacyl-tRNA synthethase (aaRS) enzymes responsible for delimiting the specificity of the genetic code. We are also data mining the simulations with machine learning methods. Preliminary analyses using graph theoretical approaches reveal communities of synchronized movements preferentially associated with the secondary structure at the C-termini of the loop regions. These analyses have extended to the sampling of metaconsensus metabolic enzymes and could unravel hidden evolutionary links between the dynamics and functions of proteins.

INTRODUCTION

Protein loops are unstructured regions that play an important role in defining the function and structural stability of proteins [1]. Loops are not completely disordered. They endow an individual protein with characteristic motions that are vital to the protein's function. Protein dynamics, along with backbone flexibility, have been found to be strongly conserved [2]. While the regular secondary structure is highly rigid, the irregular loop counterparts are major contributors to molecular flexibility. The correlation between protein dynamics and function is demonstrated by how certain non-homologous enzymes with similar molecular functions tend to display similar motions [2]. Also, the molecular functions of proteins have been attributed to a combination of functional activities enabled by an evolutionarily conserved set of "elementary functional loops" (EFLs) that associate to form active and regulatory sites in the molecules [3]. Here we investigate this structure-function paradigm using evolutionary dynamics. For this purpose, biophysical variables are selected as candidates for functional annotation of protein domain structures and their loop components [4]. The impact of both evolution and biophysics on the makeup of proteins may provide deeper insight into gains and losses of structural domain features of biological macromolecules [5]. The goal is to decipher the drivers of molecular evolution by reconstructing the evolutionary past. Our work has the potential to build evolutionary trajectories necessary for biomolecular engineering, synthetic biology, and translational medicine. For example, understanding molecular evolution can help uncover a rationale for "engineered" metabolic pathways or can be used to develop better viral vaccines by studying how mutations affect the structure and dynamics of capsid proteins [6].

METHODS & RESULTS

Previous work using Blue Waters revealed that molecular flexibility plays a major role in defining evolutionary constraints acting upon a protein [7]. A protein loop in our study refers to both the unstructured loop structure and the bracing secondary structures that define the return of the polypeptide backbone (Fig. 1). We have simulated 87 loops associated with the domains of aminoacyl-tRNA synthetase (aaRS) enzymes using NAMD 2.9 simulation with force field parameters specified by CHARMM36 files, each for a duration of 10 nanoseconds. These loops have been annotated using gene ontology (GO) molecular functions and classified using the ArchDB "Density Search" system.

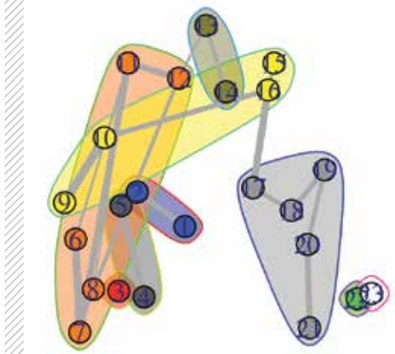

FIGURE 2: All-residue network of 1B7Y_B_408 with emphasis on community structures.

The molecular dynamics simulations were analyzed globally by computing parameters such as radius of gyration and root mean square deviation (RMSD). Also, local parameters were assessed by calculating the root mean square (RMS) fluctuations, exploring parameter variability with principal component analyses (PCA) and plotting community behavior in networks of motions. As expected, the residues in the unstructured parts of the loops possess the highest values of RMS fluctuation. The PCA is projected onto the candidate loop structure to ascertain the type of associated motion. The communities in our network analyses (Fig. 2) consist of residues that have synchronized movements during the simulation, based on residue cross-correlations (Fig. 2). Remarkably, the betweenness values plotted from the resulting networks showed high values for residues in the C-terminus secondary structure for the majority of structures that were examined. Future analyses will focus on uncovering patterns in the community structures, as well as motions, thereby dwelling deeper into the structure, function and dynamics of protein loops. These analyses will be followed by mapping patterns along an evolutionary timeline of loop-embedding structural domains generated with robust phylogenomic reconstruction methods [8].

WHY BLUE WATERS

Blue Waters has been instrumental in enabling the study of the structure-function protein interplay at the core of the origin of the genetic code. With the help of Blue Waters, we have successfully

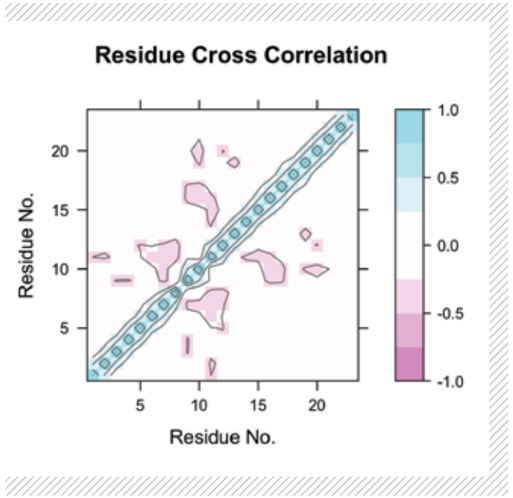

FIGURE 3: Dynamical Cross Correlational Map of the protein residue backbone of 1B7Y_B_408.

simulated a total of about 1,500 nanoseconds (1.5 milliseconds) worth of simulations involving 87 protein loop classifications associated with aaRS domains delimiting the specificities of the genetic code. Loop molecular dynamic simulations of these multi-domain proteins have yielded impressive results which have laid the **groundwork** for analyses involving single-domain metabolic metaconsensus enzymes selected using comparative bioinformatics techniques grounded in sequence, structure, and metabolic reactions.

NEXT GENERATION WORK

Our preliminary experiments indicate an intricate set of patterns among structure, function, and dynamics of proteins. These patterns have the potential to **uncover** evolutionary drivers that may shed light on a basic yet confounding phenomenon in nature: does structure dictate function or vice versa? We aim to expand our analysis to bigger datasets concerning proteins in signaling networks. Also, we are interested in performing machine learning analyses of dynamic simulation datasets to dissect distinct molecular dynamic patterns along an evolutionary timeline of protein domains.

PUBLICATIONS AND DATA SETS

Mughal, F., G. Caetano-Anollés and F. Gräter. Mining the evolutionary dynamics of protein loop structure and its role in biological functions. *Blue Waters Annual Report* (Urbana, Illinois, 2015), pp.130-131.

Caetano-Anollés, G., F. Gräter, C. Debès, D. Mercadante, and F. Mughal. The dynamics of protein disorder and its evolution: Understanding single molecule FRET experiments of disordered proteins. *Blue Waters Annual Report* (Urbana, Illinois, 2014), pp. 100-101.

IMPROVING THE ACCURACY OF DRUG PERMEABILITY CALCULATIONS

Allocation: Illinois/25.0 Knh
PI: Christopher Chipot[1]
Co-PI: Jeffrey Comer[2]

[1]University Illinois at Urbana-Champaign
[2]Kansas State University

EXECUTIVE SUMMARY

The inhomogeneous solubility-diffusion model has provided a convenient framework for understanding membrane permeation by drug molecules. This model shows the relationship between the resistance to permeation in the direction normal to the membrane to the position-dependent diffusivity of the drug and the one-dimensional free-energy profile underlying its translocation from the bulk aqueous phase to the interior of the lipid environment. For the **first time**, we provide a model for membrane permeation of a drug that, in stark contrast with the solubility-diffusion model, does not assume a lack of long-range correlations in time and space. Our model allows for better understanding of permeation dynamics for molecules exhibiting subdiffusive behavior on the characteristic timescales of their permeation. Our simulations suggest that this subdiffusive behavior is a result of permeation being governed by the spontaneous formation of voids within the membrane, which leads to intermittent large displacements of a permeant that is otherwise nearly immobile.

INTRODUCTION

In the search of novel therapeutic agents, many chemical compounds able to bind a given target with very high affinity are eventually discarded on account of their cytotoxicity, propensity to associate with potassium channel human Ether-à-go-go-Related Gene (hERG), or poor bioavailability. Predicting these properties at an early stage of drug discovery, upstream from costly organic syntheses and clinical trials, is desirable. One possible avenue to address high drug-attrition rates [1] consists in quantifying the ability of the substrate to traverse lipid membranes spontaneously, for instance, in the gastrointestinal tract, and reach the targeted protein in an adequate amount. A consistent theoretical model of the lipid membrane permeation process is essential for linking the physicochemical properties of drug candidates to their adsorption and distribution.

METHODS & RESULTS

The goal of this research is to understand a question central to drug discovery, namely how a drug spontaneously translocates across the biological membrane to reach its designated target. A model that has pervaded the field over the past twenty years is the so-called solubility–diffusion model of passive membrane permeation of small molecules [2]. In this model, the diffusion of the permeant is ordinarily assumed to obey the conventional Smoluchowski diffusion equation, which describes classical diffusion of particles on an inhomogeneous free-energy and diffusivity landscape. However, this equation cannot accommodate subdiffusive behavior [3], which has long been recognized in other aspects of lipid bilayer dynamics, including lateral diffusion of individual lipids. Using large-scale molecular dynamics simulations of permeation events in a fully hydrated lipid bilayer performed on Blue Waters, we show that subdiffusive behavior is present in the transverse diffusion of a series of alcohols through a pure membrane, remaining relevant on timescales approaching the typical permeation time. We find that a model based on a fractional-order differential equation appropriately describes the motion of the permeant on timescales ranging from 1 picosecond to 1 nanosecond, which cannot be replicated by a single conventional Smoluchowski model. Multiple approaches indicate that the mean squared displacement within the bilayer, in the absence of a net force, depends on time as a power law, namely $t^{0.7}$, in contrast with the conventional model where this dependence is strictly linear. Our molecular dynamics simulations bring to light an unexpected phenomenon, linking subdiffusion to the formation of transient voids that spontaneously appear within the hydrophobic region of the bilayer and allow rare, but large displacements of the permeant, which is otherwise virtually immobile. The results of this investigation, which reweaves the fabric of the physical principles underlying membrane permeation by drug molecules, have been reported in a research article recently submitted for publication [4].

WHY BLUE WATERS

Blue Waters was essential to perform a very large series of independent molecular-dynamics simulations of a membrane assembly in a time-bound fashion.

FIGURE 2:
Illustration of a cellulosome domain acting over cellulose fibers. Cellulosomes are highly-efficient molecular machines that can degrade plant fibers. Addapted from Cover of Schöler et. al. [10].

CELLULOSOME STRUCTURE DETERMINATION EMPLOYING ATOMISTIC SIMULATIONS COMBINED TO EXPERIMENTAL ASSAYS

Allocation: Illinois/680Knh
PI: Isaac Cann[1]
Co-PI: Rafael C. Bernardi[1]
Collaborators: Klaus Schulten[1], Edward Bayer[2], Hermann Gaub[3], and Michael Nash[3]

[1]University of Illinois at Urbana-Champaign
[2]Weizmann Institute
[3]Ludwig Maximilian University

EXECUTIVE SUMMARY

Cellulosomes are multi-enzyme complexes that target the deconstruction of cellulose and hemicellulose in anaerobic cellulosome-containing bacteria. Briefly, in cellulosome assembly, a large noncatalytic polypeptide called the scaffoldin, embedded with various cohesins (Coh), anchors dockerin (Doc)-containing enzymes through Coh–Doc interactions. The precision of the Coh–Doc interaction allows the addition of different catalytic cellulases and hemicellulases onto the scaffoldin that may or may not be bound to another domain attached to the cell wall. Cellulosomes' ability to efficiently degrade plant-cell-wall biomass allows them to be used in the second-generation biofuel industry, which aims to use agricultural waste to produce ethanol. Furthermore, the recent discovery of cellulosomal bacteria in the lower gut of humans is paradigm-shifting as it has allowed demonstration of the capacity to degrade both hemicellulose and cellulose, at least in some humans. Our project employs molecular dynamics simulations, complementing single-molecule and biochemistry experiments, to characterize the structure of cellulosomes.

INTRODUCTION

Bacteria play a key role in the second-generation biofuel industry since their cellulolytic enzymes, used for plant-cell-wall degradation, are employed in the production of these advanced biofuels. Also, symbiont bacteria greatly influence human health and play a significant role in pathogenesis, disease predisposition, physical fitness, and dietary responsiveness. Here we are investigating key processes underlying bacterial activity, namely, plant fiber metabolism. Specifically, we are examining the structure and function of cellulosomes, the highly cooperative macromolecular complex that is central to this metabolic process in some bacteria.

Cellulosomes are multi-enzyme complexes that target the deconstruction of cellulose and hemicellulose in anaerobic cellulosome-containing bacteria. Integration of cellulosomal components occurs via highly ordered protein–protein interactions among three major components. Briefly, in cellulosome assembly, a large noncatalytic polypeptide called the scaffoldin, embedded with various Cohs, anchors Doc-containing enzymes through Coh–Doc interactions (Fig. 1). Specificity of the cohesin–dockerin interaction allows incorporation of different catalytic cellulases and hemicellulases onto the scaffoldin that may or may not be bound to another domain tethered to the cell wall. Cellulosome assembly promotes the exploitation of enzyme synergism because of spatial proximity and enzyme-substrate targeting [1].

Combined with biochemical and single molecule experiments, we employed molecular dynamics (MD), steered MD (SMD) and generalized simulated annealing (GSA) [2] simulations on Blue Waters. Utilizing QwikMD [3], our new intuitive "point and click" graphical interface connecting visual MD (VMD) [4] and nanoscale molecular dynamics (NAMD) [5], we are studying the detailed mechanism of cellulase complexes, in particular, cellulosomes. Using stochastic search algorithms connected to molecular dynamics tools, we are building the **first** comprehensive structure of a cellulosome. Employing GSAFold/NAMD we were already able to obtain the structure of a cellulosome scaffoldin and, using Blue Waters, we are working to determine the structure of a whole cellulosome complex, including enzymatic domains. We expect that a complete model of cellulosome's structure will shed light on the mechanism that allows these enzymatic complexes to be highly efficient.

METHODS & RESULTS

As mediators in the interactions between comparatively large bacterial cells and cellulose particles, scaffolds, and carbohydrate-binding molecule (CBM) domains are critical cellulosomal components (Fig. 2). As many cellulosomal habitats (for example, cow rumen) exhibit strong flow gradients, shear forces will accordingly stress bridging scaffold components mechanically *in vivo*. Protein modules located at stressed positions within these networks are likely to be preselected for high mechanostability, as demonstrated by our group [6]. However, thus far little is known about cellulosome structure as they are formed by interactions of Doc and Coh that are connected by flexible linkers made of proteins with just a few to more than a thousand amino acids. The many flexible linker regions allow for very complex structural dynamics [7].

Using stochastic search algorithms coupled to NAMD we can generate thousands of different structure conformations for the cellulosome. GSA[2] analysis shows that the different linkers between Coh (and a CBM) in cellulose-integrating protein A (CipA) scaffoldin assume a different number of more stable conformations. Small angle X-ray

FIGURE 1:
Illustration of a complete cellulosomal structure. The scaffoldin (yellow) can be attached by a specific Doc-Coh interaction to a cell-anchoring domain (in orange). Another Doc-Coh interaction is responsible for attaching the enzymatic domains (red) to the scaffoldin. Addapted from Cann et. al. [1].

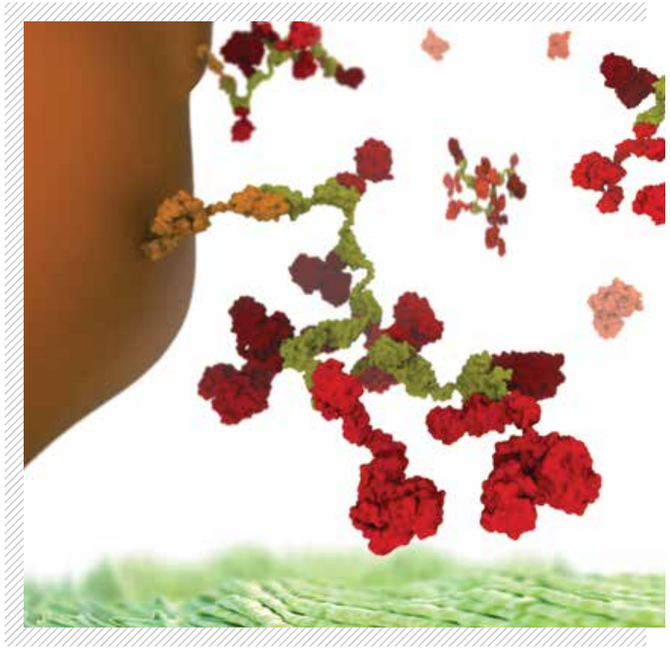

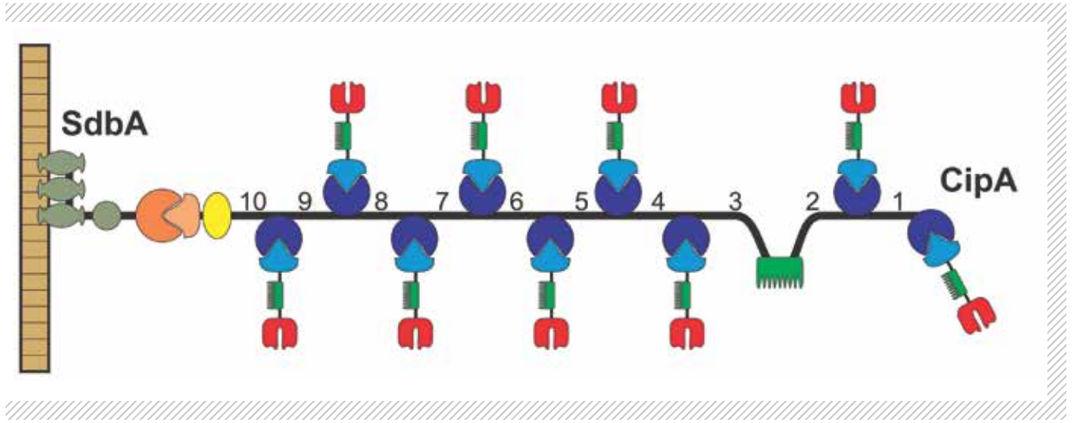

FIGURE 3: Organization of Clostridium thermocellum cellulases and hemicellulases in the SdbA/CipA cellulosome. The C. thermocellum scaffoldin (CipA) contains one CBM (Green) and nine type I cohesins (Dark Blue) and thus organizes a multiprotein complex with nine enzymes (Red). The C-terminal type II dockerin (Pink) domain of CipA binds specifically type II cohesin domains (Orange) found in cell-surface proteins. The CipA linkers already studied using GSAFold/NAMD integration are numbered.

scattering (SAXS) analysis has previously shown that three conformations are observed for linker 10 (Fig. 3). GSAFold is capable of predicting these three conformations and all the other conformations for CipA. To perform this analysis, 20,000 conformations were obtained per linker and clustered. Combined, these linker conformations would give us 1043 CipA conformations. From clustering, we reduce this number to 3888 structures that were obtained and also subjected to a cluster analysis that gave rise to the five most significant structures.

Following well-established protocols for large macromolecular systems [8,9], and using one of the CipA conformations that we obtained using GSAFold, we built a **first** model of an entire cellulosome structure. MD simulations are now employed to study the quaternary structure stability.

WHY BLUE WATERS

Investigating the structure and functional processes of large enzymatic complex machineries, such as the cellulosomes, is only possible on petascale computing resources, such as Blue Waters. Structures obtained using enhanced sampling techniques, such as GSA, are only reliable if thousands of conformations (models) are predicted. Employing GSA for the numerous linkers of the cellulosome is a well-suited task for the large-scale parallel architecture of Blue Waters.

NEXT GENERATION WORK

Our primary goal is to obtain a clear picture of the cellulosome structure at work. For that, long molecular dynamics simulations of different cellulosomes, some of them with hundreds of millions of atoms, will have to be performed. To investigate the enzymatic mechanism in the context of the cellulosome, hybrid quantum mechanics (QM)/molecular mechanics simulations will have to be performed using multiple QM regions that require massive computer power. Such complex study might only be feasible in a few years, requiring pre-exascale and exascale systems.

PUBLICATIONS AND DATA SETS

Schoeler, C., et al., Ultrastable cellulosome-adhesion complex tightens under load. *Nat. Commun.* 5 (2014), p. 5635.

Bernardi, R.C., M.C.R. Melo, and K. Schulten, Enhanced Sampling Techniques in Molecular Dynamics Simulations of Biological Systems. *Biochim. Biophys. Acta.*, 1850 (2015), p. 872.

UNDERSTANDING BIOMOLECULAR STRUCTURE AND DYNAMICS BY OVERCOMING BARRIERS TO CONFORMATIONAL SAMPLING

Allocation: NSF PRAC/2.00 Mnh
PI: Thomas Cheatham[1]
Co-PI: Adrian Roitberg[2], Carlos Simmerling[3], and David Case[4]
Collaborators: Darrin York[4], and Shantenu Jha[4]

[1]University of Utah
[2]University of Florida
[3]Stonybrook University
[4]Rutgers University

EXECUTIVE SUMMARY

Large ensembles of independent molecular dynamics, running optimized AMBER code on Blue Waters' GPUs, enable full sampling of the conformational ensemble of biomolecules, including DNA helices, RNA tetranucleotides, and RNA tetraloops. This allows detailed validation and assessment of enhanced sampling approaches and biomolecular force fields and provides detailed insight into biomolecular structure, dynamics, interactions, and function. The ensemble simulations currently being performed are possible only on computational hardware with large numbers of GPUs. While today our simulations are pushing the state of the art, such large simulations will become routine within a few years. The even larger and more powerful parallel resources available in the near future will enable molecular dynamics simulations to probe more relevant biological time scales (milliseconds to seconds) and to study larger biomolecular assemblies more completely.

INTRODUCTION

Biomolecular simulation—although known as a powerful tool for probing the structure, dynamics, interactions, and functions of proteins and nucleic acids for over 40 years—is really coming of age thanks to access to large-scale computational resources such as Blue Waters. Not only can simulations be applied to larger biomolecular assemblies, but for modest sized biomolecules the community has demonstrated the ability to fold proteins *de novo* and to fully sample the conformational distributions of various nucleic acid motifs. A challenge is parallel scaling since, for a fixed system size, adding additional cores does not increase performance. To overcome this, the community has moved toward application of ensemble methods and application of various enhanced sampling methodologies that couple together independent molecular dynamics (MD) simulation engines. AMBER, a suite of programs for biomolecular simulation whose latest version, AMBER 16, was released in April 2016, has been highly optimized for use on GPUs. The optimized GPU code, and the ensembles that are

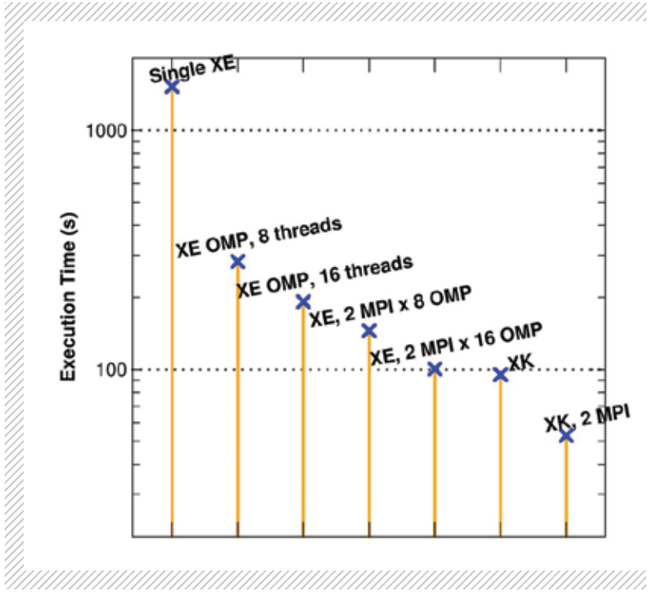

FIGURE 1: Relative performance of CPPTRAJ on different nodes when determining the 965 closest solvent molecules out of 15,022 to 4,143 solute atoms from a 2,000 frame MF trajectory (no imaging) using various parallelization modalities, including CUDA on the XK nodes.

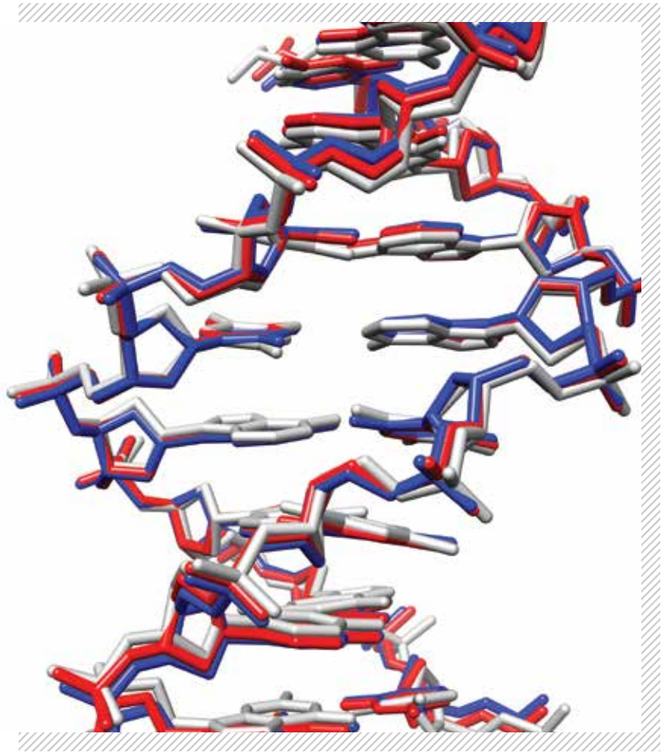

produced, provide a powerful means to assess and validate the available force fields and to apply these methodologies to give novel biological insight into protein and nucleic acid structure and dynamics. Improving the codes, methods, and force fields while also pushing ensemble methods to their limits are critical research activities since these tools are being used by an ever-increasing pool of researchers throughout the world.

METHODS & RESULTS

The coupling together of independent molecular dynamics simulations into loosely coupled ensembles to enhance conformational sampling is an increasingly used modality for applications in biomolecular simulation. If you peruse any recent journal in the field, including *The Journal of Chemical Theory and Computation, The Journal of Computational Chemistry,* and *The Journal of Physical Chemistry*, among others, you will see multiple publications developing and applying ensemble methods. With a variety of methodological variations and names ranging from replica-exchange MD, metadynamics, swarms, and Markov state modeling to constant pH MD and lambda dynamics, all of the approaches promise more efficient means to explore structural ensembles, free energy pathways, and kinetics. Although these techniques are promising, there is no free lunch, since as the size of the biomolecule increases, sampling a complete ensemble takes longer and longer. Various methods attempt to speed the process through applications of enhanced sampling in different degrees of freedom. However it is often difficult to verify claims of efficiency, especially with method variants implemented into vastly different code bases. Therefore, we have explored making our ensemble data available to the community online (http://amber.utah.edu) to allow other researchers to directly compare our results to results obtained using different ensemble approaches and to assess relative convergence and efficiency.

The combination of large ensemble methods with the availability of many fast GPUs on Blue Waters leads to an explosion in data. In order to process vast amounts of data efficiently, the AMBER MD trajectory analysis code CPPTRAJ has been modified to include multiple levels of parallelism, aided by a Petascale Application Improvement Discovery (PAID) collaboration with Blue Waters. Specifically, CPPTRAJ now implements four levels of parallelism. MPI parallelism over ensemble instances (with sorting of data across all ensembles), MPI parallelism over reading/writing of trajectory files in a given ensemble, OpenMP parallelism for time-intensive analyses, and most recently GPU parallelization of time-intensive analyses involving calculations of a large number of distances. Relative performance is shown in Fig. 1.

Key results of our ensemble approaches are described in greater detail in the listed publications and range from accurate modeling of magnesium-dependent conformational changes in RNA, correct modeling of proteins binding and modulating DNA structure, and optimization, assessment, and validation of nucleic acid force fields. This includes fully converging the conformational and dynamic distribution of the Dickerson-Drew dodecamer with average structures over milliseconds of aggregated MD data less than 0.8 Å from the published experimental structures as shown in Fig. 2.

WHY BLUE WATERS

Within the NSF ecosystem of computational resources, Blue Waters is the GPU-optimized resource with a sufficiently large set of GPUs to allow ensembles on the 300-3,000 scale (assuming a single ensemble instance per GPU). Our team has shown the ability to converge the conformational ensemble of an RNA tetraloop with multidimensional replica exchange; using 360 GPUs, this requires about 2-3 microseconds of MD simulation per ensemble instance, or approximately five to 10 days of continuous MD simulation on those resources.

NEXT GENERATION WORK

Ensemble-based biomolecular simulation methods will continue to evolve in terms of their generality and power. With next-generation computational resources, the community will be able to not only study larger biomolecular systems but also to more fully sample and converge the accessible conformational space of these systems. While at present we can aggregate to milliseconds of effective sampling, to reach biological time scales we still need orders of magnitude greater sampling (to seconds and beyond) in simulations that likely will require multiple GPUs or accelerators to reach system sizes of hundreds of thousands to millions of atoms and beyond. The kind of ensemble analyses we are currently undertaking, made possible by Blue Waters, will be effectively routine by 2019-20 and will enable the larger community while we push the edge of what is possible on future petascale resources.

PUBLICATIONS AND DATA SETS

http://amber.utah.edu

Galindo-Murillo, R., J.C. Garcia-Ramos, L. Ruiz-Azuara, T.E. Cheatham, III, and F. Cortes-Guzman. Intercalation processes of copper complexes in DNA. *Nuc. Acids Res.* 43 (2015) pp. 5364-5376.

Bergonzo, C., N. Henriksen, D.R. Roe, and T.E. Cheatham, III. Highly sampled tetranucleotide and tetraloop motifs enable evaluation of common RNA force fields. *RNA* 29 (2015) pp. 1578-1590.

Bergonzo, C. and T.E. Cheatham, III. Improved force field parameters lead to a better description of RNA structure. *J. Chem. Theory Comp.* 11 (2015) pp. 3969-3972.

Bergonzo, C., K.B. Hall, and T.E. Cheatham, III. Stem-loop V of Varkud satellite RNA exhibits characteristics of the Mg^{2+} bound structure in the presence of monovalent ions. *J. Phys. Chem. B* 119, (2015) pp. 12355-12364.

Robertson, J.C. and T.E. Cheatham, III. DNA backbone BI/BII distribution and dynamics in E2 protein-bound environment determined by molecular dynamics simulation. *J. Phys. Chem. B* 119 (2015) pp. 14111-14119.

Galindo-Murillo, R., D.R. Davis, and T.E. Cheatham, III. Probing the influence of hypermodified residues within the $tRNA_3^{Lys}$ anticodon stem loop interacting with the A-loop primer sequence from HIV-1. *Biochemica Biophys. Acta* 1860 (2016) pp. 607-617.

Dissanayake, T., J. Swails, M. Harris, A.E. Roitberg, and D. York. Interpretation of pH-activity Profiles for Acid-Base Catalysis from Molecular Simulations. *Biochemistry* 54 (2015) pp 1307-1313.

Hopkins, C., S. LeGrand, R.C. Walker, A.E. Roitberg. Long Time Step Molecular Dynamics through Hydrogen Mass Repartitioning. *J. Chem. Theory Comp.* 11 (2015) pp. 1864-1874.

Alvarez, L., A. Louis Ballester, A.E. Roitberg, D. Estrin, S.-R. Yeh, M. Marti, L. Capece. Structural study of a flexible active site loop in human indoleamine 2,3-dioxygenase and its functional implications. *Biochemistry* 55 (2016) pp. 2785-2793.

Nguyen, H., A. Perez, S. Bermeo, and C. Simmerling. Refinement of generalized Born implicit solvation parameters for nucleic acids and their complexes with proteins. *J. Chem. Theory Comp.* 11 (2016) pp. 3714-3728.

Maier, J.A., C. Martinez, K. Kasavajhala, L. Wickstrom, K.E. Hauser, and C. Simmerling. ff14SB: Improving the Accuracy of Protein Side Chain and Backbone Parameters from ff99SB. *J. Chem. Theory Comp.* 11 (2015) pp. 3696-3713.

Lai, C.-T., H.-J. Li, W. Yu, S. Shah, G.R. Bommineni, V. Perrone, M. Garcia-Diaz, P.J. Tonge, and C. Simmerling. Rational Modulation of the Induced-Fit Conformational Change for Slow-Onset Inhibition in Mycobacterium tuberculosis InhA. *Biochemistry* 54 (2015) pp. 4683-4691.

FIGURE 2 (LEFT): Overlap of the average structures, omitting hydrogens and the two terminal base pairs on each end from nearly milliseconds of aggregated MD simulation data from 100 independent 11 microsecond length MD simulations (omitting the first 2 microseconds) with the parmbsc1 (blue) and AMBER ff15 (or ol15, red) force fields compared to the PDB average structure from 1NAJ (gray).

PREDICTING PROTEIN STRUCTURES WITH PHYSICAL MOLECULAR SIMULATIONS

Allocation: NSF PRAC/5.00 Mnh
PI: Ken Dill[1]
Co-PI: Alberto Perez[1]
Collaborators: Emiliano Brini[1], Joseph Morrone[1], Lane Votapka[1], and Cong Liu[1]

[1]SUNY Stony Brook University

EXECUTIVE SUMMARY

Ab initio protein folding has been a computational challenge for the last 50 years. We have developed a highly efficient platform called MELD, running on graphics processing units (GPUs) that allow the folding of protein structures in weeks of simulation time. We are using Blue Waters in a **worldwide** blind protein folding event involving ~200 scientific groups independently to assess the value of the methodology (CASP, Critical Assessment of Structure Prediction). CASP operates under very strict timelines; for four months, protein sequences are released each day, for which we have three weeks to predict the 3D structure of the protein. We are currently testing the capacity to fold proteins up to 200 residues, twice as large as has been previously possible *ab initio*.

INTRODUCTION

A long-standing grand challenge in computational biology is determining if we can use computers to find out the native structure of proteins given their sequence. This lead to IBM's effort in the 1990s with BlueGene and more recently DE Shaw's Anton supercomputer [1], which produced some of the fastest folding proteins. This approach is extremely computationally demanding and does not scale to larger and slower folding proteins.

We have developed MELD (Modeling Employing Limited Data) as a Hamiltonian and temperature replica exchange method and plugin to the program OpenMM. It has many biological applications (folding, docking, mechanisms, etc.) and runs very efficiently on GPUs. MELD's differentiating factor is the ability to incorporate noisy, sparse and ambiguous data through a Bayesian inference approach—increasing the performance to obtain protein structures by five orders of magnitude. We use coarse physical insights (e.g. proteins have hydrophobic cores) with very low signal to noise ratio. We have already produced three high accuracy structure predictions—but much more is needed to shift the field from a purely bioinformatics approach to physics-based simulations. We are participating in CASP, a blind competition extending four months, with close to 200 participating groups. This competition is the perfect scenario to witness the real life performance of MELD in the context of all the state-of-the-art methods. However, MELD as a platform goes beyond folding, and we are concurrently using Blue Waters to calculate relative binding free energies of peptides folding upon binding and to identify the most favorable oligomerization conformation for protein dimers. Our goal is to breach several milestones: (1) Fold longer proteins than previously possible *ab initio*, (2) Obtain binding free energies of peptides to proteins using MELD to do flexible ligand/protein binding and (3) Obtain binding conformations for protein dimers.

The second grand challenge is comparing the binding of several peptides to the MDM2 gene involved in cancer therapy. Traditional methods cannot obtain relative binding free energies because they take a very long time to converge. We run simulations in which two peptides are competing to bind a protein—enforcing that at any time one peptide should be in a reference state far from the protein and the other near it.

METHODS & RESULTS

We are using MELD to carry out Hamiltonian and replica exchange molecular dynamics simulations. In MELD we introduce data to guide simulations which limit the conformational space accessible to simulations, and inside those regions compatible with the data, it is physics that drive the sampling.

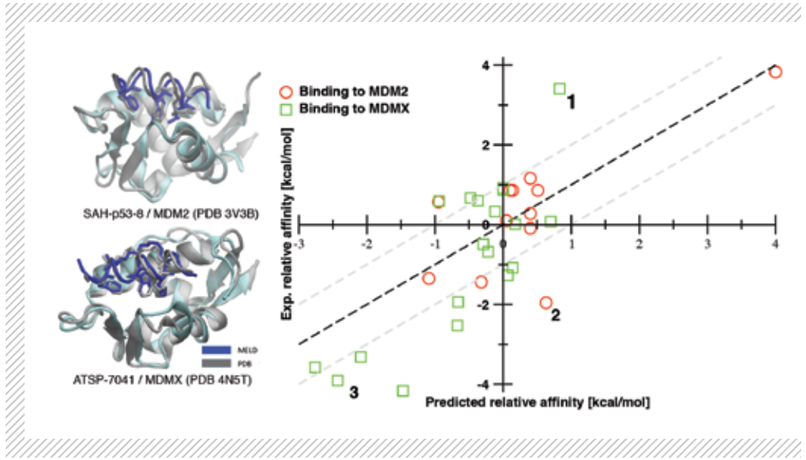

FIGURE 1: Comparison of computed vs experimental relative free energies for the binding of different peptides to target proteins MDM2 and MDMX involved in cancer pathways. The left panel shows the predictions coming from MELD on top of the experimental structure.

The biggest advantage of MELD is that the user specifies what data should be trusted [2]. The simulations then optimize two different problems: finding structures that are most compatible with the physics and the data. Solving these two problems together gives a five orders of magnitude speedup over using physics alone.

We have produced over 95 atomistic detailed protein models for proteins ranging from 80 to 230 amino acids within CASP [4]—far beyond the standard 100 residue limit for *ab initio* modeling. The CASP event continues, and we are projecting over 500 predictions by the end of the summer. The biggest advantage over database methods is that MELD provides meaningful populations which allow us to identify native conformations. We are the only physics-based group in this competition as the tight CASP deadlines are not possible with MELD's efficient sampling protocols and GPU computing.

To approach the challenge of comparing the binding of peptides to the MDM2 gene, we observe binding of one or the other at different times using MELD simulations. This protocol yields relative binding free energies in agreement with experiments (Fig. 1). This flexible receptor-flexible docking is needed for developing new drugs in which flexibility is important. In doing these two studies together (folding and binding) we are showing a pipeline that can go all the way from genomic sequencing to folding to binding, unveiling new possibilities for drug design.

WHY BLUE WATERS

With Blue Waters, we can tackle simulations of proteins that would require years of simulation time in weeks—producing high accuracy atomistic detailed protein models. In addition, we are running a time-sensitive blind prediction, in which multiple proteins need to be simulated independently. This would be impossible without Blue Waters resources.

NEXT GENERATION WORK

Our goal through the increased power of the next-generation Track-1 system is to reach the average protein size in the human body, which is 300 residues long for *ab initio* structure prediction—3 times longer than the current state of the art. Membrane proteins, the target for most pharmaceutical drugs, remain outside of our scope. We are developing a new methodology to tackle them but will need extreme supercomputer capability. We are preparing the proof of concept with Blue Waters and expect to do membrane structure prediction on a future Track-1 system.

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

Allocation: GLCPC/383 Knh
PI: Peter Freddolino[1]
Co-PI: Morteza Khabiri[1]
Collaborator: Arttu Jolma[2]

[1]University of Michigan Medical School
[2]University of Toronto

EXECUTIVE SUMMARY

Transcription factors (TFs) and other DNA-binding proteins shape the behavior of all cells, coordinating appropriate gene expression patterns in response to internal or external cues. For any particular transcription factor, maps of the binding affinity for different DNA sequences must be obtained through laborious and expensive experiments. Using the massive computing resources available through Blue Waters, we are pursuing a strategy to computationally map the DNA-binding affinity landscapes of several human transcription factors. Through comparison with experimental results on the same systems, we will validate and refine computational protocols for allowing reliable *in silico* determination of TF affinity landscapes, obtain completely novel insight into the structural basis for these affinity landscapes, and catalog the effects of the binding of different transcription factors on DNA structure, which appears likely to play a key role in the interplay between different transcription factors regulating the same gene *in vivo*.

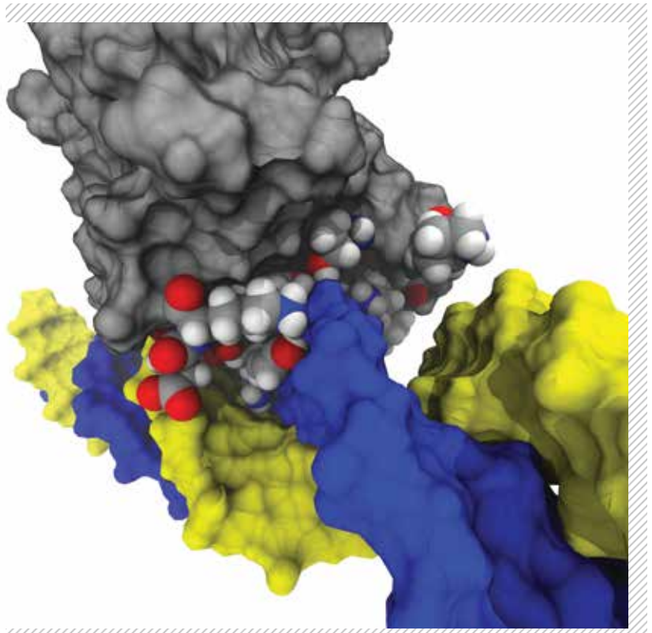

FIGURE 1: Structure of the transcription factor ELK1 (grey) bound to DNA (yellow, blue); the key residues on the protein involved in binding are shown outside of the surface.

INTRODUCTION

Transcriptional regulation is driven in large part by the action of TFs and other DNA binding proteins that either recruit or inhibit the recruitment of RNA polymerase; thus, to understand and predict the behavior of transcriptional regulatory networks, it is necessary to know the landscape of binding affinities of each transcription factor to all possible sequences of DNA. While experimental methods have been developed to measure these landscapes (e.g., protein binding microarrays (PBMs) [1] or HT-SELEX experiments [2]), these experiments remain expensive, labor intensive, and provide no structural insight into the nature of the protein-DNA complex or the specific interactions governing affinity landscapes.

The need for a **more efficient method** to obtain sequence affinity landscapes, insight into the structural basis of these landscapes, and information on the structural effects of different transcription factors on DNA all argue for the large-scale application of molecular modeling to study the affinity landscapes of DNA-binding proteins. Prior efforts to predict TF binding affinity landscapes using approximate methods performed well in identifying the native binding site for a particular protein, but poorly in efforts to more broadly map the binding affinity of a given protein for a variety of sequences (reviewed in [3]). Extremely promising preliminary work using all-atom simulations with explicit solvent (and thus requiring far fewer approximations) has shown near chemical accuracy (average absolute error of 0.31 kcal/mol to experimental values) for all possible single base pair perturbations of the native target site for the zinc-finger transcription factor Zif268 [4].

Using the massive computing resources provided by Blue Waters, we are applying similar atomistic free energy simulations to map the complete sequence affinity landscapes for a set of four carefully chosen human transcription factors for which direct comparisons with experimental results are possible. Our computational protocol requires long equilibrium simulations of the system with each of the DNA sequences being considered. Thus, we will obtain information on the binding landscape of the protein and gain insight into the structural basis for specificity of different transcription factors by allowing detailed analysis of the structural changes associated with mutations in the target sequence. We will use the availability of direct experimental comparisons to **benchmark** the accuracy of our methods, and test force field modifications to enhance performance.

METHODS & RESULTS

Building on previous results that showed accurate calculations of protein-DNA binding free energies for a small number of cases [4], we are applying the Crooks-Gaussian intersection (CGI) method [5] to calculate the free energy changes for base pair substitutions in the binding site of the transcription factors of interest. The method requires calculations of very long equilibrium simulations of the protein-DNA complex and the DNA alone for each of two sequences to be compared, followed by many short simulations morphing the system between the two sequences. We will perform the free energy calculations for all possible single nucleotide perturbations of the consensus binding site for the transcription factor of interest, and subsequently pursue additional mutations through a tree-like approach (suggested by [4]) in which at each layer of mutations, only those which have not been shown to strongly inhibit binding will be considered in the next round.

Our results to date illustrate that the free energy changes from *in silico* calculations on the protein-DNA complex can indeed reproduce the consensus binding sequence of transcription factors obtained from HT-SELEX data, and additionally show that the ensemble of thermally accessible protein-DNA interaction conformations in the bound state is far broader than what might be inferred from crystallographic structures. Our findings on the latter point may be particularly relevant to future efforts in developing streamlined prediction or design of protein-DNA interfaces.

WHY BLUE WATERS

The computational work described here requires the capability to bring huge numbers of nodes efficiently together to run dozens of simulations of independent trajectories using graphics processing unit-accelerated molecular dynamics software, and then for each such trajectory to perform more than 100 short follow-up simulations using central processing unit-only code for the free energy calculation. The hybrid architecture of Blue Waters has been ideal for these applications, providing us with the most efficient possible environment for each portion of our workflow, and allowing us to make progress on huge numbers of mutational calculations simultaneously.

NEXT GENERATION WORK

Recent findings by our collaborators and others have demonstrated that when two or more transcription factors bind to nearby sites on DNA, they can alter each other's sequence affinity landscapes, both through direct protein-protein interactions and the effects of binding on DNA structure. We will use the Track 1 system to simulate two cases each of protein-mediated and DNA-mediated binding site modulation, with an aim toward identifying the biophysical basis for the observed changes in sequence specificity.

ELUCIDATING THE MOLECULAR BASIS OF CHARGE SELECTIVITY IN PENTAMERIC LIGAND-GATED ION CHANNELS

Allocation: Illinois/132 Knh
PI: Claudio Grosman[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Pentameric ligand-gated ion channels (pLGICs) are a unique superfamily of ion channels containing both cation- and anion-selective channels. While charge selectivity has been studied in pLGICs for many years experimentally, the minimal number of mutations needed to alter charge selectivity does not appear to be consistent among all members of the superfamily. In recent years, high-resolution structures of both cation- and anion-selective pLGICs in multiple conformational states have become available. Therefore, we set out to understand the molecular mechanisms that underlie anion selectivity. We tested experimentally generated hypotheses about the origin of charge selectivity with the goal of providing an encompassing model that can explain charge selectivity in a wide variety of pLGICs. To do this, we used Blue Waters to calculate free energy profiles of ions moving through the transmembrane region of various pLGICs. We found that charge selectivity is determined by the electrostatics of the transmembrane pore domain.

INTRODUCTION

Ion channels passively allow ions to diffuse in and out of a cell. However, many ion channels allow only specific ions to cross the membrane. One way an ion channel can discriminate among ions is based on the formal charge of an ionic species, known as charge-selectivity. Cation-selective ion channels allow only positively charged cations to move across the membrane, whereas anion-selective ion channels allow only negatively charged ions to cross the membrane. Charge-selectivity is highly important for ion channel biophysics because in adult neurons, when cations are allowed to pass through the cell, the electrical signal is said to be excitatory because it enhances the likelihood of an action potential by depolarizing the cell and therefore the propagation of the neuronal signal. However, when anions are allowed to pass, the electrical signal is inhibitory because it hyperpolarizes the cell, making it more difficult for an action potential to occur.

The pentameric ligand-gated ion channel (pLGIC) superfamily is particularly well-characterized. These channels open in response to neurotransmitters such as acetylcholine, serotonin, and GABA and are important for regulating neuronal signals in the body. These channels are also targets for important classes of drugs such as anesthetics and benzodiazepines. Furthermore, they are the only superfamily of neurotransmitter-gated ion channels that contains members that are highly cation-selective, such as the serotonin and acetylcholine receptor, and other members that are highly anion-selective, such as the glycine and GABA receptor. This feature makes pLGICs a model system for understanding the molecular basis for how ion channels are able to selectively catalyze the conduction of ions based on the formal charge of the ion.

METHODS AND RESULTS

Ion-selectivity has been studied experimentally for years using patch-clamp electrophysiology. While these experiments are able to observe patterns in ion-selectivity and elucidate mutations needed to change selectivity of a particular channel, electrophysiology does not have the resolution necessary to explain the molecular underpinnings of charge-selectivity. To understand charge-selectivity at atomic and molecular scales, computer simulations are necessary. To accomplish this task, free energy calculations are performed on a variety of member of the pLGIC superfamily. Utilizing NAMD, umbrella sampling simulations are performed on anion- and cation-selective ion channels in order to compute the free energy profiles for both cations and anions. In doing this, one can ascertain differences in ion conduction between the cation- and anion-selective channels. Furthermore, mutations to these channels can be done *in silico* to observe how the energy landscape for both types of ion changes, and these observations can be directly compared to experimental results. The comparison with experimental results is powerful because the two techniques are highly complementary. Experimental evidence provides verification for the observations seen in simulations, while simulations provide a highly detailed atomistic picture of the ion permeation process.

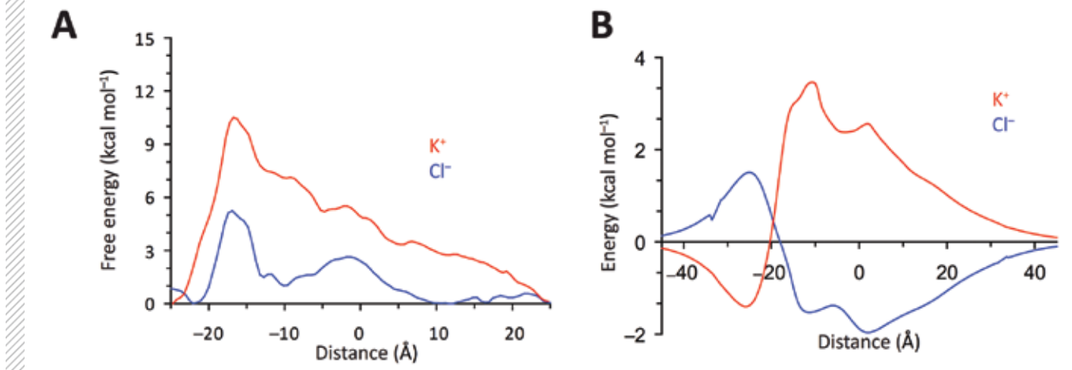

FIGURE 1: Energetics of ion permeation through GlyR: (A) PMF of potassium and chloride permeation through GlyR (PDB 3JAF) M1-M3 transmembrane segment using umbrella sampling molecular dynamics simulations. The transmembrane segments are centered about 0; positive numbers are toward the extracellular end of the protein and negative numbers are toward the intracellular portion. (B) Electrostatic ion-protein solvation energies for potassium and chloride, calculated using APBSmem. The distance axis corresponds to the same axis as in (A).

Figure 1A shows a potential of mean force plot generated using umbrella sampling for an anion-selective member of the pLGIC superfamily. This channel is clearly anion selective because the major barrier is much larger when the cation, potassium, is permeating through the channels relative to the anion, chloride. Figure 1B shows an electrostatic calculation for both a potassium and chloride along the axis of permeation. The electrostatic calculation shows the channel stabilizing the passing anion by creating an energetic well and destabilizing the passing cation by creating a barrier. Together, these calculations show that the charge selectivity of pLGICs is due to the electrostatics of the protein.

WHY BLUE WATERS

Ion permeation, especially in absence of a membrane potential, can be a slow process. Therefore, computational methods such as umbrella sampling have been devised in order to make the process of computing the energy landscape of permeation much more efficient. To do free energy calculations for ion permeation using all-atom molecular dynamics simulations, a large computational resource is needed. To make the calculation computationally efficient, umbrella sampling is used where the reaction coordinate that describes the permeation of an ion through the channel is discretized into many small windows that can be run independently. Umbrella sampling allows each umbrella window to be run as an independent simulation, which speeds the process of calculating free energies but requires a large number of processors to run each simulation efficiently.

NON-BORN-OPPENHEIMER EFFECTS BETWEEN ELECTRONS AND PROTONS

Allocation: Blue Waters Professor/24.0 Knh
PI: Sharon Hammes-Schiffer[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

The quantum mechanical behavior of nuclei plays an important role in a wide range of chemical and biological processes. The inclusion of nuclear quantum effects and non-Born-Oppenheimer effects between nuclei and electrons in computer simulations is challenging. Our group has developed the nuclear-electronic orbital (NEO) method for treating electrons and select nuclei quantum mechanically on the same level using an orbital-based formalism. We have used Blue Waters to perform NEO calculations on systems in which all electrons and one proton are treated quantum mechanically. Also, we have developed a methodology to study the nonadiabatic dynamics of photo induced proton-coupled electron transfer (PCET) reactions. The electronic potential energy surfaces are generated on-the-fly, and grid-based methods are utilized to calculate the nuclear wave function for the transferring proton. We have applied this methodology to a hydrogen-bonded phenol-amine complex in solution.

INTRODUCTION

The inclusion of nuclear quantum effects such as zero-point energy and tunneling in electronic structure calculations is important in a variety of chemical systems, particularly those involving hydrogen transfer or hydrogen-bonding interactions. Moreover, nonadiabatic effects, also called non-Born–Oppenheimer effects, between electrons and certain nuclei are significant for many of these systems. In this case, the electrons cannot be assumed to respond instantaneously to the nuclear motions, and the concept of the nuclei moving on a single electronic potential energy surface is no longer valid. This type of nonadiabaticity has been shown to play a critical role in PCET reactions, which are essential for a wide range of chemical and biological processes, including photosynthesis, respiration, enzyme reactions, and energy devices such as solar cells. The development of non-Born–Oppenheimer methods to enable accurate and efficient calculations of PCET reactions will impact many scientific endeavors, from drug design to the design of more effective catalysts for solar energy devices.

METHODS & RESULTS

In the NEO approach, typically all electrons and one or a few protons are treated quantum mechanically, and a mixed nuclear-electronic time-independent Schrödinger equation is solved. We have recently proposed an ansatz with the primary goal of improving computational tractability to enable the study of larger systems of chemical interest within the NEO framework. In this approach, denoted NEO-RXCHF, only select electronic orbitals are explicitly correlated to the nuclear orbital(s), and certain exchange terms are approximated. We recently incorporated an improved integral package into the NEO code. The new integral code enables the use of a direct algorithm for the NEO-RXCHF method and provides a multiple order of magnitude speed up while significantly reducing the memory overhead. Current efforts for the NEO-RXCHF method are focused on benchmarking proton basis sets for small molecules with the goal of obtaining accurate proton wavefunctions.

The methodology for investigating the mechanism of condensed phase photoinduced PCET reactions involves on-the-fly nonadiabatic dynamics on electron-proton vibronic surfaces using surface hopping in a hybrid quantum mechanical/molecular mechanical (QM/MM) framework. The solute is treated quantum mechanically, while the solvent is modeled using an empirical force field. We performed calculations of photoinduced PCET for a hydrogen-bonded phenol-amine complex in 1,2-dichloroethane solution. Our recent calculations treated the proton quantum mechanically, whereas our previous studies of this system treated the proton classically. Our new results highlight the important role of solvent reorganization in facilitating proton transfer to an excited electronic state with significant charge transfer character. Because of the absence of proton tunneling between localized proton vibrational states, replacing hydrogen by deuterium does not change the overall rate of decay to the ground state in this system. Thus, an important conclusion of this work is that the absence of an isotope effect on the overall relaxation rate does not necessarily signify the absence of proton transfer.

WHY BLUE WATERS

The NEO-RXCHF method requires the calculation of **trillions of integrals**. Our in-house NEO code has been parallelized using the message passing interface (MPI) protocol but still requires a large number of processors. Additionally, with the implementation of a direct algorithm for the NEO-RXCHF method, the fast calculation of the integrals is even more important, as the integrals must be calculated many times during the calculation instead of a single time at the start. The speed and scalability of Blue Waters is crucial in enabling large NEO-RXCHF calculations.

For studying photoinduced PCET with a quantum mechanical treatment of the transferring proton, the computational bottleneck is the calculation of the energies, forces, and nonadiabatic couplings between the electronic states at the proton grid points (typically 24) for each molecular dynamics time step. We have used MPI to distribute these calculations over the cores of a single node on Blue Waters. The large number of trajectories required by the surface hopping algorithm are made possible by the simultaneous use of a large number of nodes on Blue Waters. Moreover, the large memory per node and disk space on Blue Waters have made it possible to run these memory-intensive trajectories and store them for subsequent analysis. Assistance from the project staff has been crucial in successful building and testing of the locally modified MOPAC code (a program frequently used in computational chemistry) on Blue Waters.

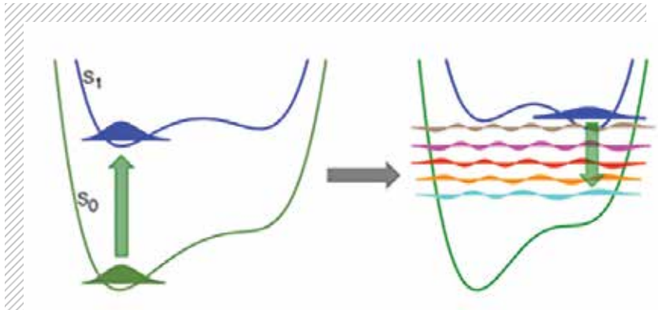

FIGURE 1: For the *p*-nitrophenylphenol--ammonia hydrogen-bonded complex in 1,2-dichloroethane solution, photoexcitation to an electronic state with significant charge transfer character is followed by solvent reorganization. The resulting change in asymmetry of the proton potential on the excited electronic state leads to proton transfer from the phenol to the amine. Decay to the ground vibronic state is dominated by vibrational relaxation on the ground electronic state.

NEXT GENERATION WORK

We hope to perform NEO-RXCHF calculations with inclusion of electron-electron correlation on larger molecular systems such as DNA base pairs on a next-generation Track-1 system. The current largest NEO-RXCHF calculation has been performed on a three-atom molecule without any electron-electron correlation, so this would represent significant progress for the NEO-RXCHF method. We also hope to study photoinduced PCET occurring on timescales longer than a few picoseconds as well as in large biomolecular systems. With a next-generation Track-1 system, it may also be possible to use an *ab initio* multiconfigurational method instead of a semiempirical multiconfigurational method in a QM/MM framework to simulate condensed phase photoinduced PCET.

PUBLICATIONS AND DATA SETS

Sirjoosingh, A., M. V. Pak, K. R. Brorsen, and S. Hammes-Schiffer, Quantum treatment of protons with reduced explicitly correlated Hartree-Fock approach. *J. Chem. Phys.*, 134 (2015), p. 214107.

Brorsen, K. R., A. Sirjoosingh, M. V. Pak, and S. Hammes-Schiffer, Nuclear-electronic orbital explicity correlated Hatree-Fock approach: Restricted Basis Sets and Open-Shell Systems. *J. Chem. Phys.*, 134 (2015), p. 214108.

Goyal, P., C. A. Schwerdtfeger, A. V. Soudackov, and S. Hammes-Schiffer, Proton quantization and vibrational relaxation in nonadiabatic dynamics of photoinduced proton-coupled electron transfer in a solvated phenol-amine complex. *J. Phys. Chem.* B, 120 (2016), pp. 2407-2417.

ICE AND WATER

Allocation: Blue Waters Professor/200 Knh
PI: So Hirata[1]
Collaborators: Soohaeng Y. Willow[1] and Michael A. Salim[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Systematic *ab initio* many-body approximations to electronic Schrödinger equations have transformed molecular sciences. Thanks to Blue Waters, they are now routinely applicable to a certain class of condensed matter: molecular crystals, amorphous solids, and liquids, enabling predictive simulations at such theoretical levels as *ab initio* many-body perturbation and coupled-cluster theories for a wide range of their structural, thermodynamic, and spectroscopic properties. This class of matter includes nature's most important and abundant solids and liquids: ice and liquid water. Here, we performed *ab initio* second-order many-body perturbation (MP2) calculation of liquid water to predict its structural, thermodynamic, response, and spectral properties. We also computationally reproduced the thermal contraction of ice-Ih (the hexagonal crystal form of ordinary ice) at low temperatures followed by expansion at higher temperatures, anomalous volume isotope effect, and pressure-induced amorphization to glass-like high-density amorphous phase.

FIGURE 1: The infrared and Raman spectra of liquid water. Reproduced from [1].

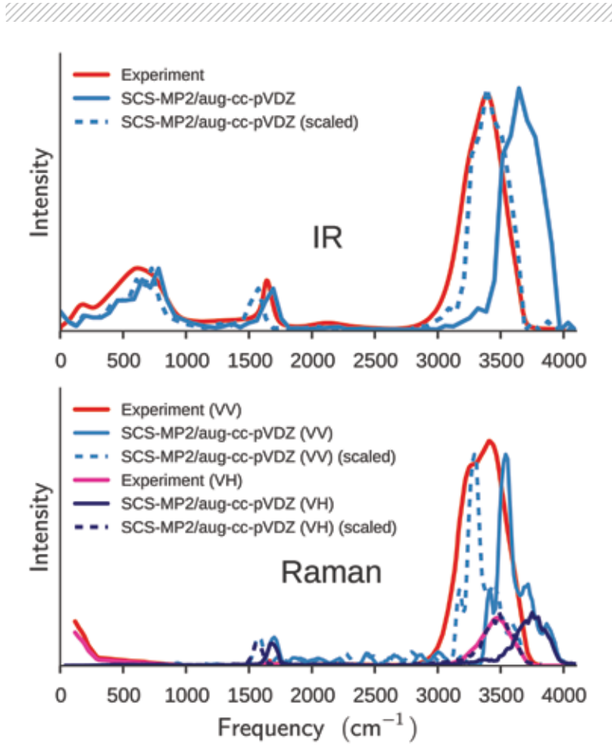

INTRODUCTION

As described more fully in [1], the properties of ice and liquid water have the most decisive influence on everything from climate to geology to biology on Earth. What may be surprising is that there are some unsolved mysteries and unsettled controversies surrounding these properties. For instance, unlike most other solids, ice-Ih contracts thermally at low temperatures before it expands at higher temperatures. Ice-Ih expands upon heavy-isotope (deuterium) substitution, which is the opposite behavior from normal solids and is said to have the anomalous volume isotope effect (VIE). Furthermore, ice-Ih undergoes a pressure-induced amorphization to become high-density amorphous (HDA) ice, which is distinct from low-density amorphous (LDA) ice. There seems perennial disagreement among researchers about the average number of hydrogen bonds per molecule and the mean dipole moment in liquid water.

Using Blue Waters, we performed direct applications of systematic *ab initio* many-body electronic structure theories, which go far beyond empirical force fields or density-functional theories that have dominated condensed-phase thermodynamic simulations, for a variety of properties of ice and liquid water. They mark the **beginning of a new era** of condensed-matter simulations with a systematic path to exactness.

METHODS & RESULTS

Our *ab initio* simulation of liquid water [2] is based on molecular dynamics (MD) using on-the-fly atomic forces calculated fully quantum mechanically in each MD step by highly scalable embedded-fragment MP2 method.

Our simulation accurately predicts the oxygen-oxygen radial distribution function of liquid water, i.e., its average structure, and shows that the distribution of the hydrogen-bond number is strongly peaked at four, confirming "the standard picture" of liquid water with an average hydrogen-bond number of 3.8, while disagreeing with "the string theory" of liquid water that asserts that each water molecule has an average of only two hydrogen bonds in a chain or a ring. It also reproduces the general intensity profiles of the observed infrared and parallel- and perpendicular-polarized Raman spectra extremely well (Fig. 1), including the Raman noncoincidence effect.

Ice-Ih, unlike other solids, contracts upon heating at low temperatures (< 70° K), and then expands at higher temperatures (solid curves in Fig. 2). Upon deuterium substitution for hydrogens, it expands by 0.09% at 0° K (the anomalous VIE; solid curves in Fig. 2). These are quantum effects of lattice vibrations in anharmonic potentials, which require high-power methods to explain computationally. Our study [3] reproduces the thermal contraction at low temperatures quantitatively and the thermal expansion at higher temperatures qualitatively (dots in Fig. 2). The study confirms the origin of the thermal contraction as the volume-collapsing effect of hydrogen-bond bending modes, whereas the subsequent thermal expansion is caused primarily by the hydrogen bond-stretching modes.

The same MP2 calculation gives qualitatively varied predictions about the anomalous VIE, depending on calculation details, although the one shown in figure 2 reproduces it quantitatively. The extreme sensitivity to VIE is traced to the fact that net VIE at 0° K is the result of delicate cancellation of competing (i.e., volume-expanding and contracting) effects of *all* fundamental phonons that exist in the zero-point state.

Compression of ice-Ih at slightly below 0° C results in a phase transition to liquid water because the latter is denser than the former; the thermodynamic phase boundary then has a negative slope. In 1984, Mishima et al. [4] extrapolated the negative-slope melting curve to lower temperatures, predicting and then observing a pressure-induced transition from ice-Ih to glass-like disordered HDA phase at 1.0 GPa and 77° K.

With increasing pressure, MP2 geometry optimization of ice-Ih, starting from an initial crystalline structure, becomes more difficult at pressures greater than 2.35 GPa. By 3 GPa, the optimized crystalline structure never converges to a minimum-energy crystalline structure, but instead, leads to a loss of symmetry and long-range order with a 15% reduction in volume relative to the ambient-pressure value. Figure 3 contrasts the optimized crystalline structure at 0 GPa and a nonconverged, semi-amorphous structure at 3 GPa. Concomitant with the difficulty in optimizing the geometry, we observe pressure-induced red-shifting of acoustic phonons, starting at around 2 GPa. Our calculated thermodynamic transition pressure between the two phases is 1.6 GPa at 0° K, which is consistent with experimental values. These constitute the first *ab initio* computational detection of the pressure-induced amorphization of ice-Ih.

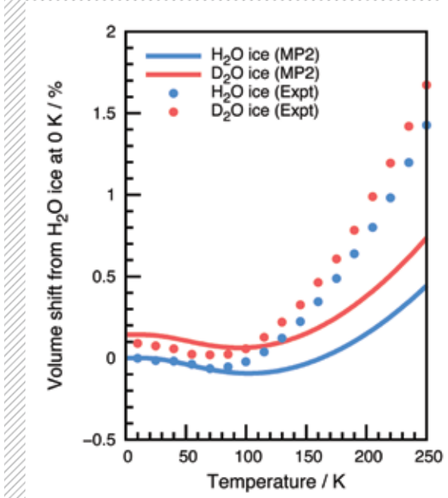

FIGURE 2: Temperature-dependence of the volume of H2O and D2O ice-ih. Reprinted [3].

WHY BLUE WATERS

A first-principles quantum-mechanical calculation of an infinitely extended solid, not to mention liquid, was previously unthinkable. Blue Waters has now **made these calculations a reality** using a combination of computing power and the algorithmic breakthrough (embedded-fragmentation) that exposes scalability with both system and computer sizes.

NEXT GENERATION WORK

We will fully develop a software system that will allow routine applications of predictive *ab initio* all-electron quantum-mechanical methods to a whole range of properties of any molecular solids and molecular liquids on supercomputers. Thanks to Blue Waters, it will no longer be necessary to rely on empirical potentials or density-functional approximations.

PUBLICATIONS AND DATA SETS

Hirata, K. et al., *Ab initio* ice, dry ice, and liquid water. *Fragmentation: Toward Accurate Calculations on Complex Molecular Systems*, Ed by Mark S. Gordon, in review.

Salim, M.A., S. Y. Willow, and S. Hirata, Ice Ih anomalies: Thermal contraction, anomalous volume isotope effect, and pressure-induced amorphization. *J. Chem. Phys.*, 144, (2016), doi: 10.1063/1.4951687

Willow, S.Y., M. A. Salim, K. S. Kim, and S. Hirata, *Ab initio* molecular dynamics of liquid water using embedded-fragment second-order many-body perturbation theory towards its accurate property prediction. *Scientific Reports*, 5 (2015), doi:10.1038/srep14358

Willow, S.Y., et al., Why is MP2-water "cooler" and "denser" than DFT-water? *J. Phys. Chem. Lett.* 7 (2016), p. 680-684, doi: 10.1021/acs.jpclett.5b02430

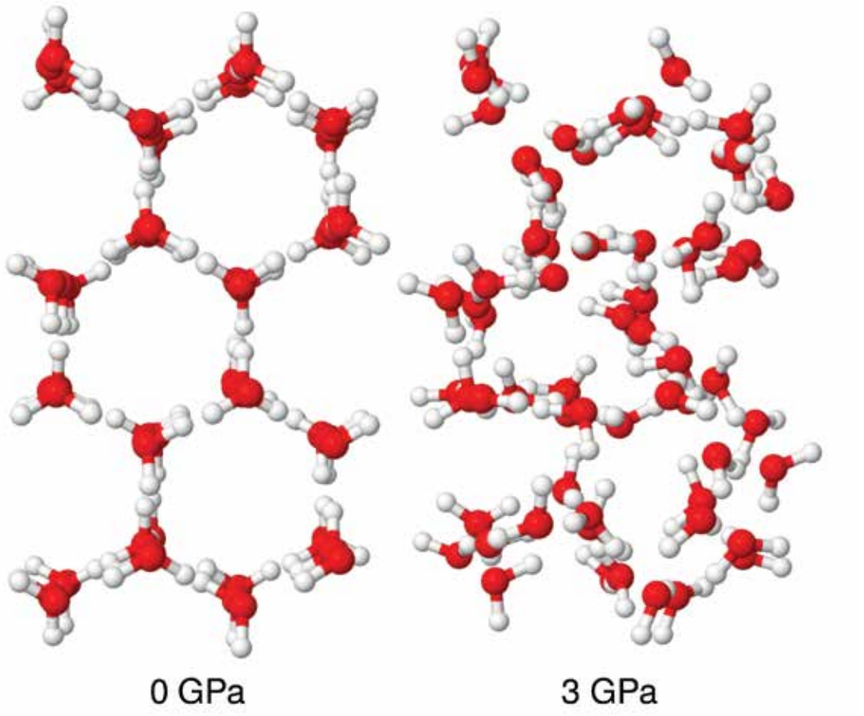

FIGURE 3: Optimized (left) and partially optimized (right) structures of ice-Ih and ice-HDA. Reprinted from [3].

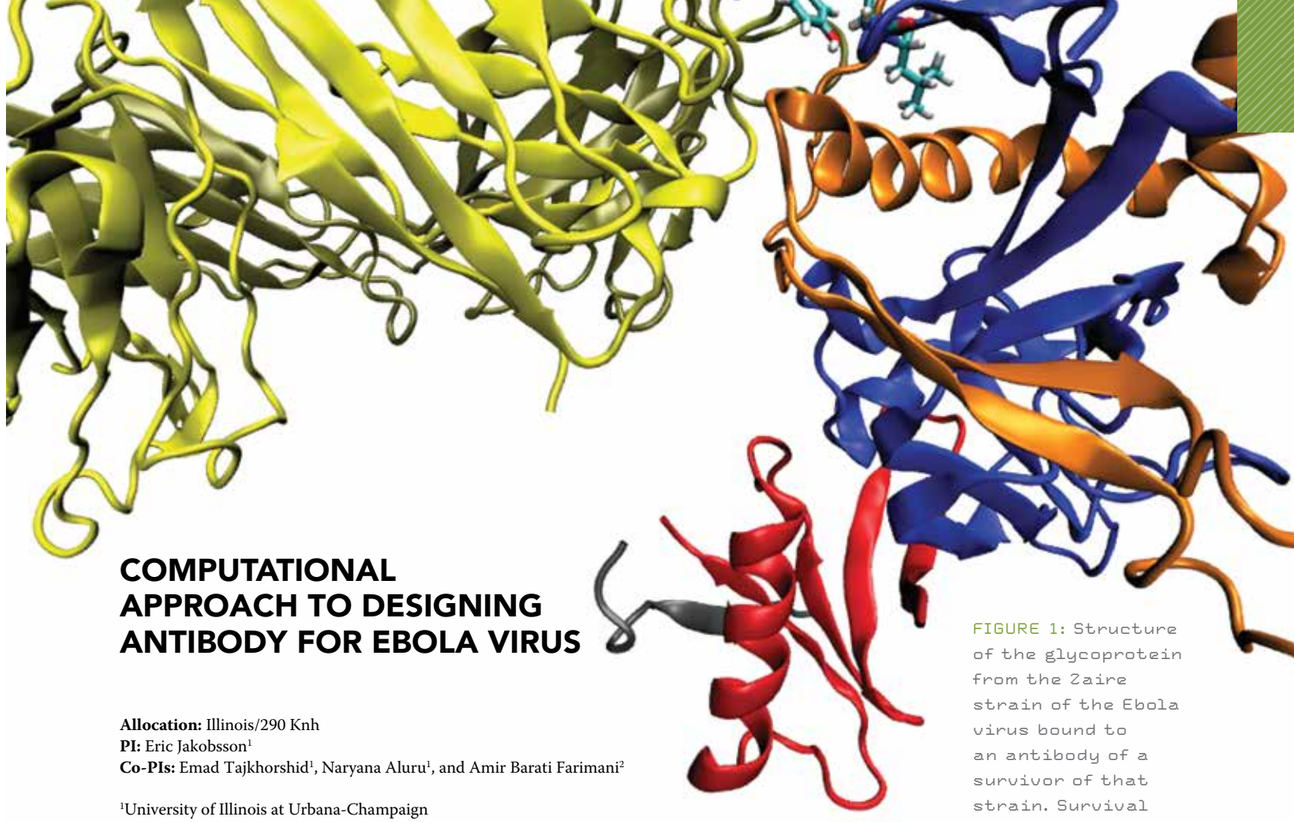

COMPUTATIONAL APPROACH TO DESIGNING ANTIBODY FOR EBOLA VIRUS

FIGURE 1: Structure of the glycoprotein from the Zaire strain of the Ebola virus bound to an antibody of a survivor of that strain. Survival was because of the effective binding of the glycoprotein to the virus pictured. The glycoprotein is essential to the virus' ability to enter the host cell.

Allocation: Illinois/290 Knh
PI: Eric Jakobsson[1]
Co-PIs: Emad Tajkhorshid[1], Naryana Aluru[1], and Amir Barati Farimani[2]

[1]University of Illinois at Urbana-Champaign
[2]Stanford University

EXECUTIVE SUMMARY

Our work on Blue Waters successfully demonstrated the feasibility of computational design of synthetic antibodies against evolving Ebola infections. We simulated multiple cycles of viral mutation and redesign of a synthetic antibody to counter the mutations successfully and restore high-affinity binding of the antibody to the virus. Viral mutations were selected by random walk theory biased according to the statistical propensity for amino acid substitution. Trial substitutions for redesign were selected according to the statistical propensity for forming favorable interfaces. The success of the redesign was evaluated by using molecular dynamics (MD) to compute viral protein-antibody binding energy. Blue Waters provided the essential computational power to do the many simulations to test the ability to redesign successfully, and we feel this approach should be extendable to other viruses. In combination with experimental sequencing and structure determination, our approach should enable rapid design and redesign of synthetic antibody therapy in response to rapidly evolving viral challenges.

INTRODUCTION

The ability to produce antibodies specific to predefined biomolecular targets was a landmark development in biological research and potential therapy [1]. At the core of this work was engineering at the cellular level, in particular induced cell fusion to produce hybrid cells. A fundamental advance was to extend the engineering to the molecular level, including the engineering of chimaeric antibodies [2]. It should be possible to engineer antibodies against viruses, specifically, the coat glycoprotein that is an essential component of the entry mechanism for Ebola and many viruses into the host cell [3]. Nature has afforded us with a proof-of-concept for such engineering, by providing us with definable sequences and structures for Ebola glycoprotein bound to antibodies that enabled the host to survive the disease [4]. On the other hand, nature makes our task harder by enabling the virus to evolve in such a way as to neutralize the effect of the antibody [5]. Sequences and structures provide an outline of how the evolutionary arms race proceeds between Ebola virus glycoprotein vs. antibodies from the host immune system for multiple glycoprotein-

FIGURE 2: RMSD (magnitude of structural fluctuations, a measure of structural instability) derived from MD simulations when Sudan strain antibody is matched with Sudan virus, when Sudan antibody is matched with Zaire virus, when Zaire strain antibody is matched with Zaire virus, and when Zaire antibody is matched with Sudan virus. This image shows proof-of-concept that molecular dynamics is sufficiently sensitive to distinguish between effective binding (that can overcome the infection by preventing viral entry into cells) and ineffective binding.

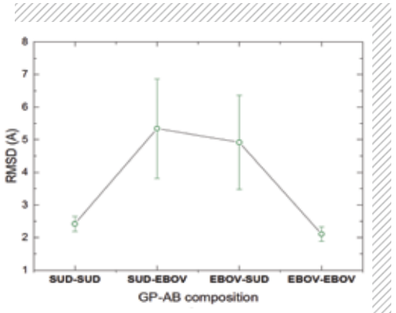

antibody complexes [6]. However, this description of the glycoprotein-antibody competition does not in itself lead to a predictive model for how the virus will evolve and what change in the antibody will be effective against the evolved viral protein.

METHODS & RESULTS

We began by considering how to construct a predictive model for how the virus is likely to evolve, and what alterations in the sequence of a binding region of the antibody would most effectively counter the viral mutation(s) and restore the ability of the antibody to bind the glycoprotein. To predict likely mutations, we used existing statistical data on the likelihood of particular substitutions, as embodied in a "substitution matrix" in which each element corresponds to a relative probability of an amino acid substitution [7]. To predict effective responses to viral mutations, we used existing statistical data on amino acids that interact favorably at protein-protein interfaces [8]. Finally, we used MD simulations of the mutated glycoprotein-antibody complex to test the statistical prediction by computing the effects of the postulated mutations [9].

The starting points for the simulations were structures of Ebola glycoprotein complexed with antibody fragments that were known to prevent infection successfully. We then mutated the glycoprotein, which invariably resulted in degradation of binding energy between glycoprotein and antibody. Following that, we used the databases of favorable amino acid interactions to make educated guesses as to mutations on the antibody. In the majority of cases, we were able to re-engineer the antibody to bind the viral protein as well or better than the wild type. In those cases where we did not succeed in doing that, we believe that we would have succeeded if provided with more computertime.

WHY BLUE WATERS

We could not have done the project without the sheer computational power of Blue Waters.

NEXT GENERATION WORK

Our long-term goal is to establish our approach as a standard for the design and redesign of synthetic antibodies against viral infections so that humankind can prevail in the evolutionary arms race against evolving viral pathogens.

PUBLICATIONS AND DATA SETS

Farimani, A.B., et al., Computational approach to designing antibody for Ebola virus. *Biophyss J.*, 110:3 (2016)537a.

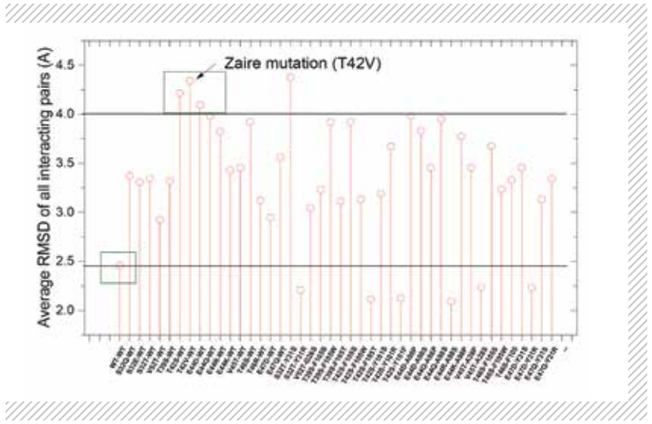

FIGURE 3: RMSD for selected mutants of the Sudan virus and mutated antibodies. The symbols that fall below 2.5 angstroms are actually binding better than wild type. Antibody redesign succeeded by that criterion for approximately half of the mutant glycoproteins.

UNVEILING ALLOSTERIC PATHWAYS IN ION CHANNELS

Allocation: NSF PRAC/5.60 Mnh
PI: Michael L Klein[1]
Co-PI: Vincenzo Carnevale[1]

[1]Temple University

EXECUTIVE SUMMARY

Transient receptor potential (TRP) channels are central to environmental sensation in animals, fungi, and unicellular eukaryotes. All known TRP channels are nonselective cation channels that open in response to a wide array of factors. Clarifying how TRP channels convert physical and chemical stimuli from the environment into the allosteric signals underlying channel activation is key to understanding how they control cell excitability in both physiological and pathological conditions. Their relevance in the molecular pathways that mediate pain makes them promising targets for novel classes of analgesics (medicines that relieve pain). Building on the structural information made recently available for transient receptor potential cation channel subfamily V member 1 (TRPV1), thanks to a series of cryo-electron microscopy (CryoEM) experiments, we performed free energy (metadynamics) simulations on models of TRPV1 embedded in a lipid bilayer. Harnessing the computation capabilities of Blue Waters, we explored several pathways of activation and characterized ion channel conductance and selectivity. Our calculations reveal a **novel** mechanism for sensing temperature and osmolality.

INTRODUCTION

A fit cell must perceive and comprehend the conditions of its inner and outer worlds, integrating diverse and transitory physicochemical stimuli into concerted cellular decisions. For this reason, the membranes of even the simplest bacteria are studded with ion channel proteins that detect cellular conditions and translate them into electrochemical information via gated ionic conduction [1].

In eukaryotes, the complexity of cellular life has taken this requirement to its apex. Accordingly, natural selection has elaborated on the ion channel, producing an impressive array of polymodal cellular sensors, the TRP channels [2]. All TRP channels detect multiple physicochemical stimuli, with some overlap among the eight extant TRP subfamilies. However, the response to each stimulus varies substantially from channel to channel, presumably dictated by heterogeneous and subfamily-specific intra- and extracellular domains [3]. Indeed, since their divergence from the voltage-gated potassium (Kv) channel superfamily over a billion years ago, TRP channel proteins have maintained a tetrameric six-transmembrane (6-TM) architecture and little else. The (TRPV1) or vanilloid receptor 1 is a polymodal mammalian nociceptive integrator [4] abundantly expressed in the free nerve endings of primary pain-sensing afferent Aδ and C fibers [5]. Structurally, the TRPV1 channel is a homotetramer, symmetrically organized around a solvent exposed central pore. Each subunit is formed by six transmembrane helices (S1–S6) with the channels' N- and C-termini located in the intracellular medium [6].

TRPV1 is activated by a wide range of proinflammatory and proalgesic mediators [7]; including temperatures above 43°C, external pH, bradykinin, anandamide, arachidonic acid metabolites, jellyfish and spider toxins, vanilloid and others. The scope of the TRPV1 pharmacological spectrum [8-10] is mainly in the area of analgesics: novel painkillers could be either TRPV1 agonists or antagonists. Moving forward toward the rational drug design of TRPV1 modulators requires a basic understanding of how known ligands trigger the closed to open transition in TRPV1.

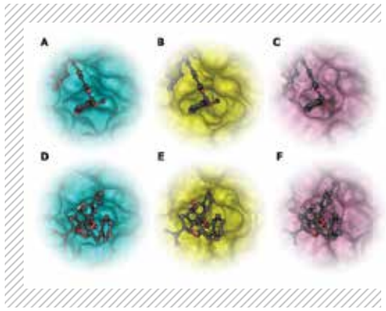

FIGURE 1: Optimal docking poses of capsaicin and resiniferatoxin in the vanilloid binding site of TRPV1. Shown are: A) capsaicin in the apo structure; B) capsaicin in the TRPV1-capsaicin complex; C) capsaicin in TRPV1-resiniferatoxin complex; D) resiniferatoxin in the apo structure; E) resiniferatoxin in the TRPV1-capsaicin complex; F) resiniferatoxin in the TRPV1-resiniferatoxin complex. Adapted from ref. (Elokely, 2015)

We investigated the detailed molecular mechanism of activation of TRPV1. We found that dynamically controlled solvation of hydrophobic protein pockets is the key determinant of activation. We are now in the process of extending this investigation to address the relevant issue of how this structural transition is triggered by environmental stimuli such as temperature, osmotic pressure and binding of PIP_2 lipids.

METHODS & RESULTS

We performed extensive multi-microsecond molecular dynamics (MD) simulations on several TRPV1 systems, varying the initial setup and the equilibration protocol for the purpose of characterizing stable conformational states close to the experimental structures.

Binding mode of agonists

Initially, we focused on the binding mode of agonists, an issue that was not solved by the cryoEM investigation. We used the information contained in the experimental electron density maps to determine the binding mode of capsaicin and resiniferatoxin. In collaboration with the lab of Tibor Rohacs at Rutgers New Jersey Medical School, we validated the theoretical predictions via mutagenesis experimentally. By these results, we characterized the structural rearrangements entailed by binding of each ligand to the pocket (Fig. 1).

Open to close transition

We then started MD simulations of the capsaicin-bound state in a standard POPC (1-palmitoyl-2-oleoyl-sn-glycero-3-phosphocholine) bilayer to explore the mechanism of channel opening. At the end of the 750 ns MD trajectory, we analyzed the root-mean-square deviation (RMSD) from the initial capsaicin-bound structure and from the open structure. Surprisingly, we found that two subunits out of four diverged from the initial state and relaxed to the open state. Consistently, after approximately 400 ns of simulation, we observed a sudden increase in the pore radius and a simultaneous hydration of the so-called hydrophobic gate. We thus compared the radius and the hydration profile of the resulting open state to those of the closed state and found significant differences (Fig. 2). By performing metadynamics simulations, we finally ascertained that the conformational state obtained from the simulation is conductive for sodium (Fig. 2C).

WHY BLUE WATERS

The project relied crucially on a quantitative description of complex processes occurring on time-scales of several microseconds in large membrane-protein assemblies with a typical size of approximately 300,000 atoms. The system is constituted by an ion channel, a model lipid bilayer, and an electrolyte solution and was simulated under different conditions, i.e. in presence or in absence of a ligand. The capabilities of Blue Waters turned out to be key to the success of this computationally intensive project.

NEXT GENERATION WORK

The complexity behavior of ion channels relies in part on the peculiar environment of the lipid membranes present in neurons. The insight developed during this project will enable computational modeling of the response of a channel in a neuron-like environment; ultimately, the goal is to develop a computational framework to quantitatively address the following question: how does the nervous system transduce and transmit the information about noxious environmental stimuli?

PUBLICATIONS AND DATA SETS

Elokely, K., et al., Understanding TRPV1 activation by ligands: Insights from the binding modes of capsaicin and resiniferatoxin, *PNAS,* 113:2 (2016), pp. E137-E145, doi:10.1073/pnas.1517288113

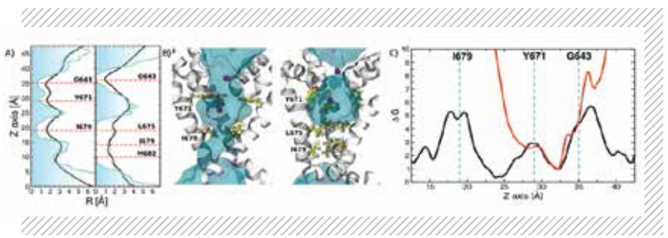

FIGURE 2: Open and closed states of TRPV1. A) The radius of the pore (black line) is shown along with the water density profile (blue shading) for the open (left) and closed (right) conformation. B) Three-dimensional shape of the pore (blue shading) is shown together with pore-lining residues. C) Free-energy profile for the permeation of a sodium ion as calculated form metadynamics for the open (black) and closed (red) confomations. Figure adapted from ref. (Kasimova, 2015).

RIBOSOME BIOGENESIS IN REPLICATING CELLS

Allocation: Illinois/606 Knh
PI: Zaida Luthey-Schulten[1]
Co-PIs: Tyler Earnest[1] and John Cole[1]
Collaborator: Thomas Kuhlman[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Ribosomes—the cellular machines responsible for making proteins—are complexes of nucleic acids and proteins. The ways in which the living cell manages their production and assembly is of deep biological importance and is highly conserved among all bacteria. Here we extend a recent spatially resolved whole-cell model of ribosome biogenesis in cells of a fixed volume [1] to include the effects of DNA replication, growth, and cell division [2]. Cell cycle parameters describing the replication schedule were obtained by analyzing single cells from a series of *Escherichia coli* strains with fluorescently labeled genes distributed evenly throughout the chromosome. For cells with a two hour generation time, replication was initiated 42 minutes into the cell cycle and completed after an additional 42 minutes. All of the biological processes of replication, transcription, translation, and ribosome assembly are described regarding reaction-diffusion master equations and solved stochastically using Lattice Microbes v2.3a software package [3-5].

INTRODUCTION

In bacteria, ribosomes account for approximately one-fourth of the cellular dry mass and the majority of the total RNA [1]. The ribosomes' role in protein synthesis couples them to essentially every process within the cell. It can be tempting to think of the bacterial cell as a finely tuned machine for building ribosomes. Previously, we published a model of ribosome assembly in the cell based on kinetic and thermodynamic data [2-3]. Although unprecedentedly complete, the model did not account for some of the most basic functions of the cell—DNA replication, growth, and cell division. Duplication of the chromosome affects the copy number distribution of gene products since the rate of transcript production effectively doubles [4]. However, to correctly capture the shape of the distribution the behavior over the interval where the mRNA copy number relaxes to the new steady state under the doubled gene dose must be accounted for [8-9]. The changing cell volume also impacts the copy number dynamics due to the dependence of reaction rates on volume. To account for this, we have augmented this model by deterministically modeling cell growth and DNA replication using experimental data from our collaborator, Thomas Kuhlman. *E. coli* follows a simple three-period cell cycle based on the timing of DNA replication and cell division (Fig. 1). Using estimates of the durations of these periods as well as the mean cell length at division computed from experiments performed by our collaborator, we extended our model to include gene replication and cell growth through deterministically modifying the cell geometry, operon locations, and copy numbers over the course of the cell cycle.

METHODS & RESULTS

Previously, we constructed a kinetic model of the biogenesis of the ribosomal small subunit (SSU), including the transcription of rRNA and mRNA coding for the ribosomal proteins and transcription and translation of mRNA [1]. This model was embedded in a spatial model of *E. coli*, derived from cryo-electron tomograms [10], which takes into account the geometry of the cell and its compartmentalization into four distinct regions: extracellular, membrane, cytoplasm and nucleoid (Fig. 1). This system was simulated on Blue Waters using the Lattice Microbes v2.3a software package, which includes an improved algorithm for computing reaction propensities [6] (Fig. 3).

Our simulated cell was allowed to grow from 2.4 to 4.8 μm over its 120-minute cell cycle. Genes are replicated in the simulation at times determined from the cell cycle parameters and their position in

the genome. Compared to a fixed-volume version of this model (without growth and cell division), the initial and final species counts are practically identical for all protein types, ribosomal subunits, and translating ribosomes. However, since the volume expansion of the cell proceeds at the same rate as the new species are created, the cell can maintain a constant concentration of ribosomes and ribosomal protein. Finally, the importance of a stochastic, spatially resolved representation is emphasized yet again due to the limited copy number and lifetimes of the assembly intermediates.

WHY BLUE WATERS

Our simulations are computationally intensive and require high-performance graphics processing unit (GPU) accelerators. The stochastic nature of these simulations requires multiple realizations of the cell cycle to properly investigate the statistical variation of the cell's phenotype. These simulations are I/O intensive: on average, 20 GB of data is generated for a single realization of a cell cycle. This demands both high-capacity storage as well as high bandwidth I/O. Blue Waters provided the technical resources necessary to study this system. The recent introduction of a Blue Waters Python environment (BWPY) by the staff allowed us to streamline our development/simulation cycle by providing an up-to-date and complete Scientific Python stack optimized for Blue Waters. Since our code relies on modern Python features, this relieved us from the task of building our Python environment optimized for a supercomputing environment.

NEXT GENERATION WORK

In the next generation Track-1 systems, we hope to build a highly realistic *in silico* bacterium that will incorporate spatially resolved signaling, transcription, replication, translation and metabolic networks which include chemical species at densities spanning many orders of magnitude. The computational challenge of simulating the numerous metabolites and reactions needed to describe these processes will require the use of multiple nodes as well as multiple GPU accelerators per node for a realization of a single cell. Next-generation GPUs allow significantly faster time-to-solution (Fig. 2) and will allow the study of longer timescales while collecting the requisite sample size for proper statistical analysis. This challenge will necessitate the continual development of our Lattice Microbes Message Passing Interface code and implementation of hybrid deterministic–stochastic treatments of species allowing for a density scale separation, which will avoid the use of computationally intensive stochastic treatments unnecessary for highly abundant species.

PUBLICATIONS AND DATA SETS

Earnest, T.M., et al., Toward a Whole-Cell Model of Ribosome Biogenesis: Kinetic Modeling of SSU Assembly, *Biophys. J.*, 109:6 (2015), pp. 1117–1135. doi:10.1016/j.bpj.2015.07.030

Earnest, T.M., et al., Ribosome biogenesis in replicating cells: Integration of experiment and theory. *Biopolymers*, (2016). doi:10.1002/bip.22892

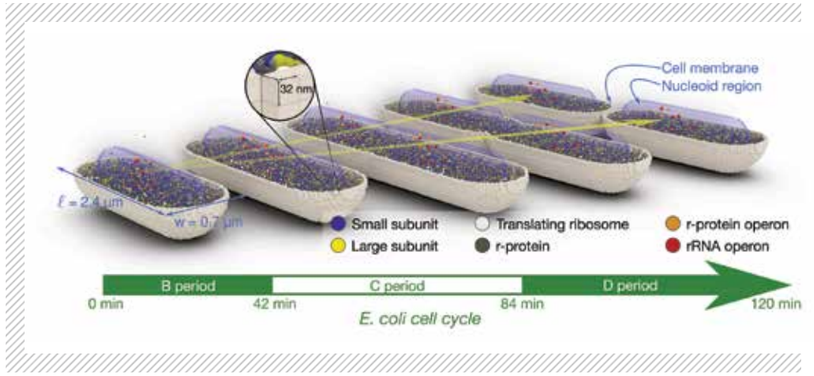

FIGURE 1: Schematic of whole-cell simulations of ribosome biogenesis in slow growing *E. coli*. The simulation is performed on a 32 nm lattice of dimensions 32×32×192, with four compartments representing the extracellular, cell membrane, cytoplasm, and nucleoid regions. The cell cycle is divided into three periods: the time after cell division but before the initiation of DNA replication (B period), the period of DNA replication (C period), and the lag time between the termination of replication and completion of cell division (D period). Using experimental estimations of the duration of these periods, we are able to schedule the replication of the ribosomal operons (red and orange spheres) during the cell cycle. These operons are placed along the long axis of the cell based on their location in the genome. The cell grows from 2.4 µm to 4.8 µm over 120 minutes, increasing its particle count from ~50,000 to ~100,000, while new ribosomal operons are added at times derived from the cell cycle parameters and the genomic distance from the origin of replication, which are moved along the long axis of the cell such that the operons are found in the same relative positions in the daughter cells as in the mother cell (yellow arrow).

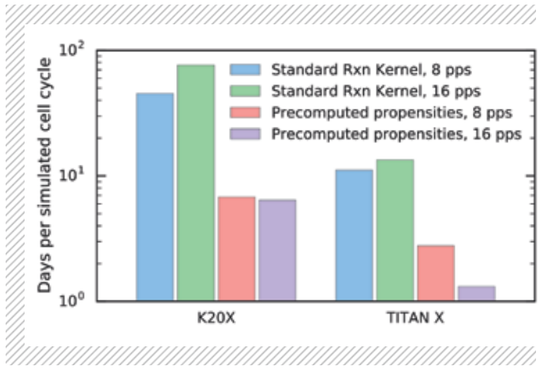

FIGURE 2: Compute time necessary to simulate a 120-minute cell cycle. The standard reaction kernel computes the reaction propensity for each possible reaction, whereas the precomputed propensities kernel uses precomputed values and loops over the number of particles per site. The latter method is significantly faster due to the unusually large number of reactions (800) in the model. By increasing the total number of particles from 8 to 16 per lattice site, a speed-up is achieved in spite of the increased time necessary to transfer the lattice data. Since this model can attain high particle densities, the greater lattice capacity reduces the rate of overflows, which occur when a site has exceeded its capacity and the new particle must be placed in a nearby site. An approximate 6× speed-up switching from the Kepler architecture (K20X, 6 days) to Maxwell (TITAN X, 1 day) was observed.

CUSTOM GENOTYPING CHIP FOR AFRICAN POPULATIONS

Allocation: Illinois/250 Knh
PI: Liudmila Sergeevna Mainzer[1]
Collaborators: Gerrit Botha[2], Victor Jongeneel[1], Ayton Meintjes[2], Nicola Mulder[2], Gloria Rendon[1]

[1]University of Illinois at Urbana-Champaign
[2]University of Cape Town

EXECUTIVE SUMMARY

Our goal was to aid in the design of a cost-effective genotyping chip that would capture the genetic diversity in populations of African origin, including African-Americans. This collaborative project was part of the H3Africa Custom Chip Design task force led by Zane Lombard at the University of Witwatersrand and Debo Adeyemo of the National Human Genome Research Institute. This work will enable the identification of genetic variation specific to African populations, improve understanding of the links between genotype and disease in people of African origin, and extend the principles of personalized medicine to these underserved populations. It will also permit deeper study of African genetic diversity, bringing important insights into the history and evolution of humans in general.

INTRODUCTION

Much of what is known about the genetics of diseases is based on people with European ancestry. The Consortium for Human Heredity and Health in Africa—H3Africa—aims to change that by promoting health research that takes into account the genetic diversity of African populations. Genomic variation plays a large role in disease predisposition and drug response. Thus, it is important to develop tools for genomic variant discovery specifically in people of African descent, who have been underrepresented in many worldwide genetic diversity measurement projects. The Genome Analysis Working Group from the H3Africa consortium partnered with H3ABioNet and the Wellcome Trust Sanger Institute to construct a genotyping chip to test for genomic variants found specifically in African populations. The chip will be a tool for rapid and inexpensive genotyping of individuals, to aid in the studies of human evolution and to identify genomic bases of disease. The data include publicly available sequence data from the 1000 Genomes Project and over 2,000 low-depth whole genome sequences from the Sanger Institute's African Genome Variation Project. An additional set of 348 samples was sent for deep sequencing (30X) at the Baylor College of Medicine thanks to a funding supplement from the National Institutes of Health. Blue Waters was used to perform genomic variant calling analysis on this set of 348 deeply-sequenced whole human genomes.

METHODS & RESULTS

Scientists at the University of Cape Town led by Nicola Mulder, together with Dr. Manj Sandhu's team at the Wellcome Trust Sanger Institute, developed the computational workflow to extract genomic variants from the 348 samples sequenced at Baylor. The workflow followed best practices [1], with additional steps ensuring quality control, robustness and mechanisms of recovery from failure added by collaborators at the University of Illinois. This workflow was instantiated on Blue Waters and used to analyze the data in six batches, ranging from eight to 87 samples. Production was completed in 51 days, producing genomic variant calls on the entire dataset of 348 individuals, and those calls are now being used to make the final design for the genotyping chip. The final product, the chip itself, will be used in biomedical research throughout the world.

A number of data management issues had to be resolved to make this possible. The Blue Waters team engaged in debugging data transfers from Baylor to Illinois and from Illinois to South Africa. The data-transfer challenges we faced led directly to learning and evaluating community data transfer tools, such as bbftp. Ultimately all data transfers were successfully completed using Globus, but our work to test the functionality and performance of additional tools will provide useful alternatives, should the need arise.

We also demonstrated, in a production-grade project, the capability of Blue Waters to **conduct high-throughput analysis** of human genomes. Extensive benchmarking data were collected, and the computational workflow was hardened with many quality control steps to ensure delivery of correct results. The code is posted on GitHub, to be shared with the community: https://github.com/HPCBio/BW_VariantCalling.

WHY BLUE WATERS

Several hundred deeply sequenced human samples is a lot of genetic data, which have a large disk footprint and take a long time to analyze. The available capacity of the University of Cape Town's computing cluster was not sufficient to process all of the data in the required time. Blue Waters enabled them to process this wealth of data in a timely manner, using nearly 250,000 node-hours, even with extensive use of backfill and accuracy discounts. The total disk footprint was nearly 600 TB, while only 200 TB total was available on the Cape Town cluster.

The Blue Waters team has been **instrumental** to the success of this project. **Extremely high uptime**, rapid resolution of failures (in one case under 20 minutes) and close collaboration on data transfers between the United States and South Africa have made the Blue Waters support team an **integral** part of the project. Without their involvement and professionalism, this work would have been very difficult to complete. In our case, the team was as important as the compute, storage, and networking resources.

INSTRUMENTING HUMAN VARIANT CALLING WORKFLOW AT SCALE

Allocation: Illinois/619 Knh
PI: Liudmila Sergeevna Mainzer[1]
Collaborators: Arjun Athreya[1], Subho Banerjee[1], Ravishankar K. Iyer[1], Victor C. Jongeneel[1], Volodymyr Kindratenko[1], and Zachary Stephens[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Whole genome sequencing and analysis are becoming part of the clinical standard of care. President Obama's 2015 Precision Medicine Initiative included genomics as an inextricable component in development of medical treatment and prevention strategies. Understanding the associated computational challenges is necessary in order to plan and construct the computing infrastructure that will support very high-throughput analyses in regional genomic sequencing centers across the country. The information we obtained in this project will help make such design recommendations, from relatively small clusters to large supercomputers. Specifically, we addressed problems associated with workflow scheduling, job management and recovery, file distribution, auto-archiving, and workflow scalability. We identified, documented, and resolved the bottlenecks associated with the large number of small files created by the workflow, saturated I/O bandwidth for part of the workflow, and potential for unbalanced data load on the file system. The resultant codes are available on GitHub at https://github.com/HPCBio/BW_VariantCalling.

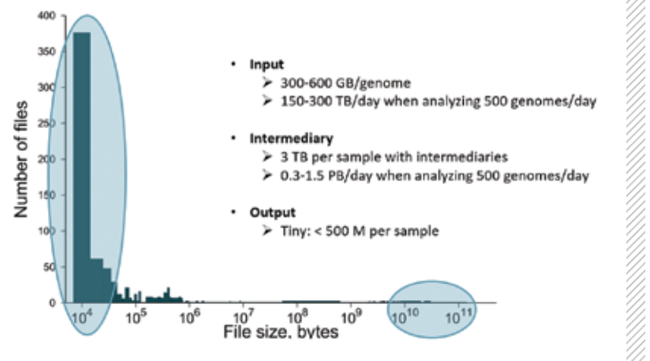

FIGURE 1: File size distribution for all the intermediate and output files produced in a variant calling workflow on one sample. Inset shows total disk footprint per sample, and per batch of 500 samples.

INTRODUCTION

Human variant calling (genotyping) is a process of searching for the differences between an individual's genome and that of the average in a human population. It is widely used to diagnose diseases and study human diversity and evolution. Genotyping every patient arriving at a given hospital is likely to become routine within our lifetime. This will require powerful computational facilities. Genotyping every newborn in Illinois would require analysis of ~500 genomes daily. At this scale, the standard workflow accepted by the community would use thousands of nodes in parallel and have computational bottlenecks affecting performance and reliability of the system. Our goal in this project was to identify such bottlenecks, construct workarounds, and come up with recommendations for the kind of computational/analysis systems that could handle this high throughput.

METHODS & RESULTS

We set up the "standard practices" workflow [1] and tested a number of tools [2,3] to shorten the wall time required to complete computation of each genome. We were able to accomplish a total run time of ~24 hours to process a single whole human genome, sequenced at the depth of 50X, using 25 nodes in parallel. In addition, we benchmarked the CPU, RAM, and I/O utilization across the workflow using Perfsuite [4], Cray Profiler [5], OVIS [6], Valgrind [7], and our own software. The data suggest two classes of problems are associated with genomic variant calling at-scale: workflow management and data management.

Workflow management: The workflow consists of multiple steps, each taking substantial time. It is therefore common practice to separate the steps into a series of independent jobs. When analyzing hundreds of genomes simultaneously, the generated deluge of jobs causes significant challenges for resource management and job scheduling systems. Indeed, on a system as large as Blue Waters, submitting more than 2,000 to 3,000 jobs at a time causes longer than desired scheduler planning cycles. For a big portion of the workflow, that would only accommodate analysis of 100 genomes, while we were targeting 500 genomes for simultaneous processing. To solve this problem, we worked with the Blue Waters team to incorporate a job launcher, which serves as an MPI wrapper around our OpenMP software. This adaption has been fully tested and works well.

Data management: Variant calling is a big-data workflow. When used on 500 deeply-sequenced, whole genome samples, it can create hundreds of thousands of files that use over a petabyte of disk space (Figure 1). The vast majority of these data are generated in the form of small files. Thankfully, we found that the Lustre file system on Blue Waters placed files evenly across disks, preventing a problematic load imbalance. Handling large numbers of files can strain the metadata servers and result in uneven performance. We worked with the Blue Waters team to create a parallel packaging utility, similar to TAR that could efficiently bundle files, thus preventing such issues.

Re-sorting intermediary files is common in a number of places along the workflow. The fastest software for it is Novosort [8], which involves two phases: sorting data in individual fragments, then merging the sorted fragments to produce the final output file. Ideally, the algorithm keeps all the intermediary fragments in RAM. When the available RAM is insufficient, it will write the fragments to disk, which can create enough I/O to saturate the bandwidth of the network routers and object storage servers. How many genomes can be feasibly analyzed in a single batch, before these bottlenecks begin to affect performance? We conducted Novosort scalability studies on three filesystems of different sizes (Figure 2) and concluded that the Blue Waters scratch filesystem can handle up to 1,000 genomes analyzed in parallel, without detrimental consequences. This is twice the required throughput for daily genotyping of all newborns in Illinois. The team reported the resultant I/O activity at the incredible rate of 4 TB/sec, likely due to re-reading from the Lustre cache, which functioned to compensate for the lack of RAM on the compute nodes (64 GB instead of the desired 256 GB).

Another I/O bottleneck resides in the alignment step of the workflow, when the algorithm mapping the sequencing reads against the reference attempts to access the input data on disk. Doing so on hundreds of samples at once engages the entire filesystem, forcing the processes to compete for disk access. We were able to overcome this bottleneck by striping the input data (width 3) and staggering the analysis in groups of 100 genomes. Configured this way, a massively parallel analysis of 500 genomes (30X coverage) completes in about 30-35 hours (Figure 3).

WHY BLUE WATERS

Our target problem size, 500 genomes, would use half the nodes and 1/20th of the file system daily on Blue Waters. The sheer size of the system made it possible for us to even consider such a problem. Our work demonstrated that Blue Waters can sustain **very high throughput** of genomic analyses without much impact on other users, and will remain very attractive for large genomics projects. The Blue Waters support staff have been instrumental in helping us figure out and eliminate issues with computational performance.

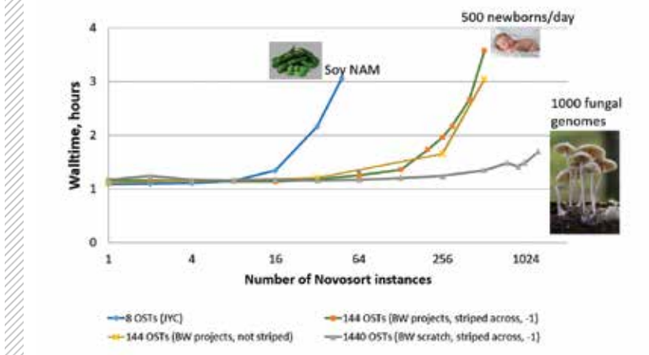

FIGURE 2: Scalability of Novosort on filesystems of different sizes. Images illustrate example problems feasible on each filesystem: 42 Soybean Nested Association Mapping (NAM) parents on the Blue Waters test cluster JYC, 500 newborns per day on /projects and up to 1,000 samples on /scratch (e.g. the 1,000 fungal genomes project). Legend lists the filesystem size as the number of object storage targets (OSTs).

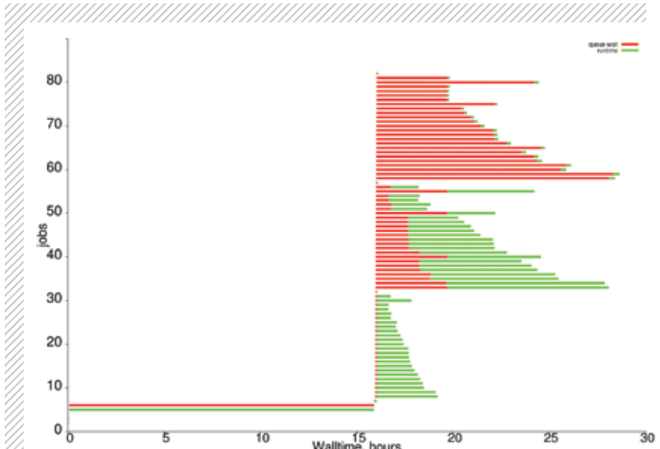

FIGURE 3: A Gantt chart displaying job execution in a single 100-genome batch. Each bar is a job. Red denotes waiting in the queue due to a job dependency, and green denotes the time of active execution. Five such batches staggered within a few minutes of each other permit 500 genomes (30X) to run in ~30-35 hours.

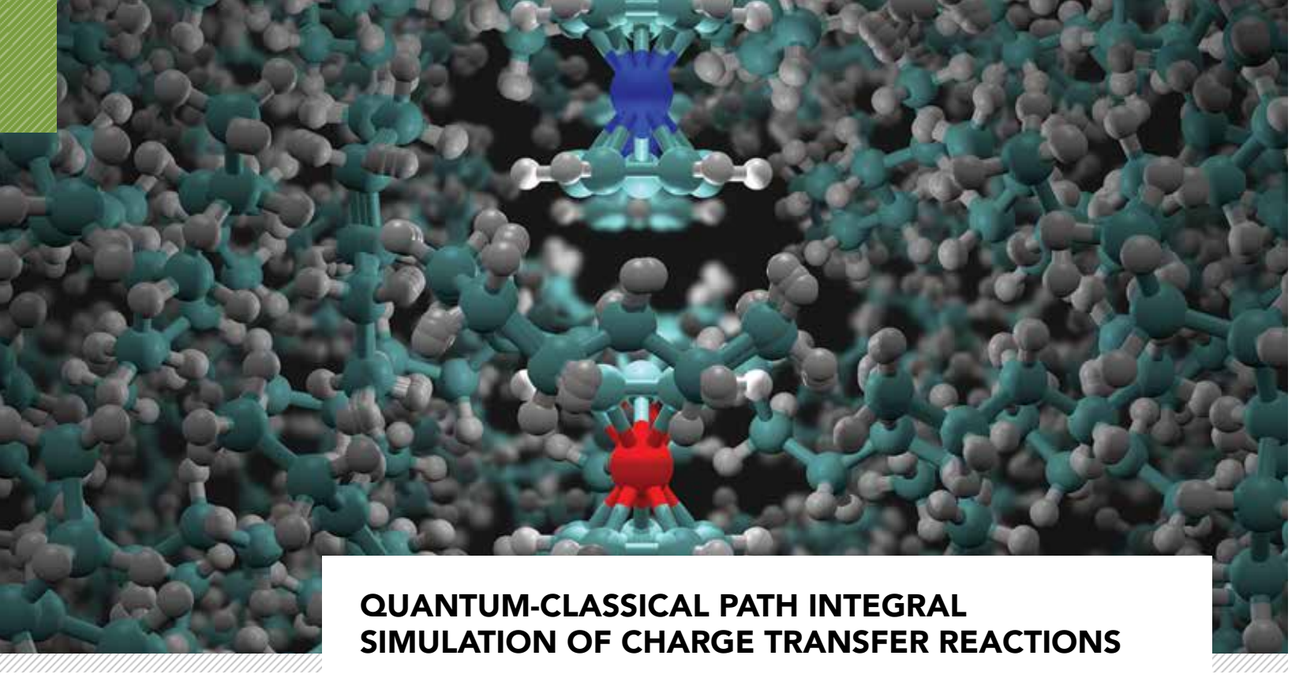

FIGURE 1: A snapshot of the simulation after approximately 0.4 ps showing the solvent configuration arising from the superposition of three quantum-classical paths.

QUANTUM-CLASSICAL PATH INTEGRAL SIMULATION OF CHARGE TRANSFER REACTIONS

Allocation: Illinois/113 Knh
PI: Nancy Makri[1]
Co-PI: Peter L. Walters[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Quantum mechanical effects play an essential role in chemical and biological processes. However, it is impossible, as well as unnecessary, to treat every particle involved in these processes quantum mechanically. Our group has recently developed a rigorous quantum-classical path integral (QCPI) methodology [1-4] in which a few of the particles are treated by full quantum mechanics, and the effects of the remaining particles are captured via classical trajectories. The QCPI approach is free of *ad hoc* assumptions, allowing a faithful description of the interaction between the quantum and classical particles. Its implementation on Blue Waters provides a detailed picture of charge transfer in solution with **unprecedented** accuracy.

INTRODUCTION

Quantum mechanical effects are prominent in many chemical and biological processes, yet they present a major challenge to simulation. Quantum mechanics is non-local in nature, which leads to an exponential scaling of the computational cost with the number of interacting particles. Thus, the computational effort required for simulating these processes, which contain many thousands of particles, with quantum mechanics is astronomically large.

For many processes of interest, quantum mechanical effects are only vital in the treatment of a small number of particles (e.g., those corresponding to a transferring charge). The remaining particles (solvent molecules or biological medium) can be adequately described via Newtonian dynamics. While this partitioning makes the simulations computationally feasible, the interaction between the classical and quantum partitions must be treated with care. The traditional Schrödinger formulation of quantum mechanics (which is based on delocalized wave functions) is incompatible with Newtonian trajectories (which are local in space). In the Schrödinger formulation, the interaction between the quantum and classical partitions cannot be obtained without introducing severe approximations.

METHODS & RESULTS

The Makri group's aim has been the development of rigorous quantum-classical formulations based on Feynman's path integral formulation of quantum mechanics [5]. The trajectory-like nature of the Feynman paths leads naturally to combined quantum-classical treatments. Unlike wave functions, Feynman paths are local in space. Thus these combined treatments are free of approximations. Recent work has described a QCPI methodology, which incorporates these ideas as well as several advances in the understanding of decoherence (loss of coherence) processes [6]. QCPI treats a small subsystem by full quantum mechanics, while the effects of the environment are captured via standard molecular dynamics (MD) procedures. Two widely used MD packages, NAMD and LAMMPS, are used to yield trajectories subject to forces obtained using the coordinates of the charged particle, which are specified by the given path.

The current project involves the **first** implementation of QCPI to the simulation of charge transfer in condensed media, which is a key feature of many chemical and biological processes. Using traditional quantum methods to simulate the dynamics accurately in such cases is practically impossible. However, with QCPI, these simulations become feasible. Our recent work has focused on two prototypical charge transfer reactions. The first involves the transfer of an electron between the ferrocene-ferrocenium charge transfer pair in solution (hexane, hexene and various mixtures of the two). The second is the proton transfer reaction for the phenol-amine complex in methyl chloride [7]. We have accurately simulated the dynamics of the charge transfer and obtained the time evolution of the state populations for the ferrocene-ferrocenium pair in hexane, and for the phenol-amine complex in methyl chloride. These were carried out at the full atomistic level and represented the **first** of their kind.

Additionally, we performed the same simulations with the solvent replaced by an effective bath of harmonic oscillators. This was done to help quantify the validity of linear response for such systems [8].

The QCPI simulations shed light on the complex interplay between the classical solvent and the quantum system at an atomistic level. Figure 1 shows the delocalizing effect that the quantum electron transfer has on the classical hexane solvent. It is apparent that even for solvents that are weakly coupled to the charge transfer, such as hexane, the system's influence on the solvent cannot be ignored.

Upon completion of the present phase of this work, we will obtain quantitative results for the dynamics of the ferrocene-ferrocenium charge transfer in hexene and in hexane-hexene mixtures. This will allow for a greater understanding of how stronger system-solvent coupling impacts the dynamics of electron transfer in solution. Additionally, we will be able to further test the validity of linear response in these less hospitable regimes.

WHY BLUE WATERS

The QCPI formulation is well suited to a decomposition based on multi-level parallelism, and Blue Waters provides the ideal platform for its implementation. Specifically, the set of system paths is distributed across nodes; one processor within each node is assigned to the quantum mechanical calculations, while the other performs supporting trajectory computations. Moreover, because the trajectories are independent and generally relatively short, it is possible to assign a single trajectory to each core within a given processor while maintaining computational efficiency. This multi-level approach has the benefit of minimizing communication time while maximizing concurrent processing since related classical and quantum-mechanical calculations are performed within the same node, where inter-processor communication should be much faster than if the information were more widely distributed.

NEXT GENERATION WORK

With the power of next generation of supercomputers it should be possible to use QCPI to accurately simulate the dynamics of complex charge transfer reactions containing multiple quantum particles and possessing long lived quantum coherence.

PUBLICATIONS AND DATA SETS

Walters, P. L., and N. Makri, Quantum-classical path integral simulations of ferrocene-ferrocenium charge transfer in liquid hexane. *J. Phys. Chem. Lett.*, 6 (2015), 4959–4965, doi:101021/acs.jpclett.5b02265

Walters, P. L., and N. Makri, Quantum-classical path integral simulations of ferrocene-ferrocenium charge transfer in solution. *251th ACS National Meeting & Exposition*, San Diego, Calif., March 13-17, 2016.

ALLOSTERIC SELECTIVITY AND DRUG BINDING PATHWAY OF μ-OPIOID RECEPTORS

Allocation: NSF PRAC/2.45 Mnh
PI: Vijay Pande[1]

[1]Stanford University

EXECUTIVE SUMMARY

We used Blue Waters' powerful graphics processing units (GPU) and central processing units (CPU) to find the allosteric binding sites of the μ-Opioid Receptor (μOR) and also the path opiates take to bind to the orthosteric site. Many important analgesics relieve pain by binding to μOR, and are therefore from a clinical perspective among the most important classes of G Protein Coupled Receptors (GPCRs). The mechanism of opiate binding and the selectivity of μOR are largely unknown. In this study, we discovered the allosteric site responsible for the attraction and selection of opiates. Using Markov State Models, we unveiled the pathway of opiates in binding to the orthosteric site. Our results have important implications for designing **novel** analgesics.

INTRODUCTION

The most powerful analgesic and addictive properties of opiates are mediated by the μOR. Since this receptor is primarily responsible for the effects of opium, the μOR is one of the oldest drug targets for the discovery of analgesics [1]. The μOR activation results in signaling through the heterotrimeric G protein Gi, resulting in analgesia and sedation. The activity studies of μOR have revealed that subtle changes in ligand structure can convert an agonist into an antagonist, so there is a general philosophy within the GPCR drugs where distinct pharmacophores are responsible for efficacy (message) or selectivity (address) [1]. In spite of the great studies in the past few years on the conformational changes triggered due to drug binding, the origin of the selectivity and how a certain receptor is selective to specific drugs are largely unknown. Does this selectivity originate from the binding pocket or are there other significant selectivity sites that prescreen the drug? In this regard, another fundamental question is the binding dynamics and the path that the drug travels from the extracellular to the binding pocket. We still do not know how drug diffuses through a highly tortuous cavity to arrive at the binding pocket.

To address the above questions, we used molecular dynamic (MD) simulations and the state of the art post-processing software MSMBuilder and machine learning algorithm tICA (Time-Structure Based Independent Component Analysis) developed in the Pande Lab. First, we found an allosteric site responsible for the selectivity of μOR through monitoring binding affinity of different ligands. Second, we tried to understand the pathway from the allosteric to orthosteric site. Since ligands need to undertake specific orientations to permeate through the complex and highly-energetic barrier cavity of the receptor, conventional long-trajectory MD simulations starting from allosteric site can't help us to unravel the binding pathway.

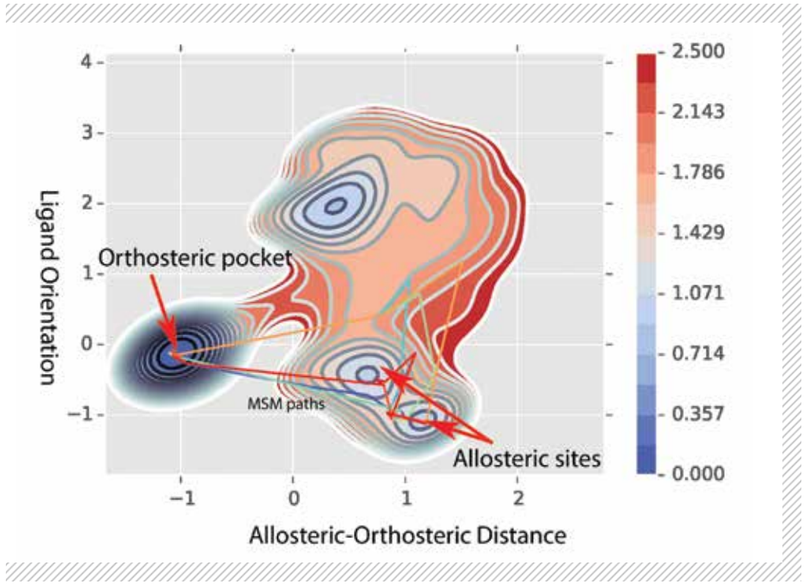

FIGURE 1: Weighted free energy and MSM path for the ligand to move from allosteric to orthosteric site. Ligand takes specific orientation to reach the orthosteric site. There are two possible orientations for the ligand at allosteric site.

METHODS & RESULTS

To tackle the challenges mentioned in the introduction, we randomly generated 560 initial ligand positions and orientations as the starting points of the simulations. Each simulation was run for 200 ns to allow enough time for the relaxation of the ligand to equilibrium position/orientation, as well as to traverse metastable potential wells. These seeds were created in, as well as between, the allosteric and orthosteric sites. We analyzed the trajectories using tICA to find the most important reaction coordinates (Fig. 1). MSMs were also built on these trajectories to find the most populated states and their connectivity (Fig. 2).

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 like Folding@home. Also, the availability of NAMD simulation package on Blue Waters has particular advantages for adaptive sampling and restrained equilibrations.

NEXT GENERATION WORK

Our future work will focus on the activation/deactivation mechanisms of μOR. Since the nature of conformational changes in the activation pathway of these receptors is very subtle, and delicate and encompasses microsecond to millisecond timescales, we will take advantage of Blue Waters to run thousands of MD simulations to be able to shed light into this drug-attractive receptor. Our ultimate goal is to create a unified model of receptor activation, from μOR initial interaction with opiates to receptor activation to Gi protein interaction.

PUBLICATIONS AND DATA SETS

Dodani, S. C., et al., Discovery of a regioselectivity switch in nitrating P450s guided by molecular dynamics simulations and Markov models, *Nature Chem.*, 8:5 (2016), pp. 419-425.

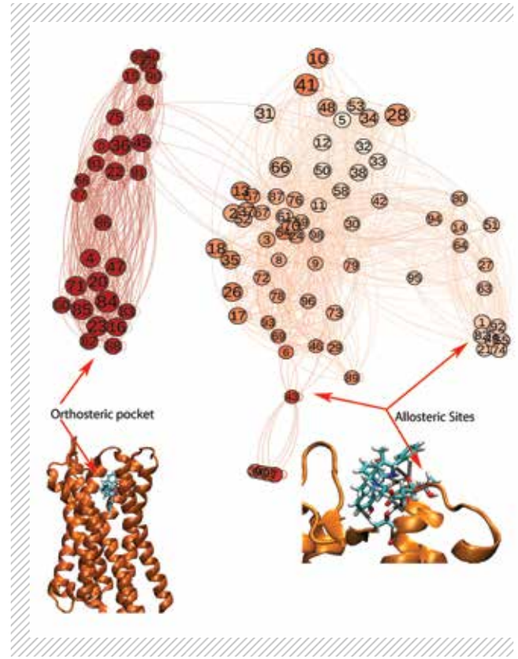

FIGURE 2: Network of MSM states depicting the allosteric and orthosteric sites. The larger spheres represent more populated states. The corresponding ligand position and orientation with respect to the receptor is demonstrated.

COMPUTATIONAL INVESTIGATION OF DROUGHT-RESISTANCE IN PLANTS

Allocation: Illinois/300 Knh
PI: Diwakar Shukla[1]
Collaborator: S. C. Huber[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

With the Earth's rising population and changes in the global climate, the biggest challenge facing humanity will be meeting future food and energy needs. While photosynthetic plants are our principal source of food and biofuels, we still know little about how plants adapt to environmental stresses. Plant kinases and phosphatases have been identified as the key signaling enzymes involved in regulation of photosynthetic efficiency and response to external stresses, but the molecular understanding of these stress and energy signaling enzymes remains elusive. Using Blue Waters, we investigated the conformational dynamics of plant receptor-like kinases involved in Brassinosteroid signaling and conducted a methodological study using evolutionary information to guide protein simulations. We have identified several key plant kinases whose activity can potentially be regulated via conformational engineering of αC-helix. These kinases are involved in key processes related to plant growth and development, such as regulating nutrient transport and drought-tolerance.

INTRODUCTION

Increase in demand for our primary foodstuffs is outstripping increase in yields, an expanding gap that indicates large potential food shortages by mid 21st century [1, 2]. This comes at a time when yield improvements are slowing or stagnating as the approaches of the Green Revolution reach their biological limits [3]. With the threat of global climate change and frequent occurrence of extreme weather events such as droughts, the task of producing sufficient food and biofuels is expected to become even more challenging in the future [4]. Plants respond to changing environmental conditions by translating extracellular signals (typically in the form of small molecules such as hormones) into appropriate intracellular responses. Cell-surface receptor-like kinases (RLKs) play a major role in extracellular sensing and transmitting the information into the cytoplasm for downstream signaling [5, 6]. In plants, the receptor-like kinases play key roles in regulating growth and development, protection against pathogens, and reproductive success in generating seeds and fruits and hindering premature abscission. Therefore, a quantitative molecular-level understanding of these plant-signaling processes is fundamental to future food and energy security. We focus on one of the best-characterized Leucine-rich repeat (LRR) RLK in plants, Brassinosteroid-insensitive--1 (BRI1), the receptor for plant steroid hormones called Brassinosteroids, which are crucial for plant growth.

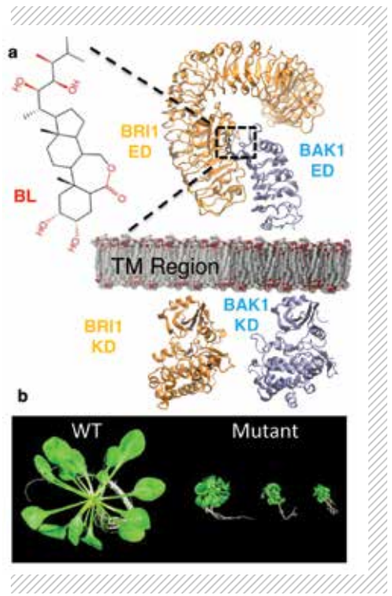

FIGURE 1: a) Crystal structures of the extracellular domains (ED) bound to brassinolide (BL) and the cytoplasmic kinase domains (KD) of AtBRI1 and AtBAK1 RLKs. b) Arabidopsis thaliana wild-type (WT) and dwarf plants from the BL deficient mutants.

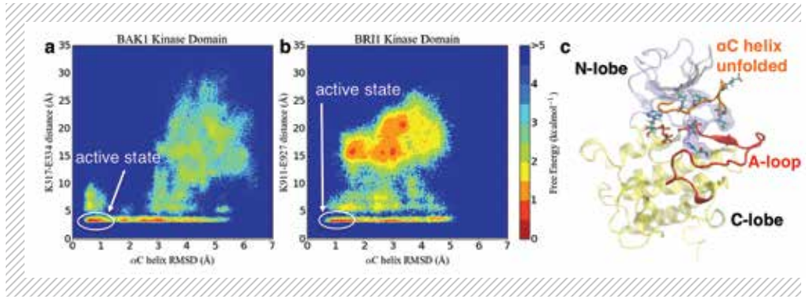

FIGURE 2: Free energy landscape of fully phosphorylated, ATP-bound kinase domain as a function of αC helix RMSD from the crystal structures and K317 side-chain amine nitrogen to E334 side-chain carboxyl carbon distance for a) BAK1 b) BRI1 kinase. c) Structure of BRI1 kinase domain with a αC-helix in unfolded state.

METHODS & RESULTS

The key aims of this project were to computationally investigate Brassinosteroid signaling from hormone reception to regulation of the BR activity using Blue Waters. Specifically, we have now performed extensive molecular simulations of the BAK1 and BRI1 receptor kinase domains starting from the available crystal structures. The simulation data was used to build Markov state models (MSMs) of kinase dynamics by clustering the structurally similar conformations into states and obtaining the interconversion rates between these states from the simulated trajectories. Such detailed thermodynamic (stability of a particular conformation of protein or a protein-protein complex) and kinetic (rate of interconversion among different conformations) information about the underlying conformational free energy landscape has provided new insights into the regulation of BAK1 and BRI1 kinase function.

The crystal structures of the fully phosphorylated BAK1 and BRI1 kinase domain show a folded αC-helix conformation. However, simulations reveal completely unfolded conformations of the αC-helix. The folded conformation of the αC-helix is a critical feature of the active state as the catalytically important K317-E334 h-bond is disrupted in the unfolded state. Fig. 2 shows regions with an unfolded αC-helix (high αC root mean square deviation (RMSD) from the folded structure) and a broken K-E h-bond. A simulation snapshot of the BRI1 kinase domain with an unfolded αC-helix is also shown (Fig. 2c). These results indicate the presence of catalytically incompetent conformations in the ensemble of kinase domain in its fully phosphorylated state. We have also performed CD experiments on BAK1 and BRI1 kinases to validate the computational prediction of the unstable αC-helix region. Finally, we have performed *in silico* mutagenesis to design BAK1 and BRI1 kinase domains with stable a αC-helix, thereby enhancing the catalytic activity of kinases. Our investigations reveal that Somatic Embryogenesis Receptor Kinase (SERK) family kinases show large differences in their αC-helix unfolding propensity, which could provide new avenues for conformational control of kinase activity of individual SERK-family kinases.

WHY BLUE WATERS

Slow conformational transitions are difficult to observe by running simulations on commodity hardware. Powerful resources like Blue Waters are required to study such complex biological processes in full atomistic detail and over long timescales. Blue Waters provides thousands of GPUs that are used for parallel molecular dynamics simulations to perform Markov state model-based adaptive sampling of conformational energy landscape of proteins. Blue Waters increases the overall compute performance by **several orders of magnitude** (in terms of the real time required for simulation).

PUBLICATIONS AND DATA SETS

Shukla, D. *et al.*, Conformational Heterogeneity of the Calmodulin Binding Interface. *Nature Communications* 7, 2016, doi:10.1038/NCOMMS10910

Shukla, S., Z. Shamsi, A. Moffett, B. Selvam and D. Shukla. Application of Hidden Markov Models in Biomolecular Simulations. *Methods in Molecular Biology*, In press, 2016.

Moffett, Z. Shamsi, and D. Shukla. Guiding functional dynamics and dimerization of proteins using evolutionary couplings. Submitted, *Nature Communications*, 2016.

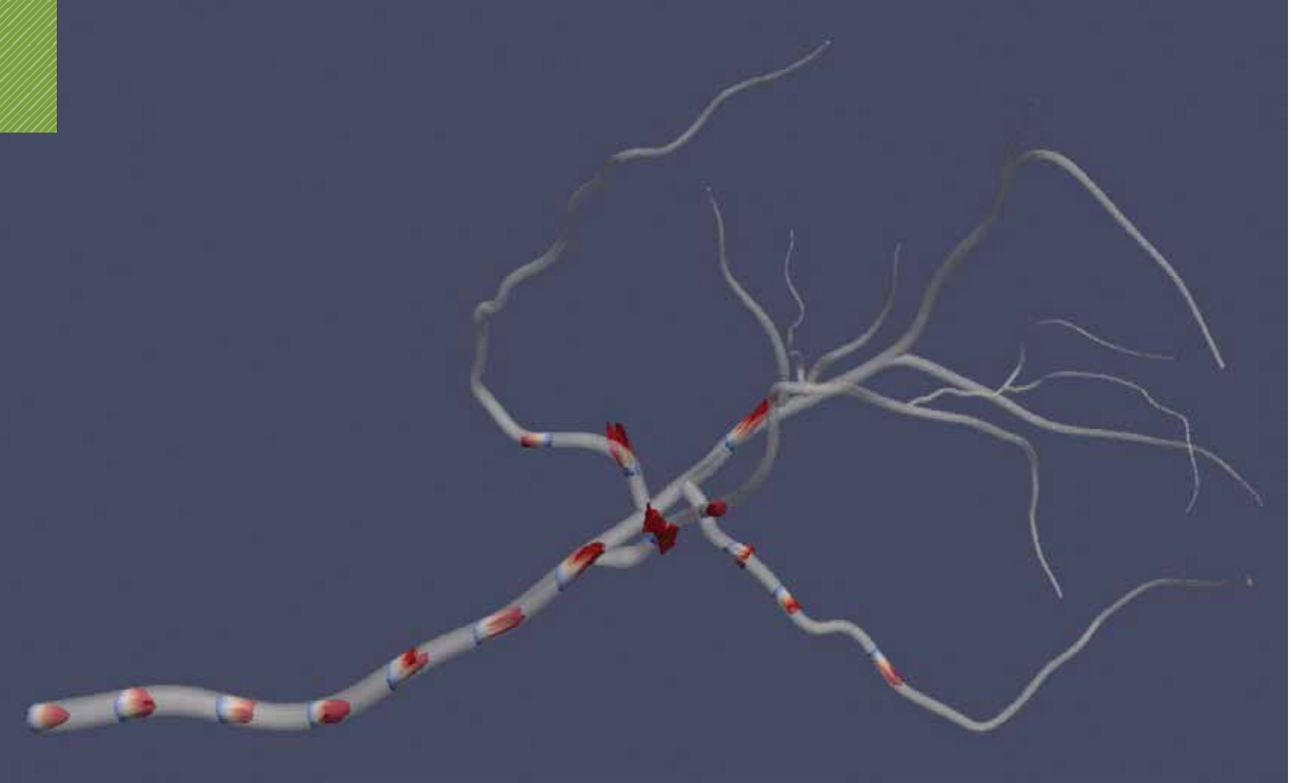

VARIATIONAL MULTISCALE METHOD FOR BLOOD FLOW SIMULATION IN PATIENT-SPECIFIC ARTERIAL GEOMETRIES

FIGURE 3: Instantaneous flow field in the arterial tree.

Allocation: Illinois/200 Knh
PI: Arif Masud[1]
Co-PI: JaeHyuk Kwack[1]
Collaborators: Soonpil Kang[1] and Lixing Zhu[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

In this project, we extended and applied our variational multiscale method [1-7] for non-Newtonian blood flow modeling to account for material and geometrical uncertainty in patient-specific models shown (Figs. 1, 2). Blood is modeled as a non-Newtonian shear-rate dependent viscous fluid, and artery walls are modeled as nonlinear viscoelastic material with relaxation time of the arterial tissue.

Significant results from these studies are: (i) A correlation is found between high G forces and local reduction of pressure and flow rate to the **brain**, an issue of considerable concern in high speed roller coasters as well as in gravity-induced loss of consciousness (G-LOC) which is a major threat to high-performance aircraft pilots and astronauts. (ii) Standing waves are found in continuous flow ventricular assist devices (VADS), where pulsatility gets minimized, and consequently, the risk factor for blood coagulation and therefore of stroke increases.

INTRODUCTION

We have developed **novel** numerical methods that are integrated with non-Newtonian constitutive models to simulate and analyze blood-artery interaction in patient-specific geometries in the cardiovascular system (Figs 1, 2). Since blood artery interaction models are mathematically involved and computationally expensive, Blue Waters was employed to explore the mathematical attributes of our new coupling scheme. Our target problem was predicting the probability of stoke in patients who have received surgical treatments for VAD insertion. One risk factor is the hemodynamic conditions, and reduced pulsatility is considered to be a major contributor to it. Local states of flow in the arterial system are also important. For example, an aneurysm in the carotid artery can provide a local spot for blood to flow slow, increasing the risk for coagulation. Consequently, accurate representation of the geometry is important, and any information related to local changes in artery diameters needs to be accounted for. Another allied investigation led us to investigate the effect of high G forces on reduced perfusion of the brain and potential loss of consciousness.

Using our codes, we investigated the effects of the stiffening of arterial walls that can occur as a function of age on the progression of hypertension or high blood pressure—a problem of great relevance as uncontrolled high blood pressure increases the risk of heart attack and stroke. We were able to carry out a set of parametric studies with larger arterial systems as shown in Figures 1 and 2. These studies helped us identify standing waves in continuous flow VADS, and we found optimal pulsatility that can disrupt the standing wave without inducing heavy mechanical beating on the flow pump. Also, to develop methods that can focus on smaller subsections of the arterial model, outflow boundary conditions for non-Newtonian models for blood were developed. This development required numerous high fidelity computations and comparison with clinical data. A third allied investigation led us to the modeling of

FIGURE 1: Side view of the patient-specific geometric model of arterial system in the brain. Also shown is one of the cruder meshes that were employed in the numerical simulations.

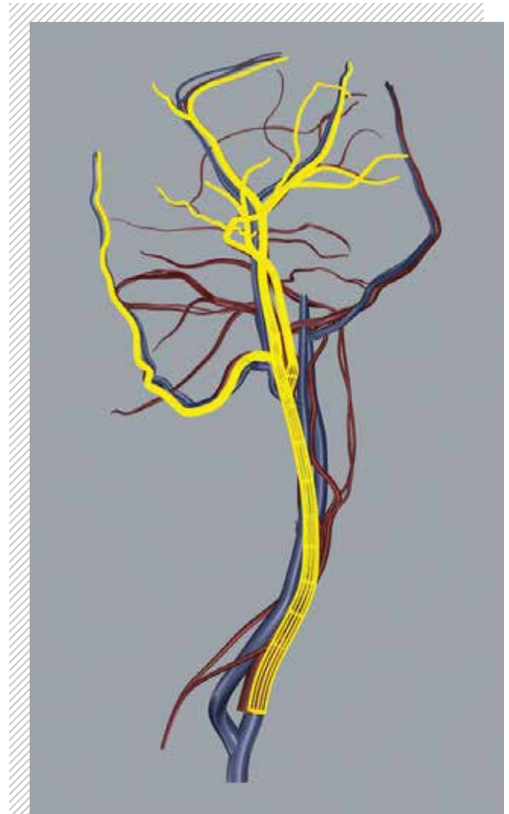

METHODS & RESULTS

Our previous work on Blue Waters allowed us to extend and verify our non-Newtonian models for blood that account for shear-rate response in patient-specific geometries. We developed hierarchical multiscale finite-element methods with local and global (coarse and fine) description of the variational formulations that result in telescopic depth in scales. The telescopic depth in scale leads to two coupled nonlinear systems, namely, the coarse-scale and the fine-scale subsystems. Fine-scale models that were extracted from the residual-driven fine-scale sub-problems were then variationally embedded in the coarse-scale description of the problem, and led to the class of hierarchical multiscale methods for non-Newtonian fluids with enhanced stabilization properties.

Our newly developed methods employed here led to substantially reduced global communications in favor of increased local computing. With a five percent increase in the cost of computation of the stiffness matrix and the residual force vector, we were able to reduce the mesh size to less than half the nodes that would otherwise be needed for equivalent engineering accuracy, thereby substantially reducing the overall cost of computation of the problem. This unique feature of our method is of tremendous benefit in massively parallel computing as it reduces communication costs across the partitioned subdomains.

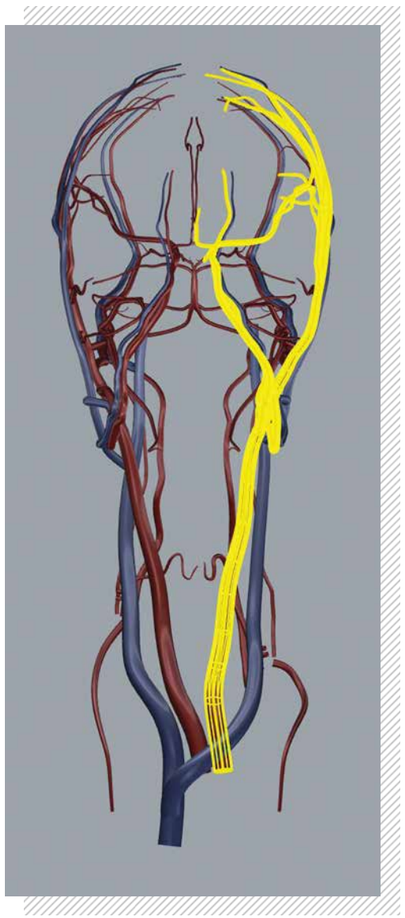

FIGURE 2: Frontal view of the patient specific geometric model of arterial system in the brain along with the superposed computational grid.

the high G forces on reduced perfusion of the brain and potential G-LOC. To summarize, following are the significant findings of our work:

1. We show that the Newtonian model for blood underestimates the wall shear-stress (WSS) by 16% in diastole, which is a non-conservative estimate from patient care perspective.

2. We show viscosity build-up in aneurisms in arteries where local ballooning effect causes a recirculation region, and viscosity build-up is significant at the end of diastole which can lead to clot formation that is caused by local stagnation.

3. We show that constant flow VADS with minimal pulsatility lead to standing waves that cause viscosity build-up which can increase the risk factor for stroke in the patients. This aspect of simulation can be used to introduce minimal pulsatility in the VAD device, thus optimizing the device for the patient, giving rise to patient-specific treatments.

4. We can calculate increased pressure fields in stiffened arteries, even when other geometric and flow conditions are held constant, thus providing insight into the effects of mechanical and rheological factors on hypertension.

5. High-performance aircraft pilots are routinely exposed to high levels of +Gz (head-to-foot) accelerations. G-LOC is a major threat to high performance aircraft pilots and astronauts. With recent advances in MRI technology and computing power, cerebral hemodynamic simulations can now be run using individual arterial geometries, and these can open the door to understanding G-LOC from an intracranial perspective.

WHY BLUE WATERS

This work relies on having many long simulations to achieve rigorous sampling of the variability in biological systems. Blue Waters was critical for both the development of **cutting-edge** software and the application of this software to perform large-scale biomechanics simulations. We found Blue Waters to be an extremely powerful and versatile computational resource that, in addition to powerful CPU and GPU hardware, provided fast interconnects that allowed us to do types of calculations that we could not have done on other platforms. Specifically, the large local memory of Blue Waters is ideally suited for our methods as we can exploit the resident memory on the processing nodes to make the macro elements "smart," reducing the size of the global problem and minimizing data communication. The Blue Waters project staff provided in-depth technical information and timely advice on the optimal deployment and performance tuning of our software.

NEXT GENERATION WORK

In the next Track-1 system we plan to extend and embed the method in a probabilistic framework for blood flow simulation in patient-specific arterial geometries, with the objective of optimization of VADs for patient-specific needs.

PUBLICATIONS AND DATA SETS

Weddell, J.C, J. Kwack, P.I. Imoukhuede, and A. Masud, Hemodynamic analysis in an idealized artery tree: Differences in wall shear stress between Newtonian and non-Newtonian blood models. *PLoS ONE*, 10(4): e0124575, 2015.

Kwack, J., S. Kang, G. Bhatt, and A. Masud, Outflow boundary conditions for non-Newtonian models for blood, (eds. K. Takizawa and Y. Bazilevs), Modeling and Simulation in Science, Engineering and Technology Book Series, Springer, 2015.

Kwack, J., A. Masud and K.R. Rajagopal, Stabilized Mixed Three-field Formulation for a Generalized Incompressible Oldroyd-B Models, *Int. J. Num. Meth. Fl.*, (2016), doi: 10.1002/fld.4287

IMPROVING THE RESOLUTION OF BRAIN BLOOD FLOW IMAGING WITH ADVANCED MRI ACQUISITIONS AND COMPUTATION

Allocation: Illinois/26.0 Knh
PI: Brad Sutton[1]
Collaborators: Alex Cerjanic[1], Joseph Holtrop[1], and Giang Chau Ngo[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Adequate blood flow in the brain is critical for supporting healthy brain function into old age. Imaging of the brain's blood flow and obtaining information about the age-related changes in the structure of blood vessels, noninvasively using magnetic resonance (MR) imaging, requires advanced approaches for measuring this small signal. The small volume fraction of blood in the brain allows for physics-based image reconstruction models to be used to improve quality and usefulness of the images. These reconstruction models can be computationally demanding, and hundreds of images may need to be acquired to estimate information about the blood flow and vessels. As part of this project, PowerGrid, a toolkit for accelerating MR image reconstructions using graphics processing units (GPUs) and distributed computing, was created in C++ and OpenACC, while also leveraging message passing interface (MPI) to distribute across multiple GPUs. Using PowerGrid, we can reconstruct full datasets of images from a patient in a time frame similar to the acquisition of the images.

INTRODUCTION

Sufficient and reactive blood flow in the brain is a critical component for the health of neurons and their supporting cells. However, advanced aging is accompanied by critical changes to the vasculature [1], including the microvasculature that is involved in exchanging nutrients and waste with tissues. Measuring changes and degradations in the microvascular architecture of the human brain is

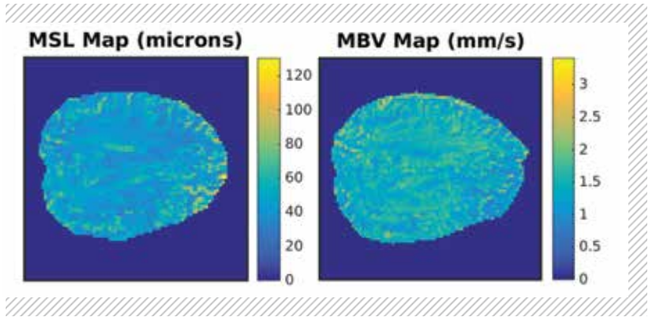

FIGURE 1: (left) Microvascular segment length (MSL) for a young subject in a slice high in the brain. (right) Microvascular blood velocity (MBV) for the same subject on the same slice. Resolution: 2mm isotropic, Scan time is 24:00 minutes total for both required scans.

limited to postmortem samples by the destructive nature of microscopy and histology used to image human brain tissue [2, 3]. Relying on destructive measurements of microvascular parameters prevents studies from assessing cognitive function of subjects, microvascular function, and observing changes longitudinally. New noninvasive, neuroimaging-based biomarkers for the state and function of the microvasculature in the brain are needed.

MR imaging provides an excellent neuroimaging platform to develop novel biomarkers for microvascular changes from age and pathology due to the flexible contrast mechanism, minimal risk profile, and broad clinical availability and applicability. Diffusion-weighted MR (DW-MR) imaging, MR imaging deriving sensitivity from the microscopic motion of water molecules, is an ideal platform for the development of novel biomarkers of blood flow in the microvasculature. The use of DW-MR imaging to characterize blood flow has been proposed with the Intravoxel Incoherent Motion (IVIM) method as introduced by Le Bihan[4]. While DW-MR imaging is exquisitely sensitive to blood motion in the microvasculature of the brain, it is difficult in practice due to low signal to noise ratio and challenges in sampling blood motion.

METHODS & RESULTS

Our team developed PowerGrid, a toolkit for accelerating iterative, model-based MR image reconstructions using GPUs and distributed memory computing. Implemented in C++, PowerGrid allows researchers and clinicians to retain the familiar structure arising from years of work developing advanced MR image reconstruction algorithms in MATLAB while leveraging high-performance computing (HPC) resources, such as Blue Waters and OpenACC.

The object-oriented structure of PowerGrid was designed to combine **state-of-the-art** image reconstruction in MR with MPI and OpenACC. This approach was necessary, as MR physicists traditionally do the majority of software development for image reconstructions inside a high-level interactive language like MATLAB.

Our use of MPI is enabled both by the code written in PowerGrid and the exploitation of natural parallelism that exists in the parallel, multi-coil imaging of modern MR scanners. These separate streams of data from each receiver coil provide a natural work unit that can be assigned to a GPU. The resulting parallelism is not complete, requiring the use of MPI communication routines inside of each global iteration of image reconstruction. Preliminary results enabled by Blue Waters and our PowerGrid project are demonstrated in the form of microvascular blood velocity (MBV), and microvessel length (MSL) maps for a single slice high in the brain of a young male volunteer are shown in Figure 1.

WHY BLUE WATERS

Blue Waters and NCSA provided excellent project support for moving our MR image reconstructions from desktop class machines to an HPC environment. Through a Blue Waters sponsored OpenACC Hackathon, our team, with hands-on assistance from two members of Blue Waters project staff, used OpenACC to accelerate the core routines of our image reconstruction utility as part of our PowerGrid project. Using PowerGrid and the Blue Waters multiple GPU support, we have shown speed up factors of up to ~11 times above the single GPU case through the use of MPI for distributed computing as shown in figure 2 for a small benchmark data set. We anticipate increased speed up factors as the dataset size and complexity increases.

Furthermore, the scale of Blue Waters supports exploring the inherent parallelism that exists in MR imaging. For each blood flow dataset, we need to reconstruct multiple 3D images from the 10-20 measurements that vary the diffusion weighting and are received from 32 parallel receiver coils [5]. This creates a large amount of data that is well-suited to parallel implementation across GPU-equipped nodes equipped with high speed interconnect. Also, the OpenACC support present in the Cray Programming Environment makes Blue Waters ideal for this work.

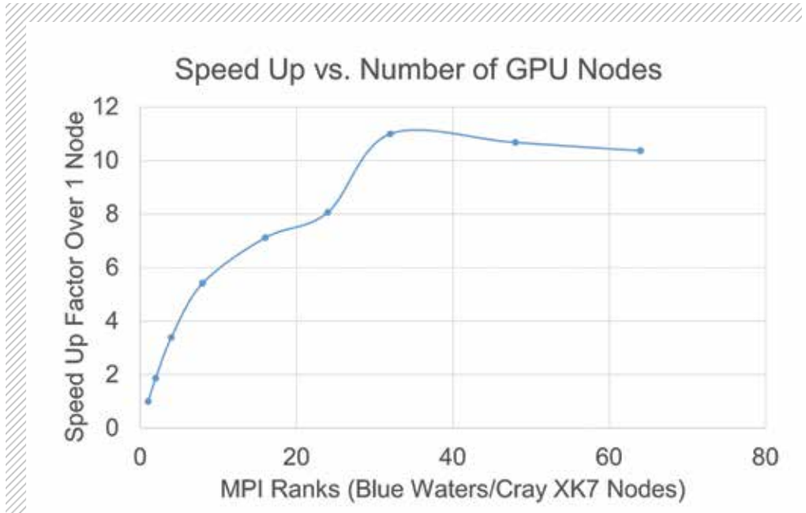

FIGURE 2: Results showing speed up versus number of MPI ranks (K20x GPUs) on Blue Waters showing peak speed up of ~11x with 32 nodes and saturation for additional ranks for a small benchmark case distributed with PowerGrid. This benchmark case represents 1/300th of a complete full brain dataset for blood flow imaging.

NEXT GENERATION WORK

Currently, it is common for MR acquisitions to use techniques to separate the encoding of a full 3D volume into smaller pieces to reduce the size of the smallest interdependent work unit or imaging volume[6]. However, recent advances in MR hardware have introduced 64 channel receivers, enabling imaging of larger volumes in less time with new data sampling patterns[7]. However, these new sampling patterns, such as a full brain 3D encoding strategy with 64-channel head coils would increase our smallest work unit by a factor of 30, presenting new demands on hardware accelerators and communication hardware.

Moving beyond larger imaging volumes, we have worked with strategies to exploit information across time or contrast weightings using 4D or higher model-based reconstructions [8, 9]. Applying these models to blood flow would easily increase the smallest work unit by a further factor of 15 – 20 for a predicted increase of 600 times the existing work unit, requiring 9,600 K20x class GPUs to reconstruct images in the same amount of time as it takes to acquire the data with a clinical MRI scanner. A next-generation Track-1 system will provide a testbed for future advanced reconstruction strategies for blood flow imaging in the human brain at high resolution at a scale not currently possible with Blue Waters.

PUBLICATIONS AND DATA SETS

Cerjanic, A., et al., PowerGrid: A open source library for accelerated iterative magnetic resonance image reconstruction. *ISMRM Annual Meeting*, Singapore, May 7th-13th 2016.

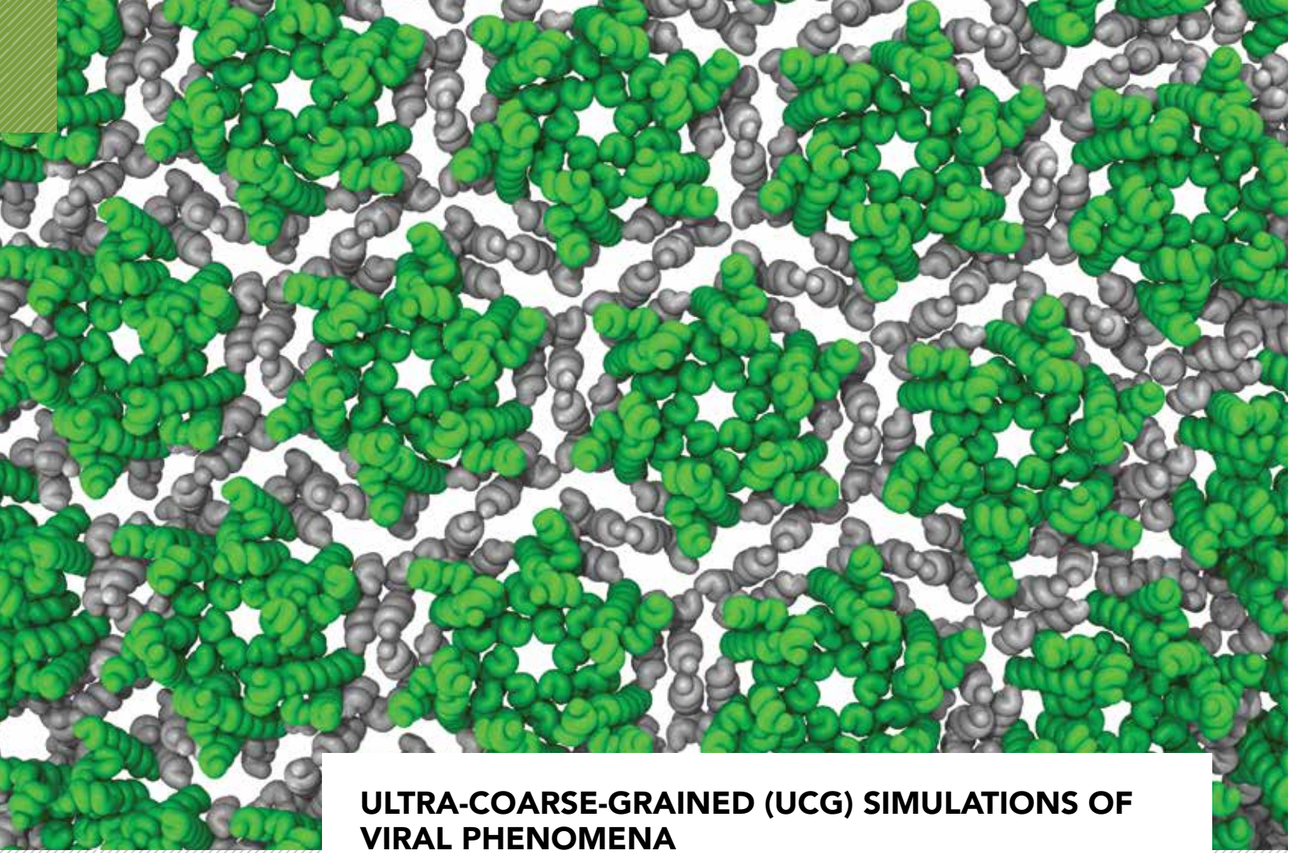

ULTRA-COARSE-GRAINED (UCG) SIMULATIONS OF VIRAL PHENOMENA

FIGURE 1: Detail of the mature capsid protein lattice of HIV-1. The N-terminal regions of each capsid protein molecule are shown in green to highlight the hexagonal lattice structure, with the C-terminal regions of each capsid protein in gray.

Allocation: NSF PRAC/2.60 Mnh
PI: Gregory A. Voth[1]
Co-PI: John M. A. Grime[1]
Collaborators: James. F. Dama[1], Barbie K. Ganser-Pornillos[2], Mark Yeager[2], Cora L. Woodward[3], and Grant J. Jensen[3]

[1]The University of Chicago
[2]The University of Virginia
[3]California Institute of Technology

EXECUTIVE SUMMARY

The maturation of HIV-1 viral particles is crucial for viral infectivity, and one important aspect of maturation is the generation of a viral capsid. The viral capsid is formed by the self-assembly of many copies of a specific capsid protein, generating a complicated lattice structure (Fig. 1) that encloses the genetic material of the virus inside the viral particle. While experimental data exists to illustrate the characteristic features of mature capsid lattice, specific details regarding the early-stage mechanisms of capsid nucleation and growth remain elusive. Using the Blue Waters, we employed "coarse-grained" (CG) models of the capsid protein to investigate the self-assembly of capsid lattice under various conditions, and also the disassembly of capsid lattice structures when a viral capsid is transferred into an uninfected cell - a process crucial for viral infectivity.

INTRODUCTION

The lifecycle of human immunodeficiency virus (HIV) requires a complicated series of morphological changes to convert the initially non-infectious "immature" viral particle into a mature and infectious form, a process referred to as maturation. During the maturation process, many large Gag protein molecules are cut apart to release smaller proteins with important contributions for viral infectivity. One such protein is the viral capsid protein (CA).

Many copies of the CA self-assemble into a cone-shaped "capsid" structure that encloses the viral RNA in a mature and infectious viral particle. Viral particles which fail to generate suitable capsid structures are non-infectious, and so the specific details of capsid formation are of interest for the design of therapeutic treatments. The inherent variability of HIV-1 viral particles makes controlled studies with conventional experiments difficult, and the extremely small length scales over which capsid assembly occurs likewise present barriers to the study of capsid self-assembly using current experimental techniques.

Computational simulations of the CA self-assembly process can, therefore, play an important role in elucidating the details of capsid self-assembly. The use of CG molecular models, where full atomic detail is replaced with simplified and efficient molecular representations, is of particular interest for relatively large-scale molecular processes such as viral capsid self-assembly.

METHODS & RESULTS

We deployed **novel** ultra-CG models [1] of the capsid protein, based directly on experimental data, to study the early stages of viral capsid nucleation and growth. The simulations were performed on a custom, highly-parallel molecular dynamics package developed specifically for this type of system using the Blue Waters platform [2]. The model indicates the importance of metastable triangular structures, formed from several capsid protein molecules, in the initial nucleation of the viral capsid: progressive addition of capsid proteins to an existing (but unstable) triangle can eventually produce a stable nucleating structure from which mature lattice growth proceeds. Our simulations suggest that this nucleating structure contains a central hexagon, surrounded by a ring of supporting triangles. Simulations were performed under various conditions to examine the influence on self-assembly of capsid protein concentration and molecular crowding.

Importantly, a protocol was also devised to examine the influence of structural flexibility in the capsid protein, with the innate flexibility of the molecule found to be critically important in controlling the self-assembly process. Previously self-assembled capsid protein structures were exposed to a process of rapid dilution under fixed molecular crowding, to model the viral capsid entry into an uninfected cell. The results of the rapid dilution studies indicate the instability of viral capsids which contain seams or holes in their surfaces, adding further data to our understanding of cellular defense mechanisms (such as tripartite motif (TRIM) protein restriction) which are involved in fighting viral infection. The data produced by our computer simulations using Blue Waters thus elucidate the critical early stages of viral capsid self-assembly, while also providing insight into the eventual fate of viral capsids as they break down at a key stage in the viral lifecycle.

WHY BLUE WATERS

The Blue Waters platform was a crucial element for the development of **cutting-edge** software algorithms and the application of these technologies to enable large-scale and dynamic biomolecular simulations. One key aspect of Blue Waters was the availability of the project staff, who provided invaluable technical insight and frequent advice to tune the performance of our software.

NEXT GENERATION WORK

For the next-generation Track-1 system, we hope to significantly extend the scope of our biomolecular simulations. Rather than performing molecular simulations of isolated aspects of the viral lifecycle, we intend to design and implement CG models capable of encompassing both the early and late stages of viral infection using essentially the same model. These models will involve an ultra-CG [1] representation of the Gag molecule which is capable of representing the aggregation of immature Gag lattice at cell membranes (and subsequent budding events) and the specific enzymatic cleavage of the Gag model during viral maturation. By applying the techniques pioneered by the UCG theoretical framework [1] and closely-related software tools [2], we thus hope to generate models of viral infection capable of rich analysis of medically relevant aspects of the viral lifecycle.

PUBLICATIONS AND DATA SETS

Grime, J., et al, Coarse-grained simulations reveals key features of HIV-1 capsid self-assembly. *z* 7 (2016) 11568. doi:10.1038/ncomms11568

ADVANCING GENOME-SCALE PHYLOGENOMIC ANALYSIS

Allocation: Illinois/158 Knh
PI: Tandy Warnow[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

The project has three main goals, each geared towards advancing the accuracy of large-scale estimation of evolutionary history. The first year has focused on the first goal of method development for multiple sequence alignment, maximum likelihood phylogeny estimation, and species tree estimation. Highlights include (a) scalable versions of BAli-Phy, a Bayesian method for statistical co-estimation of multiple sequence alignments and trees, (b) a new method (HIPPI) for classifying sequences into gene families, and (c) new supertree methods with improved accuracy and scalability. One journal paper has been published, two others have been submitted, three others are in preparation, and six projects are underway. Three graduate students and two postdoctoral fellows at University of Illinois at Urbana-Champaign participated in this project.

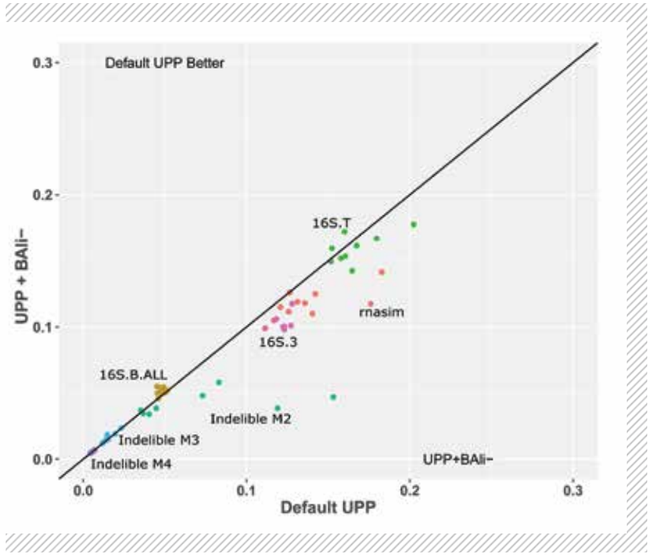

INTRODUCTION

Because phylogenies and multiple sequence alignments are the basis of many biological discoveries, improving the accuracy of methods that estimate these alignments and phylogenies will improve downstream biological inferences. These inferences range from the timing of different evolutionary events, to understanding which genes are involved in trait evolution, to how species adapt to their environments, to analyzing the human gut microbiome. All of these inferences depend on accurate gene trees (i.e., how genes evolve within the species tree) and species trees (how organisms diversified from a common ancestor), which in turn depend on accurate multiple sequence alignments. The estimation of alignments and trees is computationally difficult, as the best methods are attempts to solve NP-hard optimization problems or use Markov Chain Monte Carlo techniques. This project developed methods and software to enable highly accurate estimations of large-scale multiple sequence alignments and phylogenetic trees and thus advance scientific discovery.

METHODS & RESULTS

Our research focused on the development of methods with improved accuracy for multiple sequence alignment, large-scale maximum likelihood, species tree estimation from multiple genes, remote homology detection and protein family classification, and supertree estimation. We also used Blue Waters to perform large-scale phylogenomic and metagenomic analyses of biological datasets, in collaborations with research groups here at Illinois and around the country. These analyses gave us insight into how we can improve the methods regarding accuracy and computational performance. Many of these biological analyses involved estimating multiple sequence alignments using our new methods (PASTA [1] and UPP [2]) and estimating species trees from multi-locus datasets using maximum likelihood heuristics such as RAxML [3] or FastTree [4]. Some species tree analyses have also been based on ASTRAL-2 [5], a method we developed for estimating species trees in the presence of gene tree heterogeneity, and that is statistically consistent in the presence of incomplete lineage sorting. One of the major outcomes of this year was a study performed in collaboration with Illinois Professor Bryan White (Carle Woese Institute for Genomic Biology) and Mihai Pop (University of Maryland at College Park). Our recently published study [6] compared computational methods for defining operational taxonomic units. To perform this study, we computed a multiple sequence alignment of approximately 40,000,000 16S sequences, using UPP; this may be the largest multiple sequence alignment ever computed, and Blue Waters was essential for this study. Two Ph.D. students and one postdoctoral fellow worked on research supported by this Blue Waters allocation; the Ph.D. dissertations of the students will include research reported here.

The highlights of the method development were:

- Scalable versions of BAli-Phy [7], a Bayesian method for statistical co-estimation of multiple sequence alignments and phylogenetic trees, so that it can analyze datasets with 10,000 sequences. The improvement in scalability was obtained by integrating BAli-Phy into PASTA and UPP, two of our divide-and-conquer methods for estimating multiple sequence alignments. The existing BAli-Phy code is otherwise limited to, at most, 100 or so sequences.
- A new method (HIPPI) for classifying sequences (including short reads generated by sequencing technologies) into gene families. HIPPI is based on a new machine learning approach we have developed called "ensembles of Hidden Markov Models," which we have also used for metagenomic taxon identification and abundance profiling [8]. HIPPI improved accuracy compared to existing gene family classification methods based on BLAST [9] or single Hidden Markov Models [10].
- A new supertree method, FastRFS, which finds an optimal solution to an NP-hard optimization problem (Robinson-Foulds Supertrees) within a constrained search space. FastRFS provides improved accuracy and greatly reduced running times compared to other supertree methods.

WHY BLUE WATERS

Blue Waters was necessary for the development and testing of computationally intensive methods, which is not feasible on other platforms. Also, we were able to analyze several large-scale biological datasets using Blue Waters, which would not have been feasible using the other computational platforms that were available. Finally, the Blue Waters staff helped us port our codes and get them running, which was very helpful.

NEXT GENERATION WORK

The significant advances we are making are obtained by enabling statistical estimation methods, some employing computationally intensive MCMC techniques, to scale to large datasets. Should a next generation Track-1 system become available for us to use, we would be able to analyze the largest biological datasets coming online, including whole genome datasets with tens of thousands of species and metagenomic datasets with up to billions of reads. These analyses would transform biological and biomedical research.

PUBLICATIONS AND DATA SETS

Nguyen N., T. Warnow, M. Pop, B. White. A perspective on 16S rRNA operational taxonomic unit clustering using sequence similarity. *Nature Biofilms and Microbiome Analysis*, 2, article number 16004 (2016), doi:10.1038/npjbiofilms.2016.4

FIGURE 1 (LEFT): Scatterplot of multiple sequence alignment error rates of UPP-default compared to UPP+Bali-Phy on datasets with up to 10,000 sequences. We show alignment error of UPP run in default mode to UPP using BAli-Phy (written as UPP+BAli-Phy) on biological and simulated datasets, ranging in size from about 5000 to 10,000 sequences. On nearly all these datasets, using BAli-Phy within UPP resulted in reduced alignment error, showing that integrating BAli-Phy within UPP improved UPP's accuracy. Since Bali-Phy cannot run on these datasets due to computational limitations, this analysis also shows that integrating BAli-Phy within UPP improves BAli-Phy's scalability.

BIG DATA ON SMALL ORGANISMS: GENOME-SCALE MODELING AND PHENOTYPIC PREDICTION OF ESCHERICHIA COLI

Allocation: NSF PRAC/300 Knh
PI: Ilias Tagkopoulos[1]

[1]University of California-Davis

EXECUTIVE SUMMARY

We developed semi-supervised normalization pipelines and performed experimental characterization (growth, transcriptional, proteome) to create a consistent, quality-controlled multiomics compendium for *Escherichia coli* with cohesive metadata information. We then used this resource to train on Blue Waters a multi-scale model that integrates four omics layers to predict genome-wide concentrations and growth dynamics. Large scale simulations and subsequent validation led to several interesting results. First, the genetic and environmental ontology reconstructed from the omics data was found to be substantially different and complementary to the genetic and chemical ontologies. Second, the integration of all layers led to the predictor with the highest performance, although the increase in accuracy was incremental. We have validated 16 new predictions by genome-scale transcriptional profiling. This work constitutes the largest omics-based simulation and demonstrates how large high-performance computing infrastructure, big data and novel computational methods can lead to an integrative framework to guide biological discovery.

INTRODUCTION

Predicting microbial behaviors through data-driven computational tools is key to a) understand how microbial behaviors emerge and b) engineer the organisms in a fast, robust and accurate manner for biomedical and biotechnological applications. Realistic models of bacterial cells can lead to new paradigms related to bacterial physiology as well as the rational redesign of genomes with desired behavior. For this task, multi-scale modeling and visualization tools are of paramount importance as they provide insight of the systems under study.

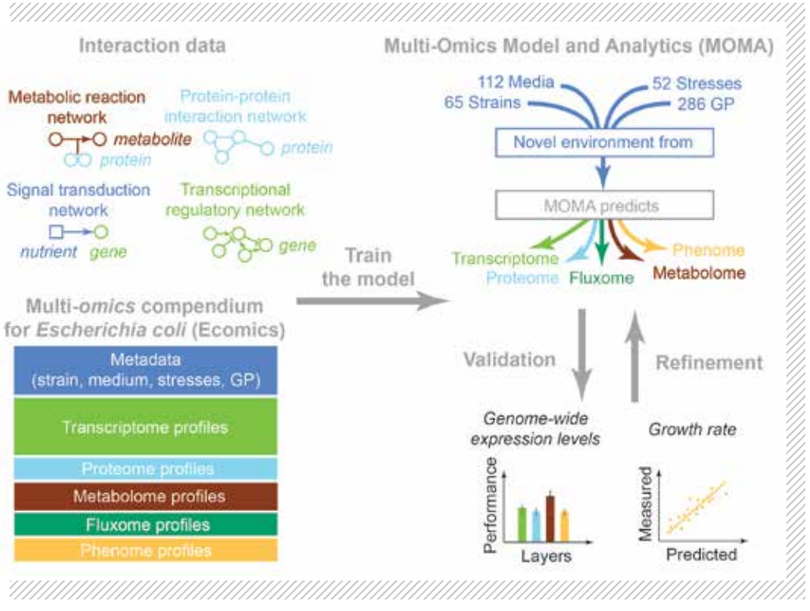

FIGURE 1: Dataset and Methodology

METHODS & RESULTS

This work had two parts, the generation of a consistent omics compendium and the development of the integrative model that uses it as a training set. As such, we first constructed Ecomics, a normalized, well-annotated, multi-omics database for *E. coli*, developed to provide high-quality data and associated meta-data for performing predictive analysis and training data-driven algorithms. This compendium houses 4,389 normalized expression profiles across 649 different conditions. We then proceeded to create the Multi-Omics Model and Analytics (MOMA) platform, an integrated model that learns from the Ecomics and other available network data to predict genome-wide expression and growth.

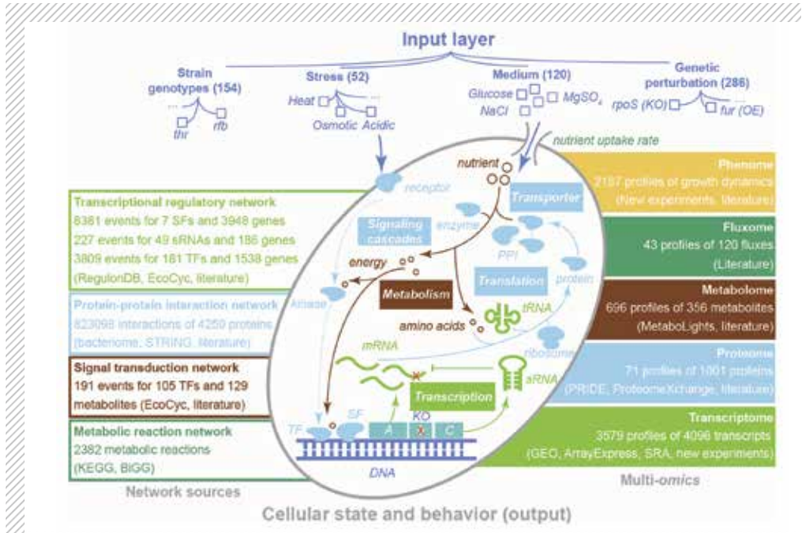

FIGURE 2: Layers of biological organization captured and information used to train the model.

WHY BLUE WATERS

Due to the complexity and size of the datasets (18 million points from various platforms and molecular species), the computational complexity of the algorithms that range from constrained regression to artificial neural networks, **Blue Waters was critical for the completion of the large-scale simulations** we performed.

NEXT GENERATION WORK

Our first publication appeared in *Nature Communications*. We have developed a deep learning method for proteome reconstruction, which has been tested already in Blue Waters and will be integrated into the model. Aside from algorithmic improvements, we aspire to move to population simulations of this model to investigate and predict population-level phenomena, similarly what we have done before, with simpler models [1,2]

PUBLICATIONS AND DATA SETS

Kim, M., N. Rai, V. Zorraquino, and I. Tagkopoulos, Multi-omics integration accurately predicts cellular state in unexplored conditions for Escherichia coli, *Nat. Commun.* 7:13090 EP (2016), doi:10.1038/ncomms13090.

THE COMPUTATIONAL MICROSCOPE

Allocation: NSF PRAC/19.6 Mnh
PI: Klaus Schulten[1]
Co-PIs: James C. Phillips[1] and John E. Stone[1]
Collaborators: Peijun Zhang[2], Angela Gronenborn[2], Tatyana Polenova[3], Christopher Aiken[4], Rebecca C. Craven[5], and Eva Nogales[6,7]

[1]University of Illinois at Urbana-Champaign
[2]University of Pittsburgh
[3]University of Delaware
[4]Vanderbilt University
[5]Penn State University College of Medicine
[6]University of California, Berkeley
[7]Howard Hughes Medical Institute

EXECUTIVE SUMMARY

In the petascale era, computational biology has redefined itself by connecting atomic-level descriptions of biological systems with cellular architecture and behavior. Molecular dynamics (MD) simulations serve as the computational microscope that empowers scientists with a tool to elucidate the dynamics and the underlying physical and chemical mechanisms of cellular processes. Due to its complementary role to experimental observation, MD unveils cellular organelles at atomic resolution.

INTRODUCTION

Recent developments in hybrid experimental methods, based on the revolutionary advance of electron microscopy, have led to previously unimaginable information on the cell-level structures that computational modeling requires for solidly-based descriptions. Very fortunately, computational modeling can play a significant role in hybrid method structure analysis [1]. First, the accuracy of computational modeling has drastically increased, such that results from computational studies today often exhibit astounding agreement with observation where available. Second, a broad range of so-called sampling methods based on statistical mechanical concepts is developed so that the biologically functional timescale in living cells can be covered by these simulations. The projects in this report leverage computational methods to combine structural data from multiple sources of differing resolutions (X-ray crystallography of individual proteins, medium-resolution cryo-EM (electron microscopy) of multi-protein systems, and low-resolution cryo-EM tomography of subcellular organelles) yielding atomic-resolution structural models of structures on the order of up to 100 nm in size.

METHODS & RESULTS

Solving the atomic-level structure of the mature HIV capsid [2] allowed us to study for the **first time** dynamic and structural properties of this multi-protein complex crucial for the biomedical community. Recently, experimental-computational studies demonstrated the influence of the human protein, Cyclophilin-A (CypA), in the dynamics of the capsid [3] and how hundreds of CypA bind to its surface [4]. Such properties may help scientists to understand better how the HIV capsid infects the host cells and could lead to new HIV therapies.

Virus infection and proliferation have been subjects of intense research for therapy development, targeting different maturation stages of the viruses. A strategy for preventing such proliferation, in particular, is to lock the viral particles in their immature, non-infectious state. Our investigation on the lattice model of the immature Rous sarcoma virus (RSV) [5], closely related to HIV, revealed a key source of dynamic and structural details presented on the recent experimental-computational characterization of the maturation process of the RSV virus [6].

The importance of simulating functional assemblies up to the level of complete capsids, opposed to isolated capsid proteins, was similarly demonstrated in a recent study. When bound to the HIV and hepatitis B virus (HBV) capsids, the drugs (PF74 on the HIV-1 and HAP1 on the HBV) caused changes not only in the vicinity of the binding sites but also on the global conformation of the capsids structure. The effect of PF74 was also evidenced on the dynamics of the processes taking place in distant regions of these structures [7].

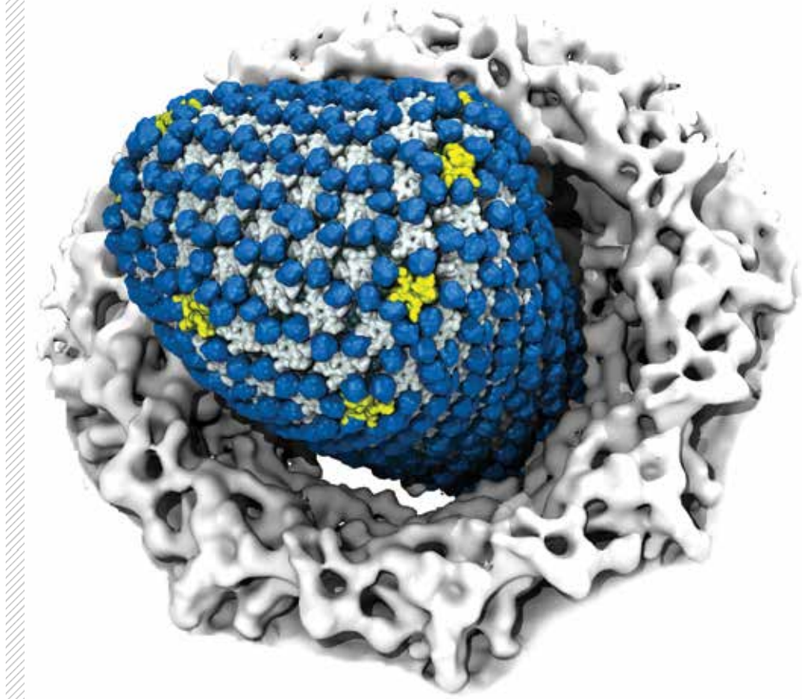

FIGURE 1: HIV-1 capsid docked to the human nuclear pore complex. HIV-1 exploits cellular pathways during infection, recruiting cyclophilin-A (blue) to go undetected in the cell.

Bacteria utilizes large, highly ordered clusters of sensory proteins, known as chemosensory arrays, to detect and respond to chemicals in their environment. Recently, we have integrated multi-scale structural data from experimental sources to computationally construct the **first atomic model** of the chemosensory array's molecular architecture [8]. Also, we identified a novel conformational change in a key signaling protein that is linked to chemotaxis function at the cellular level. This model may inspire and assist future experimental and computational studies in elucidating a general mechanistic description of signal transduction in the biological sensory apparatus.

Microtubules are a major component of the cell cytoskeleton, important for maintaining cell structure, intracellular transport, and cell division. Combining multi-scale structural data allowed us to build atomic models for microtubules in different nucleotide binding states (crucial for switch between phases of assembly and disassembly). Our MD simulations suggested important, structurally dynamic events towards the microtubules' assembly and stability. Such simulations pave the way to understand the atomic details of the assembly and disassembly phases of the microtubules as well as the effect of anticancer drugs on microtubule dynamic instability.

WHY BLUE WATERS

Without Blue Waters and other petascale computing resources, projects involving large molecular systems like HIV, RSV, HBV, chemosensory array, and

microtubule would not be possible. These molecular systems are composed by several dozens of millions of atoms and must be simulated for long periods of time (microseconds). These projects are examples of how **Blue Waters enables bold, new projects** that push the limits of what can be done with scientific computing. In our case, that means expanding molecular dynamics simulation capabilities from simulating just a few proteins to simulating full organelles.

PUBLICATIONS AND DATA SETS

Goh, B. C., et al., Atomic modeling of an immature retroviral lattice using molecular dynamics and mutagenesis. *Structure*, 23 (2015), pp. 1414-1425.

Fraiberg, M., CheY's acetylation sites responsible for generating clockwise flagellar rotation in Escherichia coli. *Mol. Microbiol.*, 95:2 (2015), pp. 231-44.

Sisneros, R., C. Malone, A. Nonaka, and S. Woosley Chemical visualization of human pathogens: the Retroviral Capsids. *SC15 Exhibition, Huston, TX: High Performance Computing, Networking Storage and Analysis.*

Lu, M., et al., Dynamic allostery governs cyclophylin A-HIV capsid interplay. *PNAS*, 112: 47 (2015), pp. 14617-14622.

Cassidy, C. K., et al., CryoEM and computer simulations reveal a novel kinase conformational switch in bacterial chemotaxis signaling. eLife, (2015), 10.7554/*eLife*.08419.

Stone, J. E., Atomic detail visualization of photosynthetic membranes with GPU-accelerated ray tracing. *Parallel Comput.* 55 (2016), pp. 17-27.

Liu, C., et al., Cyclophilin A stabilizes the HIV-1 capsid through a novel non-canonical binding site. *Nat. Commun.*, 7 (2016), pp. 10714-10724.

Stone, J. E., Sherman, W. R. and Schulten K., Immersive molecular visualization with omnidirectional stereoscopic ray tracing and remote rendering. IEEE International Parallel and Distributed Processing Symposium Workshop (IPDPSW) (IEEE Computer Society, Chicago, Illinois, May 23-27, 2016).

Stone, J. E., et. al. , High performance molecular visualization: In-situ and parallel rendering with EGL. IEEE International Parallel and Distributed Processing Symposium Workshop (IPDPSW) (IEEE Computer Society, Chicago, Illinois, May 23-27, 2016).

Perilla, J.R., et. al., All-Atom Molecular Dynamics of Virus Capsids as Drug Targets. *J. Phys. Chem. Lett.* 7 (20126), pp. 1836–1844.

Heyrana, K., et al., Contributions of charged residues in structurally dynamic capsid surface loops to Rous sarcoma virus assembly. *J. Virol*, 90:12 (2016), pp. 5700-5714.

Stone, J. E., Early experiences porting the NAMD and VMD molecular simulation and analysis software to GPU-accelerated OpenPOWER platforms. International Workshop on OpenPOWER for HPC (ISC High Performance conference, Frankfurt, Germany, June 23, 2016).

THE RECYCLING MACHINERY OF THE CELL

Allocation: Illinois/500 Knh
PI: Klaus Schulten[1]
Co-PIs: Till Rudack[1] , Lela Vukovic[2]
Collaborator: Wolfgang Baumeister[3]

[1]University of Illinois at Urbana Champaign
[2]University of Texas at El Paso
[3]Max Plank Institute of Biochemistry

EXECUTIVE SUMMARY

While waste recycling became popular in our daily life more recently, living cells have mastered recycling of their protein content since their very beginning. Recycling of unneeded protein molecules in cells is performed by a molecular machine called 26S proteasome, which cuts these proteins into smaller pieces for reuse as building blocks for new proteins. Proteins that need to be recycled are labeled by tags made of poly-ubiquitin protein chains. The 26S proteasome machine recognizes and binds to these tags, pulls the tagged protein close, then unwinds it, and finally, cuts it into pieces.

Despite its substantial role in the cell's life cycle, the proteasome's atomic structure and function remain elusive. Employing a combination of computational techniques implemented using nanoscale molecular dynamics (NAMD) along with cryo-electron microscopy (EM) data, we obtained an atomic structure of the human 26S proteasome and investigated the mechanism underlying substrate recruitment and unfolding.

FIGURE 1: The recycling system of the cell. The 26S proteasome is the key player of the human protein recycling system. The first structure of the human 26S proteasome obtained through integrative modeling utilizing Blue Waters will lead to breakthroughs in understanding its detailed function and will play a pivotal role in the development of the 26S proteasome as a drug target for molecular disease therapies.

INTRODUCTION

Recycling of proteins by degradation is vital for a variety of essential cellular processes, including protein quality control, cell cycle regulation, adaptive immune response, and apoptosis. The 26S proteasome is responsible for the vast majority of regulated intracellular protein degradation and is an important drug target for multiple diseases, including cancer, neurodegenerative diseases, and immuno-inflammatory disorders. The 26S proteasome is an adenosine triphosphate (ATP) hydrolysis driven 2.5 MDa molecular machine that recruits, unfolds, and degrades poly-ubiquitin tagged proteins through a complex interaction clockwork of over 60 known protein subunits (Fig. 1).

Despite its substantial role in the cell's life cycle, the proteasome's atomic structure and function remained elusive. However, recent developments in hybrid experimental methods based on the revolutionary advance of electron microscopy, together with improvements in real space refinement methods [1], have led to previously unimaginable

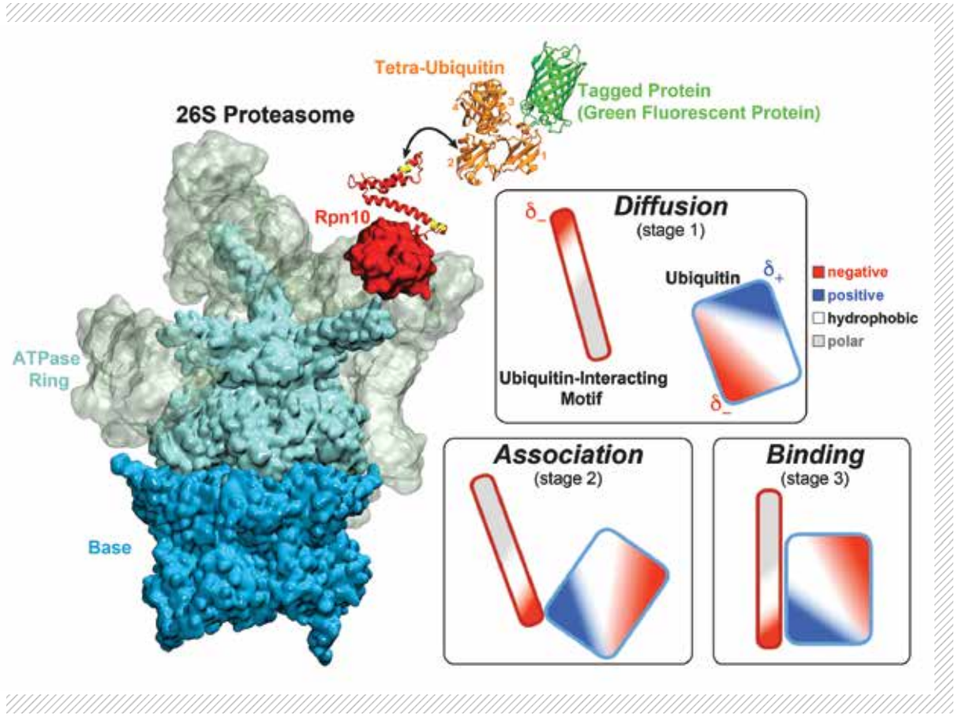

FIGURE 2: Recognition of the tetra-ubiquitin tag by the 26S proteasome in three stages: In stage 1 of the recognition process conserved complementary electrostatic patterns of Rpn10 and ubiquitins guide protein association; stage 2 induces refolding of Rpn10 and tetra-ubiquitin tag; stage 3 facilitates formation of hydrophobic contacts between the tag and Rpn10.

structural information on large macromolecular complexes [2].

In this report, we present the **first** atomic structure of the human 26S proteasome with bound nucleotides based on a 3.9 Å resolution cryo-EM density. It was obtained by integrative modeling combining molecular dynamics flexible fitting (MDFF) with *de novo* structure prediction algorithms [3]. Furthermore, utilizing molecular dynamics (MD), steered molecular dynamics (SMD) simulations, and generalized simulated annealing (GSA) techniques, we revealed the underlying mechanism of poly-ubiquitin tagged protein recognition and recruitment through the 26S proteasome [4].

METHODS & RESULTS

To build an accurate atomic model of the human 26S proteasome we followed the strategy established for large macromolecular complexes [1,2]. We first built comparative models of the human 26S proteasome subunits based on the yeast 26S proteasome and extended structurally-unresolved segments with *de novo* modeling. Next, we refined the structure according to the density in real space by combining MDFF with *de novo* structure prediction as well as Monte Carlo-based backbone and side chain rotamer search algorithms in an iterative manner. The features observed in the resulting structure are important for coordinating the proteasomal subunits during substrate processing. One of the key novel features is that in the predominant state of the heterohexameric ATPase motor of the proteasome there is one adenosine diphosphate (ADP) and five ATP-bound. The structure of the ADP-bound subunit is distinct from the other five subunits most notably in the pore loop region, which is known to be a key region interacting with the substrate during unfolding. Detailed investigations of the unfolding process by MD simulations and path sampling techniques are now possible and planned in future.

Furthermore, we employed MD, SMD, and GSA simulations on Blue Waters, utilizing QwikMD [5], our new intuitive "point and click" graphical interface connecting visual molecule dynamics (VMD) and NAMD. QwikMD significantly simplifies the setup, execution, and analysis of NAMD simulations from laptop computers up to supercomputers. Our simulations elucidated the process of how ubiquitin-tagged proteins are recognized by the partially disordered flexible arm of the ubiquitin receptor Rpn10 of the 26S proteasome (Fig. 2). Rpn10 consists of a globular domain that binds to the 26S proteasome and a flexible arm that contains ubiquitin-interacting motifs (UIM). Previous studies identified that ubiquitin binds the UIMs through hydrophobic interactions; however, poly-ubiquitin mostly assumes closed forms in solution, in which the hydrophobic patches are protected. Our results reveal that hydrophobic UIMs of Rpn10, which bind to hydrophobic patches of ubiquitin, are likewise mostly protected prior to binding. Therefore, refolding of both tetra-ubiquitin and Rpn10 to reveal hydrophobic patches is a necessary step of the binding process. Since it is unlikely that protected hydrophobic patches can initiate refolding and efficiently search for each other, our results indicate that electrostatic interactions are responsible for the initial stage of the mutual recognition between UIMs and Rpn10.

The combination of MDFF with Monte Carlo-based rotamer search algorithms enabled us to obtain a structural model based on high-resolution (< 4.0 Å) cryo-EM densities and obtain new insights into substrate recognition and unfolding. The obtained atomic structure permits further computational studies of proteasomal function and will serve as a starting point for future structure-guided drug discovery that will assist to develop the proteasome as a drug target.

WHY BLUE WATERS

Investigating the functional processes of large protein machinery, such as the proteasome (3.5 M atoms), which occur on the millisecond timescale, is **only possible on petascale** computing resources, such as Blue Waters.

De novo structure prediction and Monte Carlo-based sampling algorithms are only efficient if thousands of models are predicted. Employing these algorithms for the numerous subunits of the proteasome is a well-suited task for the large-scale parallel architecture of Blue Waters.

NEXT GENERATION WORK

We would like to investigate how the motor action of ATP hydrolysis in the proteasome is driving the subunit reorganization during the functional cycle of substrate processing, which requires extensive sampling of long trajectories.

In order to exploit the proteasome's role as a drug target, high-throughput screening of thousands of drugs that possibly interact with the proteasome might be feasible by the year 2020 and would require pre-exascale computing for such a complex macromolecular machinery as the 26S proteasome.

PUBLICATIONS AND DATA SETS

Schweitzer, A., et al., The structure of the human 26S proteasome at a resolution of 3.9 Å. *PNAS*, (2016).

Zhang Y., et al., Recognition of poly-ubiquitins by the proteasome through protein re-folding guided by electrostatic and hydrophobic interactions. *J. Phys. Chem. B*, (2016). DOI: 10.1021/acs.jpcb.6b01327

SOCIAL SCIENCE, ECONOMICS, & HUMANITIES

SOCIAL SCIENCE

ECONOMICS

STATISTICS

A COMPUTATIONAL MODEL FOR CAUSAL INFERENCE VIA SUBSET SELECTION

Allocation: Illinois/50.0 Knh
PI: Wendy K. Tam Cho[1]
Co-PI: Yan Y. Liu[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Researchers in all disciplines desire to identify causal relationships. Randomized experimental designs isolate the treatment effect and thus permit causal inferences. However, experiments are often prohibitive because resources may be unavailable or the research question may not lend itself to an experimental design. In these cases, a researcher is relegated to analyzing observational data. To make causal inferences from observational data, one must adjust the data so that they resemble data that might have emerged from an experiment. The data adjustment can proceed through a subset selection procedure to identify treatment and control groups that are statistically indistinguishable. Identifying optimal subsets is a computationally complex and challenging problem but a powerful tool for discovering scientific insights in a wide variety of fields.

INTRODUCTION

The aim of the project is to design efficient computational and statistical algorithms for making causal inferences. The research builds on previously developed statistical models [1]. The proposed research advances this concept from both theoretical and applied perspectives—it furthers discrete optimization models that capture the features and subtleties of the causal inference problem and provides quantitative tools and models that address the causal structures of interest in important substantive problems in a diverse set of scholarly domains, respectively.

Enhancing the ability to make causal inferences from observational data will stimulate research in a wide variety of fields and enhance our understanding of a broad array of phenomena. In medicine or health, causal inference studies have included applications to criminality rates related to gene patterns [7], the effect of generic substitution of presumptively chemically equivalent drugs [6], in utero exposure to phenobarbital on intelligence deficits [5], and the effect of maternal smoking on birthweight [2], to name but a few. In social science, causal inference models have been applied to studies on the impact of different voting technologies [3] and the effect of affirmative action [4]. The array of potential research questions is important, diverse, and unconstrained by field of study. Additionally, the possible applications of causal inference models are limitless given a proper research design and available data.

METHODS & RESULTS

Our **novel** approach to the data matching problem proffers a paradigm change from statistical matching methods. The statistics literature focuses on identifying individual data matches with models that are dependent on a set of underlying assumptions. While we retain the same goal of post-processing observational data to resemble experimental data, we **redefine** the process to obtain the goal of covariate balance directly, bypassing the assumptions of the statistical models entirely by replacing them with a computationally challenging problem. By consolidating steps to determine the optimal subsets of the controls and treatments based on some set of balance measures, we show that important insights into many problems that have been traditionally analyzed via statistical models can be obtained by re-formulating and evaluating within a large-scale optimization framework.

Our analysis speaks to statistical frameworks with astronomical (10^{5000}) solution spaces. Here, the identified solutions need to be independent of one another. This independence requirement adds a significant challenge for standard optimization methodologies. In this vein, we designed a hybrid metaheuristic with specialized intensification and diversification protocols in the base search algorithm. We experimentally demonstrate that our diversification protocol cuts the time required to find independent solutions in half while our intensification protocol enables the identification of solutions that are difficult to find with non-collaborative processors. We extend our algorithm to the high-performance-computing realm by implementing methods for utilizing multiple processors to collaboratively hill climb, broadcast messages to one another about the landscape characteristics, diversify across the landscape, and request aid in climbing particularly difficult peaks.

For scalability, our code eliminates the costly global synchronization operation which was originally conducted for incoming message checking. With the global synchronization, the message passing time ranged from 42% on 240 cores to 72% on 960 cores. The message passing cost declined considerably with our asynchronous communication protocol (0.007% on 240 cores and 0.01% on 960 cores). Figure 1 shows that the numerical performance scales well in our weak scaling experiments. We continue to explore scalability with larger numbers of cores as the problem instances for most practical problems are vast and require a substantially greater number of processing cores to obtain satisfactory results.

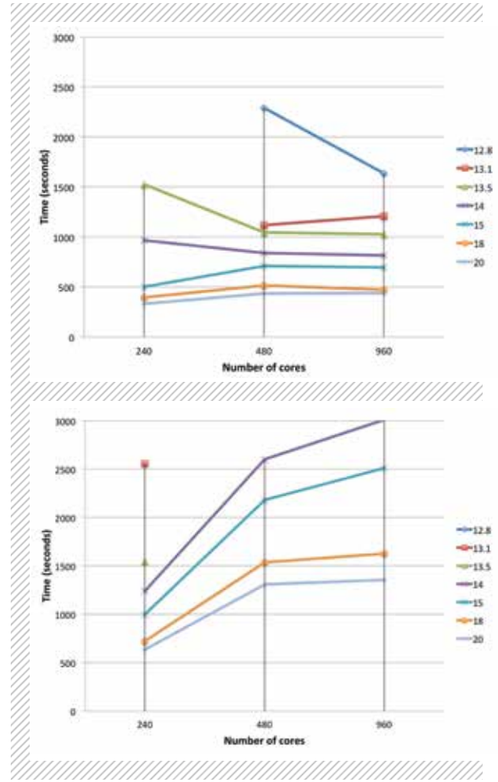

FIGURE 1: Results from Weak Scaling Experiments. Asynchronous (top) versus Synchronous (bottom) Communication.

WHY BLUE WATERS

Balance Optimization Subset Selection (BOSS)'s shift from individual matching to subset selection highlights an interesting combinatorial aspect of both the matching methodologies as well as the subset selection methodology. In particular, for even moderately sized data sets, the set of possible "solutions" is extremely large. For instance, if our control pool has 100 members, and we wish to choose a subset of size 20, there are $\binom{100}{20} = 5.359834 \times 10^{20}$ possibilities. This would be a small instance of an actual substantive problem. In a classic data set for this problem, the control pool has 16,000 members from which we choose a subset of size 185—an astronomically large problem. Given the sheer size and the non-rugged solution landscape, finding balanced subset in this solution space proves to be computationally challenging. The problem presents an extreme-scale optimization problem for which a petascale resource like Blue Waters is necessary.

PUBLICATIONS AND DATA SETS

Cho, W.K.T., and Y.Y. Liu, A parallel evolutionary algorithm for subset selection in causal inference models. Proceedings of the 2016 Annual Conference on Extreme Science and Engineering Discovery Environment. XSEDE'16: Diversity, Big Data, & Science at Scale, Miami, FL, July 17–21, 2016.

Cho, W.K.T., and Y.Y. Liu, A high-performance approach for solution space traversal in combinatorial optimization. SC15: The International Conference for High Performance Computing, Networking, Storage and Analysis. Austin, TX, November 16–19, 2015.

POLICY RESPONSES TO CLIMATE CHANGE IN A DYNAMIC STOCHASTIC ECONOMY

Allocation: GLCPC/530 Knh
PI: Lars Hansen[1]
Co-PIs: Kenneth Judd[2], William Brock[3], and Thomas Hertel[4]
Collaborators: Yongyang Cai[1], Simon Scheidegger[5], and T. J. Canann[6]

[1]University of Chicago
[2]Hoover Institution
[3]University of Wisconsin
[4]Purdue University
[5]University of Zurich
[6]University of Minnesota

EXECUTIVE SUMMARY

We extended our Integrated Assessment Model (IAM) framework, called DSICE (Dynamic Stochastic Integration of Climate and the Economy), for evaluating alternative policy responses to future climate change. The new elements allow us to use stochastic processes to model risks in both the climate and economic systems, and to examine how robust our results are to uncertainties about parameters of those stochastic systems [1].

One substantive result of DSICE is that the social cost of carbon (SCC) is substantially greater when we include economic and climate risks in analyzing the impact of climate change on the economy. Moreover, the SCC is itself a stochastic process with significant variation. Our findings show that a single parameterization of DSICE, representing the most recent empirical macroeconomic analysis of economic growth, implies a stochastic process for carbon emissions with ~10% chance of emissions exceeding all of the Intergovernmental Panel on Climate Change (IPCC) emissions scenarios.

INTRODUCTION

Climate change will have significant impacts on human life, through altered agricultural productivity, changes in the demand for heating and air-conditioning, and changes in the flora and fauna valued by humans. Integrated Assessment Models (IAMs) aim to analyze the impact and efficacies of alternative policy responses to climate change. The benchmark IAM model is DICE [2]. It is frequently used to analyze economic policies for dealing with climate change in the United States Interagency Working Group on Social Cost of Carbon [3]. However, it ignores heterogeneity in the economy as well as uncertainty in economic conditions and the impact of warming on the climate. The perfect foresight modeling choices are often excused on the ground that computational limitations make it impossible to do better.

Our work on Blue Waters clearly showed otherwise. We construct computational frameworks merging the basic elements necessary for any such analysis—how the climate reacts to anthropogenic-induced changes in the atmosphere, how those changes affect the economy, alternative policies, how individual economic agents respond to uncertainties, and how policymakers should incorporate various uncertainties into their decisions.

METHODS & RESULTS

We are building a computational framework to address the issues discussed above. Cai, Judd, and Lontzek have used Blue Waters to develop a computational model, called DSICE [1], that allows shocks to the economic and climate systems and uses specifications for agent preferences consistent with empirical evidence of how much people are willing to pay to reduce risks about the future.

In 2015, Cai, Judd and their collaborators published two papers applying DSICE with Blue Waters support in prestigious journals: *Nature Climate Change* [4] and *PNAS* [5]. They found that the risk arising from potential tipping points substantially raises the optimal level of emission control, meaning that a low rate of discount should be used for climate damage when evaluating climate policies in the presence of tipping points.

In 2016, Cai and collaborators extended DSICE to study the impact of multiple interacting tipping points with Blue Waters support [6]. The article examines the solution of a dynamic stochastic programming problem with ten continuous state variables and five discrete state variables.

For policy analysis on climate change with multiple sectors, our goal is to solve competitive equilibrium and dynamic stochastic problems in high dimensions. In the past year, Cai and Judd used Blue Waters to develop a nonlinear certainty equivalent approximation method (NLCEQ) to solve them efficiently and in parallel [7].

To implement effective policy, ideas need to be complemented by analysis of what is possible, given the limitations placed on policy by social institutions. In the case of climate change, the key challenge is to find international agreements that sovereign governments can support. This is a problem in game theory; more precisely, one needs to find solutions using Nash equilibria in the dynamic interactions of various countries and their policy choices. These interactions can be modeled as supergames. Cai, Judd, and Yeltekin used Blue Waters to develop the algorithms that can solve supergames with states [8]. Cai, Hertel and Judd have also examined how the world demand for food, oil, and timber interacts to alter land use patterns [9].

WHY BLUE WATERS

Solving these high-dimensional dynamic stochastic problems is time intensive, but parallelizing dynamic programming methods allowed us to solve them efficiently [10]. In our benchmark example, we used up to 84,000 cores and it scaled almost linearly; a serial computation would take about 77 years [1]. We also developed NLCEQ to solve even higher-dimensional problems in parallel, e.g., a 400-dimensional dynamic stochastic problem [7]. Another algorithm, developed for solving a dynamic game with multiple players and many choices in Blue Waters [8], would take decades without parallelism.

NEXT GENERATION WORK

Our goal is to merge critical sectors, such as agriculture and forestry, with an energy balance climate model within a dynamic optimization model where policies are chosen to maximize an objective that accounts for different income groups and generations and reflects future uncertainty. DSICE has provided the beginning of this framework. Our next generation work will extend DSICE with Blue Waters to perform research on policy analysis with Bayesian learning, robust decision making, multi-sector integrated assessment of climate and economy under uncertainty, and geoengineering.

PUBLICATIONS AND DATA SETS

Lontzek, T.S., Y. Cai, K.L. Judd, and T.M. Lenton (2015). Stochastic integrated assessment of climate tipping points indicates the need for strict climate policy. *Nat. Cli. Change*, 5 (2015), pp. 441–444.

Cai, Y., et al., Environmental tipping points significantly affect the cost-benefit assessment of climate policies. *PNAS*, 112:15 (2015), p. 4606–4611.

Cai, Y., T.M. Lenton, and T.S. Lontzek, Risk of multiple climate tipping points should trigger a rapid reduction in CO2 emissions. *Nat. Cli. Change*, 6 (2016), p. 520–525.

GRADUATE FELLOWS

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

Maureen T. Brooks, University of Maryland Center for Environmental Science
2015-2016 Graduate Fellow

RESEARCH SUMMARY

The macroalgae commonly known as "gulf weed," *Sargassum fluitans* and *Sargassum natans*, are keystone species in the Atlantic Ocean and the Gulf of Mexico, where they serve as habitat and forage for a diverse floating ecosystem [1, 2]. Satellite observations [3] indicate a seasonal cycle of *Sargassum* distribution with high abundances in the Gulf of Mexico in the spring and early summer, and in the Sargasso Sea in fall and winter, however, the drivers of this seasonal pattern are not well understood. Recent changes in *Sargassum* biomass and extent [4] coincide with increasing reports of wash-ups on beaches in Africa, the Caribbean, and the Gulf of Mexico which negatively impact fishing and tourism. I have used a coupled modeling approach to understand the drivers of *Sargassum* distribution across temporal and spatial scales.

To simulate ocean circulation in the North Atlantic, I have generated a Hybrid Coordinate Ocean Model (HYCOM) [5] domain at 1/12° resolution which is forced using the National Centers for Environmental Prediction Climate Forecast System Reanalysis for the years 1992-2010. This domain is coupled to a biogeochemical model adapted from the work of Fennel [6] that includes phytoplankton, zooplankton, and variable nutrient stoichiometry relevant to *Sargassum* growth in oligotrophic conditions. Also coupled to this model system are a Lagrangian particle model and an individual-based *Sargassum* growth model developed for this study.

The addition of the *Sargassum* individual-based model resulted in a threefold increase in percent areal match with observations compared with particle advection from the physical model alone. Over the seasonal cycle at the basin scale, the coupled model system can generate a *Sargassum* distribution that is consistent with observations (Fig. 1) from a randomly distributed particle initial condition. Typically the Sargasso Sea is considered the critical *Sargassum* habitat in the north Atlantic, however analysis of particle locations and growth rates highlights the importance of the tropics as a reservoir for the *Sargassum* population. Connectivity matrices generated using the Lagrangian particle model show high retention of *Sargassum* in the Sargasso Sea, as expected, but also show moderate retention in the western Gulf of Mexico. These retentions suggest that the western Gulf of Mexico region may contain a seed population of *Sargassum* that remains local while a fraction of its biomass gets advected out towards the central gyre.

The combination of the circulation model and the *Sargassum* growth model has also allowed for examination of mortality and export. This model system shows a hot spot of *Sargassum* mortality and sinking that is consistent with observations of *Sargassum* on the sea bed at the western boundary of the Sargasso Sea [7]. Physical advection tends to aggregate *Sargassum* in the Sargasso Sea with less than a 30% probability of escape to better growth conditions at timescales of one year.

At spatial scales less than 100km and temporal scales less than one month, mesoscale eddies can influence *Sargassum* transport as well as nutrient conditions via differential vertical velocities.

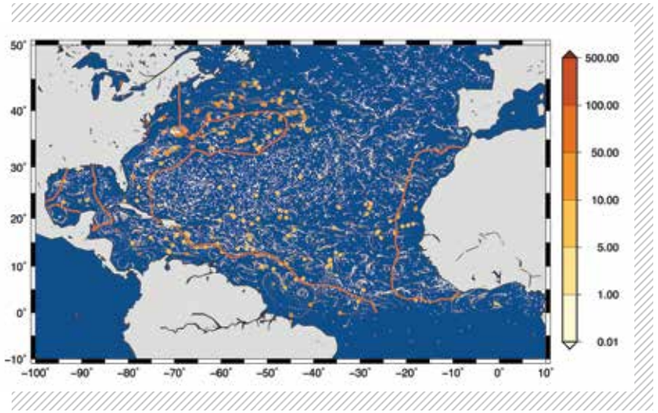

FIGURE 1: Modeled Sargassum distribution in November. Particle color indicates Sargassum biomass normalized to initial condition, orange contours are regions of high Sargassum biomass observed via satellite.

Preliminary Lagrangian coherent structure analysis in the Gulf of Mexico at the start of the *Sargassum* growing season illustrates the difference between the relatively quiescent western Gulf of Mexico and the eddy field associated with the Loop Current in the east (Fig. 2). Because particles tend not to cross Lagrangian coherent structures, these features may present a boundary that helps explain the high retention of *Sargassum* in the western Gulf of Mexico.

Investigation of the links between mesoscale features and *Sargassum* is ongoing. I will examine transects of nutrient concentrations and *Sargassum* location and growth across eddies and fronts in the tropics, the Gulf of Mexico, and along the Gulf Stream wall, and attempt to link these features with the large-scale *Sargassum* distribution. Application of these results will help forecast costly *Sargassum* wash-up events more effectively. I also plan to conduct comparisons of these results with those from a coarse resolution model system to gain insight into the spatial and temporal resolutions at which online coupling is necessary for individual-based models.

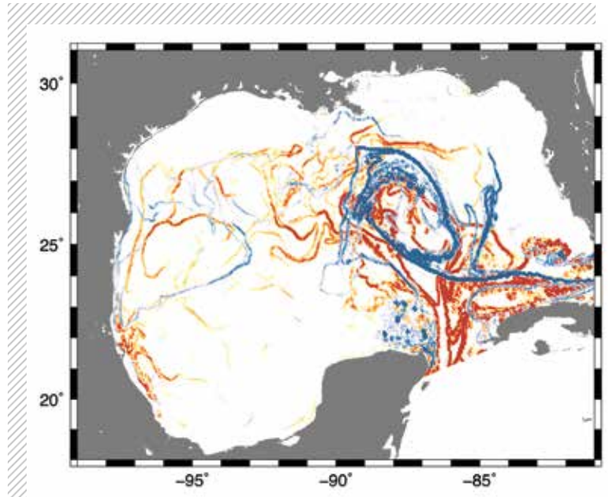

FIGURE 2: Lagrangian coherent structures (LCS) in the Gulf of Mexico at the start of the Sargassum growing season (mid-May). At short time scales Sargassum will aggregate along lines of attracting LCS (blue) and be exported from regions of repelling LCS (red).

WHY BLUE WATERS

Access to Blue Waters has been critical to the success of this project. Modeling ocean circulation at this scale and resolution has a high computational cost, and the capabilities of Blue Waters have allowed me to couple both physics and biogeochemistry. The expert staff at NCSA have been extremely responsive to my questions and needs and have helped me grow from a novice into a proficient user with plans to continue using high-performance computing to study physical-biological interactions across marine ecosystems.

PUBLICATIONS AND DATA SETS

Brooks, M. T., V. J. Coles, R. R. Hood, and J. F. R. Gower, Linking satellite observations with coupled bio-physical models of *Sargassum*. *AGU/ASLO/TOS Ocean Sciences Meeting*, New Orleans, LA, February 21-27, 2016.

Maureen T. Brooks is in the third year of her Marine-Estuarine Environmental Sciences Ph.D. studies at University of Maryland. She expects to graduate in the spring of 2018 and wishes to pursue a career at the interface between applied science and resource management.

"Ideally, I would like to conduct applied research that will better inform marine resource policy. The Blue Waters Fellowship has opened doors for me by providing firsthand experience with the high-performance computing resources necessary for cutting-edge science in oceanography. This experience has directed me to choose a career that will allow me to help ensure that advances in scientific understanding are translated into the best management of our shared natural resources."

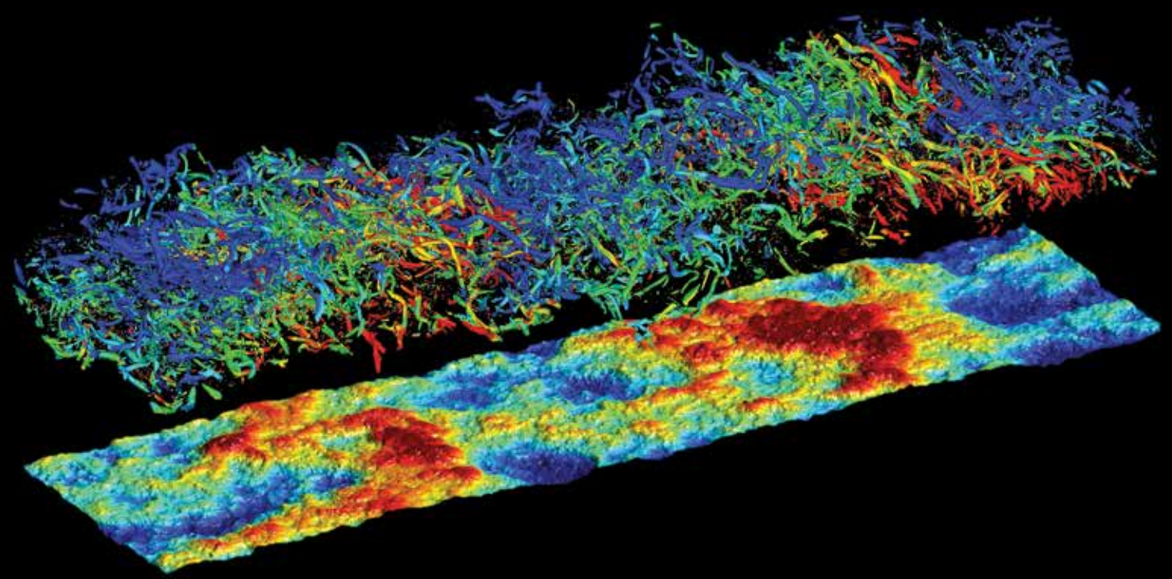

LARGE EDDY SIMULATIONS OF AERO-OPTIC DISTORTIONS

Edwin Matthews, University of Notre Dame
2014-2015 Graduate Fellow

RESEARCH SUMMARY

The distortion of an optical beam caused by turbulent air surrounding a projecting or receiving aperture, known as aero-optics, is a major impediment to applications of airborne optical systems for communication, imaging, targeting, and directed energy systems. When an initially planar optical wavefront is transmitted through a compressible turbulent flow, it is distorted due to the non-uniform speed of light resulting from density fluctuations. These distortions can greatly affect the fidelity and coherence of an optical beam as it propagates to far distances, rendering some airborne optical systems effectively unusable.

In the past year, research using large-eddy simulations (LES) of compressible flows to investigate aero-optic phenomena has progressed on two fronts. First, the numerical investigation of a hemisphere-on-cylinder optical turret at the experimental Reynolds number of 2.3 million has continued to advance. To better converge statistical quantities, the runtime of the finest resolution simulation has been doubled and over 90 TBs of optical and flow field data have been collected to study the aero-optic environment. Also, a coarse mesh simulation of the same configuration was calculated to investigate the sensitivity of the solution to grid resolution. By analyzing nearly 400 optical viewing angles calculated from the simulation database, our knowledge about the effects of turbulent wake structure, viewing angle, and aperture size on optical distortions found in a common beam transmission platform has greatly improved.

In addition, a parallel solver to compute data-driven decompositions like Proper Orthogonal Decomposition (POD) and Dynamic Mode Decomposition (DMD) was developed to explore the dynamics of the turret flow and aero-optics. Their computation requires the calculation of the economy singular value decomposition (SVD) of a large dense matrix, requiring the memory and I/O capabilities of Blue Waters. POD calculations of the density field using the solver have scaled well up to 1,024 XE nodes using 2,000 time instances each composed of 200.5 million data points.

Finally, to study the spectral behavior of aero-optic phase distortions, high-fidelity LES of weakly compressible mixing layers with different initial conditions were computed. The spectral behavior of aero-optical distortions is of fundamental interest to the field, helping to understand the role of different length scales of turbulence in aero-optic phenomena and aiding in the development of adaptive-optic methods. Optical results from the simulations were compared with theoretical expressions derived for the spectral behavior of optical phase distortions and confirmed a direct relationship to the spectral behavior of density fluctuations. Density fluctuations can then be shown to be dependent on the conditions of the top and bottom streams of the mixing layer. Depending on flow conditions, simulation results show that small-scale density fluctuations can be strongly dependent on temperature fluctuations, pressure fluctuations, or a combination of both.

WHY BLUE WATERS

A unique aspect of Blue Waters has been the ability to generate massive simulation databases, move them quickly between near-line storage to scratch space, and then efficiently read and write the data in post-processing. Capacity to handle data in large volumes combined with the quick queue times allows for a much higher level of productivity compared to other systems. Blue Waters can run through tens of terabytes of data in post-processing, moving all of the data needed in hours, not days. This aspect has enabled our research to progress quickly, taking an idea for data processing through to realization within a day.

PUBLICATIONS AND DATA SETS

Mathews, E., K. Wang, M. Wang, and E. J. Jumper, LES of an aero-optical turret flow at high Reynolds number. Proceedings of the 54th AIAA Aerospace Sciences Meeting, AIAA SciTech, San Diego, California, January 4-8, 2016.

Mathews, E., K. Wang, M. Wang, and E. J. Jumper, LES prediction and analysis of the aero-optical environment around a 3-D turret. 68th Annual Meeting of the APS Division of Fluid Dynamics, Boston, Massachusetts, November 22-24, 2015.

Mathews, E., K. Wang, M. Wang, and E. J. Jumper, Numerical investigation of aero-optical distortions over a hemisphere-on-cylinder turret with gaps, AIAA Aviation, Dallas, Texas, June 22-26, 2015.

FIGURE 1 (LEFT): Coherent turbulence structures in a weakly compressible mixing layer are visualized by isosurfaces of the q-criterion and colored by fluid density. An initially planar optical wavefront propagating through the mixing layer will be distorted by changes in the refractive index due to turbulent density fluctuations. The result is an aberrated wavefront as shown on the surface under the turbulent structures.

Edwin Matthews is in his fifth year as a Ph.D. student studying Aerospace Engineering at the University of Notre Dame. Edwin plans to finish his degree in 2017 and wishes to pursue a career in industry or at a national laboratory.

"The use of high-fidelity computational fluid dynamics methods, like LES and hybrid LES/RANS, is becoming more commonplace in the design of engineering fluid systems. Blue Waters has allowed me to gain experience with these methods using a world-class high-performance computing system while learning to manage large volumes of data and to program with scalability in mind."

IMPROVING CHECKPOINT RESTART WITH LOSSY COMPRESSION

Jon Calhoun, University of Illinois at Urbana-Champaign
2014-2015 Graduate Fellow

RESEARCH SUMMARY

Checkpoint restart is a fundamental component of HPC applications and is used to recover the application from system failure and to extend execution beyond a single time allocation. Checkpointing to the parallel file system is prohibitively expensive, but is often unavoidable due to its non-volatility. Multi-level checkpointing schemes have been developed to utilize the memory hierarchy to improve checkpointing times, therefore reducing expensive file system operations. The next generation of HPC systems such as ANL's Aurora and ORNL's Summit are expected to increase the amount of total system memory by 5-9x over current 10 and 27 petaflop machines. However, the file system bandwidth will remain around 1 TB/s. Failure to consistently scale file system bandwidth can have negative effects on HPC applications. Applications that are scalable on current systems may not be scalable on future systems as file system operations will become a larger fraction of runtime. As machine size continues to scale, the system's mean time between failure (MTBF) is expected to decrease, resulting in more checkpoints being taken [1] resulting in a further reduction in time spent on computation. Reducing data movement is key for efficient usage of HPC systems.

Data compression techniques can be used to reduce dataset size and come in two forms: lossless and lossy. Lossless compression schemes provide small compression factors for floating-point data sets [2]. Lossy compression schemes [3, 4], can provide significantly higher compression factors at the expense of a small but controllable error added to the data. For lossy compression to be used in HPC applications, understanding this compression error is paramount. Selection of this compression error is naively a trial and error based approach and requires detailed knowledge of what is simulated to judge acceptance. Expressing the compression error in mathematical terms that would be familiar to application scientists will help adoption by providing a generic methodology to interpret the compression error.

Many HPC applications approximate the solution of partial differential equations (PDE) or ordinary differential equations (ODE) by using a numerical method. This numerical method approximates the solution to a given accuracy dependent on the spatial discretization of the problem. The accuracy of the simulation as expressed by the numerical methods and spatial discretizations can be leveraged in the selection of the lossy compression error tolerance. For example, if we know our simulation is accurate to 1e−4, then a numerical solution to this problem, u^h, is indistinguishable from a slightly perturbed solution, $\tilde{u}^h = u^h + \epsilon$, if $\epsilon < 1e{-}4$. This observation creates an upper bound on the compression error tolerance such that restart from a lossy checkpoint is indistinguishable from that of a lossless checkpoint. If we seek higher compression factors by setting $\epsilon < 1e{-}4$ (adding in error), we can assign compression error tolerances to the accuracy of related numerical methods, e.g. varying order Runge-Kutta methods. This selection methodology is outlined in Fig. 1.

We test our compression error selection methodology using plasma physics code PlasComCM. PlasComCM solves the Navier-Stokes equations using Runge-Kutta 4. Our problem simulates flow past a fixed cylinder and is accurate to 1.8e−5. We use SZ version 0.5.14 [3] with compression error tolerance ϵ=1e−5 and restart from a lossy compressed checkpoint every 5000 time-steps. We plot the inf-norm of the error

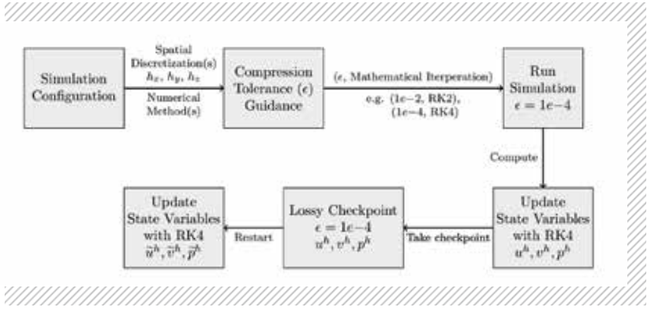

FIGURE 1: Overview of the method to select lossy compression error tolerance. Compression error tolerance can be related to accuracy of numerical methods by knowing the simulation's spatial discretization.

in each state variable against time (Fig. 2). Error in the state variables never exceeds the accuracy of the simulation's mathematics. There is an accumulation of error with each additional checkpoint, but it decreases due to non-periodic boundaries allowing the perturbed flow to exit the domain. Propagation and removal of error in the momentum variable, between the time-step after the first lossy restart until the time step just before the second restart is observed (Fig. 3). Initially, error propagates until around time 310, after which the error begins to exit the domain. Exploitation of boundary conditions and physical properties yields knowledge for selecting an optimal compression error. For problems that attenuate error, it may be possible to compress to a tolerance above the accuracy without affecting the outcome of the simulation.

Lossy compression is an exciting new research direction that shows great promise in reducing the volume of data transmitted by HPC applications. Compared with standard lossless compression, lossy compression can achieve higher compression factors, but at the expense of adding error into the simulation. The results discussed above, and our ongoing research, suggest this error can be interpreted mathematically. Thinking of lossy compression error in this way shifts its application from a trial and error process to the well-understood domain of understanding numerical error.

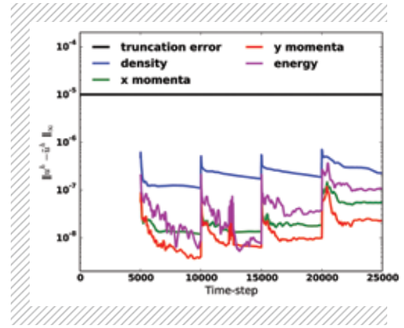

FIGURE 2: Inf-norm between numerical solution u^h and compressed numerical solution $\tilde{u}^h$. There is no error before time-step 5000, as the simulation has yet to be restarted. At each restart, we see a large spike in the error, but error remains less than the simulation's accuracy (truncation error). This error is partially removed by our selection of boundary conditions.

WHY BLUE WATERS

Blue Waters represents a balanced HPC system where fast compute nodes are complimented by I/O bandwidth of more than 1 TB/s. This **advantage**, combined with the heterogeneous nature of the XK nodes, makes Blue Waters an **ideal** system for testing next generation optimizations for checkpoint restart.

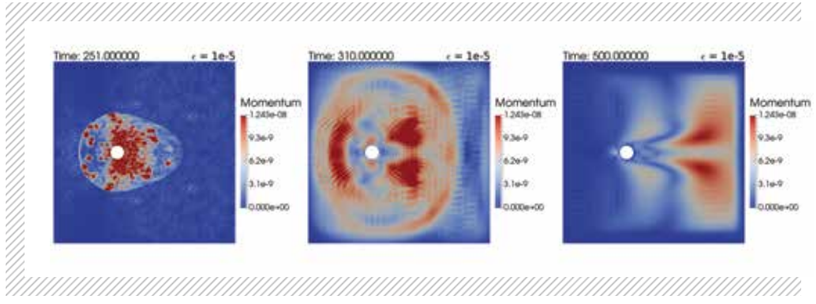

FIGURE 3: Location of error in the domain after first restart from a lossy compressed checkpoint, time-steps 5001-9999. After restart, error is confined near the cylindrical object (left plot). As time evolves, error is propagated though the domain (center plot), and eventually exits the domain due to our selection of boundary conditions (right plot).

John Calhoun is in the fourth year of his Computer Science Ph.D. studies at the University of Illinois at Urbana-Champaign. He expects to graduate in August 2017 and aspires to teach at a university or work at a national laboratory.

"Working on the Blue Waters system has given me valuable experience on a flagship HPC system. Without access to such a system, it would be difficult to test the new resilience ideas I am investigating. The research that I am conducting on Blue Waters serves as a stepping-stone to what I will accomplish post-graduation."

SOLVATION THERMODYNAMICS OF THE PROTEIN BACKBONE: IMPLICATIONS FOR COLLAPSE AND AGGREGATION

Justin Drake, University of Texas Medical Branch
2015-2016 Graduate Fellow

RESEARCH SUMMARY

Intrinsically disordered regions (IDRs) in proteins assume a diverse ensemble of highly dynamic, flexible structures. The development of drugs to treat IDR-mediated diseases requires extensive knowledge of the thermodynamic and structural properties of these disordered regions. Here, we wish to understand if the thermodynamic mechanisms that drive the proper folding of well-structured proteins similarly promote the collapse of IDRs to their native state as well as the aggregation of IDRs in a pathological state. Using molecular simulations and computational free energy methods, we calculated the solvation free energy, enthalpy, and entropy of oligoglycine, a protein backbone model, as a function of chain length. We find that the entropic penalty upon solvating successively longer oligoglycines does not explain the experimentally observed aggregation of short oligoglycines and collapse of longer ones. Rather, we propose that favorable peptide-peptide interactions outcompete favorable peptide-solvent interactions in a concentration and length-dependent manner.

For much of the 20th century, it was widely held that the well-defined three-dimensional structure of a protein dictates its cellular function. While true for a large group of proteins, at the turn of the century, seminal papers [1, 2] opened the door to an entirely new class of proteins that rely on highly dynamic, flexible (i.e. disordered) regions to carry out their function. These IDRs found within proteins play a major role in the protein signaling networks and regulation, and, unsurprisingly, have been implicated in many diseases [2]. If we are to be as successful in targeting drugs to IDRs as we are in targeting well-structured proteins, we need to have an extensive, complete understanding of the structural properties of IDRs and how they enable protein function. IDRs assume a diverse ensemble of structures that range from completely collapsed to highly extended depending, in part, on their sequence composition and preference to interact with the solvent as well as their length or number of amino acids [3]. However, the biophysical mechanisms that drive the collapse or extension of IDRs are not well understood.

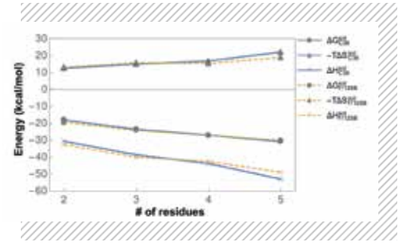

FIGURE 1: Solvation free energy (circle), enthalpy (x), and entropy (triangle) of successively longer oligoglycines calculated using free energy perturbation with the C36 (blue) and ff12SB (orange) force fields.

To uncover these mechanisms, we study the solvation thermodynamics of successively longer oligoglycine models ($Gly_{2\text{-}5}$), where the subscript denotes the number of glycine residues. Oligoglycine is an ideal protein backbone model [4] and is also found in many IDRs [5]. We use computational free energy methods to calculate the solvation free energy (ΔG^{sol}), entropy (ΔS^{sol}), and enthalpy (ΔH^{sol}) of $Gly_{2\text{-}5}$ with the commonly used CHARMM36 (C36) [3] and Amber ff12SB [7] force fields. We find that ΔG^{sol} is negative (i.e. favorable) and continues to decrease with the length of oligoglycine. The decrease in ΔG^{sol} with chain length is driven by a large, negative ΔHsol that is only moderately counterbalanced by an unfavorable entropic component, $-T\Delta S^{sol}$ (Fig. 1). Interestingly, experiments and simulations show that short oligoglycines (~Gly_5) become increasingly insoluble (aggregate) with chain length and long oligoglycines (>Gly_{10}) collapse in highly dilute solutions [4, 8, 9], but the mechanism that drives the aggregation of short oligoglycines and collapse is unclear. However, our results and those of others [3, 6, 9] suggest that favorable, non-specific peptide-peptide interactions outcompete the still favorable peptide-solvent interactions and that these weak peptide-peptide interactions are amplified in a length and concentration dependent manner. Whereas the initial collapse and folding of well-structured proteins are widely accepted as an entropically driven process, it may not play as dominant a role in the collapse of IDRs. Lastly, we find that despite C36 and Amber ff12SB generating different structural ensembles of various oligoglycine chains [10], both yield remarkably similar solvation thermodynamic profiles with respect to chain length (Fig. 1). Attempting to relate a thermodynamic mechanism to the structural properties of highly disordered polypeptides predicted by different force fields may be problematic.

WHY BLUE WATERS

To calculate solvation free energy, enthalpy, and entropy of our various oligoglycine models, we used a type of computational free energy method, termed "alchemical free energy perturbation (FEP)," that relies on molecular dynamics sampling in explicit solvent. While considered one of the most accurate methods to calculate free energy it is computationally demanding yet reasonably parallelizable. Our project required a total of 1,784 simulations for an aggregate simulated time of 92.3 microseconds. The high-throughput and ability to run a large number of simulations concurrently on Blue Waters enabled us to tackle such a large-scale project in a tractable amount of time.

PUBLICATIONS AND DATA SETS

Drake, J.A., and B.M. Pettitt, Solvation thermodynamics of oligoglycine with respect to chain length and flexibility: implication for aggregation and collapse. *Biophysical Journal*, (2016), dx.doi.org/10.1016/j.bpj.2016.07.13

Drake, J.A., and B.M., Pettitt, The protein backbone: how structural and thermodynamic properties scale with length. *251st American Chemical Society National Meeting and Exposition* (Philadelphia, PA, March 2016).

Drake, J.A., and B.M. Pettitt, The protein backbone: how structural and thermodynamic properties scale with length. *Sealy Center for Structural Biology and Molecular Biophysics 21st Annual Structural Biology Symposium* (Galveston, TX, April 2016).

Drake, J.A., and B.M. Pettitt, Protein backbone thermodynamics: Towards estimating conformational entropy. *Sealy Center for Structural Biology and Molecular Biophysics 21st Annual Structural Biology Symposium* (Galveston, TX, May 2015).

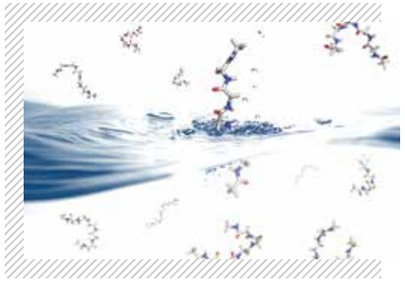

FIGURE 2: Artistic depiction of the physical process of the transfer of oligoglycine from the gas phase to aqueous solution (i.e. solvation thermodynamics).

Justin Drake is in his fifth year as a Ph.D. student in Biochemistry and Molecular Biology at University of Texas Medical Branch. After graduation, he would like to continue to use high-performance computing (HPC) to investigate complex biological systems either in academia or industry.

"Specifically, I am interested in continuing my research into the biophysical mechanisms that allow structural disorder to persist in proteins and how disease mutations alter the biophysical properties of disordered regions," says Drake. "The Blue Waters Graduate Fellowship has not only played a pivotal role in helping me achieve my doctorate but it has also exposed me to a variety of fields and topics related to HPC that I would have otherwise not experienced as part of my degree plan. I believe this breadth of knowledge will give me a competitive edge as I continue to the next stages of my career. In the future I also hope to further bridge the gap between biology and computational sciences and advocate how both fields can influence each other's further progression."

THE TRIBOELECTRIC CHARGING OF VOLCANIC MATERIALS

Joshua Mendez, Georgia Institute of Technology
2015-2016 Graduate Fellow

RESEARCH SUMMARY

When fluidized, insulating granular materials readily become electrified as individual grains undergo collisions. Although contact and frictional electrification processes—collectively termed *triboelectrification*—have been studied for millennia, the mechanisms that drive the exchange of charge between contacting surfaces remain obscure [1]. Perhaps the most dramatic demonstrations of granular charging are the brilliant displays of lightning during violent volcanic eruptions [2]. The industrial sector routinely monitors the hydrodynamics of granular materials non-intrusively using electrostatic probes, electric field mills, and capacitive tomography [3]. In principle, similar methodologies could be widely applied to volcanic systems. Indeed, the observed variability in electrical behavior between eruptions and between volcanoes suggests that the generation and separation of charge is modulated by specific eruption parameters, such as the properties of the ejected materials (both ash and volatiles), environmental conditions, and the energy of the eruption. Thus, understanding the coupling between eruption dynamics and electrical activity, much of which can be studied remotely, may yield information about the internal dynamics of an eruption that would otherwise be opaque to observation [4].

Most work on the electrification of volcanic plumes has focused on studying macro-scale effects, namely the detection and processing of lightning signals or changes to the ambient electric field. However, because triboelectrification is inherently a microscopic phenomenon involving the interaction of minute ash particles, developing robust tools to probe inside volcanic clouds requires that we build constitutive relationships that link often elusive, nanoscale charge transfer processes to evident meter- and kilometer-scale discharges and electric fields. While we have learned much about the changing behavior of ash from experimental efforts, determining how electrification is effected by the kinetics of the system (which include particle collision frequency, collisional energies, and aggregation rates) can still present a challenge without specialized equipment.

To complement our experimental work and answer outstanding questions, I developed a novel discrete element model (DEM) that includes a charging module and interparticle electrostatic force module. The model implements the trapped electron model first described by [5], to account for charging between chemically similar particles (Fig. 1). Specifically, this mechanism proposes that charging is driven by electrons trapped in unfavorable energy states on one surface that transfer into empty low-energy states on another surface when they come into contact. A particularity of this process is that if asymmetric contact is involved, such as small particles colliding with larger grains, the smaller particles become negative and the larger ones acquire a positive charge. As particles charge, their dynamics may become influenced by electrostatic forces. My model integrates a formulation for polarizable, dielectric spheres [6]. At large distances, the electrostatic force can be approximated by that between two point charges. However, at close ranges, the effects of mutual polarization can radically change the magnitude and sign of the forces between grains. Indeed, there are given combinations of particle size and charge that generate attractive forces even if the particles carry charge of the same sign (Fig. 2)! The model reproduces the bimodal charging behavior observed in a number of experiments [7]. Using the charging rates obtained from the DEM allows us to start exploring charging behavior in a large-scale Eularian-Eularian simulations of volcanic plumes, to reproduce the lightning storms present in large volcanic eruptions. As a case study, we will use the 2009 Redoubt eruption. Applicable to study other systems, my model is also being used to explore morphological peculiarities of the "plastic" sand dunes of Saturn's moon Titan [8], structures perhaps molded by electrostatic forces.

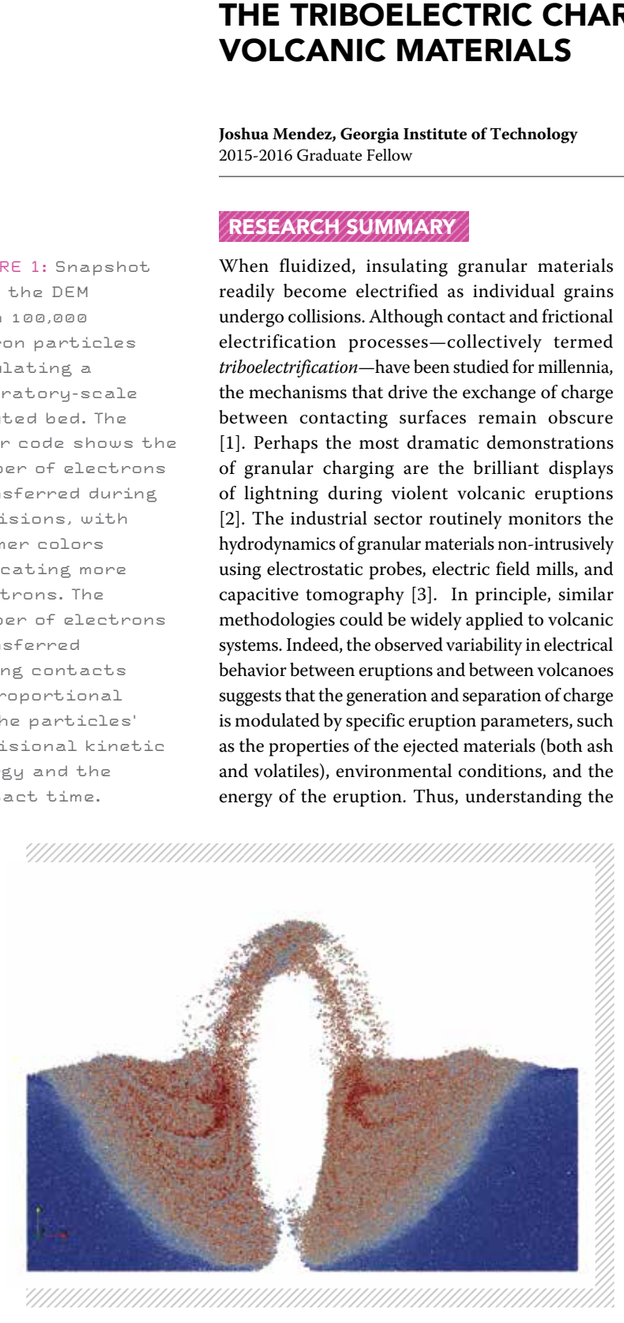

FIGURE 1: Snapshot from the DEM with 100,000 micron particles simulating a laboratory-scale spouted bed. The color code shows the number of electrons transferred during collisions, with warmer colors indicating more electrons. The number of electrons transferred during contacts is proportional to the particles' collisional kinetic energy and the contact time.

WHY BLUE WATERS

Discrete element models have the advantage that the interactions between particles can be resolved directly. However, even simulating a physically small, laboratory setup often involves tracking several hundred-thousands to millions of particles, exceeding the abilities of conventional desktop machines or small clusters. These simulations and would not have been viable without a machine like Blue Waters.

While I have done computing work during much of my academic career, most of the work prior to receiving the Blue Waters Fellowship involved the programming of embedded systems. One of the processors with which I am most familiar runs on a 1 Mhz clock, has little more than 1 kB of Flash memory and 32 bytes of ram, costs around $1, and is about 3 mm on a side. Starting to run code on one of the **world's most powerful** machines was intimidating at times. The Blue Waters staff, particularly my point-of-contact, rapidly helped to dissipate any uncertainty—resolving compilation issues or jumping in to improve the code's performance. As someone who is relatively new to HPC, I think they are one of the most valuable parts of this fellowship and have convinced me to use HPC resources for the rest of my career.

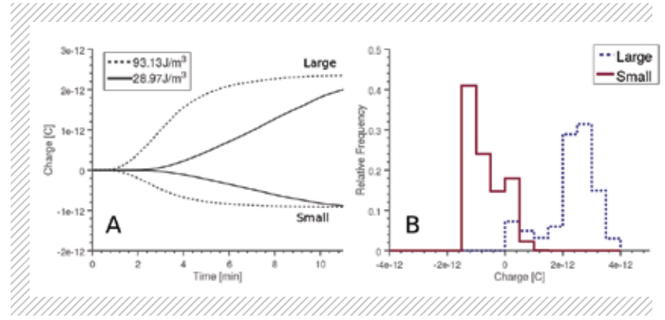

FIGURE 2: Output from the spouted bed simulation. Implementing the trapped electron model, the simulation reproduces the bi-polar charging behavior observed in many natural and man-made systems. Specifically, smaller particles are observed to charge negatively while larger ones charge positively.

PUBLICATIONS AND DATA SETS

Méndez Harper, J.S. and Josef Dufek (2016), The effects of granular dynamics on the triboelectric charging of volcanic ash: Experiments and numerical simulations. *Proceedings of the 2016 Meeting of the Electrostatics Society of America*, Purdue University, USA.

Joshua Méndez Harper is a fifth-year Ph.D. student at the Georgia Institute of Technology, working with advisor Josef Dufek. He enjoys "projects that straddle disciplines," and hopes to continue pursuing multidisciplinary work as either a faculty member at a public research university or at a national laboratory.

"This fellowship has not only brought key aspects of my research into different, clearer light, it has also guided me toward avenues that I considered inaccessible a year ago," he says. "Furthermore, the code I have developed during my fellowship has garnered interest among researchers beyond the confines of my primary field of study, opening new and exciting possibilities for future interdisciplinary work."

DISTRIBUTED ALGORITHMS FOR POWER SYSTEM MONITORING AND CONTROL

Ariana Minot, Harvard University
2015-2016 Graduate Fellow

RESEARCH SUMMARY

Tomorrow's electric power grid will have many components distributed over a large geographic area, necessitating distributed, scalable solutions to monitor and control the future grid. Our work on Blue Waters focuses on developing distributed algorithms for the security-constrained optimal power flow problem (SC-OPF). SC-OPF is a large-scale, nonlinear optimization problem that economically balances supply and demand under certain physical and operational constraints. It must be solved within five minutes or faster to operate the grid in a responsive manner to contingencies. The main computational burden is solving a series of large, sparse linear systems. In our work, we develop specialized solvers to exploit the sparsity pattern induced by the underlying network structure and physical laws of power networks. By using Blue Waters, we can push the boundary of the size of the largest SC-OPF problem that is solvable within required time limits.

Increasing power demands, aging infrastructure, and growth in renewable energy production necessitate new strategies for grid operations. To achieve reliable operations in this setting, it is important to shift away from centralized control paradigms to distributed approaches that 1) allow for quicker solution times by decomposing the problem to be solved in parallel and 2) avoid communication bottlenecks that result from all data and measurements being sent to a centralized location (Fig. 1). Since power systems are interconnected systems, decoupling such problems is challenging.

Our work on Blue Waters focuses on developing distributed second-order algorithms for the SC-OPF problem in power systems. SC-OPF is an optimization problem that seeks the amount of power to be generated from each power plant to minimize costs under certain physical and operational constraints. A fast, robust solution to SC-OPF remains an important challenge in power system operations [1].

SC-OPF is a large-scale, nonlinear optimization problem that ideally can be solved within five minutes or less to operate the grid in a responsive manner to contingencies (i.e. equipment failures). We explore the use of primal-dual interior point methods to solve SC-OPF. The main computational burden of this method is solving a series of large, sparse linear systems. In our work, we develop specialized solvers to exploit the sparsity pattern induced by the underlying network structure and physical laws of power networks (Fig. 2). In particular, we adapt domain decomposition methods, typically used for partial differential equation solvers, since our problem has a similar structure. Namely, different contingencies are loosely coupled only via the power generation decision variables. Furthermore, the physical laws governing the power grid (e.g. Kirchhoff's laws) can be written as a set of equations involving only local and neighboring quantities. By adapting domain decomposition techniques, we can parallelize across different contingencies and power decision variables the main computational bottleneck of our algorithm, which serially requires 98% of runtime. If grid operators adopt a parallelized strategy, important improvements in runtime can be made, which is critical for real-time operations.

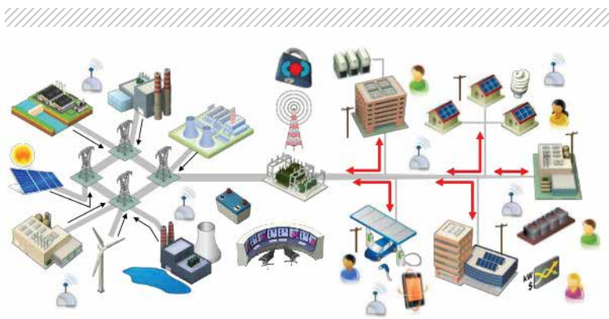

FIGURE 1: Tomorrow's power grid will have many components distributed over a large geographic area on both the supply and demand side. Rather than processing all information at the central control center, distributed solutions will provide scalable, large-scale coordination of the many active points in the grid. Image courtesy of Electric Power Research Institute (EPRI).

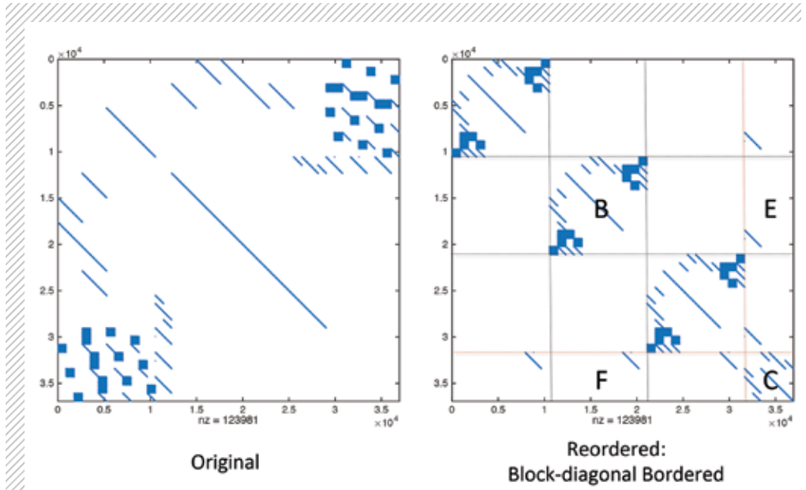

FIGURE 2: The structure of the system matrices involved in the security-constrained optimal power flow can be reordered so that it is amenable to domain decomposition techniques.

WHY BLUE WATERS

Working with Blue Waters allows us to measure the level of parallelization required to achieve runtimes needed for real-time operation (less than five minutes). Using Blue Waters, runtimes can be tested on realistic systems with hundreds to thousands of interconnections and with hundreds to thousands of contingencies. The speed of computation limits the number of contingencies, or fault scenarios that power grid operators can ensure against when solving for the optimal operating point since the result of the computation is needed within a short time-frame for operational use. By using Blue Waters, we can push the boundary of the size of the largest SC-OPF problem that is solvable within required time limits.

I would like to acknowledge the NCSA staff for their extensive help in building the GridPACK software [2] on Blue Waters. GridPACK is a software package developed by Pacific Northwest National Laboratory that provides a framework for developing power grid applications on HPC platforms.

PUBLICATIONS AND DATA SETS

Minot, A., N. Li, and Y. Lu. A Parallel Primal-Dual Interior Point Method for Optimal Power Flow. *Power Systems Computation Conference*, Genoa, Italy, June 20-24, 2016.

Ariana Minot is in her fifth year of Ph.D. studies in Applied Mathletics at Harvard University. She plans to graduate in May 2017 and hopes to continue advancing research at the intersection of power grid engineering and scientific computing.

"After graduation, I plan to pursue a research position at a national laboratory or industry laboratory related to the electric power grid," Minot says, "Blue Waters has given me the opportunity to develop algorithms on high performance computing (HPC) platforms and learn from the wider community of Blue Waters users."

PARTICLE ACCELERATION IN LASER-DRIVEN MAGNETIC RECONNECTION

Samuel Totorica, Stanford University
2015-2016 Graduate Fellow

RESEARCH SUMMARY

Magnetic reconnection is a fundamental plasma process that converts magnetic field energy into plasma energy through the breaking and rearrangement of magnetic field lines [1]. The energy released as the magnetic field changes topology drives plasma flows, heats the plasma, and accelerates particles to high energies. Reconnection is ubiquitous in magnetized plasmas, occurring in systems in space physics, astrophysics, and the laboratory. It is believed to play a key role in frontier problems in physics including the origin of cosmic rays and has important implications for applications with a societal benefit such as space weather and nuclear fusion energy.

Observations of high-energy astrophysical sources such as pulsar wind nebulae, gamma-ray bursts, and jets from active galactic nuclei often indicate rapid energy dissipation and efficient particle acceleration, and the environment of these systems is commonly thought to be a magnetized plasma. As an efficient mechanism for dissipating magnetic energy in a plasma, reconnection is a promising candidate for producing the non-thermal particle distributions inferred to be present in these systems. However, the efficiency of reconnection in accelerating non-thermal particles and its dependency on plasma conditions remains poorly understood, and it is currently an active area of research to determine whether reconnection can account for the astrophysical observations.

As a result of the national inertial confinement fusion program, high energy laser facilities have been developed that can ablate extremely hot and dense plasmas when focused onto solid targets. The conditions produced are in a regime where ohmic dissipation is negligible, enabling the use of scaling laws to study astrophysical phenomena. Experiments performed at these facilities study astrophysical processes, several of which have investigated magnetic reconnection. The expansion of two nearby ablated plasma bubbles can drive reconnection between self-generated or externally-imposed magnetic fields. Several features of reconnection have been observed in these experiments, however, none of these previous studies have addressed non-thermal particle acceleration, an important signature of reconnection for connecting to systems in astrophysics.

The goal of our research is to use first principles particle-in-cell (PIC) simulations to investigate particle acceleration from magnetic reconnection can be studied in laser-driven plasma experiments using the state-of-the-art PIC code OSIRIS [2]. The results of our simulations clearly indicate that laboratory experiments can play an important role in the study of particle acceleration by reconnection. For current experimental conditions, we show that non-thermal electrons can be accelerated to energies up to two orders of magnitude larger than the initial thermal energy (Fig. 1). The non-thermal electrons gain energy primarily by direct acceleration from the reconnection electric field near the X-points (Fig. 2), while particle injection into the reconnection layer and escape from the finite system establishes a distribution of energies that resembles a power-law spectrum. Energetic electrons can also become trapped inside the plasmoids (magnetic islands) that form in the current layer and gain additional energy from the electric field arising from the motion of the plasmoid. Based on our findings, we provide an analytical estimate of the maximum electron energy and threshold condition for observing suprathermal electron acceleration in terms of experimentally tunable parameters. These results pave the way for laser-driven plasmas as a new platform for the experimental study of particle acceleration induced by reconnection, which could help illuminate the role reconnection plays in explosive phenomena associated with space and astrophysical plasmas.

Also, we are developing a **novel** algorithm that holds potential for reducing noise and computational expense in PIC simulations. Interpreting the simulation particles as the vertices of an unstructured mesh that traces the evolution of the distribution function in phase space enables a discretization using deformable phase space volume elements rather than fixed-shape particles [3]. This new perspective retains fine-scale structure in the distribution function and may reduce the number of simulation particles required. We are currently using this method as a novel post-processing technique (Fig. 3), and future work will involve implementing it directly in PIC simulations to reduce noise and unphysical artifacts.

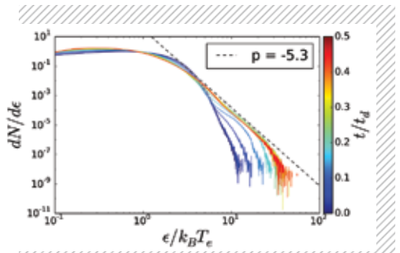

FIGURE 1: Temporal evolution of the electron energy spectrum for a simulation with typical laboratory conditions.

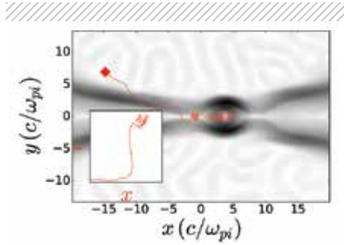

FIGURE 2: Example trajectory of an energetic electron over the magnitude of the magnetic field. Inset shows the energy of the electron as a function of position.

WHY BLUE WATERS

Simulations used for quantitative comparison with laser-driven plasma experiments must bridge the multiscale physics, from fluid dynamics to kinetic microscopic processes. These computationally demanding simulations can only be performed using the cores, memory, and communication performance of large-scale resources like Blue Waters. The fast response time of the NCSA staff to technical issues allowed for maximum productivity on the machine. I plan to continue to use high power computing systems to study problems in plasma physics and develop improved plasma simulation methods with an eye towards next generation machines.

PUBLICATIONS AND DATA SETS

Totorica, S. R., T. Abel, and F. Fiuza, Nonthermal electron energization from magnetic reconnection in laser-driven plasmas. Phys. Rev. Lett., 116 (2016), p. 095003.

Samuel Totorica is in his fourth year of the Physics Ph.D. program at Stanford University. He plans to graduate in spring 2018 with the goal of continuing research in computational plasma physics in a national research center.

"[In such a setting] I could have access to large-scale computer resources, contribute to projects with societal benefit, and interact closely with experimentalists," says Totorica, "I hope to continue to develop novel plasma simulation methods, and apply them to support nuclear fusion efforts as well as explore astrophysical scenarios and how they can be tested in the laboratory. Access to Blue Waters has allowed me to gain experience and proficiency in performing large scale plasma simulations at an early stage in my career, and the high impact research it enabled me to perform has helped me to establish myself in the field of plasma physics."

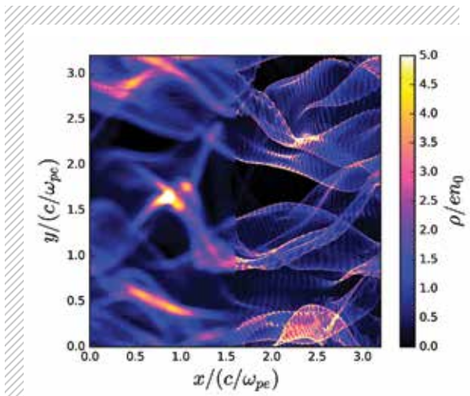

FIGURE 3: Charge density from a 2D simulation of the Weibel instability, rendered using standard CIC particles (left) and the new method (right).

REDUCING THE COMPUTATIONAL COST OF COUPLED CLUSTERY THEORY

Sara Kokkila Schumacher, Stanford University
2015-2016 Graduate Fellow

RESEARCH SUMMARY

Electronic structure methods are capable of describing the electron arrangement in molecules. In order to accurately describe the absorption spectra of molecules, it is necessary to use electronic structure methods. Unfortunately, few electronic structure methods are capable of efficiently calculating the excited states of molecules. The equations-of-motion second-order approximate coupled-cluster singles and doubles (EOM-CC2) method is an accurate electronic structure method for studying the absorption of light by molecules. [1,2] However, this approach costs $O(N^5)$ where N is the number of basis functions used to represent a molecule. In order to improve the efficiency of this method, the tensor hypercontraction (THC) approximation is introduced to reduce the computational cost to $O(N^4)$. [3, 4, 5]

Graphical processing units (GPUs) have proven to be useful for accelerating quantum chemistry codes. [6] In order to further improve the efficiency of THC-EOM-CC2, we developed a GPU accelerated version of the grid based THC-EOM-CC2. Unlike other GPU implementations of coupled-cluster theory, which often focus on using GPUs only to accelerate the most computationally demanding step, the algebraic flexibility of THC can be used to form contributions to THC-EOM-CC2 "on-the-fly" on the GPU. Because the THC formation allows us to represent high order tensors as a product of lower order tensors, there are many alternative ways to implement THC-CC2 and THC-EOM-CC2. In order to parallelize these methods, we developed THC-CC2 and THC-EOM-CC2 algorithms that focus on blocking over a grid point index for a series of intermediates. This allows for parallelization at multiple levels. Figure 1 demonstrates the general parallelization scheme for THC-CC2 where Parts 1-5 each represent a series of intermediates formed through Compute Unified Device Architecture (CUDA) kernels and NVIDIA CUDA Basic Linear Algebra Subroutines (cuBLAS). The THC-EOM-CC2 method can be parallelized in a similar manner.

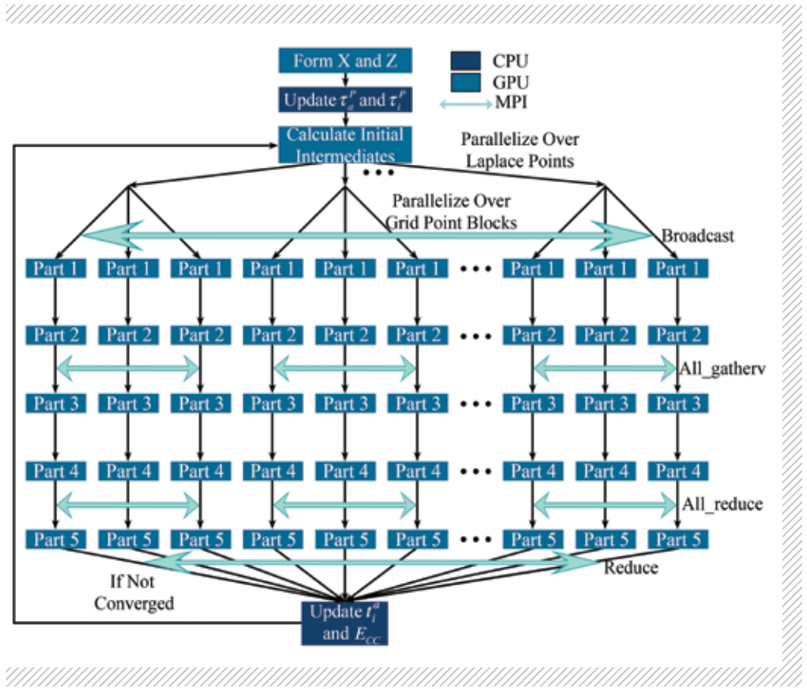

FIGURE 1: Parallelization outline for the GPU-based THC-CC2 algorithm. The five main parts represent a series of linear algebra and tensor algebra routines performed on GPUs.

Timings for the GPU accelerated THC-CC2 and THC-EOM-CC2 approaches are shown in Figure 1. This shows that the recently developed GPU accelerated algorithms do achieve $O(N^4)$ cost for both the ground state calculation (THC-CC2) and the excited state calculations (THC-EOM-CC2). These timings were computed on a single GPU on a single node of Blue Waters.

We have already shown that the THC approach can be used to reduce the computational cost of different electronic structure methods. This work is the first to demonstrate how one can take advantage of the algebraic flexibility offered by THC to develop massively parallel GPU accelerated THC based electronic structure methods. This approach can be used to extend the applicability of electronic structure methods to larger chemical systems and to rapidly compute energies of sampled geometries from dynamics simulations. We expect this to lead to the development of other THC based electronic structure methods in different quantum chemistry software packages.

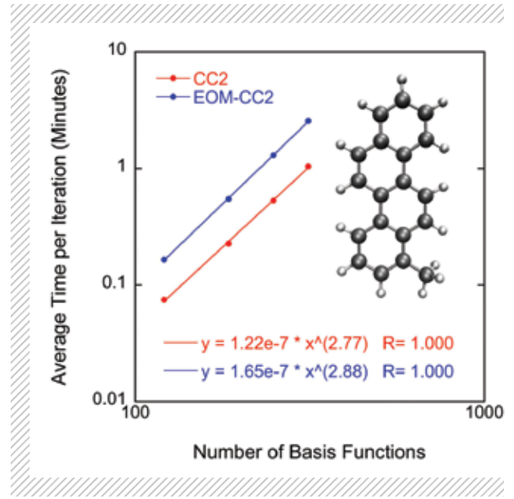

FIGURE 2: Timings of THC-CC2 and THC-EOM-CC2 iterations for a series of methylated polycyclic aromatic hydrocarbons. These timings were performed on a single node of Blue Waters. Each computation was performed in triplicate and the averaged timings are shown here.

WHY BLUE WATERS

Access to Blue Waters allowed for the development of the MPI enabled GPU accelerated THC-EOM-CC2 approach. With Blue Waters, we were able to rapidly develop and test different THC-EOM-CC2 implementations. As our first foray into combining MPI and GPU parallelization approaches, the Blue Waters project staff offered helpful advice during the design process. Additionally, we are using the THC-EOM-CC2 code to study the absorption of light by different protein chromophores. This requires sampling many different geometries and evaluating the THC-EOM-CC2 energies. The use of Blue Waters is necessary to compute enough geometries, ground state energies, and excited state energies to accurately describe the absorption spectra of protein chromophores.

PUBLICATIONS AND DATA SETS

Kokkila Schumacher, S.I.L, E. G. Hohenstein, R. M. Parrish, L.-P. Wang, and T. J. Martínez, Tensor Hpercontraction Second-Order Møller-Plesset Perturbation Theory: Grid Optimization and Reaction Energies, *J. Chem. Theory Comput.*, 11, pp. 3042-3052 (2015).

Sara Kokkila Schumacher graduated fall 2016 from Stanford with a Ph.D. in Chemistry, where she worked with advisor Todd Martinez. She is currently a postdoctoral researcher in high performance computing at IBM. She ultimately aims to become a professor where she can develop a course that introduces high-performance computing to science majors.

ENERGETIC DYNAMICS OF A ROTATING HORIZONTAL CONVECTION MODEL OF THE SOUTHERN OCEAN WITH SURFACE BUOYANCY AND WIND FORCING

Varvara Zemskova, University of North Carolina at Chapel Hill
2015-2016 Graduate Fellow

RESEARCH SUMMARY

The global ocean is on average stably stratified, which makes vertical mixing play an important role in distributing heat and nutrients. Vertical mixing is essential because it supplies dissolved oxygen from the surface to the deep waters and brings nutrients that result from bacterial decomposition back to the surface, where they can be used by organisms. If the ocean-atmosphere interactions change, from both directly from heat exchange and indirectly from of winds at the ocean surface, the amount of vertical mixing in the ocean will be affected. Quantifying the changes in vertical mixing is one of the important questions in oceanography today.

The Southern Ocean, part of the global ocean between the Antarctic continent and 60°S latitude, is an ideal place to study the relative contributions of differential surface heating and surface wind stress to ocean mixing. The Antarctic Circumpolar Current, which is driven by eastward winds, goes around Antarctica unobstructed by continents, unlike the currents in other ocean basins. Because of such predominantly unidirectional flow, it can be modeled in a periodic domain that decreases computational demand without sacrificing completeness of the problem. This region also experiences significant differential heating and cooling at the surface; dense water formation occurs near the Antarctic continent where the surface waters are cooled, while surface waters in the mid-latitudes receive more heat and become less dense. This problem can be studied in the form of horizontal convection, a model of a fluid flow forced by differential surface buoyancy forcing. The simulations are run on Blue Waters using SOMAR [1], which is a finite element method adaptive mesh refinement code that solves incompressible Navier-Stokes equations with Boussinesq approximation and added rotation rate consistent with the Southern Hemisphere.

For this study, several combinations of surface and wind boundary conditions and rough bottom topography conditions are considered. All cases are forced to the top surface with a prescribed density distribution matching the average meridional surface density distribution in the Southern Ocean. The simulation with only buoyancy forcing is taken as a base state, and simulations with different magnitude of wind stresses and rough or smooth bottom topography are compared to the base state. All simulations are run to a steady state before terms in the energy cycles, such as the generation of kinetic and potential energies and transfers between energy reservoirs, and vertical mixing can be quantified. From the preliminary results, even a small magnitude of wind forcing shifts the energy paradigm of the simulation into an energetic regime observed in the ocean, in which the amount of vertical mixing can be estimated from the amount of surface generation of available potential energy due to differential buoyancy forcing [2]. The rough topography enhances the amount of vertical mixing compared with the base state; however, the competing effect of wind forcing and bottom topography is still to be investigated.

WHY BLUE WATERS

The computational resources of Blue Waters were essential for this research because I was able to run simulations that have sufficiently large resolution and aspect ratio of the domain to capture forward and inverse energy cascades. Such simulations are very computationally costly, and the allocation on Blue Waters allows me to compare results from multiple simulations with varying boundary conditions. In addition to having many mesh points to allow for sufficient resolution, the simulations must also be run for months before converging to a statistical steady state. It is particularly important for the simulations with rough bottom topography where grid cells become irregularly shaped and additional resolution is required. While for the calculations of the energy budget I only need to keep the files for the last few time steps, intermediate time steps are also necessary for determining whether the simulation has converged and for making videos of the evolution of the simulation. The large storage space allocated by the Blue Waters is essential for making this part of the project feasible for running multiple simulations at the same time. With my continuing allocation I plan to extend this project to analyze the sensitivity of mixing and energetics in the Southern Ocean to changing boundary conditions that either correspond with estimated past climate conditions or future climate predictions.

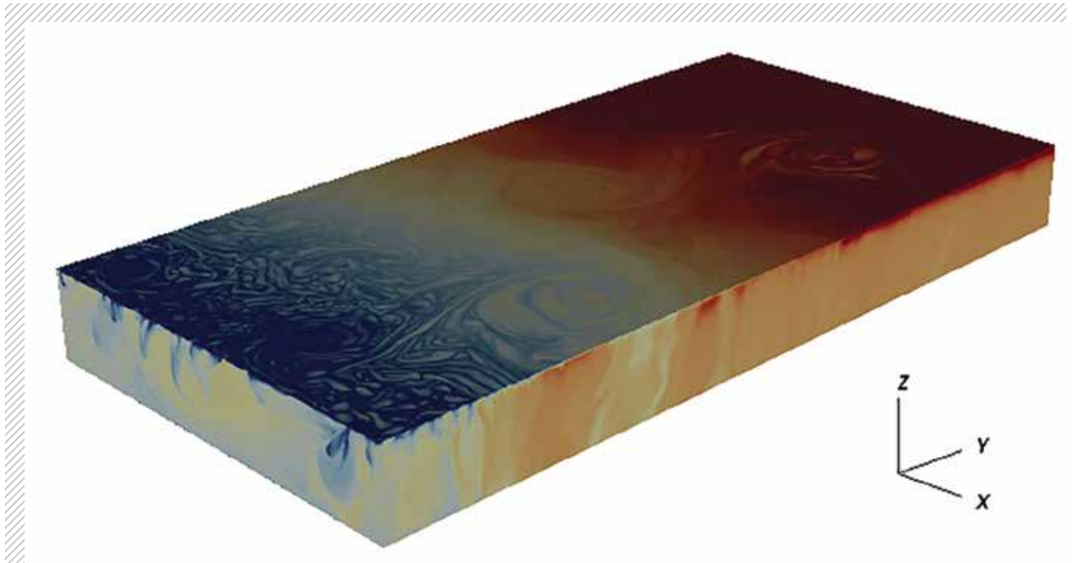

FIGURE 1: The density distribution for the simulation with realistic wind and buoyancy forcing in the Southern Ocean (north pointed into the page). Blue colors indicate cooler, denser fluids, which are sinking near the pole and red colors are warmer, less dense waters. Many eddies and complex vertical structures can be observed.

Vervara Zemskova works with advisors Brian White and Alberto Scotti and expects to complete her Ph.D. in geoscience from the University of North Carolina at Chapel Hill in May 2018. Upon graduation, she would like to continue research in an academic setting and combine her background in oceanography with the experience in supercomputing.

"My primary interest lies in small-scale physical oceanography and how such small scale mechanics affect the biological aspects in the ocean. Because continuous ship measurements of small scale dynamics are both difficult and expensive, my main focus is understanding these dynamics through numerical modeling."

HIGH ACCURACY RADIATIVE TRANSFER IN CLOUDY ATMOSPHERES

Alexandra L. Jones, University of Illinois at Urbana-Champaign
2014-2015 Fellow

RESEARCH SUMMARY

One of the most important roles clouds play in the atmosphere is in redistributing radiative energy from the sun. Given the ubiquity of cloud coverage, it is imperative that we get the interactions between clouds and radiation correct if we want to accurately predict and observe weather and climate. However, radiative transfer in the atmospheric sciences is generally modeled crudely because of the perceived computational expense. Evidence of a bias due to these crude assumptions has been seen in observed properties from satellites as well as modeled cloud properties.

A model that treats broadband integration and 3-D radiative transfer in a highly accurate and unbiased way is needed to quantify the bias in the simpler models ubiquitously used. This model will serve as a previously nonexistent standard of comparison for other similar models, and provide accuracy bounds for simpler models and parameterizations attempting to capture 3-D effects at lower computational cost. Such a model was not publicly available prior to this project. So, one was developed that uses Monte Carlo methods to capture the 3-D transfer of radiation and sample at high resolution the broad range of the electromagnetic spectrum. Unlike the direct approach to solving the radiative transfer equation, the Monte Carlo approach has the potential to be embarrassingly parallel since the random samples are independent from one another. The overarching 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.

A monochromatic, 3-D Monte Carlo community solar radiative transfer model was further developed to include terrestrial emissions in addition to solar sources of radiation. That model was then further developed to include integration over the electromagnetic spectrum to produce the broadband 3-D model discussed above. In addition to these two models, several other products have resulted so far and will be made available to the community. This includes databases of high-spectral resolution, radiative properties of earth's gaseous atmosphere and liquid water clouds. These are the largest and highest resolution publicly available databases of their kind. The tools and workflow to create and subset them will also be made available. This data can be mined to update the decades old broadband parameterizations of cloud radiative properties that are still in wide use today, for example. Each product has been thoroughly vetted for accuracy. The results of these tests will be made available for reproduction by other scientists to test these models or their own. Finally, the first few idealized experiments with long heritage in the literature have been conducted to provide the **first set of benchmark simulation results** that can be used to evaluate other models.

The results shown in the figure demonstrate the impact accounting for 3-D radiative transfer has on the heating rates within a cubic cloud in a broadband simulation of thermal emissions from a cloud and a black surface. The left panel shows a cross section through a plane parallel cloud, which in this context means radiation cannot enter or escape through the sides of the cloud, only the top or bottom and therefore can make no contribution to heating or cooling of the cloud edges. The right panel shows a cross section through a cloud with the same physical properties but a finite horizontal extent, meaning that radiation is free to enter or leave through the cloud sides and can contribute to cloud side heating or cooling. The heating rates at the top and bottom of the plane parallel cloud are within 0.025 K/day of the heating rates at the horizontal center of the finite cloud. However, the cooling at the finite cloud's edges is unaccounted for in the simulation of the plane parallel cloud. In time-integrated simulations of cloudy atmospheres, this cooling that occurs when radiation is allowed to exit through cloud sides is unaccounted for in the temperature at cloud edge, which could ultimately impact the evolution of modeled-cloud physical properties and cloud scale dynamics.

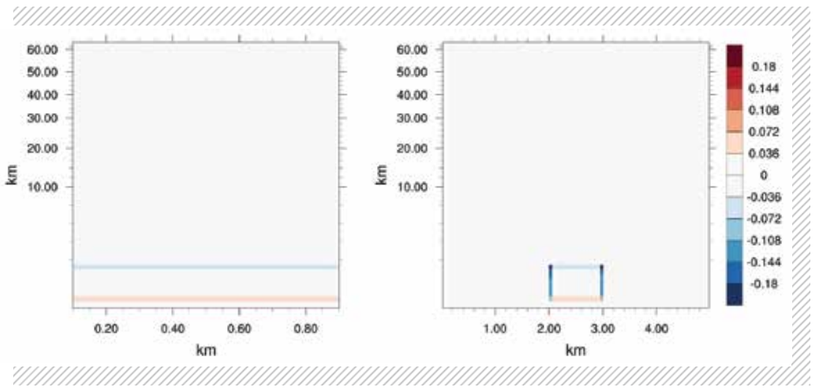

FIGURE 1: Vertical cross section of long wave heating rate K day^{-1} through the center of an isothermal (280 K), homogeneous (effective radius=8.8 microns, liquid water content=0.32583 gm^{-3}) cloud set in a vacuum above a non-scattering surface emitting as a blackbody according to its temperature (290.38 K). The left panel shows a plane parallel cloud whose horizontal extent is effectively infinite and vertical extent is 1km. The right panel shows a 1 km x 1km x 1km cloud set in a 5 km x 5 km x 66 km domain.

WHY BLUE WATERS

Access to debugging and profiling tools such as CrayPat and DDT allowed the development process to progress in a streamlined fashion. Access to a point of contact at SEAS helped me find tailored solutions for problems that otherwise would have delayed progress by weeks. The quick responsiveness of the Blue Waters staff through the JIRA ticket system allowed for limited interruption in progress when small issues or questions arose. My experience as a Blue Waters graduate fellow has been invaluable to my professional development. I hope to make use of Blue Waters for the rest of its lifetime.

Alexandra L. Jones completed her Ph.D. in Atmospheric Science at the University of Illinois at Urbana-Champaign in January 2016. She is currently a postdoctoral researcher in climate science at the Cooperative Institute for Climate Science, a collaboration between Princeton University and the National Oceanographic and Atmospheric Administration's (NOAA) Geophysical Fluid Dynamics Laboratory (GFDL).

"My work on Blue Waters allowed me to hit the ground running, utilizing GFDL's access to another Cray supercomputer, Gaea. I can see how my computational science knowledge and skills, gained during my time as a Blue Waters Graduate Fellow, are enabling me to make faster progress," she says. "My work parallelizing and making the scientific software needed more efficient will enable me to expand the scope of the project and the amount of data I can include as my career continues, I would like to remain at the cutting edge of high performance computing and atmospheric science, whether as a researcher or an advocate."

RIGOROUS QUANTUM-CLASSICAL SIMULATION OF ELECTRON TRANSFER IN A BACTERIAL PHOTOSYNTHETIC REACTION CENTER

Thomas Allen, University of Illinois at Urbana-Champaign
2015-2016 Graduate Fellow

RESEARCH SUMMARY

Photosynthesis forms one of the core biological energy cycles responsible for the maintenance of life on Earth and has existed in some form since the earliest days of evolution, persisting today in bacteria, algae, and green plants. The molecular apparatus of photosynthetic energy conversion exhibits many similarities between different prokaryotic families, and even plant photosystems share some common motifs with the reaction centers of purple bacteria [1]. Moreover, the efficiency of these systems in trapping and converting light to other forms of energy is extremely high and has served to inspire human attempts to collect and store solar energy [2]. As a result, an improved understanding of the processes involved in natural photosynthesis has implications for the past, in understanding the origin and development of life on Earth, as well as the potential to inform the future with a focus on clean, renewable energy sources.

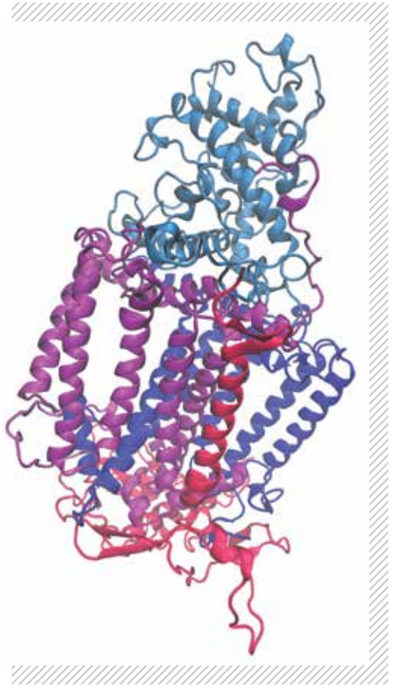

FIGURE 1: A visualization of the protein backbone of the reaction center in *Bcl. viridis*. The L, M, and H chains are shown in dark blue, purple, and red, respectively. The cytochrome unit is shown in light blue at the top.

Theoretical and experimental progress in understanding and characterizing the molecular processes of photosynthesis has the potential for major impact on the development of clean energy technologies. However, aspects of the theoretical side of this work are hampered by the fact that the electron dynamics induced as part of the photon capture event require a quantum mechanical description for accurate treatment. It also requires a fully-quantum description of the surrounding reaction center protein, and solvent molecules would be too computationally expensive to complete. Several approaches have been developed to work around this issue, but the simplest have major shortcomings when applied to chemical problems [3], and others introduce *ad hoc* elements which limit their rigor and accuracy. Recently, progress has been made on a rigorous approach to decoupling the study of a quantum mechanical system from its surrounding classical environment, in the form of the quantum-classical path integral (QCPI) [4,5], and this method is now computationally efficient enough to apply to biological problems.

We are in the process of applying a highly parallel version of our QCPI code, coupled to the LAMMPS molecular dynamics package [6] to simulate the first steps of electron transfer in the photosynthetic reaction center of the purple bacterium *Blastochloris viridis*. Our initial results from the molecular dynamics compare well to previous studies [7,8], but extend them significantly in that the behavior of the entire protein complex is included in the dynamics, and the simulation takes place at fully-atomistic resolution. The output of this work will have promise not only as a source of insight into the molecular mechanisms of biological charge transfer, but also as the first direct test of the linear response approximation in a protein system. It will be a landmark on the path toward truly rigorous simulation methods of large biochemical systems.

WHY BLUE WATERS

The expert staff and extensive computational resource available through Blue Waters was essential in this research program. The staff helped to convert the old parameterization of reaction center cofactors by Treutlein *et al.* [9] into a more modern format. They provided additional assistance with improving these parameters where our more complete dynamical simulations indicated potential problems, increasing our confidence in the results of the quantum-classical calculations.

From a computational standpoint, the high degree of parallelism and the large per-node memory available on Blue Waters is critical in performing our calculations within a reasonable time frame. This is because the QCPI method requires not only a path summation but also an average over many initial conditions. Together, these represent significant computational expense not feasible in a serial approach. However, the QCPI formalism is amenable to a message packing interface (MPI)+MPI decomposition which dramatically reduces running time and exhibits very favorable parallel scaling, making the approach feasible for atomistic simulation of realistic systems provided many cores are available. Additionally, the large memory capabilities available on Blue Waters ensure that all of the dynamical information required by each node can be stored locally, saving costly disk access and providing additional computational savings.

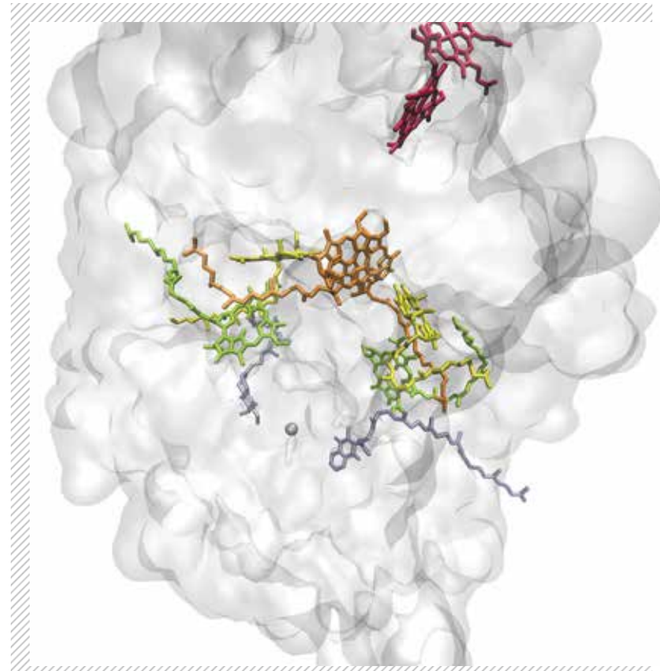

FIGURE 2: A close-up of the cofactors in the L and M chain region where the first stages of electron transfer occur. The most important participants in these early stages are the special pair of bacteriochlorophylls (shown in orange), the accessory bacteriochlorophylls (yellow), and the bacteriopheophytins (green).

Thomas Allen intends to continue pursuing research in a postdoctoral position upon completion of the University of Illinois Urbana-Champaign Chemical Physics Ph.D. program.

"Based on the charge transfer work I have performed using Blue Waters, I have been offered a position with Peter Rossky at Rice University, " he says. "The research program there is focused on the understanding and application of charge transfer models and mixed quantum-classical simulation to problems in materials science, with an emphasis on solar energy conversion. I will be starting at this position in Fall 2016, and I look forward to continuing to pursue the research threads which Blue Waters has made possible for me to study."

PROCESSING TRILLION EDGE GRAPHS IN DISTRIBUTED MEMORY

George M. Slota, The Pennsylvania State University
2014-2015 Graduate Fellow

RESEARCH SUMMARY

Graphs are a mathematical construct representing discrete entities or objects (vertices) and some form of link or interaction between them (edges). Human social interactions, the Internet, and the neural structures of the brain are just a subset of modern, real-world and extreme-scale datasets that are representable as graphs. The irregularity, scale, and complexity of such graphs present a high level of challenge for domain experts to use computational resources to extract useful insight into these graphs. As a result, many graph processing frameworks have recently been introduced with the goal to simplify the analysis of real-world graphs on commodity hardware. However, these popular frameworks lack scalability to modern massive-scale datasets, require specialized hardware to run, or simply cannot outperform optimal serial code [1]. Our previous work focused on overcoming these barriers from the ground up, developing general approaches for graph analytic optimization that are highly performant from a single node to a small cluster to a large and powerful system such as Blue Waters. There has been other prior research aimed at abstracting graph algorithms themselves, including as linear algebraic operations [2] or into a nested loop structure [3]. As such, our ongoing work has identified several other key data storage and communication abstractions that allow the straightforward implementation of broad classes of graph analytics. These findings will enable domain scientists to study graphs at a larger scale and with more complex algorithms than has previously been possible.

Our current work considers graph algorithms running on distributed systems such as Blue Waters and processing on an in-memory distributed graph representation. Here, each MPI task owns only some subset of a large graph, iteratively performing computation on its subset and communicating the results to neighboring tasks. By analyzing the computation and communication patterns typical of a number of graph algorithms, we noticed several generalizable similarities. There are two primary ways in which algorithm-specific data is updated. Either these updates are pulled by a given vertex to update some data associated with this vertex or these updates are pushed by a given vertex for its neighbors to use while computing updates to their data. In addition to this push/pull pattern, there is also a difference in the sizes of the update and work sets, where these sets can be composed of vertices for each iteration. The sets are either variable or fixed, depending on the algorithm. Overall, we identified four distinct processing patterns into which many graph algorithms fall (push-variable, push-fixed, pull-variable, pull-fixed) and created optimized outlines for these patterns. Using these patterns, we then implemented several common graph analytic algorithms. We observe high performance of our implementations, despite the generalized approach. Our implementations are noted to scale to graphs of over a trillion edges and over ten billion vertices while running on up to 8,192 nodes and 131,000 cores of Blue Waters. We can end-to-end process graphs of this scale in mere minutes on Blue Waters, including I/O, preprocessing, and output. Ongoing work aims to simplify our approach and methods further to enable its use among domain experts and graph analysts.

George Slota accepted a temporary staff position at Sandia National Labs upon graduation in May 2016. In Fall 2016 he started as an assistant professor in the Computer Science Department at Rensselaer Polytechnic Institute.

"The Blue Waters program has helped me achieve my career goals," he says, "by allowing me to focus on my research during my Ph.D. as well as enabling my access to a large-scale system that my research required."

WHY BLUE WATERS

Access to the Blue Waters systems greatly benefited this research for several reasons. The high performance I/O filesystem greatly accelerated end-to-end processing times for experiments on our large test datasets. We often observed over 50GB/s read time across a moderate number of nodes, even during periods of high concurrent usage, which enabled test data of terabytes in size to be read in seconds. This minimized node-hour waste when running high numbers of parametric tests with quick turnaround times. Blue Waters also represents the state-of-the-art in overall intra-node and inter-node communication and computational performance, which allows it to serve as a good representable testbed for generalizing our methods to be run on other current and future systems.

PUBLICATIONS AND DATA SETS

Slota, G. M., S. Rajamanickam, and K. Madduri, A Case Study of Complex Graph Analysis in Distributed Memeory: Implementation and Optimization. *Proceedings of the 30th IEEE International Parallel and Distributed Processing Symposium* (IPDPS16), (IEEE, Chicago, Illinois, May 23-27, 2016).

Slota, G. M., and K. Madduri, Parallel Color-coding. *Parallel Computing*, 47 (2015), pp. 51-69.

Slota, G. M., S. Rajamanickam, and K. Madduri, High Performance Graph Analytics on Manycore Processors. *Proceedings of the 29th IEEE International Parallel and Distributed Processing Symposium* (IPDPS15), (IEEE, Hyderabad, India, May 25-29, 2015).

SENSITIVITY OF SIMULATED URBAN-ATMOSPHERE INTERACTIONS IN OKLAHOMA CITY TO URBAN PARAMETERIZATION

Larissa Reames, University of Oklahoma
2015-2016 Graduate Fellow

RESEARCH SUMMARY

Earth's population is becoming increasingly concentrated in urban areas, with nearly two-thirds of the world's population expected to live in such areas by 2050. As the number of people within cities grows, it is becoming more important to understand, and to be able to predict correctly, the interactions between urban environments and the atmosphere. Many previous investigations have used numerical simulations alongside observations to investigate how and why land-atmosphere interactions differ from rural to urban areas. However, most of these simulations were performed over large, dense urban areas such as Tokyo, Beijing, or New York City. Hence, is not clear if the parameters used to represent urban areas in weather prediction models can appropriately handle cities that are largely suburban, like many of those in the Great Plains of the United States. Cities can a have a significant effect on precipitation distribution [1], lightning intensity [2], winter weather frequency [3], and flooding [4], and proper simulation of all urban areas is critical to proper modeling (and thus prediction) of various forms of hazardous weather.

Using the Weather Research and Forecasting (WRF) [5], a community mesoscale numerical weather prediction model, Reames investigated urban-atmosphere interactions in Oklahoma City, OK—a city typical of the Great Plains. The simulation was run for seven days and to properly resolve complex urban structure, all simulations were run on a 500-m horizontal grid over a 200-km x 200-km grid to measure average model performance. Also, to resolve the atmospheric boundary layer, 120 vertical grid points were used, with 20 in the lowest 1.5 km above ground. In all, 19.2 million points were integrated over 604,800 time steps for each simulation. One simulation with no urban area was produced to serve as a comparison point for the other urban simulations. To compare against the control run, one run using a simple land surface model to parameterize the urban surface was performed. Finally, two simulations using the more complex single-layer urban canopy model (SLUCM) parameterization to represent urban areas, each with different urban densities, were run.

The results from these simulations suggest that the land surface model is more appropriate for parameterization of cities with a large portion of suburban area. Observations suggest that wind speeds in the urban area are 20 to 30% slower than in surrounding rural areas, likely due to obstruction by buildings. While the simple land surface model reproduces this feature well, both urban canopy model runs do not. Also, when the urban canopy model was provided with a more realistic representation of the urban density of Oklahoma City, resulting in less dense urban areas over most of the city, the results were less accurate when compared to observations. In particular, while the nighttime urban-rural temperature difference (i.e. the "urban heat island") produced by the denser SLUCM was somewhat lower than observations, the less dense SLUCM run had a much less intense, and hence even less realistic, nocturnal urban heat island.

WHY BLUE WATERS

The scale of these simulations made access to Blue Waters critical. The domain size and time scale of each simulation would have made these simulations impractical using other accessible resources. The next step of this research will be to simulate the interaction of plains cities with supercellular thunderstorms. This future endeavor will involve over 150 simulations on scales similar to those described here, making future access to Blue Waters vital.

Larissa Reames is in her fourth year of the Meteorology Ph.D. program at the University of Oklahoma and plans to graduate in December 2016. Upon completion, she will begin a postdoctoral research position at the National Oceanic and Atmospheric Administration's National Severe Storms Laboratory where she will perform numerical modeling and model analysis.

"This type of position is the kind I would like to hold for the rest of my career," Reames says, "My work on Blue Waters involving numerical simulations and data analysis have allowed me to showcase my scientific and computational abilities, both of which are critical for holding research positions involving numerical modeling."

TRANSFORMATIVE SCIENCE

The Blue Waters project provides the national open research community with the computational power and data analysis capability that have become increasingly vital to virtually all areas of inquiry and discovery. With the power of the Blue Waters integrated computing and data system, researchers are pushing the boundaries of their disciplines and achieving breakthrough results in a wide range of fields, including particle physics, molecular bioscience, astronomy and astrophysics, earth sciences, and materials research. Figure 1 shows the diverse disciplines using the Blue Waters system between June 2015-May 2016 (Project Year 3 – PY3). In PY3, 1,269 members of 338 projects ran over 2.7 million jobs consuming over five billion core hour equivalents on Blue Waters. Since Blue Waters has different types of nodes (XE and XK), we use the term "core hour equivalents" to help compare node hours to core hours, which may be more familiar. There are two ways to use the AMD Interlagos process, one mode is as 16 floating point cores, each with one integer core. The other mode is as 32 integer cores sharing 16 floating point cores. We define core hour equivalents to be the number of node hours multiplied by 32. The average job size, based on node hours consumed, utilized 768 nodes or over 24,000 integer core equivalents. In addition to the compute power of Blue Waters there is an incredible data analysis capability evidenced by the ability of science teams to read and write in excess of 10PBs of data in a 24-hour time period multiple times during the year.

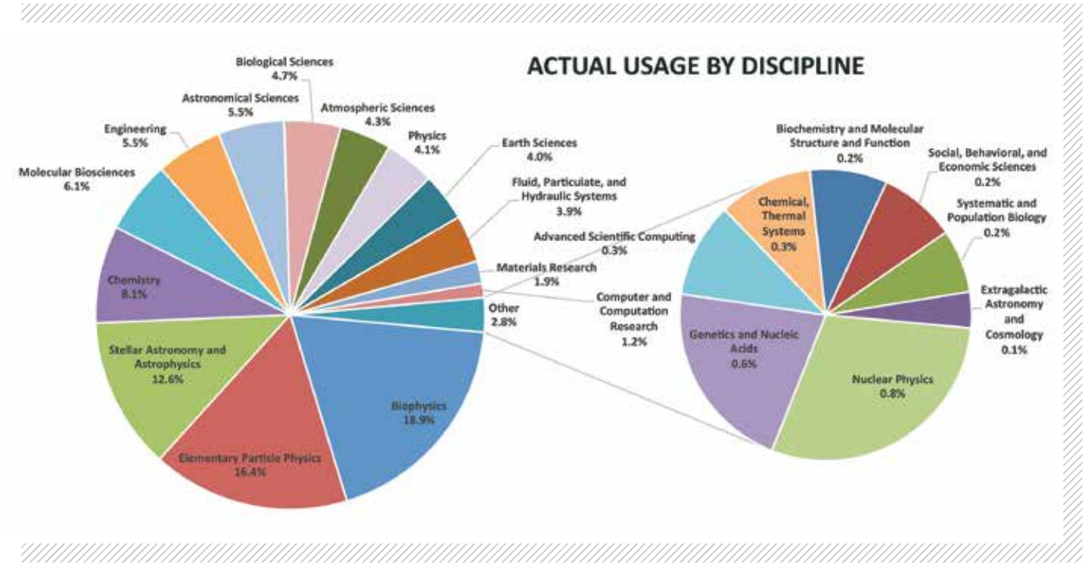

FIGURE 1: Distribution of annual actual usage by discipline June 2015-May 2016.

From the beginning of the project, the primary goal of Blue Waters has been to enable transformative science and engineering research by giving research teams access to a large number of powerful compute nodes and large data capacity. The target is to have at least 40 percent of the system's capacity dedicated to research challenges that cannot be addressed on smaller-scale systems. These "capability" jobs are defined as those that are either large using 512-4,528 XE nodes, (i.e. 16,384 to 144,896 integer AMD cores) or 64-845 XK nodes or very large, using more than 4,528 XE nodes, (i.e. greater than 144,896 integer AMD cores – equivalent to 72,448 core-modules), or more than 845 XK nodes. Large jobs are equal to the largest computations possible on many other compute resources, and very few systems other than Blue Waters can support very large jobs. The project consistently surpassed this 40 percent benchmark over the past year, with at least 50 percent and sometimes greater than 60 percent of computational time being devoted to the "capability" jobs that allow researchers to push the boundaries of their disciplines and to achieve results that would not be possible without Blue Waters (see Table 2). Several science teams are regularly running jobs that use 2,048, 4,096 and 8,192 nodes on Blue Waters.

- *"On Blue Waters, what would be a tremendous calculation extending over months or an entire year on another system took a week. Blue Waters entirely changed our concept of what is practical and what important research it is sensible to undertake."* **Paul Woodward, University of Minnesota**

METRIC	TARGET	6/1/2015-8/31/2015	9/1/2015-11/30/2015	12/1/2015-2/29/2016	3/1/2016-5/31/2016	PY3 6/1/2015-5/31/2016
The majority of computational time goes to capability jobs	40% of all computational time will go to capability jobs	59.0% Exceeds Expectations	66.9% Exceeds Expectations	63.0% Exceeds Expectations	50.0% Exceeds Expectations	60.1% Exceeds Expectations

TABLE 1: Blue Waters substantially exceeded its metric in regards to enabling large and very large "capability" jobs.

Many researchers who use Blue Waters have told us that their work would simply not be possible without this Track-1 resource:

- *"The simulations required to conduct our research simply cannot be accomplished in a timely manner on less powerful machines. Without Blue Waters, iNEMO research is, at best, hampered, and at worst, impossible. Blue Waters allows us to get results in a day, rather than weeks on other computer systems."* **Gerhard Klimeck, Purdue University**
- *"Without Blue Waters…neither the HIV nor the chromatophore project would be possible. … These projects are examples of how Blue Waters enables bold, new projects that push the limits of what can be done with scientific computing."* **the late Klaus Schulten, University of Illinois at Urbana-Champaign**

Support from the Blue Waters team is essential in enabling research partners to perform extreme-scale research and to achieve breakthrough results. Key features of the Blue Waters user support model include:

- Traditional Partner Consulting. Standard requests for assistance with porting, debugging, allocation issues, and software requests are handled through the Jira ticket system at help+bw@ncsa.illinois.edu.
- Advanced Application Support. Requests beyond and the above traditional consulting requests can be made via the ticket system. These requests are reviewed and evaluated by the Blue Waters Project Office (BWPO) for breadth, reach, and impact for the project and the community at large.
- Major research teams (for example, those receiving allocations through the NSF PRAC program) are assigned an individual Point of Contact (PoC) from the Blue Waters Science and Engineering Application Support (SEAS) group. Initial contact with a science team starts with a questionnaire, followed by meetings to gather information on project goals, approaches, status, challenges, requirements, and concerns. The assigned PoC works with the science team on issues such as tuning, modeling, I/O and optimizing application codes as well as standard service requests.
- Advanced Application Support activities for these major science teams have PoCs participate in code restructuring, re-engineering or redesign, such as implementing GPU functionality via OpenACC or alternatives to MPI collective operations. Such work is tracked via a coordinated work plan that is developed between the PoC and the science team

to clearly indicate the scope and scale of the work involved, including milestones and deliverables. Work plans are reviewed and approved by the Blue Waters Project Office.

- Support for workflows, data movement and visualization are provided by the appropriate Blue Waters teams and can be requested by sending email to help+bw@ncsa.illinois.edu or entering a Service Request.

The Blue Waters project monitors the time required to respond to and resolve service requests, which include traditional "trouble tickets" as well as requests for help, expanded assistance, added services and suggestions. Table 2 demonstrates how Blue Waters met or exceeded expectations for key service metrics over the past year.

Many users have mentioned the crucial role of Blue Waters support staff in enabling their work:

- *"One key aspect of the Blue Waters facility was the availability of the Blue Waters staff, who provided invaluable technical insight and frequent advice in order to tune the performance of our software."* **Greg Voth, University of Chicago**
- *"The expert staff at NCSA have been extremely responsive to my questions and needs and have helped me grow from a novice into a proficient user with plans to continue using high- performance computing to study physical-biological interactions across marine ecosystems."* **Maureen Brooks, Blue Waters Graduate Fellow**

TABLE 2: Metrics for service request response and resolution time.

METRIC	TARGET	6/1/2015-8/31/2015	9/1/2015-11/30/2015	12/1/2015-2/29/2016	3/1/2016-5/31/2016	PY3 6/1/2015-5/31/2016
Service requests are recorded and acknowledged in a timely manner	95% of partner service requests are acknowledged by a human-generated response within four working hours of submission	97% of partner service tickets had a human response within four business hours **Meets Expectations**	97% of partner service tickets had a human response within four business hours **Meets Expectations**	99% of partner service tickets had a human response within four business hours **Exceeds Expectations**	97% of partner service tickets had a human response within four business hours **Meets Expectations**	97% of partner service tickets had a human response within four business hours **Meets Expectations**
Most problems are solved within a reasonable time	80% of partner service requests are addressed within three working days, either by - resolving them to the partner's satisfaction within three working days, or, - for problems that will take longer, by informing the partner how the problem will be handled within three working days (and providing periodic updates on the expected resolution)	80.5% of partner service requests were resolved within three business days **Meets Expectations**	82% of partner service requests were resolved within three business days **Meets Expectations**	88% of partner service requests were resolved within three business days **Exceeds Expectations**	87% of partner service requests were resolved within three business days **Exceeds Expectations**	85% of partner service requests were resolved within three business days **Exceeds Expectations**

The Petascale Application Improvement Discovery (PAID) program is another way hat the Blue Waters project supports advanced application development and improvement work. PAID forges collaborations between Improvement Method Enablers (IMEs) and science and engineering research teams, helping them to create and implement technologies that improve application performance. [A detailed report on the PAID program can be found elsewhere in this volume.]

Additionally, the annual Blue Waters Symposium provides a excellent opportunity for cross-pollination between research teams and across disciplines. Attendees report that the facilitation of networking, interaction and discussion is a particular strength of the symposium and that the event provides useful resources and information.

The Blue Waters project also aims to develop the next generation of computational researchers by supporting education and training activities for undergraduate, graduate students and professional staff.

- Each year an elite group of PhD students from a diverse range of domains and from institutions across the country are awarded Blue Waters Graduate Fellowships. These competitive fellowships include both financial support and access to the Blue Waters system to substantially advance the fellows' computational research. To date, 26 young scholars have participated in the program, which over three years will award more than $1.3 million in direct financial support and over 50 million core equivalent hours to support graduate research. The 26 Fellows are from 20 states and 22 institutions, including six institutions from EPSCoR states and one that is a minority serving institution (MSI).

In order to ensure that the fellows are able to effectively utilize the power of the Blue Waters system, each fellow is paired with a member of the Blue Waters Science and Engineering Application Support (SEAS) group; these "points of contact" assist the fellows with essential activities such as porting and tuning codes and with any support requests. In the December 2015 review of the Blue Waters project by the National Science Foundation, the review panel pointed to this in depth POC model as a "best practice" for fellowship and education programs.

Some of the impressive research results achieved by the 2015-2016 Blue Waters Graduate Fellows are showcased in this annual report. Many fellows also highlighted the impact that Blue Waters and their fellowship experience will have on their research and their future career path:

- *"The Blue Waters Fellowship has opened new doors for me by providing me first-hand experience with the high-performance computing resources necessary for cutting-edge science in oceanography."* **Maureen Brooks**
- *"The Blue Waters Graduate Fellowship has not only played a pivotal role in helping me achieve my doctorate but it has also exposed me to a variety of fields and topics related to HPC that I would have otherwise not experienced as part of my degree plan. I believe this breadth of knowledge will give me a competitive edge as I continue to the next stages of my career."* **Justin Drake**

Approximately 20 students are selected each year for the Blue Waters Undergraduate Internship program, which includes an intensive two-week workshop covering the fundamentals of high-performance computing and engages undergraduate students in year-long projects in which they apply HPC to problems in science, mathematics, and engineering with guidance from faculty mentors. To date, the Blue Waters project has funded 99 interns from 64 unique institutions in 27 states, DC, and Puerto Rico. Of the 99 interns, 42 are women and/or minorities. Twenty-six of the 64 institutions were in EPSCoR regions and 17 were classified as minority serving institutions.

External evaluation has found that both interns and their advisors find the program beneficial. Some participants have reported altering their educational and career path or embarking on a path with greater confidence thanks to the internship experience. For example:

"The Blue Waters Internship Program was a life-changing event. The program has given me the experience necessary to work in HPC after graduation. In addition, I have made extremely valuable connections in the supercomputing and HPC world that I know will help me in the future."

In another aspect of the Blue Waters engagement with students is the fact that from the beginning of the project, **42.0 percent** of the people using Blue Waters were students, which accounts for **1023** total student users (**881** graduate students, **125** undergraduate students, **and 17 High School students**). Furthermore at least **350** users (**14.4 percent**) can be identified as post docs and other early career researchers.

The Project also funds, in partnership with XSEDE, the content in HPC University (http://hpcuniversity.org), which is focused on developing educational materials for and about petascale computing and analysis. This year, the HPC University had 67,533 petascale learning modules and training materials downloaded.

Because of excellent overall project management, the Blue Waters project is pleased that the National Science Foundation has approved continued operations of the Blue Waters project through at least March 2019. With this one-year extension of system operations, the project will continue to serve the research needs of the scientific community and as well as the undergraduate and graduate education activities. The dedicated Blue Waters team at NCSA looks forward to continuing productive collaborations with our science and engineering partners as they continue to innovate and discover.

SCIENCE AND ENGINEERING TEAM ADVISORY COMMITTEE

PETASCALE COMPUTING RESOURCE TEAMS

- **Paul Woodward**, Physics and Astrophysics, University of Minnesota (SETAC Chair)
- **Susan Bates**, Oceanography & Climate Dynamics, NCAR
- **Manuela Campanelli**, Physics and Astronomy, Rochester Institute of Technology
- **Tom Cheatham**, Chemistry, University of Utah
- **David Ceperley**, Physics and Material Science, University of Illinois at Urbana-Champaign
- **Tiziana Di Matteo**, Physics and Cosmology, Carnegie Mellon University
- **Said Elghobashi**, Mechanical and Aerospace Engineering, University of California, Irvine
- **Tom Jordan and Philip J. Maechling**, Seismology, University of Southern California
- **Paul Morin**, Earth Sciences, University of Minnesota
- **Brian O'Shea**, Physics and Astronomy, Michigan State University
- **Nikolai Pogorelov**, Space Science, University of Alabama, Huntsville

GREAT LAKES CONSORTIUM FOR PETASCALE COMPUTATION TEAMS

- **Ben Rogers**, Information Technology Research Services, University of Iowa

INDUSTRY TEAMS

- **Rick Arthur**, General Electric Global Research, Advanced Computing

UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN TEAMS

- **Athol Kemball**, Astronomy, University of Illinois at Urbana-Champaign

The Science and Engineering Team Advisory Committee (SETAC) brings together a diverse group of scientists and engineers who represent the various Blue Waters the science and engineering research teams using Blue Waters and the breadth of research powered by the Blue Waters system. The committee provides guidance and assessment to help the Blue Waters project deliver the best possible performance and services that will in turn assist research teams in achieving breakthrough results.

The SETAC makes recommendations on technical directions, strategies, and management while identifying potential challenges for petascale applications. As users themselves, the SETAC members also provide advice for solving common issues that arise from moving applications to Blue Waters and from system software at scale.

The SETAC convenes three to four times per year and is available to the Blue Waters Project to provide guidance and advice as needed throughout the year. Members are nominated and typically serve two-year terms.

In its 2016 annual report, published in July 2016, the SETAC commended the Blue Waters team "on their excellent operation of the system, which is making possible an exceptional level of scientific achievement by the participating teams." In particular, the committee noted that Blue Waters' "very large disk file system...and the nearline tape system represent special resources for the scientific community" and that the recent addition of a "big-data" service track "has the potential to add significantly to the impact of the system."

Among other recommendations, the Committee members strongly agreed that the Petascale Application Improvement Discovery Program (PAID) program should continue in some form for as long as possible.

Finally, in its annual report the SETAC noted the necessity of developing a plan to provide the computing and data capabilities that will be needed by the academic and open research communities in the near future to drive scientific discovery forward. "As computational scientists, we need to plan for such developments, to adapt and enhance our codes, and we strongly desire to enjoy the benefits of the greater power of future systems."

OTHER BLUE WATERS PROJECTS

Arrika, Kenza
Allocation Type: Graduate Fellow

Beltran, Chris
GPU-Accelerated and Monte Carlo Based Robust Optimization for Spot Scanning Proton Therapy
Allocation Type: Industrial

Bernholc, Jerzy
Petaflops simulation and design of nanoscale materials and devices
Allocation Type: NSF PRAC General

Bhatele, Abhinav
Tracking the Influenza Pandemic in the Continental United States using EpiSimdemics
Allocation Type: Innovation and Exploration

Brunner, Robert J.
Applying Deep Learning on Time Series Astronomical Data
Allocation Type: Illinois Exploratory

Chaudhuri, Santanu
Material Discovery using Evolutionary Algorithms: Finding the Missing Ruthenium-based Ternary Phases for Breakthrough Applications in Energy Technologies
Allocation Type: Illinois General

Dyadechko, Vadim
Allocation Type: Industrial

Fischer, Paul
Numerical methods and software for computational fluid dynamics (CFD), nek5000
Allocation Type: Illinois Blue Waters Professor

Jha, Shantenu
Collaborative Research: The Power of Many: Scalable Compute and Data-Intensive Science on Blue Waters
Allocation Type: NSF PRAC

Johnson, Curtis
Reconstruction of High-Resolution, Quantitative MRI Images in a Clinical Setting
Allocation Type: Illinois Exploratory

Karimabadi, Homayoun
Enabling Breakthrough Kinetic Simulations of the Magnetosphere via Petascale Computing
Allocation Type: NSF PRAC

Kemball, Athol
Pixel-level uncertainty quantification, optimized for Blue Waters
Allocation Type: Illinois Blue Waters Professor

Khalili-Araghi, Fatemeh
Paracellular Transport Mechanism in Tight Junctions
Allocation Type: GLCPC General

Lee, Jong
Numerical simulation analysis for grid convergence and code validation studies using the NASA CRM wing-body-nacelle-pylon configuration
Allocation Type: Illinois General

Lele, Sanjiva
Shock-induced turbulent multi-material mixing
Allocation Type: NSF PRAC

Ricker, Paul
Effects of Active Galaxy Feedback on the Evolution of Galaxy Clusters
Allocation Type: Illinois Exploratory

Tajkhorshid, Emad
Atomic Resolution Description of the Transport Cycle in Neurotransmitter
Allocation Type: Illinois General

Wang, Jue
Allocation Type: Industrial

Woosley, Stan
Type Ia Supernovae
Allocation Type: NSF PRAC

REFERENCES

SPACE SCIENCE

Balsara, Dinshaw

[1] Balsara, D.S., Wave Propagation in Molecular Clouds, *Ap.J.* 465 (1996) pp. 775-794.

[2] Meyer, C., D.S. Balsara, B. Burkhart, and A. Lazarian, Observational diagnostics for two-fluid turbulence in molecular clouds as suggested by simulations, *Mon. Not. R. Astron. Soc.*, 439 (2014), pp. 2197-2210.

[3] Burkhart, B., et al., Alfvenic turbulence beyond the ambipolar diffusion scale, *Astrophys. J.*, 805 (2015) pp. 118-126.

[4] Garain, S., D. S. Balsara, and J. Reid, Comparing Coarray Fortran (CAF) with MPI for several structured mesh PDE applications. *J. Comput. Phys.*, 297 (2015), pp.237 253, doi:10.1016/j.cp.2015.05.020.

[5] Balsara, D.S., T. Amano, S. Garain, J. Kim, High order accuracy divergence-free scheme for the electrodynamics of relativistic plasmas with multidimensional Riemann solvers. doi: 10.1016/j.jcp.2016.05.006

Cox, Donna

[1] Turk, Matthew J. *et al.*, yt: A Multi-code Analysis Toolkit for Astrophysical Simulation Data. *The Astrophysical Journal Supplement*, 192:1 (2011), article id 9.

Gammie, Charles

[1] Balbus, S., and J. Hawley, A powerful local shear instability in weakly magnetized disks. I - Linear analysis. *ApJ*, 376 (1991), pp. 214-233.

[2] Davis, S., J. Stone, and M. Pessah, Sustained magnetorotational turbulence in local simulations of stratified disks with zero net magnetic flux. *ApJ*, 713 (2010), pp. 52-65.

[3] Bodo, G., et al., On the convergence of magnetorotational turbulence in stratified isothermal shearing boxes. *ApJ Lett.*, 787:1 (2014), L13, 4 pp.

[4] Kestener, P., and S. Fromang, Compressible magneto hydrodynamics simulations on a GPU cluster. astrophysics application: High resolution MRI-driven turbulence. *Astronum 2012*, Big Island, Hawaii, USA, June 25-29, 2012.

[5] Fromang, S., and J. Papaloizou, MHD simulations of the magnetorotational instability in a shearing box with zero net flux. I. The issue of convergence. *Astron. Astrophys.*, 476:3 (2007), pp. 1113-1122.

Lentz, Eric J.

[1] Bruenn, S. W., et al. The axisymmetric *ab initio* core-collapse supernova simulations of 12-25 M stars. *Astrophys. J. Lett.*, 767 (2013), L6.

[2] Lentz, E. J., et al. Three-dimensional core-collapse supernova simulated ssing a 15 M progenitor. *Astrophys. J. Lett.*, 807(2015), L31.

[3] Woosley, S. E., and A. Heger, A. Low mass title, *Astrophys. J.*, 810 (2015), p. 34.

[4] Melson, T., H.-T. Janka, and A. Marek, Neutrino-driven supernova of a low-mass Iron-core progenitor boosted by three-dimensional turbulent convection. *Astrophys. J. Lett.*, 801(2015), L24

Nobel, Scott C.

[1] Zilhao, M., and S. Noble, Dynamic fisheye grids for binary black hole simulations. *Classical Quant. Grav.*, 31:6 (2014), 065013.

[2] Noble, S., et al., Circumbinary magneto-hydrodynamic accretion into inspiraling binary black holes, *Astrophys. J.*, 755:1 (2012), 51.

[3] Zilhao, M., and S. Noble, and M. Campanelli, and Y. Zlochower, Resolving the relative influence of strong field spacetime dynamics and MHD on circumbinary disk physics, *Phys. Rev. D*, 91:2 (2015), 024034.

[4] Paczynski, B., A model of accretion disks in close binaries, *Astrophys. J.*, 216:1 (1977), p. 822-826.

[5] Papaloizou, J.; J. Pringle, Tidal torques on accretion discs in close binary systems, *Mon. Not. R. Astron. Soc.*, 181 (1977), p. 441-454.

[6] Artymowicz, P.; S. Lubow, Dynamics of binary-disk interaction. 1: Resonances and disk gap sizes, *Astrophys. J.*, 421:2 (1994), p. 651-667.

[7] Kale, L., and S. Krishnan, Charm++: Parallel Programming with Message-Driven Objects. In *Parallel Programming using C++*, MIT Press, (1996), pp 175-213.

Makwana, Kirit

[1] Cargill, P., Theories of heating and particle acceleration in the solar corona. *Advances in Space Research*, 26:11 (2001), 1759.

[2] Kasper, J., B. Maruca, M. Stevens, and A. Zaslavsky, Sensitive Test for Ion-Cyclotron Resonant Heating in the Solar Wind. *Phys. Lett. Rev.*, 110 (2013), 091102.

[3] Mignone, A., et al., PLUTO: A Numerical Code for Computational Astrophysics. *ApJ Supplement Series*, 170 (2007), pp. 228-242.

[4] Bowers, K., et al., Ultrahigh performance three-dimensional electromagnetic relativistic kinetic plasma simulation. *Physics of Plasmas*, 15 (2008), 055703.

O'Shea, Brian

[1] Peng, Y.-j. et al. Mass and environment as drivers of galaxy evolution in SDSS and zCOSMOS and the origin of the Schechter function. *Astrophys. J.*, 721 (2015), pp. 193–221

[2] Ivezic , Z. et al. The Milky Way tomography with SDSS. II. Stellar metallicity., *Astrophys. J.*, 684 (2008), pp. 287–325

[3] Putman, M. E., J.E.G. Peek, and M.R. Joung, Gaseous galaxy halos. *Ann. Rev. Astro. Astrophys.*, 50 (2012), pp. 491–529

[4] Tumlinson, J. et al., The COS-Halos Survey: Rationale, design, and a census of circumgalactic neutral hydrogen. *Astrophys. J.*, 777 (2013), p. 59.

[5] Bryan, G. et al., ENZO: An adaptive mesh refinement code for astrophysics. *Astrophys. J. Suppliments*, 211 (2014), p. 19.

Pogorelov, Nikolai

[1] Burlaga, L. F. and N.F Ness, Interstellar Magnetic Fields Observed by Voyager 1 beyond the Heliopause. *Astrophys. J. Lett.*, 795:1 (2014), L19.

[2] Gurnett, D. A., S.W. Kurth, L.F. Burlaga, and N.F. Ness, In situ observations of Interstellar plasma with voyager 1. *Science*, 341:6153 (2013), pp. 1489-1492.

[3] McComas et al., Local interstellar medium: Six years of direct sampling by IBEX. *Astrophys. J. Suppl.*, 220:2 (2015), p. 22.

[4] Zhang, M., P. Zuo, N. Pogorelov, Heliospheric influence on the anisotropy of TeV cosmic rays. *Astrophys. J.*, 790:1 (2014), 5.

Shapiro, Stuart

[1] Abbott, B. P. et al., Observation of gravitational waves from a binary black hole merger. *Phys. Rev. Lett.*, 116 (2016), 061102.

[2] Paschalidis, V., M. Ruiz, and S. L. Shapiro, Relativistic simulations of black hole—neutron star coalescence: the jet emerges. *ApJ*, 806 (2015), L14.

[3] Etienne, Z. B., V. Paschalidis, and S. L. Shapiro, General relativistic simulations of black hole—neutron star mergers: Effects of tilted magnetic fields. *Phys. Rev. D*, 86 (2012), 084026.

[4] Rezolla et al., The missing link: Merging neutron stars naturally produce jet-like structures and can power short Gamma-Ray Bursts. *ApJ*, 732 (2011), L6.

[5] Kiuchi et al., High resolution numerical-relativity simulations for the merger of binary magnetized neutron stars. *Phys. Rev. D*, 90 (2014), 041502.

[6] Etienne, Z. B., Y. T. Liu, and S. L. Shapiro, Relativistic magnetohydrodynamics in dynamical spacetimes: A new AMR implementation. *Phys. Rev. D*, 82 (2010), 084031.

GEOSCIENCE

DiGirolamo, Larry

[1] NRC, Earth Science and Applications from Space: National Imperatives for the Next Decade and Beyond, (2007), The National Academies Press, Washington DC.

[2] Wolfe, R., and H. Ramapriyan, Scaling the pipe: NASA EOS Terra data systems at 10, (2010), *IEEE International. Geoscience and Remote Sensing Symposium (IGARSS), 2010 IEEE International*, Honolulu, HI. p.1300-1303.

[3] Liang, L., L. Di Girolamo, and W. Sun, Bias in MODIS cloud drop effective radius for oceanic water clouds as deduced from optical thickness variability across scattering angles, *J. Geophys. Res. Atmos.*, 120, (2015), doi:10.1002/2015JD023256.

[4] Zhao, G., et al, Regional changes in Earth's color and texture as observed from space over a 15-year period (2016), *IEEE Trans. Geosci. Remote Sens.*, doi:10.1109/TGRS.2016.2538723.

Garcia, Marcelo

[1] Bulle, H., Untersuchungen ber die geschiebeableitung bei der spaltung von wasserlufen (in German), *V.D.I Verlag, Berlin, Germany, Tech. Rep.*, (1926).

[2] Giosan, L., J. Syvitski, S. Constantinescu, and J. Day, Protect the world's deltas, *Nature*, 516 (2014), pp. 31–33.

[3] Kenney, M.A., et al., Cost analysis of water and sediment diversions to optimize land building in the Mississippi river delta, *Water Resour. Res.*, 49 (2013), pp. 3388–3405.

[4] Dutta, S., et al., Application of computational fluid dynamic modelling to improve flow and grit transport in Terence J. O'Brien Water Reclamation Plant, Chicago, Illinois, *J. Hydrl. Res.*, 52:6 (2014), pp. 759-774.

[5] Lee, S., et al., Direct numerical simulation of transitional flow in stenosed carotid bifurcation, *J. Biomech.*, 41 (2008), pp. 2551–2561.

[6] Fischer, P.F., J.W. Lottes, and S.G. Kerkemeier, Nek 5000 webpage, (2008), http://nek5000.mcs.anl.gov

[7] Deville, M., P. Fischer, and E. Mund, High-Order Methods for Incompressible Fluid Flow, *Cambridge, England: Cambridge University Press*, (2002).

Gregg, Patricia

[1] Gregg, P.M., S.L. de Silva, and E.B. Grosfils, Thermomechanics of shallow magma chamber pressurization: Implications for the assessment of ground deformation data at active volcanoes. *EPSL*, 384 (2013), pp. [100-108].

[2] Grosfils, E.B., et al., Elastic models of magma reservoir mechanics: A key tool for understanding planetary volcanism, in Volcanism and Tectonism Across the Inner Solar System, T. Platz, et al., Editors. 2015, *Geological Society, London, Special Publications*. p. 139-158.

[3] Hickey, J., J. Gottsmann, and P. Mothes, Estimating volcanic deformation source parameters with a finite element inversion: The 2001–2002 unrest at Cotopaxi volcano, Ecuador. *J Geophys. Res.*, 120 (2015), pp. 1473-1486.

[4] Battaglia, M. and D.P. Hill, Analytical modeling of gravity changes and crustal deformation at volcanoes: The Long Valley caldera, California, case study. *Tectonophysics*, 471:1-2 (2009), pp. 45-57.

[5] Evensen, G., Data Assimilation: The Ensemble Kalman Filter, Second Edition. (2009) Heidelberg: Springer, ISBN: 978-3-642-03710-8.

[6] Gregg, P.M. and J.C. Pettijohn, A multi-data stream assimilation framework for the assessment of volcanic unrest. *J Volc. Geotherm. Res.*, 309 (2016), pp. 63-77.

[7] Pritchard, M.E. and P.M. Gregg, Geophysical Evidence for Silicic Crustal Melt in the Continents: Where, What Kind, and How Much? *Elements*, 12 (2016), pp. 121-127.

Jordon, Thomas H.

[1] Bielak, J. et al., The ShakeOut earthquake scenario: verification of three simulation sets. *Geophys. J. Int.*, 180:1 (2010), pp. 375–40.

[2] Taborda, R., and J. Bielak, Ground-motion simulation and validation of the 2008 Chino Hills, California, earthquake. *Bull. Seism. Soc. Am.*, 103:1 (2013), pp. 131–156.

[3] Ely, G. P., S.M Day, and J.B. Minster, A support-operator method for 3-D rupture dynamics. *Geophys. J. Int.*, 177:3 (2009), pp. 1140–1150, doi:10.1111/j.1365-246X.2009.04117.x.

[4] Shi, Z., and S.M. Day, Rupture dynamics and ground motion from 3-D rough-fault simulations, Journal of Geophysical Research: *Solid Earth*, 118:3 (2013), pp. 1122–1141, doi:10.1002/jgrb.50094.

Kumar, Praveen

[1] Houghton, J. T., Climate change 2001: The scientific basis: contribution of Working Group I to the third assessment report of the Intergovernmental Panel on Climate Change. Cambridge, U.K.; New York: Cambridge University Press.

[2] Drewry, D. T., et al., Ecohydrological responses of dense canopies to environmental variability: 1. interplay between ver- tical structure and photosynthetic pathway. *J. Geophys. Res.*, 115 (2010) G04022.

[3] Le, P. V. V., P. Kumar, D.T. Drewry, and J.C. Quijano, A graphical user interface for numerical modeling of acclimation responses of vegetation to climate change. *Comput. Geosci.*, 49 (2012), pp. 91 – 101, doi:http://dx.doi.org/10.1016/j.cageo. 2012.07.007.

[4] Le, P. V. V., P. Kumar, A.J. Valocchi, and H. V. Dang, GPU-based high-performance computing for integrated surface– subsurface flow modeling. *Environ. Model. Software*, 73 (2015), p. 1 – 13. doi:http://dx.doi.org/10.1016/j.envsoft.2015.07. 015.

Lasher-Trapp, Sonia

[1] Dawe, J.T., and P. H. Austin, Interpolation of LES Cloud Surfaces for Use in Direct Calculations of Entrainment and Detrainment. *Mon. Wea. Rev.*, 139 (2011), pp. 444-456.

Morin, Paul

[1] Noh, M.J., and I. M. Howat, Automated stereo-photogrammetric DEM generation at high latitudes: Surface Extraction with TIN-based Search-space Minimization (SETSM) validation and demonstration over glaciated regions. *GIScience Remote Sens.*, August (2015), pp. 1–20, doi:10.1080/15481603.2015.1008621.

[2] White House, Using new elevation data to explore the Arctic, Available from: https://www.whitehouse.gov/blog/2015/09/02/using-new-elevation-data-explore-arctic

Rauber, Robert

[1] Giorgi, F., Simulation of regional climate using a limited area model nested in a general circulation model. *J. Climate*, 3:9 (1990), pp. 941-963.

[2] Mearns, L. O., F. Giorgi, L. McDaniel, C. Shields, Analysis of daily variability of precipitation in a nested regional climate model: comparison with observations and doubled CO_2 results. *Global Planet. Change*, 10:1 (1995), pp. 55-78.

[3] Small, R.J., et al., A new synoptic scale resolving global climate simulation using the Community Earth System Model. *J. Adv. Modeling Earth Systems*, 6 (2014), pp. 1065–1094.

[4] Wang, J., V.R. Kotamarthi, High-resolution dynamically downscaled projections of precipitation in the mid and late 21st century over North America. *Earth's Future*, 3:7 (2015), pp. 268-288.

[5] Liang, X.-Z., et al., Regional Climate-Weather Research and Forecasting Model (CWRF). *Bull. Amer. Meteor. Soc.*, 93 (2012), pp. 1363-1387.

[6] Reynolds, R. W., et al., Daily high-resolution blended analyses for sea surface temperature. *J. Climate*, 20 (2007), pp. 5473-5496.

Simpson, Jamesina

[1] Shea, M. A., *et al.*, Solar proton events for 450 years: The Carrington event in perspective. *Advances in Space Research*, 38:2 (2006), pp. 232-238.

[2] Odenwald, S., The 23rd Cycle: Learning to Live with a Stormy Star, Columbia University Press (2002).

[3] Severe Space Weather Events – Understanding Societal and Economic Impacts, A Workshop Report, National Research Council of the National Academies, Washington D.C., May (2008).

[4] Simpson, J. S. and A. Taflove, A review of progress in FDTD Maxwell's equations modeling of impulsive sub-ionospheric propagation below 300 kHz. *IEEE Transactions on Antennas and Propagation*, 55:6 (2007), pp. 1582-1590.

[5] Simpson, J. S., Current and future applications of 3-D global Earth-ionosphere models based on the full-vector Maxwell's equations FDTD method. *Surveys in Geophysics*, 30:2 (2009), pp. 105-130.

Sriver, Ryan

[1] Gent , P. et al., The Community Climate System Model Version 4, *J. Climate*, 24 (2011), pp. 4973-4991,doi: 10.1175/2011JCLI4083.1

[2] Small, J. R., et al., A new synoptic scale resolving global climate simulation using the Community Earth System Model, *J. Adv. Model. Earth Syst.*, 6:4 (2014), doi:10.1002/2014MS000363.

Um, Junshik

[1] Um, J., and G. M. McFarquhar, Formation of atmospheric halos and applicability of geometric optics for calculating single-scattering properties of hexagonal ice crystals: Impacts of aspect ratio and ice crystal size. *J. Quant. Spectrosc. Radiat. Transfer*, 165 (2015), pp. 134-152, doi:10.1016/j.jqsrt.2015.07.001.

West, Matthew

[1] Dockery, D., and A. Pope, Epidemiology of acute health effects: summary of time-series studies. *Particles in our air: Concentrations and health effects*, (1996) p. 123-147.

[2] Stocker, T. F., et al., Technical summary. *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press, (2013) p. 33-115.

[3] Riemer, N., West, M., Zaveri, R. A., and R.C. Easter, Simulating the evolution of soot mixing state with a particle-resolved aerosol model. *J. Geophys. Res.: Atmos.*, 114:D9 (2009).

[4] Zaveri, R. A., et al. Overview of the 2010 carbonaceous aerosols and radiative effects study (CARES). (2012).

PHYSICS & ENGINEERING

Ackerman, David

[1] Jackson, N. E. et al., Conformational order in aggregates of conjugated polymers. *J. Am. Chem. Soc.*, 137 (2015), pp. 6254–6262.

[2] Mickiewicz, R. A. et al., Enhancing the potential of block copolymer lithography with polymer self-consistent field theory simulations. *Macromolecules*, 43 (2010), pp. 8290–8295.

[3] Hur, S.-M., A. L. Frischknecht, D. L. Huber, and G. H. Fredrickson, Self-consistent field simulations of self- and directed-assembly in a mixed polymer brush. *Soft Matter*, 7 (2011), pp. 8776-8788.

[4] Kratky, O., G. Porod. Röntgenuntersuchung gelöster Fadenmoleküle. *Rec. Trav. Chim*, 68 (1949), pp. 1106-1123.

[5] Jiang Y., J. Z. Y. Chen. Influence of chain rigidity on the phase behavior of wormlike diblock copolymers. *Phys. Rev. Lett.*, 110 (2013), 138305.

Allain, Jean

[1] Facsko, S., et al., Formation of Ordered Nanoscale Semiconductor Dots by Ion Sputtering. *Science*, 285:5433 (1999), pp. 1551-1553.

[2] Sadana, D. K., et al., Recrystallization of amorphous gallium arsenide by ion beams. *Appl. Phys. Lett.*, 44:3 (1984), pp. 301-303.

[3] Kaijaks, N., Ion-scattering spectroscopy of III-V semiconductor surfaces. PhD thesis, Warwick University (2000).

[4] Bracht, H., et al., Large disparity between gallium and antimony self-diffusion in gallium antimonide. *Nature*, 408:6808 (2000), pp. 69-72.

[5] Chroneos, A., and H. Bracht, Concentration of intrinsic defects and self-diffusion in GaSb. *J. Appl. Phys.*, 104:9 (2008), p. 093714.
[6] Norris, S. A., et al., Molecular dynamics of single-particle impacts predicts phase diagrams for large scale pattern formation. *Nat. Commun.*, 2 (2011), p. 276.
[7] Norris, S. A., et al., Crater functions for compound materials: A route to parameter estimation in coupled-PDE models of ion-bombardment. *Nucl. Instrum. Meth. Phys. Res. B*, 318B (2014), pp. 245-252.

Aluru, Narayana

[1] Wells, D.B., M. Belkin, and J. Comer, A. Aksimentiev. Assessing Graphene Nanopores for Sequencing DNA. *Nano Lett.*, 12 (2012), pp. 4117-4123.
[2] Branton, D., et al., The Potential and Challenges of Nanopore Sequencing. *Nat. Biotech*, 26 (2008), pp. 1146-1153.

Bodony, Daniel

[1] Hwu, W.-M. (Ed.), *Heterogeneous System Architecture: A new compute platform infrastructure*, Morgan Kaufmann, Waltham, MA, USA, 2016.
[2] Gregory, K. and A. Miller. *C++ Amp: Accelerated Massive Parallelism With Microsoft Visual C++*, Microsoft Press, 2012.
[3] Adve, V. *et al., The LLVM Compiler Infrastructure*, http://www.llvm.org.
[4] Harris, M., *Inside Pascal: NVIDIA's Newest Computing Platform* https://devblogs.nvidia.com/parallelforall/inside-pascal, June 19, 2016.

Ceperley, David

[1] McMahon, J. M., M. A. Morales, C. Pierleoni, and D. M. Ceperley, The properties of hydrogen and helium under extreme conditions. *Rev. Mod. Phys.*, 84 (2012), pp. 1607-1653.
[2] Morales, M. A., C. Pierleoni, E. Schwegler, and D. M. Ceperley, Evidence of a liquid–liquid phase transition in hot dense hydrogen. *Proc. Nat. Acad. Science*, 107 (2010), pp. 12799-12803.
[3] Knudson, M. D., et al., Direct observation of an abrupt insulator-to-metal transition in dense liquid deuterium. *Science*, 348 (2015), p. 1455.
[4] Zaghoo, M., A. Salamat, and I. Silvera, Evidence of a first-order phase transition to metallic hydrogen. *Phys. Rev.*, B 93 (2016), p. 155128.
[5] Pierleoni, C., *et al.*, Liquid–liquid phase transition in hydrogen by coupled electron–ion Monte Carlo simulations. *Proc. Nat. Acad. Science* (US,) 113 (2016), pp. 4953-4957.

Clark, Bryan

[1] Pal, A., and D. Huse, Many-body localization phase transition. *Phys. Rev. B*, 82 (2013), p. 17441, doi:10.1103/PhysRevB.87.134202
[2] Basko, D. M., I. L. Aleiner, and B. L. Altshuler, Metal–insulator transition in a weakly interacting many-electron system with localized single-particle states. *Annals of Physics*, 321:5 (2006), pp. 1126-1205, doi:10.1016/j.aop.2005.11.014.

Curreli, Davide

[1] Khaziev R., and D. Curreli, Ion energy-angle distribution functions at the plasma-material interface in oblique magnetic fields, *Physics of Plasmas*, 22, 043503 (2015), DOI: http://dx.doi.org/10.1063/1.4916910
[2] Hockney, R.W., and J.W Eastwood, Computer Simulation Using Particles, CRC Press, 1988
[3] Andruczyk, D., D. N. Ruzic, D. Curreli, J. P. Allain and the HIDRA Team, HIDRA: Hybrid Illinois Device for Research and Applications, *Fusion Science and Technology*, 68(3), 497-500 (2015)
[4] http://www.mcs.anl.gov/project/fastmath-scidac-institute

Draayer, Jerry

[1] Dytrych, T., et al., Collective Modes in Light Nuclei from First Principles. *Phys. Rev. Lett.*, 111 (2013) 252501.
[2] *LSU3shell* code, Louisiana State University, available under the *GNU General Public License* at the git repository, http://sourceforge.net/ projects/lsu3shell

Glotzer, Sharon

[1] http://glotzerlab.engin.umich.edu/hoomd-blue/
[2] Anderson, J. A., C. D. Lorenz, A. Travesset, General purpose molecular dynamics simulations fully implemented on graphics processing units, *J. Comp. Phys.*, 227:10, (2008), pp. 5342-5359.
[3] Anderson, J. A., M. E. Irrgang, S. C. Glotzer, Scalable Metropolis Monte Carlo for simulation of hard shapes, *Comput. Phys. Commun.*, 204 (2016), pp. 21-30.
[4] Feigelson, R. S., Crystal Growth through the Ages: A Historical Perspective, in Handbook of Crystal Growth, (2015), pp. 1-83.
[5] Engel, M., P. F. Damasceno, C. L. Phillips, S. C. Glotzer, Computational self-assembly of a one-component icosahedral quasicrystal, *Nat. Mater.*, 14 (2015), pp. 109-116.
[6] Glaser, J., *et al.*, Strong scaling of general-purpose molecular dynamics simulations on GPUs. *Comput. Phys. Commun.*, 192 (2015), pp. 97-107.

Jasiuk, Iwona

[1] www.nof.org
[2] Weiner, S., Traub, Bone structure: from angstroms to microns. *FASEB*, 6 (1992), pp. 879–85.
[3] Niebur, G.L., et al., High-resolution finite element models with tissue strength asymmetry accurately predict failure of trabecular bone. *J. Biomech.*, 33:12 (2000), pp. 1575-1583.
[4] Stölken, J.S., J.H. Kinney, On the importance of geometric nonlinearity in finite-element simulations of trabecular bone failure. *Bone*, 33:4 (2003), pp. 494-504.
[5] Koric, S., Gupta, A, Sparse matrix factorization in the implicit finite element method on petascale architecture. *Comput. Methods Appl. Mech. Eng.* 302 (2016), pp. 281-292.

Klimeck, Gerhard

[1] Steiger, S., M. Povolotskyi, H. Park, T. Kubis, and G. Klimeck, NEMO5: A Parallel Multi-scale Nanoelectronics Modeling Tool, *IEEE Transactions on Nanotechnology*, 10 (2011), pp. 1464-1474.
[2] Fonseca, J. E., T. Kubis, M. Povolotskyi, B. Novakovic, A. Ajoy, G. Hegde, H. Ilatikhameneh, Z. Jiang, P. Sengupta, Y. Tan, and G. Klimeck, Efficient and realistic device modeling from atomic detail to the nanoscale. *Journal of Comp. Electronics*, 12:4 (2013), pp. 592-600.

Koric, Seid

[1] Vázquez, M., G. et al., Alya: Multiphysics engineering simulation toward exascale. *J. Comput. Sci.*, 2016, in Press, doi:10.1016/j.jocs.2015.12.007
[2] S. Govert, et al., Turbulent combustion modelling of a confined premixed methane/air jet flame including heat loss effects using tabulated chemistry, *Appl. Energy*, 156 (2015), pp. 804-815.
[3] Casoni, E., et al., Alya: computational solid mechanics for supercomputers. *Arch. Comput. Method. E* , 22:4 (2014), pp. 557-576.

Laskowski, Gregory

[1] Arts, T., and M. L. Rouvroit, Aerothermal Performance of a Two-Dimensional Highly Loaded Transonic Turbine Nozzle Guide Vane: A Test Case for Inviscid and Viscous Flow Computations, *ASME J.* Turbomachinery, 1992, Vol. 114.
[2] Saumweber, C., A. Schulz, and S. Wittig, Free-stream turbulence effects on film cooling with shaped holes, *ASME J. Turbomachinery*, 125(1) (2003) pp. 65-73.
[3] Gaitonde, D. V., and M. R. Visbal, High-order schemes for Navier-Stokes equations: algorithms and implementation into FDL3DI, Technical Report AFRL-VA-WP-TR-1998-3060, Air Force Research Laboratory, Wright-Patterson AFB, Ohio (1998).
[4] Wang, Z. J., K.J. Fidkowski, R. Abgrall, F. Bassi, D. Caraeni, et al., High-Order CFD Methods: Current Status and Perspective, *Int. J. Numer. Meth. Fluids*, 72 (2013), pp. 811-845.
[5] Laskowski, G. M., J. Kopriva, V. Michelassi, S. Shankaran, U. Paliath, et al., Future directions of high-fidelity CFD for aerothermal turbomachinery research, analysis and design, 46th AIAA Fluid Dynamics Conference, AIAA Aviation (2016).
[6] Talnikar, C., G.M. Laskowski, and Q. Wang, Unsteady adjoint pressure loss over a gas turbine inlet guide vane, submitted to 2016 IGTI, GT2016-56689, Seoul, South Korea.

Luitz, David

[1] Basko, D. M., I.L Aleiner, and B.L. Altshuler, Metal–insulator transition in a weakly interacting many-electron system with localized single-particle states." *Annals of Physics* 321 (2006), pp. 1126–1205.
[2] Deutsch, J. M., Quantum statistical mechanics in a closed system. *Phys. Rev. A* 43, 2046–2049 (1991).
[3] Srednicki, M. Chaos and quantum thermalization. *Phys. Rev. E*, 50 (1994), pp. 888–901.
[4] Pal, A., and D.A. Huse, D., Many-body localization phase transition. *Phys. Rev. B* 82 (2010), 174411.
[5] Luitz, D. J., N. Laflorencie, and F. Alet, Many-body localization edge in the random-field Heisenberg chain. *Phys. Rev. B* 91 (2015), 081103.
[6] Bar Lev, Y., G. Cohen, and D.R. Reichman, Absence of diffusion in an interacting system of spinless Fermions on a one-dimensional disordered lattice. *Phys. Rev. Lett.*, 114 (2015), 100601.
[7] Luitz, D. J., N. Laflorencie, and F. Alet, F., Extended slow dynamical regime close to the many-body localization transition. *Phys. Rev. B* 93 (2016), 060201.
[8] Rigol, M., V. Dunjko, and M. Olshanii, Thermalization and its mechanism for generic isolated quantum systems. *Nature*, 452 (2008), pp. 854–858.
[9] Yu, X., Luitz, D. J. and Clark, B. K. "Bimodal entanglement entropy distribution in the many-body localization transition", arXiv:1606.01260 [http://arxiv.org/abs/1606.01260]

Mori, Warren B.

[1] Corde, S., et al., Multi-gigaelectronvolt acceleration of positrons in a self-loaded plasma wakefield, *Nature*, 524 (2015), p. 442.

[2] Davidson, A., Ph. D. Dissertation UCLA 2016.

[3] Lemos, N. et al., Self-modulated laser wakefield accelerators as x-ray sources. *PPCF* (2016) 58, 034018.

Schleife, Andre (1)

[1] Hlawacek, G., V. Veligura, R. van Gastel, and B. Poelsema, Helium ion microscopy. *J. Vac. Sci. Technol. B*, 32:2 (2014), 020801, doi: 10.1116/1.4863676.

[2] Schleife, A., E. W. Draeger, Y. Kanai, and A. A. Correa, Plane-wave pseudopotential implementation of explicit integrators for time-dependent Kohn-Sham equations in large-scale simulations. *J. Chem. Phys.* 137:22, (2012), 22A546, doi: 10.1063/1.4758792.

[3] Schleife, A., Y. Kanai, and A.A. Correa, Accurate atomistic first-principles calculations of electronic stopping. *Phys. Rev. B.*, 91(2015), 014306, doi: 10.1103/PhysRevB.91.014306

Schleife, Andre (2)

[1] Snaith, H., Perovskites: The emergence of a new era for low-cost, high-efficiency solar cells. *J. Phys. Chem. Lett.*, 4:21 (2013), pp. 3623-3630, doi: 10.1021/jz4020162.

[2] Onida, G., L. Reining,, and A. Rubio, Electronic excitations: density-functional versus many-body Green's-function approaches. *Rev. Mod. Phys.*, 74:2 (2002), pp. 601-659, doi: 10.1103/RevModPhys.74.601.

[3] Bechstedt, F., K. Seino, P. Hahn, and W. G. Schmidt, Quasiparticle bands and optical spectra of highly ionic crystals: AlN and NaCl. *Phys. Rev. B*, 72 (2005), 245114, doi:10.1103/PhysRevB.72.245114.

[4] Saidaminov, M. I. et al., High-quality bulk hybrid perovskite single crystals within minutes by inverse temperature crystallization. *Nat. Commun.*, 6:7586 (2015), doi:10.1038/ncomms8586.

[5] Juarez-Perez, E. J. et al., Photoinduced giant dielectric constant in lead halide perovskite solar cells. *J. Phys. Chem. Lett.*, 5 (2014), 2390, doi:10.1021/jz5011169

[6] Shirayama, M. et al., Optical transitions in hybrid perovskite solar cells: Ellipsometry, density functional theory, and quantum efficiency analyses for $CH_3NH_3PbI_3$. *Phys. Rev. Applied*, 5 (2016), 014012, doi:10.1103/PhysRevApplied.5.014012.

Subramaniam, Shankar

[1] Sun, B., S. Tenneti, and S. Subramaniam, Modeling gas-solid heat transfer using particle-resolved direct numerical simulation. *Int. J. Heat Mass Transfer*, 86 (2015), pp. 898–913.

[2] Tenneti, S., and S. Subramaniam, Particle-resolved direct numerical simulation for gas–solid flow model development. *Annu Rev Fluid Mech.*, 46 (2014), pp. 199-230.

[3] Sun, B., Modeling Heat and Mass Transfer in Reacting Gas-solid Flow Using Particle-resolved Direct Numerical Simulation. PhD thesis, Iowa State University, 2016.

[4] Sun, B., S. Tenneti, S. Subramaniam, D.L. Koch, Pseudo-turbulent heat flux and average gas-phase conduction during gas-solid heat transfer: flow past random fixed particle assemblies. *J. Fluid Mech.*, (2016), doi:10.1017/jfm.2016.290

Sugar, Robert

[1] Bazavov, A., et al., Leptonic decay-constant ratio fK/fpi from lattice QCD with physical light quarks. *Phys. Rev. Lett.* 110 (2013),172003 [arXiv:1301.5855].

[2] Bazavov, A., et al., Determination of |V_{us}| from a lattice-QCD calculation of the K -> pi l nu semileptonic form factor with physical quark masses, *Phys. Rev. Lett.* 112 (2014), 112001 [arXiv:1312.1228].

[3] Bazavov, A., et al., Charmed and light pseudoscalar meson decay constants from four-flavor lattice QCD with physical light quarks, *Phys. Rev.*, D90 (2014) 074509 [arXiv:1407.3772].

[4] Bazavov, A., et al., Decay constants fB and fBs from HISQ simulations. *Proceedings of Science* (Lattice 2015), 14 -18 July 2015, Kobe, Japan arXiv:1511.02294

[5] Clark, M.A., R.G. Edwards and B. Joo, IPDS '14, Proceedings of the A Framework for Lattice QCD Calculations on GPUs, F.T. *IEEE 28th International Parallel and Distributed Processing Symposium*, pp. 1073-1082, IEEE Computer Society, Washington, DC.,Winter, 2014

[6] Dudek, J.J., R.G. Edwards, C.E. Thomas, and D.J. Wilson. Resonances in coupled pi K - eta K scattering from quantum chromodynamics. *Phys. Rev. Lett.* 113 (2014), 182001, [arXiv:1406.4158 [hep-ph]].

[7] Wilson, D.J., et al., Coupled pi-pi, K K scattering in P-wave and the rho resonance from lattice QCD, *Phys. Rev.* D92 (2015), 094502, [arXiv:1507.02599 [hep-ph]].

[8] Briceno, R.A., et al., The resonant pi gamma -> pi pi amplitude from Quantum Chromodynamics, *Phys. Rev. Lett.* 115 (2015), 24200, [arXiv:1507.06622 [hep-ph]].

[9] Bolton, D.R., R. A. Briceno, and D. J. Wilson, Connecting physical resonant amplitudes and lattice QCD, *Phys. Lett. B* 757 (2016), pp. 50-56. arXiv:1507.07928 [hep-ph]].

Thomas, Brian

[1] Steel Statistical Yearbook 2014, World Steel Association, Brussels, Belgium, 2014.

[2] Singh, R., B. G. Thomas, and S. P. Vanka, Large eddy simulations of double-ruler electromagnetic field effect on transient flow during continuous casting. *Metall. Mater. Trans. B*, 45B (2014), pp. 1098-1115.

[3] Cho, S-M., B G. Thomas, and S-H. Kim, Transient two-phase flow in slide-gate nozzle and mold of continuous steel slab casting with and without double-ruler electro-magnetic braking. *Metall. Mater. Trans. B*, (2016) pp. 1-19.

[4] Cho, S-M., S-H. Kim, and B. G. Thomas, Transient Fluid Flow during Steady Continuous Casting of Steel Slabs: Part I. Measurements and Modeling of Two-phase Flow. *ISIJ Int.*, 54:4 (2014), pp. 845-854.

[5] Jin, K., B.G. Thomas, and X. Ruan, Modeling and measurements of multiphase flow and bubble entrapment in steel continuous casting., *Metall. Mater. Trans. B* , 47B:1 (2016), pp. 548-565. DOI: 10.1007/s11663-015-0525-5

[6] Jin, K., S. P. Vanka, B.G. Thomas, and X. Ruan, Large eddy simulations of the effects of double-ruler electromagnetic braking and nozzle submergence depth on molten steel flow in a commercial continuous casting mold. *TMS Annual Meeting, CFD Modeling and Simulation in Materials Processing Symposium*, Nashville, TN, Mar. 14-18, 2016, TMS, Warrendale, PA, 2016, pp. 159-166.

[7] Yan, X., A. Jonayat, B. G. Thomas, Thermal-fluid model of meniscus behavior during mold oscillation in steel continuous casting, *TMS Annual Meeting, Nashville*, TN, Mar. 14-18, 2016, TMS, Warrendale, PA, 2016, pp. 181-186

Yeung, P.K.

[1] Frisch, U. *Turbulence. The Legacy of A.N. Kolmogorov*. Cambridge University Press (1995).

[2] Sawford, B.L., Turbulent relative dispersion. *Annu. Rev. Fluid Mech.*, 33 (2001), pp. 289-317.

[3] Gotoh, T. and P.K. Yeung, Passive scalar transport: a computational perspective. In *Ten Chapters in Turbulence*, Cambridge University Press (2013).

[4] Gotoh, T. Hatanaka, S. and Miura, H.Spectral compact difference hybrid computation of passive scalar in isotropic turbulence. *J. Comput. Phys.*, 231 (2012), pp. 7398-7414.

COMPUTER SCIENCE & ENGINEERING

Deng, Jia

[1] Cisco visual networking index: Forecast and methodology, 2013-2018. http://www.cisco.com/c/ en/us/solutions/collateral/service-provider/ip-ngn-ip-next-generation-network/white_paper_ c11-481360.pdf, 2014.

[2] Mnih, V. *et al.*, Playing Atari with deep reinforcement learning, arXiv preprint arXiv:1312.5602 (2013).

[3] Krizhevsky, A., I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. *NIPS*. (2012)

[4] Szegedy, C., W. Liu, Y. Jia, P. Sermanet, S. Reed, et al., Going deeper with convolutions. *CVPR*. (2015)

[5] Russakovsky, O., J. Deng, H. Su, J. Krause, S. Satheesh, et al., Imagenet large scale visual recognition challenge. *International Journal of Computer Vision* 115(3) (2015) 211–252

[6] Chao, Y-W., Z. Wang, Y. He, J. Wang, J. Deng. HICO: A Benchmark for Recognizing Human-Object Interactions in Images. *ICCV*. (2015)

[7] Xu, C., C. Xiong, J. J. Corso, Streaming hierarchical video segmentation, *European Conference on Computer Vision* (Springer, Firenze, Italy, Oct. 7-13, 2012), pp. 626–639.

[8] Xu, C., J. J. Corso, Actor-action semantic segmentation with grouping process models, *IEEE Conference on Computer Vision and Pattern Recognition* (IEEE Computer Society, Las Vegas, Nevada, June 26-July 1, 2016).

[9] Kuhn, H., The Hungarian method for the assignment problem. *Nav. Res Logist. Q.*, 2 (1955), pp. 83-97.

[10] Koopmans, T. and M. Beckmann, Assignment problems and the location of economic activities. *Econometrica*, 25:1 (1957), pp. 53-76.

Gurel, Levent

[1] Gürel, L., and Ö. Ergül, Hierarchical parallelization of the multilevel fast multipole algorithm (MLFMA), *Proc. IEEE*, 101:2 (2013), pp. 332–341.

[2] Ergül, O, and L. Gürel, *The Multilevel Fast Multipole Algorithm (MLFMA) for Solving Large-Scale Computational Electromagnetics Problems*. Wiley-IEEE Press, 2014. (ISBN: 978-1-119-97741-4)

[3] Velamparambil, S., and W. C. Chew, Analysis and performance of a distributed memory multilevel fast multipole algorithm, *IEEE Trans. Antennas Propag.*, 53:8 (2005), pp. 2719–2727.
[4] Gürel, L., and Ö. Ergül, Fast and accurate solutions of extremely large integral-equation formulations discretised with tens of millions of unknowns. *Electron. Lett.*, 43:9 (2007) pp. 499–500.
[5] Kirk, D.B., and W. W. Hwu, *Programming Massively Parallel Processors*. Burlington, MA, USA: Morgan Kaufmann, 2010.

Nagi, Rakesh

[1] Adams, W., et al., A level-2 reformulation-linearization technique bound for the Quadratic Assignment Problem. *Eur. J. Oper. Res.*, 180:3 (2007), pp. 983-996.
[2] Burkard, R., S. Karisch, and F. Rendl, QAPLIB - A Quadratic Assignment Problem Library. *J. Global Optim.*, 10:4 (1997), pp. 391-403.

Sanders, William

[1] Krishna, V.D, G. A. Weaver, and W. H. Sanders, PCA-based method for detecting integrity attacks on advanced metering infrastructure. *Proceedings of QEST 2015*, Springer International Publishing.
[2] Krishna, V.D., K. Ravishankar, K. Iyer and W.H. Sanders. ARIMA-based modeling and validation of consumption readings in power grids. *Proceedings of CRITIS 2015*, Springer International Publishing.
[3] Hyndman, R.J., and Y. Khandakar, Automatic time series forecasting: The forecast package for R7. *J. Stat. Soft.*, 27:3 (2008), p. 1–22.

Schutt-Aine, Jose

[1] J. E. Schutt-Aine, "Latency Insertion Method for the Fast Transient Simulation of Large Networks," *IEEE Trans. Circuit Syst.*, 48, pp. 81-89, (2001).
[2] K. S. Yee, " Numerical solution of initial boundary value problems involving Maxwell's equations in isotropic media", *IEEE Trans. Antennas Propagat.*, 14, pp. 302-307 (1966).

BIOLOGY, CHEMISTRY, & HEALTH

Aksiementiev, Aleksei

[1] Gibcus, J. H. and J. Dekker, The hierarchy of the 3D genome. *Molecular Cell*, 49 (2013), pp. 773–82, doi:10.1016/j.molcel.2013.02.011.
[2] Dekker, J., M. A. Marti-Renom, and L. A. Mirny, Exploring the three-dimensional organization of genomes: Interpreting chromatin interaction data. *Nat. Rev. Genet.*, 14 (2013), pp. 390–403, doi:10.1038/nrg3454.
[3] Ngo, T. T. M., Q. Zhang, R. Zhou, J. Yodh, and T. Ha, Asymmetric unwrapping of nucleosomes under tension directed by DNA local flexibility. *Cell*, 160 (2015), pp. 1135–1144, doi:10.1016/j.cell.2015.02.001.
[4] Jenuwein T., and C. Allis, Translating the histone code. *Science*, 293 (2001), pp. 1074–80, doi: 10.1126/science.1063127.

Amaro, Rommie

[1] Sharma, M. *et al.*, Insight into the Mechanism of the Influenza A Proton Channel from a Structure in a Lipid Bilayer. *Science*, 330:6003 (2010), pp. 509-512.
[2] Gamblin, S. J., and J. J. Skehel, Influenza Hemagglutinin and Neuraminidase Membrane Glycoproteins. *J. Biol. Chem.*, 285:37 (2010), pp. 28403-28409.
[3] Baum, E. Z. *et al.*, A point mutation in influenza B neuraminidase confers resistance to peramivir and loss of slow binding. *Antiviral Res.*, 59:1 (2003), pp. 13-22.
[4] Le, Q. M. *et al.*, Avian flu: isolation of drug-resistant H5N1 virus. *Nature*, 437:7062 (2005), pp. 1108.
[5] LeGoff, J. *et al.*, I223R Mutation in Influenza A(H1N1) pdm09 Neuraminidase Confers Reduced Susceptibility to Oseltamivir and Zanamivir and Enhanced Resistance with H275Y. *PLoS One*, 7:8 (2012), pp. e37095.
[6] Harris, A. *et al.*, Influenza virus pleiomorphy characterized by cryoelectron tomography. *Proc. Natl. Acad. Sci. U. S. A.*, 103:50 (2006), pp. 19123-19127.
[7] Durrant, J. D., and R. E. Amaro, LipidWrapper: an algorithm for generating large-scale membrane models of arbitrary geometry. *PLoS Comput. Biol.*, 10:7 (2014), pp. e1003720.

Hirata, So

[1] Hirata, K. et al., *Ab initio* ice, dry ice, and liquid water. In *Fragmentation: Toward Accurate Calculations on Complex Molecular Systems*, Ed by Mark S. Gordon, in review.
[2] Willow, S.Y., M. A. Salim, K. S. Kim, and S. Hirata, *Ab initio* molecular dynamics of liquid water using embedded-fragment second-order many-body perturbation theory towards its accurate property prediction. *Scientific Reports*, 5 (2015), doi:10.1038/srep14358.
[3] Salim, M.A., S. Y. Willow, and S. Hirata, Ice Ih anomalies: Thermal contraction, anomalous volume isotope effect, and pressure-induced amorphization. *J. Chem. Phys.*, 144, (2016), doi: 10.1063/1.4951687).
[4] Mishima, O., L. D. Calvert, and E. Whalley, Melting ice at 77° K and 10 kbar: A new method of making amorphous solids. *Nature* 310 (1982), p. 393–395.

Jakobsson, Eric

[1] Köhler, G., and Cesar Milstein, Continuous cultures of fused cells secreting antibody of predefined specificity. *Nature* 256 (1975), pp. 495-497.
[2] Riechmann, L., M. Clark, H. Waldmann, and G. Winter. Reshaping human antibodies for therapy. *Nature* 332:6162 (1988), pp. 323-327.
[3] Takada, A., et al., A system for functional analysis of Ebola virus glycoprotein. *PNAS*, 94:26 (1997), pp.14764-14769.
[4] Jeffrey L.E., et al., Structure of the Ebola virus glycoprotein bound to an antibody from a human survivor. *Nature* 454:7201 (2008), pp. 177-182.
[5] Corti, D., and A. Lanzavecchia. Broadly neutralizing antiviral antibodies. *Annu. Rev. Immunol.*, 31 (2013), pp. 705-742.
[6] Murin, C. D. et al., Structures of protective antibodies reveal sites of vulnerability on Ebola virus. *PNAS*, 111:48 (2014), p. 17182-17187.
[7] Henikoff, S., and J.G. Henikoff, Amino acid substitution matrices from protein blocks. *PNAS*, 89:22 (1992, pp. 10915-10919.
[8] Moreira, I.S., P. A. Fernandes, and M. J. Ramos, Hot spots—A review of the protein–protein interface determinant amino-acid residues. *Proteins: Struct., Funct., Bioinf.* 68:4 (2007), p. 803-812.
[9] Karplus, M., and J. A. McCammon, Molecular dynamics simulations of biomolecules. *Nat. Struct. Mol. Biol.*, 9:9 (2002), pp. 646-652.

Kline, Micheal

[1] Hille, B., *Ion Channels of Excitable Membranes*. 3rd Ed. Sinauer Associates, Inc. Sunderland, MA (2001).
[2] Clapham, D.E., TRP channels as cellular sensors. *Nature*, 426 (2003), p. 517–524.
[3] Owsianik, G., D. D'hoedt, T. Voets and B. Nilius, Structure–function relationship of the TRP channel superfamily. In *Reviews of Physiology, Biochemistry and Pharmacology*, S.G. Amara, eds. Springer Berlin Heidelberg, (2006), pp. 61–90.
[4] Caterina MJ, et al. (1997) The capsaicin receptor: a heat-activated ion channel in the pain pathway. *Nature*, 389:6653 (1997), pp. 816-824.
[5] Menigoz, .A., and M Boudes, The expression pattern of TRPV1 in brain. *J. Neurosci.*, 31:37 (2011), pp. 3025-13027.
[6] Zhang F, et al., Identification of a tetrameric assembly domain in the C terminus of heat-activated TRPV1 channels. *J. Biol. Chem.*, 286:17 (2011), pp. 15308-15316.
[7] Holzer, P., The pharmacological challenge to tame the transient receptor potential vanilloid-1 (TRPV1) nocisensor. *Br. J. Pharmacol.*, 155:8 (2008), pp. 1145-1162.
[8] Nilius, B., TRP channels in disease. *BBA Mol. Cell. Bio. Lett.*, 1772:8 (2007), pp. 805-812.
[9] Vriens, J., G. Appendino, and B. Nilius, Pharmacology of vanilloid transient receptor potential cation channels. *Mol. Pharmac.*, 75:6 (2009), pp. 1262-1279.
[10] Nilius, B. and A. Szallasi, Transient receptor potential channels as drug targets: from the science of basic research to the art of medicine. *Pharmacol. Rev*. 66:3 (2014), pp. 676-814.

Luthey-Schulten, Zaida

[1] Earnest, T.M., et al., Toward a whole-cell model of ribosome biogenesis: kinetic modeling of SSU assembly, *Biophys. J.*, 109:6 (2015), pp. 1117–1135.
[2] Earnest, T.M., et al., Ribosome biogenesis in replicating cells: Integration of experiment and theory. *Biopolymers*, (in press).
[3] Roberts, E., J. Stone, and Z. Luthey-Schulten, Lattice microbes: High-performance stochastic simulation method for the reaction-diffusion master equation. *J. Comput. Chem.*, 34:3 (2013), pp. 245–255.
[4] Hallock, M., et al., Simulation of reaction diffusion processes over biologically-relevant size and time scales using multi-GPU workstations. *Parallel Comput.* 40:5–6 (2014), pp. 86–99.
[5] Peterson, J., M. Hallock, J. Cole, and Z. Luthey-Schulten, A problem solving environment for stochastic biological simulations. *PyHPC '13: Proceedings of the 3rd Workshop on Python for High-Performance and Scientific Computing* (ACM, Denver, CO, USA, November 18, 2013).
[6] Hallock, M., and Z. Luthey-Schulten, Improving reaction kernel performance in Lattice Microbes: particle-wise propensities and run-time generated code. *15th IEEE International Workshop on High Performance Computational Biology* (IEEE, Chicago, IL, USA, May 22–23, 2016).
[7] Mulder, A. M., et al., Visualizing ribosome biogenesis: parallel assembly pathways for the 30S subunit. *Science*, 330:6004 (2010), pp. 673–677.

[8] Jones, D.L., R. C. Brewster, and R. Phillips, Promoter architecture dictates cell-to-cell variability in gene expression. *Science*, 346:6216 (2014), pp. 1533–1536.
[9] Peterson, J. R., et al., Effects of DNA replication on mRNA noise. *PNAS*, 112:52 (2015), pp. 15886–15891.
[10] Roberts, E., et al., Noise Contributions in an Inducible Genetic Switch: A whole-cell simulation study. *PLoS Comput. Biol.*, 7:3 (2012). e1002010.

Makri, Nancy

[1] Lambert, R., and N. Makri, Quantum-classical path integral: I. Classical memory and weak quantum nonlocality. *J. Chem. Phys.*,137 (2012), 22A552, doi:10.1063/1.4767931.
[2] Lambert, R., and N. Makri, Quantum-classical path integral. II. Numerical methodology, *J. Chem. Phys.*, 137 (2012), 22A553, doi:10.1063/1.4767980.
[3] Banerjee, T., and N. Makri, Quantumclassical path integral with self-consistent solvent-driven propagators. *J. Phys. Chem.*, 117 (2013), 13357, doi:10.1021/jp4043123.
[4] Walters, P. L., and N. Makri, Iterative quantum-classical path integral with dynamically consistent state hopping. *J. Chem. Phys.*, 144 (2016), 044108, doi:10.1063/1.4939950.
[5] Feynman, R. P., Space-time approach to non-relativistic quantum mechanics. *Rev. Mod. Phys.*, 20 (1948), 367, doi:10.1103/RevModPhys.20.367.
[6] Makri, N., Exploiting classical decoherence in quantum dissipative dynamics. *Chem. Phys. Lett.*, 593 (2014), 93, doi:10.1016/j.cplett.2013.11.064.
[7] Azzouz, H., and D. Borgis, A quantum molecular-dynamics study of proton transfer reactions along asymmetrical H bonds in solution. *J. Chem. Phys.*, 98 (1993), 7361, doi:10.1063/1.464727.
[8] Makri, N., The linear response approximation and its lowest order corrections: an influence functional approach. *J. Phys. Chem.*,103 (1999), 2823, doi:10.1021/jp9847540.

Mainzer, Liudmila Sergeevna (1)

[1] DePristo, M., *et al.*, A framework for variation discovery and genotyping using next-generation DNA sequencing data. *Nature Genetics*, 43 (2011), pp. 491-498.

Mainzer, Liudmila Sergeevna (2)

[1] DePristo, M., *et al.*, A framework for variation discovery and genotyping using next-generation DNA sequencing data. *Nature Genetics*, 43 (2011), pp. 491-498.
[2] Tarasov A., Vilella A.J., Cuppen E., Nijman I.J., Prins P. Sambamba: fast processing of NGS alignment formats. *Bioinformatics*, 31:12 (2015), pp. 2032-4.
[3] Faust G.G., Hall I.M., SAMBLASTER: fast duplicate marking and structural variant read extraction. *Bioinformatics*, 30:17 (2014), pp. 2503-2505.
[4] Kufrin R., PerfSuite: An Accessible, Open Source Performance Analysis Environment for Linux. *6th International Conference on Linux Clusters: The HPC Revolution 2005*. Chapel Hill, NC. April 2005.
[5] http://docs.cray.com/books/S-2315-50/html-S-2315-50/z1055157958smg.ht
[6] Showerman, M., *et al.*, (2014) Large Scale System Monitoring and Analysis on Blue Waters using OVIS. *Proceedings of Cray User Group Meeting (CUG 2014)*
[7] http://valgrind.org/docs/manual/manual.html
[8] http://www.novocraft.com/ accessed

Masud, Arif

[1] Kwack, J, and A. Masud, A three-field formulation for incompressible viscoelastic fluids, *Int. J. Eng. Sci.*, 48 (2010), pp. 1413-1432.
[2] Masud, A., and J. Kwack, A stabilized mixed finite element method for the incompressible shear-rate dependent non-Newtonian fluids: Variational multiscale framework and consistent linearization. *Comput. Meth. App. Mech. Eng.*, 200 (2011), pp. 577-596.
[3] Masud A., and R. Calderer, "A variational multiscale method for incompressible turbulent flows: bubble functions and fine scale fields" *Comput. Meth. App. Mech. Eng.*, 200 (2011), pp. 2577-2593.
[4] Anand, M., J. Kwack, and A. Masud, A new generalized Oldroyd-B model for blood flow in complex geometries", *Int. J. Eng. Sci.*, 72 (2013), pp. 78-88.
[5] Calderer, R., and A. Masud, Residual-based variational multiscale turbulence models for unstructured tetrahedral meshes, *Comput. Meth. App. Mech. Eng.*, 254 (2013), pp. 238-253.
[6] Masud, A., and R. Calderer, Residual-based turbulence models for moving boundary flows: Hierarchical application of VMS and three-level scale separation", *Int. J. Num. Meth. Fl.*, 73 (2005), pp. 284-305.
[7] Kwack, J., and A. Masud, A stabilized mixed finite element method for shear-rate dependent non-Newtonian fluids: 3D benchmark problems and application to blood flow in bifurcating arteries, *Comput. Mech.*, 53 (2014), pp. 751-776.

Pande, Vijay

[1] Huang, W. J., et al., Structural insights into mu-opioid receptor activation. *Nature*, 524:7565 (2015), pp. 315-321.

Schulten, Klaus (1)

[1] Goh, B. C., et al., Computational methodologies for real-space structural refinement of large macromolecular complexes. *Annu. Rev. Biophys.*, 45 (2016), pp. 253–278.
[2] Perilla, J. R., et al., Molecular dynamics simulations of large macromolecular complexes. *Curr. Opin. Struct. Biol.*, 31 (2015), pp. 64–74.
[3] Schweitzer, A., et al., The structure of the human 26S proteasome at a resolution of 3.9 Å. *PNAS*, (2016).
[4] Zhang Y., et al., Recognition of poly-ubiquitins by the proteasome through protein re-folding guided by electrostatic and hydrophobic interactions. *J. Phys. Chem. B*, (2016). DOI: 10.1021/acs.jpcb.6b01327.
[5] Ribeiro, J. V., et al., QwikMD - Integrative Molecular Dynamics Toolkit for Novices and Experts. *Sci. Rep.*, 6 (2016), 26536.

Schulten, Klaus (2)

[1] Goh, B. C, et.al., Computational Methodologies for Real-Space Structural Refinement of Large Macromolecular Complexes. *Annu. Rev. Biophys.* 45 (2016), 062215–011113.
[2] Zhao, G., et al., Mature HIV-1 capsid structure by cryo-electron microscopy and all- atom molecular dynamics. *Nature*, 497:7451 (2013), pp. 643-646.
[3] Lu, M., et al., Dynamic allostery governs cyclophylin A-HIV capsid interplay. *PNAS*, 112: 47 (2015), pp. 14617-14622.
[4] Liu, C., et al., Cyclophilin A stabilizes the HIV-1 capsid through a novel non-canonical binding site. *Nat. Commun.*, 7 (2016), pp. 10714-10724.
[5] Goh, B. C., et al., Atomic modeling of an immature retroviral lattice using molecular dynamics and mutagenesis. *Structure*, 23 (2015), pp. 1414-1425.
[6] Heyrana, K., et al., Contributions of charged residues in structurally dynamic capsid surface loops to Rous sarcoma virus assembly. *J. Virol*, 90:12 (2016), pp. 5700-5714.
[7] Perilla, J.R., et. al., All-Atom Molecular Dynamics of Virus Capsids as Drug Targets. *J. Phys. Chem. Lett.* 7 (2016), pp. 1836–1844.
[8] Cassidy, C. K., et al., CryoEM and computer simulations reveal a novel kinase conformational switch in bacterial chemotaxis signaling. *eLife*, (2015), 10.7554/eLife.08419.

Skukla, Diwakar

[1] Ray, D. K., N. Ramankutty, N. D. Mueller, P. C. West, and J. A. Foley, Recent patterns of crop yield growth and stagnation, *Nature Communication*s, vol. 3, no. 1293, 2012.
[2] Tilman, D., C. Balzer, J. Hill, and B. L. Befort, Global food demand and the sustainable intensification of agriculture, *PNAS*, USA, vol. 108, no. 50, pp. 20260– 20264, 2011.
[3] Long, S. P., A. Marshall-Colon, and X.-G. Zhu, Meeting the global food demand of the future by engineering crop photosynthesis and yield potential, *Cell*, vol. 161, no. 1, pp. 56–66, 2015.
[4] Wheeler, T. and J. von Braun, Climate change impacts on global food security, *Science*, vol. 341, no. 6145, pp. 508–513, 2013.
[5] De Smet, I., U. Voss, G. Jurgens, and T. Beeckman, Receptor-like kinases shape the plant," *Nature Cell Biology*, vol. 11, no. 10, pp. 1166–1173, 2009.
[6] Zhu, J.-Y., J. Sae-Seaw, and Z.-Y. Wang, Brassinosteroid signalling, *Development*, vol. 140, no. 8, pp. 1615–1620, 2013.

Sutton, Brad

[1] Fang, H.C.H, In *Neurobiology of Aging*, Raven Press, New York, ed. 1, 1976, pp. 155–166.
[2] Cassot, F., et al., A novel three-dimensional computer-assisted method for a quantitative study of microvascular networks of the human cerebral cortex. *Microcirculation*, 13 (2006), pp. 1–18.
[3] Lauwers, F., et al., Morphometry of the human cerebral cortex microcirculation: General characteristics and space-related profiles. *NeuroImage*. 39 (2008), pp. 936–948.
[4] Le Bihan, D., et al., Separation of diffusion and perfusion in intravoxel incoherent motion MR imaging. *Radiology*. 168 (1988), pp. 497–505.
[5] Pruessmann, K.P., M. Weiger, M. B. Scheidegger, and P. Boesiger, SENSE: Sensitivity encoding for fast MRI. *Magn. Reson. Med.*, 42 (2009), pp. 952–962.
[6] Holtrop, J.L., and B. P. Sutton, High spatial resolution diffusion weighted imaging on clinical 3 T MRI scanners using multislab spiral acquisitions. *J. Med. Imaging*, 3.2 (2016), 023501-023501.
[7] Deng, W., B. Zahneisen, and V. A. Stenger, *Magn. Reson. Med.*, in press, doi:10.1002/mrm.25863.
[8] Fu, M., et al., High-frame-rate full-vocal-tract 3D dynamic speech imaging. *Magn. Reson. Med.*, (2016), doi:10.1002/mrm.26248.

[9] Fu, M., et al., High-resolution dynamic speech imaging with joint low-rank and sparsity constraints. *Magn. Reson. Med.*,73 (2014), pp. 1820–1832.

Tagkopoulos, Ilias

[1] Carrera, J., et al., An integrative, multi-layer, genome-scale model reveals the phenotypic landscape of Escherichia coli", *Mole. Syst. Biol.*, 10:7 (2014), p. 735.
[2] Mozhayskiy, V. and I. Tagkopoulos, Guided evolution of in silico microbial populations in complex environments accelerates evolutionary rates through a step-wise adaptation, *BMC Bioinformatics*, (2012)
[3] Tagkopoulos, I., Y. Liu, and S. Tavazoie, Predictive Behavior within microbial genetic networks, *Science*, 320(2008), pp.1313-7.

Voth, Gregory

[1] Dama, J., et al., The Theory of ultra-coarse-graining. 1. General Principles. *J. Chem. Theory Comput.* 9 (2013) pp. 2466-2480. doi:10.1021/ct4000444
[2] Grime, J., and G. Voth, Highly scalable and memory efficient ultra-coarse-grained molecular dynamics simulations. *J. Chem. Theory Comput.* 10 (2013) pp. 423-431. doi:10.1021/ct400727q

Warnow, Tandy

[1] Mirarab, S., et al., PASTA: Ultra-large multiple sequence alignment for nucleotide and amino-acid sequences. *J. Comput. Biol.*, 22:5 (2014), pp. [377-386].
[2] Nguyen, N., S. Mirarab, K. Kumar, and T. Warnow. Ultra-large alignments using phylogeny-aware profiles. *Genome Biol.*, 16.1:124 (2015), pp. [1-15].
[3] Stamatakis, A. RAxML-VI-HPC: maximum likelihood-based phylogenetic analyses with thousands of taxa and mixed models. *Bioinformatics*, 22:21(2006), pp [2688–2690].
[4] Price, M., P. Dehal, and A. Arkin. FastTree 2–approximately maximum-likelihood trees for large alignments. *PLoS One*, 5:3(2010), p. [9490].
[5] Mirarab, S., and T. Warnow. ASTRAL-II: coalescent-based species tree estimation with many hundreds of taxa and thousands of genes. *Bioinformatics*, 31:12 (2015), pp. [i44–i52].
[6] Nguyen, N., T. Warnow, M. Pop, and B. White. A perspective on 16S rRNA operational taxonomic unit clustering using sequence similarity. *Nature Biofilms and Microbiome Analysis*, 2, article number 16004 (2016), doi:10.1038/npjbiofilms.2016.4.
[7] Redelings, B., and M. Suchard. Joint Bayesian estimation of alignment and phylogeny. *Systemic Biology*. 54:3(2005), pp. [401–418].
[8] Nguyen, N., et al., TIPP: taxonomic identification and phylogenetic profiling. *Bioinformatics*, 30:24(2014), pp. [3548–3555].
[9] Altschul, S., et al., Basic local alignment search tool. *J. Mol. Biol.*, 215:3(1990), pp. [403–410].
[10] Eddy, S. Profile hidden Markov models. *Bioinformatics*, 14:9(1998), pp. [755–763].

SOCIAL SCIENCE, ECONOMICS, & HUMANITIES

Hansen, Lars

[1] Cai, Y., K.L. Judd, and T.S. Lontzek, The social cost of carbon with economic and climate risks, *J. Polit. Econ.*, arXiv preprint arXiv:1504.06909, (2015)
[2] Nordhaus, W.D., *A Question of Balance: Weighing the Options on Global Warming Policies*. Yale University Press. (2008)
[3] SSC, IWG, *Social Cost of Carbon for Regulatory Impact Analysis under Executive Order 12866*. United States Government. Appendix 15a, Washington DC. (2010).
[4] Lontzek, T.S., Y. Cai, K.L. Judd, and T.M. Lenton, Stochastic integrated assessment of climate tipping points indicates the need for strict climate policy. *Nat. Cli. Change*, 5 (2015), pp. 441–444.
[5] Cai, Y., et al., Environmental tipping points significantly affect the cost-benefit assessment of climate policies. *PNAS*, 112:15 (2015), p. 4606–4611.
[6] Cai, Y., T.M. Lenton, and T.S. Lontzek (2016). Risk of multiple climate tipping points should trigger a rapid reduction in CO2 emissions. *Nat. Cli. Change*, 6 (2016), p. 520–525.
[7] Cai, Y., K.L. Judd, and J. Steinbuks, *A Nonlinear Certainty Equivalent Approximation Method for Dynamic Stochastic Problems*. No. w21590. National Bureau of Economic Research (2015).
[8] Yeltekin, S., Y. Cai, and K. L. Judd, Computing equilibria of dynamic games. *Oper. Res.*, in review
[9] Cai, Y., et al., The effect of climate and technological uncertainty in crop yields on the optimal path of global land use. *J. Env. Econ. Mgmt.*, World Bank Policy Research Working Paper 7009, in review.
[10] Cai, Y., K.L. Judd, G. Thain, and S. Wright, Solving dynamic programming problems on a computational grid. *Comput. Econ.*, 45:2 (2015), pp. 261–284.

GRADUATE FELLOWS

Allen, Thomas

[1] Heathcote, P., P. Fyfe, and M. Jones, Reaction centres: the structure and evolution of biological solar power. *Trends Biochem. Sci.*, 27:2 (2002), pp. 79-87.
[2] Kalyanasundaram, K., and M. Graetzel, Artificial photosynthesis: biomimetic approaches to solar energy conversion and storage. *Curr. Opin. Biotech.*, 21:3 (2010), pp. 298-310.
[3] Tully, J. C., Theories of the Dynamics of Inelastic and Reactive Processes at Surfaces. *Annu. Rev. Phys. Chem.*, 31:1 (1980), pp. 319-343.
[4] Lambert, R., and M. Nancy, Quantum-classical path integral. I. Classical memory and weak quantum nonlocality. *J. Chem. Phys.*, 137:22 (2012), pp. 22A552.
[5] Lambert, R., and M. Nancy, Quantum-classical path integral. II. Numerical Methodology. *J. Chem. Phys.*, 137:22 (2012), pp. 22A553.
[6] Plimpton, S., Fast Parallel Algorithms for Short-Range Molecular Dynamics. *J. Comp. Phys.*, 117 (1995), pp. 1-19.
[7] Marchi, M., J. N. Gehlen, D. Chandler, and M. Newton, Diabatic Surfaces and the Pathway for Primary Electron Transfer in a Photosynthetic Reaction Center. *J. Am. Chem. Soc.*, 115:10 (1993), pp. 4178-4190.
[8] Makri, N., E. Sim, D. E. Makarov, and M. Topaler, Long-time quantum simulation of the primary charge separation in bacterial photosynthesis. *Proc. Natl. Acad. Sci. USA*, 93:9 (1996), pp. 3926-3931.
[9] Treutlein, H., et al., Molecular Dynamics Simulation of the Primary Processes in the Photosynthetic Reaction Center of Rhodopseudomonas viridis. in *The Photosynthetic Bacterial Reaction Center: Structure and Dynamics*, (1988), pp. 139-150.

Brooks, Maureen

[1] Fine, M. L., Faunal variations on pelagic *Sargassum. Mar. Biol.*, 7 (1970), pp. 112-122.
[2] Hoffmayer, E. R., et al., Larval and juvenile fishes associated with pelagic *Sargassum* in the northcentral Gulf of Mexico. *50th Gulf Caribb. Fish. Inst.*, (2005), pp. 259-269.
[3] Gower, J. F. R. and S. A. King, Distribution of floating *Sargassum* in the Gulf of Mexico and the Atlantic Ocean mapped using MERIS, *Int. J. Remote Sens.*, 32:7 (2011), pp. 1917-1929.
[4] Gower, J., E. Young, and S. King, Satellite images suggest a new *Sargassum* source region in 2011. *Remote Sens. Lett.*, 4:8 (2013) pp. 764-773.
[5] Chassignet E., et al., US GODAE: Global Ocean Prediction with the HYbrid Coordinate Ocean Model (HYCOM). *Oceanography*, 22:2 (2009) pp. 64-75.
[6] Fennel, K., R. Hetland, Y. Feng, and S. DiMarco, A coupled physical-biological model of the Northern Gulf of Mexico shelf: model description, validation and analysis of phytoplankton variability. *Biogeosciences*, 8:7 (2011) pp. 1881-1899.
[7] Schoener, A., and G. T. Rowe, Pelagic Sargassum and its presence among the deep-sea benthos. *Deep-Sea Res.*, 17 (1970), pp. 923-925.

Calhoun, Jon

[1] Cappello, F., et al., Toward exascale resilience. *Int. J. of High Perf. Comput. App.*, (2009), doi: 10.1177/1094342009347767
[2] Son, S.W., et al., Data compression for the exascale computing era - survey. *Supercomp. Front. Inno.*, 1:2 (2014), pp. 76-88 .
[3] Di, S., and C. Franck. Fast Error-bounded Lossy HPC Data compression with SZ. In *Proceedings of the 2016 IEEE International Parallel and Distributed Processing Symposium*, IPDPS '16., Washington, DC, USA, 2016. IEEE Computer Society.
[4] Lindstrom, P., and M. Isenburg. Fast and efficient compression of floating-point data. *IEE Trans. Vis. Comput. Graph.*, 12:5 (2006), p. 1245–1250.

Drake, Justin

[1] Wright, P.E., and H.J. Dyson, Intrinsically unstructured proteins: re-assessing the protein structure-function paradigm. *J. Mol. Biol.*, 293 (1999), pp. 321–331.
[2] Dunker, A.K., et al., Intrinsic protein disorder in complete genomes. *Genome Inform. Workshop Genome Inform.*, 11 (2000), pp. 161–171.
[3] Uversky, V.N., C. Santambrogio, S. Brocca, and R. Grandori, Length-dependent compaction of intrinsically disordered proteins. *FEBS Lett.*, 586 (2012), pp. 70–73.
[4] Auton, M., and D.W. Bolen, Additive transfer free energies of the peptide backbone unit that are independent of the model compound and the choice of concentration scale. *Biochemistry (Mosc.)*, 43 (2004), pp. 1329–1342.
[5] Bondos, S.E., X.-X. Tan, and K.S. Matthews, Physical and Genetic Interactions Link Hox Function with Diverse Transcription Factors and Cell Signaling Proteins. *Mol. Cell. Proteomics*, 5 (2006), pp. 824–834.

[6] Best, R.B., et al., Optimization of the Additive CHARMM All-Atom Protein Force Field Targeting Improved Sampling of the Backbone φ, ψ and Side-Chain χ_1 and χ_2 Dihedral Angles. *J. Chem. Theory Comput.*, 8 (2012), pp. 3257–3273.
[7] Case, D.A., et al., Amber 12. (2012).
[8] Teufel, D.P., et al., Backbone-Driven Collapse in Unfolded Protein Chains. *J. Mol. Biol.*, 409 (2011), pp. 250–262.
[9] Karandur, D., R.C. Harris, and B.M. Pettitt, Protein collapse driven against solvation free energy without H-bonds. *Protein Sci. Publ. Protein Soc.*, 25 (2016), pp. 103–110.
[10] Drake, J.A., and B.M. Pettitt, Force field-dependent solution properties of glycine oligomers. *J. Comput. Chem.*, 36 (2015), pp. 1275–1285.

Kokkila Schumacher, Sara

[1] Geertsen, J., M. Rittby, and R. J. Bartlett, The equation-of-motion coupled-cluster method: Excitation energies of Be and CO. *Chem. Phys. Lett.*, 164, pp.57-62 (1989).
[2] Christiansen, O., H. Koch, and P. Jørgensen, The second-order approximate coupled cluster singles and doubles model CC2. *Chem. Phys. Lett.*, 243, pp. 409-418 (1995).
[3] Hohenstein, E.G., R. M. Parrish, and T. J. Martínez, Tensor hypercontraction density fitting. I. Quartic scaling second- and third-order Møller-Plesset perturbation theory. *J. Chem. Phys.*, 137, pp. 044103:1-10 (2012).
[4] Hohenstein, E.G., S. I. L. Kokkila, R. M. Parrish, and T. J. Martínez, Quartic scaling second-order approximate coupled cluster singles and doubles via tensor hypercontraction: THC-CC2. *J. Chem. Phys.*, 138, pp. 124111:1-10 (2013).
[5] Hohenstein, E.G., S. I. L. Kokkila, R. M. Parrish, and T. J. Martínez, Tensor Hypercontraction Equations-of-Motion Second-Order Approximate Coupled Cluster: Electronic Excitation Energies in O(N4) Time. *J. Phys. Chem. B*, 117, pp. 12972-12978 (2013).
[6] Ufimtsev, I.S., and T. J. Martínez, Quantum Chemistry on Graphical Processing Units. 1. Strategies for Two-Electron Integral Evaluation. *J. Chem. Theory Comput.*, 4, pp. 222-231 (2008).

Mendez, Joshua

[1] Lacks, D. J., and R. M. Sankaran (2011), Contact electrification of insulating materials, *J. Phys. Appl. Phys.*, 44(45), 453001, doi:10.1088/0022-3727/44/45/453001.
[2] Thomas, R. J., P. R. Krehbiel, W. Rison, H. E. Edens, G. D. Aulich, W. P. Winn, S. R. McNutt, G. Tytgat, and E. Clark (2007), Electrical Activity During the 2006 Mount St. Augustine Volcanic Eruptions, *Science*, 315(5815), 1097–1097, doi:10.1126/science.1136091
[3] Forward, K. M., D. J. Lacks, and R. M. Sankaran (2009), Charge Segregation Depends on Particle Size in Triboelectrically Charged Granular Materials, *Phys. Rev. Lett.*, 102(2), 028001, doi:10.1103/PhysRevLett.102.028001.
[4] James, M. R., S. J. Lane, and J. S. Gilbert (2000), Volcanic plume electrification: Experimental investigation of a fracture-charging mechanism, *J. Geophys. Res. Solid Earth*, 105(B7), 16641–16649, doi:10.1029/2000JB900068.
[5] Lowell, J., and W. S. Truscott (1986), Triboelectrification of identical insulators. II. Theory and further experiments, *J. Phys. Appl. Phys.*, 19(7), 1281, doi:10.1088/0022-3727/19/7/018.
[6] Bichoutskaia, E., A. L. Boatwright, A. Khachatourian, and A. J. Stace (2010), Electrostatic analysis of the interactions between charged particles of dielectric materials, *J. Chem. Phys.*, 133(2), 024105, doi:10.1063/1.3457157.
[7] Méndez Harper, J., and J. Dufek (2016), The effects of dynamics on the triboelectrification of volcanic ash, *J. Geophys. Res. Atmospheres*, 2015JD024275, doi:10.1002/2015JD024275.
[8] Méndez Harper, J.S. et al., 2015. Triboelectric charging of Titan dune grains: Effect on sediment transport. *LPSC* 2015. Abstract 1637.

Minot, Ariana

[1] Cain, M., R. O'Neill, and A. Castillo, History of Optimal Power Flow and Formulations. *Federal Energy Regulatory Commission*, (2012), pp. 1-36.
[2] Palmer, B. et al. "GridPACK Overview" https://www.gridpack.org/wiki/images/7/71/GridPack_Overview_v4.pdf

Reames, Larrissa

[1] Shepherd, J. M., H. Pierce, and A. J. Negri, Rainfall modification by major urban areas: Observations from spaceborne rain radar on the TRMM satellite. *J. Appl. Meteor.*, 41:7 (2002), pp. 689– 701.
[2] Tan, Y. B., L. Peng, Z. Shi, and H. R. Chen, Lightning flash density in relation to aerosol over Nanjing (China). *Atmos. Res.*, 174-175 (2016), pp. 1–8.
[3] Changnon, S. A., Urban modification of freezing-rain events. *J. Appl. Meteor.*, 42: 6 (2003), pp. 863– 870.
[4] Smith, J. A., M. L. Baeck, J. E. Morrison, P. Sturdevant-Rees, D. F. Turner-Gillespie, and P. D. Bates, The regional hydrology of extreme floods in an urbanizing drainage basin. *J. Hydrometeorol.*, 3: 3 (2002), pp. 267–282.
[5] Skamarock, W. and J. Klemp, A time-split nonhydrostatic atmospheric model for weather research and forecasting applications. *J. Comput. Phys.*, 227:7 (2008), pp. 3465–3485.

Slota, George

[1] McSherry, F., I. Michael, and D. G. Murray, Scalability! but at what COST?. *15th Workshop on Hot Topics in Operating Systems, HotOS XV, USENIX*, Kartause Ittingen, Switzerland, May 18-20, 2015).
[2] Mattson, T., et al., Standards for graph algorithm primitives. arXiv preprint arXiv:1408.0393 (2014).
[3] Slota, G. M., S. Rajamanickam, and K. Madduri, High Performance Graph Analytics on Manycore Processors. *Proceedings of the 29th IEEE International Parallel and Distributed Processing Symposium* (IPDPS15), (IEEE, Hyderabad, India, May 25-29, 2015).

Totorica, Samuel

[1] Zweibel, E. G., and M. Yamada, Magnetic reconnection in astrophysical and laboratory plasmas. *Annu. Rev. Astron. Astrophys.*, 47 (2009), pp. 291–332.
[2] Fonseca, R. A., et al., Exploiting multi-scale parallelism for large scale numerical modelling of laser wakefield accelerators. *Plasma Phys. Controlled Fusion*, 55 (2013), p. 124011.
[3] Abel, T., O. Hahn, and R. Kaehler, Tracing the dark matter sheet in phase space. *Mon. Not. R. Astron. Soc.*, 427 (2012), pp. 61–76.

Zemskova, Varvara

[1] Santilli, E., and A. Scotti, The stratified ocean model with adaptive refinement (SOMAR). *J. Comput. Phys.*, 291 (2015), pp. 60-81. doi:10.1016/j.jcp.2015.03.008
[2] Zemskova, V.E., B.L. White, and A. Scotti, Available Potential Energy and the General Circulation: Partitioning Wind, Buoyancy Forcing, and Diapycnal Mixing. *J. Phys Oceanogr.*, 45:6 (2015), pp.1510-1531. http://dx.doi.org/10.1175/JPO-D-14-0043.1

INDEX

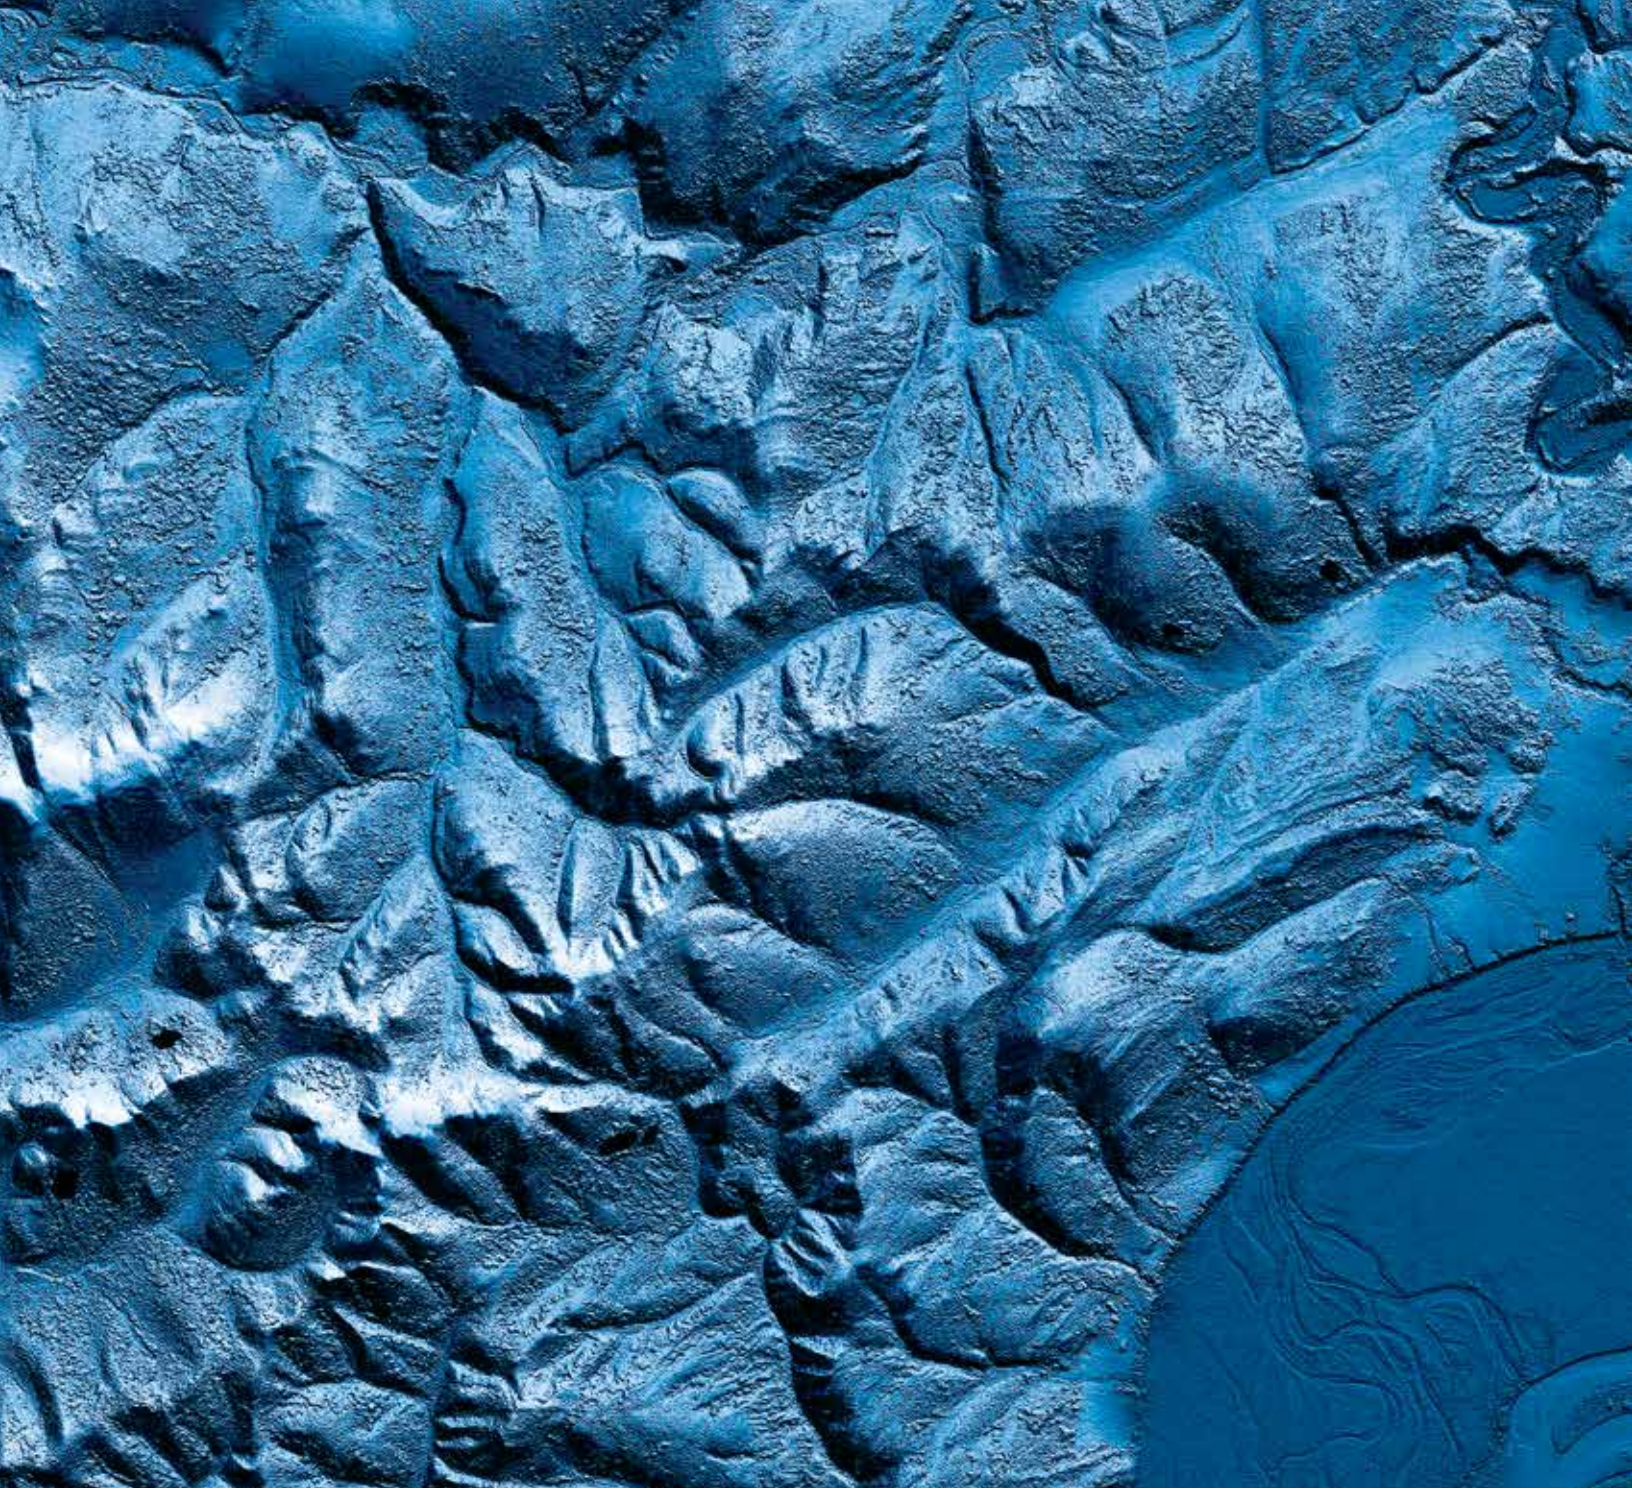

NCSA

ILLINOIS
UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

Blue Waters is supported by the National Science Foundation