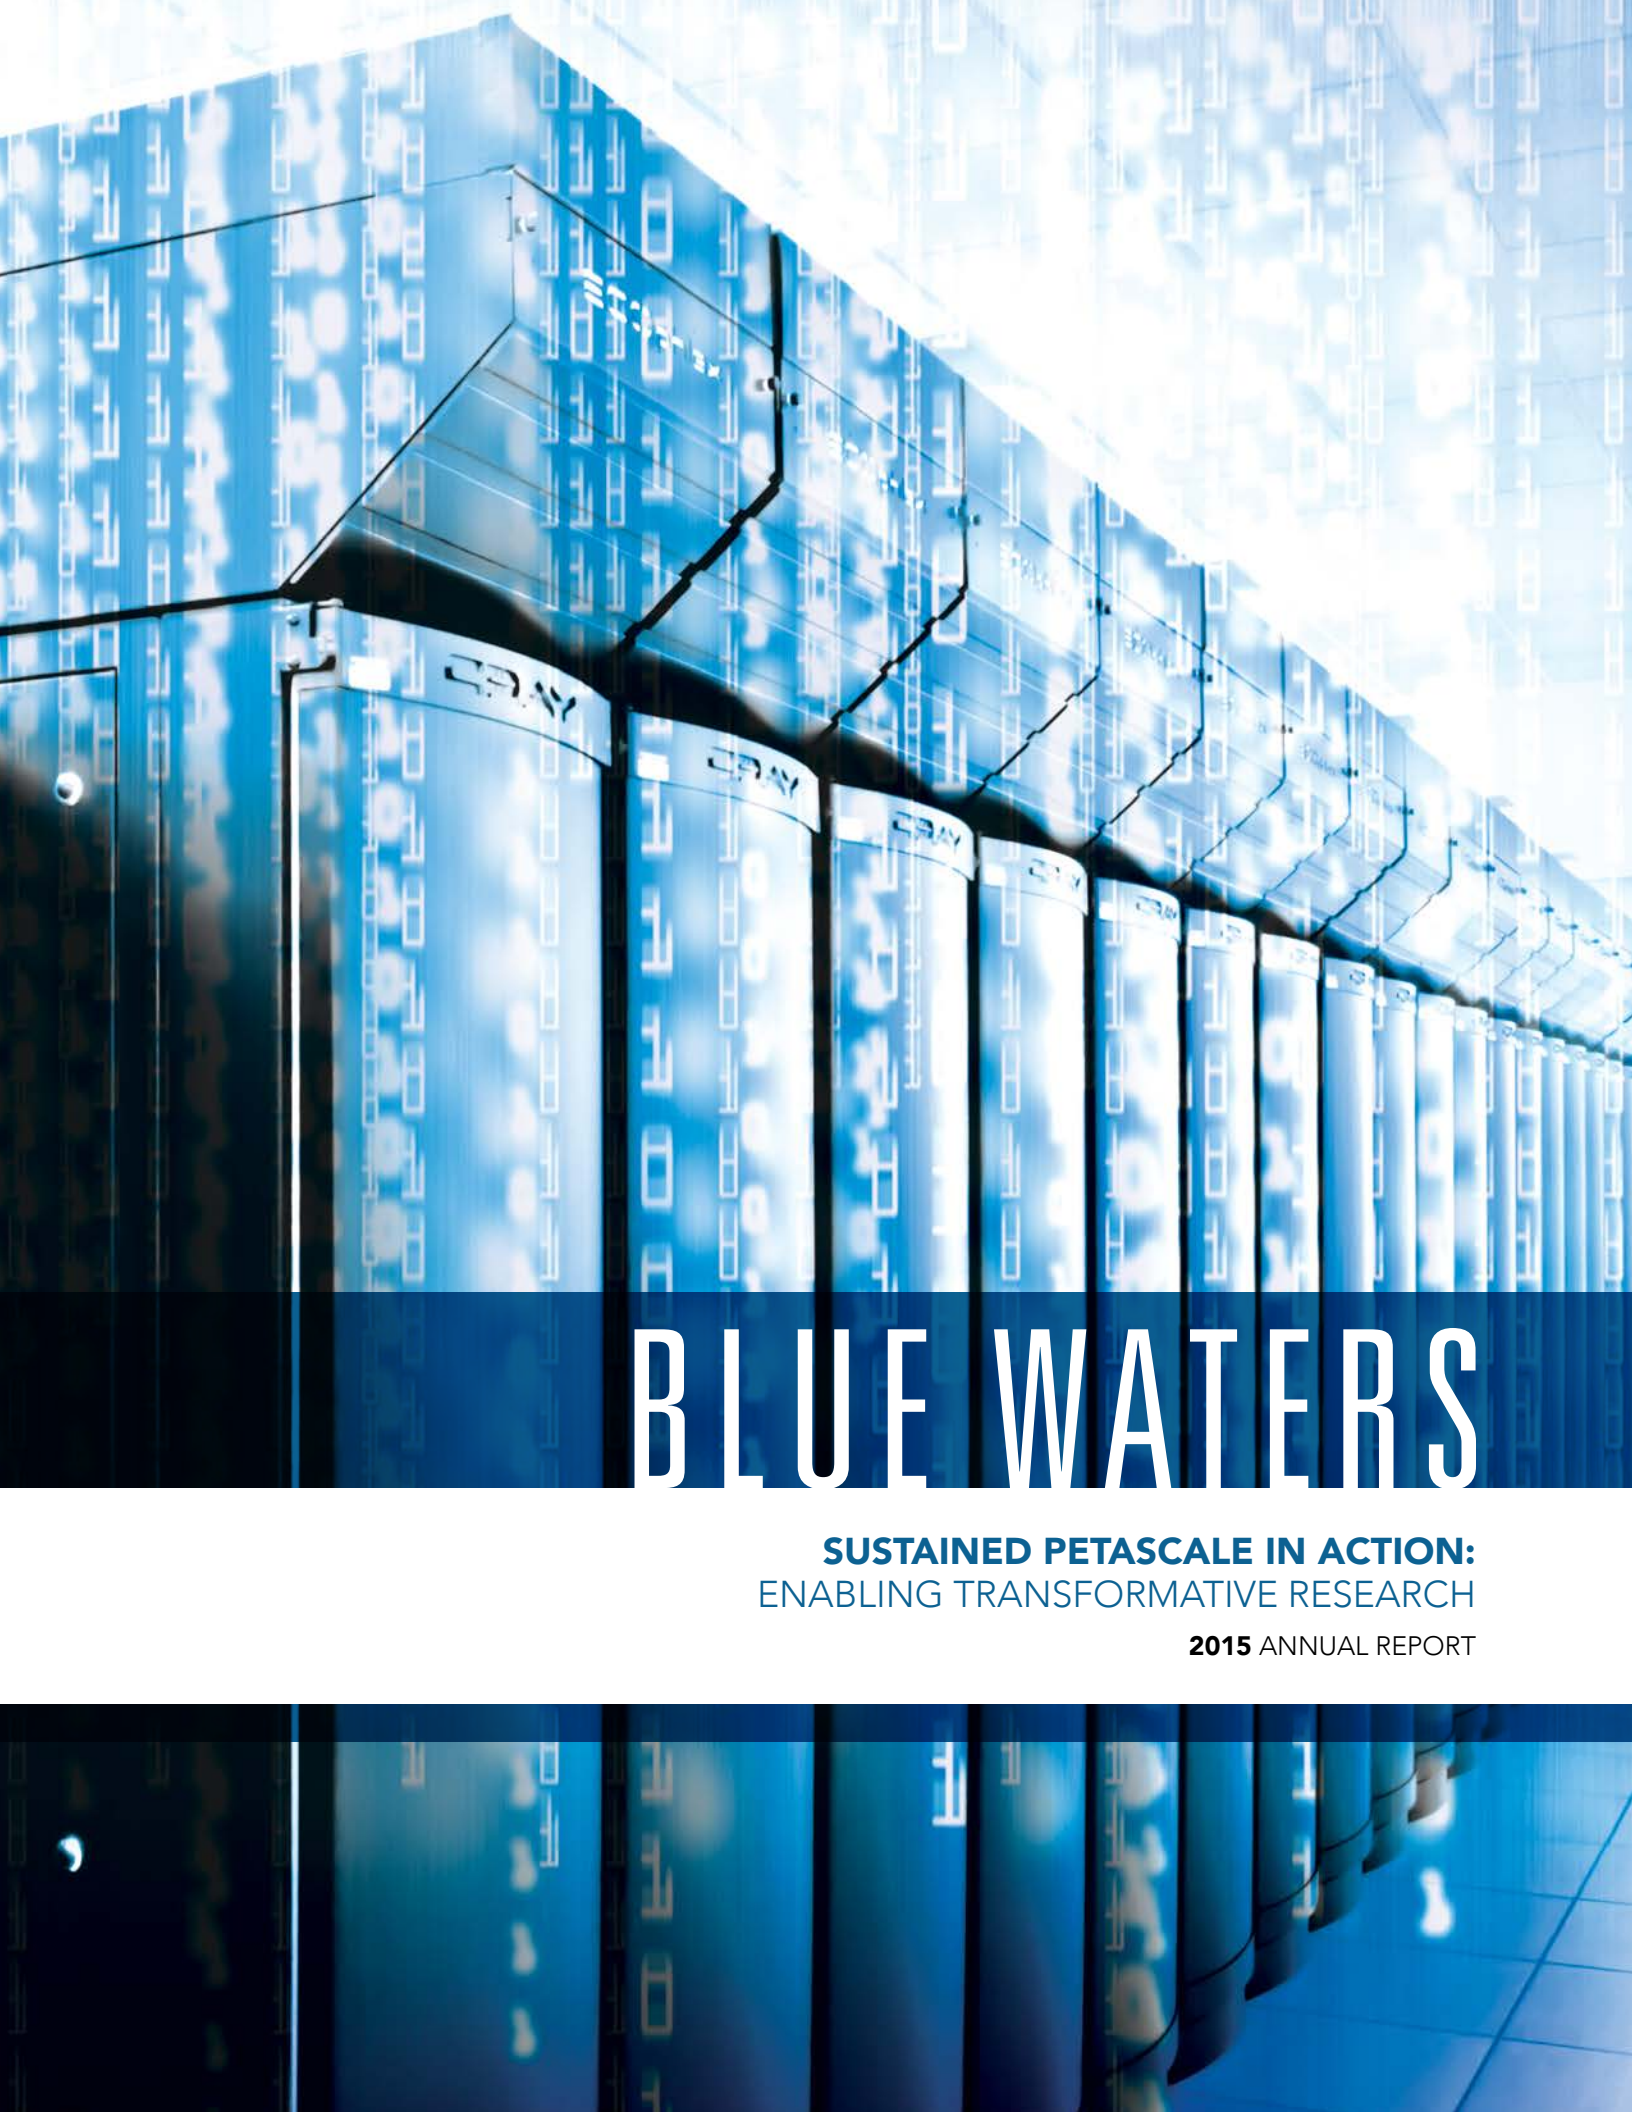

BLUE WATERS

SUSTAINED PETASCALE IN ACTION:
ENABLING TRANSFORMATIVE RESEARCH

2015 ANNUAL REPORT

BLUE WATERS

SUSTAINED PETASCALE IN ACTION:
ENABLING TRANSFORMATIVE RESEARCH

2015 ANNUAL REPORT

Editor
Nicole Gaynor

Art Director
Paula Popowski

Designer
Steve Duensing

Editorial Board
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 2015 Blue Waters Symposium presentations.

ISBN 978-0-9908385-3-1

THE MASSIVE BLUE WATERS SUPERCOMPUTER—CAPABLE OF PERFORMING QUADRILLIONS OF CALCULATIONS EVERY SECOND AND OF WORKING WITH QUADRILLIONS OF BYTES OF DATA—ENABLES BREAKTHROUGH DISCOVERY FOR RESEARCHERS ACROSS THE UNITED STATES, FROM BETTER UNDERSTANDING AND PREDICTION OF TORNADOES AND EARTHQUAKES, TO DESIGNING NEW MATERIALS, TO INVESTIGATIONS OF ALZHEIMER'S, HIV, AND CANCER.

TABLE OF CONTENTS

4 A MESSAGE FROM BILL KRAMER
6 BLUE WATERS SYMPOSIUM 2015
8 CADENS
10 COMMUNITY ENGAGEMENT & EDUCATION
12 BLUE WATERS EPSCOR STATES
14 EXTENDED ABSTRACTS
- 14 *space science*
- 38 *geoscience*
- 66 *physics & engineering*
- 120 *computer science & engineering*
- 126 *biology, chemistry & health*
- 176 *social science, economics, & humanities*
- 186 *graduate fellows*

208 MEASURING BLUE WATERS
214 BLUE WATERS PROFESSORS
216 OTHER BLUE WATERS PROJECTS
218 REFERENCES
232 INDEX

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 BILL KRAMER

Welcome to the Blue Waters Annual Report for 2015.

We are very excited about this report that documents many of the accomplishments and insights that over 150 science teams have achieved this year using Blue Waters for their modeling and simulation as well as the data science efforts. Many research teams from the first full service year continued to use Blue Waters to work with their data and their simulation, but we also found a number of new teams coming to use Blue Waters for their greatest challenges. This further increased the diversity of work on Blue Waters and expanded the areas of science and engineering to almost every realm. Some of these teams are doing projects that are national priorities at the highest levels such as research to understand the Ebola virus and to re-image the entire polar region of our planet. It turns out that Blue Waters may be the only system able to accomplish these things in a timely enough manner.

As I write this message, we have had 1410 science and engineering partners from 163 institutions and organizations in 33 states plus Puerto Rico using Blue Waters, the National Science Foundation's most powerful system for computation and data analysis. In just two years, Blue Waters has provided over 10 billion core hour equivalents to the national open research community. We looked at the federal funding (by PI and co-PI) and Blue Waters is supporting over $500 billion of funded research by providing computer processing time, data storage, and advanced intellectual services such as software and visualization assistance to help researchers make the most effective use of Blue Waters.

As you look just a few pages past this message, you will find exciting summaries of the research teams' work as well as other information. I encourage you to look for important phases in these reports such as *"never possible before," "first of a kind,"* and *"order of magnitude(s)."* But that is only part of the story. Without Blue Waters, these research problems would have taken far longer to complete, or maybe even been impossible to study.

We improved our systems and services this year for better use of Blue Waters and to make our science teams more productive. Perhaps the biggest improvement was to implement automated Topology Aware Scheduling on the system so applications can express the desired configuration of nodes to reduce their communication overhead and have the system provide that to the jobs in a very timely manner. As I write, we are finishing an upgrade of the wide area networking and cyber protection to multiple 100 Gbps pathways, connecting to all the major research networks. This will provide faster and more consistent network data transfers. This past year we also created and implemented the Blue Waters Data Sharing Pilot Service, to help teams and communities that want to share their data and result files with others. This service is designed to allow sharing of petabytes of data, a feature not easily found elsewhere.

Later in this report, you will find a special section about the project's positive engagements with people and organizations in EPSCoR states. But this is only part of the story about education and developing the world's best and most competitive workforce. For example, in the 2014-15 service year covered in this report, almost 60 percent of the partner users on Blue Waters were students and postdocs. Last year we started the Blue Waters Graduate Fellowships for PhD students who are using highly parallel computing as part of their research. We selected 16 students and they are making tremendous progress in their research projects. More importantly, almost all of the graduate fellows now believe they will stay involved with HPC after graduating. You will also learn more about the Blue Waters Undergraduate Internships, with 21 interns funded in the last year (3 from minority-serving institutions, 6 from institutions in EPSCoR jurisdictions). They join the 22 interns we have funded previously.

A current total of 52 students at a cost of almost $2 million have had funded research and development projects through the Blue Waters Project. Blue Waters Education Allocations (over 60M core hour equivalents) continue to increase in use and popularity. Blue Waters funded (with the XSEDE program) the creation of course content for the HPC University where 30 Blue Waters-developed undergraduate course modules were created for use in high school and college classes with peer reviewed content. There have been over 20,000 downloads and since August 1, 2014, between ~5-600 visits per day, >7,000 pages accessed daily; 700,000 pages accessed in all. Blue Waters started the Virtual School for Computational Science and Engineering (with our sub-awardee University of Michigan) that started with one or two distributed summer workshops and has trained over 2,000 people. In the last two years, we expanded this program to include semester long, credit bearing, distributed classes on extreme scale topics such as highly parallel computing, accelerator computing and advanced data and visualization. Last, but not least, you will see information about the new CADENS award associated with Blue Waters and XSEDE to bring science and engineering understanding to the general public through very high quality documentaries based on real science results. The first production, Solar Superstorms, is now released and relies on the impressive work of several Blue Waters teams.

The Blue Waters Project has taken a leadership role in an increasing set of activities within the government and alongside it regarding the debate about the need for continuing government investments in computation and data analysis resources and skills, working with many of the science and engineering teams using Blue Waters. We continued working with community software projects, both open source and other consortiums, such as OpenSFS, HPSS, and Hackathons. We have also increased working with partners in the U.S. (XSEDE and the Great Lakes Consortium for Petascale Computing) and internationally with such activities as the Joint Laboratory for Extreme Scale Computing and the Global Initiative to Enhance @scale and distributed Computing and Analysis Technologies to address grand challenge problems around the world.

Finally, I would like to acknowledge and thank all the people and organizations that help Blue Waters have such an amazing impact. First, many thanks to the Blue Waters staff who literally work day and night to make Blue Waters the system of choice for the nation's open science petascale research. Our staff is amazing in addition to being extremely talented people and dedicated to the project. Next, our supporters and funders at the National Science Foundation, in particular our program manager Irene Qualters, the state of Illinois, and the University of Illinois, without whom there would not be a Blue Waters. Also, the people who give us help, advice, and support such as the members of the Science and Engineering Technical Advisory Committee, our colleagues in peer centers and facilities and all our vendor and science partners.

We hope you enjoy reading all the wonderful advances Blue Waters helped make happen. Please let us know if you have any questions and please do not wait for next year to see what is happening on Blue Waters. Visit our portal at bluewaters.ncsa.illinois.edu to see up to the minute information.

With great respect,

Dr. William T.C. Kramer
Blue Waters Project Director
and Principal Investigator

BLUE WATERS SYMPOSIUM 2015

RESEARCHERS SHARE SUCCESSES AND METHODS AT 2015 BLUE WATERS SYMPOSIUM

The 2015 Blue Waters Symposium, held May 10-13 at Oregon's beautiful Sunriver Resort, brought together leaders in petascale computational science and engineering to share successes and methods.

Around 130 attendees, many of whom were Blue Waters users and NCSA staff who support their work, enjoyed presentations on computational advances in a range of research areas—including sub-atomic physics, weather, biology, astronomy, and many others—as well as keynotes from innovative thinkers and leaders in HPC.

Over the three days of the symposium, 58 science teams from across the country presented on their work on Blue Waters.

"This event showcases the importance of a computational resource of this size," said NCSA Director Ed Seidel. "Thanks to Blue Waters, detailed simulations of complex astrophysical phenomena, HIV, earthquake events, and industrial engineering processes are being done, leading to major scientific breakthroughs or new products that cannot be achieved any other way."

Seidel also reminded attendees of the importance of what is next.

"Experiments and observation systems are now generating unprecedented amounts of data," he said. "A new type of highly integrated environment must be developed where computing, experiment, and data services will need to be developed together."

TALK OF TRACK-1

William T. C. Kramer, director and principal investigator of the Blue Waters project, gave the opening keynote on the need for sustained petascale (and more) computing and data analysis. Recently, Kramer participated in two community BrainstormHPCD workshops to identify the on-going requirements among open science and engineering communities for future high performance computational and data analysis (HPCD) resources and services. He also briefly highlighted some of the significant science successes realized in year two of Blue Waters.

Irene Qualters, the National Science Foundation (NSF) Division of Advanced Cyberinfrastructure director, spoke later that same day on where she sees the future of supercomputing heading and the paths it may take along the way.

Satoshi Matsuoka, professor at the Global Scientific Information and Computing Center of Tokyo Institute of Technology (GSIC), believes demands for extreme computing and huge data processing are leading to a future with an inevitable convergence of the two infrastructures.

"Supercomputing is not just about providing the hardware resources; it's about the software, the consulting, the libraries, the applications, and it is even about maintaining the open source. You get all that with Blue Waters and the cloud lacks pretty much everything except the low-level hardware," he explained. "That is why I believe you would find cloud vendors are eager to talk convergence."

Steve Scott, senior vice president and chief technology officer for Cray, wrapped up the symposium discussion of the future of big compute with his keynote on programming and technology for the next decade.

In the past few decades simulation has become more and more important, and now every area of science has problems they can only address with simulation; Blue Waters gives them one of the most powerful tools in the world to do their work, as Scott emphasized.

"We need to start teaching programming as a parallel exercise from day one, so students learn parallel programming first, versus learning serial programming and then later trying to convert it," said Scott. "They need to be told what is expensive versus what is not expensive. Adding two numbers, not expensive; moving data from one place to another, expensive."

NCSA plans to continue these conversations and connections next year at the 2016 Blue Waters Symposium and lead large-scale computing through the petascale and beyond.

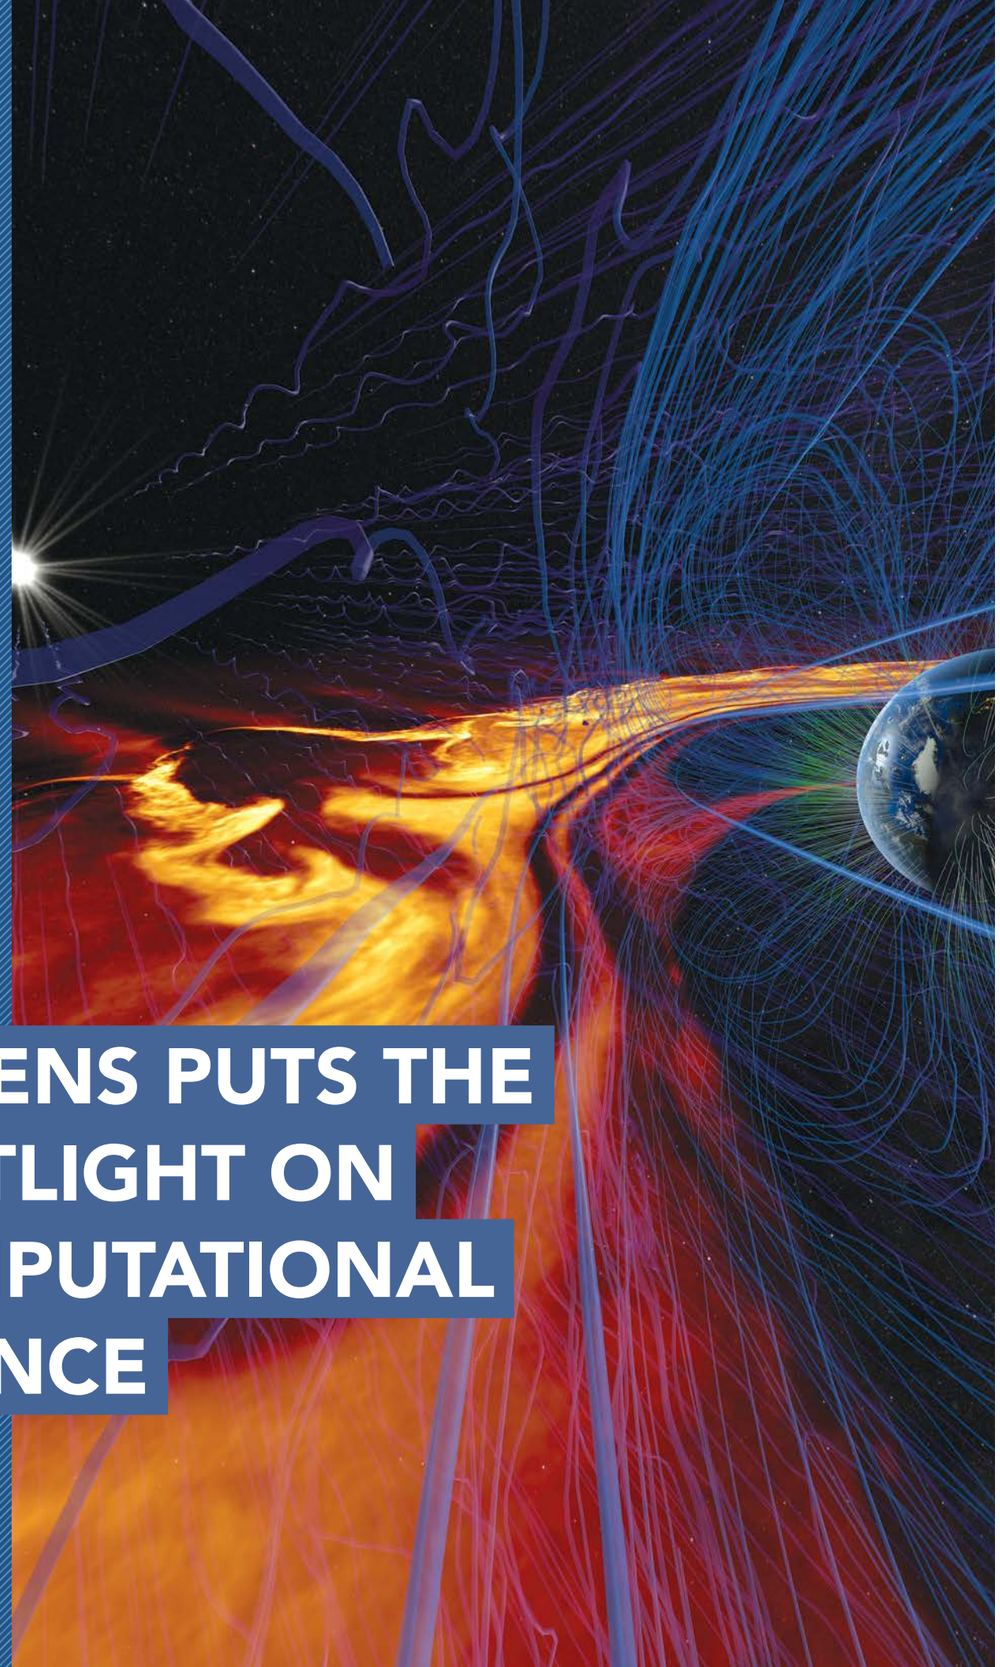

CADENS PUTS THE SPOTLIGHT ON COMPUTATIONAL SCIENCE

In fall 2014, the National Science Foundation (NSF) awarded $1.5 million to an effort to raise awareness and appreciation of the key role that computational and data-driven research plays in discovery. The Centrality of Advanced Digitally ENabled Science (CADENS) has already released the first in a series of planned science documentaries—"Solar Superstorms."

"The ability to generate and analyze giant data sets has produced a paradigm shift across all fields of inquiry. Compelling visualizations of these data will connect with broad audiences and provide insight into a wide range of natural phenomena," says CADENS principal investigator Donna Cox, director of the Advanced Visualization Lab (AVL) at the National Center for Supercomputing Applications (NCSA) and professor in Art and Design at the University of Illinois at Urbana–Champaign. AVL has created cinematic data visualizations seen by millions of people at museums, planetariums, and IMAX theaters.

"Our hope is that through visualization and visual storytelling, the project will help educate the public about what computational science is and how it impacts our lives," said Rudolf Eigenmann, a program director in the advanced cyber-infrastructure division at NSF who oversees the award.

CADENS first production is a 24-minute, high-resolution documentary about the dynamics of the Sun. Narrated by acclaimed actor Benedict Cumberbatch, "Solar Superstorms" asks the question, "What can cause our normally benign sun to erupt in such fury that it can threaten the world's power and technological infrastructure?" To answer that question, this innovative show takes audiences into the inner workings of our star. It follows the path of hot magnetized gas from deep inside the Sun, along its tangled journey through the Sun's churning outer layers, and on to explosive magnetic eruptions so powerful they can affect the Earth.

Telling this complex story involved collaboration with researchers and visualization experts across the country, in addition to the core CADENS team that includes NCSA's AVL, Thomas Lucas Productions, Spitz Creative Media, Blue Waters project director William Kramer, and John Towns, leader of the National Science Foundation's Extreme Science and Engineering Discovery Environment (XSEDE) project.

Included in the documentary is a visualization of data produced by Homa Karimabadi, Mahidhar Tatineni and Vadim Roytershteyn of the University of California, San Diego, who used Blue Waters to simulate the solar wind interacting with Earth's magnetic field during a powerful solar storm. Such large disturbances can result in loss of communication satellites and blackouts.

John Wise, Georgia Tech, used Blue Waters to simulate the early universe (from ~20 million to 500 million years after the Big Bang), showing the formation of the first generations of stars as their light ionizes much of the space around them. For the visualizations in the documentary, three features were targeted: the dense filaments along which stars form; the large regions of gas that become ionized and heated as new stars begin emitting lots of ultraviolet light, and later cool off after those stars' lives end; and the heavier elements mixed into the gas after those dying stars explode as supernovae, forming small bubbles of high "metallicity density." A second supernova simulation examines the life and death of a single first-generation star, looking in detail at the processes that mix the heavy elements into their environment.

"Solar Superstorms" debuted June 30, 2015, at the Louisiana Art & Science Museum in Baton Rouge and July 4 at Montpellier Méditerranée Métropole in France. The show will roll out to more than a dozen planetariums and science centers around the world, including Planetarium Hamburg (Germany), the State Museum of Pennsylvania (Harrisburg), Tellus Science Museum (Cartersville, Georgia), and the Eugenides Planetarium (Athens).

CADENS will produce two more high-resolution digital documentaries for giant screen fulldome theaters and nine high-definition documentaries for online distribution via YouTube, Hulu, and other outlets.

The CADENS call for participation remains open. If a researcher generates, analyzes, or visualizes data and depends upon access to advanced cyberinfrastructure, the research and/or visualizations can be considered for inclusion in a future CADENS project. For more information, visit http://www.ncsa.illinois.edu/enabling/vis/cadens.

"We hope," said Tom Lucas, "to open people's eyes to whole new ways of seeing the world, and to show the passion of scientists on the verge of profound discoveries."

COMMUNITY ENGAGEMENT & EDUCATION

EDUCATION AND OUTREACH PROGRAMS PREPARE THE NEXT GENERATION OF COMPUTATIONAL RESEARCHERS

The Blue Waters system provides the computational science and engineering community with an effective, scalable, sustained petascale computing platform—a system that meets the current needs of the research community as well as providing a window to future computing technologies.

Achieving the full potential of the Blue Waters system, with all of its advanced technology components, requires additional knowledge on the part of current and future computational scientists and engineers. For example, they may need to know how to adapt and scale code to petascale systems, create complex workflows, manage large data collections, conduct large-scale scientific visualizations, and other related topics. NCSA is committed to working closely with the science and engineering community to enable them to take full advantage of the extraordinary capabilities of Blue Waters, including preparing the next generation of researchers and educators. These efforts are aimed at providing researchers with skills to use emerging HPC systems over the next decade to advance all fields of study.

Blue Waters has established a community engagement effort to prepare current and future generations of computational experts through a combination of training sessions, workshops, fellowships and internships, outreach to the community, and dissemination of teaching and learning materials for use by the community. The community engagement effort is broadening the impact of the Blue Waters project well beyond the community of research teams using Blue Waters.

Education Allocations - Education allocations are available upon request by the community to support training events, workshops, courses, and other innovative mechanisms for preparing current and future generations in the applications of petascale systems to advance scientific discovery.

Workshops - Blue Waters conducts workshops and webinars for all levels of users throughout the year covering a wide range of topics tailored to making effective use of the Blue Waters system.

Virtual School of Computational Science and Engineering – The VSCSE supports interested faculty in providing semester-long courses, offered on the web, to allow students to participate at multiple institutions across the country. The courses are offered as for-credit graduate level college courses, including a syllabus with learning outcomes, 40 hours of instruction, reading assignments, homework and exercises, and assessment of learning progress.

Student Internship Program - The Blue Waters Internship Program is designed to immerse undergraduate and graduate students in research projects that are aligned with using Blue Waters or XSEDE resources. The students are provided with introductory in-depth training to ensure they are familiar with the Blue Waters environments, and then complete a year-long research internship.

Repository of Learning Materials – Blue Waters provides access to education and training materials developed by practitioners to foster the development of a broader, well-educated community able to conduct computational science and engineering research using petascale technologies, resources, and methods via HPC University (www.hpcuniversity.org) and the Computational Science Education Reference Desk (http://nsdl.cserd.org).

Blue Waters Graduate Fellowship Program – The Blue Waters Graduate Fellowship program is modeled on the NSF Graduate Fellowships (GRFP). The program supports graduate students engaged in a program of study and research that is relevant to the utilization of Blue Waters resources in support of their research. Preference will be given to candidates engaged in a multidisciplinary research project that combines disciplines such as computer science, applied mathematics and computational science applications.

Blue Waters Graduate Fellowships provide Ph.D. students with a year of support, including a $38,000 stipend, up to $12,000 in tuition allowance, an allocation of up to 50,000 node-hours on the powerful Blue Waters petascale computing system, and funds for travel to a Blue Waters-sponsored symposium to present research progress and results.

Ten individuals were selected during the spring of 2014 as Fellows during the 2014-2015 period. These very exceptional students completed their year-long fellowship this summer. They have given talks and presented posters showcasing their research at the 2015 Symposium, and they attended the SC14 Conference in November to broaden their knowledge of HPC and their exposure to the applications of HPC in academia, government, and industry. The fellows have bonded with one another and benefitted from getting to know each other as well as the Blue Waters staff at NCSA.

The following findings of the fellows program have been reported by the external evaluators of the program:

- The fellows agreed that this fellowship program allowed them to find intersections between their science fields and HPC community.
- Fellows agreed that because of the financial support, the program allowed for independence in their research area.
- Fellows reported that the Blue Waters staff (point-of-contact person) was very responsive, cooperative, and helpful.

A second cadre of six fellows was selected during the spring of 2015. They attended the 2015 Blue Waters Symposium, and began their year of research in August 2015.

Blue Waters looks forward to recruiting a third cadre of fellows with a new call for applications in the fall of 2015.

BLUE WATERS IMPACT ON EPSCOR STATES

The Blue Waters project is committed to ensuring access for researchers across the United States in alignment with the National Science Foundation's Experimental Program to Stimulate Competitive Research (EPSCoR), which aims to strengthen research and education across the nation. The Blue Waters project facilitates broader participation through outreach to institutions and researchers across the nation, including participants in EPSCoR states.

- Blue Waters supports research in 17 of the 28 EPSCoR states, as well as Puerto Rico.
- 25 percent of science projects that use Blue Waters involve researchers and/or students from EPSCoR states, engaging and training more than 250 individuals so far.
- Blue Waters provides the equivalent of more than $29.3 million of benefits in terms of system time and funding to researchers, students, and others from EPSCoR state institutions.
- Over 1.4 billion core hours went to research teams and year-long undergraduate internships and graduate fellowships for 14 students at EPSCoR institutions.

This level of investment will increase over the next three years of the Blue Waters project.

In addition to projects using the high-performance computing capabilities of Blue Waters, the University of Illinois and NCSA:

- Conduct education and training programs to provide access to Blue Waters through workshops, classes, and other programming.
- Fund a year-long undergraduate internship program and a graduate fellowship program
- Created and support the Virtual School for Computational Science and Engineering, which provides training opportunities for students across the country.

Individuals from institutions in EPSCoR states are well represented in these programs, accounting for about 20 percent of the graduate fellows and undergraduate interns.

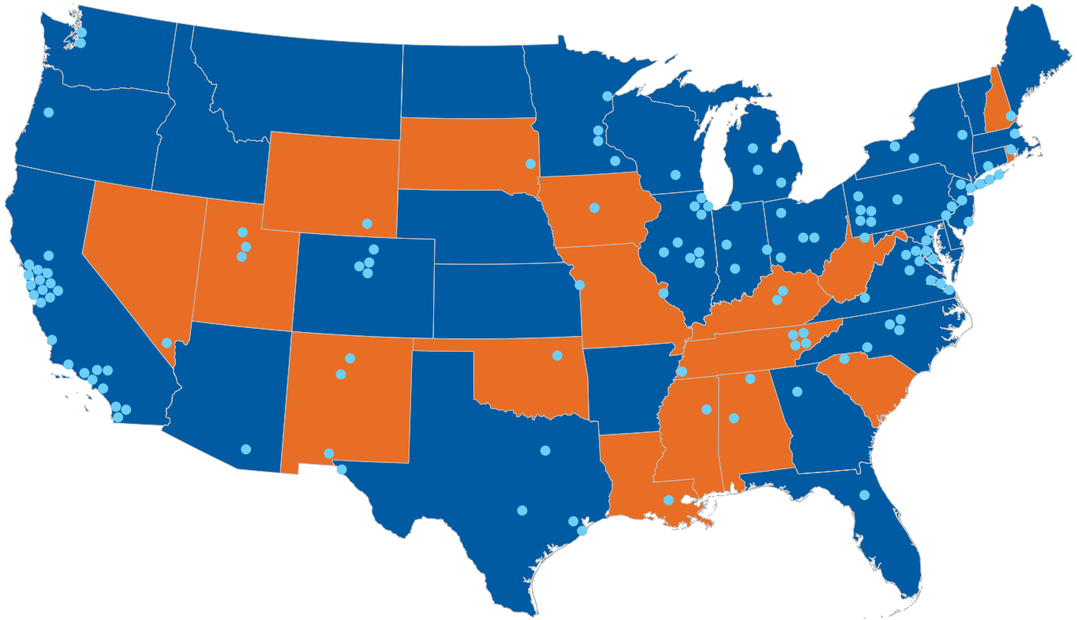

MAP OF PRINCIPAL INVESTIGATOR INSTITUTIONS.
Orange states indicate EPSCoR states with at least one Blue Waters Project principal investigator. Puerto Rico is not shown. The dots show the institutions for principal investigators leading one or more projects.

SPACE SCIENCE

ASTRONOMY

ASTROPHYSICS

COSMOLOGY

HELIOPHYSICS

16 *The First Galaxies and Quasars in the BlueTides Cosmological Simulation*

18 *Enabling Breakthrough Kinetic Simulations of the Magnetosphere*

22 *Core-Collapse Supernovae through Cosmic Time*

24 *Simulating Magnetized Plasma Turbulence from Macro to Micro Scales*

26 *Exploring the First Generations of Galaxies*

28 *Modeling Heliophysics Phenomena with a Multi-Scale Fluid-Kinetic Simulation Suite*

30 *Relativistic Jet Formation in Black Hole–Neutron Star Mergers*

32 *Ab-Initio Models of Solar Activity*

34 *3-D Simulations of i-Process Nucleosynthesis in the Early Universe*

36 *Three-Dimensional Core-Collapse Supernova Simulations*

THE FIRST GALAXIES AND QUASARS IN THE BLUETIDES COSMOLOGICAL SIMULATION

Allocation: NSF PRAC/2.625 Mnh
PI: Tiziana Di Matteo[1]
Co-PIs: Yu Feng[1,2], Rupert Croft[1]
Collaborators: Nick Battaglia[3], Steven Wilkins[1,4], Ananth Tenneti[1]

[1]Carnegie Mellon University
[2]University of California, Berkeley
[3]Princeton University
[4]University of Sussex (U.K.)

EXECUTIVE SUMMARY:

Numerical simulations are now the dominant tool in theoretical cosmology. Our team has led the development of codes adapted to petascale supercomputers to study how supermassive black holes and galaxies formed, from the smallest to the rarest and most luminous. Our BlueTides simulation of nearly 1 trillion particles is the largest-ever simulation with full physics (hydrodynamics, star formation, black holes) and targeted the early universe of galaxies and quasars. The simulation agreed well with observations from the Hubble Space Telescope when the universe was only 5% of its current age. It covers a sky area 300 times larger than the largest survey with Hubble, allowing predictions of what Hubble's successors will see. Among the millions of galaxies that formed in the simulation were large disk galaxies as massive as the Milky Way and the first quasars; these objects will be exciting targets for the frontier of observations.

INTRODUCTION

The ability to model the complex physics of structure formation in the universe is increasing, driven by advances in both computer power and software. To continue to support observational and theoretical discovery, we need highly detailed simulations that follow the evolution of the matter, energy, and radiation in model universes from the Big Bang to the present.

In the coming decade, a new generation of astronomical instruments, all in the billion-dollar class, will observe the universe during the period of the first stars and quasars. Petascale computing resources running the newest codes are the theoretical equivalent to these programs and maximize the return on comparable observational and experimental studies. Using our simulation, we hope to make predictions for what should be seen in future observations and surveys as well as to imitate future surveys.

METHODS & RESULTS

The actions of gravity, hydrodynamics, forming stars, black holes, molecular gas, inhomogeneous ionizing radiation, and more were included in the BlueTides cosmological simulation [1]. We carried out a full-machine run on Blue Waters with our newly developed MP-Gadget. The simulation aims to understand the formation of the first quasars and galaxies from the smallest to the rarest and most luminous, and the role of these processes in the reionization of the universe.

Computation

BlueTides followed the evolution of 0.7 trillion particles in a large volume of the universe (600 co-moving Mpc on a side) over the first billion years of the universe's evolution with a dynamic range of 6 orders of magnitude in space and 12 in mass. This makes BlueTides by far the largest cosmological hydrodynamic simulation ever run.

We improved the cosmological code Gadget, now MP-Gadget, to make it suitable for the petascale and beyond:

- Threading efficiency: We replaced global critical sections with per-particle (per-node) spin locks. The code is now scalable to any number of threads and allows us to use fewer domains, which reduces the complexity of domain decomposition and inter-domain communication.
- Mesh gravity solver [2]: Domain decomposition was an issue with the huge Fourier transforms (up to $16{,}384^3$) needed to carry out large-scale mesh computation of the gravitational force. We improved the speed by a factor of 10 by switching from a slab decomposition to a "pencil" decomposition, which allows work to be distributed among 20,000+ XE6 nodes on Blue Waters.
- MP-sort [3]: We implemented and published with Blue Waters staff a new sorting algorithm that exchanges each data item exactly once (O(N) communication). It is only limited by the network bandwidth [4].

We also improved the physical modelling in MP-Gadget:

- A new pressure–entropy SPH (Smoothed Particle Hydrodynamics) formulation [9] replaced the old density–entropy formulation, which suppressed phases mixing in a non-physical way.
- A patchy reionization model from [10] introduced a UV field based on a predicted time of reionization at different spatial locations in the simulation.
- Improved stellar feedback models to include mass/halo dependent supernova wind efficiency (also metal cooling and H2 mass fraction) [11].

Science Results

The core scientific data set of BlueTides summed to 2 PB, consisting of 20 snapshots (47 TB each) of the properties of mass elements used in the simulation. We performed on-the-fly and offline data analysis to extract scientific results from this data set. Recent work included predictions about the first galaxies, first quasars, and their contribution to reionization.

First galaxies: We applied observational selection algorithms (SourceExtractor) to the simulated sky maps and created catalogs of millions of galaxies. This uncovered a striking and unexpected population of large Milky-Way-sized disk galaxies present when the universe was 5% of its present age [5].

First quasars: We computed statistical measures of the population of galaxies and quasars at different redshifts, which were consistent with early data from the Hubble Space Telescope and made predictions for upcoming surveys. We also can make predictions for the massive black holes that assemble at these early epochs and have masses 100 million times that of the Sun [7].

Reionization from galaxies and quasars: We found that the luminosity functions of the first galaxies and quasars indicate that a high escape fraction is required if galaxies dominate the ionizing photon budget during reionization. With smaller galaxy escape fractions, a large contribution from the faint quasars is feasible [8].

WHY BLUE WATERS?

Previous calculations have either used a small simulation volume or large particle masses, both of which resolve only large galaxies. BlueTides used more than 20,000 nodes and about 90% of the memory on Blue Waters and included both a large simulation volume and small particle masses to capture the rarest and brightest objects, the first quasars.

Blue Waters staff was critical to our success. They helped us run our full-system job in a timely fashion, helped with MPI+OpenMP development of the Gadget code, and assisted with fine-tuning file handling using HDF and Lustre. They also aided development of some radical new ways to process petabytes of simulation output.

Our simulations blaze a trail for future calculations. The next Track-1 system will make detailed 14-million-year simulations (early universe to today) possible. This would facilitate detailed comparisons between our simulated galaxy and quasar population and predictions that are relevant to the large-scale upcoming observational projects.

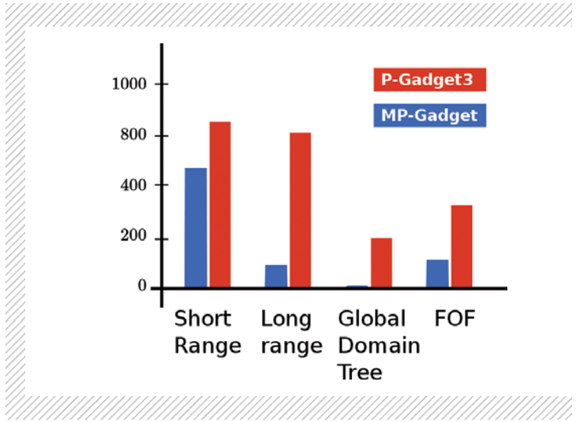

FIGURE 1: Comparing the wall time in a single step in MP-Gadget and P-Gadget in the BlueTides simulation. Even with more physics modules, the short-range force in MP-Gadget still spends less time per step than P-Gadget3. The long range force module in MP-Gadget out-performs P-Gadgets by a factor of 8. The global domain builder a factor of 10 and the galaxy catalogue (FOF) builder a factor of 2.

PUBLICATIONS

Feng, Y., et al., The formation of Milky Way-mass disk galaxies in the first 500 million years of a cold dark matter universe. *Astrophys. J. Lett.*, 808:1 (2015), L17, doi:10.1088/2041-8205/808/1/L17.

Croft, R., T. Di Matteo, N. Khandai, and Y. Feng, Petascale cosmology: simulation of structure formation. *Comput. Sci. Eng.*, 17:2 (2015), pp. 40–46, doi:10.1109/MCSE.2015.5.

Feng, Y., M. Straka, R. Croft, and T. Di Matteo, Sorting At Scale on Blue Waters. *Cray User Group Conf. 2015*, Chicago, Ill., April 26–30, 2015. [Best Paper finalist]

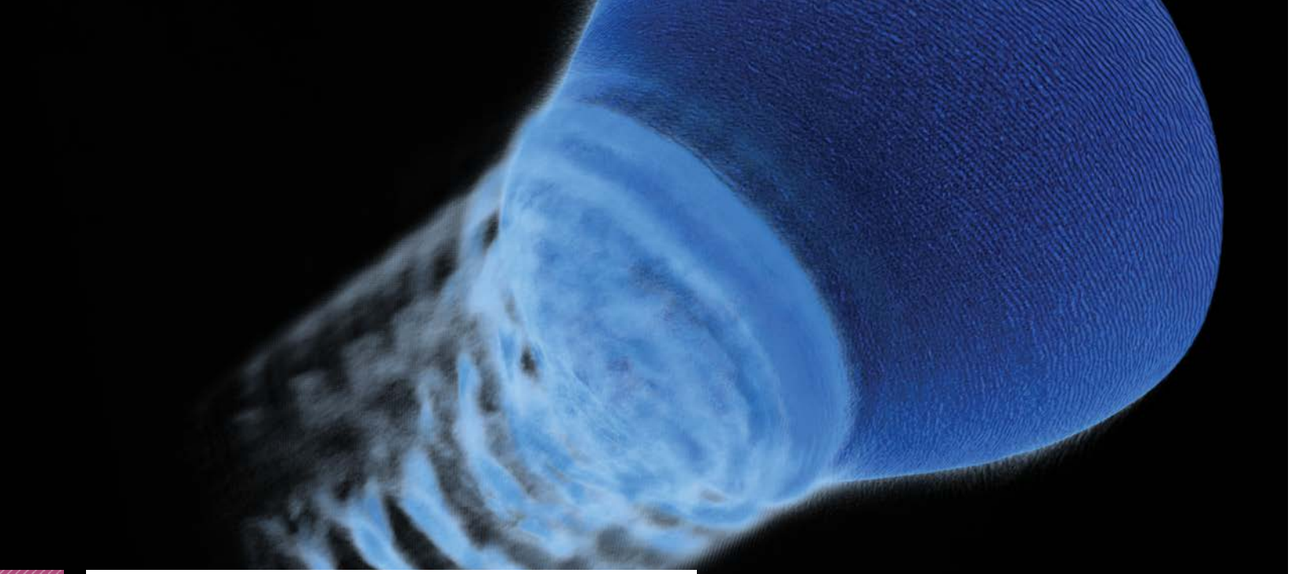

ENABLING BREAKTHROUGH KINETIC SIMULATIONS OF THE MAGNETOSPHERE

Allocation: NSF PRAC/11.1 Mnh
PI: Homayoun Karimabadi[1,2]
Collaborators: Vadim Roytershteyn[2,3], William Daughton[4], Kai Germaschewski[5], William Matthaeus[6], Minping Wan[6], Mahidhar Tatineni[1], Amit Majumdar[1]

[1]University of California, San Diego
[2]SciberQuest
[3]Space Science Institute
[4]Los Alamos National Laboratory
[5]University of New Hampshire
[6]University of Delaware

EXECUTIVE SUMMARY:

We present recent results from large-scale kinetic plasma simulations aimed at advancing the basic understanding of phenomena such as magnetic reconnection and plasma turbulence, and the degree to which they affect global dynamics of the Earth's magnetosphere.

INTRODUCTION

The ability to understand and forecast the behavior of the Earth's magnetosphere is becoming increasingly important as our society relies more and more on space-based technologies such as telecommunications, Global Positioning System, and others. The Earth's magnetosphere is a protective "bubble" formed when our planet's internal magnetic field interacts with the solar wind, a stream of plasma originating on the Sun. The magnetosphere responds dynamically to variations in the relatively steady solar wind and to more transient, sometimes explosive events on the Sun, such as coronal mass ejections. Many of the processes determining the behavior of the Earth's magnetosphere are very complex, involving plasma physics and often coupling global, large-scale dynamics to microscopic processes occurring on short spatial and temporal scales associated with plasma kinetic effects.

The goal of this project is to investigate this cross-scale coupling for two of the most ubiquitous plasma processes that are essential to the dynamics of the magnetosphere and the solar wind: magnetic reconnection and plasma turbulence. Magnetic reconnection is a process of rapid merging and breaking of the magnetic field embedded into plasma. Often explosive in nature, magnetic reconnection enables rapid release of magnetic energy accumulated in the system and is ultimately behind phenomena like magnetospheric sub-storms and solar eruptions. But reconnection can also operate in a more steady fashion in the solar wind and in the magnetosphere (for example, enabling solar wind plasma entry into the magnetosphere when conditions are favorable).

Similarly to fluid turbulence, turbulence in plasmas is a ubiquitous process that is thought to provide a crucial mechanism for energy, momentum, and mass transport in many systems across the universe. For example, the solar wind is known to be highly turbulent and many theories invoke energy input due to dissipation of turbulence to explain anomalous temperature profiles in the solar wind. Due to availability of in situ measurements provided by a fleet of spacecraft, the solar wind is an excellent laboratory for studying the basic physics of plasma turbulence.

While magnetic reconnection and plasma turbulence are two distinct physical processes, in large-scale systems they often interact and must be studied simultaneously. Moreover, in high-temperature rarified plasmas typical of those encountered in space, both turbulence and magnetic reconnection involve kinetic processes that depend on the details of interaction between the Earth's electromagnetic field and individual plasma particles (electrons and ions). Extreme separation of scales between typical spatial and temporal scales associated with kinetic physics and those associated with global dynamics is a formidable computational challenge and requires the largest available computational resources.

METHODS & RESULTS

The project utilized particle-in-cell (PIC) kinetic plasma simulations using two complementary approaches. Fully kinetic simulations that describe all plasma species using a Vlasov–Maxwell (or Vlasov–Boltzmann) system provided essentially a first-principles description of the plasmas of interest. Such simulations were conducted primarily using VPIC, a high-performance relativistic plasma simulations code [1], and focused on local modeling of magnetic reconnection and turbulence in various environments [2–8]. In the hybrid model, relatively heavy ions were still described kinetically while electrons were modeled as a massless fluid, allowing much larger systems to be considered in the situations where the dominant kinetic effects are associated with ions. The hybrid simulations for this project used simulation code H3D [9] and focused primarily on modeling global interaction of the solar wind with a magnetosphere, as illustrated in fig. 1. Below, we briefly discuss two of the topics considered in this project.

One of the most interesting results that emerged from recent global hybrid simulations was the coupling between bow shock, turbulence, and magnetic reconnection. While many individual pieces of physics had long been studied using spacecraft observations, analytical theory, and small-scale simulations, the degree to which such coupling affects global dynamics had only become apparent recently [10]. The majority of the existing results were obtained using 2D simulations that severely restricted the geometry of the plasma flow around magnetosphere. Blue Waters enabled the first 3D simulations, which provided a multitude of exciting new details (e.g., fig. 1).

The studies of plasma turbulence conducted on Blue Waters primarily focused on understanding mechanisms responsible for dissipation of turbulence in high-temperature rarified plasmas where Coulomb collisions are too infrequent to provide the required microscopic dissipation (see example in fig. 2). Under such conditions, the so-called collective modes provide the dissipation. The three most commonly discussed alternatives are damping associated with wave-like oscillations, dissipation associated with coherent structures sheets (e.g., current sheets, which are localized narrow regions where the magnetic field is highly sheared), and stochastic damping originating from the motion of individual ions in a turbulent electromagnetic field. In many plasmas of interest, these three basic processes are likely to operate simultaneously, but their relative efficiency may depend on the parameter regime. Fully kinetic simulations provide a unique model that is capable of seamlessly describing all of the relevant processes, which is essential to understanding how they interact with each other and which may be dominant under given conditions. This in turn will allow researchers to discriminate between various theoretically conceived scenarios regarding the role of turbulence in such systems as the solar corona and the solar wind.

WHY BLUE WATERS?

Due to the extreme separation of scales, the problems that need to be solved require the largest available computational resources.

FIGURE 1 (BACKGROUND): Structure of the ion foreshock in 3D global hybrid simulation of the Earth's magnetosphere. Perturbations generated in the foreshock by ions reflected from the bow shock were convected back into the shock, got amplified, and led to both small-scale turbulence and global perturbations in the region behind the shock (called the magnetosheath).

In additional to providing the necessary computational time, Blue Waters is notable for its advanced infrastructure for data management and analysis.

PUBLICATIONS

Guo, F., Y.-H. Liu, W. Daughton, and H. Li, Particle acceleration and plasma dynamics during magnetic reconnection in the magnetically-dominated regime. *Astrophys. J.*, 806 (2015), 167, doi:10.1088/0004-637X/806/2/167.

Makwana, K. D., V. Zhdankin, H. Li, W. Daughton, and F. Cattaneo, Energy dynamics and current sheet structure in fluid and kinetic simulations of decaying magnetohydrodynamic turbulence. *Phys. Plasmas,* 22 (2015), 042902, doi:10.1063/1.4916492.

Roytershteyn, V., H. Karimabadi, and A. Roberts, Generation of magnetic holes in fully kinetic simulations of collisionless turbulence. *Philos. Trans. A,* 373:2041 (2015), doi:10.1098/rsta.2014.0151.

Wan, M., et al., Intermittent Dissipation and Heating in 3D Kinetic Plasma Turbulence. *Phys. Rev. Lett.,* 114 (2015), 175002, doi:10.1103/PhysRevLett.114.175002.

Guo, F., H. Li, W. Daughton, and Y.-H. Liu, Formation of Hard Power Laws in the Energetic Particle Spectra Resulting from Relativistic Magnetic Reconnection. *Phys. Rev. Lett.,* 113 (2014), 155005, doi:10.1103/PhysRevLett.113.155005.

Jara-Almonte, J., W. Daughton, and H. Ji, Debye Scale Turbulence within the Electron Diffusion Layer during Magnetic Reconnection. *Phys. Plasmas,* 21 (2014), 032114, doi:10.1063/1.4867868.

Karimabadi, H., et al., The link between shocks, turbulence, and magnetic reconnection in collisionless plasmas. *Phys. Plasmas,* 21 (2014), 062308, doi:10.1063/1.4882875.

Liu, Y.-H., J. Birn, W. Daughton, M. Hesse, and K. Schindler, Onset of reconnection in the near magnetotail: PIC simulations. *J. Geophys. Res.,* 119 (2014), pp. 9773–9789, doi:10.1002/2014JA020492.

FIGURE 2: Current density in a 3D fully kinetic simulation of decaying plasma turbulence. This snapshot illustrates formation of several types of coherent structures, including planar current sheets, electron-scale magnetic holes, and flux ropes.

CORE-COLLAPSE SUPERNOVAE THROUGH COSMIC TIME

Allocation: NSF PRAC/6 Mnh
PI: Eric J. Lentz[1,2]
Co-PIs: Stephen W. Bruenn[3], W. Raphael Hix[1,2], O. E. Bronson Messer[2], Anthony Mezzacappa[1,2]

[1]University of Tennessee, Knoxville
[2]Oak Ridge National Laboratory
[3]Florida Atlantic University

EXECUTIVE SUMMARY:

We study the explosive deaths of massive stars, supernovae, and their contribution to the evolution of the elemental content of the universe. Core-collapse supernovae (CCSNe) are tightly coupled multi-physics events without natural symmetry and require physically and spatially detailed 3D simulations to resolve. The progenitor stars of CCSNe vary in mass, heavy element composition, rotation, and other parameters that affect how the explosion develops, or if it develops at all, and the elemental abundances of the ejecta. Over the course of our project we will compute CCSN models that broadly cover the range of masses and compositions representative of massive stars throughout the history of the universe. In prelude to that survey, we are conducting studies with Blue Waters that will assess the impact of spatial resolution on our simulations.

INTRODUCTION

The inner iron core of a massive star (with mass more than about eight solar masses) collapses in less than 1 second to a (hot or proto) neutron star about 100 km across, releasing about 10^{53} ergs (10^{46} J = 100 B) of energy primarily in the form of neutrinos. When the collapse is halted, the outer core rebounds for a shock that stalls above the newly formed proto-neutron star (PNS). Neutrino heating of the in-falling material located just outside the PNS creates a layer of buoyantly unstable fluid that rises convectively and after several hundred milliseconds drives an outward shock wave. Once launched, about a full simulated second is required to build the typical 1 B explosion energy [1] and most of the element production [2].

CCSNe are important sources of many of the elements and are important contributors to the evolution of chemical abundances in galaxies and new stars. Variations of the initial mass and composition (built up through previous generations of stars, including supernovae) not only impact the strength of explosions, but the relative and absolute abundances of heavy elements in the ejecta. No 3D simulation with an *ab initio* explosion (as opposed to a parameterized one) has been evolved long enough to compute the required element synthesis.

METHODS & RESULTS

To capture the required physics in CCSNe we use a code we constructed call Chimera. It includes fluid dynamics, multi-pole self-gravitation with spherical general relativity corrections, radial neutrino transport with a full set of neutrino–matter interactions, a dense nuclear equation of state, and a nuclear network to track nuclear burning in the outer layers and ejecta.

To date, our group has worked mostly with 2D simulations. In a recent study covering a range of masses [1], we obtained appropriately energetic explosions from all four tested progenitors (12–25 solar masses). These simulations also clearly demonstrated that much longer simulations were required to obtain basic supernova properties like explosion energy and ejecta abundances. While the shock is revived by neutrino heating and fluid instabilities about 250 ms after bounce, development of the full explosion energy can take a full second more [1], which has a large impact on the required length of the runs. The conclusion of nuclear burning and the sorting of ejecta from non-ejecta take similarly long periods [2]. Our primary goal with Blue Waters is to move from simple 2D explosions to full 3D that also cover the range of stellar initial abundances.

Our first full-physics and well-resolved 3D model (180 x 180 zones in angle; 540 in radius) showed a delay of shock revival (explosion) of about 100 ms relative to its 2D counterpart [3] (or rather we should say the 2D model revived 100 ms earlier as the imposed symmetry of 2D is artificial). Fig. 1 depicts the heated and convective fluid behind the shock at 250 ms after core bounce with numerous rising convective plumes (yellow-green and red) pressing against the shock upon which the inner part of the star continued to accrete. The accreting matter passed in streams between the plumes and released gravitational binding energy that contributed to the heating. The 2D model had only a couple of large plumes with a single stream and began to explode at this epoch. During the delay in the 3D model, the plumes grew in angular size and this seemed to be important for the initiation of explosion [3]. These effects will be explored more thoroughly when we have completed more simulations.

The effects of resolution on 3D CCSN simulations are largely unknown and are our current focus. Insufficient resolution of the plumes and accretion streams could lead to excess dissipation of the accretion before it can contribute to the heating, truncate the turbulent cascade, and numerically stiffen plumes so they grow to large scales more quickly. To explore these issues we are running reduced-resolution models on Blue Waters and are engaged in other tests of 3D resolution. This will help maximize the usefulness of our large survey of simulations.

WHY BLUE WATERS?

Well-resolved 3D simulations of CCSNe require more than a thousand nodes for up to 1,000 hours of wall-clock time to fully develop the explosion (32,000 core hours). Blue Waters is in a rare class of machines that can sustain the required throughput of simulations at this size to fully evolve even a few 3D supernova models in a year and provide the needed infrastructure for data analysis and handling. Future Track-1 systems will continue to accelerate discovery in this area

FIGURE 1: Volume rendering of entropy in heated region of 15 solar-mass star 250 ms after core bounce, about 100 ms before the explosion initiates [3]. Blue corresponds to the shock position and yellow-green and red to progressively more heated material in rising convective plumes. [Credit: Mike Matheson, Oak Ridge Leadership Computing Facility. Animated version available at ChimeraSN.org]

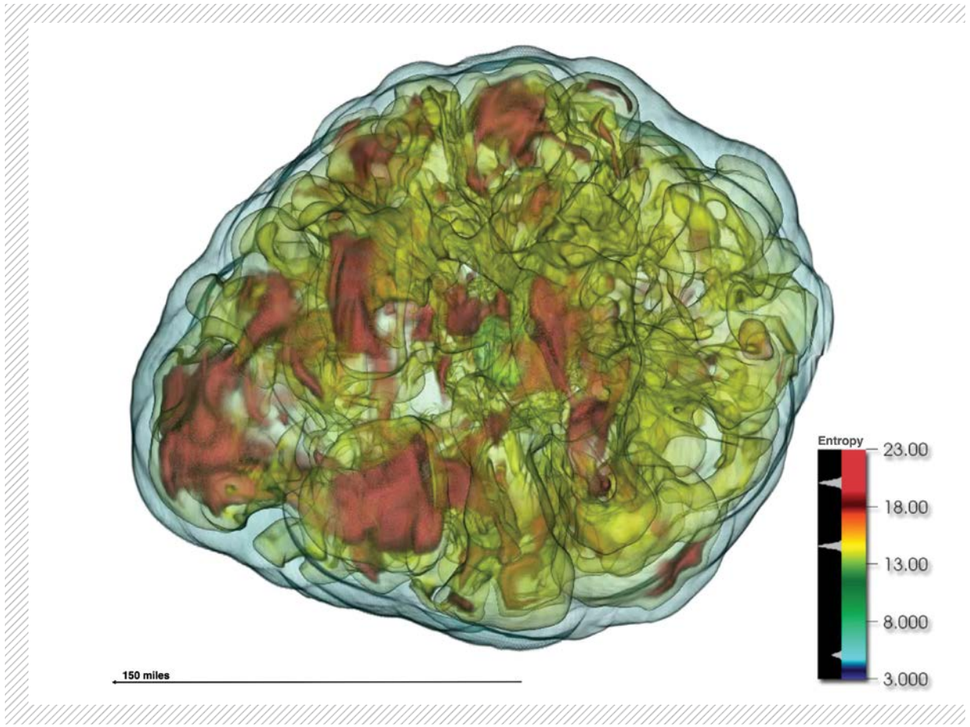

SIMULATING MAGNETIZED PLASMA TURBULENCE FROM MACRO TO MICRO SCALES

Allocation: GLCPC/0.56 Mnh
PI: Kirit Makwana[1]
Co-PIs: Fausto Cattaneo[1], William Daughton[2], Hui Li[2]
Collaborator: Vladimir Zhdankin[3]

[1]University of Chicago
[2]Los Alamos National Laboratory
[3]University of Wisconsin-Madison

EXECUTIVE SUMMARY:

Magnetized plasmas are ubiquitous in space, astrophysical, and laboratory environments. Energy is injected in 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 ultimately is converted into heat. We simulate this process using both magnetohydrodynamic (MHD) and kinetic (particle-in-cell) codes. 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. A scan over the plasma beta parameter is performed to reveal significant non-thermal particle energization at low beta. This is linked with reconnecting current sheets in the simulation.

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, in planetary magnetospheres, in 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 used the magnetohydrodynamics (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. For understanding of this process, kinetic simulations are required. In this project, we simulated the turbulent cascade of energy using both MHD and kinetic codes. This tells us whether MHD is the correct limit of plasma turbulence at macro scales. It also tells us 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 advanced our understanding of dissipation in plasma turbulence.

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 for both codes was an ensemble of superimposed shear Alfven waves. 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 the 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 showed that MHD simulations of turbulence produce reliable energy dynamics.

At smaller scales, important differences appear between the two descriptions. We observed formation of current sheets in both the codes, as shown in Figure 1. 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 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). We are currently investigating the link of magnetic reconnection in current sheets with this non-thermal particle energization.

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 high-performance computing, and Blue Waters specifically, are now capable of reproducing MHD results from first principle, particle-in-cell simulations. This opens up the small scale physics of energy dissipation to 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 planetary magnetospheres. This 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. The VPIC code we used for the kinetic simulations is part of the NCSA Blue Waters Sustained Petascale Performance (SPP) suite. NCSA and Cray improved compiler optimization of loops not already using optimized vector compiler intrinsic functions, optimized to eliminate extra data copies, added FMA4 compiler intrinsic functions to improve performance, and used Cray I/O buffering functionality. Our simulations used a 1024x1024x1024 cell domain with 3.758 billion particles.

The Blue Waters project staff was also helpful in the data analysis and visualization. They helped us in using Paraview to visualize the data in 3D. Figure 1 shows an example of 3D visualization of the current sheets that form in our simulation.

The next Track-1 system can help us in simulating plasmas with even lower beta, thereby giving significant non-thermal particle energization. This is more attractive for explaining the abundance of non-thermal particle distribution functions observed in astrophysical plasmas.

PUBLICATIONS

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.

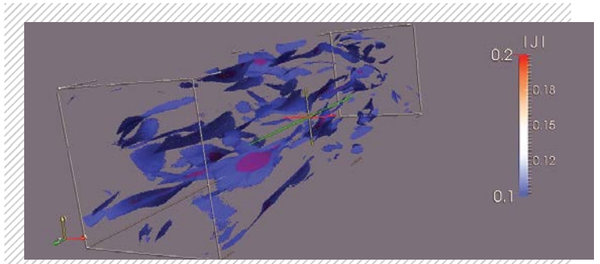

FIGURE 1 : The large-scale Surfaces of constant current density in the rectangular simulation domain. The simulation domain is elongated along the background magnetic field. The two colors are for two current density values. We can clearly see the sheet like structure of current density isosurfaces. The sheets have three typical dimensions. They are longest in the direction parallel to background magnetic field. Perpendicular to the field they have an intermediate width, and very small thickness.

EXPLORING THE FIRST GENERATIONS OF GALAXIES

Allocation: NSF PRAC/7.8 Mnh
PI: Brian W. O'Shea[1]
Co-PI: Michael L. Norman[2,3]
Collaborators: James Bordner[3], Pengfei Chen[2], Dan Reynolds[4], Devin Silvia[1], Britton Smith[5], John Wise[6], Hao Xu[2]

[1]Michigan State University
[2]University of California at San Diego
[3]San Diego Supercomputing Center
[4]Southern Methodist University
[5]University of Edinburgh
[6]Georgia Institute of Technology

EXECUTIVE SUMMARY:

We are investigating the earliest stages of cosmological structure formation—namely, the transition of the universe from a dark, empty place to one filled with stars, galaxies, and the cosmic web. In investigating the "cosmic dark ages," we focused on three specific topics: the transition between metal-free and metal-enriched star formation that marks a fundamental milestone in the early epochs of galaxy formation; the evolution of the populations of high-redshift galaxies that will form Milky-Way-like objects by the present day; and the reionization of the universe by these populations of galaxies. All of these problems required simulations with extremely high dynamic range in space and time, complex physics (including radiation transport and non-equilibrium gas chemistry), and extremely large simulation volumes. We used the Enzo code, which has been modified to scale to large core counts on Blue Waters, the only machine available that satisfies the heavy data and communication needs of this best of breed study.

INTRODUCTION

The mechanisms that control the formation and evolution of galaxies are poorly understood. This is doubly true for the earliest and most distant galaxies, where observations are limited and often indirect. It is critical to understand the properties of the first generations of galaxies since they were responsible for reionizing the universe, for dispensing large quantities of metal into the low-density intergalactic medium (and thus fundamentally altering the behavior of plasmas in the future cosmological structure), and as the sites of formation and early growth for the supermassive black holes that are found at the center of every massive galaxy at the present day. Understanding the first generations of galaxy formation—the goal of this research project—is critical to the interpretation of current and future surveys of both the most distant galaxies that can be observed and also of our own Milky Way. More generally, this work helps to answer the philosophical questions "where do we come from?" and "how did we get here?"

METHODS & RESULTS

Our simulation tool of choice is the Enzo code [1,2], an open-source, community-developed software platform for studying cosmological structure formation. Enzo allowed us to include all of the critical physical components needed to study galaxy formation—gravity, dark matter dynamics, fluid dynamics, the microphysics of plasmas, and prescriptions for star formation and feedback—and to do so using a tool that can scale to large numbers of CPUs. All analysis was done with the YT code [3,4].

Using Blue Waters, we successfully modeled the formation of the first generation of metal-enriched stars in the universe and showed that there are several possible formation modes for these stars. These included self-enrichment of the halo where the first primordial star formed and pollution of neighboring halos by the supernova remnant. In addition, we showed that the presence of dust (which may form in the ejecta of the first supernovae) can have a critical effect on metal-enriched star formation, directly resulting in additional cooling and the formation of additional molecular hydrogen. Extra hydrogen further increases cooling rates and may cause additional fragmentation and lower-mass stars.

We have also modeled the evolution of large populations of galaxies in the early universe in several different large-scale environments. We find that if these Population III stars form massive black hole/stellar binary systems, they are likely to be prodigious emitters of X-ray radiation. This radiation both heats and ionizes the intergalactic medium, in some cases to 10^4 Kelvin! This may be important for predicting the topology of the 21 cm neutral hydrogen signal, which will be detectable by low-wavelength radio arrays in the coming years. We also show that radiation from early star formation can help suppress the collapse of gas in neighboring halos, delaying star formation and causing the galaxy luminosity function to be strongly suppressed at lower luminosities.

WHY BLUE WATERS?

To properly model the earliest galaxies, our simulations required complex physics—most importantly, radiation transport and non-equilibrium gas chemistry—and extremely high spatial and temporal dynamic range. Furthermore, large simulation volumes (and thus many resolution elements) were needed in order to model enough galaxies to be able to adequately compare theory to observations in a statistically meaningful way. Taken together, this required 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 machine in the U.S. available to the academic community that fits all of these requirements.

Given the expected increase in machine size for the next Track-1 system, we will be able to expand our research in several dimensions: we will be able to increase our simulation spatial resolution by at least a factor of two and mass resolution by a factor of ten, which will enable us to explore even fainter galaxies than we are currently examining; we will be able to extend our simulations to later times, thus making a wider variety of useful predictions (i.e. we will be able to predict galaxy properties for both ground- and space-based telescopes, as opposed to solely space-based telescopes); and we will be able to include additional physics, such as active galactic nucleus feedback, that is not currently in our models.

PUBLICATIONS

Ahn, K., et al., Spatially Extended 21 cm Signal from Strongly Clustered UV and X-Ray Sources in the Early Universe. *Astrophys. J.*, 802 (2015), 8, doi:10.1088/0004-637X/802/1/8.

O'Shea, B. W., J. H. Wise, H. Xu, and M. L. Norman, Probing the Ultraviolet Luminosity Function of the Earliest Galaxies With the Renaissance Simulations. *Astrophys. J.*, 807:1 (2015), L12, doi:10.1088/2041-8205/807/1/L12.

Chen, P., J. H. Wise, M. L. Norman, H. Xu, and B. W. O'Shea, Scaling Relations for Galaxies Prior to Reionization. *Astrophys. J.*, 795 (2014), 144, doi:10.1088/0004-637X/795/2/144.

Xu, H., K. Ahn, J. Wise, M. Norman, and B. W. O'Shea, Heating the Intergalactic Medium by X-Rays from Population III Binaries in High-redshift Galaxies. *Astrophys. J.*, 791 (2014), 110.

Xu, H., J. Wise, and M. L. Norman, Population III Stars and Remnants in High-redshift Galaxies. *Astrophys. J.*, 773 (2013), 83.

FIGURE 1 (BACKGROUND): The gas ejected from the supernova created by the first star to form in a cosmological volume as it is leaving the dark matter halo where the star formed. The supernova remnant preferentially expands into the void regions near the halo rather than the filaments, since this is the path of least resistance, and fills the intergalactic medium with metal-enriched gas. Simulations first published in Smith et al. (Astrophys. J., submitted).

MODELING HELIOPHYSICS PHENOMENA WITH A MULTI-SCALE FLUID-KINETIC SIMULATION SUITE

Allocation: NSF PRAC/0.94 Mnh
PI: Nikolai Pogorelov[1]
Collaborators: Matthew Bedford[1], Raymond Fermo[1], Tae Kim[1]

[1]University of Alabama in Huntsville

EXECUTIVE SUMMARY:

We investigated physical phenomena occurring when the solar wind (SW) interacts with the local interstellar medium (LISM). These problems include (1) issues related to the mixing of the SW and LISM plasma at the heliospheric interface defined by the heliopause (the boundary of the heliosphere, the spherical region around the Sun that is filled with solar magnetic fields and the outward-moving solar wind consisting of protons and electrons) in particular due to the heliopause instability and magnetic reconnection at its surface; (2) the influence of the heliosphere on the observed anisotropy of teraelectronvolt (TeV) galactic cosmic rays; (3) the dynamic effect of non-thermal ions, and (4) the influence of time-dependent phenomena on the energetic neutral atom flux observed by the Interstellar Boundary Explorer (IBEX) space mission. Our simulations are critical for the explanation of Voyager and IBEX measurements pertinent to the SW–LISM interaction, and Tibet, Milagro, Super-Kamiokande, IceCube, and ARGO-YGB observatories' measurements of the cosmic ray anisotropy.

INTRODUCTION

Voyager 1 (V1) and Voyager 2 (V2) spacecraft crossed the heliospheric termination shock in December 2004 and August 2007, respectively [6,7]. After 37 years of historic discoveries, Voyager 1 started sampling the local interstellar medium (LISM) while Voyager 2 is approaching the heliopause (HP), a tangential discontinuity separating the solar wind (SW) from the LISM. V1 and V2 acquire *in situ* information about the local properties of the SW plasma, energetic particles, and magnetic field at the heliospheric boundary [8]. In addition to the thermal component, the SW plasma has a non-thermal component represented by pickup ions (PUIs). These are born when thermal SW ions exchange charge with the LISM neutral atoms. On the other hand, IBEX is measuring line-of-sight integrated fluxes of energetic neutral atoms (ENAs) in different energy bands [4]. Since most ENAs are created during charge exchange between hot PUIs and LISM neutral atoms, they bear the plasma properties of the region in which they are created.

The combination of IBEX and Voyager observations gives us a unique opportunity to investigate physical phenomena at the heliospheric interface and the transport of galactic cosmic rays into the heliosphere. Additionally, by fitting a narrow "ribbon" of an enhanced ENA flux discovered by IBEX and a TeV cosmic ray anisotropy measured at Tibet, Milagro, Super-Kamiokande, ICECube, and ARGO-YGB observatories (see [10] for an extensive list of references) we can constrain the properties of the LISM at large distances from the heliopause.

METHODS & RESULTS

To address the problems described above, we solved the equations of ideal magnetohydrodynamics (MHD) equations coupled with the kinetic Boltzmann equation describing the transport of neutral atoms. In a less strict, but very efficient 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 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 built on the Chombo reference framework.

During the second year of our Blue Waters allocation, we (1) performed high-resolution simulations of the heliopause instability [1,5] and identified areas of possible magnetic reconnection near the heliopause crossed by Voyager trajectories (fig. 1), (2) analyzed the heliotail flow for different LISM conditions and proved [10] that the observed TeV cosmic ray anisotropy may be explained by the LISM magnetic field distortion by the heliosphere, and (3) investigated the effect of non-thermal ions on time-dependent plasma distributions [3]. The results are published in six papers and reported at 14 (11 invited) scientific meetings. By addressing the basic physical phenomena occurring at the interface of the heliosphere and LISM, our results are of substantial importance for 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?

Our simulations are computationally intensive. This is especially true in the following typical situations: (1) Neutral atoms are modeled kinetically and we need of the order of 10^{12} particles in our Monte Carlo simulations. These simulations require particle splitting, multiple grids, and careful coupling with the MHD module, and cannot be performed on other smaller resources. They also produce multiple data sets exceeding 0.5 TB each. Hybrid parallelization becomes obligatory. (2) Computational region sizes are very large, as in the case of long-tail simulations. (3) Very deep adaptive mesh refinement (AMR) is necessary near magnetic reconnection sites. This will be our grand challenge in the final year of our allocation and help explain Voyager data. Additionally, we will perform parametric simulations of the heliospheric tail to fit IBEX and TeV cosmic ray anisotropy observations.

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.

Borovikov, S. N., and N. V. Pogorelov, Voyager 1 near the Heliopause. *Astrophys. J.,* 783 (2014), L16, doi:10.1088/2041-8205/783/1/L16.

Heerikhuisen, J., E. J. Zirnstein, H. O. Funsten, N. V. Pogorelov, and G. P. Zank, The effect of new interstellar medium parameters on the heliosphere and energetic neutral atoms from the interstellar boundary. *Astrophys. J.,* 784 (2014), 73.

Kim, T. K., et al., MHD heliosphere with boundary conditions from a tomographic reconstruction using interplanetary scintillation data. *J. Geophys. Res.,* 119:10 (2014), 7981.

Zhang, M., P. Zuo, and N. Pogorelov, Heliospheric influence on the anisotropy of TeV galactic cosmic rays. Astrophys. J., 790 (2014), 5.

Pogorelov, N., S. N. Borovikov, J. Heerikhuisen, and T. K. Kim, Time-dependent processes in the sheath between the heliospheric termination shock and the heliopause. in *Numerical Modeling of Space Plasma Flows: ASTRONUM-2013,* 488 (2013),pp. 167–177. [eds. N.V. Pogorelov, E. Audit, G.P. Zank]

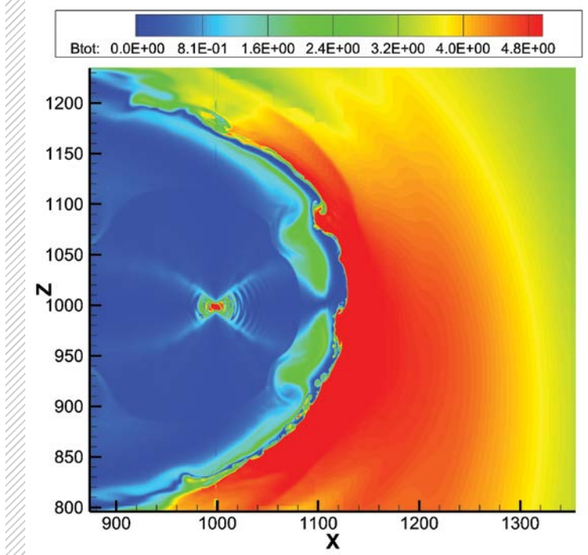

FIGURE 1: The magnetic field distribution at the interface between the SW and LISM exhibits both Rayleigh Taylor instability of the heliopause due to charge exchange between ions and neutral atoms and signatures of tearing mode instability typical of magnetic reconnection. The simulation was performed with the local AMR ratio of 32.

RELATIVISTIC JET FORMATION IN BLACK HOLE–NEUTRON STAR MERGERS

Allocation: NSF PRAC/2.1 Mnh
PI: Stuart L. Shapiro[1]
Co-PI: Vasileios Paschalidis[2]
Collaborators: Ryan Lang[1], Milton Ruiz[1]

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

EXECUTIVE SUMMARY:

Merging black hole–neutron star binaries will be prominent sources for advanced gravitational-wave detectors. They are also thought to produce relativistic jets and serve as engines that power short-hard gamma-ray bursts. Simultaneous observation of gravitational waves and gamma rays from these systems is the holy grail of "multi-messenger astronomy." Until now, no self-consistent calculation existed in full general relativity that starts from the compact binary inspiral and demonstrates that jets can be launched after the tidal disruption of the neutron star. We performed ideal magnetohydrodynamic simulations of black hole–neutron star systems in full general relativity and showed, for the first time, that they can indeed launch incipient jets if the neutron star is initially endowed with a dipolar magnetic field extending from its interior well into its exterior.

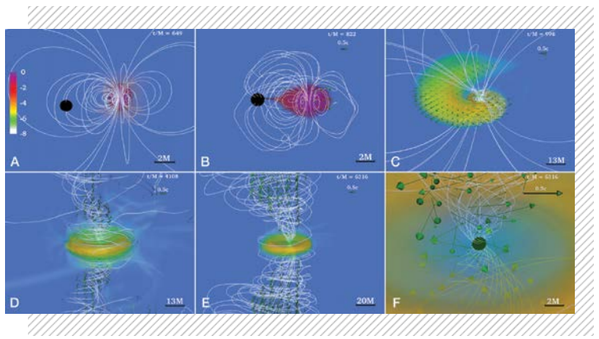

INTRODUCTION

On the centennial anniversary of Einstein's theory of general relativity, we are on the verge of directly detecting one of its most remarkable predictions: gravitational waves (GWs). The inspiral and merger of compact binaries—binaries with black hole (BH), neutron star (NS), or white dwarf companions—are among the most promising sources of GWs. Many of these sources likely also generate electromagnetic (EM) radiation counterparts to the GWs. Detecting both GW and EM radiation from the same cosmic source will be a major advance in "multi-messenger astronomy."

Gamma-ray bursts (GRBs) were first discovered in 1967, and theorists have been working to explain them since. The current best model for short bursts (sGRBs), those with duration less than two seconds, is the merger of a NS with a companion NS or BH. These systems are thus excellent candidates for multi-messenger detection. In order to verify the binary–sGRB 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 black hole–neutron star (BHNS) case.

METHODS & RESULTS

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

Using our latest adaptive-mesh refinement GRMHD code (see, e.g., [1]) we performed the first study of merging magnetized BHNS binaries [2,3]. In these simulations, the NS was seeded with strong magnetic fields confined to the NS interior. Many sGRB models require relativistic outflows and magnetic field collimation in the merger remnant, but these features were not observed in these initial simulations. Instead, they showed that interior-only initial magnetic field configurations inevitably led to a final magnetic field in the disk that was almost purely toroidal. Toroidal magnetic fields cannot launch a jet. Accretion of poloidal magnetic flux is necessary to launch and sustain jets from BH accretion disk systems [4].

Real neutron stars are expected to be endowed with dipole magnetic fields that extend from the interior well into the exterior, as required by current theories of pulsars. In addition to being more realistic, this initial magnetic field configuration is also more likely to provide the necessary conditions for launching a relativistic jet. When the magnetic field is allowed to extend into the NS exterior, poloidal magnetic field lines attached to fluid elements thread the BH prior to tidal disruption. Following disruption, while the magnetic field in the disk winds up in a predominantly toroidal pattern, a strong poloidal component is also amplified, threading the low-density debris.

To study this configuration, we performed simulations of a BHNS binary initially on a quasicircular orbit [5]. The binary mass ratio was 3 to 1, the BH had initial spin $a/M = 0.75$, and the NS was modeled as an irrotational, initially unmagnetized polytrope. We evolved the hydrodynamic and metric fields until two orbits formed prior to tidal disruption, at which point the NS was seeded with a dynamically weak dipole magnetic field that extended from the stellar interior well into the exterior (fig. 1, upper left). To ensure reliable evolution of the exterior field and to properly mimic the conditions that likely characterize the exterior magnetosphere, we also 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 tidal disruption of the NS and subsequent formation of a magnetized accretion disk (fig. 1, upper center and right).

We found that at ~100 ms following the onset of accretion, magnetic field winding above the remnant BH poles built up the magnetic field sufficiently to launch a mildly relativistic, collimated outflow—an incipient jet (fig. 1, bottom row). The duration of the accretion and the lifetime of the jet was ~0.5 seconds, which is consistent with the typical duration of sGRBs. The luminosity was 10^{51} erg/s, also consistent with observed sGRB values. Our simulations were the first self-consistent calculations in full general relativity that provided theoretical corroboration that mergers of BHNS systems can launch jets and be the central engines that power sGRBs.

WHY BLUE WATERS?

By adding OpenMP support to our MPI-based code, scalability on multi-core machines improved greatly. With the Blue Waters interconnect and processors, our hybrid OpenMP/MPI code exhibited 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 to our code's performance, making Blue Waters uniquely capable for tackling the astrophysical problems we want to address.

Our undergraduate research team also used Blue Waters to make visualizations (e.g., fig. 1) and movies of our simulations with the VisIT software. Recently, we created the first 3D movies using VisIT on Blue Waters, in collaboration with Professor Donna Cox and her group.

PUBLICATIONS

Paschalidis, V., M. Ruiz, and S. L. Shapiro, Relativistic simulations of black hole–neutron star coalescence: the jet emerges. *Astrophys. J.*, 806 (2015), L14, doi:10.1088/2041-8205/806/1/L14.

Etienne, Z. B., et al., Improved moving puncture gauge conditions for compact binary evolutions. *Phys. Rev. D*, 90 (2014), 064032, doi:10.1103/PhysRevD.90.064032.

Gold, R., et al., Accretion disks around binary black holes of unequal mass: GRMHD simulations of postdecoupling and merger. *Phys. Rev. D*, 90 (2014), 104030, doi:10.1103/PhysRevD.90.104030.

Gold, R., et al., Accretion disks around binary black holes of unequal mass: GRMHD simulations near decoupling. *Phys. Rev. D*, 89 (2014), 064060, doi:10.1103/PhysRevD.89.064060.

Ruiz, M., V. Paschalidis, and S. L. Shapiro, Pulsar spin-down luminosity: Simulations in general relativity. *Phys. Rev. D*, 89 (2014), 084045, doi:10.1103/PhysRevD.89.084045.

FIGURE 1 (LEFT): Snapshots of the rest-mass density, normalized to its initial maximum value (log scale), at selected times. Arrows indicate plasma velocities and white lines show the magnetic field lines. Panels d-f highlight the system after an incipient jet launched.

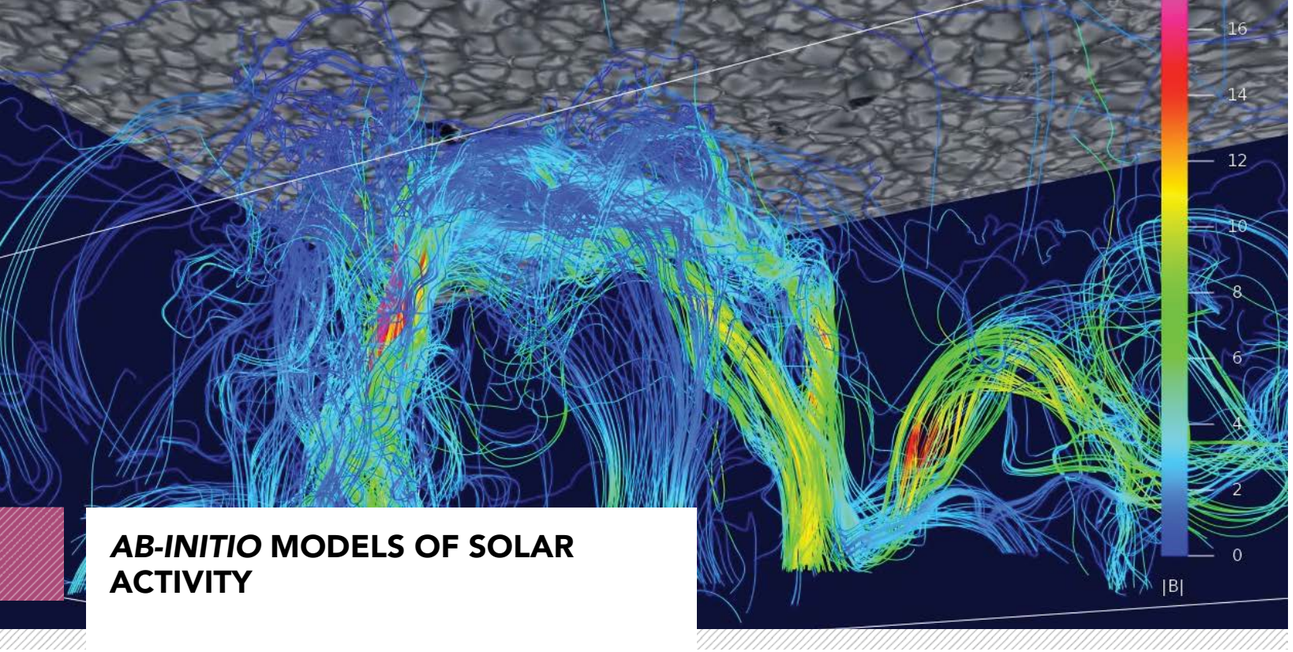

AB-INITIO MODELS OF SOLAR ACTIVITY

Allocation: NSF PRAC/6.5 Mnh
PI: Robert Stein[1]
Co-PIs: Aake Nordlund[2], Mats Carlsson[3], William Abbott[4], Bart De Pontieu[5]
Collaborators: Viggo Hansteen[3], Troels Haugbolle[2], Juan Martinez-Sykora[5], Tiago Pereire[3], Jesper Schou[6]

[1]Michigan State University
[2]Niels Bohr Institute, University of Copenhagen
[3]University of Oslo
[4]University of California, Berkeley
[5]Lockheed Martin Solar and Astrophysics Laboratory
[6]University of Göttingen

EXECUTIVE SUMMARY:

The objective of this project is to understand how solar magneto-convection powers the Sun's activity, heats the chromosphere and corona, and accelerates charged particles. To achieve this we have begun modeling the emergence of dynamo-generated magnetic fields through the solar surface. The first task was to thermally and dynamically relax a model of solar magneto-convection from the top of the photosphere to a depth of 30 megameters (Mm)—one million meters—in order to use results from the deep convection zone dynamo and flux emergence calculations as boundary conditions for the surface magneto-convection simulations. A second task was to improve the code performance by implementing MPI decomposition in all three dimensions rather than just two.

INTRODUCTION

Our project attempts to model the following chain of processes that connects the inner workings of the Sun to the Earth and space weather:

- The Sun controls Earth's weather and space weather via radiation, coronal mass ejections into the solar wind, and energetic particles originating from solar active regions.
- In turn, the interaction of magnetic fields, convection, and radiation influence the Sun.
- How magnetic fields emerge through the photosphere of the Sun and are shuffled around by convective motions governs chromospheric and coronal heating and determines the generation of flares and coronal mass ejections.
- The solar convective dynamo influences the behavior of magnetic fields at the solar surface. New magnetic flux emerging in active regions interacts with existing fields to release huge amounts of energy when the fields reconnect. This heats the local coronal environment to many millions Kelvin and can produce flares and coronal mass ejections.

METHODS & RESULTS

An anelastic, global dynamo code is used to provide bottom boundary conditions for the 3D compressible, local thermodynamic equilibrium radiation, magneto-hydrodynamic code STAGGER, which models the upper (15% geometrically and 75% of the scale height) solar convection zone. These results in turn provide lower boundary conditions for the non-equilibrium 3D radiation, magneto-hydrodynamic code BIFROST to model the chromosphere and corona.

A major effort by co-PI Nordlund this past year was to enable MPI decomposition in the vertical and horizontal directions for radiation transfer and dynamics in STAGGER in order to allow larger blocks of data and reduce communication. The STAGGER code is used in many applications (star formation, turbulence, as well as stellar and solar convection) so this work can benefit many other research teams.

An ongoing effort is to obtain a thermally and dynamically statistically relaxed initial magneto-convection model from the solar surface down to 30 Mm depth. This is necessary to use the results of a global dynamo simulation as a bottom boundary condition for the spatial and temporal evolution of the magnetic field and possibly also for vertical velocity. The global simulation of the upper boundary is near 20 Mm depth and it is necessary to use values away from any boundary effects. The thermal relaxation time is several days at these depths, so an initial state close to thermal equilibrium is essential. An earlier extended model was found to be too far from thermal equilibrium to relax in a reasonable computation time. A new extended model was constructed extending a relaxed 20 Mm depth model in hydrostatic equilibrium. The density and internal energy were then adjusted to make entropy constant both on horizontal planes and with depth in the extended layers. Because the temperature and density fluctuations are very small fractions of their mean values at large depths a new table was constructed with finer resolution in density and energy per unit mass at large depths than near the surface.

WHY BLUE WATERS?

To model the emergence of magnetic flux through the solar surface requires simulating the time evolution of magneto-convection for many hours of solar time. To obtain results in a feasable time requires as many processors as possible. Currently, only Blue Waters provides a substantial number of useable processors. The relaxation of magneto-convection down to a depth of 30 Mm is running on 32,000 and 64,000 cores on Blue Waters.

PUBLICATIONS

Stein, R., and A. Nordlund, Simulations of Magnetic Flux Emergence. *40th COSPAR Scientific Assembly,* Moscow, Russia, August 2–10, 2014.

Granules, Flux Emergence and Subsurface Structure of Sunspots. *Max Planck Institute for Solar System Research colloquium,* Göttingen, Germany, January 14, 2015.

Flux Emergence and Active Region Formation. *NORDITA invited talk,* Stockholm, Sweden, March 10, 2015.

FIGURE 1 (BACKGROUND): Magnetic field lines formed into a loop by the up and down motions of convecting plasma. The loop apex in ascending plasma moves toward and then through the visible solar surface, leaving behind legs in downflowing plasma that produce dark pores and Sunspots because convective transport of energy to the surface is inhibited.

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

Allocation: PRAC/3.0 Mnh
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]University of Minnesota, Minnesota Institute for Astrophysics
[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:

The scale and speed of Blue Waters enable 3D simulations of brief events in the evolution of stars that can have profound impacts upon their production of heavy elements. Hydrogen ingestion flash events have been identified as potential sites for the origin of anomalous abundance signatures in stars that formed in the early universe. We have simulated H-ingestion in a young white dwarf star, Sakurai's object, and 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 earlier 1D modeling of this phenomenon.

INTRODUCTION

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 late stages of evolution of the first generations of 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 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 be dragged down as the result of convective boundary mixing. In order to understand the H-ingestion flashes, as well as the evolution of many other types of stars, it is critically important to be able to quantitatively simulate this convective boundary mixing.

At sufficient depth within the convection zone, the entrained hydrogen can react with the abundant 12C to form 13N, with a significant release of energy. The 13N decays, with a half-life of about 10 minutes to produce 13C and more energy. At the very high temperatures near the helium-burning shell, this 13C reacts with the remaining helium to produce oxygen and a free neutron at a very high rate. Neutrons produced in this way are then captured by trace concentrations of heavy elements in the gas to build an entire series of progressively heavier nuclei. This material can be expelled from the star along with its outer envelope as it forms a planetary nebula.

METHODS & RESULTS

Using 440,000 processor cores on Blue Waters, we simulated a very low metallicity AGB star of the early universe on a grid of 3.6 billion cells for more than 9 million time-steps. The simulated time is nearly two 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 30 to more quickly traverse a slow ingestion period leading up to the flash.

Figure 1 shows six images of the local volume fraction of entrained hydrogen plus helium gas. We have cut away the front hemisphere of the star and look into the far hemisphere. We make the gas of the star's center transparent, whereas unprocessed hydrogen plus helium gas above the convection zone can be identified by its continuous red color. The gas that is otherwise visible consists of mixtures of the convection zone gas and the unprocessed hydrogen and helium mixture pulled from above the convection zone. The hydrogen concentration increases as the color changes from dark blue to aqua, then to white, yellow, and red.

The first image shows the star during the gradual buildup of entrained hydrogen in the upper part of the convection zone. Rising plumes of hotter, buoyant gas reaching the top of the convection zone spread horizontally along that upper surface. When they encounter other spreading flows, they are forced to turn downward, and it is here, along lanes separating convection cells, that buoyant hydrogen gas becomes entrained in the downward flow. These lanes appear dark blue because we see the low hydrogen concentrations in the gas first looking from the center of the star. Where the gas of the convection zone scrapes along the upper surface in a nearly horizontal flow, we see easily through the very thin layers of low hydrogen concentrations to the white, yellow, and red of higher hydrogen concentrations above.

In the second image, we see these higher concentrations mainly in cross-section around the edges of the spherical convection zone top. A layer of gas enriched with hydrogen has accumulated at the top of the convection zone. At the lower left, we see high hydrogen concentrations in red and yellow, where a wave traveling around the top surface of the convection zone is peeling off a slice of the hydrogen-helium gas above. This is the characteristic sign of a global oscillation of shell hydrogen ingestion, which we call a GOSH. It is a result of unstable hydrogen burning that we have discovered in this work by adopting a global, 3D description of the flow in contrast to previous models in 1D approximation. These waves rushing back and forth can raise the hydrogen ingestion rate by two orders of magnitude or more, which is expected to have significant impact on heavy-element production.

In the image at the middle left, the layer of buoyant, hydrogen-rich gas near the top of the convection zone impedes convective motions, and further hydrogen ingestion proceeds more slowly. But heat continues to be deposited into the convection zone from helium burning and from hydrogen burning at the base of this upper layer, and buoyant plumes punch into the upper layer of hydrogen-rich gas (middle right image) and bring that gas down to the high temperatures, where it burns violently. The resulting bursts of huge, localized energy release cause violent updrafts that set off a more powerful GOSH, seen in the bottom two images.

WHY BLUE WATERS

The sustained petaflops capability of Blue Waters makes it possible to perform 3D simulations that allow us to explore and understand phases of stellar evolution that are long in duration compared to a stellar explosion but short compared to those behaviors that are well approximated by 1D codes.

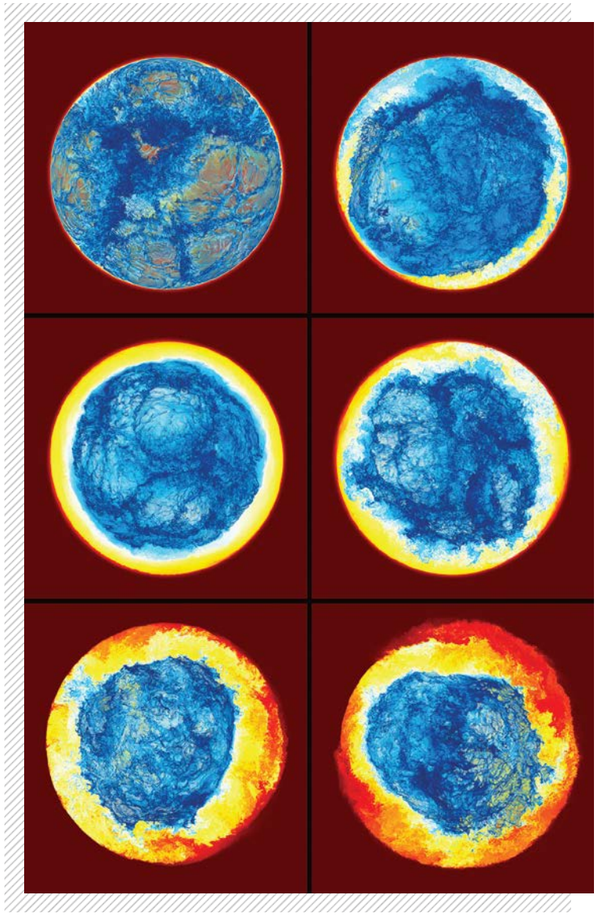

FIGURE 1: Entrainment of hydrogen-rich gas into the helium shell flash convection zone of a very low metallicity star and the subsequent development of the Global Oscillation of Shell Hydrogen ingestion (GOSH).

PUBLICATIONS

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.

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.

THREE-DIMENSIONAL CORE-COLLAPSE SUPERNOVA SIMULATIONS

Allocation: PRAC/NSF 3.2 Mnh
PI: Adam Burrows[1]

[1]Princeton University

EXECUTIVE SUMMARY:

Core-collapse supernovae are the most common type and occur when the iron core of a massive star collapses to a neutron star. Releasing its gravitational binding energy as neutrinos, the protoneutron star (PNS) for a few seconds outshines the rest of the observable universe. Capturing a small fraction of this energy is thought to power the explosion of at least the most frequent supernovae, though detailed calculations proving this paradigm are still lacking. We conduct three-dimensional radiation/hydrodynamic simulations of core-collapse supernovae with the goal of determining the mechanism of explosion, using the newly developed and tested code Fornax, incorporating state-of-the-art microphysics and methodologies.

INTRODUCTION

Viewed as a nuclear physics laboratory, core collapse supernovae produce the highest densities of matter and energy in the modern universe. They thus probe the nuclear and particle physics of matter at super-nuclear density, high temperature, and at extremes of isospin. They are also responsible for creating most of the elements in nature, many of which are produced as radioactive precursors, both neutron-rich (r- process) and neutron-poor (explosive nucleosynthesis). The neutrino signal they emit carries information about the nuclear equation of state, and the strength of their explosion is sensitive to how both the neutrinos and ultra-dense matter are treated. Supernovae, thus, probe the same sort of physics as the Facility for Rare Isotope Beams, the Argonne Tandem Linac Accelerator System, and low-energy runs at Brookhaven's Relativistic Heavy Ion Collider, the Facility for Antiproton and Ion Research in Damstadt, Germany, CERN's SPS Heavy Ion and Neutrino Experiment, and Russia's Nuclotron-based Ion Collider Facility. Our project employs our new state-of-the-art radiation/hydrodynamics code Fornax to simulate the physics of core collapse and explosion in three spatial dimensions. This effort supports NSF's experimental nuclear physics program by exploring nucleosynthesis in astrophysical explosions, the properties of the neutrino, and the equation of state and phases of dense nuclear matter. A solution to the core-collapse supernova problem would benefit ongoing efforts of observers and instrument designers in the United States and around the world engaged in projects to determine the origin of the elements, measure gravitational waves (LIGO), and interpret laboratory nuclear reaction rate measurements in light of stellar nucleosynthesis.

METHODS & RESULTS

Fornax is a new code for multidimensional, self-gravitating, radiation hydrodynamics that is second-order accurate in space and time and was designed from scratch with the core-collapse supernova problem in mind. The code solves the equations of compressible hydrodynamics with an arbitrary equation of state, coupled to the multigroup two-moment equations for neutrino transport and the Poisson equation for gravity. The equations are solved on a fixed Eulerian grid, enhanced by a mechanism for static mesh derefinement in the inner zones. All of the these transport related terms are treated explicitly in time and are updated with the same time-step as the hydrodynamics. After core bounce, explicit time integration is not only simpler and generally more accurate, it is also faster than globally coupled time-implicit transport solves that are typically employed in radiation hydrodynamics methods. The latter allows us to avoid angular Courant limits at the center, while maintaining accuracy and enabling us to employ the useful spherical coordinate system natural for the supernova problem. We have spent much of the first phase of our Blue Waters effort developing, testing, and running Fornax. Our project has received a modest extension, and we hope to have 3D simulation results to report in 2016.

WHY BLUE WATERS

It is not possible to explore this central problem in nuclear astrophysics without supercomputer resources such as Blue Waters. Its Cray architecture comports very nicely with our code architecture, and the large number of cores available enables the exploitation of our code's efficient parallelism. On a small, Beowulf-like cluster, we estimate that one calculation would take a few decades to perform, while requiring only one month on Blue Waters, using one-quarter of the machine.

GEOSCIENCE

WEATHER

CLIMATE

GEOLOGY

ENVIRONMENT

40 *Global Observations of Cloud Microphysics through Terra Data Fusion*

42 *Large Eddy Simulation of Sediment Transport and Hydrodynamics at River Bifurcations*

44 *Earthquake System Science Research*

46 *Simulating Cumulus Entrainment: A Resolution Problem, or Conceptual?*

48 *Design and Management of Satellite Assets to Advance Space-Based Earth Science*

50 *Location-Specific Space Weather Hazards to Electric Power Grids Calculated on a Global Scale*

52 *Scaling the CESM to Ultra-High Resolutions for Analyzing Tropical Cyclone–Climate Feedbacks*

54 *Petascale modeling of convective storms under climate change and variability*

56 *Applicability of Geometric Optics Method for Calculations of Single-Scattering Properties of Atmospheric Ice Crystals*

58 *Computational Mineral Physics*

61 *Using Petascale Computing Capabilities to Address Climate Change Uncertainties*

64 *Long-Track EF5 Tornado Simulation*

GLOBAL OBSERVATIONS OF CLOUD MICROPHYSICS THROUGH TERRA DATA FUSION

Allocation: Blue Waters Professor/0.24 Mnh
PI: Larry Di Girolamo[1]
Collaborators: Alexandra Jones[1], Lusheng Liang[2], Guangyu Zhao[1]

[1]University of Illinois at Urbana-Champaign
[2]Science Systems and Applications, Inc. and NASA Langley Research Center

EXECUTIVE SUMMARY:

The Moderate Resolution Imaging Spectroradiometer (MODIS) is our most advanced satellite instrument for long-term, global observations of effective radius (Re) of the cloud drop size distribution. Here we show that the fusion of data collected by MODIS and the Multi-angle Imaging SpectroRadiometer (MISR), both on NASA's Terra satellite, can be used to characterize the space-time varying biases in the MODIS Re product. The biases are determined to range from 3 to 11 μm for marine liquid water clouds, which, when removed, provides a very different understanding of the distribution of liquid cloud drop sizes in our atmosphere compared to what previously was determined from the original MODIS data. This greatly advances the utility of the Terra data for weather and climate research.

INTRODUCTION

Clouds cover about 68% of our planet. They are one of the most interconnected components of the Earth system, playing a key role in the Earth's hydrological cycle, regulating the incident solar radiation field more than any other atmospheric variable, and acting as the most important greenhouse constituent in our atmosphere. As such, they modulate a wide range of processes on Earth. The Intergovernmental Panel on Climate Change (IPCC) affirms that the role of clouds remains the leading source of uncertainty in anthropogenic climate change predictions. In addition, the role of cloud microphysics and cloud-radiation interactions in the timing and intensity of weather events remains an active area of research.

To make headway in reducing uncertainty in weather and climate predictions, the World Meteorological Organization and the IPCC defined a list of Essential Climate Variables (ECVs) requiring global observations (http://www.gosic.org/gcos). This list includes a wide range of cloud properties, most of which are primarily derived from satellites. While significant advancements have been made over the past four decades in retrieving cloud properties from space, most still do not meet accuracy requirements for climate research [1].

Here, we focus on Re for liquid water clouds over the ocean. Re is an ECV with a target accuracy of 5% or 0.5 μm. This accuracy has not yet been met, largely because the remote sensing algorithms used in the retrieval of Re assume clouds and the radiative boundary conditions to be horizontally homogeneous. This greatly simplifies the radiative transfer to 1D. However, a simple look at clouds shows that they are not horizontally homogeneous over a wide range of scales. This leads to a bias in the retrieved Re when the algorithms are applied to structured (3D) clouds. The bias co-varies with the underlying structure of the cloud field and the sun-view geometry, leading to an inability to disentangle true space-time variability in cloud drop sizes from variability in the biases caused by the 1D assumption [2]. Based on 3D radiative transfer simulations applied to a small number of synthetic 3D cloud fields, it has been suggested that errors as high as 100% are possible [3]. The extent to which the results from these simulations are globally representative in nature remains unknown–until now.

METHODS & RESULTS

MODIS provides a global dataset of Re based on the spectral signatures that it measures and a 1D radiative transfer formulation of the inverse problem. Also on Terra is MISR, which samples the angular anisotropy of upwelling scattered sunlight. As fully described in [4], spectral signatures from MODIS and angular signatures from MISR were used synergistically to provide upper and lower bound estimates on MODIS Re-biases in the face of 3D radiative effects.

Figure 1 shows an example of our results: the January zonal-mean Re and bias. The results show that the zonal variations in the bias meet the expectation that regions with more cumulus (structured) cloud contribution to the total cloud samples (e.g., ~10°N and ~35°S) will have the largest bias in Re due to the greater deviations from the 1D assumption used in the MODIS retrievals. Quantitatively, the biases are large. The sample zonal mean values from MODIS Re range from about 13 to 18 μm, while the bias-corrected values range from 7 to 12 μm. The latter is in much better agreement with in situ observations of marine liquid water clouds [5]. In addition, latitudinal gradients in zonal-mean Re are also very different between the original and bias-corrected products. Given these large spatially varying biases, implications for studies that use the original MODIS products–say for retrieving liquid water content [6] and cloud drop number concentration [7], evaluating cloud microphysical parameterizations in climate models [8], and in characterizing cloud-aerosol interactions [9]–need to be drawn out.

Our characterization of the space-time varying biases in the MODIS Re product advances the utility of this important dataset for weather and climate research and paints a radically different picture on the distributions of liquid cloud drop sizes in our atmosphere compared to what was previously determined from the original MODIS data. It is also an excellent demonstration of what can be accomplished through MISR and MODIS data fusion. With more than 1 PB of Terra data and counting, such data fusion projects call for high-performance computing facilities and enhanced national infrastructures for data distribution. Our work on Re has provided an excellent use case for improving this infrastructure, as highlighted below.

WHY BLUE WATERS

The Terra data is > 1 PB and growing. Key advantages of using Blue Waters for access, usage, and distribution of Terra data fusion products are that the Terra data and processing are local, with access and sharing that are global. 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 optimum framework for large-scale processing, analytics, and mining of the entire Terra record. In addition, the Blue Waters staff provides expertise critically needed to optimize workflows. For example, Blue Waters staff worked closely with NASA to make substantial improvements to NASA's cyberinfrastructure for transferring MISR data from its data center to Illinois.

PUBLICATIONS

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.

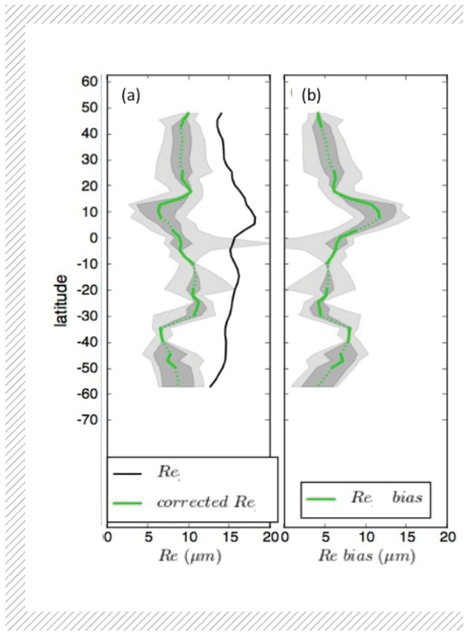

FIGURE 1 (RIGHT): (a) MODIS-Re (black) and bias-corrected Re (green). The dark shade represents the region between the upper and lower bounds of Re, and the light shade represents the region that is one standard deviation of the sample mean below the lower bound and above the upper bound of Re. The green line connects the midpoints of upper and lower bounds of Re. (b) Biases in MODIS-Re.

LARGE EDDY SIMULATION OF SEDIMENT TRANSPORT AND HYDRODYNAMICS AT RIVER BIFURCATIONS

Allocation: Illinois/50 Knh
PI: Marcelo H. Garcia[1]
Co-PI: Paul Fischer[1]
Collaborator: Som Dutta[1]

[1]University of Illinois at Urbana–Champaign

EXECUTIVE SUMMARY:

A fundamental morphological element present in most river systems is a bifurcation. A diversion is a special kind of bifurcation in which one of the post-bifurcation channels continues along the direction of the original un-bifurcated channel. Previous studies have shown that a larger percentage of near-bed sediments at diversions tend to enter the new channel, even in cases where the opposite trend is exhibited by the water discharge. This study attempts to bolster fundamental understanding of this and related phenomena, such as secondary flows and vorticity-driven sediment transport.

Our simulations resolved all the relevant turbulent eddies of the flow, which not only provided an accurate description of the dynamics of the flow, but also helped model sediment transport accurately when coupled with a Lagrangian particle model for the sediment. We ran computational fluid dynamics simulations using the incompressible Navier–Stokes solver Nek5000. The maximum bulk Reynolds (Re) number of the completed simulations was 7,000; more simulations with higher Re and different diversion angles are planned.

INTRODUCTION

Most river systems contain one or more bifurcations, where a river divides into two channels that each carries part of the water and sediment. Most bifurcations form naturally, like the dendritic networks in deltas, the in-stream bifurcations in braided rivers, etc. Bifurcations have also been built for various engineering purposes, like river connectivity, flood protection, etc. The engineered bifurcations are more commonly referred to as diversions. In a diversion, a lateral channel comes out from the side of the original channel, which continues along the original stream path.

In order to understand the division of flow and near-bed sediment between the main and the diverted channels, H. Bulle conducted the first systematic study of the division of near-bed sediment between bifurcating channels [1]. Through his laboratory experiments, Bulle showed that the near-bed sediment discharge distribution tends to favor the lateral channel, even in cases where water discharge favors the original channel; this nonlinear phenomenon is known as the Bulle Effect.

Interest in fluvial diversions has spiked in the last decade as old diversion channels clog with sediment and engineers look to diversions as a possible means to maintain deltas in the face of rising sea levels [2]. For example, several designs are being evaluated to mitigate potential loss of coastal land and delta degradation in the Lower Mississippi River in Louisiana [3]. Finding the optimal location and layout of the diversions will be more efficient if the fundamental mechanism that causes the Bulle Effect is well understood. On the other side of the world, the Indian government plans to interconnect several rivers [4], many of which carry huge sediment loads that may be redistributed by fluvial diversions.

Apart from improving fundamental understanding of the fluid dynamics at river diversions and its implications on morphological evolution, the current study will also help improve numerical models of bifurcations used for field-scale simulations. A better understanding of the fundamental mechanism behind the Bulle Effect will help shed light on the related phenomenon of vorticity-driven sediment transport, which affects both natural and man-made systems [5].

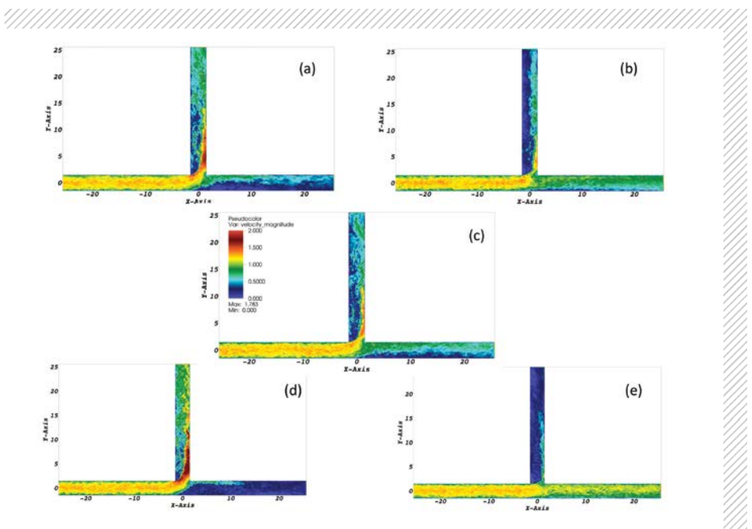

FIGURE 1: Velocity magnitude at a height of 0.75 from the bottom (where channel height is 1.0) for a bulk Reynolds number of 7,000, for different ratios of flows going through each channel: in percentages, (a) 65-35, (b) 35-65, (c) 50-50, (d) 85-15, and (e) 15-85.

METHODS & RESULTS

The simulations had two parts: the fluid portion that was modeled using the 3D incompressible Navier–Stokes equation with the large eddy simulation (LES) approach, and the sediment transport that was modeled as Lagrangian particles [7, 8]. In the LES approach the relevant scales of the flow are resolved and a subgrid stress model simulates the subgrid scales [9]. The simulations are being conducted using Nek5000, an open-source, highly-scalable eddy-resolving incompressible Navier–Stokes solver based on the spectral element method (SEM) [10]. The SEM combines the accuracy of spectral methods and the flexibility of numerical methods based on local approaches, like the finite elements method [11].

The first group of cases that was simulated for the current study was for a 90° diversion angle and bulk Reynolds number of 7,000. Five simulations with different water discharge ratios between the two channels (in percentages: 50-50, 65-35, 35-65, 85-15, and 15-85) were conducted using about 224 million computational points. Velocity magnitudes at a height of 0.75 from the channel bottom (where channel height is 1.0) were plotted for the simulated cases (fig. 1). A distinct and sustained high-velocity zone hugging the right side of the diverted channel was observed in all the cases. The high-velocity zone was expected, but it was surprising to see the relative lack of mixing between the high-velocity and low-velocity zones in the diverted channel, especially for cases where flow in the diverted channel was less than the main channel. The simulations were able to successfully capture the two separation zones in the flow in both channels, though the separation zone in the main channel was more persistent than expected. The structure of the flow clearly indicated the cause of the Bulle Effect. These preliminary results also show how HPC can be used to accurately simulate eddy-resolving river flows at an unprecedented scale.

WHY BLUE WATERS?

The current study pushes the limit of the scale at which eddy-resolving numerical simulations have been used to study complex multi-phase river mechanics. The simulated cases herein required 224 million computational points, and this number can go up to a billion for a Reynolds number of the order 10^4–10^5. Additional computing time would be required to model the sediment using Lagrangian particle tracking; thus, researchers need petascale facilities in order to complete the simulations in a realistic amount of time. Access to the next generation of Track-1 systems will allow us to conduct eddy-resolving simulations of environmental flows at scales and with complexity similar to that of nature. This will allow us to better understand the underlying mechanisms of different environmental flows and help improve predictions of different natural processes.

EARTHQUAKE SYSTEM SCIENCE RESEARCH

Allocation: NSF PRAC/6.3 Mnh
PI: Thomas H. Jordan[1,2]
Co-PIs: Jacobo Bielak[3], Yifeng Cui[4], Kim Bak Olsen[5]
Collaborators: Scott Callaghan[1], Ewa Deelman[2], David Gill[1], Edward Field[6], Robert Graves[6], Gideon Juve[2], Naeem Khoshnevis[7], Philip J. Maechling[1], Kevin Milner[1], Doriam Restrepo[8], William Savran[5], Zheqiang Shi[5], Patrick Small[2], Ricardo Taborda[7], Karan Vahi[2], Kyle Withers[5]

[1]Southern California Earthquake Center (SCEC)
[2]University of Southern California
[3]Carngie Mellon University
[4]San Diego Supercomputing Center
[5]San Diego State University
[6]United States Geological Survey (USGS)
[7]University of Memphis
[8]Universidad EAFIT Medellin Colombia

EXECUTIVE SUMMARY:

SCEC's multi-disciplinary research team is using Blue Waters to develop high-resolution computational models of earthquake processes and to calculate physics-based probabilistic ground motion forecasts for selected urban areas in the United States. The U.S. Geological Survey (USGS), through its National Seismic Hazard Mapping Project (NSHMP), currently uses empirical probabilistic seismic hazard analysis (PSHA) to promote seismic safety engineering and disaster preparedness across the United States, including California. PSHA is the scientific basis for many engineering and social applications including performance-based design, seismic retrofitting, resilience engineering, insurance rate setting, disaster preparation and warning, emergency response, and public education. SCEC's research goal is to develop physics-based models for the urban regions of California that are more accurate than the empirical NSHMP standard. Our long-term goal is to extend the more accurate physics-based PSHA across the full bandwidth needed for seismic building codes, simulating ground motions at frequencies up to 10 Hz. We are working with the USGS to integrate SCEC computational results into NSHMP products that will benefit end-users of earthquake information including scientists, engineers, and the public.

INTRODUCTION

Over the last decade, SCEC's earthquake system science research program has developed detailed earth structural models and HPC software needed to perform realistic, physics-based earthquake simulations. During that time, SCEC has verified the computational readiness and scalability on Blue Waters for dynamic rupture, deterministic earthquake wave propagation, and probabilistic seismic hazard simulation codes.

METHODS & RESULTS

Over the last year, SCEC used Blue Waters to conduct earthquake system science computational research in two main areas. First, we integrated advanced physics into deterministic earthquake wave propagation software. Second, we calculated the first 1 Hz CyberShake hazard model, doubling the maximum frequency of the previous 0.5 Hz CyberShake hazard model completed in 2014.

We integrated more realistic physics into our high-performance earthquake simulation software to model frequency-dependent attenuation [1], free-surface topography [2], and non-linear yielding effects [3]. All of these effects become increasingly important when simulating high-frequency ground motions.

In March 2015, SCEC's GPU-based, CUDA-language, high-performance wave propagation software received the NVIDIA Global Impact Award [4–6]. We integrated this high-performance GPU-based code into our physics-based probabilistic seismic hazard workflows. This enabled SCEC to make use of highly efficient GPUs for our floating-point-intensive processing.

Starting in April 2015, our SCEC team ran the first 1 Hz physics-based Los-Angeles-area seismic hazard model. SCEC used Blue Waters, together with Titan (Oak Ridge Leadership Computing Facility, OLCF), over seven weeks to produce this hazard calculation for scientists and civil engineers. These recent results, called CyberShake Study 15.4, combined the large number of GPU nodes on Titan with both GPU and CPU nodes on Blue Waters to reduce the time-to-solution for a CyberShake model calculation from months to weeks. This CyberShake calculation represents a collaboration that includes the University of Southern California, NSF Track-1 facilities, and Department of Energy Leadership HPC centers working together and with scientists and engineers on socially important, broad-impact, HPC-based research.

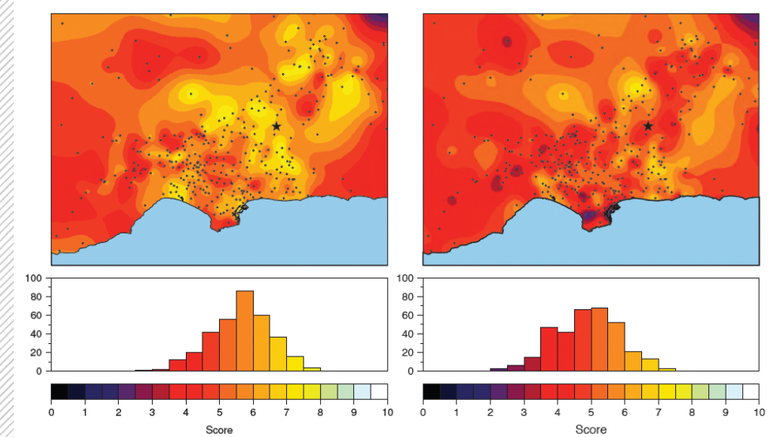

FIGURE 1: Maps of the Los Angeles region that compare SCEC simulated ground motions to observed ground motions using two different 3D Earth structure models as input to the 1 Hz ground motion simulation (left: CVM-S4.26, right: CVM-H). Red/purple indicates a poor match and yellow indicates a good match using a goodness-of-fit method. The star indicates the epicenter of the 2008 M5.4 Chino Hills earthquake and dots show ground motion recordings for the event. [Credit: Ricardo Taborda, University of Memphis; Jacobo Bielak, Carnegie Mellon University]

WHY BLUE WATERS?

SCEC computational research continues to expand in multiple dimensions, including geographical range, higher resolution, and more time steps, so new simulations require more computational, memory, and storage resources.

SCEC computational needs continue to grow because individual earthquake simulations do not "solve" a problem when run once. In many cases, seismic hazard calculations involve use of fault models and earth structure models. When these structural models are updated, simulations need to be run again, with new inputs, in order to produce updated seismic hazard estimates. Re-running simulations with new inputs increases our need for computational time.

Deterministic ground motion simulation techniques work well at 1 Hz and below in many areas. The maximum necessary simulated frequency for global earthquake simulations is about 1 Hz. The maximum necessary simulated frequency for seismic hazard analysis is above 10 Hz. Ground motion modelers must implement improved physics and improved code performance in order to increase the maximum simulated frequency of wave propagation codes to the required higher frequencies. Due to the large node count and significant memory per core on Blue Waters, SCEC researchers were able to perform ground motion simulations on Blue Waters at 4 Hz frequencies using two different computational methods. Validation of these new simulation capabilities will involve simulating well-recorded historic California earthquakes and comparing the simulated ground motions against the recorded ground motions [7].

SCEC's earthquake system science computational research uses a broad range of codes and system capabilities. Blue Waters provides the broad range of system capabilities we need in our computation research including CPUs with a significant amount of RAM as well as many highly efficient GPU nodes. Blue Waters also provides very large online disk storage that is very valuable for storing temporary intermediate files during long-duration calculation. The Blue Waters software stack and computing policies also provide very valuable support for SCEC's large-scale scientific workflows that rely on scientific workflow management tools and remote job submission support.

Blue Waters helps reduce the time-to-solution for SCEC's long-duration CyberShake calculations to manageable levels. SCEC is currently collaborating with a civil engineering group on CyberShake hazard model development. The project requires a six-month update cycle, so

scientists and engineers meet every six months to review the updated CyberShake hazard model. Blue Waters enables SCEC to meet this schedule because a full CyberShake simulation can run within one review cycle.

SCEC plans to expand use of its most advanced modeling from southern California to more regions. Several specific goals, including a statewide California CyberShake seismic hazard model, will require the computing capabilities of the next generation of Track-1 systems.

PUBLICATIONS

Roten, D., and Y. Cui, GPU-Powered Simulations of Seismic Waves in Nonlinear Media. *GPU Technology Conference,* San Jose, Calif., April 4–8, 2015.

Isbiliroglu, Y., R. Taborda, and J. Bielak, Coupled Soil-Structure Interaction Effects of Building Clusters During Earthquakes. *Earthquake Spectra,* 31:1 (2015), pp. 463–500, doi:10.1193/102412EQS315M.

Poyraz, E., H. Xu, and Y. Cui, Application-specific I/O Optimizations on Petascale Supercomputers. *Procedia Comput. Sci.,* 29 (2014), pp. 910–923, doi:10.1016/j.procs.2014.05.082.

Wang, F., and T. H. Jordan, Comparison of probabilistic seismic hazard models using averaging-based factorization. *Bull. Seismol. Soc. Am.,* 104:3 (2014), pp. 1230–1257, doi:10.1785/0120130263.

Xu, H., et al., Aftershock sequence simulations using synthetic earthquakes and rate-state seismicity formulation. *Earthquake Sci.,* 27:4 (2014), pp. 401–410, doi:10.1007/s11589-014-0087-7.

SIMULATING CUMULUS ENTRAINMENT: A RESOLUTION PROBLEM, OR CONCEPTUAL?

Allocation: BW Professor/0.245 Mnh
PI: Sonia Lasher-Trapp[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY:

Understanding and predicting the rate at which cumulus clouds deplete their water content by the entrainment of dry air, affecting their development, longevity, and ability to precipitate, has been elusive. Simulations performed on Blue Waters are enabling us to investigate if the problem lies in our ability to represent the smallest scales of turbulence in our models, or if our underlying conceptual models of the physics of entrainment are flawed, or both.

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. Entrainment is the process by which clouds bring dry air from outside the cloud inward. It can initially favor precipitation formation; eventually, it dilutes the cloud and encourages its demise. A long-standing problem in meteorology has been to reproduce how quickly entrainment dilutes a cumulus cloud. Currently, all models fail. This has long been assumed to be a problem of inadequate spatial resolution, where the smallest scales of turbulence must be parameterized and their effects are improperly represented. It could also result from a fundamental problem in our conceptual understanding of the entrainment process. Our goals are to test both possibilities.

METHODS & RESULTS

We've run numerous simulations at relatively coarse (50 m) resolution in order to see how entrainment in a single cumulus cloud differs due to basic physical parameters, such as the strength of the cloud forcing, the size (width) of the cloud, and the amount of wind shear (the change in wind speed with height) in the atmosphere surrounding the cloud. We've also developed tools to quantify the entrainment that is occurring in the simulated clouds as they grow in time. As predicted from laboratory and theoretical models of thermals, narrower clouds are diluted by entrainment more quickly, helping to validate the cloud model results. However, contrary to theoretical models, weaker clouds, i.e. those having weaker updrafts, also appear to be diluted more quickly than stronger clouds. This finding is perplexing, and we continue to explore related parameter spaces to determine if the result is physically viable or due to a computational artifact in cloud models. Future comparison of the simulation results with aircraft observations of real cumulus clouds will also be key to understanding these findings.

WHY BLUE WATERS

Our Blue Waters allocation is essential for testing the resolution-dependency of the entrainment process and in particular for determining the sizes of the eddies that are most critical to represent in simulations of cumulus entrainment. Blue Waters–with its huge number of nodes, its high speed, and its large storage capability for high-resolution model output and analysis–allows us to push the spatial scale limit much farther than in the past. We intend to increase the resolution to as high as 5 m (over domain sizes of 10 km or greater) in order to understand any computational issues related to cumulus entrainment, in addition to improving our knowledge of the underlying physics. The computational demands of these large simulations quickly supersede the limits of most computers

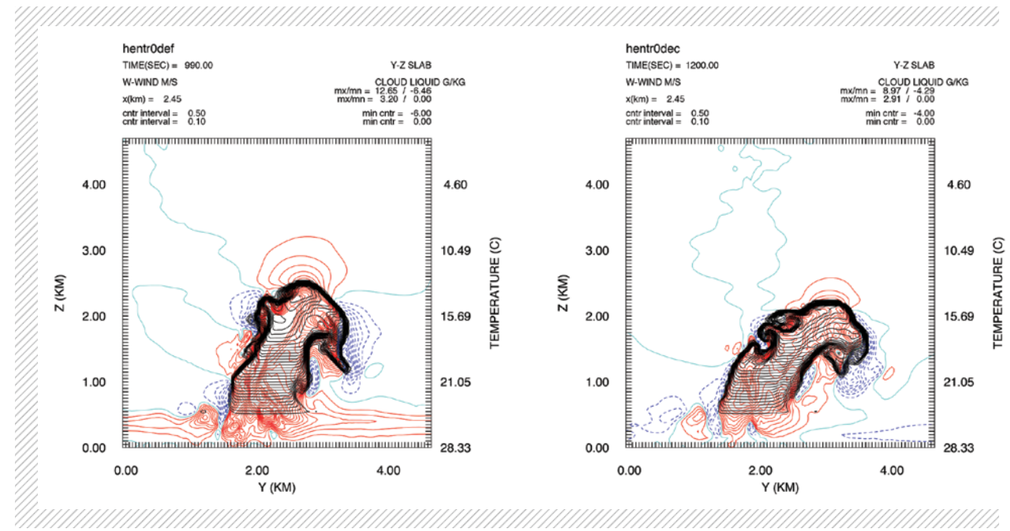

FIGURE 1: Vertical cross-sections through two cloud simulations, showing the difference in vertical configurations of the clouds with respect to initial cloud forcing. Vertical velocities are contoured in increments of 0.5 m s-1, where red indicates upward motion and blue indicates downward motion. Black contours indicate cloud water mass in increments of 0.1 g per kg of air. The left panel shows results for stronger cloud forcing, producing a more upright cloud, and the right panel shows results for weaker cloud forcing, where the cloud is more susceptible to leaning due to the ambient wind flow (not shown). The leaning cloud appears to entrain more dry air than the former, for similar cloud top heights.

DESIGN AND MANAGEMENT OF SATELLITE ASSETS TO ADVANCE SPACE-BASED EARTH SCIENCE

Allocation: NSF PRAC/4.53 Mnh
PI: Patrick M. Reed[1]
Co-PIs: Matthew Ferringer[2], Eric F. Wood[3]

[1]Cornell University
[2]The Aerospace Corporation
[3]Princeton University

EXECUTIVE SUMMARY:

This project is advancing a petascale planning framework that is broadly applicable across space-based Earth observation systems design. We made substantial progress toward three transformative contributions: (1) we are the first team to formally link high-resolution astrodynamics design and coordination of space assets with their Earth science impacts within a petascale "many-objective" global optimization framework; (2) we successfully completed the largest Monte Carlo simulation experiment for evaluating the required satellite frequencies and coverage to maintain acceptable global forecasts of terrestrial hydrology (especially in poorer countries); and (3) we are initiating an evaluation of the limitations and vulnerabilities of the full suite of current satellite precipitation missions including the recently approved Global Precipitation Measurement (GPM) mission. This work will illustrate the tradeoffs and consequences of the GPM mission's current design and its recent budget reductions.

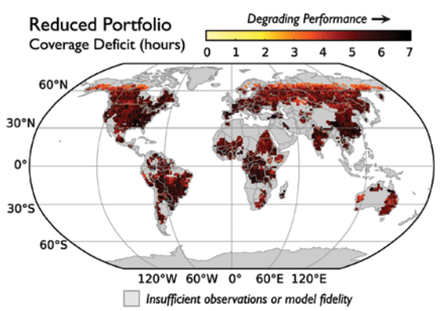

FIGURE 1: Global map of coverage deficits that emerge if the current four precipitation missions beyond their design life are lost (i.e. the reduced portfolio). Coverage deficits represent worst-case periods of missing data.

INTRODUCTION

Our team is bridging state-of-the-art innovations in satellite constellation design optimization and ensemble-based global terrestrial water and energy prediction to directly clarify satellite systems' value to managing floods and droughts. More broadly, our team is contributing highly flexible petascale space-based observation design tools that can serve a critical role for realizing the integrated global water cycle observatory long sought by the World Climate Research Programme, which to date has eluded the world's space agencies. Our research is critical for the scientific and space agency communities to overcome current computational barriers to transform the optimization of future satellite constellation architectures for delivering high-fidelity data to a broad array of applications. Similarly, we envision that there is a broad array of scientists and users whose future activities will draw upon the project's scientific findings and generated data. As examples, the water-centric stakeholder community desperately requires improved monitoring and assessment of the water cycle for improved decision making related to flooding, droughts, and food and energy security.

METHODS & RESULTS

We can categorize our project accomplishments to date within three foci: (1) scalable many-objective design optimization benchmarks, (2) advances in the use of high-fidelity astrodynamics simulation to permit passive control (i.e. minimum-energy satellite constellations), and (3) benchmark the effects of reduced satellite-based observation frequencies of precipitation on global drought and flood forecasting.

1. With respect to many-objective design evaluation, we completed the largest and best benchmark in terms of search quality and scalability for our team's underlying optimization algorithms. The results were made possible by the Blue Waters friendly user period access. At 524,288 cores, our search approached theoretically ideal performance. These results are the best benchmark ever attained for the challenge problem of focus and provide a strong foundation for our future tradeoff analyses.

2. In the context of passive control, our preliminary results focused on the patented four-satellite "Draim" constellation. Our Draim results revealed that carefully optimizing an initial orbital geometry to exploit natural perturbations (e.g., effects of Sun, Moon, etc.) yields passively controlled satellite constellations that maintain near-continuous global coverage performance as a function of elevation angle. This minimizes propellant and station keeping requirements to dramatically reduce mission costs while increasing mission duration. The Draim constellation represents a stepping stone to the more complex suite of global precipitation missions that will require the analysis of more than ten satellites.

3. We are one of the first teams to show how limits in satellite-based precipitation observations propagate to uncertainties in surface runoff, evaporation, and soil moisture at distinctly different locations globally. Our results were based on the Variable Infiltration Capacity (VIC) global macroscale land surface model at 1.0° spatial resolution. For each realization of the VIC ensemble, each model grid cell's satellite precipitation is resampled at different temporal resolutions and then run through the VIC land surface model. Our results suggested differing effects of spatial and temporal precipitation sampling on each water cycle component. For example, convection plays a dominant role in the tropics and sampling will highly impact the measured precipitation. However, plant transpiration is impacted less by the intensity and frequency of storms than the sufficiency of the total precipitation. These insights have direct relevance to water security concerns in terms of floods and droughts.

WHY BLUE WATERS

In simple terms, the scale and ambition of our computational experiments require that we have the ability to compress years of computational work into minutes of wall-clock time to be feasible. Additionally, our applications are extremely data intensive, so Blue Waters' high core count and high memory have been fundamental requirements for our work. Our initial Blue Waters allocation results required over 4.5 million node-hours and the systematic processing of approximately 2 PB of model output to support preliminary contributions to the areas of global hydrology, massively parallel many-objective optimization, and high-fidelity astrodynamics design. The global hydrologic ensemble represents a new level dataset that will be of broad interest in a variety of Earth science and engineering applications. Our satellite design trade-off analysis clarified how quickly we deviate from the "best-case" observation frequencies, with limits on spending, limits in international coordination, neglect of hydrologic objectives, and the simplified astrodynamics simulations currently used.

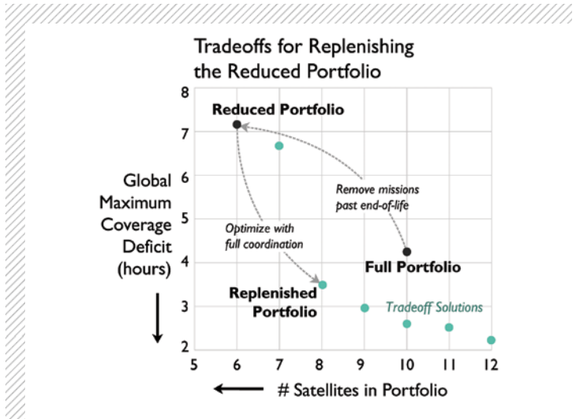

FIGURE 2: Tradeoffs between satellite count and coverage deficit for the full portfolio, the reduced portfolio, and a highlighted replenished portfolio. With coordination in planning, the replenished portfolio's eight satellites outperform the ten satellites of the full portfolio.

PUBLICATIONS

Chaney, N., J. Herman, P. Reed, and E. F. Wood, Flood and Drought Hydrologic Monitoring: The Role of Model Parameter Uncertainty. *Hydrol. Earth Syst. Sci. Discuss.*, 12 (2015), pp. 1697–1728.

Reed, P., N. Chaney, J. Herman, M. Ferringer, and E. Wood, Internationally Coordinated Multi-Mission Planning Is Now Critical to Sustain the Space-based Rainfall Observations Needed for Managing Floods Globally. *Environ. Res. Lett.*, 10 (2015), 024010, doi:10.1088/1748-9326/10/2/024010.

Ferringer, M., et al., A Framework for the Discovery of Passive-Control, Minimum Energy Satellite Constellations. Space 2014 AIAA/AAS Astrodyn. Specialist Conf. 2014, San Diego, Calif., August 4–7, 2014.

Reed, P., and D. Hadka, Evolving Many-Objective Water Management to Exploit Exascale Computing. *Water Resources Res.*, 50 (2014), pp. 8367–8373, doi:10.1002/2014WR015976.

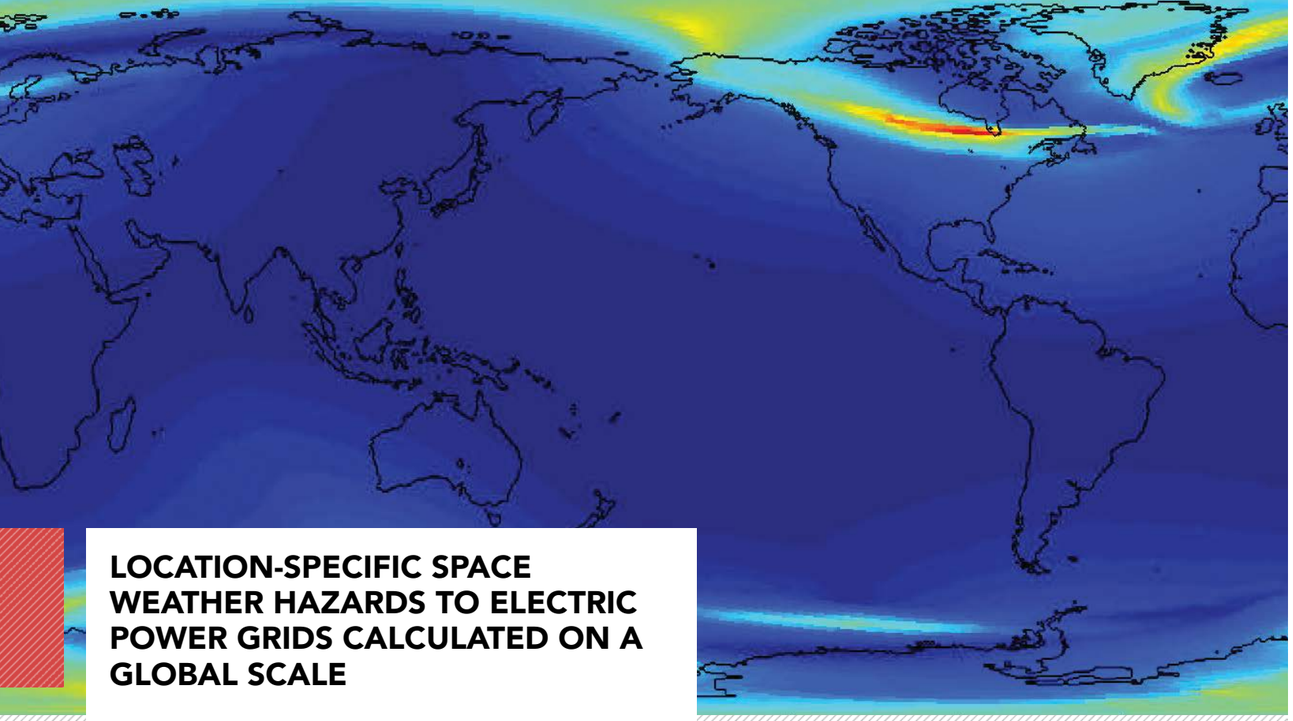

LOCATION-SPECIFIC SPACE WEATHER HAZARDS TO ELECTRIC POWER GRIDS CALCULATED ON A GLOBAL SCALE

Allocation: PRAC/2.3 Mnh
PI: Jamesina J. Simpson[1]
Collaborators: Bach Nguyen[1], Chigomezyo Ngwira[2]

[1]University of Utah
[2]Catholic University of America

FIGURE 2 (BACKGROUND): Electric field source amplitude for March 17, 2015, at 13:34 UT as calculated by BATS-R-US available through the Community Coordinated Modeling Center (CCMC) hosted by NASA Goddard Space Flight Center.

EXECUTIVE SUMMARY:

The historical record suggests that extreme space weather is likely to impact the Earth again in the future. However, modern electrotechnologies will be affected by space weather to a much larger degree than in the past. We are using a global Maxwell's equations model of the Earth-ionosphere waveguide to calculate location-specific space weather hazards to electric power grids. Specifically, we are calculating and analyzing electromagnetic field behavior during a recent geomagnetic storm in March 2015. Blue Waters permits us to account 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. Blue Waters also allows us to calculate and analyze ground-level electromagnetic fields values over time-spans of hours at microsecond time resolution (as required by our algorithm). 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. The largest documented geomagnetic storm in 1859 [1] caused telegraph operators communicating over 100 km wire lines to experience electric shocks, some nearly fatal [2]. 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 electrotechnologies. 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].

METHODS & RESULTS

The goal of the proposed work 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 the principal investigator over the past 12 years (e.g. [4, 5]). These models are based on the finite-difference time-domain (FDTD) method. FDTD is a time-domain and grid-based approach that permits us to account for such details as the Earth's complete topography, oceans, variable composition of the lithosphere, as well as the variable ionospheric composition and disturbances according to time, 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. Figure 1 illustrates a planar cut of the 3D FDTD grid as seen from constant radial coordinate. Figure 2 illustrates an example snapshot of the disturbed ionospheric electric fields during the March 2015 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. 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 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 of 2.4 km and higher (previously, our highest grid resolution was 40 km). Achieving these high resolutions has been challenging because, referring to 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 sources and their variation in both time and space than previously possible (at resolutions down to 1 km globally in space). 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. As such, the FDTD-calculated results may be instrumental for protecting individual power grids substations.

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

So far at higher spatial resolutions, our model assumes a relatively simple exponential ionospheric profile. The next Track-1 system would allow us to model a more realistic magnetized ionospheric plasma that accounts for the geomagnetic field and variable electron and/or ion content of the ionosphere.

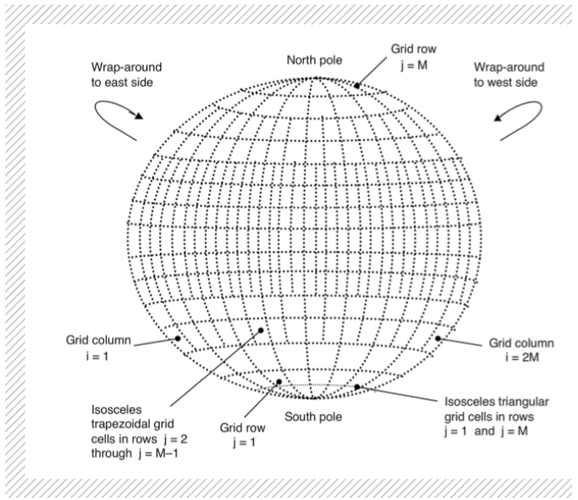

FIGURE 1: Layout of the 3D FDTD grid as seen from a constant radial coordinate.

SCALING THE CESM TO ULTRA-HIGH RESOLUTIONS FOR ANALYZING TROPICAL CYCLONE–CLIMATE FEEDBACKS

Allocation: Illinois/0.05 Mnh
PI: Ryan L. Sriver[1]
Collaborator: Hui Li[1]

[1]University of Illinois at Urbana–Champaign

EXECUTIVE SUMMARY:

The goal of this project is to optimize code performance and test the scalability of the Community Earth System Model (CESM) to ultra-high resolutions to examine connections between tropical cyclones and climate. We are conducting benchmarking simulations for different model coupling configurations (atmosphere only, ocean only, and fully coupled) and grid resolutions, including a 0.25° atmosphere and 0.1° ocean. This is the first time the ultra-high resolution version of the coupled model is being tested for studying tropical cyclones. The computational expense of the simulations is considerably outside the capabilities of traditional HPC resources, which is why we are using Blue Waters. Ultimately, this work will provide the critical next step in settling the decade-long debate about the importance of tropical-cyclone-induced feedbacks to Earth's climate.

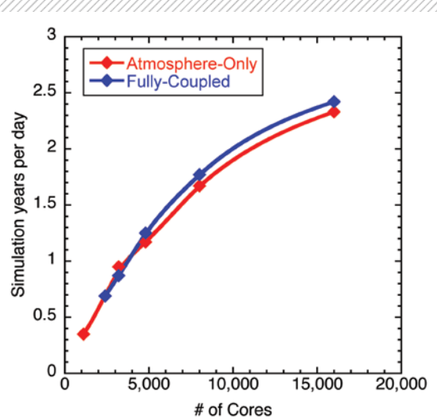

FIGURE 1: Preliminary scaling results of CESM for the atmosphere-only model (red curve) and the fully-coupled model (blue curve).

INTRODUCTION

Tropical cyclones are rare weather events, yet they consistently rank among the world's deadliest and costliest natural hazards. The frequency, intensity, and spatial distribution of tropical cyclones depend on large-scale environmental factors, such as tropical ocean temperatures and the structure of atmospheric winds. Changes in these factors due to climate variations can alter tropical cyclone activity on regional to global scales. However, tropical cyclones also actively contribute to the dynamics of Earth's climate system through complex ocean-atmosphere interactions that alter ocean temperature patterns and influence circulations within the atmosphere and ocean. These connections are poorly understood and largely missing from today's generation of Earth system models used to study climate variability, yet they may be fundamentally important to improving projections of future climate change. New earth system modeling strategies utilizing coupled ocean-atmosphere configurations, capable of resolving realistic tropical cyclone circulations and ocean-atmosphere interactions, are critical to developing a complete understanding of the relationship between tropical cyclones and climate.

METHODS & RESULTS

The Community Earth System Model (CESM) is a comprehensive global climate model that consists of atmosphere, land, ocean, and sea ice components that are connected via a central coupler that exchanges state information and energy fluxes between the components [1]. It represents the leading edge of community-wide efforts in global climate modeling and is considered a state-of-the-art earth system model. We are currently testing the scalability of CESM to ultra-high resolutions useful for analyzing tropical-cyclone-induced climate feedbacks (fig. 1). The simulations feature different model configurations, including ocean-only, atmosphere-only, and fully-coupled component sets. We are conducting these benchmarking tests at ultra-fine-scale grid resolutions (0.25° atmosphere coupled to 0.1° ocean) that are capable of simulating realistic tropical cyclones and the associated impacts on the upper ocean. These configurations are not typically used for tropical cyclone studies, due in part to the challenges of computational demand, coupling strategies, and the large volume of information produced by the high-resolution version of the coupled model. This work represents the cutting edge of current earth system modeling initiatives.

Preliminary results from high-resolution CESM simulations [2] have shown that the model generally simulates realistic tropical cyclones with the correct frequency, intensity, seasonality, and spatial distribution on global scales (fig. 2). However, the influence of these small-scale intense events on large-scale climate is largely unexplored, in particular for climate processes related to the ocean. We are using the benchmarking and scalability analysis described above as the basis for new model experiments examining how tropical cyclones affect upper-ocean energy budgets, circulation patterns, and important climate processes such as El Niño, ocean heat uptake and warming (due to enhanced ocean mixing), and large-scale ocean heat transport. In addition, we are running the model using several different grid resolutions to test sensitivities of the climate responses to changes in spatial resolution. This work will provide scientific insight into the effects of tropical cyclones on the upper ocean and will enable better understanding of the feasibility and computational demands of ultra-high resolution coupled modeling approaches to analyze connections between tropical cyclones and climate, as well as feedbacks that can influence our interpretations and projections of future climate change.

WHY BLUE WATERS?

The computational demand of the modeling efforts described above is considerably outside the capabilities of traditional HPC resources. Blue Waters provides unique resources capable of overcoming challenges due to large core-hour demands, high-frequency I/O, and post-processing of model output. Understanding the physical relationship between tropical cyclones and climate and assessing the societal and economic impacts under climate change represent a grand challenge to the Earth system modeling community. Blue Waters may provide the computational resources necessary to solve this problem.

PUBLICATIONS

Hui, L., and R. L. Sriver, Modeled Sensitivity of the Upper-Ocean Response to Tropical Cyclones in the Northwestern Pacific Using a Fully-coupled Climate Model with Varying Ocean Resolution. *Amer. Geophys. Union Fall Mtg.*, San Francisco, Calif., December 15–19, 2014.

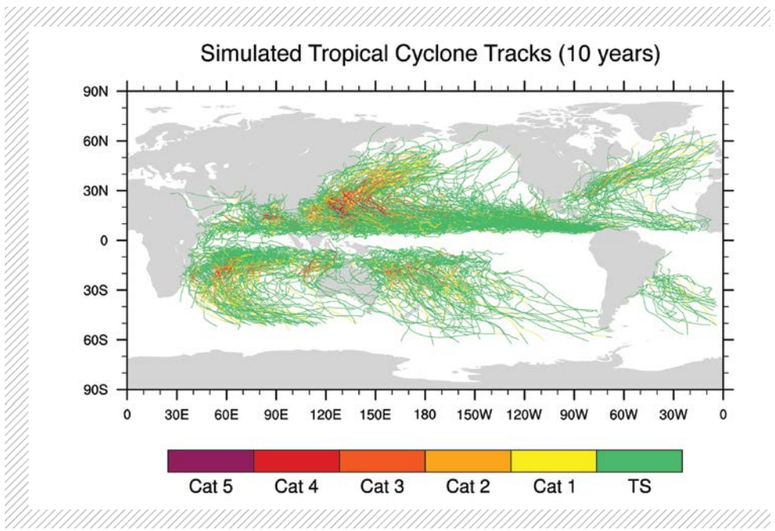

FIGURE 2: Simulated tropical cyclone tracks analyzed from ten years of CESM model output [2]. These tracks are the basis for new runs exploring the impacts of tropical cyclones on the upper ocean and the related climate impacts.

PETASCALE MODELING OF CONVECTIVE STORMS UNDER CLIMATE CHANGE AND VARIABILITY

Allocation: Blue Waters Professor/0.245 Mnh
PI: Robert J. Trapp[1]

[1]University of Illinois at Urbana–Champaign

EXECUTIVE SUMMARY:

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

We concluded thus far that a more intensely rotating tornadic storm and, by extension, a more intense tornado, is one potential effect of PGW. This is due largely to the PGW-modified thermodynamics. However other PGW modifications of the convective environment precluded storms from forming: the combined effects of increased convective inhibition and decreased forcing for convection led to a failure of convection initiation in many of the ensemble members. Additional simulations using other GCMs will help us identify the predominant effect.

INTRODUCTION

Climate change assessments persistently show uncertainty with regards to how human-enhanced greenhouse gas concentrations and the associated global radiative forcing might affect the frequency and intensity of local, high-impact thunderstorms. Part of the challenge is that such storms—and the attendant tornadoes, hail, damaging straight-line winds, lightning, and localized flooding—have spatial scales that fall well below the effective resolution of typical global and even regional climate models. Modeling approaches such as dynamical downscaling have addressed this resolution issue, but generally their applications thus far have been unconcerned with reproducing historical thunderstorms and tornadoes and therefore how these events might be projected in the future.

METHODS & RESULTS

The "pseudo-global warming" (PGW) 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 (GCM) output were used to constrain the Weather Research and Forecasting (WRF) model simulations of these events at high resolution. Comparison of an ensemble of these simulations with control simulations (CTRL) facilitated assessment of PGW effects.

We have concluded thus far that a more intensely rotating tornadic storm and, by extension, a more intense tornado, is one potential effect of PGW. This is due in part to the PGW-enhanced convective available potential energy, and in spite of the PGW-reduced environmental wind shear; neither of these PGW modifications were sufficient to significantly change the storm morphology. However, other PGW modifications precluded storms from forming: the combined effects of increased convective inhibition and decreased forcing for convection led to a failure of convection initiation in many of the ensemble members. In general, the PGW-modified thermodynamics had the largest overall impact on the realization of the tornadic storms. Additional simulations using other GCMs will help us generalize these results.

WHY BLUE WATERS?

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

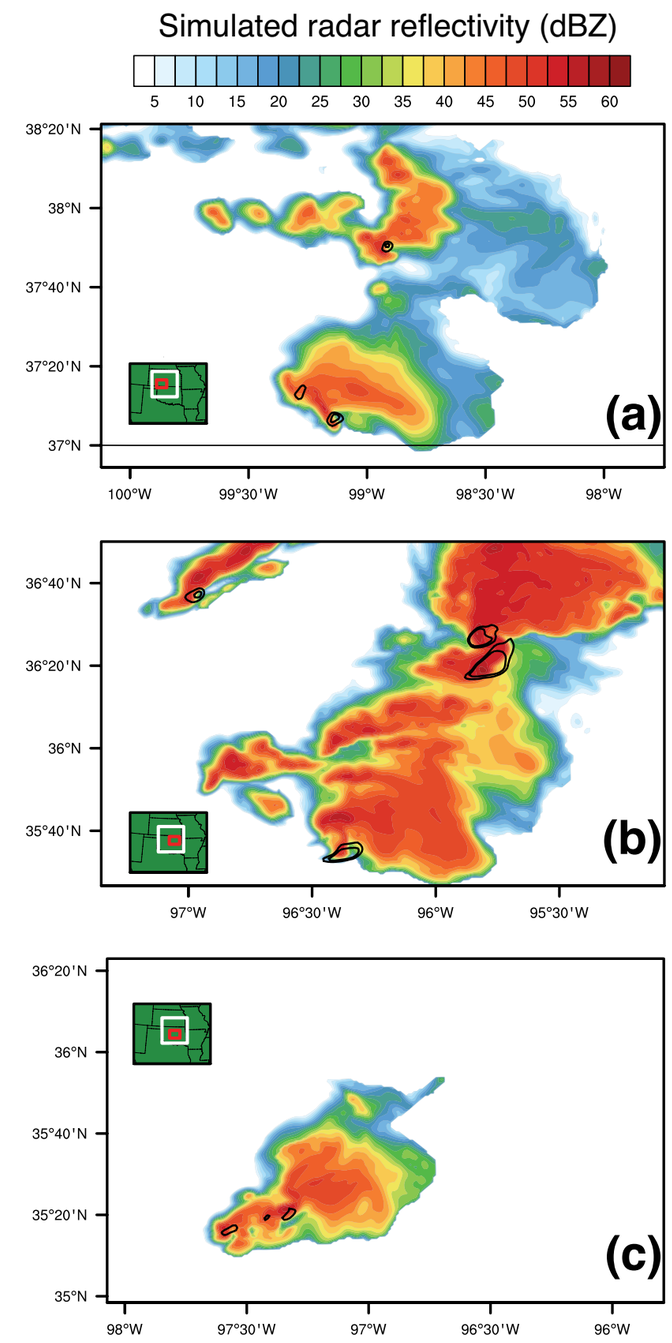

FIGURE 1: Simulated radar reflectivity (dBZ; color fill) for the control simulations of the: (a) EF-5 Greensburg, Kansas, tornadic storm of 4 May 2007, (b) EF-5 Norman, Oklahoma, tornadic storm of 10 May 2010, and (c) EF-4 Shawnee, Oklahoma, tornadic storm of 19 May 2013. Black contours are of updraft helicity (450 and 900 m^2s^2). The insets show the geographical areas described by nested computational domains d01 (black outline), d02 (white outline), and d03 (red outline).

APPLICABILITY OF GEOMETRIC OPTICS METHOD FOR CALCULATIONS OF SINGLE-SCATTERING PROPERTIES OF ATMOSPHERIC ICE CRYSTALS

Allocation: Illinois/0.75 Mnh
PI: Junshik Um[1]
Collaborator: Greg M. McFarquhar[1]

[1]University of Illinois at Urbana–Champaign

EXECUTIVE SUMMARY:

Ice clouds exhibit a cooling impact on Earth's climate by reflecting solar radiation and a warming impact through absorption of infrared radiation and thermal emission towards Earth's surface. Ice clouds consist of non-spherical ice crystals with various shapes and sizes. To determine the influence of ice clouds on Earth's radiation budget and hence climate, knowledge of how they scatter and absorb radiation, henceforth called their single-scattering properties, is required. A geometric optics method (GOM), an asymptotic solution that is appropriate when a particle is much larger than the wavelength of incident light, has been widely used to calculate the single-scattering properties of ice crystals. However, the exact size range of applicability of GOM has not been well established and should be determined through comparison against numerically exact methods.

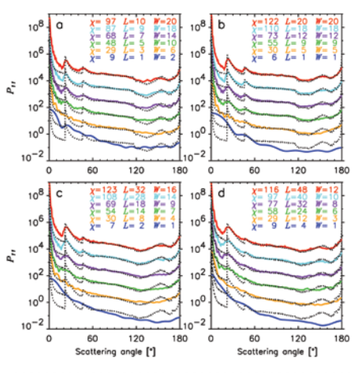

FIGURE 1: Scattering phase function (P^{11}) of hexagonal crystals with AR of (a) 0.5, (b) 1.0, (c) 2.0, and (d) 4.0 using ADDA (color solid lines) and GOM (black dot lines). Different colors indicate different crystal sizes with corresponding L, W, and χ embedded in each panel. In order to distinguish lines, 10^i is multiplied to each of the original values from small to large size, where i=0, 1, 2, 3, 4, and 5.

In this study, single-scattering properties of ice crystals were calculated using both GOM and a numerically exact method to determine the range of applicability of GOM.

INTRODUCTION

Ice clouds exhibit a cooling impact on Earth's climate by reflecting incoming solar radiation and a warming impact through absorption of thermal emission from the Earth and lower atmosphere and emission of infrared radiation at colder temperatures. The relative importance of these two effects depends on the optical thickness of ice clouds and has large impacts on the terrestrial climate system and climate feedbacks. Despite the importance of ice clouds and their large spatial coverage and temporal frequency, their representation in numerical models and assumptions used in remote sensing retrieval algorithms have large uncertainties caused mainly by the non-spherical shape of ice crystals and their wide variety of shapes and sizes.

A geometric optics method (GOM) has been exclusively used to calculate the scattering properties of non-spherical ice crystals. However, a GOM is an asymptotic solution and its range of applicability has not been well established and should be determined through comparison against numerically exact (though computationally expensive) methods, such as the discrete dipole approximation (DDA). In this study, the range of applicability of GOM was determined, and the impacts of crystal size and shape on the applicability of GOM were also quantified. Since the required computing time and memory for numerically exact methods rapidly increase with particle size, knowledge of the range of applicability of a GOM method can save computing time.

METHODS & RESULTS

In order to determine the applicability of GOM, the single-scattering properties of randomly oriented hexagonal ice crystals were calculated using the Amsterdam DDA (ADDA) [1] and conventional GOM [2] at a non-absorbing wavelength λ=0.55 μm. Observations of naturally occurring ice crystals were used to determine the range of morphological features of the crystals used in the calculations [3]. A width (W) of up to 20 μm and a length (L) of up to 48 μm of hexagonal ice crystals with aspect ratios (AR=L/W) of 0.5, 1.0, 2.0, and 4.0 were thus used in the simulations. These crystals correspond to a volume-equivalent-sphere size parameter (χ) of up to 123, where χ is defined as the ratio of the perimeter of a sphere that has the same volume of a non-spherical particle to the wavelength of incident light. Since there is no size effect on the scattering calculations using GOM at λ=0.55 μm, only one simulation using the GOM was performed for each AR.

Fig. 1 shows that the calculated scattering phase functions (P11) of hexagonal crystals using ADDA approached those calculated using GOM more closely as the size of the crystal increases for all four ARs. Better agreement between the P11 calculated using ADDA and those using conventional GOM is also shown for large hexagonal crystals with compact shapes (e.g., AR=1.0) compared to those with oblate (e.g., AR=0.5) and prolate (e.g., AR=4.0) shapes. The most notable differences in P11 are in the forward scattering region with the differences caused by errors in Babinet's principle that was used in the GOM for calculation of diffraction.

Fig. 2 shows that the differences between ADDA and GOM simulations for parameters that describe the directional scattering of radiation (the asymmetry parameter g) and the extinction efficiency (Qext) became smaller as the crystal size increased. The errors in the conventional GOM compared to ADDA were ~1.5% (7.0%) for g (Qext) of hexagonal crystals with χ=90 for all ARs, whereas they were ~1.2% (3.3%) for hexagonal crystals with χ=100. The GOM simulations (i.e. P11, g, and Qext) for hexagonal crystals with compact shapes (i.e. AR=1.0) were closer to the ADDA simulations for smaller sizes compared to the simulations for oblate and prolate hexagonal crystals. This indicates that the lower size limit of applicability of GOM depends on the crystal shape (i.e. habit) and on its morphological features (e.g., aspect ratio).

Since methods that directly solve Maxwell's equations (e.g., DDA) are computationally expensive for large particles, while GOM is a fast and flexible method to calculate scattering properties of non-spherical particles, inevitably GOM will continue to be used for scattering calculations. This work provides guidance on where errors in GOM are small enough that its use is appropriate. The wide variety of shapes, sizes, and habits of ice crystals mean the development of more databases on scattering properties will continue to be needed.

WHY BLUE WATERS?

A conventional GOM is appropriate when a particle is much larger than the wavelength of incident light. Thus, numerically exact methods (e.g., DDA and T-matrix) are required for the calculations of scattering properties of particles with small size parameters, and errors of GOM in this region should be quantified through comparison with numerically exact methods. The required computing time and memory for numerically exact methods rapidly increase with particle size and large computing resources, like Blue Waters, were thus required to complete this study.

PUBLICATIONS

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, Formation of atmospheric halos by hexagonal ice crystals. 14th Conference on Atmospheric Radiation, Boston, Mass., July 7–11, 2014.

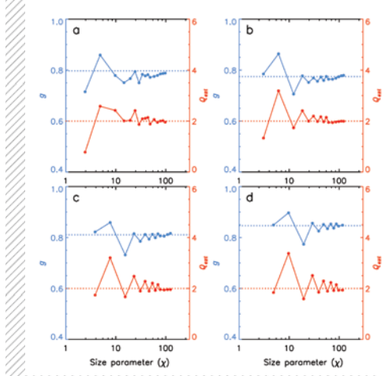

FIGURE 2: The asymmetry parameter (g, blue) and extinction efficiency (Q_{ext}, red) of hexagonal crystals with AR of (a) 0.5, (b) 1.0, (c) 2.0, and (d) 4.0 using ADDA (solid lines with circles) and GOM (dot lines).

COMPUTATIONAL MINERAL PHYSICS

Allocation: GLCPC/0.35 Mnh
PI: Renata M. Wentzcovitch[1]
Collaborators: Tao Sun[2], Dong-Bo Zhang[3], Han Hsu[4], Mehmet Topsakal[1], Gaurav Shukla[1], Kanchan Sarkar[1], Michel Marcondes[1]

[1]University of Minnesota, Twin Cities
[2]University of Chinese Academy of Sciences
[3]Computational Science Research Center, Chinese Academy of Engineering Physics, Beijing
[4]National Central University, Taiwan

EXECUTIVE SUMMARY:

Mineral physics is one of three branches of geophysics (the others being geodynamics and seismology). Geophysics advances by close cooperation between these fields. Thus, mineral physics focuses on mineral properties that are needed to interpret seismic data or are essential for geodynamics simulations. To be useful, mineral properties must be investigated across a wide range of pressure, temperature, and chemical composition. Chemical composition is complex, encompassing at least five major oxide components and tens of solid phases. Today, first-principles computations address these challenges. Minerals' properties are fundamental to elucidate a planet's state, and atomistic studies of these complex materials are fundamental to understanding their properties. We will use Blue Waters to perform first-principles calculations of unprecedented magnitude and scope in mineral physics.

INTRODUCTION

The goal of high pressure mineral physics is to increase our knowledge of the materials that make up the Earth and other planets [1]. For many problems, first-principles theory [e.g., 2–6] is the only practical method of investigation available. First-principles methods have been addressing materials problems important to understanding the present state and evolution of the Earth and other planets [4,7–14]. However, computation of full thermal elastic anomalies caused by spin crossovers at high temperatures and pressures has not been realized.

Today's first-principles calculations need databases of equilibrium thermal elastic properties of minerals and their assemblages to model Earth and other planets [1]. The challenge to computations is to perform these highly demanding calculations efficiently in a large number of phases with variable compositions. For example, the chemical composition of the terrestrial mantle includes at least five major oxide components (MgO, SiO_2, Al_2O_3, FeO, CaO) and tens of solid phases (fig. 1).

Our work is motivated and informed by central questions in Earth and planetary sciences. It focuses on building highly realistic models that account for all physical and chemical variations in the system (e.g., chemical composition, pressure, temperature, etc.).

We perform calculations on tens of mantle phases using primarily quasi-harmonic (QHA) computations. However, molecular dynamics (MD) or hybrid QHA/MD calculations are performed in the case of anharmonic phases or at temperatures near melting. As far as elasticity is concerned, a rigorous treatment of solid solutions is unnecessary for mineral aggregates. These results will address questions such as:

- What is the seismic signature in the deep mantle produced by the sinking lithospheric plates made of mid-ocean ridge basalt (MORB), harzburgite, and peridotite? Conversely, to what extent can mantle heterogeneities be attributed to seismic contrasts between these assemblages and "normal" pyrolite mantle [19]?
- What are the temperatures, chemistries, and mineral assemblages of the two large, low-shear velocity provinces in the deep mantle [21]? Did they result of accumulation of MORB enriched material?
- What seismic signature do spin changes produce in iron in lower mantle minerals [e.g., 10,16–17]?
- What are the phase boundary and velocity contrast produced by the post-perovskite transition [8,11,20] in MORB, harzburgite, peridotite, and pyrolite? This information must be available before one can interpret seismic observations of the D" region, the ~300 km above the CMB.

METHODS & RESULTS

We used standard and extended/hybrid density functional theory (DFT)-based methods as implemented in the Quantum ESPRESSO software [e.g., 2–6]. These DFT calculations combined QHA and MD to address thermal elastic properties.

Thermodynamics methods

Thermodynamic states were investigated using QHA [22] and MD. QHA and MD are complementary approaches, the former from low to intermediate temperatures ($T < \sim 2/3\ T_{melting}$) and the latter from intermediate to $T_{melting}$ and beyond. Both methods are CPU intensive because they involve statistical sampling of phase space. In MD this translates into large supercells and long runs. By contrast, the QHA used in conjunction with first principles vibrational density of states [3,5] is much more efficient for calculations corresponding to very large numbers of atoms.

We routinely performed calculations corresponding to 10^4 atoms [4,7,9,15]. However, these were high throughput calculations requiring typically 10^2 to 10^3 medium-size calculations with $\sim 10^2$ atoms. In particular, thermoelastic constants, $c^T_{ijkl}(P,T)$, a central topic in our project, was one of the most CPU intensive calculations we performed. These jobs could be executed concurrently in single or distributed compute nodes. Ref. [6] summarizes a few of these calculations [4,7,15,18]. They were

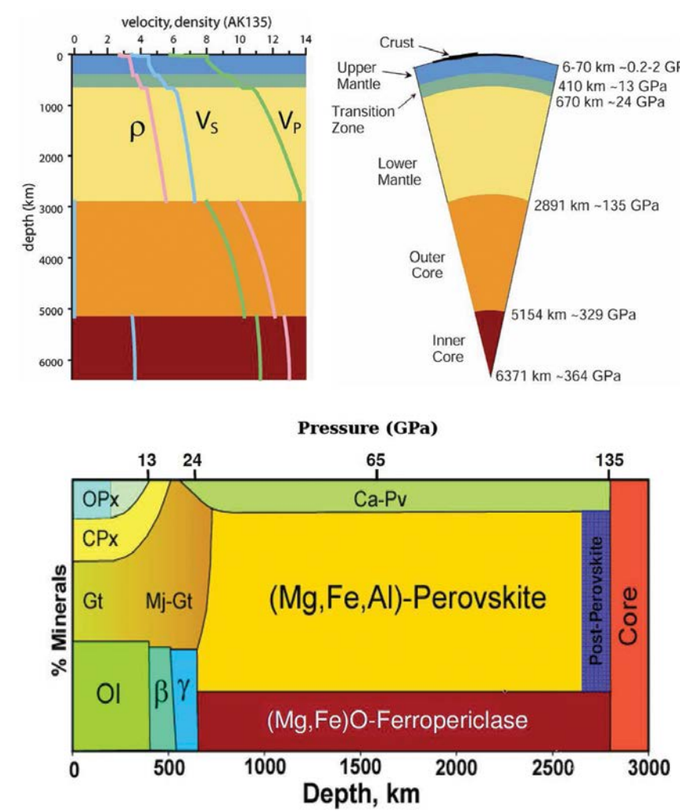

FIGURE 1: Earth's seismic discontinuities (top) [59] and mantle mineralogy (bottom; e.g., [60]).

carried out by submitting bag-of-task (BoT) [14] type jobs to single or distributed systems using the VLab service-oriented cyberinfrastructure [12–14].

On Blue Waters, these calculations were performed using BoTs or MPI directives, "image" parallelization (1st level of parallelization). For a typical system with 20 atoms (e.g., $MgSiO_3$ perovskite or post-perovskite [7,15]), the best performance was achieved with 128 cores per job. Thus, a full thermodynamic calculation could be performed in less than one minute, or a calculation of thermoelastic constants could be performed in a couple of minutes.

WHY BLUE WATERS?

The major reason for using Blue Waters is the number of cores available. We developed wrappers around several modules of the Quantum ESPRESSO software that allow us to use the system for high-throughput calculations. As explained above, phase space sampling for thermodynamics calculations is intrinsically a high-throughput problem.

A future Track-1 system will allow us to expand the dimension of phase space currently sampled. Minerals are solid solutions and a full investigation of phase equilibrium in these solids involves sampling of atomic configurations. We have just started addressing this issue. We are exploring techniques to do maximally efficient sampling of atomic configurations. Nevertheless, multiple configurations (~10^1–10^2) need to be explored in each of up to ~10^4 calculations.

We have essentially investigated all major end-member phases and simple versions of the main mineral solid solutions of the mantle using only a few few atomic configurations. While elastic properties of crystalline aggregates are well described by using few atomic configurations, phase diagrams containing binary (or multi-phase) loops need accurate free energies, which can only be achieved through careful sampling of atomic configurations. An example of this problem can be seen in fig. 1, where OPx and CPx phases "dissolve" into Gt (garnet) within a broad pressure range (i.e. 5 GPa < P < 18 GPa). A detailed description of this phenomenon requires a considerably larger number of (smarter) calculations.

PUBLICATIONS

Lacerda, M. M., and R. M. Wentzcovitch, Hybrid ab-initio/experimental high temperature equations of state: the NaCl pressure scale. *J. Appl. Phys.*, 117 (2015), 215902, doi:10.1063/1.4921904.

Shukla, G., et al., Thermoelasticity of (Mg,Fe) SiO3 perovskite. *Geophys. Res. Lett.*, 42:6 (2015), pp. 65–77, doi:10.1016/S0031-9201(96)03213-X.

Wu, R., et al., The Atomic and Electronic Structure of Black Phosphorus. *J. Vac. Sci. Technol. A*, 33 (2015), 060604, doi:10.1116/1.4926753.

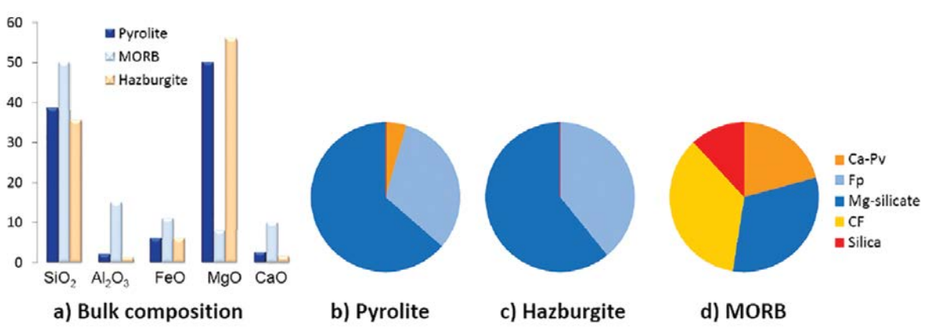

FIGURE 2: (a) Chemical compositions of mantle rocks [85-87]. Mineralogy (in mol%) of (b) pyrolite [88], (c) harzburgite, and (d) basalt [89] in the lowermost mantle. Mg-Silicate: $MgSiO_3$ perovskite and/or post-perovskite, Fp: ferropericlase, Ca-Pv: CaSiO3 perovskite, and CF: calcium-ferrite-type phase.

USING PETASCALE COMPUTING CAPABILITIES TO ADDRESS CLIMATE CHANGE UNCERTAINTIES

Allocation: NSF PRAC/5.02 Mnh
PI: Donald J. Wuebbles[1]
Co-PIs: Zachary Zobel[1], Warren Washington[2], Thomas Bettge[2], Julio Bacmeister[2], Kevin Reed[2], Xin-Zhong Liang[3], Chao Sun[3]

[1]University of Illinois at Urbana–Champaign
[2]National Center for Atmospheric Research
[3]University of Maryland, College Park

EXECUTIVE SUMMARY:

This collaborative research between the University of Illinois, the National Center for Atmospheric Research (NCAR), and the University of Maryland uses Blue Waters to address key uncertainties in numerically modeling the Earth's climate system and accuracy in analyses of past and projected future changes in climate at a level that would be impossible without petascale computing. Our studies used the latest, most advanced versions of the Community Earth System Model (CESM) and two versions of NCAR's Weather Research and Forecasting Model (WRF and CWRF) for high-resolution regional climate analyses. These model runs put us on the pathway for major international leadership in high-resolution climate modeling studies.

INTRODUCTION

This collaborative research used Blue Waters to address key uncertainties in numerically modeling the Earth's climate system and accuracy in analyses of past and projected future changes in climate. Our studies used the latest versions of the Community Earth System Model (CESM), the Weather Research and Forecasting model (WRF), and the Climate-WRF (CWRF).

We explored the effect of external forcings (e.g., concentrations of greenhouse gases) on globally averaged temperature, focusing on uncertainties associated with the representation of processes and interactions between clouds, aerosols, and radiative transfer in the models and how these uncertainties influence the climate sensitivity.

Our second objective aimed to evaluate CESM with different model dynamical cores and the effects of much higher horizontal (10–30 km) and vertical resolution. There is considerable evidence that increased resolution leads to better simulations of both transient small-scale activity (e.g., eddies and waves) and large-scale features of the mean climate.

We expect that results from our studies will be an integral part of the scientific analysis of climate change for the major international climate assessments (e.g., the follow-on to the Assessment Report 5, AR5, 2014) of the Intergovernmental Panel on Climate Change (IPCC), which will likely influence U.S. and international policies regarding the effects of human activities on the Earth's climate, and provide supplementary data for use in various climate impact assessments at regional to local scales.

METHODS AND RESULTS

CESM

After tuning the physics and testing model settings of (the atmospheric portion of) CESM, we began a series of simulations that will support studies for the next Climate Model Intercomparison Project 6 (CMIP6), which is important to the next IPCC assessment of climate change.

Recent results from simulations that ran 1979 to the present suggested that the global number of tropical storms and hurricanes per year will decrease in a warming climate. However, the maximum intensity seemed to increase, meaning more major hurricanes (category 4 and 5). In the North Atlantic basin the RCP8.5 scenario showed a decrease in the number of tropical cyclones and hurricanes. The selection of a dynamical core can have a significant impact on tropical cyclone intensity and frequency even in the presence of similar climatology and large-scale environments [1]. For example, CAM5 with the spectral element (SE) core produced stronger cyclones, and therefore more hurricanes and major hurricanes per year, than CAM5 with the finite volume (FV) core (figs. 1–2). The exact causes for these differences is an area of continued work.

Compared to the out-of-box model, we improved the performance of CAM5 from 1.3 years/day to 2.5 years/day. Pat Worley (ORNL) and Ryan Mokos (NCSA) helped remove bottlenecks in MPI calls, while John Truesdale

(NCAR) rebalanced the PE layout among the model components. These and future modifications could improve model performance on other machines as well.

WRF at 12 km Resolution

WRF model was used to downscale the GFDL and the Hadley Centre model results, both of which provide ~1° (~100 km) resolution, to 12 km horizontal resolution in one historical simulation and two climate projections using different greenhouse gas scenarios. High-resolution model runs can resolve mesoscale processes and provide better detail of extreme weather, such as heat waves and convective thunderstorms. For instance, at 12 km resolution, long convective storms lead to the majority of flash floods and severe weather. A lower resolution model would have trouble correctly resolving these types of storms. When we compared the GFDL and WRF to the North American Regional Reanalysis (NARR), a gridded dataset derived from observed precipitation, we saw obvious improvements in resolution especially along coasts and in mountainous parts of the country when using WRF (fig. 3).

CWRF Uncertainty Analyses

We also used the regional CWRF [2] to examine uncertainties in the treatment of cloud, aerosol, and radiative transfer processes. We defined a subset of computationally feasible size from the Cloud-Aerosol-Radiation (CAR) [3] ensemble system that can capture the observed climate characteristics, especially those relevant to climate variability and extreme events. We ran an ensemble of 24 CWRF physics configurations, all driven by ERA-Interim reanalysis data. The results showed that mean surface temperatures and summer precipitation over both land and ocean differed dramatically among CWRF configurations despite the same lateral boundary conditions from ERI. Some configurations had overall small biases while others produced much larger errors.

We plan to complete the retrospective integration during 1979–2014 for all 24 configurations to ensure statistically robust results, especially for diagnosing interannual variations and extreme events. Subsequent diagnostic analysis of this initial ensemble output will allow us to form a better strategy to choose a more representative (and likely larger) ensemble to conduct CWRF climate change projections driven by CCSM4.

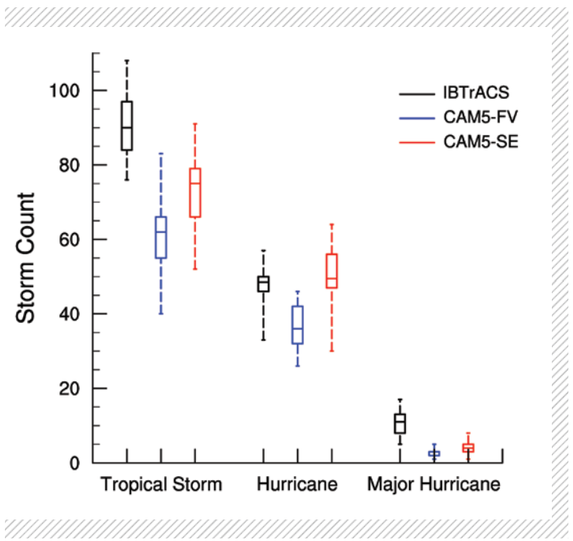

FIGURE 1: Distribution of the global number of tropical cyclones that reach tropical storm, hurricane, and major hurricane (cat. 4-5) strength per year for the CAM5-FV and CAM5-SE simulations 1980-1999. The boxes show the 25th, 50th, and 75th percentiles; values outside the box represent extrema (i.e. years with the most and least number of TCs). IBTrACS observations are included as a reference.

WHY BLUE WATERS?

The major objectives of the project addressed long-recognized issues associated with model resolution and understanding climate sensitivity. The unique super-ensemble approach used in the climate sensitivity studies and the advanced dynamical model cores used in the high-resolution studies both required petascale capabilities.

CAM5-SE required five wall-clock days to run a decade of model time on Blue Waters using thousands of cores. Several of these decade-long tests were required to determine an optimal set of model parameters, and then we used those settings to simulate a century or longer. To complete all of these model runs at high resolution and in a timely manner required petascale computing resources. Similarly, the high-resolution studies using WRF and CWRF also required a system like Blue Waters.

The next Track-1 system will allow us to run century-long ensembles of high-resolution climate models to quantify and reduce uncertainty. The climate modeling community plans to further increase the ocean/sea ice resolution from 1° to 0.1° and atmospheric resolution from 0.25° to 0.125° and beyond. This would improve representation of ocean eddies, cloud processes, and other important details of the Earth system.

PUBLICATIONS

Reed, K. A., et al., Impact of the dynamical core on the direct simulation of tropical cyclones in a high –resolution global model. *Geophys. Res. Lett.*, 42:9 (2015), pp. 3603–3608, doi:10.1002/2015GL063974.

Wuebbles, D., W. Washington, G. Meehl, and T. Bettge, High-Resolution Earth System Global and Regional Modeling for Climate Assessment and Policymaking Require Advanced Computing Infrastructure. *NRC Meeting on Computations for Science and Technology*, Dallas, Texas, January 21–23, 2015.

Wuebbles, D. J., et al., High-resolution climate simulations using Blue Waters. (*Blue Waters Annual Report*, University of Illinois, Urbana, Ill., 2014).

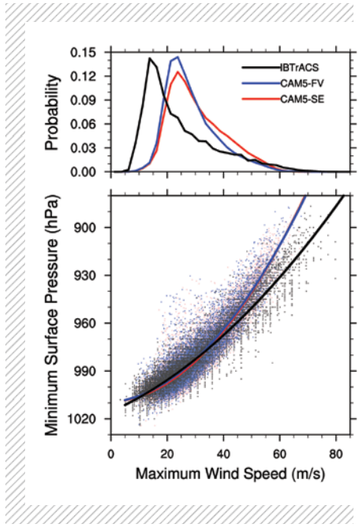

FIGURE 2: Intensity distribution shown as a (top) probability density function of the maximum wind speed and (bottom) minimum surface pressure vs. maximum wind speed relationship with quadratic least squares fit (solid lines) for the CAM5-FV and CAM5-SE simulations 1980-1999. IBTrACS observations are included as a reference.

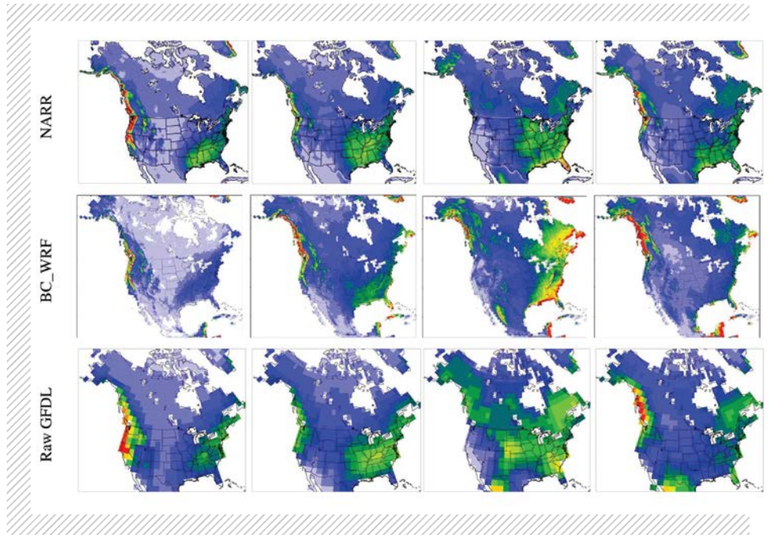

FIGURE 3: WRF simulated, GFDL-ESM2G simulated, and NARR gridded 10-year mean precipitation in winter, spring, summer, and fall (left to right; mm/day).

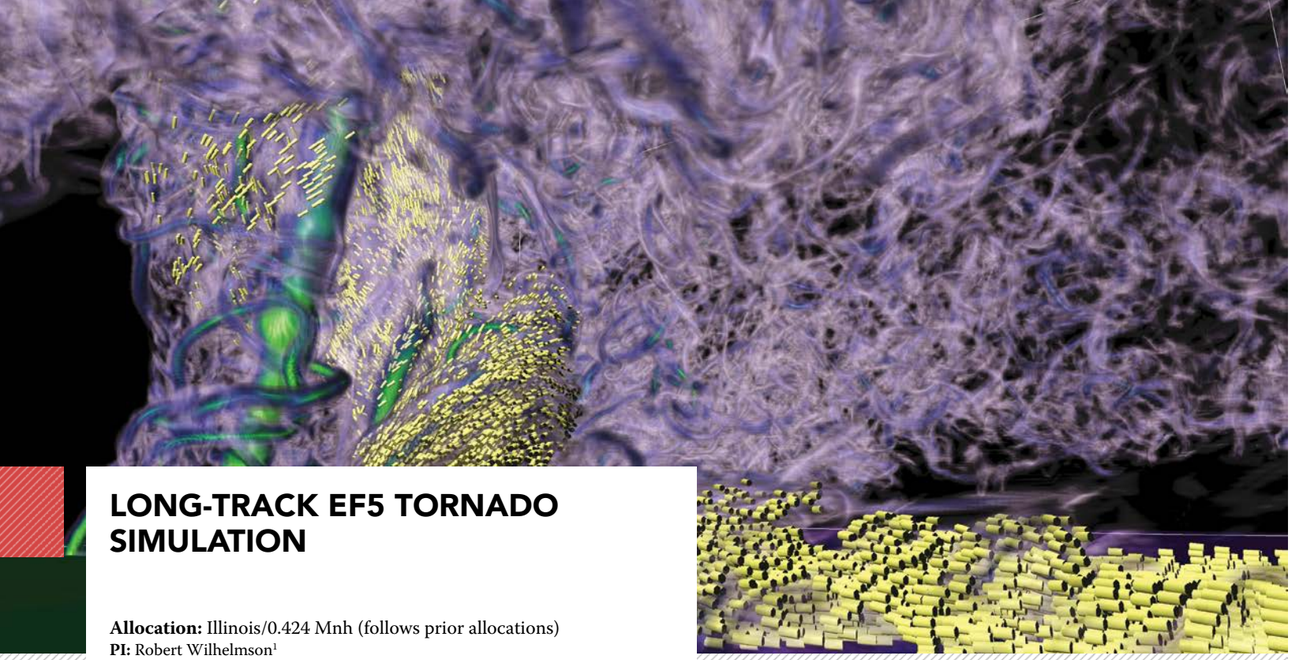

LONG-TRACK EF5 TORNADO SIMULATION

Allocation: Illinois/0.424 Mnh (follows prior allocations)
PI: Robert Wilhelmson[1]
Co-PI: Leigh Orf[2]
Collaborators: Bruce Lee[3], Cathy Finley[3]

[1]University of Illinois at Urbana–Champaign
[2]Central Michigan University
[3]Windlogics

EXECUTIVE SUMMARY:

Tornadoes are among nature's most destructive forces. The strongest, longest-lived tornadoes are spawned from supercell thunderstorms; however, only a small fraction of supercells produce such tornadoes. The infrequency of these devastating storms in nature is mirrored in simulations of supercells, making the numerical study of such storms a challenge. Utilizing state-of-the-art modeling and visualization software, a breakthrough simulation of a real long-track EF5 tornado embedded within a supercell thunderstorm was conducted and visualized with unprecedented fidelity [1]. The actual tornado which we based the simulation on produced winds exceeding 300 miles per hour, was on the ground for nearly two hours and covered a distance of 80 miles.

INTRODUCTION

We seek to understand the inner workings of supercell thunderstorms that produce long-track, devastating tornadoes, and what differentiates such thunderstorms from those that produce weak tornadoes or no tornadoes. Such understanding is necessary in order to better forecast the behavior of supercell thunderstorms as they occur in nature. If specific features identified in simulated storms are found to exist in nature in only the most devastating supercells, remote sensing technologies such as radar and satellite may be able to detect such features. This information could then guide forecasters in providing accurate, targeted warnings while also reducing the false alarm rate, increasing safety and reducing costs. Ultimately, the goal of this research is to save lives by giving anyone who is in the path of one of these storms enough time to get out of harm's way.

METHODS & RESULTS

A state-of-the-art cloud model, CM1 [2], simulated the supercell that produced the long-track EF5 tornado currently being analyzed. The model was first modified to efficiently manage the massive amount of data produced in the simulation (on the order of 100 TB). The model was then initialized with environmental conditions very similar to those in which a supercell producing a long-track EF5 tornado occurred on May 24, 2011, near El Reno, Oklahoma [3]. Utilizing an updraft nudging technique [4] to trigger the thunderstorm and a dual-moment microphysics scheme [5], the simulated cloud grew into a powerful supercell thunderstorm. Eighty minutes into the simulation, the storm's low-level updraft strengthened in concert with the consolidation of several small vertical vortices. Minutes later, a condensation funnel was seen descending from the storm's wall cloud, and the tornado was born. Over the next thirty minutes, the tornado widened, strengthened, and became enveloped in rain. The tornado exhibited a two-celled structure, characterized by a downdraft in the tornado core straddled by two violently rotating intertwined vortices. A feature called a streamwise vorticity current (i.e. a horizontal tube of rotating air) that was tilted into the storm's powerful updraft was identified for the first time and is a primary focus of current analysis. The tornado's demise occurred very quickly and was associated with heavy rain and a strong downdraft that enshrouded the weakening tornado.

WHY BLUE WATERS?

Blue Waters has both the computational and data management infrastructure required to conduct ultra-high-resolution simulations that produce very large amounts of data. Due to the infrequency of supercells that produce long-track EF5 tornadoes in both nature and in simulations, many high-resolution simulations initialized with different environmental conditions were conducted until the storm of interest was simulated. With a slower, more modest supercomputer, it is highly likely the storm of interest never would have been simulated.

The highly technical nature of this work requires top-notch staff to assist in both keeping Blue Waters running smoothly and also in helping to produce code that runs efficiently on Blue Waters. Analysis and visualization software installed on Blue Waters is regularly maintained and upgraded, and project staff has been very responsive when we encountered issues.

It is anticipated that simulations occurring in different environmental conditions, and run at higher resolution, will be conducted on Blue Waters and future Track-1 systems. Only by simulating dozens of storms, some that produce long-track devastating tornadoes and some that do not, will meteorologists be able to better forecast such storms and issue more accurate warnings to the public.

PUBLICATIONS

Orf, L., R. Wilhelmson, and L. Wicker, A Numerical Simulation of a Long-Track EF5 Tornado Embedded Within a Supercell. 94th Am. Meteorol. Soc. Ann. Mtg., February 2–6, 2014.

FIGURE 1 (BACKGROUND): Vorticity (i.e. spin) magnitude (volume-rendered field), surface temperature perturbation (pseudo-colored field), and unsteady trajectories originating along a boundary in the storm's forward flank. The tornado is encircled by horizontally oriented vortices that originate in the storm's rear flank and ascend along the tornado's periphery. The structure indicated by the collection of trajectories is dubbed the streamwise vorticity current, a horizontally oriented tube of rotating air that is tilted vertically into the storm's rotating updraft.

PHYSICS & ENGINEERING

MATERIALS

QUANTUM

FLUIDS

NANOTECHNOLOGY

68 *Ionic Conductivity, Structural Deformation, and Programmable Anisotropy of DNA Origami in an Electric Field*

70 *Large-Scale Simulations of Biological–Synthetic Interface Materials for Sensing Applications*

72 *Comparing CAF and MPI-3 and Simulating Molecular Cloud Turbulence with Two-Fluid MHD*

74 *Reducing Jet Aircraft Noise by Harnessing the Heterogeneous XK Nodes on Blue Waters*

76 *Hydrogen under Extreme Conditions*

78 *Strongly Correlated Systems through Computation: From Bad Metals to Perfect Insulators*

80 *Next-Generation Ab Initio Symmetry-Adapted No-Core Shell Model and Its Impact on Nucleosynthesis*

82 *Direct Simulation of Dispersed Liquid Droplets in Isotropic Turbulence*

84 *Simulating Thermal Transport in Nanostructures from the Ballistic to the Diffusive Regime*

86 *High-Order Methods for Turbulent Transport in Engineering and Geosciences*

88 *Many-GPU Simulations for Soft Matter Design*

90 *New Techniques for Pixel-Level Fidelity Assessment in Interferometric Images*

92 *Atomistic Modeling of Future Nanoscale Electronic Devices with NEMO5*

94 *Sparse Matrix Factorization in Solid Mechanics and Geophysics on CPUs and GPUs*

96 *Shock-Induced Turbulent Mixing*

98 *High-Performance Computing of Hypersonic, Shock–Shock Interactions using Kinetic, Particle Approaches*

100 *Particle-in-Cell Simulations of Kinetic Effects in HED Plasmas*

102 *Mechanics of Random and Fractal Media*

104 *High-Resolution 3D Simulations of Core-Collapse Supernovae*

106 *Evolution of the Small Galaxy Population from High Redshift to the Present*

108 *Theoretical Spectroscopy for Transparent Conducting Oxides*

110 *Complexities of High-Reynolds-Number Turbulence*

112 *Dynamics of Argon Bubbles in Steel Continuous Casting with a Magnetic Field*

115 *Lattice QCD on Blue Waters*

118 *Modeling Type Ia Supernova Progenitors*

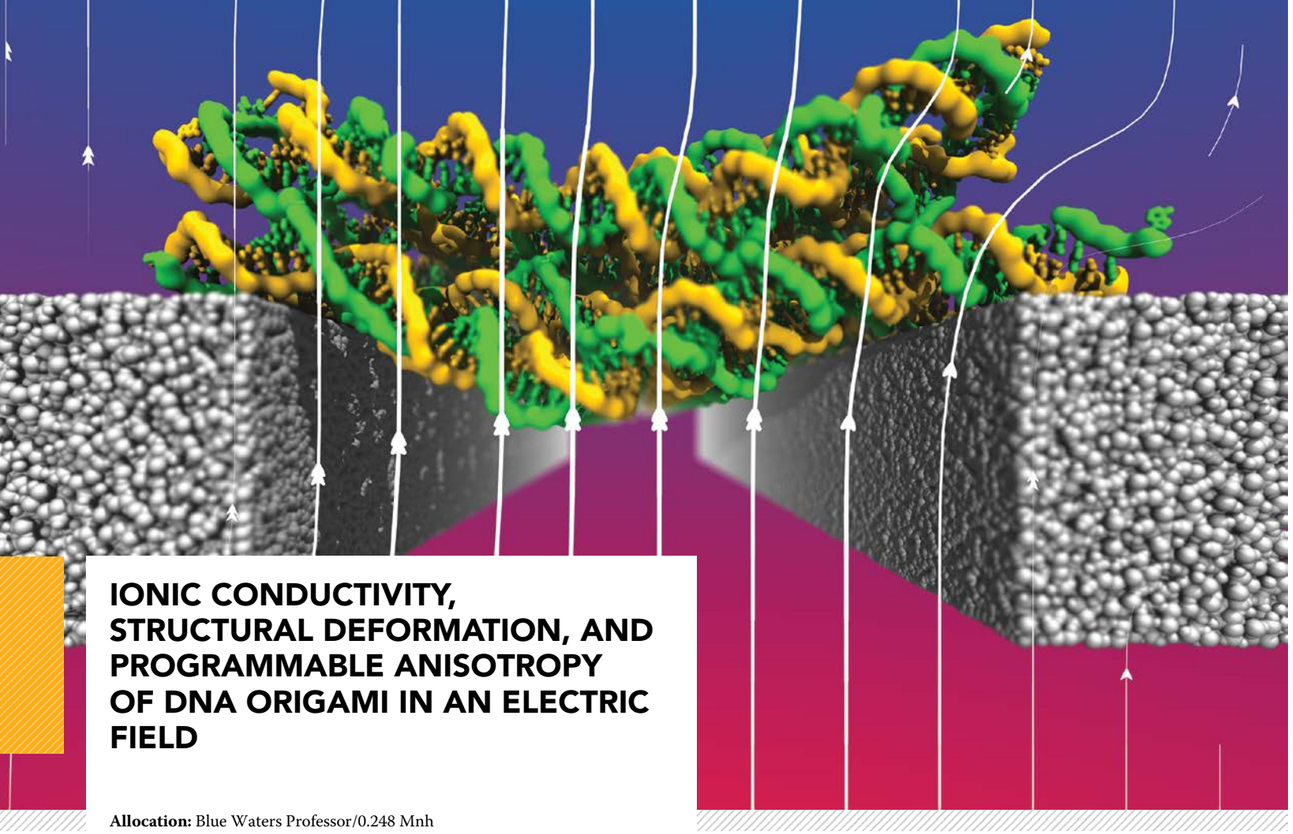

IONIC CONDUCTIVITY, STRUCTURAL DEFORMATION, AND PROGRAMMABLE ANISOTROPY OF DNA ORIGAMI IN AN ELECTRIC FIELD

Allocation: Blue Waters Professor/0.248 Mnh
PI: Aleksei Aksimentiev[1]
Collaborators: Chen-Yu Li[1], Elisa A. Hemmig[2], Jinglin Kong[2], Jejoong Yoo[1], Silvia Hernández-Ainsa[2], Ulrich F. Keyser[2]

[1]University of Illinois at Urbana–Champaign
[2]University of Cambridge, U.K.

EXECUTIVE SUMMARY:

DNA origami is an experimental technique that allows folding of a long DNA molecule into an arbitrary three-dimensional shape with sub-nanometer precision. In comparison to conventional nanomanufacturing, the DNA origami method has relatively low cost, is easy to use, and has an extremely broad range of possible applications. One such application is nanopore detection, whereby measuring the ionic current flowing through a nanopore can determine the presence and/or identity of biochemical analytes.

Using Blue Waters, our group has carried out the first all-atom molecular dynamics (MD) simulations of DNA origami nanopores. We have determined how the ionic conductivity of a DNA origami plate depends on its internal structure and environmental conditions. Single molecule experiments performed in our collaborator's laboratory (Cambridge U., U.K.) confirmed our predictions. Our work demonstrates the predictive power of the MD method in characterizing DNA origami objects and the feasibility of programming the electrical properties of nanoscale self-assembled objects.

INTRODUCTION

Nanopores have emerged as versatile tools for single-molecule manipulation and analysis. In a typical nanopore experiment, an analyte (e.g., DNA or protein) is driven through a nanometer-sized pore in an insulating membrane by applying an external electric field. The presence and, in some cases, the chemical structure of the DNA or proteins can be determined by measuring the ionic current change produced by the analyte blocking the nanopore. Solid-state inorganic materials are commonly used to create nanopore membranes. However, the dimensions and surface chemistry of such solid-state nanopores are difficult to control experimentally. A hybrid nanopore system consisting of a solid-state nanopore and a DNA origami plate can overcome the above shortcomings; however, the electrical conductivity of DNA origami objects has not been systematically characterized until now. The goal of this study was to elucidate the mechanism of ionic conductance through DNA origami objects and the means to control the conductance of such objects.

METHODS & RESULTS

To predict the structural and electrical properties of DNA origami materials, we performed molecular dynamics (MD) simulations of DNA origami plates, varying the number of DNA layers in an origami lattice, the lattice type, the nucleotide content, the orientation of the object with respect to the applied electric field, and the ionic conditions [1]. Tens of different DNA origami systems were built using a protocol our group developed [2]. The systems were simulated using the all-atom MD method and the NAMD2 package, which was optimized for Blue Waters.

Our MD simulations revealed how the current of ions through the DNA origami plate—the leakage current—depends on the internal design of the plate. First, the leakage current through the origami plate monotonically decreases as the number of DNA layers in the origami plate increases. The ionic conductivity of a DNA origami plate also depends on the lattice type of the DNA helices in the DNA origami structure (square, honeycomb, or hexagon). In addition, the plate's nucleotide content can have a nontrivial influence on its ionic conductivity. For example, a DNA origami plate made entirely of adenine (A) and thymine (T) nucleotides is about two times more conductive than a plate made entirely of guanidine (G) and cytosine (C) nucleotides.

The ionic environment (i.e. the magnitude and direction of the applied electric field) can considerably affect the ionic conductivity. Our simulations predicted that the ionic conductivity of DNA origami objects decreases as the solution concentration of Mg^{2+} increases. The simulations also revealed the mechanism of the dependence: higher amounts of Mg^{2+} lower the strength of DNA–DNA repulsion, making the structure of DNA origami more compact, reducing the ionic current through it. Experimental measurements of the ionic current through nanocapillaries blocked by different DNA origami plates and the Förster resonance energy transfer measurements directly confirmed the results of our MD simulations. The DNA origami objects were also found to have anisotropic electric properties: the objects conduct electricity better along the direction of the DNA helices than perpendicular to them. By controlling the orientation of the plates in the corresponding single-molecule experiments, our collaborators in the Keyser group (U.K.) directly confirmed this prediction as well. Finally, our MD simulation of a hybrid structure, where the DNA origami plate is placed on top of a solid-state nanogap, revealed considerable yet reversible deformations of the DNA origami structure, providing a mechanistic explanation for reduced conductance in the hybrid systems with voltage that has been observed in experiment.

The results of our study demonstrate the utility of the MD method in predicting material properties of DNA origami objects. The study has immediate consequences on the design of the next generation of DNA origami nanopore sensors, where the effect of nucleotide composition and the ionic environment will be combined to minimize the leakage current through the plates. Furthermore, the structural responses of DNA origami to the changes in ionic and electrical environments suggest possible applications of DNA origami systems as electro-mechanic switches and other components of nanofluidic electronics.

FIGURE 1 (BACKGROUND): DNA origami plate (in orange and green) on top of a solid-state nanogap. Streamlines indicate the magnitude and direction of local water flux through the hybrid nanopore system. DNA origami is a deformable material with programmable electrical properties.

WHY BLUE WATERS?

Characterizing a given material property of a DNA origami object requires a MD simulation from 100 to 1,000 nanoseconds in length. Performing many such simulations in parallel was necessary to systematically evaluate the behavior of DNA origami as a function of its design or environmental conditions. Using Blue Waters allowed us to complete the study in less than one year, which would not have been possible using other supercomputer systems.

PUBLICATIONS

Li, C., et al., Ionic Conductivity, Structural Deformation, and Programmable Anisotropy of DNA Origami in Electric Field. *ACS Nano*, 9:2 (2015), pp. 1420–1433, doi:10.1021/nn505825z.

LARGE-SCALE SIMULATIONS OF BIOLOGICAL–SYNTHETIC INTERFACE MATERIALS FOR SENSING APPLICATIONS

Allocation: Illinois/0.74 Mnh
PI: Narayana R. Aluru[1,2]

[1]University of Illinois at Urbana–Champaign
[2]Beckman Institute for Advanced Science and Technology

EXECUTIVE SUMMARY:

Sensing and detecting biological molecules is of utmost importance in DNA mapping for human health. DNA base detection using nanopore technology has paved the way towards designing high-speed, high-precision sequencing devices. The high speed of DNA translocation and low signal-to-noise ratio are two challenges in DNA nanopore technology. For the first time, we introduce the mechano-sensitive channel (MscL) for DNA sequencing. MscL can produce two distinct tension and ionic current signals, making it an attractive pore for DNA sequencing technology. Moreover, MscL slows down DNA translocation by about one order of magnitude compared to traditional biological channels used for DNA sequencing. In another project, we found that single-layer MoS_2 is a promising material for DNA detection, exhibiting distinct signals for different types of DNA bases with a high signal-to-noise ratio.

INTRODUCTION

DNA sequencing using nanopore technology evolved greatly during the last few years. Oxford Nanopore Technologies is currently fabricating a USB-stick-size device that can sequence DNA 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 signal-to-noise ratio, pore degradation due to multiple uses, the identification of a single base in real time, and the high speed of translocation [1,2]. Engineering the translocation of the DNA through a biological/synthetic nanopore is a challenging problem in biotechnology. Although a number of studies have been performed on different types of pores in terms of translocation speed and ionic current blockade, fundamental understanding of the effect of pore architecture (the material and size of the pore) on the quality of DNA sequencing and speed of translocation is still lacking. In this study, we investigate both synthetic and biological nanopores by comparing the quality of the signals obtained during DNA translocation in nanopores using molecular dynamics (MD) simulations.

METHODS & RESULTS

Recently, we simulated all-atom MD of the translocation of DNA through a mechano-sensitive channel to distinguish different DNA base types according to the mechanical response of the channel [3]. In this study, a new type of signal, namely mechanical tension (in addition to ionic current signal), was introduced for DNA sequencing. We showed that the initially closed mechano-sensitive channel of large conductance (MscL) protein pore opens for single-stranded DNA translocation under an applied electric field. As each nucleotide translocates through the pore, a unique mechanical signal was observed; specifically, the tension in the membrane containing the MscL pore was different for each nucleotide. In addition to the membrane tension, we found that the ionic current was also different for the four nucleotide types. The initially closed MscL adapted its opening for nucleotide translocation due to the flexibility of the pore. This unique operation of MscL provided single-nucleotide resolution in both electrical and mechanical signals. Finally, we also showed that the speed of DNA translocation was roughly one order of magnitude slower in MscL compared to MspA (the membrane porin from Mycobacterium smegmatis which has been shown to be suitable for DNA sequencing), suggesting MscL to be an attractive protein pore for DNA sequencing.

In another study, which gained attention through many news outlets, we simulated DNA translocation through a MoS_2 nanopore [4]. Using atomistic and quantum simulations, we found that single-layer MoS_2 is an extraordinary material (with a signal-to-noise ratio >15) for DNA sequencing by two competing technologies (i.e. nanopore and nanochannel). A MoS_2 nanopore showed four distinct ionic current signals for single-nucleobase detection with low noise. In addition, single-layer MoS_2 showed a characteristic change in the total density of states for each base. The band gap of MoS_2 was significantly changed compared to other nanomaterials (e.g., graphene, h-BN, and silicon nanowire) when bases were placed on top of the pristine MoS_2 or armchair MoS_2 nanoribbon, thus making MoS_2 a promising material for base detection via transverse current tunneling measurement. MoS_2 nanopores benefit from a pore architecture (combination of Mo and S atoms at the edge) that can be engineered to obtain the optimum sequencing signals.

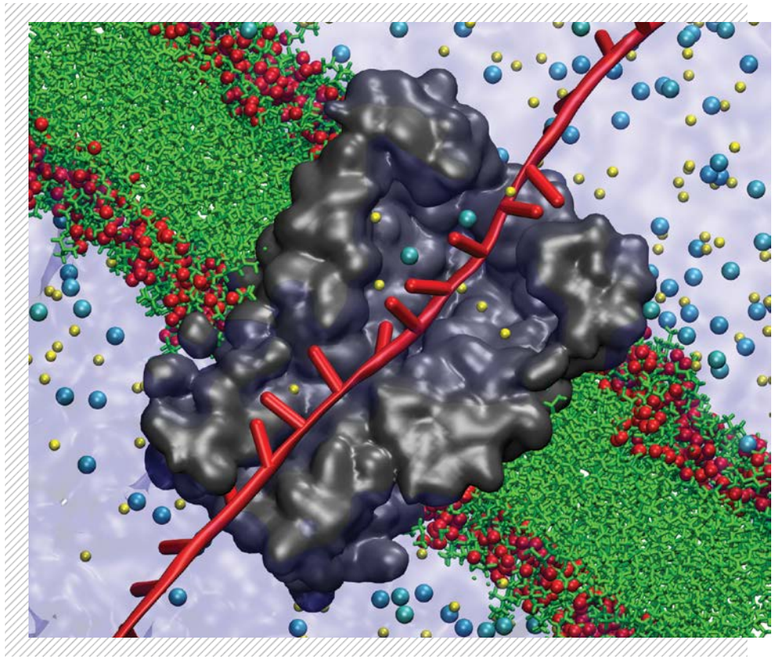

FIGURE 1: Single-strand DNA translocation through a mechano-sensitive channel.

WHY BLUE WATERS?

For both projects, we performed extensive MD simulations that involved up to 600,000 atoms. These expensive computations are not possible without a petascale supercomputer. Also, the MD package we used (NAMD) scales almost linearly with the number of cores up to 1,000 cores in our test on Blue Waters.

PUBLICATIONS

Farimani, A. B., M. Heiranian, and N. R. Aluru, Electromechanical Signatures for DNA Sequencing through a Mechanosensitive Nanopore. *J. Phys. Chem. Lett.*, 6 (2015), pp. 650–657, doi:10.1021/jz5025417.

Farimani, A. B., K. Min, and N. R. Aluru, DNA Base Detection Using a Single-Layer MoS_2. *ACS Nano.*, 8 (2014), pp. 7914–7922, doi:10.1021/nn5029295.

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

Allocation: GLCPC/0.10 Mnh
PI: Dinshaw S. Balsara[1]
Collaborators: Sudip Garain[1], John Reid[2]

[1]University of Notre Dame
[2]Rutherford Appleton Laboratory, U.K.

EXECUTIVE SUMMARY:

This project had two foci. First, we aimed to demonstrate that Coarray Fortran (CAF) offers petascale performance that is comparable to or better than the best that MPI-3 has to offer. Both CAF and MPI-3 have also been shown to outperform MPI-2 by a substantial margin. In doing and publishing this work we found efficient implementation strategies that work well for any one-sided messaging paradigm.

Second, we studied two-fluid turbulence in molecular clouds. The turbulent plasma that makes up a molecular cloud is predominantly made 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. This results in a modified turbulence. Such two-fluid turbulence has special signatures. We studied the special signatures and showed that it matches observed isophotologous lines from molecular clouds.

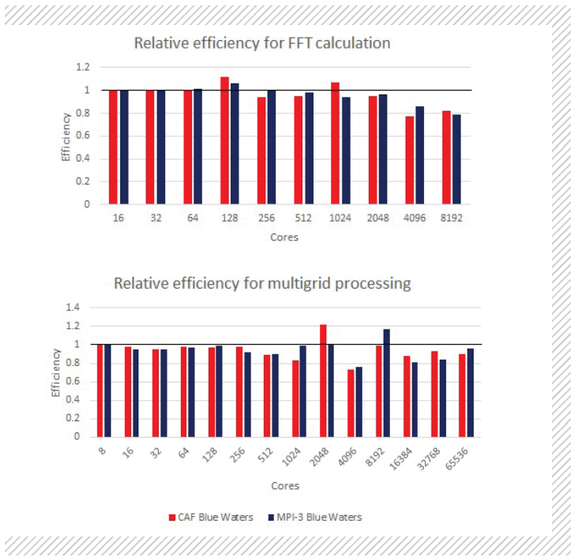

FIGURE 1: A weak scaling study for (a) fast Fourier transforms, and (b) multigrid-based partial differential equation solvers. The relative efficiency for Coarray Fortran and MPI on Blue Waters are shown in red and blue, respectively. [1]

INTRODUCTION

We are at the threshold of a new data-rich and simulation-rich era in star-formation studies. The new generation of infrared instruments will give us a new vista into chemical processes involved in star formation. Observations show that star formation takes place in a partially ionized plasma. Furthermore, Balsara [1] showed that it is very important to represent the plasma as two fluids, one made of neutral molecules and the other made of ions that are threaded by a magnetic field. Previous studies were analytical or restricted to low resolution. This research used higher model resolution to study this phenomenon. Additionally, we have NSF funding to test different highly scalable paradigms for the formation of young stars (protostars) on extremely scalable platforms.

METHODS & RESULTS

Scalability Study

In 2013–2014 we generated a very extensive comparison of Cray's Coarray Fortran (CAF) standard with MPI-3 for a range of partial differential equation applications. Both CAF and MPI-3 are novel programming paradigms and it would greatly help the community to have their capabilities documented and published. 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 carried out a weak scalability study and reported our results in Garain, Balsara, & Reid [2].

Fig. 1 shows the results of weak scaling on Blue Waters for fast Fourier transforms (FFT) and multigrid applications ranging from 8 to 65,536 cores. CAF and MPI-3 kept up with each other across the entire range of processors. We further found that both CAF and MPI-3 were more than twice as fast as MPI-2 when approaching large numbers of processors with one-sided communication. This showed the value of these novel programming paradigms when computing at scale.

In Garain, Balsara & Reid [2] we documented best practices for using CAF and MPI-3 for the community. We also showed that CAF code is much easier to write and maintain, and the simpler syntax made the parallelism easier to understand. Educational lectures on CAF were developed as part of this work and are freely available from our website [3].

Two-Fluid Turbulence Simulations

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 [4,5] and have also contributed to the theses of graduate students Meyer and Burkhart. Both have successfully landed jobs based on their computational skills.

Recent observations of differences in the linewidths between neutral and ionized tracers have led to the suggestion that ambipolar diffusion affects the turbulence in M17 [6]. Fig. 2 from Li & Houde [6] shows observed velocity dispersions as a function of length for the HCN molecule (black) and the HCO^+ ion (red). The dashed lines trace turbulent spectra. The velocity dispersion in the ions was lower than that in the neutrals on smaller scales, providing evidence for dissipation via ambipolar diffusion on small scales. The dissipation scales found by Li & Houde [6] were consistent with theoretical estimates of Alfvén wave dissipation by ion-neutral drag [2,7].

WHY BLUE WATERS

Our NSF-funded work required us to develop CAF-based adaptive mesh refinement (AMR) applications and demonstrate their scalable performance on high-end parallel supercomputers. We developed a CAF-based AMR magnetohydrodynamics application that worked extremely well with Cray's CAF compiler as well as Intel's CAF compiler. Our scientific goal was to do several two-fluid simulations of star formation, which take many more time steps (and therefore more compute time) than a single fluid simulation. Neither the scaling nor the two-fluid turbulence portions of this research would have been possible without the high-end computational resources that only Blue Waters can provide, and both are eminently suited to Blue Waters. In particular the Blue Waters technology enables very efficient one-sided communication.

PUBLICATIONS

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.jcp.2015.05.020.

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.jcp.2015.05.020

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), 2197, doi:10.1093/mnras/stt1893.

Burkhart, B., A. Lazarian, C. Meyer, D. S. Balsara, and J. Cho, Alvenic Turbulence Beyond the Ambipolar Diffusion Scale. *Astrophys. J.*, 805:2 (2015), 118, doi:10.1088/0004-637X/805/2/118.

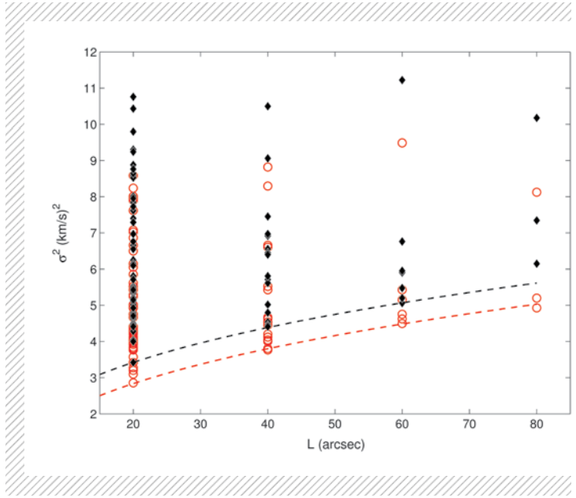

FIGURE 2: Observed velocity dispersions as a function of length for the HCN molecule (black) and the HCO^+ ion (red). The velocity dispersion in the ions was lower than that in the neutral molecules. [5]

REDUCING JET AIRCRAFT NOISE BY HARNESSING THE HETEROGENEOUS XK NODES ON BLUE WATERS

Allocation: Blue Waters Professor/0.20 Mnh
PI: Daniel J. Bodony[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY:

Reducing the jet exhaust noise from commercial and military aircraft will improve airport efficiency (by increasing the rate by which aircraft land and take-off, especially at night), reduce noise-related environmental pollution near the airport, and reduce noise-related injuries to aircraft personnel. However, the source of the jet noise defies a simple engineering description so very large simulations to accurately model the turbulence-generated sound sources are needed. The spatio-temporal multiscale nature of turbulence found in jet exhaust requires computational resources that exceed current petascale capabilities. The heterogeneous XK nodes on Blue Waters are being used as prototype hardware on which to develop a computational fluid dynamics prediction code that can harness on-node complexity and parallelism at scale in preparation for forthcoming extreme scale (e.g., exascale) computers.

INTRODUCTION

Without a guiding theory, reducing aircraft noise has been left to expensive trial-and-error experiments and, more recently, simulations. The turbulence-induced sound is generated over a very large spatial region leading to a multi-scale nonlinear fluid dynamics problem where the relevant energy is contained over six to ten decades of spatial and temporal scales. Simulations that capture the full range of spatial and temporal scales are beyond current computing capacity but remain our best hope for quieter aircraft. A paradigm shift in computational science, enabled by the XK nodes on Blue Waters, is desperately needed to advance the fields of compressible turbulence and aeroacoustics to reduce the noise from jet aircraft.

The project has impact in three critical areas. By targeting our software development to the heterogeneous XK nodes on Blue Waters we are developing power efficient, high performance computational fluid dynamics codes that can run across thousands of heterogeneous nodes and thus contribute to the emerging programming models and numerical algorithms compatible with forthcoming hardware complexity at extreme scale. Second, the scale of Blue Waters enables the largest simulations of compressible turbulent jet noise ever done that will advance the science of flow-generated sound through carefully conducted simulations and guided post-processing of the 100s-1000s of TB of data generated. Third, new engine nozzle designs will be developed that have the potential to reduce turbulent jet noise and improve the quality of life of airport communities and military personnel.

WHY BLUE WATERS

The research on Blue Waters concentrates on the heterogeneous XK nodes. When using all of the XK compute nodes the aggregate theoretical peak performance is more than 5.5 petaflops. Harnessing all the performance of the XK nodes across the entire Blue Waters system requires significant research and development of new programming approaches that (a) incorporate node-level heterogeneity and (b) can be scaled across 1,000s of nodes. Single node examples of GPU acceleration are commonplace, but codes capable of effectively using massively parallel heterogeneous machines are not.

The research contains two components. The first is focused on algorithm development to take an existing large software application that has already demonstrated scaling up to 100,000 x86 cores on ORNL's Jaguar and transform it into a CPU-GPU application that runs across all 4,000+ XK nodes with the ability to simultaneously use the x86 CPU cores and Kepler GPUs that exist on the XK nodes. The code, which solves the partial differential equations describing a compressible, viscous fluid, is currently based on MPI. To map the code onto CPU-GPU nodes, MPI is insufficient to work at the GPU level. Emerging source-to-source transformation tools that can retarget code to different architectures are needed to (a) reduce repetitious code generation tasks and to (b) abstract away many of the challenges associated with computing on a GPU device connected to a host. CPU source-to-source transformations represent a significant research task because there will be considerable inter-dependence on the memory utilization and localization, task pooling and scheduling, and bus contention on the XK.

The second component of the research utilizes the CPU-GPU enabled code on the XK nodes to conduct fundamental research on the reduction of jet noise from commercial jet aircraft engines. By building on existing research support in this area from the Air Force Office of Scientific Research and the Office of Naval Research, the heterogeneous code run across all of the XK nodes, simultaneously using the CPUs and GPUs, permits the largest, most ambitious investigation into turbulent jet noise ever conducted and allows, for the first time, detailed links to be established between the jet engine nozzle shape, the turbulent noise sources and, ultimately, the radiated sound. The simulations are conducted in two parts: first we use the immense scale and i/o capability of Blue Waters to conduct the most detailed study of how turbulent jet noise is produced over the range of jet exhaust conditions (namely, velocity and temperature) most commonly encountered in aircraft. The databases for these simulations will approach 1 PB in aggregate and provide unprecedented detail into the physics of noise generation. We will make these databases publically available. The second objective of the simulations is to design a quieter jet engine nozzle using an adjoint-based inverse design approach. At Blue Waters scale this design will be a watermark in the development of reduced jet engine noise because both the nozzle and flow will be included simultaneously. Because of the computational cost involved in adjoint-based design of fully turbulent flows, no such calculation has previously been attempted. Even at Blue Waters scale only one jet condition can be considered, and we will choose a high-subsonic dual stream jet exhaust typical of modern gas turbines, with a cool outer stream and hot inner stream. This simulation's objective is to identify what nozzle modifications are best suited for aircraft jet noise reduction. The societal impact of environmental noise pollution reduction is critical around airports, and a pacing issue for the health of military personnel operating on naval aircraft carriers. There are no computational resources available where several petaflops of computing resource can be utilized for a single calculation.

In 2015 we used the Blue Waters XK7 nodes, with NVIDIA GPGPUs. Our software tool H-MxPA, developed originally by Professor Wen-mei Hwu (University of Illinois at Urbana-Champaign), has been extended by creating a node-level runtime system that (a) permits asynchronous, simultaneous execution of the same kernel written in OpenCL on multicore CPUs and many core NVIDIA Kepler GPGPUs and (b) can discover and optimize the work distribution between streams sent to the CPU cores and GPGPUs cores, taking into account host-device transfer latency and data transfer times.

The H-MxPA runtime will continue to be hardened within the proxy application that replicates the data structure complexity and parallel environment found in the target application, PlasComCM. Continuing into 2016 we will implement the runtime in PlasComCM, verify the correctness of the results generated by code compiled with H-MxPA, and apply it to a series of increasing complex jet exhaust noise simulations using grids with up to 2 billion grid points and 10 billion degrees of freedom.

HYDROGEN UNDER EXTREME CONDITIONS

Allocation: Blue Waters Professor/0.24 Mnh
PI: David M. Ceperley[1]
Collaborator: Carlo Pierleoni[2]

[1]University of Illinois at Urbana-Champaign
[2]Universita of L'Aquila, Italy

EXECUTIVE SUMMARY:

We have completed a calculation of a first-order liquid-liquid transition in warm dense hydrogen using a quantum Monte Carlo method, Coupled-Electron Ion Monte Carlo. As hydrogen is compressed, it makes a sudden first-order transition between an insulating molecular liquid and a metallic atomic liquid. There are two conflicting experimental measurements for the transition of dense hydrogen that have been made during the past year. Our computational results lie between the two experiments.

INTRODUCTION

Hydrogen accounts for much of the visible mass in the universe. The properties of hydrogen and helium are important 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. It has long been an open question how hydrogen makes a transition from a molecular insulating state to an atomic metallic state [1]. We have developed new Quantum Monte Carlo simulation methods to treat such systems and using these methods have studied molecular dissociation in liquid hydrogen and have observed clear evidence of an extra liquid-liquid phase transition [2]. During the past year, two experimental groups [3,4] have reported observations of the transition we predicted; however, the observations do not agree with each other, differing in pressure by a factor of two. This motivated us to repeat our earlier calculations on Blue Waters, to better control the convergence and use recent improvements in methodology.

METHODS & RESULTS

We use a quantum Monte Carlo method (Coupled Electron Ion Monte Carlo) where we start with the true interaction between the electrons and protons and treat both fully quantum mechanically. In contrast to density functional calculations, all effects of electronic correlation are explicitly included. This 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 to the thermodynamic limit. With our method, we simulated hydrogen for temperatures in the range of 200K up to 5,000K, and at relevant pressures, 100GPa to 500GPa.

For temperatures below 2,000K we observe a first-order transition between an insulating molecular liquid and a more dense metallic atomic liquid. Our predicted transition pressure is 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 of the unusual properties of the molecular and atomic liquid.

The present calculations are needed to validate both our computational method and the experimental measurement. It is essential for progress in the high-pressure community to resolve 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 being 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. With more resources we will able to study systems containing elements heavier than hydrogen and helium, important not only in astrophysics but for materials.

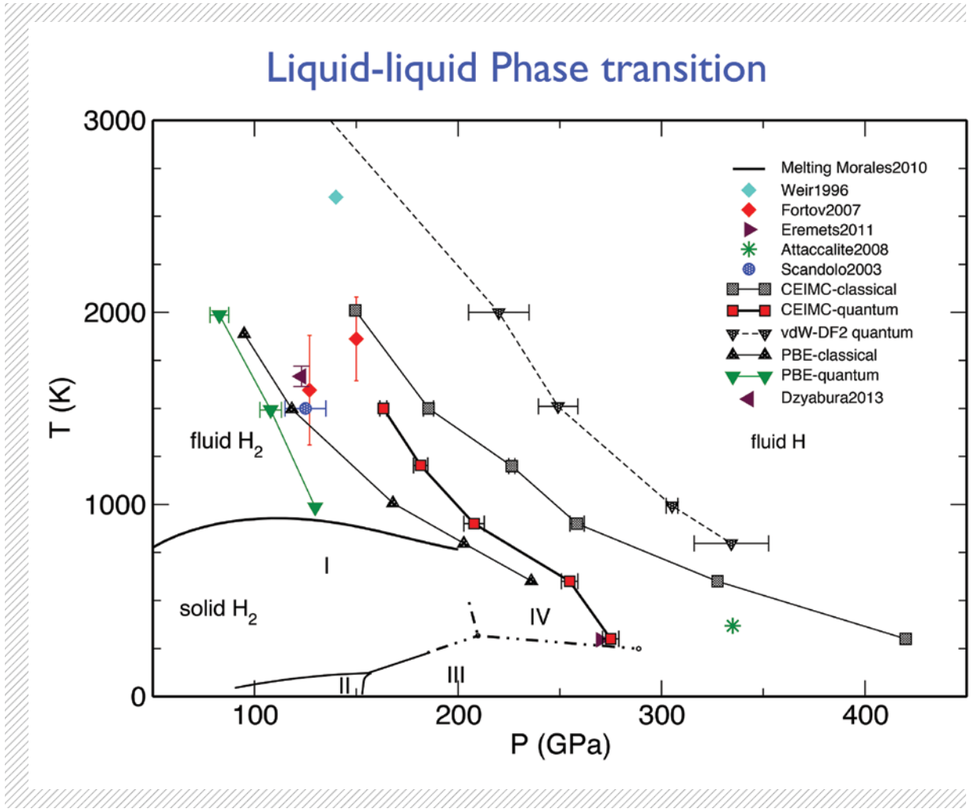

FIGURE 1: The predicted phase diagram (temperature versus pressure) of dense hydrogen. The red squares show the results from the Blue Waters calculation. The other symbols show results of other calculations or experimental measurements.

STRONGLY CORRELATED SYSTEMS THROUGH COMPUTATION: FROM BAD METALS TO PERFECT INSULATORS

Allocation: Blue Waters Professor/0.492 Mnh
PI: Bryan K. Clark[1]
Collaborators: Hitesh Changlani[1], Matthew Fisher[2], Katharine Hyatt[2], David Pekker[3], Xiongjie Yu[1]

[1]University of Illinois at Urbana–Champaign
[2]University of California, Santa Barbara
[3]University of Pittsburgh

EXECUTIVE SUMMARY:

Emergence is the principle where complicated phenomena arise from the interaction of simple rules. Two emergent phenomena in quantum mechanics include superconductivity, with its family of exotic neighboring phases, and many-body localization. It has often been said that the hardest part about understanding superconductivity is to make sense of the adjoining phases in which it is embedded; among these phases, the "bad metal," with gapless excitations but otherwise abnormal properties, stands apart [1].

Emergent phenomena in quantum mechanics can be particularly hard to understand as the computational complexity of simulating quantum mechanics scales exponentially. Using Blue Waters, we found and numerically simulated a model that supports a 2D bad metal phase. While superconductivity is the perfect conductor, the many-body localized (MBL) phase [2,3] is the exact opposite: a perfect insulator refusing to conduct even at infinite temperature. We numerically identified properties of the MBL phase, and found the algorithmic approach we developed transfers directly to a conceptual understanding of MBL. We pushed the limits of our simulations by developing improved algorithms.

INTRODUCTION

Bad Metals

In systems where you find superconductivity, you also often find the bad metal phase; understanding the bad metal phase is a likely way to make sense of superconductors. Unfortunately a simple mean field picture of the bad metal is absent and a tractable numerical model is needed. We find that fermions with a nearest-neighbor ring exchange term [4] fit the requirements.

Many-Body Localization

Quantum mechanics is usually seen at absolute zero (temperature). At high temperature, the strange nature of quantum mechanics is washed away leaving instead the classical physics we've understood for hundreds of years. One phase of matter, the many-body localized (MBL) phase, preserves the strangeness of quantum mechanics up to arbitrary temperatures. Many-body localization is formed in systems with simultaneous disorder and interactions. It is believed to form a perfect insulator and has emergent integrability. While much of the phenomenology of this phase is understood, what has been missing until now is a unifying framework in which to describe the wave-function of the MBL phase and a numerical approach for seeing the MBL phase on large systems.

New Algorithms

The limiting factor in solving much of materials and condensed matter physics is a combination of computational power and algorithmic techniques. There are two current de-facto simulation techniques for condensed matter problems: quantum Monte Carlo and tensor networks, but new ones are needed to move forward.

METHODS & RESULTS

Bad Metals

Using projector quantum Monte Carlo we computed the phase diagram of the ring exchange Hamiltonian. In addition to finding phase separation and a charge density wave, we discovered the first 2D system that convincingly demonstrates a bad metal phase—in this case an extremal d-wave Bose metal [5]. The smoking gun for this phase was the presence of two singular lines in the structure factor where it goes to zero and linear behavior in the structure factor at small momentum. Fig. 1 shows a canonical example from our simulation. Beyond finding the phase we also discovered a surprising feature—we could not identify any ordering in the spin sector, suggesting the possibility of a spin liquid riding on top of the bad-metal state.

Many-Body Localization

We devised a conceptual framework as well as a simulation approach for capturing the MBL phase through the matrix-product operator framework. A disordered non-interacting system is described by a simple product state. By turning on interactions, the product state turns into a matrix-product state where nearby states talk but distant ones have exponentially smaller effects. This defines local quasiparticles in terms of these matrices. In this language, the varied properties of the MBL phase are transparent. In addition, we have developed and used a variant of the density matrix renormalization group algorithm to capture the "finite temperature" eigenstates of the MBL phase from which we derived the properties.

New Algorithms

We found a way to combine the positive features of both quantum Monte Carlo and tensor network methods. Of particular interest is our ability to take an important algorithm that does not parallelize and use many of its features in a massively parallel way.

WHY BLUE WATERS?

Bad Metals

Unlike classical physics where single machines are able to simulate thousands of particles, simulating quantum physics is significantly more difficult. Even systems with hundreds of particles can be a computational tour de force. Blue Waters has been essential because any systems smaller than those we consider mask the physics of the bad metal phase. The quantum Monte Carlo algorithm required to simulate the bad metal phase requires 10^6 walkers and forward walking done to large imaginary time. With the next Track-1 system one could take the next step in understanding how this bad metal phase transitions into other nearby states, including the superconductor.

Many-Body Localization

Many-body localization is a phase where disorder is fundamental. Therefore, any simulations require an average over potentially thousands of disordered configurations to produce accurate statistics. Such averaging can be done only on large machines. In addition, to date there has been no reliable method of generating finite-temperature excited states. Finding the right approach required significant experimentation and each such experiment is computationally costly.

New Algorithms

Blue Waters was essential for developing these new algorithms. Not only was it necessary for testing the massive parallelization, but the benchmarking that was required to authenticate the methods required significant computational resources.

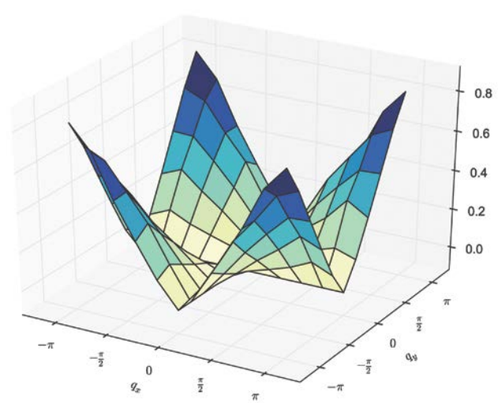

FIGURE 1: Charge structure factor of the ring-exchange Hamiltonian at density 4/10 on a 10 x 10 lattice with 10,000 walkers and tau=0.1. The singular lines show the characteristic gapless modes that are the smoking gun of the extremal DBL phase.

NEXT-GENERATION *AB INITIO* SYMMETRY-ADAPTED NO-CORE SHELL MODEL AND ITS IMPACT ON NUCLEOSYNTHESIS

Allocation: GLCPC/0.5 Mnh
PI: J. P. Draayer[1]
Collaborators: D. Lang[2], T. Oberhuber[2], J. P. Vary[3], P. Maris[3], U. Catalyurek[4], M. Sosonkina[5]

[1]Louisiana State University
[2]Czech Technical University
[3]Iowa State University
[4]The Ohio State University
[5]Old Dominion University

EXECUTIVE SUMMARY:

We carried out large-scale modeling of light and medium-mass nuclei. This included short-lived nuclei not yet accessible to experiment but key to understanding astrophysical processes, which are the foci of current and next-generation rare isotope experimental facilities.

The scale of computational challenges inherent to modeling such intricate quantum many-body systems makes Blue Waters essential for addressing long-lasting challenges in nuclear theory and experiment, as well as astrophysics. A breakthrough theoretical advance [1] coupled with the Blue Waters cutting-edge computational power opened a new region, the intermediate-mass nuclei from fluorine to argon isotopes, for first investigations with *ab initio* (i.e. "from first principles") 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.

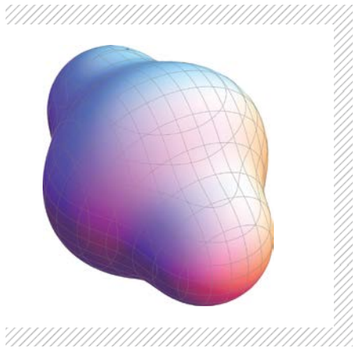

FIGURE 1: First *ab initio* results for the ^{20}Ne ground state, only feasible with the capabilities of the SA-NCSM symmetry-guided concept and the Blue Waters system.

INTRODUCTION

One of the quintessential open problems in contemporary physics research is to design a comprehensive many-body theory for modeling and predicting nuclear structure and processes starting from inter-nucleon forces that are consistent with the underlying theory of quantum chromodynamics (QCD). The ultimate goal of *ab initio* nuclear theory is to establish a bridge from the properties of many-nucleon systems down to the fundamental blocks—namely, nuclear potentials that are tied to the underlying principles of QCD. While this bridge makes the problem highly compute-intensive, it also empowers *ab initio* models with a universal character essential for modeling the co-existence of diverse nuclear substructures and predictive capabilities vital for descriptions of nuclear species far from the valley of stability. Because such nuclei are often key to understanding processes in extreme environments, from stellar explosions to the interior of nuclear reactors, first-principles nuclear models have had and will have a tremendous impact on advancing the frontiers of multiple branches of physics such as astrophysics, neutrino physics, and applied physics. While the *ab initio* nuclear structure approaches for light nuclei have achieved remarkable progress, computational challenges hindered applications to heavier systems.

METHODS & RESULTS

We developed an innovative *ab initio* nuclear structure approach (with concomitant computer code named LSU3shell), dubbed the symmetry-adapted no-core shell model (SA-NCSM) [1], that embraces the first-principles concept and capitalizes on exact as well as approximate symmetries of nuclei [2]. The *ab initio* SA-NCSM solves the time-independent Schrödinger equation as a Hamiltonian matrix eigenvalue problem using many-nucleon basis states organized according to definite shape deformation and total intrinsic spin. The main advantage of this basis stems from the fact that it is designed for the description of low-lying nuclear dynamics, thereby supporting the exclusion of configurations that contribute very little to nuclear states of interest [3,4]. The theory provides an opportunity for first-principles-guided applications, from quark/gluon considerations to nuclear structure and reactions for rare isotopes beyond the lightest "s-" and "p-shell" systems (hydrogen to carbon isotopes) [5–8]. This approach significantly reduces the dimensionality (i.e. size) of the model space from those encountered when using ultra-large basis spaces of standard *ab initio* approaches.

These theoretical advances coupled with the power of Blue Waters allowed us to advance an *ab initio* large-scale nuclear modeling initiative that proffered cutting-edge predictive capabilities for determining the structure of nuclear systems, including rare isotopes up through medium-mass nuclei that are inaccessible experimentally and fall far beyond the reach of other *ab initio* methods. Our aim is to provide nuclear structure information of unprecedented quality and scope that will further understanding of fundamental symmetries in nature that are lost in massive datasets or require exascale architectures, and to extract essential information for astrophysics (e.g., nucleosynthesis and stellar explosions), neutrino physics, and energy-related applied physics problems.

These points describe results and projected studies:

- We provided the first *ab initio* description of ^{20}Ne (an isotope of neon) and ^{18}F nuclei (a fluorine radioisotope) [9]. These are examples of open-shell isotopes in the intermediate-mass region, with complexity far beyond the reach of complementary *ab initio* methods. Following this success, we will target neon, magnesium, and silicon isotopes, especially those close to the limits of stability (at proton and neutron drip lines), providing important input to nuclear reaction studies. Such reactions in the intermediate-mass region are key to further understanding phenomena like X-ray burst nucleosynthesis or the neon-sodium and magnesium-aluminum cycles.
- We studied electron scattering off ^{6}Li (a lithium isotope) with wave functions calculated in the *ab initio* SA-NCSM [10]. Results show the efficacy of the SA-NCSM model space selection, for the first time, toward reproducing the low- and high-momentum components of the ^{6}Li ground-state density. This finding is crucial for planned studies of neutrino scattering off ^{12}C (a carbon isotope) and ^{16}O (an oxygen isotope).
- A work-in-progress focuses on one of the most challenging problems in nuclear physics: achieving an *ab initio* nuclear modeling of the first excited 0^+ state (the so-called Hoyle state) in ^{12}C, the resulting state of the essential stellar triple-alpha process. Knowing the structure of low-lying states of ^{12}C is key to modeling nucleosynthesis and stellar explosions.

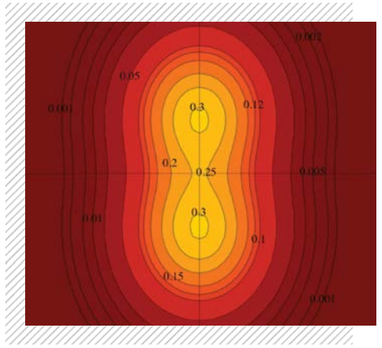

FIGURE 2: Density profile of the challenging ^{8}Be ground state from first principles, revealing two alpha subclusters. Results are obtained in the SA-NCSM utilizing only a small fraction of a large-scale model space.

WHY BLUE WATERS?

The *ab initio* nuclear structure studies are extremely compute-intensive and at the forefront of physics research. The state-of-the-art techniques of computational group theory in the SA-NCSM approach dramatically reduce the size of the problem and the associated memory requirement to hundreds of terabytes and petabytes at the cost of a major increase in computing intensity. Numerical investigations of the intermediate-mass region of the periodic table with the SA-NCSM approach are beyond the scale of available academic HPC systems.

Currently, only Blue Waters provides resources required for the *ab initio* SA-NCSM studies of medium-mass isotopes with cutting-edge accuracy. The largest production runs ran 717,600 concurrent threads on 22,425 Cray XE6 compute nodes. The next generation of Track-1 system will dramatically improve accuracy and predictive power of first-principles nuclear structure modeling while extending its reach towards even heavier isotopes.

PUBLICATIONS

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 (review article). *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. Int. Innov., 160 (2015), 20.

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.

DIRECT SIMULATION OF DISPERSED LIQUID DROPLETS IN ISOTROPIC TURBULENCE

Allocation: NSF PRAC/1.9 Mnh
PI: Said Elghobashi[1]
Collaborator: Michele Rosso[1]

[1]University of California, Irvine

EXECUTIVE SUMMARY:

The objective of our research is to enhance understanding of the two-way interactions between liquid droplets and a turbulent flow by performing direct numerical simulations (DNS). The freely moving deformable liquid droplets are fully resolved in three spatial dimensions and time and all the scales of the turbulent motion are simultaneously resolved down to the smallest relevant length and time scales. Our DNS solve the unsteady 3D Navier–Stokes and continuity equations throughout the whole computational domain, including the interior of the liquid droplet. The droplet surface motion and deformation are captured accurately by using the level set method. Discontinuous density and viscosity are smoothed out across the interface by means of the continuous surface force approach. A variable-density projection method imposes the incompressibility constraint.

INTRODUCTION

The two-way interactions between liquid droplets and turbulent flows are important in nature and engineering applications, for example cloud formation and burning of liquid fuel in internal combustion (reciprocating and jet) engines and rockets. Understanding the physical details of the vaporization and mixing processes in a turbulent flow is an essential prerequisite to understanding the chemical reaction and the eventual control/optimization of energy conversion.

METHODS & RESULTS

The freely-moving deformable liquid droplets are fully resolved in time and 3D space, and all the scales of the turbulent motion are simultaneously resolved down to the smallest relevant length and time scales. Our direct numerical simulations (DNSs) solve the unsteady 3D Navier–Stokes and continuity equations throughout the whole computational domain, including the interior of the liquid droplet. The droplet surface motion and deformation are captured accurately by using the level set method (LSM). The discontinuous density and viscosity are smoothed out across the interface by means of the continuous surface force (CSF) approach. A variable density projection method is used to impose the incompressibility constrain.

So far we focused on validating our numerical solution method. We compared our simulation results of an initially spherical liquid water droplet settling in stagnant air, under the effect of gravity and surface tension, with available experimental data. We monitored the shape changes, the terminal velocity, and the conservation of both mass and momentum. Once this validation is completed we will simulate the motion of a number of liquid droplets in isotropic turbulence. This will allow us to examine the dispersion statistics of the droplets and their effects on the carrier flow turbulence (two-way coupling).

WHY BLUE WATERS?

DNS of turbulent flows is very demanding in terms of computational power and memory. The computational grids need to be fine enough to resolve the smallest flow structures accurately; this requirement becomes more and more stringent as the Reynolds number, based on the Taylor microscale, increases. In addition, we seek an accurate time history of the flow in order to compute time-dependent statistics, thus limiting the time-step size (time interval) for advancing the solution in time. For example, DNS of single-phase isotropic turbulence at Reynolds number 300 requires a grid of $2{,}048^3$ mesh points, and about 12 hours on 65,536 processors to cover seven large eddy turnover times.

The demand for computational power is even larger if a multiphase flow (e.g., liquid droplets in air) is considered; the standard projection method for incompressible flows must be replaced by a variable-density projection method. The latter results in a variable-coefficients Poisson's equation that is not solvable by means of a fast Fourier transform, thus requiring an iterative solver. We use the multigrid-preconditioned conjugate gradient solver provided by the PETSc library. Given the requirements outlined above, Blue Waters is a necessary tool for our research.

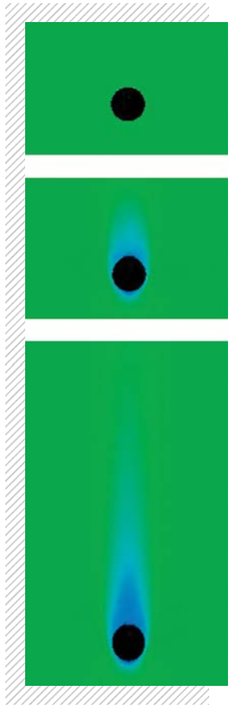

FIGURE 1: A liquid water droplet with initial diameter 0.5 mm is falling in stagnant air under the effect of gravity (=10 times gravity). The droplet's initial spherical shape changes with time due to the effects of acceleration, surface tension, and drag forces. Time (a) 0.0s; (b) 0.001 s; (c) 0.005 s. The velocity contours in the wake region are shown in blue.

SIMULATING THERMAL TRANSPORT IN NANOSTRUCTURES FROM THE BALLISTIC TO THE DIFFUSIVE REGIME

Allocation: Illinois/1.0 Mnh
PI: Elif Ertekin[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY:

With the power of the Blue Waters supercomputer, we have been carrying out large-scale molecular dynamics simulations to reveal the unusual physics that underlies thermal transport in nanostructured and low-dimensional materials. In these newly emerging materials, the physics of thermal transport remains not well understood. Thanks to Blue Waters, our work in this area has spanned simulations ranging from the ballistic to the diffusive transport regime and has enabled key developments to enhance understanding of how phonons, the primary heat carriers in solids, interact with structure and disorder at these scales. Our simulations have allowed us to answer several outstanding questions relating to (i) the possible divergence of the thermal conductivity in low-dimensional crystals as the sample size increases [1], (ii) how disorder and defects scatter phonons in low-dimensional materials [2], (iii) anomalous effects of strain on thermal transport in low-dimensions [3], and (iv) how low-dimensional superlattices can give rise to coherent phonon transport, enabling the design of nanostructured systems that serve as phonon waveguides and filters.

INTRODUCTION

Understanding heat transport is critical across a variety of applications: optimizing heat shields for rocket and jet engines, solid-state refrigeration via thermoelectrics, and enabling next-generation cooling technologies for powerful supercomputers. Heat transport at the macroscale is governed by Fourier's law and is comparatively well understood. Phonons are quantized lattice vibrations that are responsible for transporting heat through a solid. Fourier's law applies when the system size is large, relative to the mean free paths, which are typically a few hundred nanometers. If not, interesting physics arises. For example, when structure at the nanoscale is introduced, several phenomena emerge that have no analog at larger scales. This arises because the characteristic lengths are comparable to the mean free paths of the primary heat carrying phonons. At the nanoscale, phonons interact with the structure. In addition, the emergence of flexural phonon modes in low-dimensions (out of plane vibrations) introduces heat carriers with anomalous dispersion, which gives rise to unusual temperature-dependence of the thermal conductivity. The role of these carriers (scatterers of phonons, or contributors to conduction?) is now widely debated in the research community.

METHODS & RESULTS

The method we use to carry out our analysis is large-scale molecular dynamics simulations. We use highly parallelizable molecular dynamics frameworks on Blue Waters to access computational domains and simulation sizes that would otherwise be too computationally costly. In the molecular dynamics approach, the motion of each atom in the computational domain is tracked over time, and ensemble averages over the sampled configurations give access to properties of interest. In particular, we use two methods within molecular dynamics. The first is equilibrium molecular dynamics, in which a system is simulated under equilibrium conditions (without net heat current). The equilibrium fluctuations in the system contain information about the subtleties of heat transport. For instance, phonon mean free paths can be extracted from the heat flux autocorrelation function and the thermal conductivity can be determined. The second approach is non-equilibrium molecular dynamics, in which a temperature profile is established across the computational domain. The phonon flow in the system can be statistically analyzed, and linear response characteristics via the fluctuation dissipation theorem once again enable us to extract features of the thermal conductivity.

Some of the most exciting highlights of our work from the previous year are:

• Simulation of transport in two dimensional materials from the ballistic to the diffusive regime. In low dimensions, heat transport is governed by two competing factors. The first is the long-ranged, ballistic phonon carriers that give rise to increasing thermal conductivity with growing sample size; the second is the presence of flexural modes that may scatter these carriers. Figure 1 clearly demonstrates a transition from ballistic to transitional flow in two-dimensional graphene. In addition to demonstrating the transition, we obtain new physical insights into why the divergence of the thermal conductivity is suppressed: the flexural modes.

• Competition between disorder and low-dimensionality: Which one wins? For diamond nanothreads (1D materials), we were able to demonstrate the dramatic effects of disorder in 1D materials. On one hand, the thermal conductivity of a 1D material is expected to diverge linearly with increasing system length. On the other hand, disorder effects are greatest in low dimensions. How these two effects play out in a disordered, one-dimensional system is currently not well understood. Our simulations, as shown in Figure 2, show a dramatic effect: Even at the smallest degrees of disorder sampled, the divergence of the thermal conductivity becomes quickly suppressed in one dimension.

• Superlattices in two-dimensional materials: observation of coherent phonon transport. We have shown, in work published in 2014 in *Physical Review B*, that the introduction of a superlattice structure in two-dimensions allows tuning of the thermal conductivity through wave interference effects. These superlattice structures effectively demonstrate the design of thermal circuit components, such as the phonon equivalent of optical waveguides or filters.

• Anomalous strain effects in low-dimensional materials. Conventional wisdom tells us that applying strain (in 3D solids) tends to reduce the thermal conductivity of a material. We have demonstrated that in 2D, the thermal conductivity of a material exhibits an anomalous, non-monotonic dependence on the applied strain. We are able to offer physical insights for the underlying mechanisms. This work was published in 2015 in *Physical Review B*.

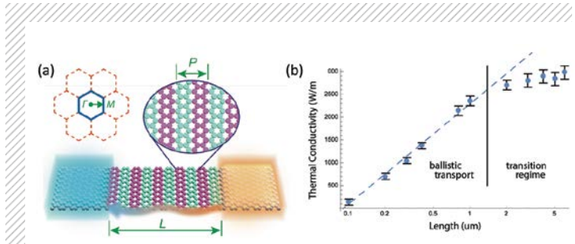

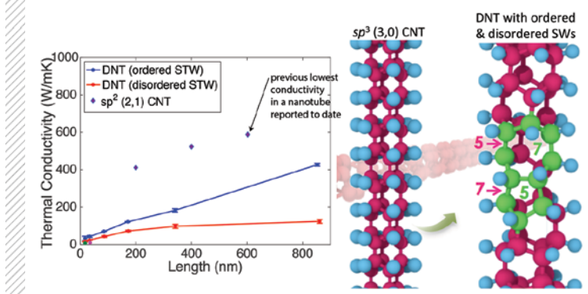

FIGURE 1 (TOP RIGHT): Non-equilibrium molecular dynamics simulations of heat transport in two-dimensional graphene carried out to large length and time scales have revealed a transition from ballistic to transitional transport. The deviation of the thermal conductivity from the logarithmic dependence demonstrates the transition clearly.

FIGURE 2 (BOTTOM RIGHT): We analyzed the competition between disorder and low-dimensionality in the materials known as "diamond nanothreads," which are the (CONT.) newest member of the carbon material family. The ordered system (blue line) shows the expected linear variation of the thermal conductivity on the length. When a small distribution of topological defects (Stone-Wales defects) is introduced (red line), the thermal conductivity quickly converges to a finite value.

WHY BLUE WATERS

The Blue Waters supercomputer has enabled us to carry out simulations of systems of several hundreds of thousands of atoms. Accessing these size scales is critical since diffusive transport in low dimensions does not emerge until system sizes approach ~10 μm or larger. Other research approaches are not capable of accessing these scales without phenomenological models or assumptions about governing transport laws. For example, we were able to demonstrate a transition in the nature of the flow from the ballistic to the transitional flow regime in two-dimensional graphene (Figure 1). This required simulations of sample size ranging from 2 to 6 μm, the largest simulations of low-dimensional heat transport carried out to date within the non-equilibrium molecular dynamics framework.

PUBLICATIONS

Zhu, T. and Ertekin, E. Phonon transport on two-dimensional graphene/boron-nitride superlattices. *Physical Review B* 90 (2014), 195209.

Zhu, T. and Ertekin, E. Resolving anomalous strain effects on two-dimensional phonon flows: the cases of graphene, boron nitride and planar superlattices. *Physical Review B* 91 (2015), 205421.

HIGH-ORDER METHODS FOR TURBULENT TRANSPORT IN ENGINEERING AND GEOSCIENCES

Allocation: Blue Waters Professor/0.491 Mnh
PI: Paul Fischer[1]
Collaborators: Som Dutta[1], Anne Goering[1], Anthony Patera[2]

[1]University of Illinois at Urbana–Champaign
[2]Massachusetts Institute of Technology

EXECUTIVE SUMMARY:

We describe simulation of two cases of turbulent flows. The first addresses sediment transport in bifurcating rivers and channels. The second provides input to system design codes for energy co-generation in industrial plants. The simulations are based on the scalable open-source code Nek5000, which employs minimally dispersive spectral element discretization. Results to date shed light on the physics of observed sediment deposition behavior in bifurcating channels.

INTRODUCTION

Turbulent transport is the principal driver for many processes in physics, engineering, geosciences, and biology. Examples include the in-fall of matter into black holes, combustion in automotive and aerospace applications, sediment and pollutant transport in rivers and oceans, and atherogenesis (formation of fatty deposits on arterial walls) in arterial blood flow. Our objective is to address these questions through direct numerical and large-eddy simulation of turbulent flow by solving the governing Navier–Stokes and associated transport equations. The open problems are as varied as the associated geometries and are challenging because of the range of scales present in turbulent flows at high Reynolds numbers (i.e. high speeds).

This project focuses on two topics from civil and mechanical engineering. The first project addresses sediment transport in bifurcating rivers. It has been observed that when a stream divides between a main branch and a side channel, a disproportionate amount of the near-bed sediment is often directed into the side channel, which can ultimately alter the flow dynamics and cause blockage of the side channel. Experiments investigating this effect date back almost a century, but the dynamics of the process have yet to be clearly identified.

The second project aims to construct fundamental building blocks for analysis of turbulent flow losses in energy co-generation systems. Patera's group at MIT developed reduced-order models for thermal-fluids systems with provable error bounds. These models reduce complex nonlinear partial differential equations to low-dimensional ordinary differential equations that can be solved on a laptop or a phone. The models, however, require high-fidelity baseline computations at a few points in the originating state space. Industrial co-generation systems involve many piping components (e.g., orifices, junctions, and elbows) and are of such complexity as to be intractable, even on the world's fastest computers. Fortunately, low-dimensional models permit assembly of the simulation results for individual piping components, each of which is within our grasp on platforms such as Blue Waters. Using a map of the response of each component, coupled with rigorous reduced-order models, it will be possible to design complex systems with confidence at low cost.

METHODS & RESULTS

Our turbulence simulations were based on the open-source spectral element code Nek5000 [1]. The spectral element method (SEM) is a domain-decomposition approach in which the solution is represented by tensor-product polynomials on individual bricks that are assembled to cover the entire domain. The bricks are typically curvilinear, which allows accurate representation of the geometry. The local tensor-product structure allows low-cost and low-storage matrix–matrix product-based operator evaluation so that high-order polynomials may be used with almost no overhead. The SEM thus yields minimal numerical dissipation and dispersion at low cost, which is ideal for simulation of turbulent flows in complex domains. Nek5000 was recognized with a Gordon Bell prize in HPC [2] and has scaled beyond one million MPI ranks.

For the sediment transport simulations, we built a sequence of meshes that followed the original experimental investigations [3] and commenced with a parameter study examining a range of flow splits and Reynolds numbers. In addition to turbulent flows, we analyze laminar cases in order to understand secondary flow mechanisms (boundary layer flows driven by external pressure gradients) that might dictate near-bed transport. Our first set of simulations investigated the effects of the channel flow split distribution (e.g., 85% in the main branch, 15% in the side) on the flow patterns downstream of a 90° bifurcation at Reynolds number Re=7,000. Fig. 1 shows the low-speed sections in the flow (blue) that yielded longer particle residence times and greater likelihood for bed deposition. The sediment transport problem requires sophisticated tracking algorithms to capture the physics of low-density particulate transport. We developed a parallel particle tracking routine that is stable for all Stokes numbers. It uses hash tables and fast generalized all-to-all exchanges to rapidly migrate particles to their host processors. We are currently testing the particulate physics on canonical problems that yield known particle distributions.

For the co-generation project, we will simulate flow through a pipe section with a hole in the side. The initial computational domain will consist of 60,000 elements of order 10, corresponding to 60 million gridpoints and 240 million degrees of freedom to be determined at every time step. The run time for each data point in the parameter space will be approximately 30 hours on 8,192 cores. In the target device, the pipe acts as a plenum for low-speed distribution of product into a large mixing tank. The first set of computations will yield pressure and flow relationships over a range of inlet/outlet pressure pairings. We will also investigate the sensitivity of the results to hole shape (taper, chamfer, etc.).

WHY BLUE WATERS?

Blue Waters provides the computational power and the relatively short queue times to quickly turn around large-scale turbulence simulations. This capability is critical, particularly in the early development stages of the project when we first start to explore resolution requirements and mesh sensitivity. The process is interactive and would be significantly hampered by slow turn-around times.

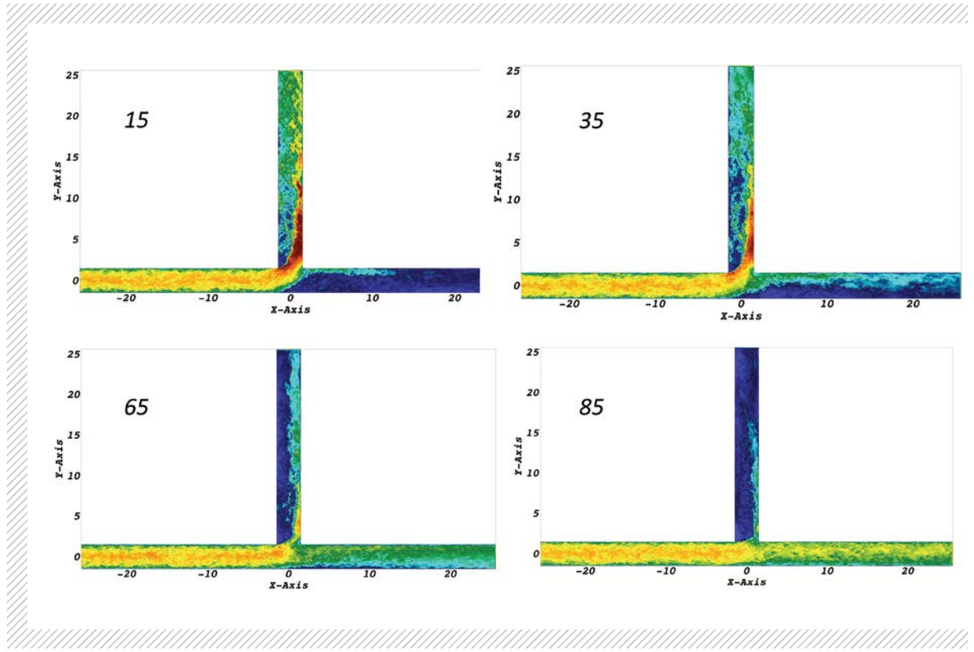

FIGURE 1: Nek5000 simulations of turbulent flow patterns at z=.05 H, Reynolds number 7,000, for main-channel flow percentages ranging from 15% to 85%, with low velocities in blue and high velocities in red. Biased velocity distributions are evident in both the main and side channels. [Simulation by Som Dutta]

MANY-GPU SIMULATIONS FOR SOFT MATTER DESIGN

Allocation: GLCPC/0.443 Mnh
PI: Sharon C. Glotzer[1]

[1]University of Michigan

FIGURE 1: Left: Strong scaling performance of a fluid of two million particles in dissipative particle dynamics on up to 1,000 GPUs, comparing HOOMD-blue and LAMMPS. Right: Strong scaling of an oleic acid brush grafted onto a substrate, in hexane solvent (not shown). Performance is shown for execution on up to 128 GPU (XK7 nodes) or CPU (XE6) nodes. A GPU/CPU node speed up of 12.5x is achieved.

EXECUTIVE SUMMARY:

We aim to leverage large-scale parallel simulations on GPUs to discover fundamental principles of how soft matter systems self-assemble and how to control the assembly process to engineer new materials. Our open-source molecular dynamics (MD) package HOOMD-blue [1], designed specifically for GPU architectures, uses completely device-resident data structures; all simulation work is performed on the GPU so the CPU merely acts as a driver for the GPU, and it offers a rich feature set for general-purpose particle-based simulations [2]. We implemented MPI spatial domain decomposition and additional code optimizations and demonstrated unprecedented performance for MD simulations running on thousands of GPUs in parallel. In addition to carrying out benchmarks, we used Blue Waters to run dissipative particle dynamics (DPD) on large systems for 50 million time-steps each to reproduce experimental findings for formation of nanofibrous microspheres and elucidate the mechanisms of self-assembly.

INTRODUCTION

Multi-GPU simulations are becoming a mainstay in scientific computing. Simulations of bulk polymeric liquids or large assemblies of anisotropic shapes involve large numbers of particles, which are best simulated using spatial domain decomposition. Such a work distribution allows the simulation of a single system on many processors—GPUs and CPUs—and results in high scaling efficiency. Often, speed-ups of several orders of magnitude can be achieved over execution on a single processor. The challenge in multi-GPU simulations is overcoming the inherent latencies of the PCIe bus, which limits communication between these extremely fast processors. Our multi-GPU enabled HOOMD-blue 1.0 contains many optimizations that were tested and benchmarked on Blue Waters. For the first time we were able to simulate self-assembly of hollow microspheres with a fibrous nanostructure and a hollow microstructure. The discovered mechanisms provide guidance on simultaneous control of nano- and micro-structure formation for synthesis of nanostructured particles, which may broadly impact biomedical and other emerging technologies.

METHODS & RESULTS

The first HOOMD releases up to 0.11.3 are highly tuned for single-GPU performance and do not support multi-GPU runs. Reducing latency is one of the biggest challenges in developing a code scalable to many GPUs. Data transferred between GPUs moves over the PCIexpress bus (PCIe), whose bandwidth (up to 16 GB/s) and latency (several μs) is much slower than on-board GPU memory (250 GB/s,~100 ns). Communicating over PCIe adds latency that is not present in single-GPU runs. In the strong scaling limit of increasing the number of GPUs, P, at constant number of particles, N, the work N/P performed by each GPU decreases to the point where it is too small to fully use the device. We addressed this challenge using the MPI protocol to support communication between different GPUs on the same node and between different nodes. Our communication routines are implemented on the GPU to reduce the data transferred over PCIe and to allow us to take advantage of CUDA-aware MPI libraries. We updated a majority of the classes in HOOMD-blue to support MPI, including the file I/O classes, integrators, pair and bond potentials, and analyzer classes. We also optimized for strong scaling on thousands of GPUs, which we achieved using a design for the neighbor list and force computation kernels based on cooperative thread arrays and an auto-tuning algorithm.

Our benchmarks on Blue Waters further establish GPUs as extremely fast engines for MD simulation compared to traditional CPU cores. GPUs not only realize an order of magnitude speed-up over current-generation CPUs, but they also scale very well using spatial domain decomposition.

This functionality is critically important in the study of real systems with interactions at the nanoscale that can dictate macroscale behavior. In our *Advanced Materials* paper, a collaborative effort with an experimental group at the University of Michigan, we studied the behavior of microdroplets with tunable porosity, which serve as implants to aid in tissue regeneration. By varying the star polymer architecture (arm number or length), we demonstrated that droplets change from solid to hollow or porous structures. These features are critically important for proper tissue adhesion. We systematically investigated the droplet formation for several lengths of polymer and arm number, each at numerous droplet sizes. These DPD simulations required over 10 million particles, divided between polymer and solvent, to approach the micron-sized length scale of smaller droplets. Our simulations would not have been possible without the scalable performance of HOOMD-blue and simultaneous access to up to 128 GPUs on Blue Waters to investigate the droplet phase diagram.

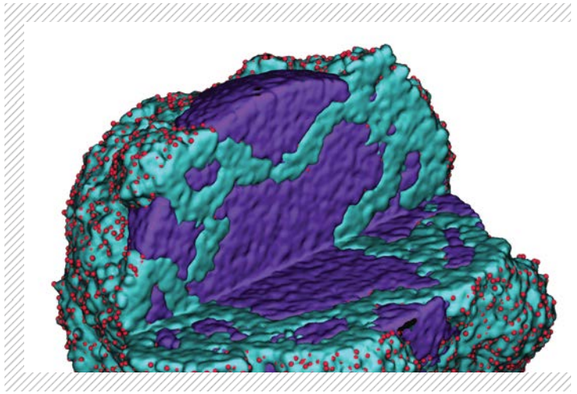

FIGURE 2: Isosurface of large-scale DPD simulation of 15-arm star-shaped polymers of varying arm length, showing spongy structure formation.

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 and thereby drastically improve the turnaround for research. We used Blue Waters to benchmark the performance of the HOOMD-blue code on up to 1,024 nodes, and the results have been published in Glaser et al [2]. Additionally, as recently reported in *Advanced Materials*, our Blue Waters allocation allowed us to run large systems on up to 128 nodes to reproduce experimental findings for structure and formation of nanofibrous microspheres and elucidate the mechanisms of self-assembly for the system.

PUBLICATIONS

Glaser J., Nguyen T.D., Anderson J.A. et al., Strong scaling of general-purpose molecular dynamics simulations on GPUs. *Comput. Phys. Commun.* 192, (2015), pp. 97-107. DOI:10.1016/j.cpc.2015.02.028.

Zhang Z., R. L. Marson, Z. Ge, S. C. Glotzer, and P. X. Ma, Simultaneous Nano- and Microscale Control of Nanofibrous Microspheres Self-Assembled from Star-Shaped Polymers, *Adv. Mater.* 27:26, (2015), pp. 3947–3952. DOI: 10.1002/adma.201501329

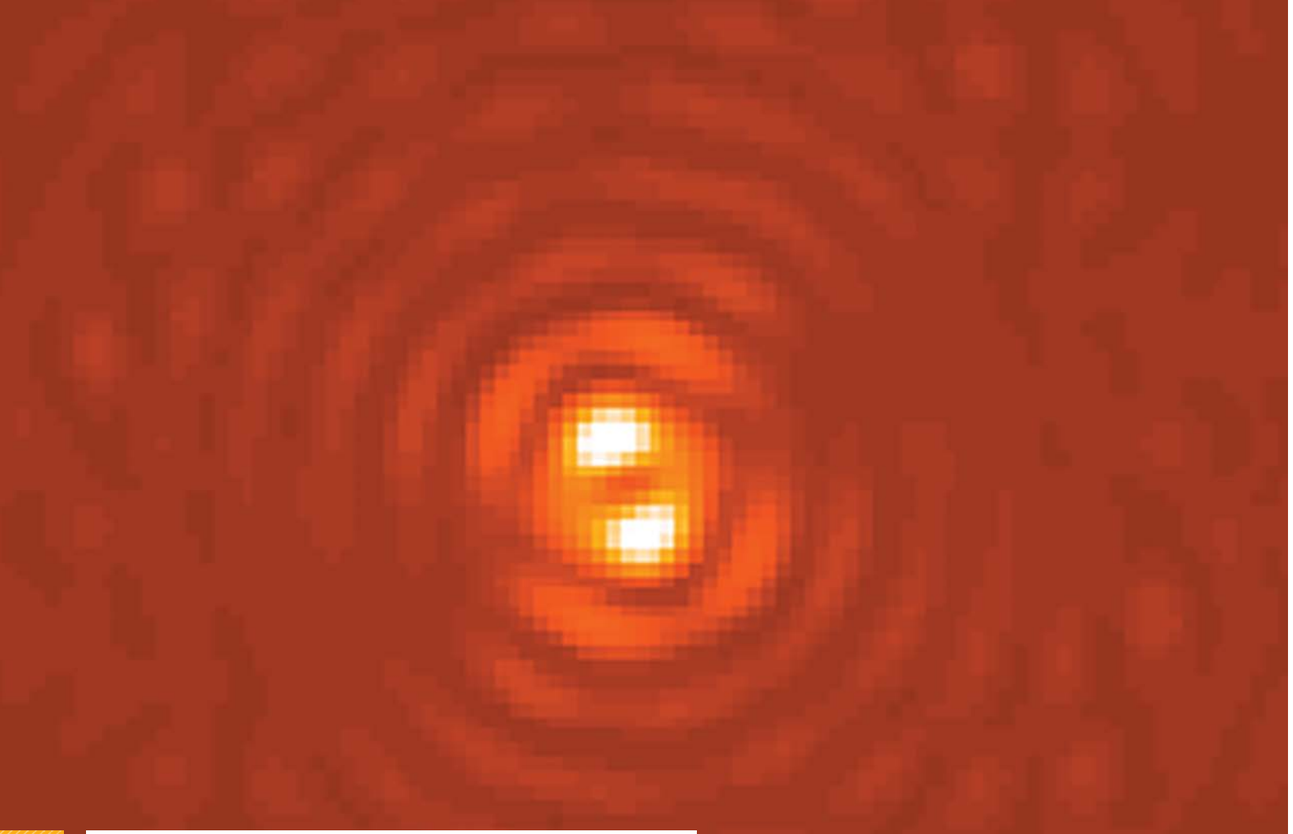

NEW TECHNIQUES FOR PIXEL-LEVEL FIDELITY ASSESSMENT IN INTERFEROMETRIC IMAGES

Allocation: Blue Waters Professor/0.28 Mnh
PI: Athol Kemball[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY:

Contemporary astronomical interferometry for modern telescope arrays and future planned instruments such as the Square Kilometer Array (SKA) present grand-challenge problems, now uniquely amenable to solution using extreme-scale computing. Data rates in this domain are rapidly escalating, and these telescopes will require extreme-scale computing within a decade. My research with Blue Waters has focused on new techniques for pixel-level fidelity assessment in interferometric images, a uniquely challenging problem, as well as investigations of other issues of relevance to the computation and data challenges faced in this discipline within the coming decade.

INTRODUCTION

Interferometry is a key observational technique that enables astronomical imaging at the highest angular resolution. In this technique, the spatial coherence of the incident radiation from the astronomical source is measured by cross-correlation or interferometric combination at all projected spatial separations sampled by the configuration of telescopes in the array. These projected separations vary over time due to Earth's rotation. Therefore, interferometry creates effective telescope apertures far exceeding any monolithic telescope diameters that could easily be physically constructed but at the cost of only partial sampling of the aperture. The measured spatial coherences are connected to the astronomical source brightness distribution on the sky by an integral equation that reduces to a Fourier transform in the limit. However, the synthesized aperture is inherently sparsely sampled in interferometry, so the inverse problem is ill-posed and requires mathematical regularization for solution. In addition, the imaging integral equation contains terms representing atmospheric propagation and instrumental calibration errors at each telescope. These terms need to be accounted for along with the unknown image brightness distribution during inversion from the measured coherence data. Our project on Blue Waters has concerned several highly challenging problems in this domain, specifically an extension of prior work concerning pixel-level fidelity estimation, as well as the exploration of data analysis in this domain at petascale.

This work is important because of escalating data rates, which scale in interferometry approximately as the square of the number of telescopes in the array multiplied by increasing detector data acquisition rates enabled by exponential advances resulting from Moore's Law. This places peta- or exascale computational demands on interferometric calibration and image formation for future interferometers such as the SKA, as these arrays will have a large number of antennas and high data-acquisition rates. These facilities represent significant community investments and need to achieve their scientific goals.

METHODS & RESULTS

Our approach on Blue Waters has been to extend interferometric fidelity estimation techniques that we have previously developed [1,2] to larger scale and over a more complete parameter range. These techniques rely on statistical resampling methods using variants of bootstrap resampling for statistically-dependent data that are intrinsically computationally very demanding. Blue Waters has allowed much larger interferometer array sizes to be considered along with richer physical models. Both are more realistic in terms of anticipated future telescopes.

This domain requires extreme-scale HPC within a decade, so this work is important to the discipline. The computational intensity and I/O and computation heterogeneity are also unusual in terms of the distribution of traditional HPC user needs. These results will also allow the community to extract greater value from significant capital investments in public telescope facilities.

WHY BLUE WATERS

The scale of Blue Waters is essential to this work due to the computational demands of the algorithmic methods used and the size of the interferometer arrays considered in these studies. Blue Waters allows fundamentally new approaches to be explored due both to its scale and balance. The system and project staff are an essential element of the success of this project as they provide a level of system support that ensures high uptime and throughput in addition to a very deep level of technical expertise. A future Track-1 system would allow significant advances in the work described here, further aiding the utilization and design of expensive capital investments in this domain.

PUBLICATIONS

Seidel, E. and Allen, G. and Freemon, M. and Gammie, C.F. and Kemball, A.J., Pennington, R. and Petravick, D., Computing and Data Challenges for Multi-Messenger Astronomy, *Exascale Radio Astronomy,* (AAS Topical Conference Series Vol. 2., Monterey, California, 30 March-4 April, 2014, Bulletin of the American Astronomical Society), Vol. 46, #3, #200.01.

FIGURE 1 (BACKGROUND): Image fidelity of a region of a reconstructed interferometric image, here showing variance per pixel over two angular coordinates on the sky.

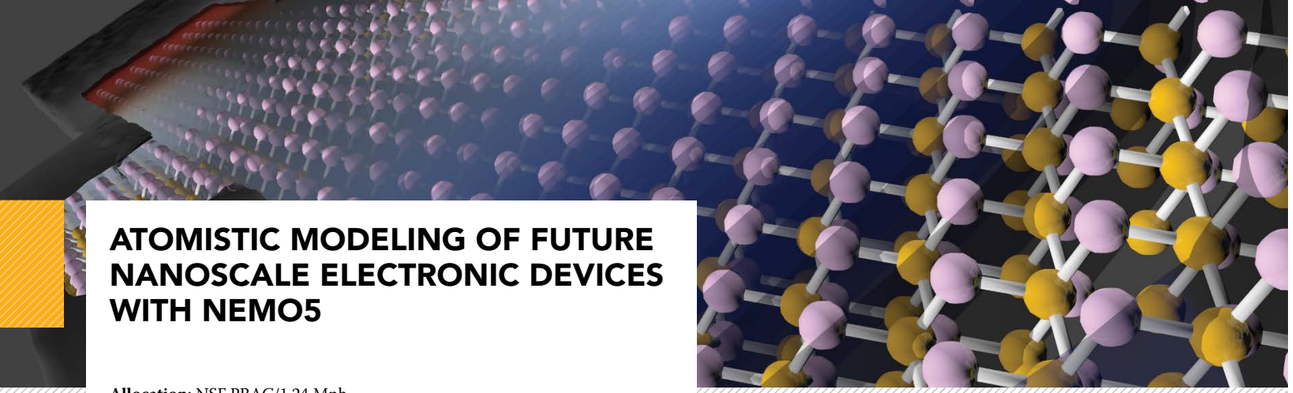

ATOMISTIC MODELING OF FUTURE NANOSCALE ELECTRONIC DEVICES WITH NEMO5

Allocation: NSF PRAC/1.24 Mnh
PI: Gerhard Klimeck[1]
Collaborators: Joerg Appenzeller[1], Fan Chen[1], Zhihong Chen[1], Jim Fonseca[1], Yu He[1], Hesameddin Ilatikhameneh[1], Luis Jauregui[1], SungGeun Kim[1], Tillmann Kubis[1], Michael Manfra[1], Saumitra Mehrotra[1], Michael Povolotskyi[1], Rajib Rahman[1], Daniel Mejia[1], Mehdi Salmani-Jelodar[1], Harshad Sahasrabudhe[1], Yaohua Tan[1], Archana Tankasala[1], Yu Wang[1], Lloyed Hollenberg[2], Sven Rogge[3], Michelle Simmons[3]

[1]Purdue University
[2]University of Melbourne
[3]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, transistor size will continue to decrease in the next ten years, but foundational issues with currently unknown technology approaches must be pursued. The number of atoms in critical device 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. The Institute for Nanoelectronic Modeling's (iNEMO) software package, NEMO5, 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 (BACKGROUND): Visualization of the atomic structure and the electrostatic potential (blue/red) in a periodic ultra-thin-body device of InAs. [Credit: Daniel Mejia]

INTRODUCTION

The U.S. is a world market leader in the semiconductor industry and produces a significant number of 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. At the same time, the end of Moore's law scaling as we know it will be reached within ten years with device dimensions expected to be about 5 nm long and about five to ten atoms in its critical active device width. Further improvements in shrinking dimensions will come only through detailed study of device designs and materials and include 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 iNEMO to address fundamental questions such as those above on a variety of semiconductor devices. Besides enabling basic research covering the fields of engineering, physics, and materials science, NEMO5 is used by leading semiconductor corporations to design future devices. The source code, binaries, and support for academic use are available through nanoHUB.org.

METHODS & RESULTS

iNEMO research with the NEMO software suite on Blue Waters encompasses work for the new International Technology Roadmap for Semiconductors (ITRS) projections, disordered transistor leads, devices for quantum computing, and transistors made of materials such as bilayer graphene and topological insulators.

- NEMO5 results on future device design parameters have been included in the ITRS since 2013. Traditional device designs exploited light effective masses to increase electron mobility. Simulations showed that overall transistor characteristics can be improved by engineering a higher effective mass in order to decrease quantum tunneling and thereby reduce current leakage when devices are in the transistor OFF state.
- As devices are shrunk for performance gains, atomistic fluctuations and associated quantum effects have stronger effects on the devices' characteristics. Research into non-ideal device leads/contacts has focused on alloy disorder and how design tuning can optimize heat flow in ITRS devices to avoid Joule heating and recycle heat into usable electric power. This research showed that failure to include realistic disordered leads in simulations overestimated transistor current by a factor of two.
- Quantum dot donor configuration interaction simulations, in close collaboration with leading experimental groups, help us understand and design two-qubit gates for semiconductor quantum computing. Blue Waters simulations elucidated charging energies and the configurations of the two-electron ground state of a single negative donor as a function of depth and external electric fields. Two-electron ground states of donor molecules matched well to experimental results and a new quantum dot design has been proposed based on the electric field dependency of the exchange coupling.
- Graphene promises improved device characteristics if a suitable bandgap can be opened to provide semiconductor behavior. Recent studies with NEMO5 have simulated transport in bilayer graphene with a tight-binding approach and dynamically controlled bandgap for calibration against experimental devices.
- Surface states of topological insulator nanowires are expected to serve as scattering-free charge conductors at room temperature. The atomistic tight binding approach showed that surface state dispersion in topological insulators is geometry dependent, a result that cannot be described by simpler approaches.

WHY BLUE WATERS?

Without Blue Waters, iNEMO research is, at best, hampered, and at worst, impossible. For instance, a simple transport calculation of a 50 nm long wire with a 3 nm diameter (about 28 nm^2 in cross section) requires around 1 TFlop/s for a single energy point using the non-equilibrium Green's Function method. Resolution of a device's characteristics requires about 1,000 energy points, and this calculation must be repeated perhaps a dozen times for a full current–voltage sweep. Computational time scales with the cube of cross sectional area (relative to the direction of the electron flow) and linearly with the length of the device. The treatment of a technically currently relevant finFET device would require an atomistic resolution of a device with a cross-section around (20x40) nm^2, which includes the core semiconductor and the surrounding gate material. Blue Waters allows researchers to get results in a day, rather than weeks as on other computer systems. Additionally, some devices, such as those in the topological insulator work, are simply too large physically (due to the number of atoms and size of related tight binding matrices) to fit in memory on smaller systems.

A future generation Track-1 system would allow 100-million-atom 3D simulations for many crystals using spin and/or classical multi-physics as well as allow for ultra-scaled transistors to be simulated with important physical phenomena such as scattering and time dependence. Transistor device environments could be included, which have an increasingly strong effect on device performance as the channel region is diminished. Statistical ensembles would also be possible to investigate process and fabrication effects on physical models and on device characteristics.

Finally, Blue Waters staff provides exemplary support and user outreach to guide system usage, help with issues as they arise, and assist with code performance and scaling.

FIGURE 2: Visualization of the atomic resolved SI-Ge alloy ultra-thin-body device with surface roughness. Colors indicate source (green), drain (orange), and channel (red) regions. The grey/black area is oxide. [Credit: Daniel Mejia]

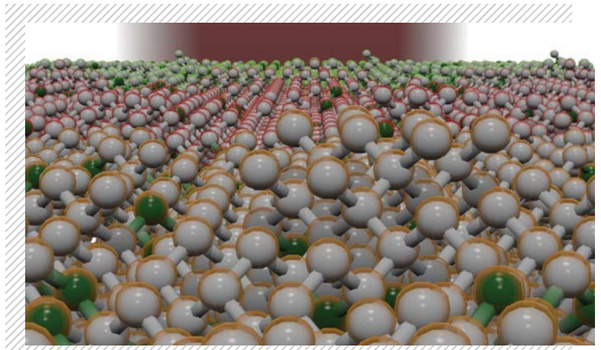

SPARSE MATRIX FACTORIZATION IN SOLID MECHANICS AND GEOPHYSICS ON CPUS AND GPUS

Allocation: Industry/60 Knh
PI: Seid Koric[1,2]
Co-PI: Anshul Gupta[3]
Collaborators: Natalia Gimelshein[4], Steve Rennich[4], Vladimir Puzyrev[5], Scott Wilkin[1,2]

[1]National Center for Supercomputing Applications
[2]University of Illinois at Urbana–Champaign
[3]IBM T. J. Watson Research Center
[4]NVIDIA Corp.
[5]Barcelona Supercomputing Center

EXECUTIVE SUMMARY:

Sparse matrix factorization is a critical algorithm in many science, engineering, and optimization applications. The performance of the massively parallel direct multi-frontal solver Watson Sparse Matrix Package (WSMP) [1] for solving large sparse systems of linear equations arising in an implicit finite element method in solid mechanics and an inversion problem in geophysics was evaluated on Blue Waters and achieved new records in sparse matrix factorizations both with CPU and GPU solvers. We performed full-scale benchmarking tests up to 65,536 cores (100 Tflop/s) on XE6 nodes using assembled global stiffness matrices and load vectors with 5 million to 40 million unknowns extracted from "real-world" commercial implicit finite element analysis (FEA) and academic geophysics codes. We also present a minimally invasive approach to the GPU acceleration of WSMP that can more than double the CPU-only numerical factorization performance and scale beneficially up to as many as 512 Cray XK7 nodes.

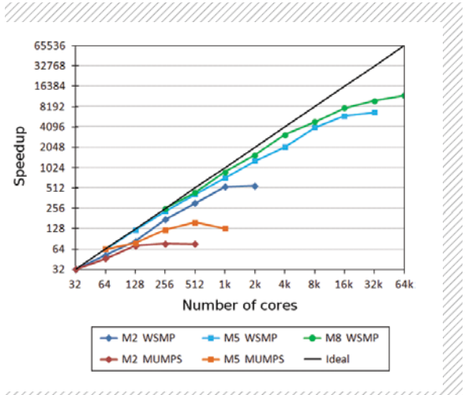

FIGURE 1: WSMP parallel speedup for geophysical tests on XE6 nodes of Blue Waters.

INTRODUCTION

Across a range of engineering fields, the use of simulation and computational models is pervasive for designing engineered systems. HPC systems play an essential role in simulations and modeling. Researchers and manufacturing teams depend on HPC to create safe cars, energy-efficient aircraft, effective communication systems, and efficient supply chain models. HPC is also necessary for large-scale geophysical simulations; in order to obtain a realistic image of a geologically complex area, industrial surveys collect vast amounts of data, making the computational cost extremely high for the subsequent simulations.

A major computational bottleneck of implicit finite modeling in solid mechanics and the inversion algorithms in geophysics is the solution of large sparse systems of linear ill-conditioned equations Ax = b in real or complex domains. There has been considerable interest in the development of numerical algorithms for solving large sparse linear systems of equations and their efficient parallel implementation on HPC systems for more than three decades. The algorithms may be grouped into two broad categories: direct methods and iterative methods. Finding and computing a good preconditioner for use with an iterative method can be computationally more expensive or often impossible, thus making implicit numerical methods with direct solvers often the only feasible methods with ill-conditioned systems. The limitation on CPU and memory requirements made the use of direct solvers uneconomical in the past, resulting in broad use of iterative solvers. The recent rise of petascale computational resources, however, has greatly increased the efficiency and practicality of using direct solvers for large sparse systems.

METHODS & RESULTS

For this study, we have chosen a few test systems resulting from large-scale geophysics electromagnetic (complex domain) and solid mechanics (real domain) problems.

Parallel speedup of the factorization of the three largest geophysics test matrices is shown in fig. 1. In these experiments, WSMP reached a performance of 97 Tflop/s on 65,536 cores of Blue Waters in the solution of the linear system with 8 million unknowns in the complex domain, or 16 million unknowns in real domain. Our scalability study credibly showed that the direct solver WSMP is prepared for even larger HPC simulations with systems having upwards of 10 million complex unknowns, which will provide high-fidelity results for the large inversion problem (example in fig. 2).

Besides running the standard CPU-only version of WSMP, we have shown that, with remarkably minor modifications to the original CPU code, the numerical factorization performance of WSMP can be accelerated on a reasonably large number of GPU-enabled XK7 nodes on Blue Waters. Modifications to the original code were mostly limited to intercepting high-level BLAS operations and handling them using a new GPU library.

The factorization performance of GPU-accelerated WSMP, so called ACCEL_WSMP, is given in fig. 3. Benchmarking with industrially derived finite-element analysis solid mechanics matrices showed that on 128 nodes, the GPU-accelerated code was two times faster than the CPU-only code using the same number of GPU or CPU nodes. A fairly high degree of acceleration was observed for lower node counts. The WSMP numerical factorization showed GPU acceleration by a factor of 2.5 to 3.5 on up to 32 nodes. Due to issues with load balancing and difficulty maintaining sufficient computational intensity to hide CPU/GPU transfers, the observed speedup was reduced at larger node counts. Work is underway to further improve the performance of the GPU code.

WHY BLUE WATERS?

Blue Waters is the only place where massively parallel sparse solver technology such as WSMP can be tested by taking full advantage of large amounts of distributed memory, tens of thousands of modern multicore processors and GPUs, and low latencies and increased bandwidth of leading interconnect network technologies. This exciting technology advancement will lead to a massive leap in terms of advances in design and manufacturing, and understanding the properties of the Earth subsurface to allow a major breakthrough in oil exploration, to name two of many applications in science and industry that can benefit from this work on current and future petascale architectures hosted at NCSA.

PUBLICATIONS

Gimelshein, N. E., A. Gupta, S. C. Rennich, and S. Koric, GPU acceleration of WSMP. GPU Technology Conference 2015, San Jose, Calif. , March 17–20, 2015.

Koric, S., Q. Lu, and E. Guleryuz, Evaluation of massively parallel linear sparse solvers on unstructured finite element meshes. *Computers and Structures*, 141 (2014), pp. 19–25, doi:10.1016/j.compstruc.2014.05.009.

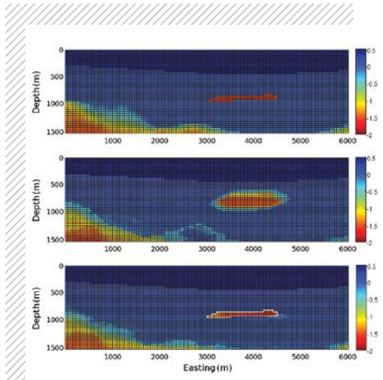

FIGURE 2: Real model (top) and the results of the unconstrained stand-alone electromagnetic inversion (middle) and constrained by the seismic results (bottom). Color represents electric conductivity on a logarithmic scale.

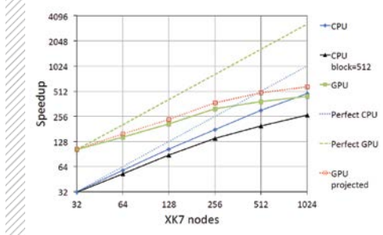

FIGURE 3: Numerical factorization speedup vs. number of XK7 nodes used for the M20 test system originating from solid mechanics finite element analysis.

SHOCK-INDUCED TURBULENT MIXING

Allocation: NSF PRAC/3.91 Mnh
PI: Sanjiva K. Lele[1]
Collaborator: Akshay Subramaniam[1]

[1]Stanford University

EXECUTIVE SUMMARY:

This project aimed to perform high-fidelity simulations of shock-induced multi-material mixing in a simple yet realistic configuration. Shock-induced acceleration of an interface between two different fluids renders the interface unstable to Richtmyer–Meshkov (RM) instability. Initial perturbations of the interface grow rapidly due to spatially varying torque exerted by the shock on the interface. Interaction of the reflected shock with a distorted interface causes a rapid breakdown of the flow into a turbulent state, which further stirs up the interface and drives mixing.

Our simulations focused on the interaction of a normal shock wave with a planar fluid interface that is at an incline to the shock wave. RM instability and shock-driven mixing is of paramount importance in many modern engineering and scientific applications: hypersonic air-breathing engines to increase mixing of fuel and oxidizer, the physics of the collapse of supernovae, transition from deflagration to detonation due to shock-flame interaction, and energy generation through inertial confinement fusion.

INTRODUCTION

The National Academy of Engineering has identified "Provide Energy from Fusion" as one of the grand challenges for engineering. Inertial confinement fusion (ICF) is a clean, sustainable source of energy and has the potential to play a huge role in addressing concerns about energy security and global climate change. The Richtmyer–Meshkov (RM) hydrodynamic instability that causes premature mixing of the capsule interface has been identified as a critical factor that limits the performance of ICF. Though the ICF problem is far more complex, predictive simulations of the hydrodynamics of shock-induced mixing is key to better design of ICF targets.

We focused on the RM instability due to a normal shock impinging on an inclined material interface (fig. 1). The simulations were coordinated with experiments at the Shock Tube & Advanced Mixing Laboratory at the Texas A&M University [4,5] and were the first to address critical issues of flow confinement by walls and resolve the effects of molecular mixing in shock-driven turbulent mixing. The proposed numerical algorithms (high-order spectrally optimized compact scheme and anisotropic solution-dependent artificial fluid properties for shock and interface capturing) and flow solver infrastructure (Miranda) are ideally suited for the problem. Conditions in the laboratory experiment were duplicated in the simulations. This will allow systematic validation of the computed results at large scales. Databases on RM-instability-driven turbulence at unprecedented spatial resolution will allow investigation of fundamental questions about turbulence physics in strongly driven transient flows and also support development of novel sub-scale closures for variable-density turbulence.

FIGURE 2: Overlay of z-vorticity and dilatation fields at t=1.5 ms. The vorticity drives the flow dynamics and the dilatation shows the shock structure that is responsible for the vorticity deposition.

METHODS & RESULTS

Due to the nature of the problem, numerical simulations of the RM instability required algorithms that treat both turbulence and discontinuities in the form of shock waves and material interfaces accurately. This is a unique challenge that is not addressed by traditional algorithms for turbulence or shock and interface capturing since the latter often employ excessive numerical dissipation for stability whereas the former try to minimize the numerical dissipation that can damp out the small scales of turbulence. The conflicting requirements make this problem extremely challenging.

We addressed this problem by using a tenth-order spectrally optimized compact difference scheme [2] in conjunction with the localized artificial diffusivity (LAD) scheme [1,3] for shock and interface capturing. The high order and spectral accuracy of the compact difference scheme preserved a high fraction of the theoretically resolvable scales in the flow without dissipating them. The LAD scheme defaulted to one with very low dissipation in smooth regions of the flow and surgically introduced numerical dissipation near regions of shocks and interfaces.

Fig. 2 shows an overlay of vorticity (in brown and green) and dilatation (in red and blue). This image shows the vorticity deposited on the interface by the shock that drives the dynamics of the problem. Complex shock patterns in blue are also seen and cause inhibition of the instability growth rate. Fig. 3 shows the mass fraction of SF6 at a later time of turbulent mixing.

WHY BLUE WATERS?

The large range of scales in flows involving the RM instability, from scales dictated by the geometry of the problem to the fine scales that govern molecular diffusion, makes this problem very demanding in terms of computational resources required. Even a simulation with one billion grid points does not adequately capture all the physical scales in the problem. Accurate depiction of all the flow features would require much higher spatial resolution. In addition, computational requirements cannot be lessened due to the lack of accurate lower-fidelity models like a subgrid-scale model for compressible multi-species wall-bounded flows. The availability of a petascale resource was critical to the success of this project.

HPC resources available under the next Track-1 system would allow us to complete more realistic simulations of the inclined interface RM mixing and with varying parameters, which would enable development of effective subgrid models for such flows in the future.

FIGURE 3 (BACKGROUND): Mass fraction of SF_6 at t=5.0 ms. The flow at this time is in a turbulent mixing regime that causes air and SF_6 to mix. Fully white is air and dark blue is SF_6. Shades in between represent mixed states.

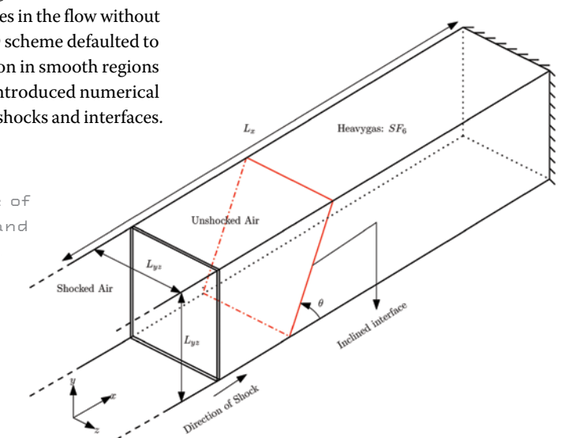

FIGURE 1: Schematic of the problem setup and initial conditions.

HIGH-PERFORMANCE COMPUTING OF HYPERSONIC, SHOCK–SHOCK INTERACTIONS USING KINETIC, PARTICLE APPROACHES

Allocation: Illinois/70 Knh
PI: Deborah A. Levin[1]

[1]University of Illinois at Urbana–Champaign

EXECUTIVE SUMMARY:

When steep gradients occur in the strong shocks of hypersonic flows, the use of continuum formulations such as the Navier Stokes equations is questionable. Instead, the direct simulation Monte Carlo (DSMC) method provides a numerical solution to the Boltzmann equation of transport, particularly for modeling thermochemical non-equilibrium shock-dominated flows. In the DSMC method, the flow is modeled by a series of collisions that occur by using simulated particles, each of which represents a large number of real molecules. The approach is computationally tractable and well suited to parallelization because the particle collisions and movement are decoupled. Recently we have implemented an approach that uses octree grids to reduce the cell size only in regions where the mean free path is small (i.e. distance travelled between particle collisions is small, so lots of particle collisions), thus reducing the required number of computational particles.

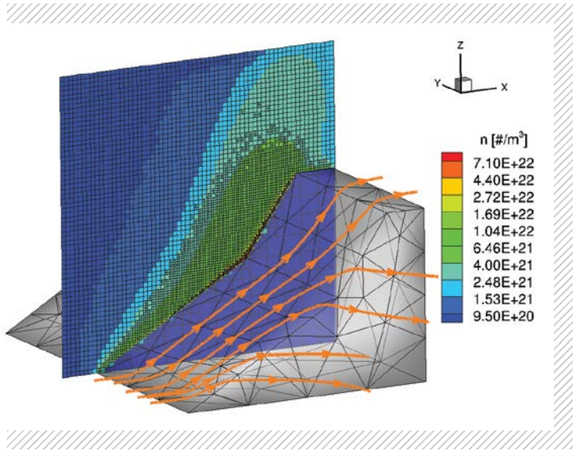

FIGURE 1: Streamlines and number density contours over a double wedge for Mach-7 flow showing increased density near the surface and the presence of 3D effects.

INTRODUCTION

In direct simulation Monte Carlo (DSMC) the flow is modeled by a series of binary collisions that occur by using simulated or numerical particles, each of which represents a large number of real molecules. The time step is selected to be small compared to the mean time between collisions. We are developing a novel, in-house DSMC code, known as scalable unstructured gas dynamics adaptive refinement (SUGAR), using the adaptive mesh refinement (AMR)/octree grids has been applied to the study of 3D ion thruster plumes [1]. The case being considered here, hypersonic flow over a double wedge, presents different challenges because a high level of spatial refinement is needed to capture shock-shock and shock-boundary layer interactions that create regions of separation and shear layers for low Knudsen number flows.

METHODS & RESULTS

The basic steps of our DSMC code involve the movement of particles through the computational domain, collisions of these particles with other particles and with surfaces, mapping of the particles onto the collision and sampling meshes, mesh refinement (when required), and sampling of the macro-parameters in each Cartesian collisional cell. Depending on the gas-wall surface interaction model, the particle is reflected specularly, diffusively, or mixed specularly and diffusively. After the movement of all the computational particles in the domain, each particle is checked to see if it is still on the same processor, has been transferred to another processor, or has exited the computational domain.

For the computational particles that are transferred to another processor, each processor creates linked lists to store the corresponding particle lists that need to be transferred. Once each processor has the information regarding which processor it must send to or receive computational particles from, communication is accomplished in a point-to-point fashion via MPI. After each processor has the list of its particles, these particles are put in the basic Cartesian cells overlying the computational domain. These Cartesian cells are refined based on criteria such as a required number of particles per cell and that the cell size is smaller than the local, spatially varying, mean free path. The SUGAR code performs AMR on two octree meshes: the collision mesh and the sampling mesh. During the collision phase, particles are selected to participate in a collision from collision cells using the appropriate physical cross section (e.g., elastic, inelastic, reaction) and their post-collision velocities are updated based on the statistical outcome. The octree grid is destroyed and recreated until it meets the refinement criteria and the flow reaches steady state.

To date we have modeled shock-dominated flows corresponding to flows over a sphere and a wedge. In figs. 1–2, the nature of the oblique and bow shocks can be observed, with the highest temperatures observed downstream of the bow shock. Three-dimensional effects were important in allowing for pressure relief at the edges of the wedge, as seen by the streamlines falling over the wedge's sides. The level of octree refinement was four in the recirculation region that formed at the intersection of the two wedges.

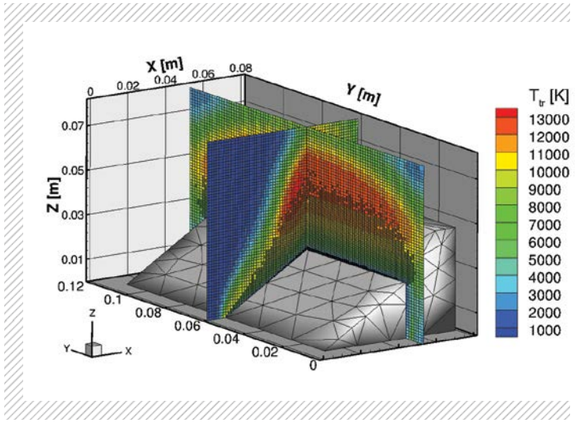

FIGURE 2: Maximum temperature (T_{tr}) decreases along the span-wise direction due to pressure-relief effects.

WHY BLUE WATERS

The starting point for modeling Edney type IV shock interactions with DSMC involves 7 billion particles (approximately 3 per collision cell) and 200 GB of physical memory. With the Blue Waters petascale facility, we have the ability to use thousands of cores to reduce the time taken for modeling a flow from a few weeks to a few days. We expect that we will generate approximately 1 TB of data per case. In order to improve the scaling performance of our new code, we are using the CPMAT profiler specifically built for CRAY systems for quick visualization.

PUBLICATIONS

Sawant, S., B. Borkut, O. Tumulku, and D. Levin, Development of an AMR Octree DSMC Approach for Shock Dominated Flows, *AIAA SciTech 2015*, Orlando, Fla., January 5, 2015. [AIAA Paper No.-2015-0070].

PARTICLE-IN-CELL SIMULATIONS OF KINETIC EFFECTS IN HED PLASMAS

Allocation: NSF PRAC/3.1 Mnh
PI: Warren B. Mori[1]
Co-PI: Frank S. Tsung[1]
Collaborator: Luis O. Silva[2]

[1]UCLA
[2]Instituto Superior Técnico, Lisbon, Portugal

EXECUTIVE SUMMARY:

Using a suite of particle-in-cell codes developed within the UCLA Plasma Simulation Group, we have been using Blue Waters to address important questions in high energy density (HED) plasmas, which are critical to the success of experiments at the SLAC National Accelerator Laboratory and the National Ignition Facility (NIF). Access to Blue Waters has allowed us to perform very large simulations in a timely manner and has provided key insights into ongoing experiments.

INTRODUCTION

The UCLA Simulation of Plasmas Group has been using particle-in-cell (PIC) simulations on parallel computers for nearly 30 years to study basic plasma science and to answer compelling science questions. This effort has involved unraveling complicated nonlinear plasma science, attempting to solve science problems of national importance, developing a robust suite of parallel PIC tools that run efficiently on a wide range of platforms, including the largest computers in the world, developing novel reduced PIC models, benchmarking these codes against theory and experiment, developing data visualization and analysis tools, and keeping abreast of relevant applied math and applied computer science research.

Currently, the research of the group and the OSIRIS consortium is focused on two key science areas, with clearly developed science questions; this research is funded by the Department of Energy (DOE) and the National Nuclear Security Administration (NNSA). The questions are:

1. Can laser plasma instabilities be controlled or even harnessed in inertial fusion plasmas?

2. Can plasma-based acceleration be the basis of new compact accelerators for use at the energy frontier, in medicine, in probing materials, and in novel light sources?

Using Blue Water resources, we have made significant progress in these areas.

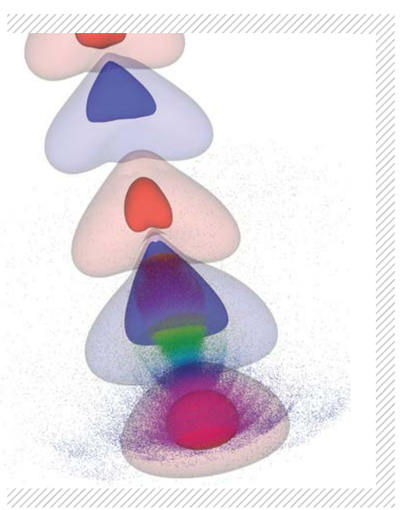

FIGURE 1: Streamlines and QuickPIC simulations of the two bunch experiment that appeared in Nature in November 2014. The figure shows the parallel electric fields and the two electron bunches. The bunch in the front is in a decelerating field and loses energy, and the trailing bunch in the back is in an accelerating field and gains energy. Using this setup, experimentalists at SLAC have both high accelerating gradient and high energy-transfer efficiency using the plasma wakefield accelerator (PWFA) concept.

METHODS & RESULTS

Our research employs the particle-in-cell (PIC) method, where Maxwell's equations are solved on a grid using current and charged densities calculated by weighing particles onto the grid; each particle is pushed to a new position and momentum using relativistically correct equations of motion. The UCLA simulation group is a leader in the development of high-performance PIC codes and was awarded a Software Infrastructure for Sustained Innovation (SI2) grant by NSF.

Using our suite of locally developed codes on Blue Waters, we have published two papers in the journal *Nature* in the past 12 months. In one paper (published in November 2014), a plasma wake-field acceleration (PWFA) with two electron bunches demonstrated high efficiency (~30%) and high-quality electron beams (with energy spread as low as 0.7%) using plasma-based accelerators.

In the second paper (published in August 2015), physicists demonstrated the acceleration of positrons using plasmas and demonstrated 5GeV energy gain using only 1.3 meters of plasma. The ability to transfer energy makes the PWFA scheme very attractive as an energy booster to an electron-positron collider. Simulations performed on Blue Waters were critical in providing key insights in these experiments.

Blue Water is also invaluable in the study of laser plasma interactions in plasma-based accelerators and laser fusion. In a paper published in *Physical Review Letters* in July 2014 we have shown that, using a low-intensity pre-pulse to build up the plasma response in front of the main pulse, the leading edge of the main pulse does not diffract and facilitates the main pulse to reach a self-guided state that remains stable for more than 10 Rayleigh lengths. These simulations use more than 150 million grids and 300 million particles for close to 2cm of plasmas. Simulations of this scale can only be performed on a Track-1 supercomputer.

Lastly, using Blue Water we have performed 3D simulations that gave insights in understanding the formation of relativistic electron rings in laser-wakefield accelerators (LWFAs). 3D simulations on Blue Waters showed that electrons trapped in the second wave bucket become defocused when it moves into the decelerating region of the second bucket. As it becomes more defocused, it moves into a small stable region behind the first bucket, which traps these electrons and forms the rings observed in the Lawrence Livermore National Laboratory experiments. The process is highly nonlinear, and insights gained from the 3D simulations performed on Blue Waters were critical in the understanding of, and ultimately the control of, ring formations in future LWFA experiments.

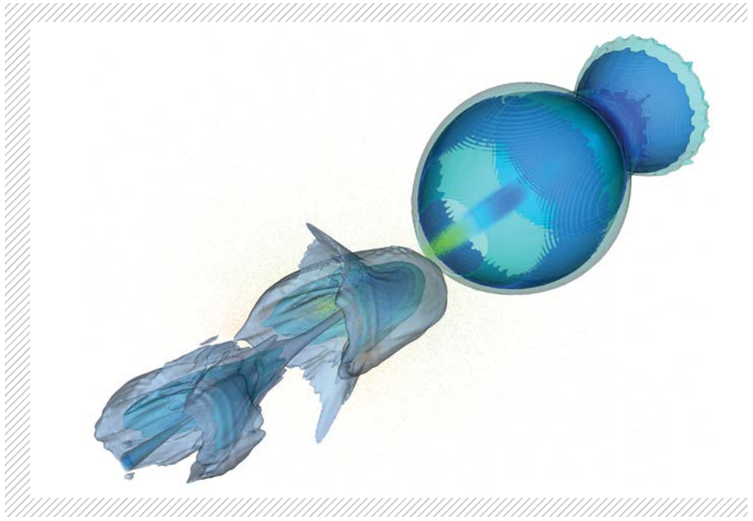

FIGURE 2: A future particle collider based on plasma wakes will need high-gradient, high-efficiency acceleration of both electrons and positrons, as seen in this computer simulation. Now positron acceleration by using wakes produced by a positron beam as it propagates through a plasma has been demonstrated. This figure shows wakefields generated by electrons (right) and positrons (left).

WHY BLUE WATERS

Having Blue Waters access has allowed us to make quantitative comparisons between simulations and experiments (like those published in *Nature*) and allowed us to perform very large simulations that cannot be done elsewhere.

PUBLICATIONS

Litos, M., et al., High-efficiency acceleration of an electron beam in a plasma wakefield accelerator, *Nature* 515:92 (2014)

Tzoufras, M.; F.S. Tsung, W.B. Mori, A.A. Sahai, Improving the Self-Guiding of an Ultraintense Laser by Tailoring Its Longitudinal Profile, *Physical Review Letters* 113:245001 (2014).

Corde, S., et al, Multi-gigaelectronvolt acceleration of positrons in a self-loaded plasma wakefield, *Nature* 524:442 (2015)

Pollock, B.B., et al, Formation of Ultrarelativistic Electron Rings from a Laser-Wakefield Accelerator, *Physical Review Letters* Volume 115, Issue 5, Article Number: 055004 (2015) (OCI-1036224)

MECHANICS OF RANDOM AND FRACTAL MEDIA

Allocation: Illinois/50 Knh
PI: Martin Ostoja-Starzewski[1]
Collaborator: Seid Koric[1,2]

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

EXECUTIVE SUMMARY:

One of the grand challenges to engineering (as identified by the National Academy of Engineering) is America's aging and failing infrastructure. A major barrier to deciding if a structure is safe or needs to be repaired or replaced is an accurate assessment of the damage state of a given load-bearing material or structure. Complex damage (crack and defect) patterns in natural and man-made materials often lead to consideration of their fractal and disordered structures. To properly understand the damage state of load-bearing materials and structures and assess the associated failure hazard, the multiscale geometry of existing damage has to be considered.

A challenge in health science that requires similar methodology is the modeling and simulation of mild traumatic human brain injury. Given a brain's complex structure, the key idea is to run computer models based on the MRI-resolved 3D images of human heads.

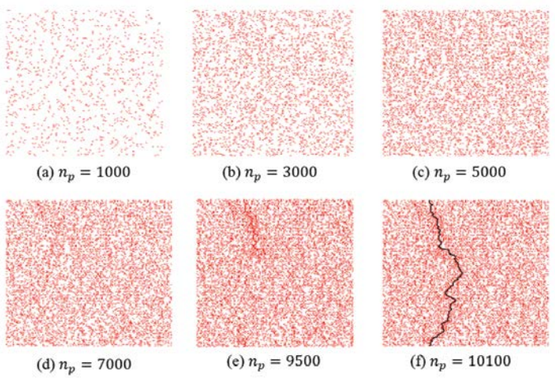

FIGURE 1: Failed springs at different loading stages. The final crack leading to macroscopic failure is shown as black line in (e). Localization of damage only took place as the peak load was reached.

INTRODUCTION

Conventional continuum solid mechanics hinges on the assumption of homogeneous material properties. However, almost all natural (inanimate and biological) and man-made materials contain some microstructure(s) and exhibit spatially distributed randomness (material defects, impurities, etc.) and even fractal patterns. As a result, in elasto-plastic materials under macroscopically uniform and monotonic loading boundary conditions, 'weaker' material grains and regions plasticize first and gradually spread in a cooperative fashion throughout the whole body. In elasto-brittle materials, nano- and microscale cracks form and then coalesce into mesoscale cracks, then grow and result in global failure of the material and structure. If we consider biomechanics of the human head under impact loading conditions, the stress waves (and possible localized damage) represent the true response of the entire brain patterns. Problems of this type can be run to greater resolution and more efficiently on Blue Waters than on any other computer.

METHODS & RESULTS

1. Scale-dependent homogenization [1] was used to estimate the effective elastic shear moduli (or effective conductivity) of 2D two-phase materials with Gaussian correlated isotropic microstructures. The microstructural randomness inherent in these materials motivates a study of homogenization from a statistical volume element to a representative volume element. This finite-size scaling, studied in terms of material responses under two uniform boundary conditions (displacement and traction), obtained at a wide range of mesoscales for a wide range of contrasts in elastic moduli of both phases. The mesoscale is defined as the ratio of domain size to the grain size. A study involving Monte Carlo sampling of many realizations of spatially correlated microstructures, each solved by finite elements, led to scale-dependent bounds and a universal normalized scaling function.

2. Large-scale lattice model simulations are computationally demanding as the system of linear equations is solved after each lattice bond fails or yields (fig. 1). The number of bonds at final failure (n_f) for elastic brittle transition scales with L as $n_f \sim O(L^{1.8})$ in 2D and $n_f \sim O(L^{2.7})$ in 3D [2–4]. Moreover, as the system approached macroscopic failure there was a critical slowdown of iterative solvers. The computational cost further increased because a significant number of realizations were required to obtain good statistics. To this end, we used algorithms involving conjugate gradient and massive parallelization to implement the elastic plastic brittle 2D and 3D code.

3. The mechanisms underlying mild traumatic brain injury need to be understood from the standpoint of stress wave patterns taking place in the brain. Fine spatial meshes are needed to properly describe such patterns and resolve the highly heterogeneous brain structure (with a fractal brain surface) [5]. This was accomplished with an MRI-based computational model, previously validated by tagged MRI and a harmonic phase imaging analysis technique on *in-vivo* human brain deformation data [6]. Computer simulations were carried out for various impacts with the cerebrospinal fluid layer explicitly modeled as a viscous or viscoelastic fluid (fig. 2).

WHY BLUE WATERS?

The objective was to introduce spatial material heterogeneity and disorder/randomness into realistic large-scale models of conductive and elastic plastic brittle man-made and natural (i.e. biological) materials, and study the growth of multiscale/fractal material systems in the presence of evolving and interacting defects/cracks or the ensuing wave patterns. In general, we needed to be able to handle very large domains in 2D and then extend the simulations to 3D. Given major computational challenges (number of degrees of freedom and, hence, memory and CPU requirements), this research can only be done on Blue Waters. At a subsequent stage, our large-scale simulations will be extended to other scenarios: anisotropic yield criteria, 2D versus 3D, coupled field (e.g., thermo-mechanical, electro-magneto-mechanical) fields, and head models with more anatomical detail.

PUBLICATIONS

Kale, S., A. Saharan, S. Koric, and M. Ostoja-Starzewski, Scaling and bounds in thermal conductivity of planar Gaussian correlated microstructures. *J. App. Phys.*, 117 (2015), 104301, doi:10.1063/1.4914128.

Kale, S., S. Koric, and M. Ostoja-Starzewski, Stochastic continuum damage mechanics using spring lattice models. *Second Int. Conf. on Damage Mechanics*, Troyes, France, July 8–11, 2015.

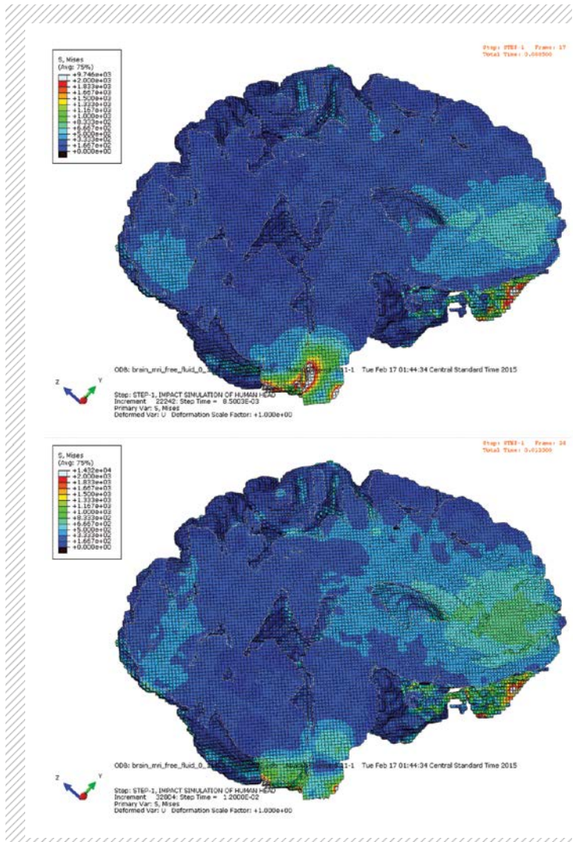

FIGURE 2: Two frames of the evolution of the von Mises stress distribution in the human brain (mid-sagittal view) during a 15 ms frontal impact (forehead) simulation; the cerebrospinal fluid layer was explicitly modeled as a Stokesian fluid.

HIGH-RESOLUTION 3D SIMULATIONS OF CORE-COLLAPSE SUPERNOVAE

Allocation: NSF PRAC/8.94 Mnh
PIs: Christian D. Ott[1], Peter Diener[2]
Co-PI: Mark Scheel[1]
Collaborators: Luke F. Roberts[1], David Radice[1], Philipp Mösta[3], Erik Schnetter[4], Roland Haas[5]

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

FIGURE 1 (BACKGROUND): Volume rendering of the specific entropy of a developing asymmetric explosion in a 3D core-collapse supernova. The volume shown is 5003 km3. Yellow contours are regions of very high entropy corresponding to outward moving bubbles/filaments of neutrino-driven turbulent convection. The bright sphere at the center is the proto-neutron star core. The time shown is 360 ms after core bounce and protoneutron star formation. The maximum shock radius at this time is 410 km.

EXECUTIVE SUMMARY:

Core-collapse supernovae are turbulent and dramatic events. Neutrinos deposit energy in the region behind the supernova shock wave and drive buoyant turbulent convection at physical Reynold's numbers up to 10^{17}. Neutrino-driven turbulence is anisotropic, mildly compressible, and only quasi-stationary. Blue Waters-enabled simulations have shown that turbulence provides an effective pressure that can amount to 50% of the total pressure behind the supernova shock wave and thus plays a crucial role in driving the explosion. Because of its importance in the explosion mechanism, it is essential to gain a full qualitative and quantitative understanding of core-collapse supernova turbulence. The results of our high-resolution reduced-physics simulations showed, perhaps surprisingly, that core-collapse supernova turbulence behaved like Kolmogorov turbulence at high resolution. Our simulations also indicated that all current simulations severely under-resolve turbulence, artificially trapping turbulent kinetic energy at large scales where it is most beneficial for launching an explosion.

INTRODUCTION

Core-collapse supernovae (CCSNe) mark the explosive ends of the lives of massive stars with mass greater than 8–10 times the mass of our Sun. There is about one supernova explosion per second in the universe. Time-domain astronomers now discover multiple CCSNe per day and there is abundant evidence linking these explosions to massive stars, yet we do not yet fully understand how the initial collapse of a massive star's core turns into an explosion. The collapse proceeds until a proto-neutron star forms and a shock wave launches. This shock wave, however, does not immediately explode the star. It stalls due to energy losses to nuclear dissociation of heavy nuclei and to neutrinos that stream away from behind the shock.

Neutrinos dominate the energetics of the CCSN problem and it has been hypothesized and shown in some simulations that the re-absorption of a small fraction (~10%) of the outgoing neutrino luminosity can help revive the stalled shock and drive an explosion. Over the last two decades it has become clear that multi-dimensional dynamics, in particular neutrino-driven turbulent convection, plays a crucial role in the explosion mechanism. More recently, our group and the Princeton group led by Prof. Adam Burrows have shown that turbulence is key to the explosion mechanism.

Simulating CCSNe is an extremely computationally intensive, but highly appealing research problem. All four fundamental forces of physics (gravity, electromagnetic, weak, and strong forces) go into CCSNe simulations, which makes them truly multi-physics. They also must be multi-dimensional (ideally 3D, since turbulence is unphysical in 2D) and multi-scale, since small-scale turbulent regions must be resolved while the entire inner regions O(10,000) km of the star must be contained on the grid.

METHODS & RESULTS

We carried out our full-physics global 3D CCSN simulations with our code *Zelmani* that is based on the open-source Einstein Toolkit [1] and the Cactus [2] framework. *Zelmani* is fully general relativistic (GR) with dynamical spacetime evolution, employs full Berger–Oliger adaptive mesh refinement (AMR), and implements GR hydrodynamics with finite-volume methods. Neutrinos are treated either with an efficient but very approximate leakage scheme or with two-moment energy-dependent radiation transport with an analytic closure relation (M1). In this M1 approach, the fluxes are handled time explicitly and only local interaction terms are treated time implicitly. This makes the code efficient and lets full adaptive mesh refinement (AMR) simulations scale well to 16,000–32,000 Blue Waters cores for typical problem sizes.

Fig. 1 shows a snapshot of our most recent 3D GR radiation-hydrodynamics CCSN simulations on Blue Waters while it was in transition to a strongly asymmetric explosion. It is to date the most highly resolved CCSN simulation with 3D radiation hydrodynamics. The entire region behind the shock was resolved with linear resolution of 1.5 km, corresponding to an angular resolution of 0.43° at a radius of 200 km, which is a factor of about five higher than previous work. We carried out a second, lower-resolution simulation with everything else fixed. This simulation transitioned to explosion about 50 ms earlier than the high-resolution simulation, indicating that low resolution indeed artificially favored explosion as we suggested previously with our 3D GR leakage simulations [3].

In order to study the resolution dependence of CCSN turbulence in more detail, we carried out reduced-physics semi-global simulations (in a 3D wedge) with the high-order GR hydrodynamics code WhiskyTHC. We reduced the physics to the bare minimum to capture qualitatively all major ingredients to the CCSN problem. With this setup, we were able to scale our simulations to 65,536 cores and follow neutrino-driven convection at an unprecedented 0.09° angular and 191 m radial resolution, more than 15 times higher resolution than the highest-resolution published simulation. Fig. 2 shows a key result of this study. As resolution increased, the flow became increasingly fine grained and large-scale plumes broke into much finer filaments and smaller plumes. Our very highest resolution resolved substantial turbulent inertial range and the turbulent cascade exhibited a Kolmogorov-type scaling. This was an important result because it had been unclear previously if CCSN turbulence should be Kolmogorov in nature or not. We also found that lower-resolution simulations artificially trapped turbulent kinetic energy at large scales due to a numerical "bottleneck" in the turbulent cascade.

WHY BLUE WATERS?

Blue Waters was essential for our CCSN simulation work; it allowed us to push our simulations to unprecedented resolution and physical detail. No other machine would allow us to carry out simulations at this scale and rate of throughput. The Blue Waters/NCSA team facilitated this throughput that resulted in rapid turn-around of science results. Future Track-1 systems would allow even higher-resolution, more physically detailed, and longer-duration CCSN simulations to more completely track the development of the explosion, explosive nucleosynthesis, and neutron star recoil. However, architectural changes between Blue Waters and a future Track-1 system will require a new approach to parallel computing. We have already started to develop a next-generation code that will be designed specifically to perform on next-generation architectures.

PUBLICATIONS

Abdikamalov, E., et al., Neutrino-driven Turbulent Convection and Standing Accretion Shock Instability in Three-dimensional Core-collapse Supernovae. *Astrophys. J.*, 808:1 (2015), 70, doi:10.1088/0004-637X/808/1/70.

Couch, S. M., and C. D. Ott, The Role of Turbulence in Neutrino-driven Core-collapse Supernova Explosions. *Astrophys. J.*, 799:1 (2015), 5, doi:10.1088/0004-637X/799/1/5.

Radice, D., S. M. Couch, and C. D. Ott, Implicit large eddy simulations of anisotropic weakly compressible turbulence with application to core-collapse supernovae. *Comput. Astrophys. Cosmol.*, 2 (2015), 7, doi:10.1186/s40668-015-0011-0.

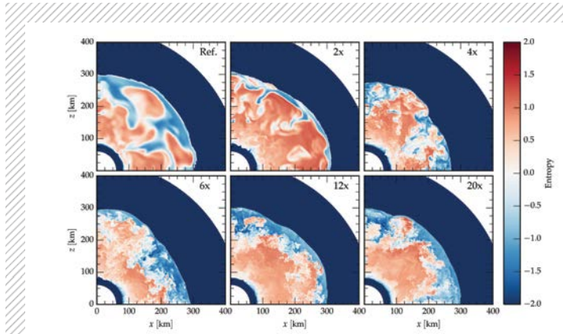

FIGURE 2: Entropy distribution (arbitrary units) in 2D slices of 3D semi-global simulations of neutrino-driven turbulent convection in core-collapse supernovae. The top left panel is the reference resolution and the remaining panels are snapshots from simulations with increased resolution as denoted in the top right of each panel. The flow structure changed dramatically with increasing resolution.

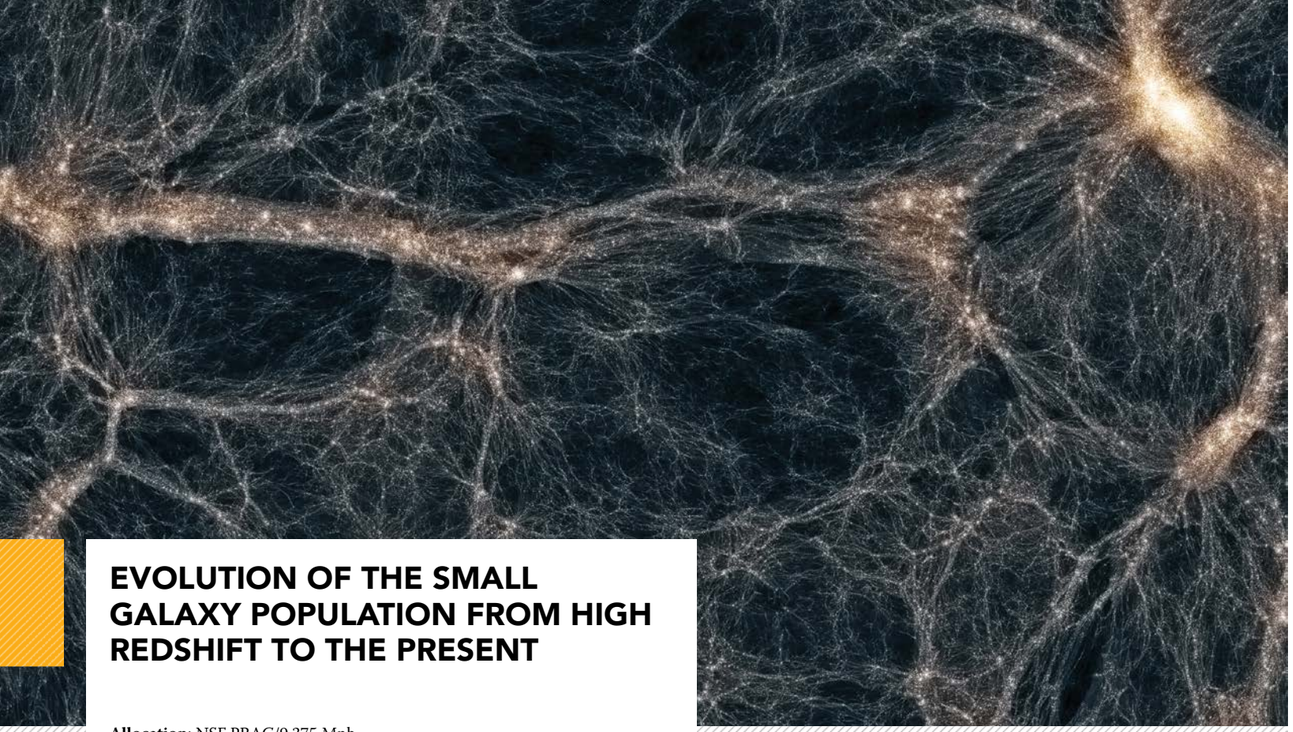

EVOLUTION OF THE SMALL GALAXY POPULATION FROM HIGH REDSHIFT TO THE PRESENT

Allocation: NSF PRAC/9.375 Mnh
PI: Thomas Quinn[1]
Co-PI: Fabio Governato[1]
Collaborators: Lauren Anderson[1], Michael Tremmel[1], Charlotte Christensen[3], Sanjay Kale[4], Lucasz Weslowski[4], Harshitha Menon[4], Gengbin Zheng[5]

[1]University of Washington
[2]University of Arizona
[3]Grinnell College
[4]University of Illinois at Urbana-Champaign
[5]National Center for Supercomputing Applications

FIGURE 1 (BACKGROUND): A slice of a 24-billion-particle simulation of the present-day dark matter in a 25 Mpc volume of the universe. The object in the upper right is roughly the size of our local group of galaxies and is resolved with several hundred million particles.

EXECUTIVE SUMMARY:

Creating robust models of the formation and 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 do this 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 million parsecs cubed volume of the universe with a physically motivated star formation/supernovae feedback model. This past year's accomplishments include using a pathfinding simulation at one-tenth the needed resolution, which we will use to study overall star formation histories and luminosity functions. This simulation is also being used to tune the feedback model to accurately model galaxy morphologies over a range of masses.

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 or less in order 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 do 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 is the overall star formation history of 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.

Preliminary results from our low-resolution simulations indicated that we can reproduce the high-redshift galaxy luminosity functions observed by HST. 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, there may be enough stellar radiation from these low-mass galaxies to completely reionize the intergalactic medium.

WHY BLUE WATERS?

The mass and spatial resolution required to reliably model galaxy morphology were set by our published resolution tests. Therefore the size of the simulation we perform was 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 on a side. This volume will not only allow us to make direct comparisons with HST surveys, but also enhance their scientific return by understanding how those surveyed galaxies will evolve to the present.

This volume and our required resolution result in a total of 24 billion particles. This is just over an order of magnitude larger than simulations that we could run on other resources to which we have access. A sustained petascale facility like Blue Waters was required for this simulation.

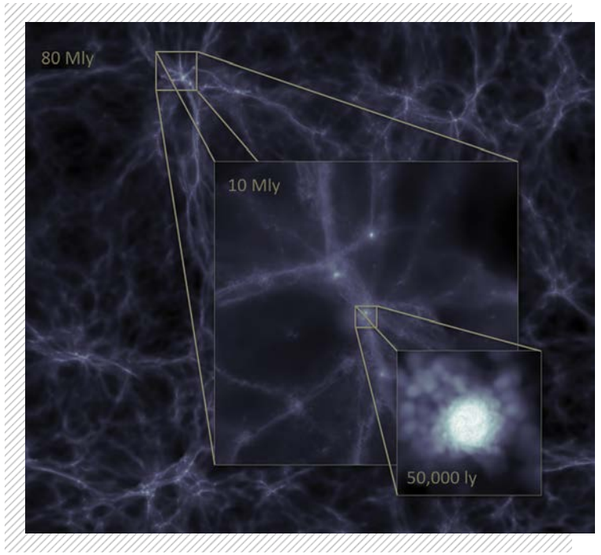

FIGURE 2: A zoom-in on one of the galaxies in a gas dynamical simulation of a 25 Mpc volume of the universe. The zoom-in shows how the disk morphology of an individual galaxy is resolved within this volume.

Nevertheless, this simulation was still a compromise. For example, if we wish to understand how high-redshift galaxies are influencing the surrounding intergalactic gas, a much larger volume is needed. Intergalactic gas is studied in absorption from observations of background quasars using the HST Cosmic Origins Spectrograph. Statistical samples of this gas require a volume of order 60 Mpc on a side, over an order of magnitude larger than our current simulation. Only with this size of simulation will we 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 for this simulation.

PUBLICATIONS

Governato, F., et al., Faint dwarfs as a test of DM models: WDM versus CDM. *Mon. Not. R. Astron. Soc.,* 448 (2015), 792, doi:10.1093/mnras/stu2720.

Menon, H., et al., Adaptive techniques for clustered N-body cosmological simulations. *Comput. Astrophys. Cosmol.,* 2 (2015), 1, doi:10.1186/s40668-015-0007-9.

THEORETICAL SPECTROSCOPY FOR TRANSPARENT CONDUCTING OXIDES

Allocation: Blue Waters Professor/0.245 Mnh
PI: André Schleife[1]
Collaborator: Joel Varley[2]

[1]University of Illinois at Urbana–Champaign
[2]Lawrence Livermore National Laboratory

EXECUTIVE SUMMARY:

Understanding the influence of dielectric screening on the electron-hole interaction and the resulting excitonic effects is a long-standing problem in computational materials science. Here, we achieve this for two transparent conducting oxides: In_2O_3 and Ga_2O_3. Both materials are very interesting in the context of transparent electronics and the semiconductor industry since they conduct electrical current while transparent. The unique combination of high-performance CPUs, large-memory configurations, and a large and fast storage system in Blue Waters allowed us to gain unprecedented insight into the physics of excitons and its impact on optical properties of these materials. We computed the most accurate theoretical optical spectra currently available for both materials and will use our results to further develop the computational framework to facilitate sustainable computational materials science and highly accurate theoretical spectroscopy for modern, complex materials that drive societal progress.

INTRODUCTION

Materials that are transparent and conducting are highly desirable for transparent electronics, photovoltaics, and optoelectronics. In this work we investigated indium oxide, In_2O_3, and gallium oxide, Ga_2O_3, since they are wide-bandgap semiconductors that show a remarkable combination of material properties. They are transparent across the visible spectrum and their electrical conductivity can be controlled over a large range, which enables a variety of novel device applications. In addition, the constituent elements are environmentally benign. Because of their properties, these oxides (amongst others) have been widely adopted as transparent conducting layers and they attract attention as stand-alone semiconductors in solar-blind photodiodes and Schottky diodes (e.g., for high-power electronics). Despite their appeal and widespread use, the influence of the electronic screening on the electron-hole interaction and the optical absorption is particularly poorly understood. This is difficult to measure directly in experiment, and the large unit cell size of In_2O_3 (bixbyite) and β-Ga_2O_3 (monoclinic) with up to 40 atoms has hampered theoretical studies, making the use of HPC necessary. Answering this question is not only important for device applications and the semiconductor industry, but it also provides insight for ongoing development of the computational framework itself.

METHODS & RESULTS

Our research is based on cutting-edge first-principles calculations within many-body perturbation theory. We computed optical absorption spectra using the Bethe–Salpeter equation framework based on single-particle energies obtained from a hybrid functional to approximate exchange and correlation. The bandgap from a hybrid-functional calculation was reproduced by a semi-local approximation to exchange and correlation and a scissor shift. This approach can meet the challenging convergence criteria necessary for optical absorption spectra.

Our theoretical spectroscopy results constitute the most accurate theoretical optical spectra available for these materials. Excitonic effects were explicitly included via the Bethe–Salpeter approach and for the entire spectral range studied here we found a pronounced influence of excitons. The strength of excitonic effects as a function of the photon energy will be explored in a follow-up project where we will extend our results to higher energies to compare with highly accurate experiments up to 40 eV. Our present work also advanced the qualitative understanding of optical anisotropy in monoclinic β-Ga_2O_3 both near the absorption edge and at high photon energies. This understanding is necessary for applications of these materials as transparent conducting oxides. From our present and future work in this context we learn how to model optical properties of technologically relevant materials (with societally important applications in photovoltaics, transparent electronics, and plasmonics) with very high accuracy using efficient parallelization and HPC.

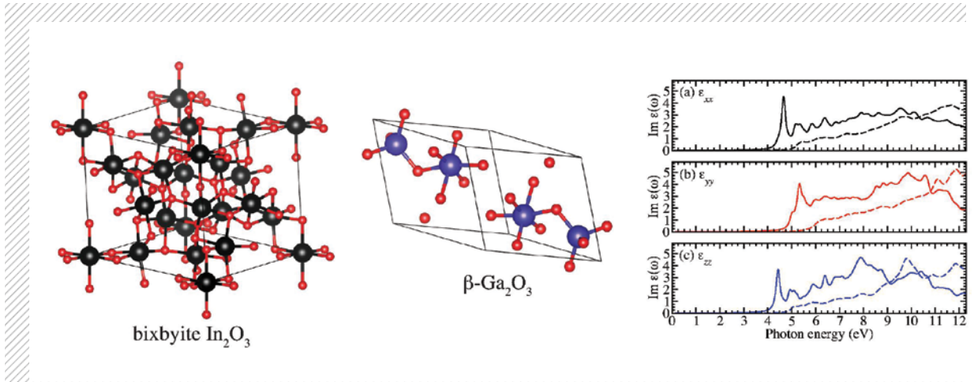

FIGURE 1: *Left and center:* Unit cells of In_2O_3 and Ga_2O_3 (In atoms: black, oxygen atoms: red, Ga atoms: blue) used in the present work. *Right:* Imaginary part of the anisotropic dielectric function (related to probability of photon absorption) of Ga_2O_3 (solid curves: excitonic and local-field effects included; dashed curves are computed without those) [1].

WHY BLUE WATERS?

For the present project, it was necessary to compute and analyze very large exciton Hamiltonian matrices. To study the optical properties of a single material in this work, we computed multiple matrices with ranks between 350k and 378k, leading to memory and storage demands between 0.5 TB and 1 TB for each individual matrix. Computing the matrix is a computationally demanding task that we accomplished using a well-parallelized Bethe–Salpeter code. The code used an ensemble run to divide the work into about 40 to 60 24-hour single-node jobs, each of which was parallelized across all cores of a given node using OpenMP. In this phase of the work, the large and fast storage capabilities of Blue Waters were important: Once the matrix was written to disk we used a time-propagation approach that scales quadratically with the rank of the matrix to compute an optical absorption spectrum. This code was parallelized using MPI; it read the matrix and distributed it into the combined memory of 16 Blue Waters nodes in less than five hours (62 MB/sec on average for the duration of 4.7 hours). The time-propagation scheme was then used to compute the optical absorption spectrum in about 6 hours.

The scenario described above was not a "one-shot" calculation. Our research requires multiple runs for each individual material and, hence, we needed a machine like Blue Waters that allows us to routinely carry out this work for multiple materials and configurations. The staff helped us with runs that needed more than the standard 24-hour wall time. Massive parallelism was not important to this work, but the entire "Blue Waters package" was. We needed high-performance CPUs that were very well integrated into large-memory nodes and connected to an excellent (large and fast) storage system. Blue Waters offered exactly this, along with outstanding availability, quick turnaround times, and a fast and very responsive support team.

We have already begun to extend this work towards much larger nanoscale systems. Those will require dealing with significantly larger matrices that lead to much more demanding requirements for memory, storage, and communication systems. For sustainable computational materials science that benefits society and, specifically, for predictive theoretical-spectroscopy techniques needed for computational design of nanoparticles for bioimaging or charge-separation in 2D materials, the next Track-1 system will be crucial.

PUBLICATIONS

Varley, J. B., and A. Schleife, Bethe–Salpeter calculation of optical-absorption spectra of In_2O_3 and Ga_2O_3. *Semicond. Sci. Tech.*, 30 (2015), 024010, doi:10.1088/0268-1242/30/2/024010.

COMPLEXITIES OF HIGH-REYNOLDS-NUMBER TURBULENCE

Allocation: NSF PRAC/9.0 Mnh
PI: P.K Yeung[1]
Co-PIs: . Majumdar[2], R.D. Moser[3], D. Pekurovsky[2], J.J. Riley[4]
Collaborators: B.L. Sawford[5], K.R. Sreenivasan[6]

[1]Georgia Institute of Technology
[2]San Diego Supercomputing Center
[3]University of Texas at Austin
[4]University of Washington
[5]Monash University, Australia
[6]New York University

EXECUTIVE SUMMARY:

We study the complexities of turbulent fluid flow at high Reynolds number, where resolution of fluctuations over a wide range of scales requires sustained petascale computation. Using 262,144 cores of Blue Waters in a favorable node-topology configuration, we have performed the first $8{,}192^3$ production simulation of isotropic turbulence ($8{,}192^3$ results in over half a trillion cells in the simulation). Both the dissipation rate and entropy, representing deformation and rotation of local fluid elements, take values as large as on the order of 10^5 times the mean at nearly the same locations in space. Regions of most extreme vorticity are not worm-like as commonly thought. These results have important implications for phenomena such as stretching of flame surfaces and preferential concentration of particles of finite inertia. An active algorithmic effort is under way to improve our ability to study the dispersion of tens of millions of fluid particles and other entities in a Lagrangian framework. Turbulent mixing at high Reynolds number will also be addressed.

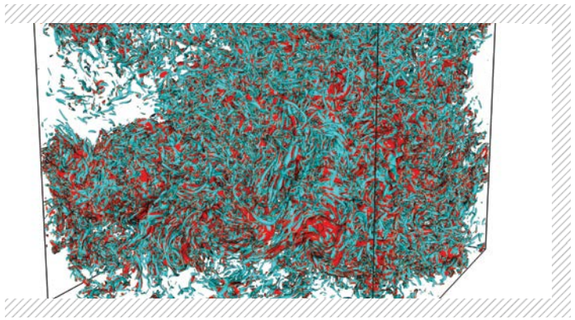

FIGURE 1: Perspective view of 3D contour surfaces at ten times the mean value of dissipation (red) and enstrophy (cyan), in a 768^3 sub-domain extracted from an instantaneous snapshot of a $8{,}192^3$ simulation. The intricate detail of small-scale turbulence is evident; many worm-like vortex filaments can be seen in cyan.

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 that are highly localized in time and space, and to use this understanding to address the effects of fine-scale intermittency [1] in applications. For example, extremely high strain rate can break a flame surface and lead to local extinction, while high rotation rate has a strong influence on the local concentration of inertial particles and vapor droplets in multiphase flows and cloud physics problems. The likelihood and the intensity of large fluctuations increase with the range of scales present, which can be expressed as a positive power of a Reynolds number, defined as the product of turbulence velocity and length scales divided by the kinematic viscosity. The study of extreme fluctuations at the small scales in high-Reynolds-number turbulence is thus of great importance.

A very important effect of turbulence is enhanced mixing and dispersion of fluid elements of distinct properties or other entities such as heat, chemical species, and pollutants carried in the flow. Predictions of rates of mixing between fuel and oxygen in jet engines and the rate of growth of pollutant plumes in the environment depend on understanding of small-scale turbulence, sometimes coupled with other physico-chemical processes. Continuing work will mostly focus on studies of turbulent mixing and dispersion [2,3], including molecular diffusivity and (in part) in a reference frame moving with the flow. In most cases, quality data at high Reynolds number with proper resolution of the small scales in space are necessary for definitive answers [4].

METHODS & RESULTS

The only technical approach capable of truly capturing extreme events localized in both time and space is direct numerical simulation (DNS) 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 codes are communication intensive, the overall performance benefits substantially from remote memory addressing and favorable network topologies. The first production simulation had 8,192 grid points in each direction (over half a trillion total points), at a Reynolds number slightly higher than recent work, and with better scale resolution as well.

The new simulation data have been analyzed from probabilistic, spatial, and temporal viewpoints. Both dissipation and enstrophy (which are quadratic measures of local strain and rotation rates) exhibited values as large as 10^5 times the mean. This was well beyond that usually called intense in the literature, where high-dissipation regions are seen as sheet-like and high-enstrophy regions as filament- or worm-like. Extreme events at this intensity likely have a spatial structure different from those at moderately intense amplitude two to three orders of magnitude lower. Indeed, 3D visualization at various viewing resolutions and thresholds (fig. 1) showed that the most extreme events of high enstrophy appeared not as narrow filaments but instead as chunky, slightly elongated regions wrapped by curved sheets of almost equally intense dissipation (fig. 2). Numerous instantaneous snapshots of velocity fields supported a robust observation that peak values of dissipation and enstrophy were nearly coincident (one or two grid points apart). These extreme events apparently maintained their intensity for about two or three small-eddy time scales. The evolution of instantaneous enstrophy was also studied though an exact transport equation.

Our results to date suggest strongly that conventional descriptions of organized vortical motions should be revisited to truly account for the effects of extreme fluctuations in the present data (which may be only partly captured if the small scales are not as well resolved). Similar considerations for extreme dissipation are also central to much-needed improvements in the modeling of effects of intermittency in turbulent combustion, pollutant transport, and other problems. We note that although extreme vents are inherently rare, they can have a first-order effect on physical phenomena of great concern to society (such as tornadoes, explosions, and extreme weather events).

WHY BLUE WATERS?

In general, a $8{,}192^3$ simulation is almost 16 times as expensive as one at $4{,}096^3$, but it is necessary since otherwise (as verified by numerical tests) we would not be able to reach sufficiently high Reynolds number and sufficient small-scale resolution at the same time. A large allocation on a multi-petaflop computer such as Blue Waters is thus vital. We have also found high-quality and dedicated support through NSF's PRAC program to be essential. In particular, Cray personnel worked on remote memory addressing, and the Blue Waters staff made arrangements for reserved partitions of up to 8,192 nodes with favorable network topology to improve productivity and time to solution. Possible future targets on the next Track-1 system include simulations at similar rigor for turbulent flows subjected to other external influences such as buoyancy, solid-body rotation, or electromagnetic forces.

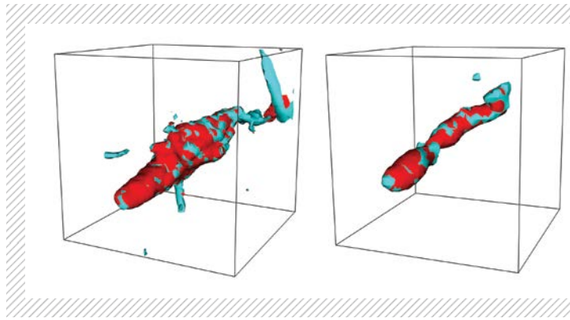

FIGURE 2: Close-ups of sub-region inside the sub-domain in fig. 1, zoomed in to the most extreme values of dissipation (red) and enstrophy (cyan): (left) 51^3 viewing box at a threshold 300 times the mean, and (right) a 31^3 viewing box at a threshold 4,800 times the mean. The global maxima are located near the bottom left in both frames.

PUBLICATIONS

Yeung, P. K., D. Buaria, and B.L. Sawford, Forward and backward relative dispersion at high Reynolds number. *Int. Workshop on Cloud Turbulence,* Nagoya, Japan, March 4–6, 2015.

Iyer, K. P., and P. K. Yeung, Structure functions and applicability of Yaglom's relation in passive-scalar turbulent mixing at low Schmidt numbers with uniform mean gradient. *Phys. Fluids,* 26 (2014), 085107, doi:10.1063/1.4892581.

Yeung, P. K., and K. R. Sreenivasan, Direct numerical simulation of turbulent mixing at very low Schmidt number with a uniform mean gradient. *Phys. Fluids,* 26 (2014), 015107, doi:10.1063/1.4861070.

Yeung, P. K., X. M. Zhai, and K. R. Sreenivasan, Turbulence at high resolution: in- tense events in dissipation, enstrophy and acceleration. *67th Ann. Mtg. Am. Phys. Soc. Division of Fluid Dynamics,* San Francisco, Calif., November 23–25, 2014.

DYNAMICS OF ARGON BUBBLES IN STEEL CONTINUOUS CASTING WITH A MAGNETIC FIELD

Allocation: Illinois/0.20 Mnh
PI: Brian G. Thomas[1]
Co-PIs: Kai Jin[1], Pratap Vanka[1], Hyunjin Yang[1], Matthew Zappulla[1], Seid Koric[2], Ahmed Taha[2]

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

EXECUTIVE SUMMARY:

This project aims to mathematically model multiphase flow in steel continuous casting in order to gain increased understanding and practical insights to improve this important commercial process. Specifically, large eddy simulations of turbulent fluid flow were conducted to investigate the dynamic motion of argon bubbles in the caster with different casting conditions such as electro-magnetic braking, in order to minimize inclusion particle entrapment. This work quantified how the oscillations of the shape and velocity of the rising bubbles can be damped with the application of a static external magnetic field. In addition, transient thermal stress models of the solidifying dendritic microstructure will be applied to investigate strain concentration and the formation of longitudinal cracks in order to understand how to avoid this problem.

INTRODUCTION

Continuous casting (CC) is used to produce more than 95% of steel in the world [1], so even small improvements can have a large impact. In this process, fig. 1 shows how molten steel flows into the mold to solidify a thin shell against the walls that is withdrawn downward at the casting speed to support the liquid pool below the mold. Most defects arise in the mold region due to the entrapment of inclusion particles into the solidifying shell and crack formation in the newly solidified steel shell.

To improve steel products, one must understand the mechanisms of defect formation and find windows of safe operation. The harsh environment makes experiments difficult, so computer simulations are an important tool. Argon gas is often injected and affects defect formation in several ways, including the attachment of inclusions to the surface of rising argon gas bubbles, so it is important to understand the complex motion and dynamics of the bubbles. The shape and motion of the bubbles are modified by applying external magnetic fields. To simulate the complex motion of argon gas bubbles rising in turbulent molten steel and their interaction with inclusion particles and external magnetic fields requires advanced computational models and computing capabilities.

METHODS & RESULTS

To better understand the behavior of argon bubbles and their interaction with inclusion particles during steel continuous casting, the motion of a single argon gas bubble rising in quiescent liquid steel under an external magnetic field was studied numerically using a volume-of-fluid (VOF) method implemented into the finite-difference fluid-flow program CUFLOW [2]. An improved algorithm for surface tension modeling, originally proposed by Wang and Tong [3], was applied to mitigate the spurious velocities generated in numerical simulation of multiphase flows with large density differences.

The computational domain of $6d \times 6d$ (section) x $16d$ (long) contained 192 x 192 x 512 (about 19 million) cells (fig. 2a). A spherical argon bubble of diameter d was initially placed at the center of the container bottom. A uniform magnetic field was applied in the x direction. The dimensionless shape and velocity were tracked with dimensionless time $t^* = t\sqrt{g/d}$.

Fig. 2 shows that rise velocity was smooth and non-oscillatory at early stages ($t^*<0.5$), especially with small bubbles. Without a magnetic field, velocity increased to 2.5 and then decreased slightly due to the inclined motion of the 3 mm bubble. Applying a transverse magnetic field of B=0.2 T lowered the rise velocities by 4%. Increasing the field strength to 0.5 decreased the rise velocities by 24% to 1.83. For a 7 mm bubble with no magnetic field, the rise velocity became oscillatory after $t^* > 1.0$ due to the varying drag force as the bubble shape expanded and contracted along different axes. With a magnetic field of 0.2 T, the oscillations eventually were dampened. Increasing the field to 0.5 T completely damped the oscillations, resulting in a steady rise velocity that was ~25% lower relative to no field.

Fig. 3 shows isosurfaces of z vorticity $\omega z^* = \pm 1$ at $t^* = 9$ for the 7 mm bubble. The alternating patterns were more complex than those behind a 3 mm bubble. With a 0.2 T field, the isosurfaces elongated in the magnetic field (x) direction because the y and z velocities in the surrounding fluid were reduced, so the flow perpendicular to the magnetic field was dampened. The magnetic-field damping effect was studied previously [5] in a driven cavity. Increasing the field strength to 0.5 T made the complex wake structure almost disappear. The front isosurface was wider (in x) near the bubble, and bundled below, where the wake spread (in y).

In related work, large eddy simulations (LES) were applied to investigate the flow in a commercial caster. Figs. 4a and 4c show instantaneous velocity magnitude in top views

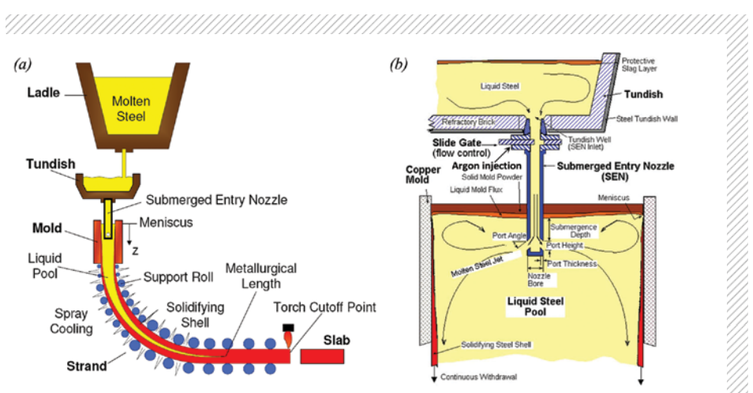

FIGURE 1: (a) Schematic of the continuous-casting process and (b) close-up view into the mold region of the caster. Argon gas is injected near the slide gate.

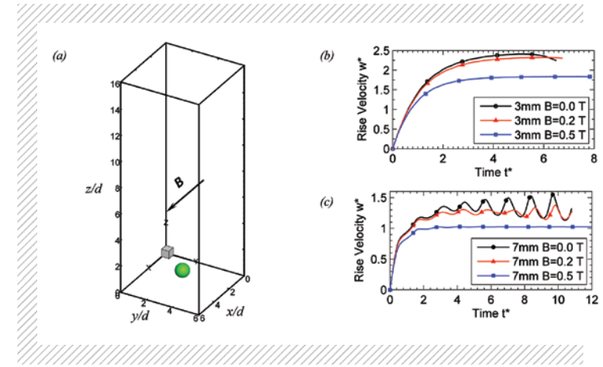

FIGURE 2: (a) Computational domain and initial bubble location, (b) rise velocity of 3 mm argon bubble, and (c) rise velocity of 7 mm argon bubble. [4]

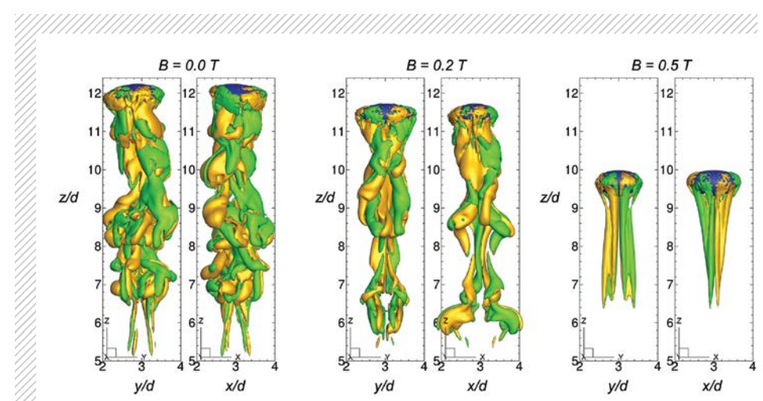

FIGURE 3: Front and side views of the bubble (blue), isosurfaces of $\omega_z^* = 1$ (yellow) and $\omega_z^* = -1$ (green) at $t^* = 9$ for 7 mm bubble. [4]

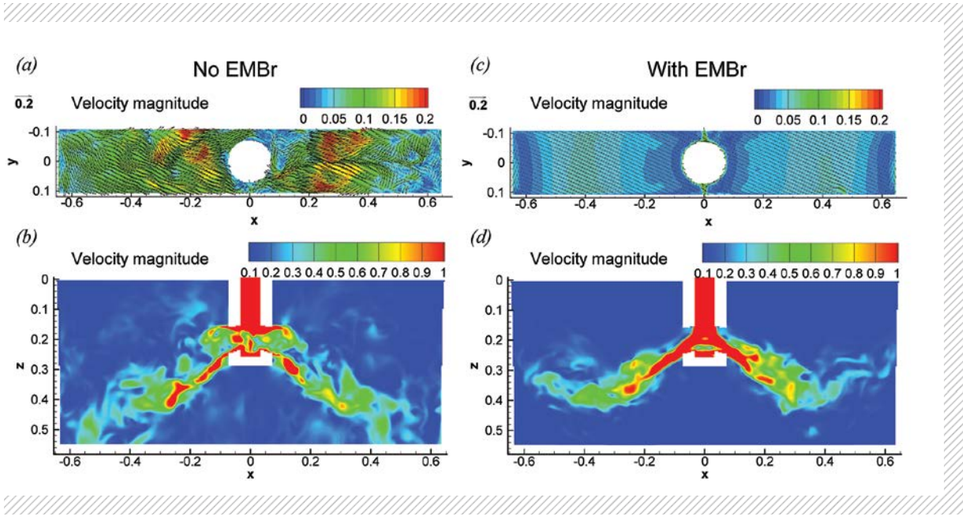

FIGURE 4: Predicted velocity in a horizontal plane near the top surface (top) and in the middle vertical plane of the caster mold (bottom) for cases without (left) and with (right) an external magnetic field.

of the caster (casting speed 1.5 m/min). Applying an EMBr field greatly lowered the velocity near the top surface, which can lead to problems of meniscus freezing in which the steel around the three-phase perimeter solidifies into detrimental hook structures that capture inclusions and bubbles. Figs. 4b and 4d show velocity in side views. Applying EMBr deflected the jets upward, but the flow in the top region was reduced by the dampening effect of the field. The steel plant that was experiencing this problem improved quality by lowering the field strength towards the top of the mold and increasing the casting speed.

WHY BLUE WATERS?

The in-house multi-GPU code CUFLOW was developed and tested on Blue Waters XK nodes and good speed-up was obtained. Less than two days were required for a 30 s LES simulation of flow in a caster domain with 14.1 million cells (based on a 100-time-step test run with an average time step size Δt=0.0005 s). Preliminary results showed that ANSYS-FLUENT also scales well on Blue Waters for this problem. To resolve turbulent flow in the real caster, complete with thousands of bubbles, is only feasible with petascale computing such as Blue Waters.

As another part of this project, transient thermal-stress models of the solidifying dendritic microstructure will be applied to investigate strain concentration and the formation of longitudinal cracks in order to understand how to avoid cracks.

PUBLICATIONS

Jin K., S. P. Vanka, and B. G. Thomas, Three-dimensional flow in a driven cavity subjected to an external magnetic field. *J. Fluids Engineer.*, 137 (2015), 071104, doi:10.1115/1.4029731.

LATTICE QCD ON BLUE WATERS

Allocation: NSF PRAC/30 Mnh
PI: Robert L. Sugar[1]
Collaborators: Alexei Bazavov[2], Mike Clark[3], Carleton DeTar[4], Daping Du[5], Robert Edwards[6], Justin Foley[3], Steven Gottlieb[7], Balint Joo[6], Kostas Orginos[8], Thomas Primer[9], David Richards[6], Doug Toussaint[9], Mathias Wagner[7], Frank Winter[6]

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

EXECUTIVE SUMMARY:

Project goals include developing highly optimized code for the study of quantum chromodynamics (QCD) to carry out calculations that will have a major impact on high-energy and nuclear physics. We have optimized and used Chroma for the simulation of Clover quarks and MILC for the simulation of HISQ quarks. Our long-term objectives with HISQ quarks 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 make precise tests of the standard model. The objective of our Clover fermion program is the determination of 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 subatomic physics. It has successfully explained a wealth of data from accelerator and cosmic ray experiments over the past 40 years. However, it has been difficult to extract many of the most interesting predictions of quantum chromodynamics (QCD), those that depend on the strong coupling regime of the theory. The only way, 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, determine a number of the fundamental parameters of the standard model, and make precise tests of the standard model. Despite the successes of the standard model, high-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, the theory of the strong interactions of subatomic physics, to the precision needed to support large experimental programs in high-energy and nuclear physics. We are using two formulations of lattice quarks. The highly improved staggered quarks (HISQ) formulation is being used to calculate fundamental parameters of the standard model, our current set of theories of subatomic physics, and to make precise tests of the standard model. 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 Cabibbo–Kobayashi–Maskawa (CKM) matrix, which are the weak interaction transition couplings between quarks. The CKM matrix elements and the quark masses are fundamental parameters of the standard model and therefore of great interest. Furthermore, a major line of research within high-energy physics has been to determine the same CKM matrix element through different processes to look for inconsistencies that would signal a breakdown in the standard model. Until now, uncertainties in the lattice calculations have limited the precision of these tests. We aim to match the precision of our calculations to that of experiments.

Our first 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 Jefferson Laboratory. 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

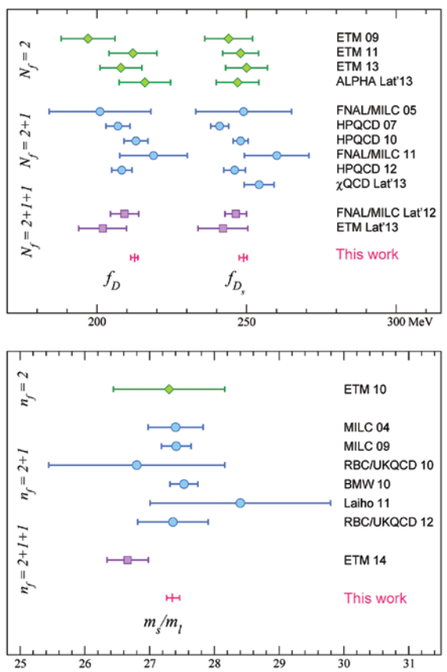

FIGURE 1: Comparison of our latest results for the leptonic decay constants f_D and f_{D_s} (top panel) and for the quark mass ratio m_s/m_l (bottom panel) with earlier lattice calculations. Results are grouped by the number of flavors of sea quarks from top to bottom: n_f=2 (green diamonds), n_f = 2+1 (blue circles), and n_f=2+1+1 (purple squares). Within each grouping, the results are in chronological order. Our new results are denoted by magenta pluses and labeled "This work".

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. Because these predictions will be made before the experiments are performed, these calculations will 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. Then 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 supercomputers available. The most computationally expensive component of the second step, the measurement routines, is to calculate the Green's functions for the propagation of quarks in the gauge configurations. For the light quarks, this calculation also requires 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 a number of mesons containing strange and charm quarks [1–6], which in turn have 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 the up, down, strange, and charm quark masses [3]. The HISQ configurations have also been used in the study of gradient flow and scale setting [7,8]. Important advances have been made in the development of code for the generation of gauge configurations and quark propagators with the Clover formulation of lattice quarks [9].

WHY BLUE WATERS

Lattice QCD calculations have made major progress in the last few years with a limited number of calculations reaching precision of a fraction of a percent and techniques in place to determine many more quantities to this level of accuracy. Such precision is needed to test the standard model and for a detailed understanding of physical phenomena controlled by the strong interactions.

The advent of petascale computers, like Blue Waters, is playing a critical role in these advances. High-precision QCD calculations are enormous undertakings that require computers of the highest capability and capacity.

QCD is formulated in the 4D space–time continuum, which is translated to a 4D lattice, or grid, for computational purposes. One must perform calculations for smaller and smaller

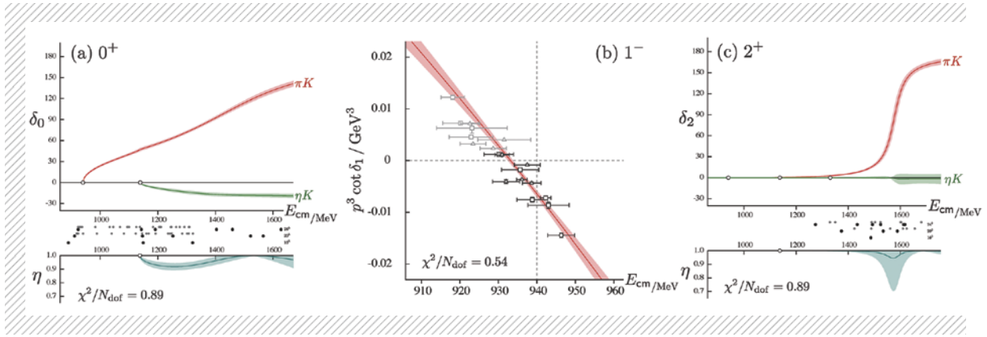

FIGURE 2: Results from [10] showing the scattering amplitudes of K-π and K-η mesons in isospin=1/2. The energy dependence of these scattering amplitudes is used to determine the resonance content of the spectrum. The left, middle, and right panels show the scattering amplitude of K-π (red), K-η (green), and K-π → K-η (lower panels) in l=0, 1, 2 waves, respectively. The points in between the panels are the energies determined from three different lattice sizes, and the curves are the parameterization of the scattering amplitudes with a resonance in $J^P = 0^+$ and 2^+, and a bound-state in $J^P = 1^-$.

lattice spacings across a fixed 4D space and extrapolate to the continuum (zero lattice spacing). Until recently, it has been too expensive to use the physical masses of the two lightest quarks (up and down), a known source of error. The computational cost grows roughly as the fifth power of the inverse of lattice spacing and also rises as the masses of the quarks decrease. Blue Waters enables us, for the first time, to carry out calculations with small lattice spacings and the masses of up and down quarks at their physical values. This has led to calculations of unprecedented precision.

PUBLICATIONS

Wilson, D. J., J. J. Dudek, R. G. Edwards, and C. E. Thomas, Resonances in coupled π K, η K scattering from lattice QCD. *Phys. Rev. D,* 91 (2015), 054008.

Bazavov, A., et al. (The Fermilab Lattice and MILC Collaborations), Charmed and light pseudoscalar meson decay constants from four-flavor lattice QCD with physical light quarks. *Phys. Rev. D,* 90 (2014), 074509.

Bazavov, A., et al. (The Fermilab Lattice and MILC Collaborations), Determination of |Vus| from a lattice-QCD calculation of the K › π l ν semileptonic form factor with physical quark masses. *Phys. Rev. Lett.,* 112 (2014), 112001.

Dudek, J. J., R. G. Edwards, C. E. Thomas, and D. J. Wilson, Resonances in coupled π K–η K scattering from quantum chromodynamics. Phys. Rev.Lett., 113 (2014), 182001.

Bazavov, A., (The MILC Collaboration), Gradient Flow Analysis on MILC HISQ Ensembles. *Proc. Sci. Lattice 2014,* 090 (2014).

Bazavov, A., et al. (The Fermilab Lattice and MILC Collaborations), Charmed and light pseudoscalar meson decay constants from HISQ simulations. *Proc. Sci. Lattice 2014,* 382 (2014).

Mohler, D., et al. (The Fermilab Lattice and MILC Collaborations), Low lying charmonium states at the physical point. Proc. Sci. Lattice 2014, 085 (2014).

Primer, T., et al., Kaon and D meson semileptonic form factors from lattice QCD. *Proc. Sci. Lattice 2014,* 374 (2014).

Winter, F. T., M. A. Clark, R. G. Edwards, and B. Joo, A Framework for Lattice QCD Calculations on GPUs. *Proc. 2014 IEEE 28th Int. Parallel Dist. Proc. Symp.,* pp. 1073–1082, doi:10.1109/IPDPS.2014.112.

Bazavov, A., et al. (The MILC Collaboration), Leptonic decay-constant ratio fK+/fπ+ from lattice QCD with physical light quarks. *Phys. Rev. Lett.,* 110 (2013), 172003.

Bazavov, A., et al. (The MILC Collaboration), Symanzik flow on HISQ ensembles. *Proc. Sci. Lattice 2013,* 269 (2013).

Gamiz, E., et al. (The Fermilab Lattice and MILC Collaborations), K semileptonic form factor with HISQ fermions at the physical point. *Proc. Sci. Lattice 2013,* 395 (2013).

Bazavov, A., (The Fermilab Lattice and MILC Collaborations), Charmed and strange pseudoscalar meson decay constants from HISQ simulations. *Proc. Sci. Lattice 2013,* 405 (2013).

DeTar, C., et al. (The Fermilab Lattice and MILC Collaborations), Charmonium mass splittings at the physical point. *Proc. Sci. Lattice 2012,* 257 (2012).

MODELING TYPE IA SUPERNOVA PROGENITORS

Allocation: NSF PRAC/0.375 Mnh
PI: Stan Woosley[1]
Co-PIs: John Bell[2], Michael Zingale[3]
Collaborators: Chris Malone[4], Ann Almgren[2], Andy Nonaka[2], Weiqun Zhang[2], Adam Jacobs[3], Max Katz[3]

[1]University of California, Santa Cruz
[2]Lawrence Berkeley National Laboratory
[3]Stony Brook University
[4]Los Alamos National Laboratory

EXECUTIVE SUMMARY:

Type Ia supernovae (SN Ia) are thermonuclear explosions of carbon–oxygen (CO) white dwarf starss in binary systems. Observations reveal diversity in both the spectra and progenitor characteristics, suggesting multiple origins. Theoretically, three different sorts of progenitor systems have been invoked: a binary star couples with a single Chandrasekhar mass white dwarf, a similar system with a lower mass dwarf, and a merging pair of white dwarfs. All have in common the eventual fusion of roughly one solar mass of carbon-oxygen (CO) to iron-group and intermediate-mass elements. Using Blue Waters, we are modeling each of these scenarios. Our most recent work involved the merging dwarf scenario.

INTRODUCTION

SN Ia are exceptionally bright explosions with an empirical relation that allows them to be used to measure cosmic distances. Both theory and observations require the progenitor to be a carbon–oxygen (CO) dwarf, but variations in how the explosion is initiated lead to different outcomes. In some cases, the explosion is highly anisotropic with significant viewing angle effects. In order to understand these fascinating explosions that are responsible for producing most of the iron on Earth as well as providing the first strong evidence for dark energy, multi-dimensional simulation is required. Subtle differences in an explosion may also affect the use of SN Ia for distance determination in an era of "precision cosmology."

METHODS & RESULTS

A suite of hydrodynamics codes has been developed and optimized specifically to study SNe Ia. Maestro [1] is a low-Mach-number hydrodynamics code that can efficiently model convection in white dwarfs during the early stages. The early stages time periods set the stage for a white dwarf's explosion and determine the outcome of the explosion. Once the explosion develops into a dynamic phase, Castro [2] solves the fully compressible equations of hydrodynamics, along with self-gravity and nuclear burning, and uses a flame model to represent the propagation of the burning front through the star. Castro can also follow the dynamics of inspiraling white dwarfs and their mutual detonation.

Last year, we used a combination of Maestro and Castro to run a Chandrasekhar mass model all the way from the early convective stage through explosion. This was one of the main goals of our Blue Waters project. This year we have focused more on the two alternate models, especially merging white dwarfs. This required substantial new coding, but we have now run the first set of test simulations for this progenitor system (fig. 1).

Both Maestro and Castro and the problem data files necessary to reproduce these simulations are available on github: https://github.com/BoxLib-Codes for other researchers.

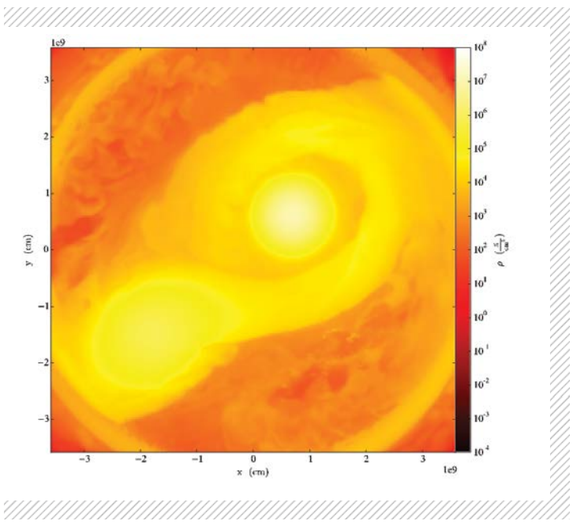

FIGURE 1: Evolution of a binary white dwarf system in the early stages of merging. This is a slice through the orbital plane of these stars. The lower-left object started as a 0.6 solar mass star, and the upper-right star started at 0.9 solar masses. The smaller star is slowly feeding its mass onto the bigger one; eventually it will be completely disrupted and the system will become completely coalesced, before possibly igniting as a Type Ia supernova.

WHY BLUE WATERS?

Blue Waters gives us the resources to study 3D systems at high resolution. All of the models considered are inherently 3D and yield outcomes that are asymmetric to varying degrees. High resolution, including several levels of mesh refinement, is required to track the nuclear burning that occurs both as a turbulent subsonic front and as detonation. In the latter case, fine resolution is needed to see the initial formation of the detonation.

PUBLICATIONS

Malone, C. M., et al., The Deflagration Stage of Chandrasekhar Mass Models For Type Ia Supernovae: I. Early Evolution. *Astrophys. J.*, 782 (2014), 11, doi:10.1088/0004-637X/782/1/11.

COMPUTER SCIENCE & ENGINEERING

OPTIMIZATION

SCALABILITY

PARALLEL FUNCTIONALITY

122 *Software Extreme Benchmarking*

124 *Algorithms for Extreme-Scale Systems*

SOFTWARE EXTREME BENCHMARKING

Allocation: Illinois/0.40 Mnh
PI: Ahmed Taha[1]

[1]National Center for Supercomputing Applications

EXECUTIVE SUMMARY:

The current allocation is a continuation for running extreme benchmarking for key computational fluid dynamics commercial codes on Blue Waters. The allocation helped us reach new "global" records in scaling two codes: Ansys-Fluent to 36,000 cores and Starccm+ of CD-Adapco to 20,000 cores.

INTRODUCTION

In service to 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 CFD codes, which are used by and of high interest to its PSP partners. The objective of the project is to explore the scaling limits of these codes to inform users of the codes' scaling capabilities for better production workflow. Two world-class CFD codes were included in the study: Fluent from Ansys,Inc., and Starccm+ of CD-Adapco.

The project was carried out (and is still running) in close collaboration between the PI (PSP-CFD Lead, Ahmed Taha), independent software vendor (ISV; Ansys and CD-Adapco) and the Blue Waters vendor (Cray).

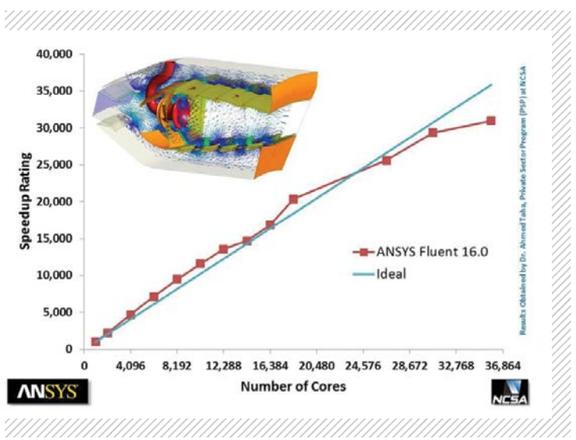

FIGURE 1: Ansys-Fluent speedup scaling curve, 830 million cells, and 36,000 cores.

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 collaboration with ISV and Cray was essential in adjusting environments.

Case study 1: After scaling Fluent to 10,000 cores towards the end of 2013 using the Ansys benchmark model of 111 million cells, Ansys provided its biggest mesh model—830 million cells. We scaled the code to 36,000 cores, the second world record the project set. The previous record, which we set in the first phase of the project in early 2014 while using the smaller 111-million-cell model, was 20,000 cores [1]. Fig. 1 presents the scaling curve for Ansys-Fluent reaching 36,000 cores with over 80% efficiency [2].

Case study 2: Starccm+. CD-Adapco provided two of its biggest mesh models: 100 million cells and 514 million cells. As shown in fig. 2, the 100-million-cell model scaled to 13,000 cores, defeating the 8,000-core record set at the HPC Center Stuttgart (HLRS), Germany. We scaled the 514-million-cell model to 20,000 cores, defeating its previous 16,000-core record achieved at HLRS (fig. 3).

Among the many benefits of this extreme scaling study are [3–4]:

1. It provided users with information on how robust the flow solvers of the tested codes are, which is presented in the high-fidelity simulations running at such massive core counts.

2. It drove the mesh count, allowing more details in the complex physics needed for more thorough design studies.

3. It showed how time to solution was 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.

4. This 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. Other leading CFD ISVs approached the project lead to run similar scaling projects for their codes. We plan to continue to present our work and records at U.S. and global conferences.

WHY BLUE WATERS?

This project would have not able to reach this success level without Blue Waters. NCSA and Blue Waters are the only place capable of doing such work where you can find massive computational resources, experienced staff, and domain experts.

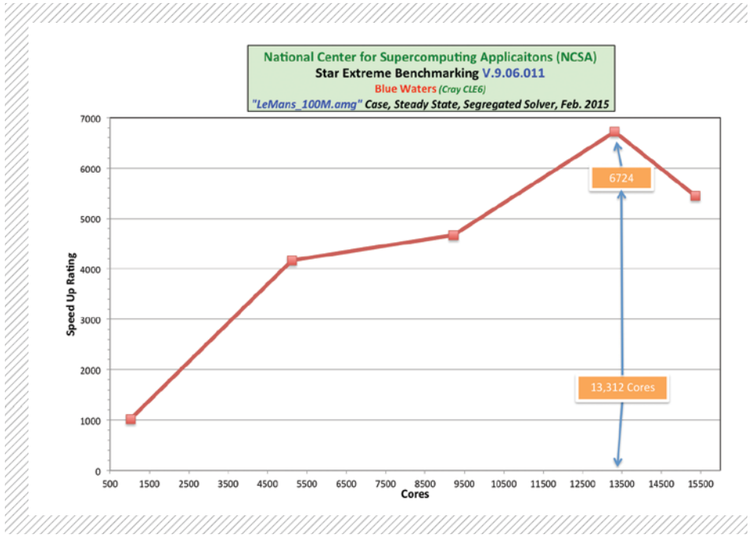

FIGURE 2: Starccm+ speedup scaling curve, 100 million cells, and 16,000 cores.

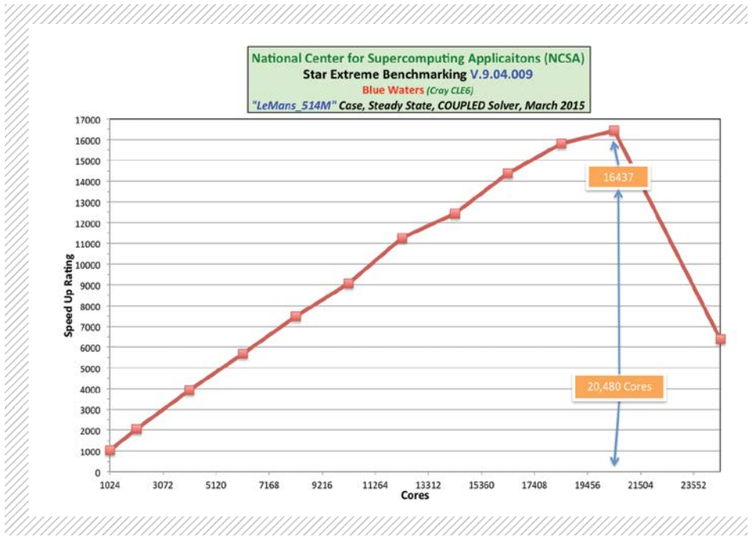

FIGURE 3: Starccm+ speedup scaling curve, 513 million cells, and 20,000 cores.

ALGORITHMS FOR EXTREME-SCALE SYSTEMS

Allocation: Blue Waters Professor/0.24 Mnh
PI: William Gropp[1]
Collaborators: Luke Olson[1], Jack Dongarra[2]

[1]University of Illinois Urbana–Champaign
[2]University of Tennessee

EXECUTIVE SUMMARY:

Continued increases in the performance of large-scale systems will come from greater parallelism at all levels. At the node level, we see this in both the increasing number of cores per processor and the use of large numbers of simpler computing elements in GPGPUs (general-purpose processing on GPUs). The largest systems must network tens of thousands of nodes together to achieve the performance required for the most challenging computations. Successfully using these systems requires new algorithms and new programming systems. This research looks at the effective use of extreme-scale systems. Over the last year, we have explored alternative formulations of a conjugate gradient method that eliminate some of the strict barrier synchronization as well as use the memory hierarchy more effectively. Other exploratory studies have begun looking at the scalability and fault tolerance of an algebraic multigrid method, scaling for large graph problems, and the benefit of lightweight intra-node balancing to scalability and performance.

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 to effective use of extreme-scale systems by applications. For example, Krylov methods such as the conjugate gradient method are used in many applications currently being run on Blue Waters (MIMD Lattice Computation—MILC—is one well-known example). Developing and demonstrating a more scalable version of this algorithm would immediately benefit those applications. Longer term, the techniques that are developed will guide the development of highly scalable applications.

METHODS & RESULTS

Early results with alternative Krylov method formulations have revealed several performance effects that can provide a factor of two or more 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 in terms of the ability to overlap Allreduce with other communication and computation. A need for sparse matrix formats for GPUs for several other areas of interest drove work in this area, though it did not use any of the Blue Waters allocation in this year; that work will be used in the upcoming year.

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 run. In addition, only Blue Waters provides a highly capable I/O system, which we plan to use in developing improved approaches to extreme-scale I/O

BIOLOGY, CHEMISTRY & HEALTH

MOLECULAR

CELLULAR

MEDICINE

BIOPHYSICS

128 *Virtual Flu, a Different Kind of Computer Virus*

130 *Mining the Evolutionary Dynamics of Protein Loop Structure and its Role in Biological Functions*

132 *Convergence and Reproducibility in Simulations of Nucleic Acids*

134 *Epistatic Interactions for Brain expression GWAS in Alzheimer's Disease*

136 *Sequence Similarity Networks for the Protein "Universe"*

138 *Elucidating the Molecular Basis of Charge Selectivity in Pentameric Ligand-Gated Ion Channels*

140 *Non-Born–Oppenheimer Effects between Electrons and Protons*

142 *Ice and Water*

144 *Effect of Confinement on Protein Dynamic and Structure*

146 *Multiscale Modeling of Bone Fracture and Strength*

148 *High-Resolution MR Elastography of the Brain in a Clinical Setting*

150 *Functional Mechanism of the Bacterial Efflux Pump*

152 *Functional Mechanics of Voltage Gated like Ion Channels*

154 *Simulating Ribosome Biogenesis in Whole Cells*

157 *Simulating Influenza Hemagglutinin Membrane Assemblies*

158 *Instrumenting Human Variant Calling Workflow*

160 *Quantum-Classical Path Integral Simulation of Proton Transfer in Solution*

162 *Blood–Artery Interactions in Patient-Specific Geometries*

164 *Understanding P53, a protein "shape shifter" responsible for many cancer tumors*

166 *Protein Structure Determination and Refinement from Raw NMR Data*

168 *Molecular Details of the Luminal Exit of Calcium Ions in the Calcium Pump of the Sarco/Endoplasmic Reticulum*

170 *The Computational Microscope*

172 *Characterizing Structural Transitions of Membrane Transport Proteins at Atomic Detail*

174 *Ultra-Coarse-Grained Biomolecular Simulations*

VIRTUAL FLU, A DIFFERENT KIND OF COMPUTER VIRUS

Allocation: NSF PRAC/6.0 Mnh
PI: Rommie E. Amaro[1,2]
Collaborators: Alasdair Steven[3], Theoretical and Computational Biophysics Group[4]

[1]University of California, San Diego
[2]National Biomedical Computation Resource (NBCR), University of California, San Diego
[3]National Institutes of Health
[4]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 and X-ray crystallography typically lack either the atomic resolution or size scaling required to answer a number of pharmacologically important questions. Additionally, "gain-of-function experiments," aimed at anticipating the genetic changes that might produce the next pandemic, could prove disastrous were artificially enhanced viral strains to escape the lab [1].

To address these concerns, we have constructed an atomic-resolution model of the entire influenza viral coat, containing 210 million atoms. We are presently performing simulations of this large-scale system in order to gain a more complete understanding of the influenza infection process. This research will allow us to explore novel opportunities for drug and vaccine development *in silico* and test how structural changes impact virulence without having to create novel strains with pandemic potential.

INTRODUCTION

We have 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 comprised of a lipid bilayer from which two spike-like glycoproteins, neuraminidase and hemagglutinin, protrude (fig. 1). When a viral particle first approaches a host cell, the hemagglutinin proteins latch onto sialic-acid molecules attached to the cell surface. Once bound to the human cell through these molecular linkers, the virus enters the cell and reproduces. Following replication, the viral progeny bud out from the cell but remain attached to its external surface by the same sialic-acid connections. Influenza's second glycoprotein, neuraminidase, 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, flu is highly adaptable and resistance to these medicines has already been documented [3–14]. 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, will provide pharmacologically relevant insights that will help us combat future pandemics. No current experimental technique is capable of providing an accurate model of the entire surface coat at the resolution needed for drug discovery; fortunately, modeling and simulation can serve as a "computational microscope" that provides the needed information.

Additional influenza studies have shown how changes in the components of the influenza surface coat affect virulence. By evolving viruses in the lab that are more infectious than the corresponding natural strains, researchers seek to anticipate future virus mutations that might lead to the next pandemic [15–25]. But the possibility that an artificially enhanced virus might escape the lab concerns some people [1]. Indeed, accidents have happened in the past. The 1977 flu pandemic may have come from a lab-preserved strain [26–28], and in 2014 the CDC accidentally contaminated a biological sample sent to external collaborators with a highly pathogenic influenza virus [29]. The latter accident was only discovered during an investigation into a similar anthrax contamination. By providing a reliable computational model of the viral surface, we hope to reduce the need for these kinds of risky experiments.

METHODS & RESULTS

Guided by experimental data, we have successfully constructed several surface-coat models. When immersed in a bath of virtual water with the appropriate electrolytes, each system contains over 200 million atoms (fig. 2). Our collaborator (Steven) used electron microscopy to identify the general shape of the influenza virus and the approximate locations of the glycoprotein spikes [30]. We then used computational methods developed in our lab to wrap this virus volume in a virtual lipid bilayer [31] and position atomic-resolution models of the glycoproteins at the appropriate locations. 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 are currently running a simulation of one of these virtual viruses on the Blue Waters supercomputer.

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 the Blue Waters NSF PRAC allocation, the current work would be impossible.

In the coming year we anticipate completing the first simulation of this large-scale system. This simulation will provide 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 binding. They will allow us to study important characteristics of the virus without having to actually create new, potentially lethal flu strains in the lab.

FIGURE 1 (BACKGROUND): The influenza surface coat. The hemagglutinin & neuraminidase glycoproteins of our model are represented in lavender and magenta, respectively. The enveloping lipid bilayer is represented in glossy pink. Credit: Jacob D. Durrant.

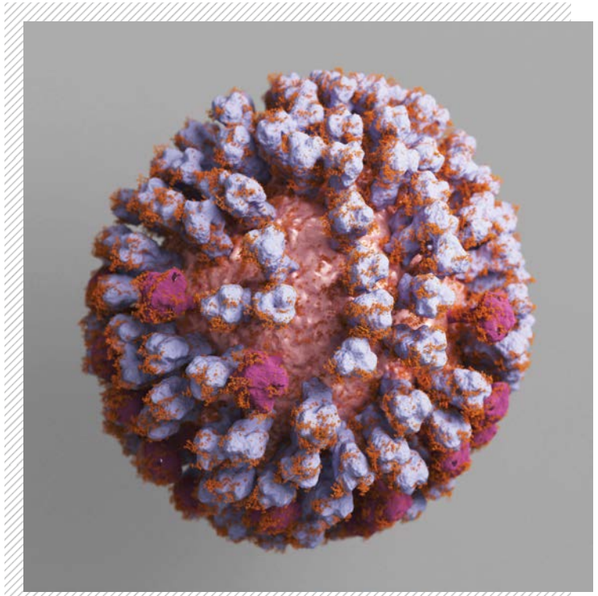

FIGURE 2: The viral surface coat, with selected neuraminidase and hemagglutinin glycoproteins rendered transparent, revealing the underlying protein backbone. This image highlights the fact that these renderings were generated from high-resolution (ultimately atomic-resolution) data. Credit: Jacob D. Durrant.

MINING THE EVOLUTIONARY DYNAMICS OF PROTEIN LOOP STRUCTURE AND ITS ROLE IN BIOLOGICAL FUNCTIONS

Allocation: Illinois/0.19 Mnh
PI: Gustavo Caetano-Anollés[1]
Co-PI: Frauke Gräter[2]
Collaborator: Fizza Mughal[1]

[1]University of Illinois at Urbana-Champaign
[2]Heidelberg Institute for Theoretical Studies, gGmbH

EXECUTIVE SUMMARY:

Flexible and unstructured regions of protein molecules introduce a source of conformational heterogeneity that is fundamental for their biological function. Here we study this heterogeneity with microsecond-scale molecular dynamics simulations using the NAMD 2.9 platform. Previous studies have shown that the dipeptide make-up of proteins delimits loop regions that are unstructured and flexible, driving the specificity and stereochemistry of the genetic code. A previous Blue Waters allocation confirmed that unstructured regions collapse quickly, but unexpectedly revealed increased intra-chain dynamics of short fragments. Our research provided unprecedented atomistic details of the dynamics of 74 loop regions of protein domains sampled along a timeline of domain history, studying collapse propensities and intrinsic fluctuations of each loop structure. Our study provides insight into how loop flexibility and disorder are linked to the genetic code and primes protein function in evolution.

INTRODUCTION

Proteins sustain life on our planet, from major biogeochemical cycles necessary for planetary stability to crucial signaling in brain activities important for cognition and behavior. Protein misfolding results in aggregation and diseases such as Alzheimer's or Creutzfeldt–Jakob. Their challenge and deregulation causes pathogenesis and cancer. Despite proteins' importance, many aspects have yet to be uncovered; among them are how and why this sophisticated molecular machinery carries biological functions, the rationale for its change in time, the mysterious origin of the "vocabulary" that shapes genetics (the genetic code) and the evolutionary drivers of the structure of proteins.

Traditionally, the structure of proteins is defined by a "regular" structure, the 3D arrangement of helical (typically α-helices) and strand (typically β-strands) segments resulting from local and long-range interactions of the protein backbone, respectively [1]. In recent years, however, there has been growing recognition that "non-regular" structural regions that connect helices and strands play more important roles at functional and regulatory levels. These intra-chain regions are compact but lack fixed internal hydrogen bonding interactions. They include loops and intrinsically disordered regions. Protein loops, which are often present in active sites of enzymes or often enable molecular mechanics, have been of particular interest [2]. Their flexibility may provide a link between protein structure and function. Information about them has been collected in numerous databases, which index their topology, supporting fold structures, and associated gene ontology definitions of molecular functions [e.g., 3,4].

METHODS & RESULTS

In previous work, we discovered that the speed of folding, which correlates with flexibility, is enhanced during the evolution of protein domains [5]. We also discovered that protein structures enriched with flexible loops appeared with the evolutionary unfolding of the genetic code [6]. Since structural flexibility is a conserved feature in the assembly of protein complexes [7], flexibility must be regarded as an emergent property of molecular evolution. Following an exploratory Blue Waters proof-of-concept allocation [8], we used the power of Blue Waters to study the dynamics of proteins loops in 74 protein domains of aminoacyl-tRNA synthetases, the enzymes responsible for the specificity of the genetic code. Advanced tools of molecular dynamics (MD) simulations using the NAMD 2.9 platform allowed gathering global parameters (RMSD, radius of gyration) and detailed mapping of motions linked to loop sequence and structure. Each equilibrium all-atom simulation required ~0.005 Mnh. Benchmarking indicates 5 Mnh are needed to cover the target of 1,000 representative loops.

Remarkably, we found a trend of reduced global RMSD and radius of gyration during the entire timeline of protein history that indicates an evolutionary tendency towards conformational order, which we would like to confirm with additional analyses. We also found dynamic heterogeneities in both loops and surrounding regular structures that could be associated with specific functions. Our research links two fields of study that have not yet interfaced in science: (1) the evolution of molecular structure and intrinsic disorder in proteins, and (2) the molecular dynamics of proteins. The former focuses on evolutionary processes spanning billions of years of biological history. The latter looks at molecular change unfolding at nanosecond to microsecond levels. This interdisciplinary exploration is expected to uncover patterns of origin and evolution of the genetic code, protein structures, and functions.

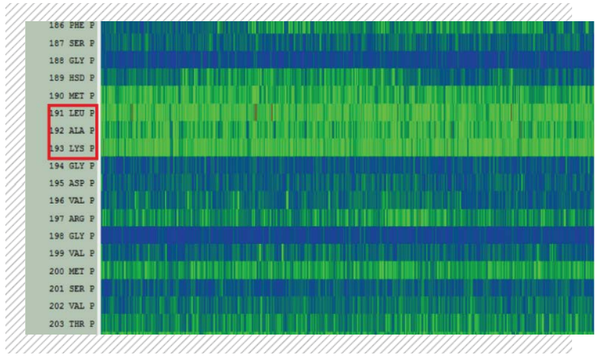

FIGURE 1: Root mean square fluctuation (RMSF) plot showing a heat map of dynamic fluctuations of each amino acid residue of loop 1JMZ_A_182, embedded in an OB-fold Myf domain, over time from left to right. Note the increased dynamics of the loop region (red square encompassing Leu, Ala, and Lys residues) and the dynamic heterogeneities of the van der Waals lock surrounding it.

WHY BLUE WATERS?

The results of our initial exploration and benchmarking exercise provided a foundation for a Blue Waters-enabled high-throughput MD simulation study of the dynamics of a massive number of protein loops, which will be indexed with evolutionary, structural, and functional information. The study will yield unprecedented atomistic details of structural and functional evolutionary constraints that are responsible for structuring both proteins and the genetic code.

PUBLICATIONS

Caetano-Anollés, G., et al., The dynamics of protein disorder and its evolution: Understanding single-molecule FRET experiments of disordered proteins. *Blue Waters Annual Report* (Urbana, Illinois, 2014), pp. 120–121.

CONVERGENCE AND REPRODUCIBILITY IN SIMULATIONS OF NUCLEIC ACIDS

Allocation: NSF PRAC/7.0 Mnh
PI: Thomas E. Cheatham, III[1,2]
Collaborators: Christina Bergonzo[2], Rodrigo Galindo-Murillo[2], Daniel R. Roe[2]

[1]Center for High-Performance Computing, University of Utah
[2]University of Utah

EXECUTIVE SUMMARY:

Over the past 20-plus years, we have developed and used biomolecular simulation methods to provide insight into nucleic acid structure and dynamics. Along the way, we have faced significant challenges in the form of force-field and sampling limitations. With access to resources like Blue Waters, we have reached the tipping point where we can fully sample ensembles of small RNA molecules and DNA helices and, with recent force-field modifications, are able to properly sample RNA conformational distributions in agreement with experiment. This opens the door for the simulations we really want to perform; rather than simple assessment and validation of simulation results, we can use these powerful biomolecular simulation methods to provide significant insight into functional RNA structure and dynamics in riboswitches, ribozymes, interfering RNA, and protein–nucleic acid interactions.

INTRODUCTION

As more is understood about the biological importance of RNA, the ability to peer into this biomolecule's dynamics on an all-atom level is of increasing interest. For example, it is known that RNA molecules (riboswitches) can detect the concentration of certain metabolites leading to structural changes that alter import or production of the metabolite. RNA is also involved in gene regulation, protein production, and many emerging functional roles. Keys to RNA's ability to function are not only its structure and dynamics, but also conformational rearrangements induced by changes in its environment. Understanding the detailed atomic structure and dynamics of RNA provides insight into how RNA functions and potentially how to modulate this function, for example through drug binding.

As atomic-resolution experimental approaches have difficulty in resolving dynamic structures or structures that populate multiple configurations, there is a need to employ alternative approaches to characterize RNA structure and dynamics. High-level molecular dynamics (MD) simulations provide an additional means to study RNA dynamics (and therefore its function) with atomic-level detail. However, for such simulations to be useful they need to not only capture all of the bio-relevant conformations, but also sample correct structure populations. The former depends on the accuracy of the parameters used during the simulation (referred to as a "force field"), while the latter is a challenging problem requiring a large amount of aggregate simulation time. These two issues are related: unless you have relatively complete sampling, it is difficult to determine if inaccuracies are the result of poor parameters or if they result from limited sampling. Until recently, the challenge in simulating RNA has been validating putative changes intended to improve RNA force fields that have been generated with limited sampling of RNA structure and dynamics. Using Blue Waters, we have been able to divorce the limitation of sampling from the dynamics and structures described by the force fields, allowing a totally unbiased look at each force field's description of both DNA and RNA [1–5] .

METHODS & RESULTS

By using multi-dimensional replica exchange MD (M-REMD), we have shown that we can reproducibly and efficiently generate effectively complete structure ensembles for a variety of RNA motifs. In previous simulations we found that a particular tetranucleotide, rGACC, sampled a wide range of conformations, though experimental nuclear magnetic resonance spectroscopy results indicated only two major structures were present [6,7]. Using Blue Waters, we have expanded testing of current force fields to include combinations of non-bonded parameter modifications with several ion and water models.

The results indicate that with improved parameters, including a modified van der Waals radii set [8] and improved water model (OPC) [9], the ensemble of rGACC structures generated by M-REMD are now in quantitative agreement with experimental nuclear magnetic resonance. Comparison between two independent runs shows the high level of convergence achievable only through use of Blue Waters.

More recent work on the UUCG tetraloop structure (which is a small helical RNA stem with a r(UUCG) sequence loop on top), where it has been problematic for all current force fields to capture the native structure, shows that including these modifications in the presence of the improved OPC water model increases the percent of native structure found by an order of magnitude. Due to the high amount of sampling required to converge these simulations, this study is ongoing.

WHY BLUE WATERS?

A key innovation enabling our work is the implementation of AMBER on NVIDIA GPUs. This results in the fastest molecular dynamics code available to date [10]. Improvements in the past year have included the ability to repartition the solute's hydrogen masses, allowing us to double the simulation time step and achieve an almost instant two-fold increase in simulated time vs. real time [11]. Combined with the M-REMD framework, this very effectively reduced, by orders of magnitude, the real-world time needed to obtain the converged ensemble. With larger ensembles and combined ensembles enabled by the next generation of Track-1 systems, we will be able to move into the detailed investigation of more bio-relevant nucleic acid structures including riboswitches, ribozymes, and various protein–nucleic interactions with full reproducibility and convergence of the results.

PUBLICATIONS

Galindo-Murillo, R., D. R. Roe, and T. E. Cheatham, 3rd, Convergence and reproducibility in molecular dynamics simulations of the DNA duplex d(GCACGAACGAACGAACGC). *Biochim. Biophys. Acta,* 1850:5 (2015), pp. 1041–1058, doi:10.1016/j.bbagen.2014.09.007.

Bergonzo, C., et al., Multidimensional Replica Exchange Molecular Dynamics Yields a Converged Ensemble of an RNA Tetranucleotide. *J. Chem. Theory Comput.,* 2014. 10:1 (2014), pp. 492– 499, doi:10.1021/ct400862k.

Galindo-Murillo, R., D. R. Roe, and T. E. Cheatham, 3rd, On the absence of intrahelical DNA dynamics on the μs to ms timescale. *Nature Commun.,* 5 (2014), 5152, doi:10.1038/ncomms6152.

Roe, D. R., C. Bergonzo, and T. E. Cheatham, 3rd, Evaluation of enhanced sampling provided by accelerated molecular dynamics with Hamiltonian replica exchange methods. *J. Phys. Chem. B,* 118:12 (2014), pp. 3543– 52, doi:10.1021/jp4125099.

EPISTATIC INTERACTIONS FOR BRAIN EXPRESSION GWAS IN ALZHEIMER'S DISEASE

Allocation: Private sector/31 Knh
PI: Nilufer Ertekin-Taner[1]
Collaborators: Mariet Allen[1]; Liudmila Mainzer[2,3]; Curtis Younkin[1]; Victor Jongeneel[2]; Thierry Schüpbach[4]; Gloria Rendon[2]; Julia Crook[1]; Julie Cunningham[5]; Summit Middha[5]; Chris Kolbert[5]; Dennis Dickson[1]; Steven Younkin[1]

[1]Mayo Clinic in Jacksonville
[2]University of Illinois at Urbana-Champaign
[3]National Center for Supercomputing Applications
[4]Swiss Institute of Bioinformatics
[5]Mayo Clinic in Rochester

EXECUTIVE SUMMARY:

It is well established that the risk for Alzheimer's disease (AD) is under substantial genetic control, and is thought to arise from multiple genetic variants. Such disease-associated variants can be identified using expression-based genome-wide association studies (eGWAS), based on the rationale that some variants will influence AD risk via their effects on brain gene expression. We hypothesize that some of the risk for AD may be due to the interaction of two or more genetic variants (epistasis). The aim of our Blue Waters project was to test for the presence of epistatic interactions that influence brain gene expression levels using data from 359 temporal cortex samples (181 AD, 178 non-AD), 223,632 SNP genotypes and ~24,526 transcripts that were measured using an expression array. We also ran an analysis on 343 cerebellum samples (173 AD and 170 non-AD). The analysis of epistatic effects in studies of this size would not be possible without the unique computing capabilities of Blue Waters. All the planned computational runs have now been completed, along with the analyses of scalability and performance.

INTRODUCTION

We have previously collected gene expression measures from pathologically confirmed AD subjects (test group) and those with non-AD pathologies (control group, Table 1) from two brain regions: temporal cortex (TCX) and cerebellum (CER). Identifying genetic variants that associate with altered gene expression levels in these subjects may pinpoint novel risk factors for AD. We investigated single genetic variants for association with these gene expression measures and found significant expression quantitative trait loci [1]. We also determined that some of the known AD risk variants likewise associate with expression of nearby gene(s) thus implicating the potential mechanism of action and the affected gene at these loci [2].

For our study on Blue waters, we hypothesized that pairs of variants may likewise influence gene expression through an interaction known as epistasis. Identifying additional genetic factors that influence AD risk can provide further insights into the pathophysiology of this disease and may have a significant impact on the development of novel therapeutic targets, identification of potential, premorbid biomarkers, and generation of in vivo disease models, much needed for pre-clinical development and testing of novel therapies.

METHODS AND RESULTS

Three groups of subjects with temporal cortex measures were analyzed: AD's-only, Non-AD's only and the two combined (AD+nonAD); the combined set only (AD+non-AD) was assessed for cerebellum (Table1, Figure 1). Prior to launching our analysis, we implemented conservative quality control measures (Figure 1) and LD pruned our dataset in order to capture the maximum genetic data whilst minimizing the multiple testing penalty, similar to a protocol described elsewhere [3].

The currently available epistasis approaches are unable to efficiently incorporate covariates into regression models. To address this we generated gene expression residuals using R for all 24,526 expression measures to account for the following key covariates: Age, Gender, #ApoEε4 alleles, PCR plate, RIN, RINsqAdj (RIN-RINmean)2 and diagnosis when appropriate, (AD=1, Non-AD=0), as described previously [4].

Three different software programs were considered for detecting the epistatic interactions: PLINK [5], EpiGPU [6] and FastEpistasis [7]. It was determined that FastEpistasis performed most optimally for testing of multiple quantitative phenotypes using the computational architecture of Blue Waters. FastEpistasis builds on the analysis paradigm used in PLINK, but is multithreaded and runs up to 75 times faster by splitting the analysis into three phases.

All the planned computational runs have now been completed, along with the analyses of scalability and performance. Results will be organized into a database to facilitate efficient filtering to identify relevant significant results. The most significant findings will be tested for validation in additional subjects using a targeted approach.

WHY BLUE WATERS

The Blue Waters supercomputer was instrumental for the success of this work. Our project simply would not have been possible without this supercomputer and its dedicated team.

First, its sheer size allows us to run all phenotypes in parallel, at the same time. While FastEpistasis is not MPI-enabled, we were still able to pack up to 32 phenotypes per compute-node. Additionally, we utilized the MPI launcher software developed by the Blue Waters staff, to run all 24,526 phenotypes in parallel in a single 787-node reservation. This configuration was used for all four experiments: TCX-AD, TCX-Control, TCX-ALL, CER-ALL. As a result, each of these tests took only 5-10 hours to run, as opposed to the two years of walltime predicted to be necessary for PLINK.

Second, the analysis presented certain challenges that required a close collaboration between the Mayo team and the Blue Waters support group:

- We have learned to monitor our jobs using resource profiling software developed by NCSA staff. This helped us detect whether jobs were progressing normally.
- FastEpistasis generated millions of files across the project, and we worked with the Blue Waters team to manage data storage and transfers in efficient ways.
- Heterogeneity of data across phenotypes also resulted in uneven walltimes within each multi-node job, and NCSA staff helped us study the impact of this property on computational cost of the jobs, so as to better plan future analyses.

Both the hardware access and staff support were required to complete this project.

Tissue, Diagnosis	N	Female (%)	Mean Age (SD)	ApoE4+ (%)	Mean RIN (SD)
TCX, AD*	181	94 (52)	73.6 (5.6)	108 (60)	6.3 (0.8)
TCX, Non-AD*	178	67 (38)	71.5 (5.6)	46 (26)	6.9 (1.0)
TCX, ALL*	359	161 (45)	72.5 (5.6)	154 (43)	6.6 (1.0)
CER, AD	173	88 (51)	73.5 (5.7)	108 (63)	7.1 (1.0)
CER, Non-AD	170	60 (35)	71.6 (5.5)	45 (26)	7.2 (0.9)
CER, ALL*	343	148 (43)	73.0 (5.7)	153 (45)	7.2 (0.9)

TABLE 1: Demographics and characteristics of the samples with expression measures. *Indicates subjects groups analyzed using Fast Epistasis. TCX: Temporal cortex, CER: Cerebellum, AD: Subjects with Alzheimer's disease, non-AD: Subjects without Alzheimer's disease, Age: Age at death, ApoE4+: Number (percent) subjects with ApoE4 allele, RIN: RNA integrity number, SD (standard deviation).

SEQUENCE SIMILARITY NETWORKS FOR THE PROTEIN "UNIVERSE"

Allocation: Illinois/0.625 Mnh
PI: John A. Gerlt[1,2]
Personnel: Katie L. Whalen[2], Daniel Davidson[1], Boris Sadkhin[1], David Slater[1]
Collaborator: Alex Bateman[3]

[1]Institute for Genomic Biology, University of Illinois at Urbana–Champaign
[2]Enzyme Function Initiative
[3]European Bioinformatics Institute, European Molecular Biology Laboratory

EXECUTIVE SUMMARY:

The Enzyme Function Initiative (EFI) is supported by the National Institutes of General Medical Sciences (U54GM093342-04). The EFI is devising strategies and tools to facilitate prediction of the *in vitro* activities and *in vivo* metabolic functions of uncharacterized enzymes discovered in genome projects. This project is enabling generation of a library of precomputed sequence similarity networks (SSNs) for all 14,831 Pfam sequence-based families, 515 Pfam sequence-based clans, and 2,737 Gene3D/CATH structure-based superfamilies in the UniProtKB protein database; the SSNs will be provided to the scientific community via a webserver [1]. At present, the UniProtKB database contains 92,672,207 nonredundant sequences (release 2015_3; 03-March-2015). We used Blue Waters to calculate the required all-by-all BLAST sequence comparisons as well as generate statistical analyses of the BLAST results. We plan to calculate the networks on a two-month refresh cycle so that the library will remain current.

INTRODUCTION

The Enzyme Function Initiative (EFI) is supported by the National Institutes of General Medical Sciences (U54GM093342) and is multi-institutional (nine academic institutions in the U.S. and the European Molecular Biology Laboratory/European Bioinformatics Institute in the U.K.) and multi-disciplinary (bioinformatics, protein production, structural biology, homology modeling, *in silico* ligand docking, experimental enzymology, microbiology, transcriptomics, and metabolomics) [2]. The goal of the EFI is to devise tools and strategies to enable prediction of the *in vitro* activities and *in vivo* metabolic functions of uncharacterized enzymes discovered in genome projects.

The UniProtKB database contains 92,672,207 sequences (release 2105_03; 04-March-2015). The functions for 547,964 entries (0.5%) have been manually curated [3]; the functions for the remaining have been assigned by automated procedures [4]. The majority of the entries were obtained from microbial genome sequencing projects, with the rationale that knowledge of the complete set of proteins/enzymes encoded by an organism will allow its biological/physiological capabilities to be understood. However, if many of the proteins/enzymes have uncertain or unknown functions, people cannot capitalize on the investments in genome projects. The EFI was conceived to meet this challenge.

Bioinformatic tools are integral to the EFI's strategies. Phylogenetic trees and dendrograms are the usual bioinformatic representations of relationships among homologous proteins. However, their construction requires structure-based sequence alignments and is computationally intensive. Recently, Babbitt described the use of sequence similarity networks (SSNs) to visualize relationships in families of homologous proteins [5]. Sequence similarities are quantitated by the BLAST bit-scores between pairs of sequences.

The EFI's goal is to provide to the biological community an on-demand library of SSNs for all 14,831 Pfam sequence-based families, 515 Pfam sequence-based clans, and 2,737 Gene3D/CATH structure-based superfamilies in the UniProtKB protein database and to update this library on a minimum two-month refresh cycle.

METHODS & RESULTS

In the first year of our Blue Waters allocation, we optimized two pieces of codes as well as the Perl scripts that control the flow of data and collect the results.

BLAST v 2.x (*blastall*) is a widely used program developed by the National Center for Biotechnology Information (NCBI). *Blastall* is not efficiently multi-threaded, so we ran as many single-threaded processes per node as there are integer cores available. The bulk of the Blue Waters CPU time is used by *blastall*.

CD-Hit is a sequence clustering algorithm [6] that we used to both generate merged datasets of input sequences and/or post-process sequences flagged as being similar to each other by *blastall*.

In the second year, we initially generated SSNs for virtually all of the Pfam families in Release 49.0 of the InterPro database/Release 2014_10 of the UniProt database. We also developed a webtool that allows members of the scientific community to download the precomputed SSNs [7].

This experience involved improvements in our data generation algorithms that reduced the computational complexity, optimized the performance, and dramatically decreased the node-hours required for the BLAST. In particular, we improved the data generation pipeline by increasing automation, clustering highly similar protein sequences, binning the sequences, and improving our data plot generation algorithms.

With these improvements, we are generating SSNs for virtually all of the Pfam families in Release 50.0 of the InterPro database/Release 2015_02 of the UniProt database. When complete, these also will be made available to the community with our webtool. As new releases of InterPro/UniProt databases are made available, we will continue to generate the library of Pfam SSNs and make these available to the scientific community.

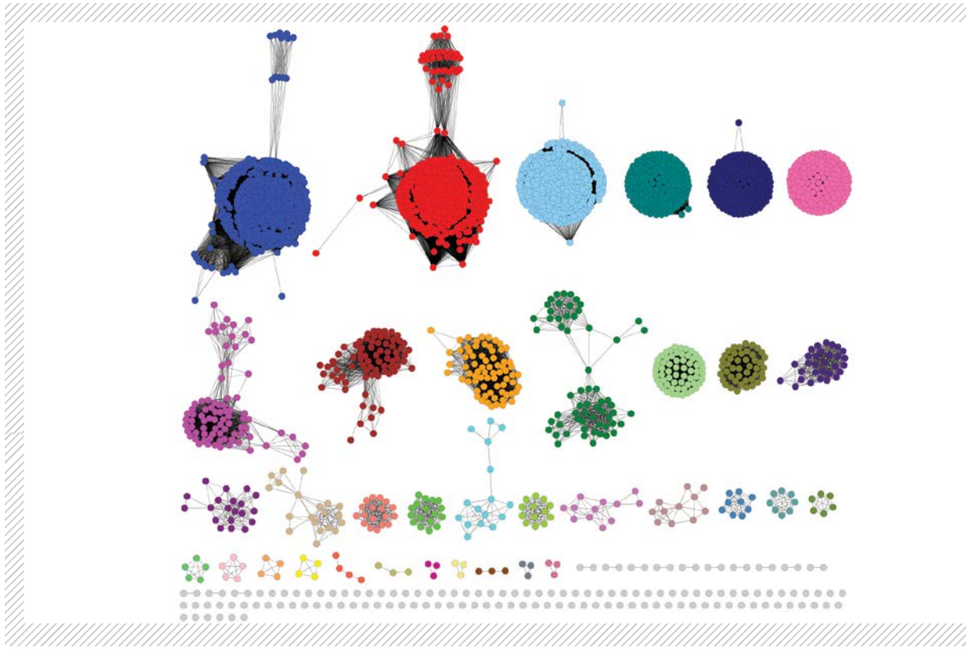

FIGURE 1: Sequence similarity network (SSN) for PF08794, the proline racemase family, displayed with a BLAST e-value threshold of 10^{-110} (50% sequence identity). The colors are used to distinguish predicted isofunctional clusters.

WHY BLUE WATERS?

The project uses an embarrassingly parallel computing model to perform the BLAST analyses and, in principle, could be run on any cluster of sufficient size. However, because of the scales of the computations used (number and sizes of Pfam families and clans and the number of sequences in each family and clan) and the time sensitivity of the production of the output relative to database updates, only a resource of the scale of Blue Waters can perform the job in a reasonable time frame.

Only Blue Waters is able to provide the ultra-large-scale computational power required to generate SSNs for Pfam-defined families and clans as the databases are updated. The sequence database is expected to continue to grow exponentially for the foreseeable future, and only Blue Waters has the necessary capacity to support these increasing needs.

ELUCIDATING THE MOLECULAR BASIS OF CHARGE SELECTIVITY IN PENTAMERIC LIGAND-GATED ION CHANNELS

Allocation: Illinois/0.132Mnh
PI: Claudio Grosman[1]

[1]University of Illinois at Urbana–Champaign

EXECUTIVE SUMMARY:

When activated, ion channels open an aqueous pore in the cell's membrane that allows passive movement of ions from one side of the membrane to the other. Because ions are charged particles, the movement of ions across the cell membrane creates a current that results in a change in the neuron's membrane potential. This change in membrane potential may cause a neuron to undergo an action potential, which propagates the electrical signal through the nervous system. However, ion channels can either promote or inhibit action potentials depending on what type of ion crosses a particular channel.

Pentameric ligand gate ion channels are unique in that they contain both anion- and cation-selective channels and therefore provide a great model to understand how particular proteins selectively allow the passage of unique ionic species. Using computer simulations, it is possible to understand the precise molecular basis for charge selectivity.

INTRODUCTION

Ion channels passively allow ions to diffuse in and out of a cell. However, many ion channels only allow 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 only allow positively charged cations to move across the membrane, whereas anion-selective ion channels only allow 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 observed to be inhibitory because it hyperpolarizes the cell, making it more difficult for an action potential to occur.

One particularly well-characterized superfamily of ion channels is known as the pentameric ligand-gated ion channels (pLGICs). These channels open in response to neurotransmitters such as acetylcholine, serotonin, and GABA and are important for regulating neuronal signals in various regions of the body. These channels are also targets for important classes of drugs such as anesthetics and benzodiazepines. Furthermore, they are highly unique in that they are the only superfamily of neurotransmitter-gated ion channels that contain members that are highly cation-selective, such as serotonin and acetylcholine receptors, and other members that are highly anion-selective, such as glycine and GABA receptors. This feature of pLGICs makes the family of proteins a model 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. Because charge selectivity is highly important for understanding how neuronal signaling can be either excitatory or inhibitory, this work has broad implications for understanding the molecular basis of neuronal processes.

METHODS & 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, we performed free-energy calculations on a variety of members of the pLGIC superfamily. Using NAMD, we completed umbrella sampling simulations 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. These simulations led us to conclude that major barriers of the free-energy landscape appear in fundamentally different locations in anion- versus cation-selective channels.

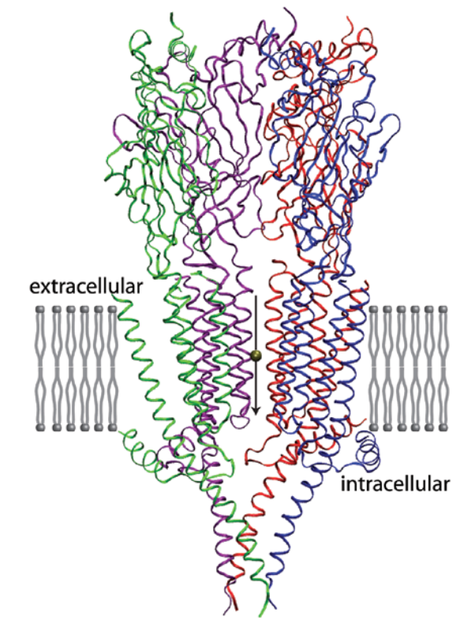

FIGURE 1: An ion, shown in gold, moves along the permeation pathway of the serotonin receptor (pdb id 4PIR), a member of the pentameric ligand-gated ion channel superfamily. One of the five subunits is removed to clearly visualize the ion, and each of the four remaining subunits is colored green, purple, red, or blue.

WHY BLUE WATERS?

Ion permeation, especially in the absence of a membrane potential, can be a rather slow process. Therefore, computational methods such as umbrella sampling have been devised in order to make the process of sampling the energy landscape of permeation much more efficient. However, while these techniques effectively use resources, the computational cost of free-energy landscapes is tremendous. In order to effectively sample the free-energy landscape for multiple ions in different channels with and without mutations, supercomputing time on Blue Waters is essential. Without Blue Waters, the project would need to be much less ambitious and smaller in scope, but with Blue Waters we are able to have a more profound understanding of the molecular basis of charge selectivity in pLGICs.

NON-BORN–OPPENHEIMER EFFECTS BETWEEN ELECTRONS AND PROTONS

Allocation: Blue Waters Professor/0.24 Mnh
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 used Blue Waters to perform NEO calculations on systems in which all electrons and one proton are treated quantum mechanically. In addition, we have developed a methodology to study the non-adiabatic dynamics of photo-induced proton-coupled electron transfer (PCET) reactions, using grid-based methods to calculate the nuclear wavefunction for the transferring proton, and 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 non-adiabaticity has been shown to play a critical role in proton-coupled electron transfer (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 nuclear–electronic orbital (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 orbitals, and certain exchange terms are approximated. Additionally, we have proposed a restricted basis set for the electronic orbitals that are explicitly correlated to the nuclear orbitals. The restricted basis set is a subset of the full electronic atomic orbital basis set, and its use results in a large decrease in the number of multi-particle integrals that must be computed and stored during a NEO-RXCHF calculation.

We performed NEO-RXCHF calculations on Blue Waters on HCN and FHF^- and compared the nuclear density of the proton to benchmarking grid-based calculations. The test calculations indicate that the NEO-RXCHF method provides accurate descriptions of the protons that are treated quantum mechanically. The use of a restricted basis set for the explicitly correlated electronic orbitals results in a speed-up of more than an order of magnitude with nearly no loss in accuracy compared to calculations performed without a restricted basis set. Current efforts are focused on extending the NEO-RXCHF method to multi-configurational systems and the implementation of a new integral code that will further speed up future NEO-RXCHF calculations.

The methodology for investigating the mechanism of condensed-phase photo-induced PCET reactions involves on-the-fly non-adiabatic dynamics on electron–proton vibronic surfaces using surface hopping in a hybrid quantum mechanical/molecular mechanical (QM/MM) framework. A semi-empirical implementation of the floating occupation molecular orbital complete active space configuration interaction method is used to obtain the electronic states of the solute, while the Fourier grid Hamiltonian grid-based method is used to calculate the vibrational states of the transferring proton. The adiabatic vibronic states are obtained by diagonalization of the vibronic Hamiltonian matrix in the basis of these electronic and vibrational states. The solvent is modeled explicitly using a QM/MM approach.

We performed test calculations with a quantum mechanical proton on a hydrogen-bonded phenol–amine complex in 1,2-dichloroethane solution. Our previous studies of this system treated the proton classically. Preliminary results support the characterization of an excited electronic state on which proton transfer is highly favored, in agreement with experimental results and results from simulations with a classical proton (fig. 1). Current efforts are directed toward obtaining experimentally testable predictions, such as hydrogen/deuterium kinetic isotope effects.

WHY BLUE WATERS

The NEO-RXCHF method requires the calculation and storage of trillions of integrals. Our in-house NEO code has been parallelized using a hybrid MPI/OpenMP protocol but still requires a large number of processors and a substantial amount of memory. The use of a large number of processors on Blue Waters is critical to speeding up the computational bottleneck of the calculation of the multi-particle integrals; even more important is the large amount of memory available on Blue Waters because, in contrast to other computer systems, a large number of nodes can be used simultaneously.

For studying photo-induced PCET with a quantum mechanical treatment of the transferring proton, the computational bottleneck is the calculation of the energies, forces, and non-adiabatic couplings between the electronic states at the proton grid points (usually 24) for each molecular dynamics time step, as well as the calculation of the non-adiabatic couplings between the vibronic states. We have used MPI to distribute these calculations over the cores of a single node on Blue Waters. The surface-hopping algorithm requires propagation of a large number of trajectories, which is made possible by the simultaneous use of a large number of nodes on Blue Waters. Moreover, the large amount of 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.

PUBLICATIONS

Sirjoosingh, A., M. V. Pak, K. R. Brorsen, and S. H. Hammes-Schiffer, Quantum Treatment of Protons with Reduced Explicitly Correlated Hartree-Fock Approach. *J. Chem. Phys.*, 142 (2015), 214107, doi:10.1063/1.4921303.

Brorsen, K. R., A. Sirjoosingh, M. V. Pak, and S. H Hammes-Schiffer, Nuclear–Electronic Orbital Explicity Correlated Hatree-Fock Approach: Restricted Basis Sets and Open-Shell Systems. *J. Chem. Phys.*, 142 (2015), 214108, doi:10.1063/1.4921304.

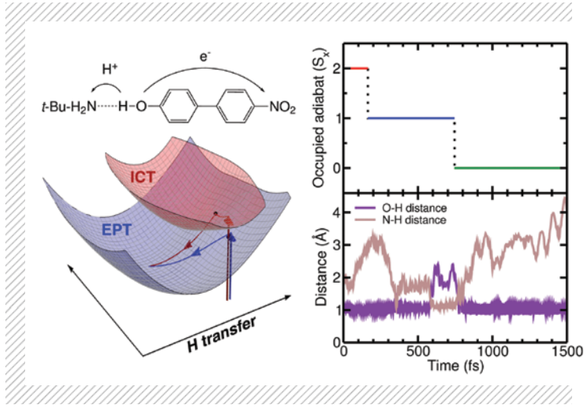

FIGURE 1: *Top left:* The hydrogen-bonded complex between p-nitrophenylphenol and t-butylamine was experimentally studied in 1,2-dichloroethane solution.

Bottom left: The two excited electronic states probed by photoexcitation of this complex.

Right: A representative surface-hopping trajectory in which the transferring proton is treated classically, initiated on S_2.

ICE AND WATER

Allocation: Blue Waters Professor/0.15 Mnh
PI: So Hirata[1]
Collaborators: Soohaeng Y. Willow[1], Michael A. Salim[1]

[1]University of Illinois at Urbana–Champaign

EXECUTIVE SUMMARY:

Predictive condensed-phase simulations at an accurate first-principles (*ab initio*) theoretical level, treating all electrons quantum mechanically, has previously been unthinkable. We performed just such simulations for a whole range of structural, dynamical, thermodynamic, and response properties of ice and liquid water, undoubtedly the most important condensed-phase systems on Earth. An algorithmic breakthrough (the embedded fragmentation technique) combined with the massive computational power of Blue Waters made these simulations possible. The predominant solid phase of water (ice Ih) displays unusual behaviors such as negative thermal expansion at low temperatures and an anomalous volume isotope effect. Our *ab initio* calculation computationally reproduced both, while also hinting at a pressure-induced amorphization. For liquid water, structures, self-diffusion coefficients, infrared and Raman spectra, and the Raman non-coincidence effect were computationally explained. The calculations revealed an atomistic detail of the hydrogen-bond network dynamics in the liquid.

INTRODUCTION

Chemistry of water in all three phases largely defines our planetary environment. Its influence is felt in subjects ranging from geology to climate, to biology and ecology, to geopolitics and history. Such a seemingly simple phenomenon—that the water volume collapses upon melting—has an immense impact on every structure and reaction found on the planetary surface and thus on every life form, but it derives from a subtle interplay of water's peculiar chemical bonding and dynamics. To this day, computationally determining some of the most delicate thermodynamic and response properties of ice and liquid water, such as melting temperature, from first principles (i.e. with no reference to any experimental data) has not been possible.

In this work, we combined the massive computational power of Blue Waters with an algorithmic breakthrough (embedded fragmentation) that made *ab initio* quantum chemistry calculations vastly faster, scalable in parallel, and thus applicable to condensed-phase systems. With it, we addressed structural, dynamical, thermodynamic, and response properties of ice and liquid water. We used *ab initio* second-order many-body perturbation theory using a Gaussian-type basis set. The theory can account for all covalent, ionic, hydrogen-bond, and dispersion interactions from first principles.

For the dominant phase of ice on Earth (ice Ih), we aimed to computationally explain its unusual behavior: the negative thermal expansion at temperatures below 70 K and the anomalous volume isotope effect (which means that heavy water has a greater molar volume than normal water). We also monitored a possible pressure-induced phase transition from a crystalline, proton-disordered (ice Ih) phase to an amorphous phase. These rather strange behaviors may well be related to the negative thermal expansion and anomalous volume isotope effect also found in liquid water.

We calculated the structure (radial distribution function, coordinate number, dipole moment), dynamics (self-diffusion coefficient, fluctuation of coordination number and dipole moment), and response (infrared and Raman spectra) properties of liquid water. Some of these properties reflect the local hydrogen-bond environment and its dynamics, which is often difficult to probe experimentally.

METHODS & RESULTS

We used the embedded fragmentation method [1]. This method first divides molecular crystal or liquid into overlapping molecular dimers embedded in the self-consistently determined electrostatic field of the crystal or liquid. Next, energies, atomic forces, force constants, dipole moments, infrared and Raman intensities, etc., are calculated for each dimer in parallel. These data are then combined to yield the corresponding properties of the bulk crystal or liquid at finite pressures and/or temperatures.

This method has been used successfully in the past, in conjunction with *ab initio* second-order many-body perturbation and/or coupled-cluster singles and doubles levels, to study the structure and spectra of ice Ih [2] and ice VIII [3] as well as the structure, spectra including pressure tuning of Fermi resonance [4], solid–solid phase transition [5], and thermal expansion [6] of dry ice (solid carbon dioxide phase I).

For ice Ih, our simulation reproduced the negative thermal expansion at low temperatures, which turned positive at higher temperatures. We found that this peculiar behavior was caused by the negative pressure dependence of the frequencies (Grüneisen parameters) of acoustic phonons, reflecting the sparse structure of the hydrogen-bond cage of ice. The anomalous volume isotope effect was reproduced correctly with only one choice of embedding field but not with another, suggesting that it is a result of a delicate balance between the competing pressure effects on vibrational and O–H stretching phonons. Furthermore, a pressure-induced volume collapse and corresponding softening of acoustic phonons were detected at a pressure similar to the one at which pressure-induced amorphization is observed experimentally. It is possible that we computationally observed a mechanical instability precursor to such a transition to the so-called high-density amorphous phase.

For liquid water, we performed probably the first molecular dynamics simulation using on-the-fly atomic forces evaluated by the *ab initio* electron-correlated molecular orbital method. The calculated radial distribution function and self-diffusion coefficient were in excellent agreement with the observed results. The simulated infrared and Raman spectra correctly predicted the shapes and widths of the O–H stretching bands, including the non-coincidence of the isotropic and anisotropic Raman components, which are known to reflect the local hydrogen-bond environment. The simulation also provided unique insights into the large fluctuation of the coordination number and the mechanism and time scale in which water molecules exchange their positions between first and second solvation shells, causing such large fluctuation. Our first-principles simulation, therefore, allowed us to study the electronic details of the structure and dynamics of liquid water with unprecedented accuracy.

WHY BLUE WATERS?

Condensed-phase applications of systematic *ab initio* electron-correlated molecular orbital theory, which goes beyond empirical force fields or even density-functional approximations, have long been unthinkable. Both algorithmic innovations and supercomputing resources at the scale of Blue Waters were essential for the project.

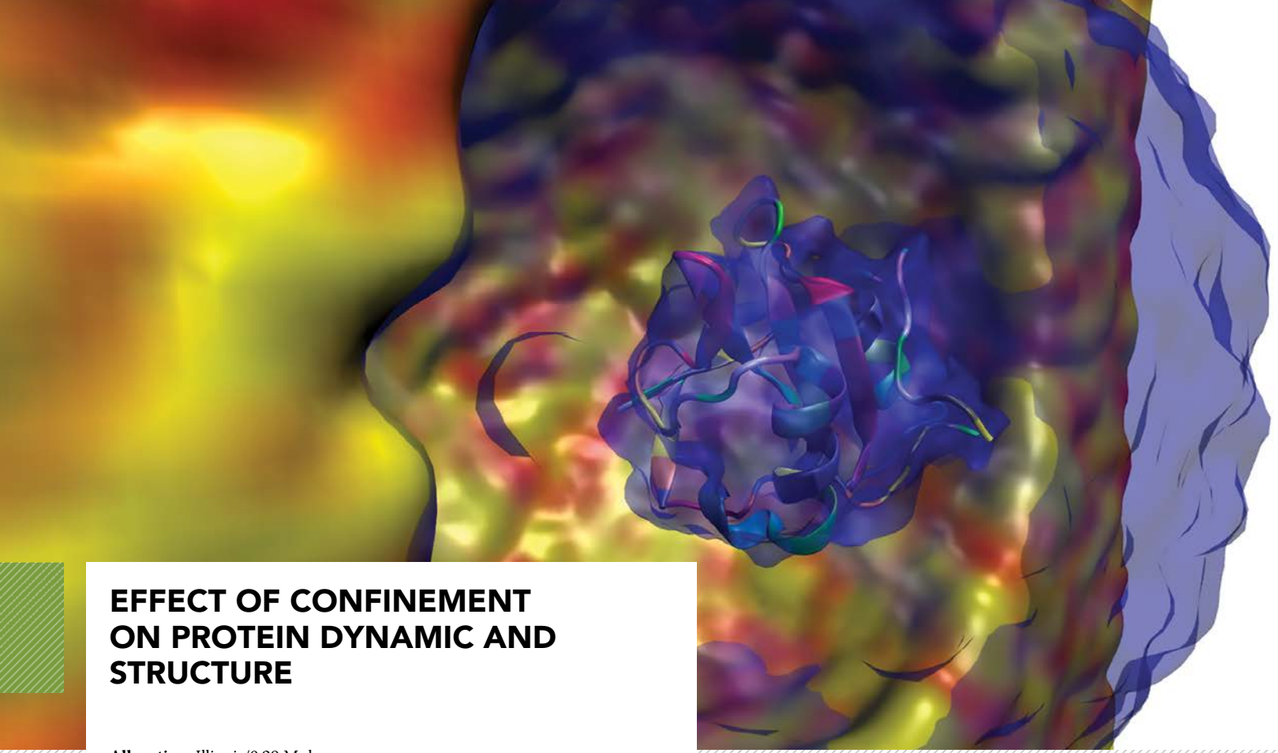

EFFECT OF CONFINEMENT ON PROTEIN DYNAMIC AND STRUCTURE

Allocation: Illinois/0.29 Mnh
PI: Eric Jakobsson[1]
Co-PIs: Narayana R. Aluru[2], Amir Barati Farimani[2], Jie Sun

[1]Beckman Institute for Advanced Science and Technology
[2]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY:

Understanding structure function relationships of proteins is a major goal of biology, biomedicine, and biotechnology. For soluble proteins, simulations have been done in an aqueous environment, essentially assuming that the molecules are in a bath of infinite size. However, proteins in cells function in a confined environment. We believe this is a factor in the many observations that protein interactions in cells differ from the same interactions in solution. In our work we confined water and proteins in silica cavities to explore the effects of the confinement on structure and dynamics. Specifically, we observed a large influence of confinement on water, bovine pancreatic trypsin inhibitor (BPTI), and on the signaling protein SHP2. For water we saw ordered ice-like structures induced at a physiological temperature where bulk water is fluid. For proteins we saw significant structural modifications and reduction in flexibility, even when the cavity was substantially larger than the largest protein dimension.

INTRODUCTION

Proteins are essential components of living cells. One of the limiting factors in our understanding of proteins is a failure, so far, to replicate the environment of the cellular interior. For soluble proteins, dynamic simulations have been done in an aqueous environment, effectively assuming that the molecules are in a bath of infinite size. However in living cells the motions of reactive proteins are constrained by surrounding molecular scaffolds. Possibly as a result of such constraints, the kinetics of soluble protein reactions in cells (*in vivo*) is often significantly different from that in aqueous environment (*in vitro*). It may be that structures are somewhat different as well.

To shed light on the structure and dynamics of proteins, we confined pure water and the proteins (BPTI, SHP2) in silica-bounded cavities.

Because of its relatively small size, biological importance (as a protease inhibitor), and ready availability, bovine pancreatic trypsin inhibitor (BPTI) has for decades served as a model system for exploring issues of protein structure and dynamics, for example the subject of the very first molecular dynamics simulation of a protein in 1977 [1].

The choice of the larger protein SHP2 for the next part of our study was due to the fact that large biologically important protein conformational changes monitored in this molecule by FRET have much more rapid kinetics *in vivo* than *in vitro* [2]. We will seek to understand this computationally.

We expect this work to make a large contribution towards bridging the gap between structure and dynamics of proteins *in vivo* and *in vitro*.

METHODS & RESULTS

In the current study we explored a native and mutant BPTI in which the mutation confers added flexibility to the protein [3]. We confirmed that our molecular dynamics in bulk replicates the reported experimental results in solution, and then explored the effects of confinement on both the native and the mutant strains.

Confinement (encapsulation in silica cavity) has a large effect on the dynamics of BPTI. The mutant of BPTI (G37A) exhibits large flexibility compared to its wild type. Both wild type and mutant BPTI 's RMSD drops when placed in confinement. The water structure is dramatically changed due to confinements and density layering is observed near silica wall. The peak density of water in density layering is augmented due to the overlap of the protein hydration shell and near-wall ordering.

WHY BLUE WATERS

Our largest simulations so far involve just under 100,000 atoms and all our simulations involve much conformational sampling because we are interested in significant deviations from crystal structure. In the next stage of the project, the systems involving the complete SHP2 protein, including fluorescent probes, will be significantly larger. These expensive computations are not possible to perform in reasonable time without a petascale supercomputer.

PUBLICATIONS

B. Farimani, N. R. Aluru, Eric Jakobsson, Effect of confinement on protein dynamic and Structure, Biophysical Society Meeting, Baltimore, MD, (2015).

FIGURE 1 (BACKGROUND): BPTI in a smaller amount of water confined in a silica cavity . The image (which is a cutaway section view) shows the protein confined in a cavity, which emulates the crowding in a cell.

MULTISCALE MODELING OF BONE FRACTURE AND STRENGTH

Allocation: Illinois/50 Knh
PI: Iwona Jasiuk[1]
Co-PI: Seid Koric[1,2]
Collaborator: Dan Dragomir Daescu[3]

[1]University of Illinois at Urbana–Champaign
[2]National Center for Supercomputing Applications
[3]Mayo Clinic, Rochester, Minn.

EXECUTIVE SUMMARY:

We are using Blue Waters to run a multiscale computational model of bone fracture and strength. Fracture is a multiscale phenomenon with cracks initiating at the atomic level. Thus, we are modeling bone as a biological material with a hierarchical structure spanning from the atomic level (crystal structure of minerals and atomic structure of collagen), to nanoscale (mineralized collagen fibril), sub-micron scale (single lamella), microscale (lamellar bone), mesoscale (cortical and trabecular bone types), and macroscale (whole bone) (fig. 1). We are conducting atomic- and nanoscale-level simulations using molecular dynamics software NAMD with CHARMM and LAMMPS and higher scales by employing finite-element software Abaqus. The inputs are taken from experiments. This will be the first experimentally based multiscale modeling of bone fracture and strength. Blue Waters is necessary because such multiscale model requires very high computational resources due to the hierarchy of bone structure and complex fracture phenomena in bone.

INTRODUCTION

We are using Blue Waters to test a novel multiscale computational approach to predict fracture and strength of normal versus osteoporotic bone. Osteoporosis is a bone disease characterized by low bone density and deterioration of bone structure leading to bone fragility and increased risk of fractures [1]. In the United States osteoporosis is a major public health threat for an estimated 44 million people. It is a silent disease with no symptoms prior to fractures and no cure, but treatments can slow its progress. Thus, early and accurate diagnosis is crucial. Currently, bone quality is assessed clinically by measuring the bone mineral density, although other factors such as bone's complex hierarchical structure (fig. 1) also contribute to bone's properties. A new approach is sorely needed for more accurate diagnosis of osteoporosis. A computational mechanics model can provide a new tool for the clinical assessment of bone.

About Bone Structure

Bone is a multifunctional biological tissue that has an ideal combination of properties when healthy: high stiffness, strength, and fracture toughness, and light weight. These superior properties are due in part to the complex hierarchical structure of bone from macroscopic (whole-bone) to atomistic levels [2] (fig. 1). At the mesoscale, one step smaller than the macroscale, bone tissue is composed of the dense cortical bone and the spongy trabecular bone. Mature human cortical bone consists of osteons embedded in an interstitial bone and surrounded by a circumferential bone, whereas the trabecular bone is made of a porous network of trabeculae. At the microscale both cortical and trabecular bones have lamellar (layered) structures formed through stacking of lamellae in different orientations. At the sub-microscale a single lamella is made of preferentially oriented mineralized collagen fibrils perforated by ellipsoidal cavities called lacunae. At the nanoscale the mineralized collagen fibril is a composite structural unit consisting of the collagen, nano-sized hydroxyapatite crystals, water, and a small amount of non-collagenous proteins. The sub-nanoscale represents the atomistic scale of bone's constituents: tropocollagen molecules and crystals.

Similarly, bone strength and fracture is a very complex phenomenon with different failure mechanisms exhibited at different structural scales (fig. 2). They include crack deflection and twist, uncracked ligament bridging, collagen fibril bridging, and constrained micro-cracking.

Furthermore, there are two types of toughening mechanisms in bone (fig. 3): 1) *intrinsic* mechanisms that act ahead of the crack tip and increase the microstructural resistance to crack initiation and growth, and 2) *extrinsic* mechanisms that act primarily behind the crack tip to inhibit crack growth by effectively reducing the crack-driving force actually experienced at the crack tip.

METHODS & RESULTS

The objective of this study is to test a complex 3D multiscale model of bone strength and fracture. Modeling starts at the atomistic level and moves up the scales to the macroscale. Atomic- and nanoscale-level simulations are done using the molecular dynamics software NAMD with CHARMM and LAMMPS, and larger scales were added using the finite-element software Abaqus. We have preliminary results from molecular dyanmics and obtained properties of collagen molecules. We are now finalizing our nanoscale model that can feed in the data for the larger-scale model.

At the sub-microscale we are modeling bone as a collection of mineralized collagen fibrils. We are generalizing an elastic-case model to use in a nonlinear case with deformations up to fracture. We have already modeled lamellar bone at the microscale and are currently simulating trabecular bone up to failure with the goal of obtaining anisotropic constitutive models for the macroscale. At the macroscale we are modeling a whole human femur.

This research is in collaboration with Dragomir Daescu from Mayo Clinic. We received the image of the whole femur bone from Mayo Clinic. We will be using this geometric model for our multiscale simulations. Our preliminary models of bone elasticity and strength are given in [3–5].

WHY BLUE WATERS?

All these factors make bone fracture and strength a very complex and highly computationally intensive problem to study. The modeling requires a multiscale approach accounting for all these different structural features and failure modes. Such a model is intractable using current high-end computers. Thus, Blue Waters is crucial for such computations. No comprehensive multiscale model yet exists in the literature due to these computational challenges and limitations.

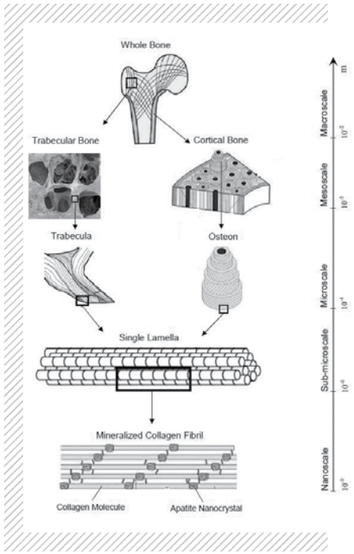

FIGURE 1: Hierarchical structure of bone.

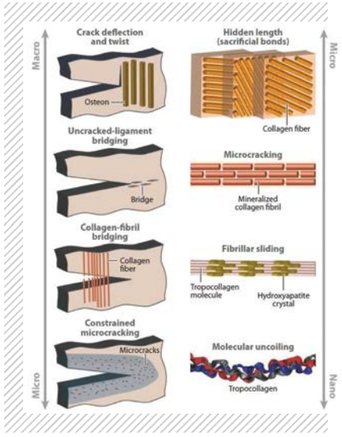

FIGURE 2: Different fracture and failure mechanisms in bone.

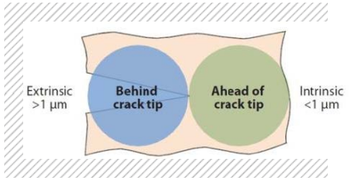

FIGURE 3: Two types of toughening mechanisms in bone.

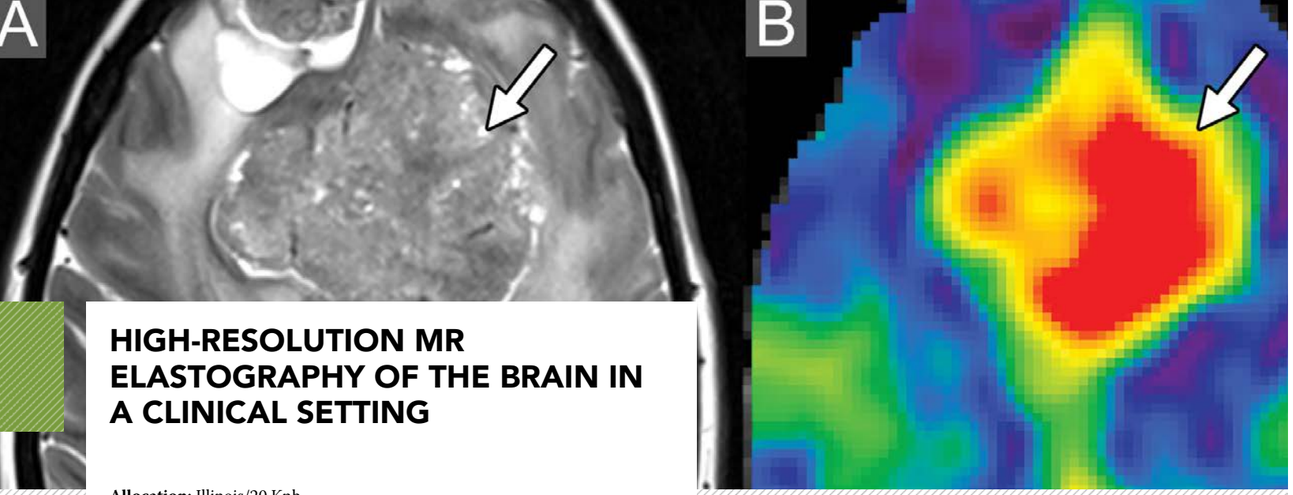

HIGH-RESOLUTION MR ELASTOGRAPHY OF THE BRAIN IN A CLINICAL SETTING

Allocation: Illinois/20 Knh
PI: Curtis L. Johnson[1]
Collaborators: Graham R. Huesmann[1,2], William C. Olivero[1,2], Bradley P. Sutton[1], Hillary Schwarb[1], Matthew DJ McGarry[3], Aaron T. Anderson[1]

[1]University of Illinois at Urbana-Champaign
[2]Carle Foundation Hospital
[3]Dartmouth College

FIGURE 1 (BACKGROUND): Imaging of a meningioma located on the falx in a 43-year-old patient. (A) T_2-weighted structural MRI with arrow highlighting location of the tumor, and (B) viscoelastic shear stiffness map obtained using high-resolution MRE demonstrating the very high stiffness of the tumor.

EXECUTIVE SUMMARY:

Our group seeks to exploit the inherent mechanical contrast in diseased tissue for diagnosing neurological conditions. This is achieved using an imaging method called magnetic resonance elastography (MRE), which allows for noninvasive characterization of brain tissue mechanics. However, obtaining accurate measures of tissue properties requires advanced imaging procedures and inversion algorithms and thus correspondingly significant computational power. We use the Blue Waters system to process the MRE images and generate viscoelastic property maps for applications in neurology, neurosurgery, and neuroscience in partnership with the Carle Foundation Hospital and the Biomedical Imaging Center at the University of Illinois' Beckman Institute. Here, we highlight clinical research in using brain tissue mechanics for the pre-surgical evaluation of intracranial tumors and the hippocampus in temporal lobe epilepsy.

INTRODUCTION

Medical diagnostics have a long history of using the mechanical properties of tissues to identify the presence of pathology through manual palpation. Pathology can result in tissue stiffness and viscosity changes up to an order of magnitude relative to healthy tissue, thus making tissue mechanical properties potentially very sensitive markers of tissue health if they can be noninvasively measured. Magnetic resonance elastography (MRE; [1]) is a technique capable of probing tissue viscoelasticity through the imaging of nondestructive shear waves and the solution of an inverse problem to determine the underlying material properties from these waves. We have developed MRE methodology that allows for brain tissue mechanics to be investigated at a high resolution—revealing neuroanatomical features previously unseen [2]—and that can potentially result in powerful clinical techniques.

The challenge in performing high-resolution MRE scans in a clinical setting is the significant computational power needed to generate the viscoelastic property maps. In order to fit such a technique in the radiological workflow, results need to be generated in minutes rather than days. Here we present clinical research results obtained using Blue Waters that significantly reduced computation time. In particular, we discuss the use of MRE in the pre-surgical evaluation of intracranial tumors and the hippocampus in temporal lobe epilepsy.

METHODS & RESULTS

All scans used a 3D multislab, multishot MRE sequence [3] for whole-brain harmonic shear wave imaging on a Siemens 3T Trio scanner. This data is the input to a finite element-based shear wave inversion algorithm to estimate viscoelastic mechanical properties, which is termed nonlinear inversion (NLI) [4]. The NLI algorithm divides the object of interest into overlapping subzones and performs an interleaved series of local subzone property updates executed in parallel and global rezoning operations in order to achieve convergence on the solution. This code is executed on Blue Waters, taking advantage of the massive parallelization available for updating all subzones simultaneously. A typical problem uses 300 subzones, each assigned to its own processor, and the memory-intensive finite-element operations require 1 GB of memory per core. Depending on the exact imaging and inversion parameters, complete computations for each dataset requires between 20 and 100 node hours.

Our initial results on clinical populations have been very promising. We have investigated patients with different types of intracranial tumors in order to characterize the mechanical properties of the lesions prior to surgery. Knowing the stiffness of tumors allows surgeons to appropriately plan their surgery based on how difficult they expect the resection to be, ultimately making the procedures safer for patients. Figure 1 shows the stiffness map of a patient with a meningioma on the falx (the fold of dura mater that descends in the fissure between the brain's two hemispheres). This tumor has a maximum stiffness of approximately 8.3 kPa—nearly three times that of normal brain tissue—and would require an extensive surgical resection procedure that could be appropriately and safely planned.

We have also explored using MRE to identify the presence of a specific form of temporal lobe epilepsy termed mesial temporal sclerosis (MTS). This condition is marked by degeneration of the hippocampus, which serves as the epileptogenic source, and we hypothesize that this results in a detectable change in the mechanical properties of the tissue. The most effective treatment for MTS is resection of the hippocampus, which can eliminate or reduce seizure activity and significantly improve length and quality of life for treated patients. However, this surgery is often not performed or is severely delayed due to the difficulty in confirming MTS diagnosis with standard imaging techniques. Our initial results suggest that the mechanical contrast observed with MRE may be a more sensitive marker for the presence and lateralization of MTS than other imaging techniques (Figure 2).

The development of reliable methods for the pre-surgical evaluation of brain tissue through mechanics represents a potentially significant advance for clinical practice. As evidenced by the results presented here, allowing neurologists and neurosurgeons access to viscoelastic property maps can greatly improve surgical treatment plans for a host of neurological conditions. The continued development of methodology to both improve the mechanical property measures and remove barriers impeding its adoption in clinical practice, such as oppressive computation time, could ultimately allow for the translation of the MRE technique into a standard patient care procedure.

WHY BLUE WATERS

Blue Waters is an integral part of the MRE efforts both at Carle and at Beckman as it provides the computational power necessary to achieve high-resolution viscoelastic property maps in an acceptable timeframe. As MRE becomes more widely accepted by physicians looking to use the inherently sensitive mechanical contrast in pathological tissue, Blue Waters will allow for images to be generated for radiological interpretation in minutes rather than days. Ultimately, we aim to develop a pipeline through a direct connection between the MRI scanner and Blue Waters to seamlessly integrate MRE with NLI in the clinical imaging workflow.

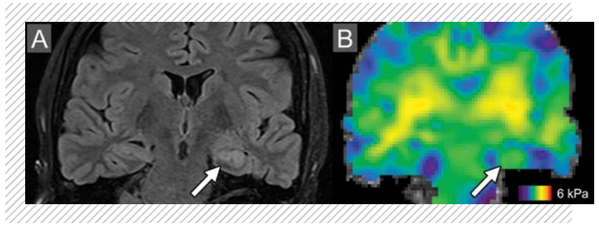

FIGURE 2: Imaging of the hippocampus in a 29-year-old patient with mesial temporal sclerosis epilepsy. (A) Coronal T_2-FLAIR MRI with arrow highlighting signal hyperintensity in left hippocampus, which is a signature of tissue pathology and the epileptogenic source. (B) Viscoelastic shear stiffness map with the epileptogenic hippocampus highlighted as a clearly stiff region.

FUNCTIONAL MECHANISM OF THE BACTERIAL EFFLUX PUMP

Allocation: GLCPC/0.43 Mnh
PI: Fatemeh Khalili-Araghi[1]
Collaborator: Arman Fathizadeh[1]

[1]University of Illinois at Chicago

EXECUTIVE SUMMARY:

Gram-negative bacteria such as *E. coli* utilize a tripartite complex system to expel toxic chemicals and antibiotics directly out of the cell. The efflux system, consisting of an inner membrane transporter, an outer membrane channel, and a fusion protein connecting the two domains, uses the proton motive force to extrude a wide variety of chemical compounds out of the cell, thus conferring resistance to a broad spectrum of antibiotics. We have built an atomic model of the tripartite complex based on the electron microscopy images of the pump and the crystal structures of its individual components, which provide a complete structure of the transport complex. Using molecular dynamics simulations, we have identified residues that may play a role in proton transport in two different pumps and determined their possible impact on the drug binding sites in the complex.

INTRODUCTION

In Gram-negative bacteria, such as *E. coli*, toxic chemicals are expelled through tripartite efflux pumps that span the inner and outer membranes [1]. The efflux system of the resistance-nodulation-division family plays a major role in the intrinsic and acquired resistance of bacteria to antibiotics. These systems provide a pathway for bacteria to export lethal concentrations of drugs and metal ions out of the cell membrane.

The complex system shown in fig. 1 is made of an energy-utilizing inner membrane transporter and an outer membrane channel that are connected by membrane fusion proteins. The inner membrane protein uses the proton motive force across the membrane to undergo conformational transitions that capture the drug from inside the cell. The substrate is then released into a tunnel that is formed by the fusion proteins and is transported all the way to the extracellular side.

The structures of individual components of the complex system have been determined by X-ray crystallography. However, it is not yet clear how the three components assemble into a tripartite system that can transport antibiotics or metal ions out of the cell. More importantly, the transport pathway of the drugs or protons through the complex and their coupling mechanism is not yet known.

METHODS & RESULTS

We used molecular dynamics simulations to construct an all-atom model of the complex system based on the crystal structures of individual protein components and a cryo-electron microscopy density map of the efflux pump from *E. Coli* [2]. The complex system involving more than a million atoms was equilibrated in explicit solvent and membrane using the program NAMD [3]. The resulting stable model provided the first atomic description of the tripartite efflux pump.

To further characterize the proton transport mechanism of the pump, we carried out extensive molecular dynamics simulations of the system for a total of several microseconds. The simulations identified possible proton transport pathways and key residues in the inner membrane protein that are involved in the process. The simulation of the complex system showed how transfer of protons through inner membrane pathways could be coupled to conformational changes of known binding sites near the fusion tunnel.

We are extending the simulations to study the transport of an antibiotic and a cancer treatment drug through the entire channel and to identify their transport pathway from the proximal binding site to the distal binding site [4] in the inner membrane protein and finally through the fusion tunnel leading to the extracellular channel. These simulations will provide insight into the underlying transport process and will determine potential barriers or binding sites along the pathway. It is imperative to include the three protein components in studying the transport process, as the presence of the fusion proteins is known to stabilize the initial drug binding sites, and in turn binding of the drug is expected to affect the opening of the extracellular channel through allosteric interactions between the three components.

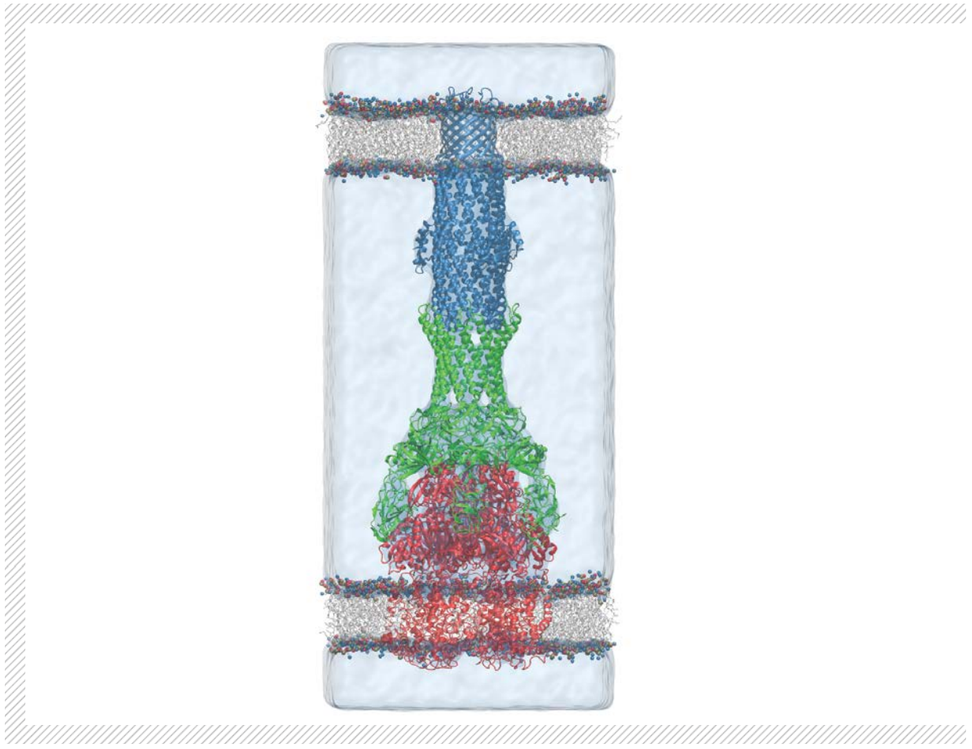

FIGURE 1: The bacterial efflux pump consisting of two lipid membranes and the tripartite assembly of AcrB/AcrA/TolC shown in red, green, and blue, respectively.

WHY BLUE WATERS?

The relatively large size of the assembled structure will make it one of the largest transport systems being studied. Obtaining meaningful statistics on the transport process are only feasible by employing the high number of cores available on Blue Waters.

FUNCTIONAL MECHANICS OF VOLTAGE GATED LIKE ION CHANNELS

Allocation: NSF PRAC/5.6 Mnh
PI: Michael L. Klein[1]
Collaborators: Christina Bergonzo[1], Rodrigo Galindo-Murillo[1], Daniel R. Roe[1]

[1]College of Science and Technology, Temple University

EXECUTIVE SUMMARY:

Ion channels are ubiquitous transducers of chemical and electrical stimuli in cells. In voltage-gated-like ion channels, conformational changes in a domain responsible for sensing a stimulus (transducer) affect the state of a gate domain (effector) that open and close, controlling the flux of permeant ions. The details of this allosteric communication are of great relevance: Pharmacological modulators of ion channels, including neurotoxins and anesthetics, are often allosteric drugs, i.e. they interfere with this coupling between the transducer and effector domains. Our calculations start to clarify the molecular underpinnings of allosteric signal propagation in voltage-gated-like ion channels. We performed molecular dynamics simulations on two classes of channels: transient receptor potential (TRP) and voltage gated cation channels. We found that, despite having similar architectures, the two classes of channels are characterized by distinct mechanical properties; accordingly each class shows unique conformational dynamics, which makes each channel responsive to a specific stimulus, i.e. ligand binding or external field.

INTRODUCTION

Ion channel proteins serve as fast sensors and transducers of chemical and electrical stimuli in cells and underpin electrical excitability. One particular ion channel family, the voltage-gated-like ion channels (VGLCs), occupy a particularly wide set of roles in the physiology of multicellular organisms. Differentiation of a single homologous transmembrane protein template has allowed an explosion of functional variability, from reporting noxious environmental conditions, to shaping the neuronal action potential, to syncing the beating of the heart. Allostery describes how certain nonlocal energetic perturbations couple to and alter the equilibrium distribution of conformations that a macromolecule visits. In ion channels, allostery provides a framework for understanding how conformational changes in domains responsible for sensing stimuli couple to channel gate domains that open and close, controlling the flux of permeant ions. Physiologically-relevant ion channel modulators, including neurotoxins and anesthetics, are often allosteric. Determining the origins of these effects in the VGLC family is crucial for controlling and correcting complex impairments in various electrically-excitable tissues. We focused on allosteric regulation of two subfamilies of the VGLC superfamily in which pore gating is regulated by different stimulus: binding of agonists and antagonists, pH and heat variation for TRP channels and membrane depolarization for voltage gated ion channels.

METHODS & RESULTS

First, we investigated whether the allosteric mechanism underlying gating is common to all TRP channels, and how this mechanism differs from that underpinning Kv channel voltage sensitivity. Thus, we performed comparative sequence analysis on large, comprehensive ensembles of TRP and Kv channel sequences. We detected sequence features that were specific to TRP channels and, based on insight from recent TRPV1 structures, we hypothesized a model of TRP channel gating that differs substantially from the one mediating voltage sensitivity in Kv channels. The hypothesized mechanism involves the displacement of a defect in the H-bond network of S6 that changes the orientation of the pore-lining residues at the hydrophobic gate (Figure 1A). To test this hypothesis and to glean additional insight into this molecular process, we performed molecular dynamics simulations of the TRPV1 channel embedded in a lipid membrane and collected trajectories spanning μs time-scales. We observed a close-to-open transition in one of the subunits and, thanks to a statistical analysis of the pattern of H-bonds, we were able to trace back the conformational transition to a collective motion compatible with the one inferred from the sequence analysis (Figure 1B-C). These results were reported in the Journal of General Physiology.

We then turned our attention to the activation mechanism of voltage gated ion channels. To investigate the activation mechanism of this class of channels, we characterized the mechanical properties of a member of this family, NavAb [1] in the context of a novel approach to identify the set of dynamical domains in proteins. We performed molecular dynamics simulations of NavAb in different conformational states and embedded in a lipid bilayer over time-scales of few μs and analyzed each trajectory. NavAb is constituted by a pore domain, assembled from the last two transmembrane helices of each monomer (conventionally referred to as S5 and S6), and four separate structural domains, each comprising the first four helices of a monomer (S1–S4), which act as voltage sensors (Figure 2A). The results from the subdivision in dynamical domains suggest an activation mechanism by which the displacement of S4 results in a motion of the linker that releases the steric hindrance exerted on the pore domain in the resting/closed state (Figure 2B-C). This observation disfavors the alternative scenario in which the linker exerts an active pulling on the pore domain. These quantitative results, reported in the journal Structure, will be valuable for designing future studies aimed at elucidating by more direct means the mechanical workings of the pore complex.

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. Each 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 an applied electrostatic field or of a ligand. The capabilities of Blue Waters were key to the success of this computationally intensive project.

PUBLICATIONS

Palovcak, E., L. Delemotte, M. L. Klein, and V. Carnevale Comparative sequence analysis suggests a conserved gating mechanism for TRP channels *Journal of General Physiology*, 146:1 (2015), pp. 37-50

Ponzoni, L., G. Polles, V. Carnevale, and C. Micheletti SPECTRUS: A Dimensionality Reduction Approach for Identifying Dynamical Domains in Protein Complexes from Limited Structural Datasets *Structure*, 23:8 (2015), pp. 1516-1525

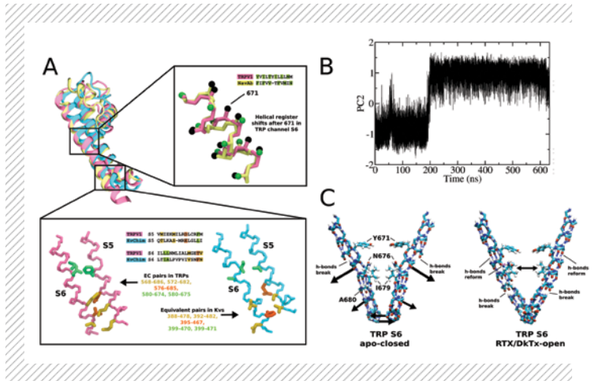

FIGURE 1: A helical distortion in the TRP channel pore.

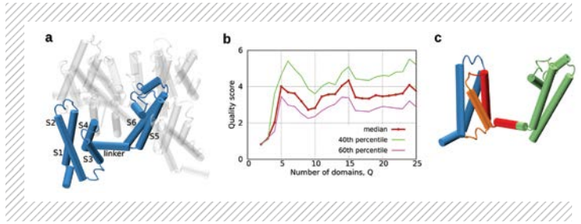

FIGURE 2: Domain Decomposition of the NavAb tetrameric ion channel.

SIMULATING RIBOSOME BIOGENESIS IN WHOLE CELLS

Allocation: Illinois/0.60 Mnh
PI: Zaida Luthey-Schulten[1,2,3,4]
Co-PIs: Tyler Earnest[1,2], Jonathan Lai[1], Ke Chen[1], Michael Hallock[1]
Collaborator: James R. Williamson[5]

[1]University of Illinois at Urbana–Champaign
[2]NSF Center for the Physics of Living Cells
[3]NIH Center for Macromolecular Modeling and Bioinformatics
[4]Beckman Institute
[5]Scripps Research Institute

EXECUTIVE SUMMARY:

Central to all life is the assembly of the ribosome: a coordinated process involving the hierarchical association of approximately 50 proteins to the RNAs forming the small and larger ribosomal subunits. Through the synthesis of data from disparate sources such as *in vitro* kinetics, cryo-electron tomography, single-particle diffusion assays, mRNA expression, and genomics data, we developed a spatially-resolved stochastic model of the biogenesis of the ribosomal small subunit in *E. coli*. Using our GPU-accelerated Lattice Microbes software [1–3] on Blue Waters, we observed biogenesis in modeled cells at timescales up to one hour. Our model reproduced the correct assembly times and predicted the spatial distribution of assembly intermediates. We will report on the results we have achieved to date and outline how further work using Blue Waters will provide us a window into cellular processes from single cells to colonies.

INTRODUCTION

Translation is the universal process that synthesizes proteins in all living cells. Central to translation is the ribosome, which itself constitutes approximately one-fourth of the bacterial dry mass. Biogenesis of the ribosome together with all cellular activities involved in translation consume a significant budget of the cell's total energy. However, the mechanism enforcing the balance between metabolism and macromolecular synthesis is yet to be fully resolved. An overarching goal of our research program is the development of large-scale whole-cell models including metabolism, gene expression, signaling, DNA replication, and cell division. Modeling the biogenesis of the ribosome is a major milestone towards the achievement of this goal. These models are crucial for our understanding of cell growth and how it is regulated in response to environmental perturbations.

The intricacy of ribosome assembly first attracted Nomura et al. [4], who originally mapped out the hierarchical dependency of the r-proteins binding to the *E. coli* 16S rRNA to form the small subunit (30S). Progress in biophysical approaches boosted our understanding of *in vitro* ribosomal self-assembly mainly for the protein-assisted dynamics of RNA folding [5–7], and the kinetic cooperativity of protein binding [8–10]. Both aspects agreed that assembly of *E. coli* proceeds through multiple parallel pathways nucleated at different positions on the rRNA, and a 5' to 3' directionality during assembly is always present. However, the order in which the r-proteins bind to assembly intermediates as inferred from thermodynamic and kinetic experiments can be contradictory, hampering our investigation of the assembly under an *in vivo* environment. Therefore, a comprehensive model that captures the topology of the protein RNA interaction network is needed to decipher the underlying rules governing the assembly of the ribosome.

METHODS & RESULTS

The stochastic *in vivo* model is constructed upon a deterministic *in vitro* assembly model. We built a chemical reaction network describing the binding of proteins to assembly intermediates using kinetic data from mass spectrometry experiments [10,11]. Since the assembly involves the binding of 20 r-proteins, naively, this would result in 2^{20} intermediate species and 20! (2.43 quadrillion) reactions. To mitigate combinatoric explosion, we use Nomura's assembly map to only include species and reactions that are allowed by the binding hierarchy. This reduces the model to 1,612 intermediates and allows for a thorough search of parameter space to fit the deterministic model to the kinetic data.

The 1,612-species model can be simplified further by removing species from the network that contribute the least to the overall assembly flux. By iterating this procedure, we were able to generate an assembly network containing

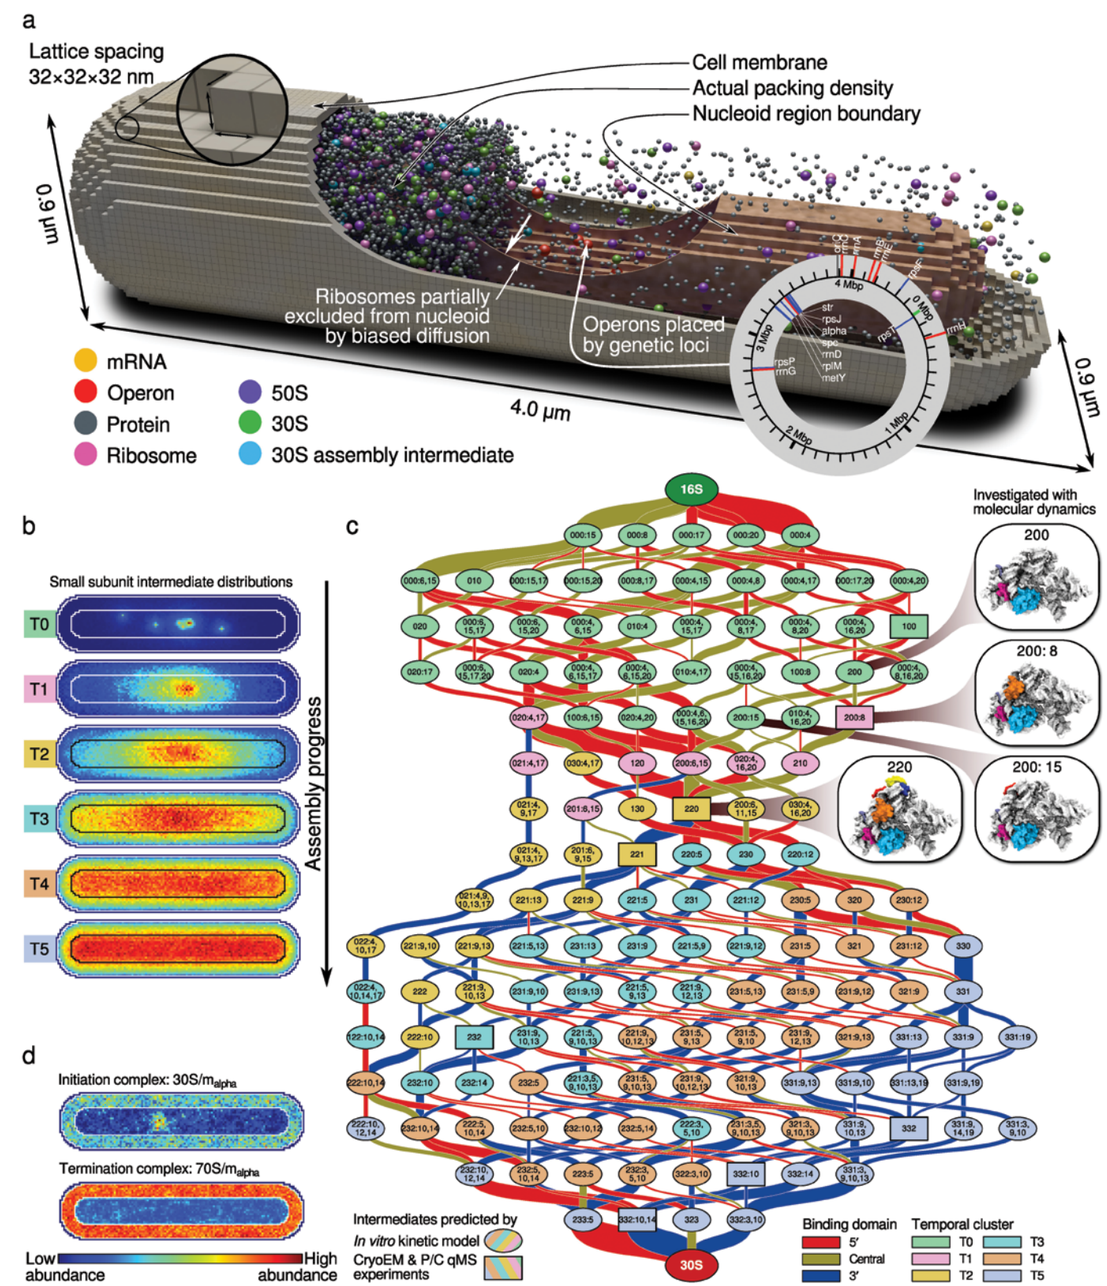

FIGURE 1: (a) Schematic of whole-cell stochastic model. (b) Heterogeneous distribution of 30S intermediates, clustered based on maximum abundance times. (c) Assembly network showing 30S intermediates as nodes and protein binding reactions as edges. Nodes are colored based on their clustering. (d) Distribution of translation initiation and termination events.

only 134 intermediates while reproducing the protein binding curves of the 1,612 parameter network up to an error of 1%. This network included the majority of the intermediates identified experimentally [11] and identified the observed 5' to 3' assembly order. We used NAMD to perform molecular dynamics simulations on these intermediates to investigate how they control the assembly process.

The cell architecture was based on slow-growing *E. coli* of dimensions 4.0 x 0.9 x 0.9 μm^3 with an inner DNA-containing nucleoid region. The cell was discretized into 32 nm^3 lattice sites. The operon sites for r-protein and rRNA were placed within the nucleoid region according to their position on the genome. Experimentally, ribosomes were observed to be excluded from the nucleoid region [12,13], so we biased the transition rates between the cytosol and nucleoid for ribosomes in order to recover this effect when DNA is omitted (fig. 1a). The diffusion constants for the mRNA, ribosomal subunits, full ribosomes, and proteins were taken from experiments [12,13].

We treated the transcription of ribosomal and messenger RNA explicitly and included the active decay of mRNA with rates derived from microarray experiments [14]. Translation of the r-proteins was modeled using explicit initiation and termination steps. The large subunit (50S) was included in the model, but participated only in assembly of the translating ribosome. Though cell division is a planned addition to this model, in the meantime we are treating the loss of species due to dilution as a simple first-order decay reaction. Finally, the assembly reactions were taken from the reduced model (fig 1c).

The resulting model contained 251 species and 1,200 reactions, which made it the most complicated Lattice Microbes model to date. We were able to optimize the code to deal with the larger species and reaction counts through the development of just-in-time code generation for the reaction kernel. Our optimizations and the preliminary just-in-time code increased the performance of Lattice Microbes, such that a 60-minute cell cycle is completed in approximately seven days.

The stochastic, spatially resolved model of ribosome biogenesis allowed us to measure the distribution of assembly times for the 30S, which matched well with experiments. The model predicted spatial heterogeneity in the distribution of assembly intermediates, with the majority of the 5' and central domain proteins bound before the intermediate diffused far from the ribosomal operon (fig. 1b). Binding of the mRNA to the 30S tends to only occur near the originating operon or in the cytoplasm following the dissolution of the translating ribosome (fig. 1d).

WHY BLUE WATERS?

Blue Waters was crucial to the investigation of this problem due to the size of system. Lattice Microbes requires NVIDIA GPUs, and was easily modified to take full advantage of the available K20X accelerators. These simulations are I/O intensive, generating trajectory files approximately 10 GB in size per trajectory, of which many are needed for a proper statistical analysis of the results. Blue Waters provided all the technical resources necessary for us to study this system.

Further work on this project will focus on the integration of other subsystems to build a highly realistic bacterium *in silico*. Using chromosome capture [15,16] and Brownian dynamics [17,18] we will build DNA dynamics into the cell with accurate loci positioning. This will eliminate the need for biased transition and allow for the development of replication. We will build metabolism and general gene expression into our model [18] to couple ribosome biogenesis with DNA dynamics. With metabolism included, exchange reactions between the cell and the environment will allow for the investigation of the system by perturbing the external nutriment conditions.

SIMULATING INFLUENZA HEMAGGLUTININ MEMBRANE ASSEMBLIES

Allocation: GLCPC/0.28 Mnh
PI: Peter Kasson[1]

[1]University of Virginia

EXECUTIVE SUMMARY:

Viruses such as influenza enter cells via a process of membrane fusion between the viral envelope and cellular membranes. Fusion is accomplished by the concerted action of specialized viral coat proteins, but how these proteins work together to achieve fusion is not known. We are using molecular dynamics simulations to understand how influenza fusion proteins interact in lipid membranes, predict structural features of this interaction, and test our understanding via biophysical experiments. Achieving this understanding requires the extensive compute power and fast networking that Blue Waters provides. Success in this endeavor will yield testable hypotheses of how fusion proteins interact and contribute to viral entry mechanisms.

INTRODUCTION

An influenza virus enters cells via a process of membrane fusion, mediated by the viral hemagglutinin protein (which is also the target of most influenza vaccines). Mutagenesis experiments have established that portions of the hemagglutinin protein that insert into host membranes are surprisingly sensitive to mutations, suggesting that protein-membrane interactions are key to viral infectivity. On the other hand, the membrane "anchor" that holds hemagglutinin in the viral envelope is critical to organizing and positioning the protein for fusion but is surprisingly robust to point mutations. The dynamics of protein-membrane interactions and loose protein assemblies in membranes are slow from an atomistic point of view but difficult to probe via experiments. Here, we have used large-scale molecular dynamics simulations to understand the formation of membrane assemblies by the hemagglutinin transmembrane "anchor" and pilot a methodology for exploring such assemblies on a larger scale.

METHODS & RESULTS

We first performed lower-resolution simulations of protein-protein assembly in the membrane. We then used these to calculate a diverse sample of assembly geometries and conducted higher-resolution, more computationally intense simulations. Each of these, a large simulation in itself, was run in parallel (50 simulations per condition) on Blue Waters. Finally, to test the determinants of protein assembly, we predicted a set of mutations to the influenza proteins that abrogated the most stable contacts in the hemagglutinin assemblies. We simulated each of these mutants both under equilibrium conditions and applying external forces to measure the strength of each of the mutated assemblies.

Our results suggest that hemagglutinin transmembrane anchors can form very stable assemblies in the viral envelope. Most interestingly, these assemblies are extremely robust to mutation: If we mutate away the most stable "core" of the assembly, it remains more or less intact. These results yield a structural model for the experimental observation that hemagglutinin function is robust to mutations of this nature; the next question we plan to address in future work is precisely how these assemblies are related to hemagglutinin protein function and how we might be able to perturb this to disable viral infection.

WHY BLUE WATERS

This work relies on having both long simulations and many simulations to achieve rigorous sampling of a complex and slow physical process. Blue Waters combines these features, as it provides many GPU-enabled nodes, each of which can be run efficiently in a few-node parallel configuration. Future Tier-1 systems that may offer more capability for running many moderate-sized simulations that could in turn exchange data efficiently will greatly accelerate both our research and many similar problems in biomolecular simulation.

INSTRUMENTING HUMAN VARIANT CALLING WORKFLOW

Allocation: Illinois/0.60 Mnh
PI: L. S. Mainzer[1,2,3]
Collaborators: Arjun Athreya[3], Subho Banerjee[3], Ravishankar K. Iyer[3], Victor C. Jongeneel[1,2,3], Volodymyr Kindratenko[1,3], Zachary Stephens[3]

[1]National Center for Supercomputing Applications
[2]Institute for Genomic Biology, University of Illinois at Urbana–Champaign
[3]University of Illinois at Urbana–Champaign

EXECUTIVE SUMMARY:

If whole-genome sequencing and analysis become part of the standard of care within the next few years, human genetic variant calling will need to be performed on hundreds of individuals on any given day. For example, genotyping every baby born in Illinois would require analysis of ~500 genomes per day. At this scale, the standard workflow widely accepted in the research and medical community will use thousands of nodes at a time and have I/O bottlenecks that could affect performance even on a world class petascale system like Blue Waters. We identified and documented 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. Now we are designing and testing tools and methods to overcome these problems.

INTRODUCTION

Human variant calling is becoming the computational tool of choice to help diagnose intractable diseases and cancers. This bioinformatics tool searches for differences between a patient's genome and that of a reference, or average, of a human population. It is likely that in the not-too-distant future, every state and major metropolitan area would produce enough human genetic sequencing data to require their own HPC facility to analyze those data. What kinds of challenges would such a computational facility face, and what preparations would be necessary to ensure good performance with sustained throughput?

METHODS & RESULTS

We set up the standard GATK-based workflow [1] and tested a number of alternative tools at every step in the computation in an effort to shorten the computation wall time per genome. In addition, we benchmarked the CPU, RAM, and I/O system across the workflow using Perfsuite [2], Cray Profiler [3], OVIS [4], Valgrind [5], and some of the software we developed. This work generated close collaborations with Cray, the Blue Waters support team, and groups on campus, and identified several performance issues, most of which have to do with data I/O.

Variant calling is a big-data workflow when used in a sustained fashion on hundreds of samples per day. The required disk space is at the petascale on a daily basis. Even if the intermediary files are removed after the workflow is complete, they still need to be created, stored, and managed for the duration, and also in case one of the steps fails. The vast majority of these data are generated in the form of small files, which creates several concerns.

Creating, reading, and writing large numbers of small files can create an I/O bottleneck if the files are not placed uniformly across the file system. We found that the Lustre file system on Blue Waters places files relatively evenly across disks, according to our measurements. We continue to explore situations that could create unbalanced file placement and introduce tools into the workflow to prevent that from happening.

Handling large numbers of files can strain the metadata servers and result in uneven node performance, which affects both variant calling itself and other users on the system. We are investigating methods to do bookkeeping for deep collections of directories to lessen the load on the metadata servers and make data handling faster.

This workflow may benefit from storage pools. Usually, HPC jobs are MPI based and require fast inter-node communication, which means they benefit from being placed onto adjacent nodes. However, the variant calling workflow consists of jobs that run independent computational tasks on their own nodes. The I/O happens between the compute nodes and storage disk, not between compute nodes themselves. Thus, we are testing the use of a wide, sparse distribution of nodes assigned to the same job across the system.

The sheer size of the data footprint can create an infrastructure bottleneck. We are automating the output data archiving and ingestion of input streams, and offloading the output streams onto the target destinations to emulate the real-world case of genome data being streamed from sequencing facilities back to the medical centers.

Re-sorting aligned files by read position or by location along the genome is a common task that has to be done in a number of places along the variant calling workflow. The fastest algorithm, to our knowledge, is Novosort [5], which involves two phases. The first phase sorts as much data as possible in memory and then writes segments of sorted records to temporary disk files. The second phase merges the sorted fragments to produce the final sorted file. This algorithm is so efficient that it saturated the peak node injection bandwidth on Blue Waters (which is 9.6 GB/sec). This could create bottlenecks at scale by saturating the network routers on the system and interfering with functioning of other users. We are now measuring the extent of this potential problem and investigating a workaround by using data disk pools and/or a staggered version of the workflow to keep the sort–merge processes from overlapping.

WHY BLUE WATERS?

This project is timely because hundreds of sequencing facilities have already opened across the nation and the deluge of human genomics data is a reality today. The HPC facilities of tomorrow need to be prepared to handle the incoming load of genomic data. We hope that our project will inform the hardware and software designers about the requirements imposed by genomics on the next generation of computing facilities. The great advantage of using Blue Waters for this work is that it combines both a state-of-the-art data system and a large number of nodes to even make these experiments possible. The Blue Waters support staff have been instrumental in helping us figure out and eliminate issues with computational performance.

QUANTUM-CLASSICAL PATH INTEGRAL SIMULATION OF PROTON TRANSFER IN SOLUTION

Allocation: Illinois/60 Knh
PI: Nancy Makri[1]
Co-PI: Thomas C. Allen[1]

[1]University of Illinois at Urbana–Champaign

EXECUTIVE SUMMARY:

Quantum mechanical effects play an essential role in chemical and biological processes. Our group has recently developed a rigorous quantum–classical path integral (QCPI) methodology [1–3] in which one particle is treated by full quantum mechanics while the effects of the environment are captured via classical trajectories. The QCPI approach is free of *ad hoc* assumptions, allowing a faithful description of interference effects and leading to correct product distributions. Its implementation on Blue Waters provides a detailed picture of proton 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. The main difficulty in the development of quantum mechanical simulation algorithms stems from the non-local nature of quantum mechanics, which leads to exponential scaling of computational effort with the number of interacting particles.

For many processes of interest, quantum mechanical effects are vital in the treatment of a small number of degrees of freedom (e.g., those corresponding to a transferring proton), while the remaining particles (solvent molecules or biological medium) could be adequately described via Newtonian dynamics. However, the traditional Schrödinger formulation of quantum mechanics (which is based on delocalized wave functions) does not lend itself to a combination with Newtonian trajectories (which are local in phase space) unless severe approximations are introduced.

METHODS & RESULTS

The Makri group has been pursuing rigorous quantum–classical formulations based on Feynman's path integral formulation of quantum mechanics [4]. The major appeal of this approach stems from the local, trajectory-like nature of the Feynman paths, which leads naturally to combined quantum–classical treatments that are free of approximations. Recent work has described a quantum–classical path integral (QCPI) methodology, which incorporates these ideas as well as several advances in the understanding of decoherence (loss of coherence) processes [5]. QCPI treats a small subsystem by full quantum mechanics, while the effects of the environment are captured via standard molecular dynamics (MD) procedures. Since all quantum interference effects and their quenching by the solvent are accounted for at the most detailed (non-averaged) level, QCPI leads to correct branching ratios and product distributions, allowing simulations of important chemical and biological processes with unprecedented accuracy.

The present project involves the first implementation of QCPI to the simulation of the proton transfer reaction for the phenol-amine complex in methyl chloride [6]. This system is considered a paradigm for proton transfer and has served as a model for many computational investigations using a variety of approximations. Accurate QCPI calculations are being performed at the full atomistic level, as well as within the linear response approximation, in which the solvent is replaced by an effective bath of harmonic oscillators. Thus, the results will also help quantify the validity of linear response on a realistic proton transfer system [7].

LAMMPS, a well-known MD package, has been combined with the QCPI software and adapted to yield trajectories subject to forces obtained with the proton coordinates specified by the given quantum path. The integrated phase along each of these trajectories, averaged with respect to the thermal distribution of the solvent, yields the desired density matrix and any corrections to more economical classical treatments.

We have simulated the dynamics of the transferring particle and obtained the time evolution of the state populations. Extensive tests are currently being performed to verify convergence. The transfer process was dominated by solvent effects, and some high-frequency vibrations of the solvent molecules were strongly coupled to the transferring particle. In addition, strongly coupled sluggish solvent modes affected the proton transfer dynamics via a quantum mechanical decoherence mechanism. The QCPI calculations quantified the significance of these contributions.

Upon completion of the present phase of the work, the QCPI calculations will provide quantitative results for the kinetics of the proton transfer process, along with a detailed understanding of the underlying mechanism(s). This will include the time scale of correlations and decoherence, the distinct roles of fast and sluggish solvent motions and associated quantum effects, as well as the importance of nonlinear solvent effects on the dynamics.

WHY BLUE WATERS?

Implementation of the QCPI methodology requires integration of a large number of classical trajectories from each initial condition sampled from the solvent density. Each of these trajectories interacts with a Feynman path of the quantum subsystem, augmenting the dynamics of the quantum particle via a phase. By exploiting the very mechanism of decoherence, we can circumvent the exponential proliferation of the number of trajectories with propagation time. Still, QCPI calculations on these systems are extremely demanding, and these calculations are not feasible on conventional computational platforms.

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.

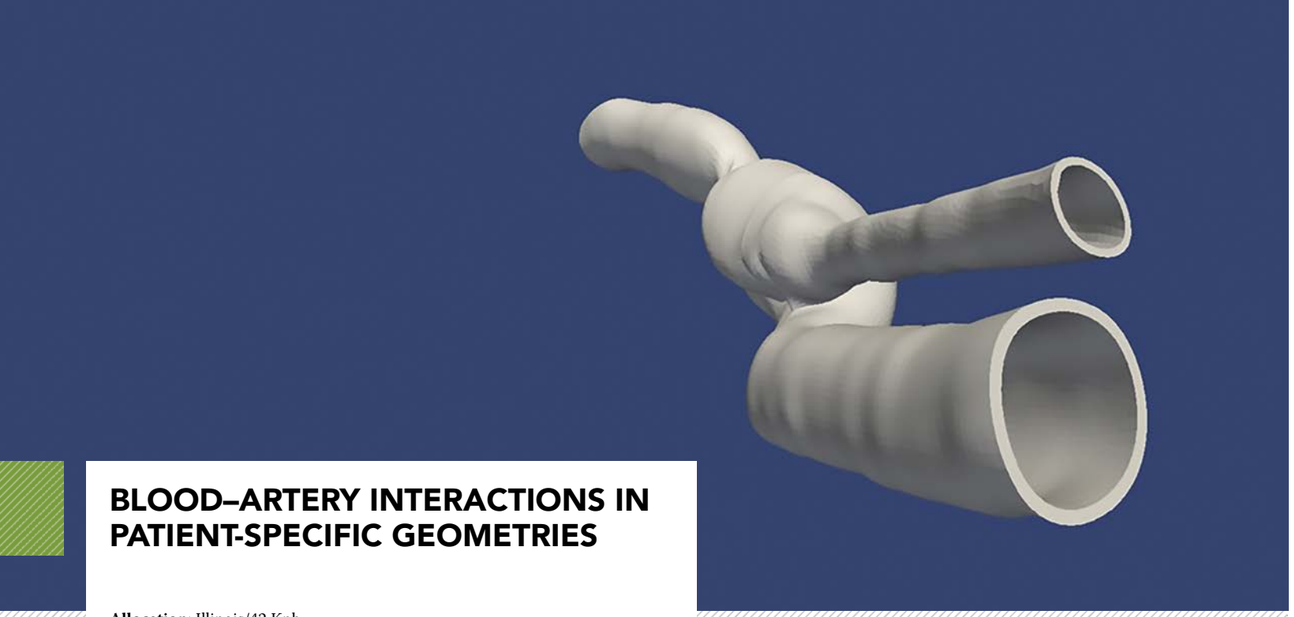

BLOOD–ARTERY INTERACTIONS IN PATIENT-SPECIFIC GEOMETRIES

Allocation: Illinois/42 Knh
PI: Arif Masud[1]
Collaborators: JaeHyuk Kwack[1], Soonpil Kang[1]

[1]University of Illinois at Urbana–Champaign

EXECUTIVE SUMMARY:

We developed highly stable and accurate blood–artery interaction models that are based on our earlier works on fluid–structure interaction (FSI) methods. In this work, 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.

Since both blood and artery walls are modeled via our newly developed multiscale finite-element methodology, the new FSI method provides a unified platform for developing advanced coupled-solution algorithms with enhanced stability properties that play a crucial role in carrying out highly transient flow simulations. Extracted flow physics highlights the effects of stiffening of the arterial walls on the progression of hypertension in patient-specific models of the carotid artery.

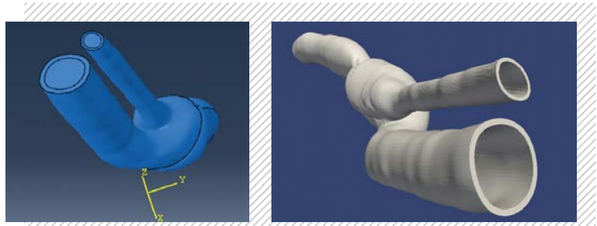

FIGURE 1: Geometry of patient-specific artery: (a) arterial subdomain, (b) full model.

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. Since blood–artery interaction models are mathematically involved and computationally expensive, Blue Waters resources were employed to further explore the mathematical attributes of our new coupling scheme. This study also helped investigate 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. It is a problem of great medical relevance as uncontrolled high blood pressure increases the risk of serious health problems, including heart attack and stroke.

METHODS & RESULTS

Based on our previous Blue Waters allocations we were able to extend and verify our non-Newtonian models for blood that account for its 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. This scale split leads to two coupled nonlinear systems, 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 it led to the class of hierarchical multiscale methods for non-Newtonian fluids with enhanced stabilization properties.

During the last year, we developed the arterial model as shown in fig. 1a, which is a patient-specific model of a carotid artery that suffers from stenosis and aneurysm. Fig. 1b shows the coupled model where the blood and artery subdomains can be seen. A coupled response of the system is shown in fig. 2, where velocity magnitude is projected on various cross sections along the length of the artery. A set of parametric studies with different relaxation times were carried out for the arterial walls and pressure parameters computed, which provided insights into the effects of mechanical material parameters on an increase in blood pressure.

From a computational and algorithmic perspective the newly developed coupled hierarchical multiscale 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 fewer 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 is of tremendous benefit in massively parallel computing as it reduces communication costs across the partitioned subdomains.

WHY BLUE WATERS?

Blue Waters' architecture, with large local memory, is ideal for our methods as we are able to exploit the local memory on the processing nodes to make the macro elements "smart," thereby reducing the size of the global problem and minimizing data communication. Ready access to Blue Waters for large-scale computing and to project staff to discuss various technical points arising in code development were key in making these major strides. Using the next Track-1 system, we plan to extend and embed our method in a probabilistic framework for blood flow simulation in patient-specific arterial geometries, with the objective of optimization of ventricular assist devices (VADs) for patient-specific needs.

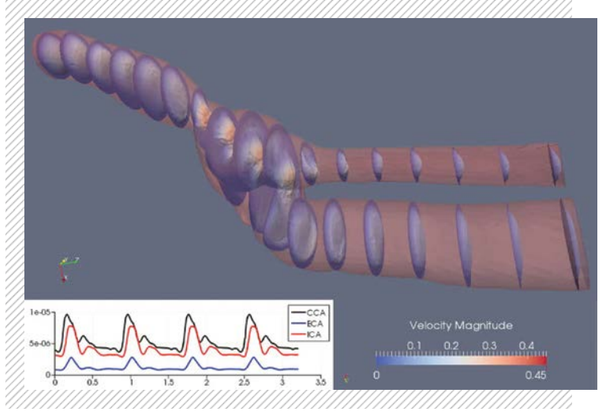

FIGURE 2A: Stiff and moderate viscoelastic carotid artery wall. Velocity profiles at various cross sections.

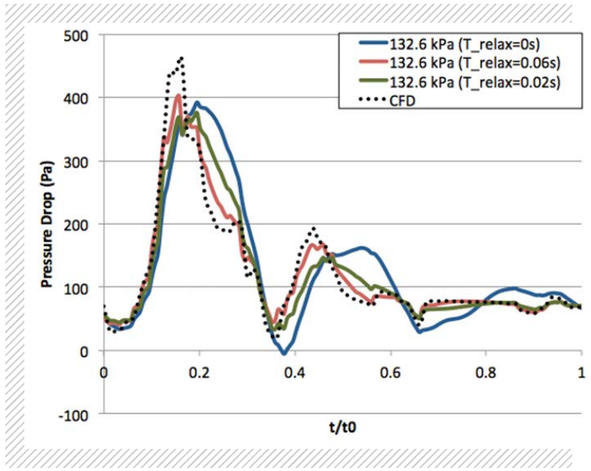

FIGURE 2B: Pressure drops between inlet and outlets.

UNDERSTANDING P53, A PROTEIN "SHAPE SHIFTER" RESPONSIBLE FOR MANY CANCER TUMORS

Allocation: NSF/2.45 Mnh
PI: Vijay S. Pande[1]

[1]Stanford University

EXECUTIVE SUMMARY:

We used Blue Water's powerful GPUs, CPUs, and network to simulate p53, a key protein involved in cancer tumors, to better understand how it interacts with other key proteins in the cell. These simulations have taken p53 studies to unprecedented timescales, hundreds of microseconds to milliseconds, and have yielded new insight into p53 assembly. Specifically, we find that p53 plasticity with multiple proteins occurs via numerous pathways, suggesting specific challenges and opportunities for small molecule therapeutic approaches.

INTRODUCTION

We used Blue Waters for biomolecular simulations that study the folding and binding of the intrinsically disordered protein p53, which is a cancer tumor suppressor. Half of all known cancer tumors involve a mutation in p53, clearly making it a key protein in cancer. Beyond disease, p53 plays a key role in apoptosis, i.e. programmed cell death, the mechanism by which cells prevent disease by "shutting down" cells that have become old and problematic. Cancer cells are, in a sense, cells that didn't get this message to shut down and instead continue to grow in a damaged state. Activation of p53 prevents tumorgenesis and maintains normal cell growth. In this work, we studied the binding of a key portion of p53 to two proteins, s100β and a member of the sirtuin family, Sir2Tm, that antagonize the p53 activation pathway.

METHODS & RESULTS

We have used cutting-edge Markov State Model sampling methods combined with the vast power of Blue Waters to access unprecedented timescales, modeling phenomena on the millisecond timescale in all-atom detail.

This study gives biophysical insight into mechanisms for the assembly and binding of disordered proteins as p53 binds in two distinct structural conformations with the different binding partners. In the most probable pathways connecting the unfolded, unbound state and the folded, bound state, we find a fly-casting mechanism (first proposed by Peter G. Wolynes and collaborators) for both proteins, illustrated in Figures 1 and 2.

An initial capture process involves broad, mostly polar, protein contacts that allows flexibility for p53 rearrangement toward the bound state. For s100β, we see that p53 can rearrange into a stable misfolded off-pathway intermediate that may contribute the bottleneck rate for finding the native state. In both proteins, disordered regions on the p53 binding partner cooperatively shuttle the p53 peptide to the native, bound state. Understanding p53 interactions with these proteins can inform development of potential inhibitors for cancer therapy, to restore p53 activation and normal cell growth cycles in tumors. We suggest that the green interfaces highlighted in Figures 1 and 2 could be used to target p53 for small molecule or peptide inhibitors to prevent p53 capture.

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.

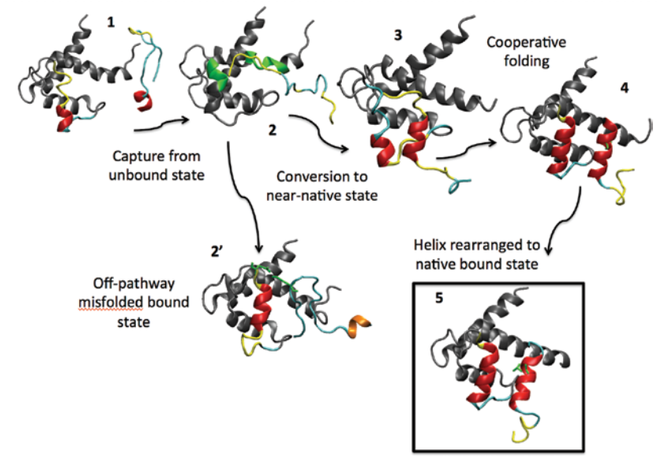

FIGURE 1: p53-s100beta predominant folding pathway. The proposed fly-casting-like mechanism involves initial capture by broad polar interfaces (green in state 2), followed by cooperative folding of p53 and Helix 3 of s100beta. p53 rearranges to the native folded, bound state, inserting L383 (green residue 5) into the hydrophobic pocket of s100beta. An off-pathway misfolded state is also identified, captured by a non-native salt bridge (green in state 2').

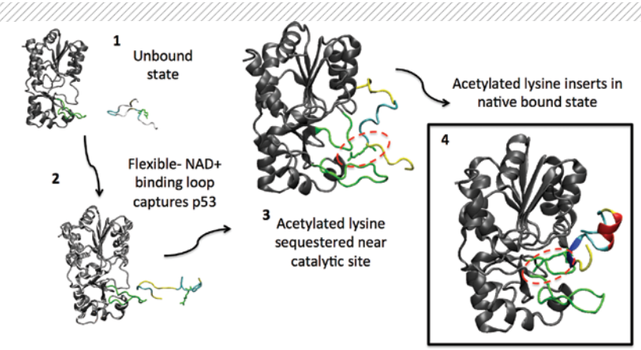

FIGURE 2: p53-sirtuin predominant folding pathway. The proposed fly-casting-like mechanism involves initial capture by flexible loops that also bind the cofactor NAD+ and involves broad polar interfaces (green in states 2 and 3). The acetylated lysine K382 (shown in red dashed circle) is shuttled via these flexible loops into the hydrophobic catalytic site, and its flanking residues form beta strands that staple two domains of sirtuin.

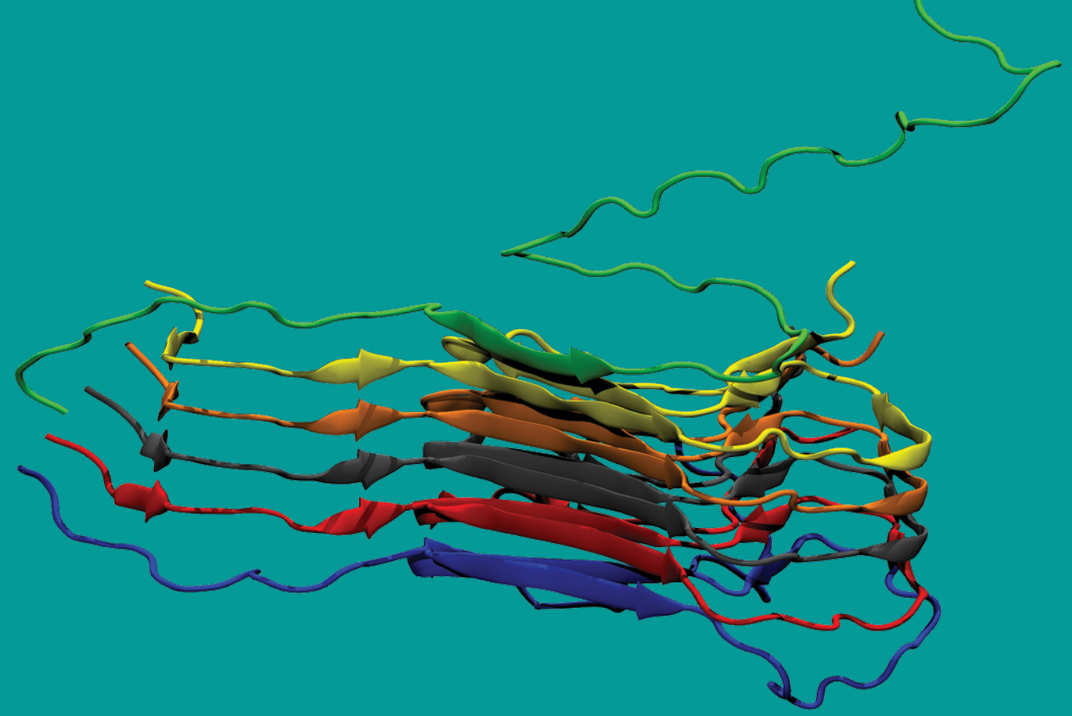

PROTEIN STRUCTURE DETERMINATION AND REFINEMENT FROM RAW NMR DATA

FIGURE 1: NAMD-GO simulation: folding and docking of monomer to alpha-synuclein fibril, 2.5 ns simulation on NAMD platform with GROMACS force field.

Allocation: Illinois/50 Knh
PI: Chad M. Rienstra[1]

[1]University of Illinois at Urbana–Champaign

EXECUTIVE SUMMARY:

Nuclear magnetic resonance (NMR) is a very powerful technique for investigating protein structure, dynamics, and interactions, and in the present project we are developing new computational methods for this purpose. One protein of special interest is alpha-synuclein, which forms protein aggregates as part of the neurodegenerative cascade of Parkinson's disease (PD) [1]. Thus, the understanding of this pathway may assist in the diagnosis and treatment of PD. We investigate the free-energy profile of alpha-synuclein aggregation using bias-exchange metadynamics, with collective variables that mimic the beta sheet structure, hydrogen bonding, and hydrophobic contacts present in the fibril, using GROMACS and PLUMED.

INTRODUCTION

Alpha-synuclein helps regulate the supply of synaptic vessels in the presynaptic terminals in the neuron, as part of the regulation of dopamine release. Misfolded alpha-synuclein can aggregate in the brain to form insoluble fibrils under pathological conditions, giving rise to disorders like Parkinson's disease [1]. The mechanism of aggregation is unknown and hence it is not clear if the protofibrils or the fibrils themselves are toxic to our brain. Using high resolution solid-state nuclear magnetic resonance (NMR), our lab recently determined the structure of the final form of the fibril. Using this structure, we aim to elucidate the pathway of aggregation by constructing a free-energy profile, including the intermediates formed during the process. This will help us determine the toxic species and eventually help identify a cure to Parkinson's disease.

METHODS & RESULTS

A trajectory on the pathway of a disordered monomer getting linked to a shorter fibril to form the final structured fibril was considered to be the model for generating free energy profile. The system was biased with GO-Potential [2] based on the final fibril structure and the NAMD-GO [3] technique was used with the GROMACS [4] force field. The nucleation was carried out in 2.5 ns.

A 20 ns equilibration simulation was performed in NAMD [6] using the CHARMM36 [7] force field on the experimental structure of the fibril to verify its stability. Next, we successfully simulated a 2.5 ns trajectory of an unfolded disordered monomer folding and docking to a shorter fibril to form a larger fibril (fig. 1). Three sharp changes in the RMSD value of the trajectory indicated that a topological barrier was present at that region in the pathway of the protein folding and docking.

As the NAMD-GO path proved to be too short for analysis of the trajectory, we simulated the reverse process of fibril unfolding with the biasing potential of the metadynamics [5] method in an explicit solvent environment. This simulation was run for 60 ns with appropriate collective variables using the NAMD2 [6] molecular dynamics code and the CHARMM36 [7] force field.

Several metadynamics [5] simulations were carried out on the initial system of a large fibril (haxamer), where a pentamer portion was kept fixed and the last monomer was made to unfold under the metadynamics [5] potential. The collective variables used were: (a) hBond [10] (calculates the hydrogen bonding present between the fibrils), and (b) coordNum [10] (calculates the amount of hydrophobic contacts present in between fibrils). Thus far, these collective variables have not been able to mimic the unfolding of the protein.

These results showed that better collective variables and technique for performing the free-energy calculations are needed. We decided to use the Bias-Exchange Metadynamics [8] technique that worked best on the GROMACS molecular dynamics code with PLUMED [9] software for analysis. The work of implementing the software and setting the simulation is in progress. Future work will include collective variables that leverage the experimental NMR shift values (via CAMSHIFT [9]).

WHY BLUE WATERS?

Blue Waters was essential to our project since we needed a supercomputing facility to run our long simulations on fairly large systems (~10,000 atoms) in a convenient amount of time. Also, the installation of required software is convenient on Blue Waters because all the prerequisites like compilers are available for easy use.

MOLECULAR DETAILS OF THE LUMINAL EXIT OF CALCIUM IONS IN THE CALCIUM PUMP OF THE SARCO/ENDOPLASMIC RETICULUM

Allocation: GLCPC/0.60 Mnh
PI: Benoît Roux[1]
Co-PI: Avisek Das[1]

[1]University of Chicago

FIGURE 1: (a) Crystal structures of calcium-bound occluded state (left, PDB ID: 3BA6) and calcium-free state with an opening toward the luminal side (right, 22BE). (b) Distances between the coordinating residues and the first calcium ion along the pathway. (c) Distances between the coordinating residues and the second calcium ion along the pathway. (d) Number of water molecules around the calcium ions along the pathway.

EXECUTIVE SUMMARY:

Sarco/endoplasmic reticulum Ca^{2+}-ATPase (SERCA) is an integral membrane protein that uses ATP hydrolysis as a source of free energy to pump two calcium ions per ATP molecule from the calcium-poor cytoplasm of the muscle cell to the calcium-rich lumen of the sarcoplasmic reticulum, thereby maintaining a ten-thousand-fold concentration gradient. Our goal is to understand the important dynamical motions of the pump via high-resolution crystal structures of several functionally relevant states along the pumping cycle. We employed an all-atom molecular-dynamics-based rare event method called string method with swarms of trajectories to study the large-scale conformational transitions of the pump. The method involves running a large number of short trajectories that communicate with each other at a regular interval. Our recent simulations on Blue Waters revealed unprecedented molecular details of an important step of the pumping cycle.

INTRODUCTION

Membrane proteins form an important class of biomolecules that are associated with the membrane dividing the inside of a cell (or a cellular compartment) and its environment. Our project is aimed at understanding the function of an integral membrane protein called sarco/endoplasmic reticulum Ca^{2+}-ATPase (SERCA) [1–3] that uses ATP hydrolysis as a source of free energy to pump two calcium ions per ATP molecule from calcium-poor cytoplasm of the muscle cell to the calcium-rich lumen of the sarcoplasmic reticulum. This process is important for relaxation of skeletal muscle, which is regulated by calcium ions. Over the past few years a number of structural studies [1–6] have provided atomic-resolution models for several important states along the pumping cycle. Two major outstanding issues are the pathways of ions from either side of the membrane to the transmembrane binding sites and a detailed description of the conformational changes that will elucidate how various parts of the protein communicate over fairly large distances in order to achieve coupled ATP hydrolysis and calcium transport.

We believe our project will have a significant impact in the field of membrane protein biophysics for several reasons. First, SERCA plays an important role in the relaxation of skeletal muscle and a close analogue, which is present in the cardiac muscle, is a therapeutic target. Therefore, understanding the molecular mechanism of the active transport process is of fundamental and biomedical interest. Second, SERCA is a member of P-type ATPase family of membrane proteins that have very similar topologies and reaction cycles. A very important member of this family is Na^{+}-K^{+}-ATPase which shares high sequence similarity with SERCA. It is conceivable that the detailed knowledge acquired by studying SERCA can be applied to understand the mechanism of Na^{+}-K^{+}-ATPase which has very few high resolution structures of relevant functional states. Third, this study will serve as a test case for evaluating success and limitations of advanced methods for studying complex conformational transitions in a biologically relevant system and will greatly benefit the entire bimolecular simulation community.

METHODS & RESULTS

Over the past three years we have been involved in simulating large-scale conformational transitions of SERCA between experimentally resolved states. The size of the system (~290,000 atoms) and the time scale involved in the large scale motions prohibited the use of brute-force unbiased molecular dynamics (MD) [7] simulations to obtain statistically meaningful information. To circumvent this challenge, the string method with swarms-of-trajectories was used to discover the optimal minimum free-energy path between the two end states [8,9]. The path, which was represented as a chain of states or images in the space of relevant collective variables, was optimized by iterating two steps: moving each image along the drift calculated from an ensemble (swarm) of short unbiased MD trajectories initiated from the image and a re-parametrization procedure that keeps all the images equidistant.

We determined a transition pathway between the calcium-bound occluded state (PDB ID: 3BA6) and a calcium-free state with an opening toward the luminal side (fig. 1a). The transition in question is responsible for release of calcium ions into the luminal solution from the transmembrane binding sites. The string was represented by 32 images and for each image we used 32 trajectories for the estimation of the drifts (a total of 1,024 copies of the system). Each iteration of the method involved 20 picoseconds of simulation and 360 iterations were performed for production calculation. Simulations were carried out using a modified version of NAMD 2.9 [10]. The last twenty iterations were used for analysis. We have monitored the coordination environments of the calcium ions along the pathway (fig. 1b-c). The large-scale movements of the cytoplasmic domains induced important motions in several transmembrane helices that disrupted the binding sites, simultaneously forming an opening toward the luminal side. As a result of that, the water from the luminal side entered the interior of the protein and interacted with the calcium ions (fig. 1d, fig. 2). The final release of the ions presumably involved protonation of several binding site residues. The water access channel from the luminal side shed some light on the putative luminal exit channel of the calcium ions (fig. 2). These results complement the findings from structural studies and provide a comprehensive picture of a crucial step in the pumping cycle.

WHY BLUE WATERS?

The string method with swarms-of-trajectories is essential for simulating large-scale conformational changes in SERCA. In order to implement this method, one needs to simulate many copies of the system that communicate with each other at a regular interval. For meaningful results, more than a thousand copies of the system and hundreds of iterations are required. To perform this calculation on SERCA (~290,000 atoms), a single job requires more than 6,000 nodes, which is more than the total node count of an entire machine for many small- to medium-sized supercomputers. Therefore, the massively parallel architecture of Blue Waters played a crucial role in the success of our project.

FIGURE 1

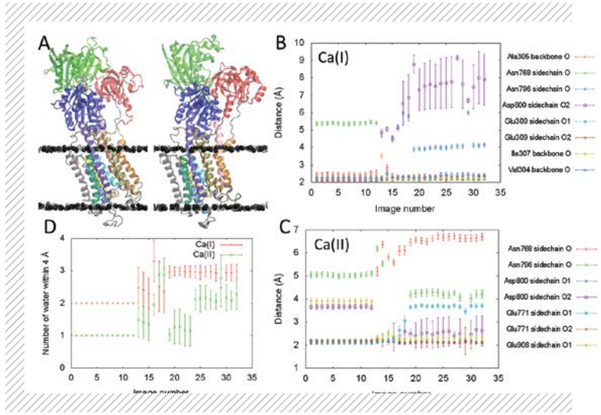

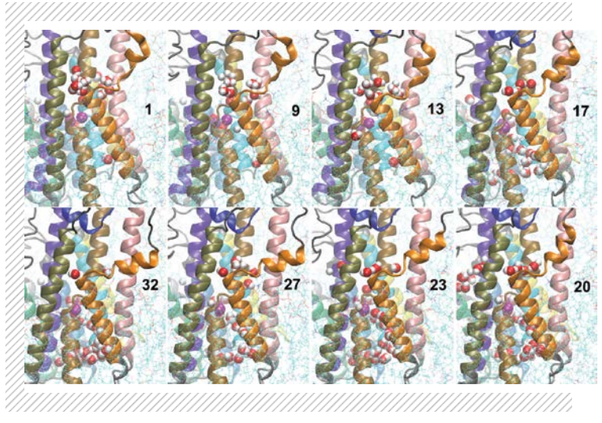

FIGURE 2: Water access pathway from the luminal side. This indicates the path ions take from the transmembrane region to the luminal side. The integers on the panels indicate the image indices of representative snapshots of the final pathway. The first and the last indices correspond to the occluded calcium-bound state and the non-occluded state with an opening toward the luminal side. The transmembrane helices are color coded as follows: M1 (orange), M2 (pink), M3 (tan), M4 (ochre), M5 (violet), M6 (cyan), M7 (dark green), M8 (sky blue), M9 (yellow), M10 (silver). Calcium ions are shown as purple spheres.

THE COMPUTATIONAL MICROSCOPE

Allocation: NSF/18 Mnh, NSF RAPID/0.55 Mnh, Blue Waters Professor/0.24 Mnh
PI: Klaus Schulten[1]
Co-PIs: James C. Phillips[1] and John E. Stone[1]
Collaborators: Rebecca C. Craven[2], Costas Maranas[3], Eva Nogales[5,6], Peijun Zhang[4]

[1]University of Illinois at Urbana-Champaign
[2]Penn State University College of Medicine
[3]Penn State University
[4]University of Pittsburgh
[5]University of California, Berkeley
[6]Howard Hughes Medical Institute

EXECUTIVE SUMMARY:

Living organisms arise from basic, inanimate matter, but do they blindly following the laws of physics? Computational biology in the petascale era offers a unique opportunity to connect atomic-level descriptions of biological systems (representing inanimate matter) with cell-scale architecture and behavior (representing animate matter). Molecular dynamics (MD) simulations can be used as the "computational microscope" to establish amazing cellular structures at the atomic level, as well as to employ those structures to elucidate the dynamics and underlying physical and chemical mechanisms of cellular processes. Such a computational approach has enjoyed a rapid acknowledgment of its complementary role to experimental observation, the latter often lacking the spatial and temporal resolution needed to truly grasp cellular processes. Here we discuss successful research on several large-scale biological systems.

FIGURE A: The first computational simulation of the immature lattice of the Rous sarcoma virus. This model provides a template for studying the maturation process of retroviruses as well as evaluating opportunities for drug targets.

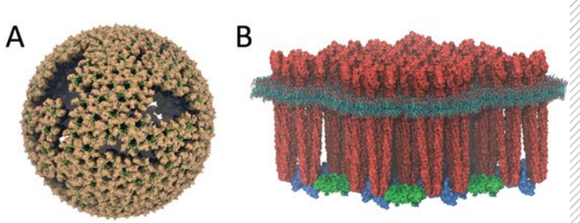

INTRODUCTION

Recent developments on hybrid experimental methods based on the revolutionary advance of electron microscopy have led to previously unimaginable information on 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. 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 are developed, enabling the biologically functional timescale in living cells to 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 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

Since solving the first atomic-level structure of the mature HIV capsid [1], we can now describe the action of host cell factors in capsid stabilization and assembly (cyclophilin A) as well as characterize the effect of a cell factor essential for disrupting the capsid and assisting in nuclear import (TRIM-family proteins). Such studies may help scientists to better understand how the HIV capsid infects the host cells and could lead to new HIV therapies.

While the mature HIV capsid is made only of capsid proteins, the immature lattice requires Gag proteins, a polyprotein essential for its assembly. To investigate the immature virus in atomistic detail, we produced the first all-atom model of the immature lattice of a retrovirus (Figure A) that is closely related to HIV, namely Rous sarcoma virus (RSV) [2]. The atomic structure of the immature RSV lattice represents a milestone for computational modeling, and the structure provides a template to study the maturation process of retroviruses as well as to evaluate opportunities for drug targets.

The recent Ebola outbreak in West Africa prompted us to develop a diagnostic tool that would reliably detect the virus in pre-symptomatic patients. To detect Ebola, we have computationally designed and optimized the first set of antibody-like prototypes. The experimental validation of these prototypes is projected to be completed by early 2016.

Bacteria use large, highly ordered clusters of sensory proteins known as chemosensory arrays (Figure B) to detect and respond to chemicals in their environment. We have integrated multi-scale structural data from experimental sources to computationally construct the first atomic model of the chemosensory array's molecular architecture. In addition, we identified a novel conformational change in a key signaling protein that is linked at the cellular level to the chemotaxis function (that is, movement in response to chemical stimulus). 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 (Figure C) in different nucleotide binding states (crucial for the switch between phases of assembly and disassembly). Our simulations suggested important structural dynamics events toward the microtubules assemble and stability. Such simulations pave the way to understand the atomic details behind the assembly and disassembly phases of the microtubules as well as the effect of anticancer drugs on microtubule dynamic instability.

Molecular motors, which travel on microtubules, are key drug targets for anti-cancer therapy. Combining novel sampling techniques with the computing power of Blue Waters, we were able to study a long timescale mechanical process (millisecond) driven by chemical energy of a ring-shaped molecular motor [3]. Such a study provides an effective approach to identify possible drug targets on protein structure.

WHY BLUE WATERS

Without Blue Waters and other petascale computing resources, projects involving large molecular systems like HIV, RSV, the chemosensory array, and microtubules would not be possible. These molecular systems are composed by tens of millions of atoms and must be simulated for long periods of time (microseconds). On the other hand, medium systems like the motor protein project (hundreds of thousands of atoms) could be simulated for a timescale of milliseconds only with the computational power delivered by Blue Waters. 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.

FIGURE B: Bacteria chemosensory array model at atomic resolution.

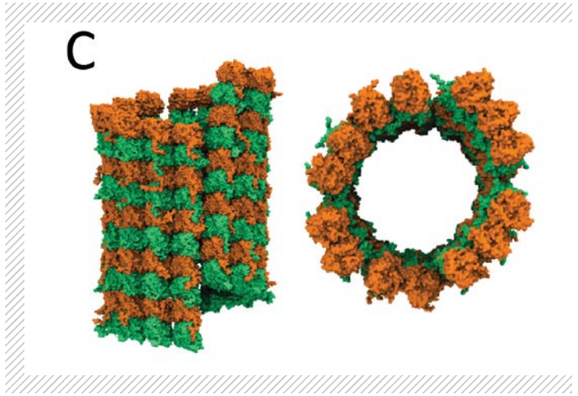

FIGURE C: Side and top view of the atomic model of a microtubule, a multi-protein structure important for cell structure, intracellular transport, and cell division.

PUBLICATIONS

Hardy, J. H., et al. Multilevel summation method for electrostatic force evaluation, *J. Chem. Theory Comput.*, 11:766-779, 2015.

Perilla, J. R., et al., Molecular dynamics simulations of large macromolecular complexes, *Current Opinion in Structural Biology*, 31:64-74, 2015

CHARACTERIZING STRUCTURAL TRANSITIONS OF MEMBRANE TRANSPORT PROTEINS AT ATOMIC DETAIL

Allocation: Illinois/0.50 Mnh
PI: Emad Tajkhorshid[1,2]

[1]Beckman Institute for Advanced Science and Technology
[2]University of Illinois at Urbana–Champaign

EXECUTIVE SUMMARY:

Membrane transporters actively and selectively pump molecules in or out of the cell. The central role of these proteins in diverse physiological processes such as metabolism and neurotransmission makes them key drug targets. Their transport mechanisms often involve large-scale conformational changes coupled to the translocation of transported species such as ions, nutrients, and drugs. However, mechanistic details of these transport processes at an atomic level remain elusive due to the limitations of both experimental and computational techniques.

We have developed a novel computational approach to reconstruct such transport cycles using ensemble-based molecular dynamics simulations within an iterative framework involving high-dimensional path-finding algorithms and free-energy calculations. Employing Blue Waters resources and our novel non-equilibrium approach, the transport cycles of several membrane transporters have been reconstructed at an atomic level and characterized thermodynamically, providing an unprecedented level of detail about the transport mechanism.

INTRODUCTION

All living organisms rely on continuous exchange of diverse molecular species across cellular membranes for their function and survival. Membrane transporters are specialized molecular devices that provide the machinery for selective and efficient transport of materials across the membrane. They actively pump their substrates across the membrane by taking advantage of different forms of cellular energy. The biological and biomedical relevance of mechanistic studies on membrane transporter proteins cannot be overstated, given their central role in a myriad of key cellular processes and their involvement in the action of a vast number of pharmaceuticals. Large-scale conformational changes are central to the mechanism of membrane transporters. During the alternating-access mode of function [1], which is needed for uphill transport of the substrate against its electrochemical gradient, the protein switches substrate accessibility from one side of the membrane to the other. The process largely relies on large-scale structural transitions between two major states: the outward-facing (OF) and inward-facing (IF) states. Given the technical challenges involved in experimental characterization of these structural phenomena, simulation studies currently provide the only method to achieve the spatial and temporal resolutions required for a complete description of the transport cycle in membrane transporters, which by any measure constitute the most important aspect of the structural biology of these molecular machines.

METHODS & RESULTS

Large-scale structural transitions of membrane transporters cannot currently be characterized at an atomic level using conventional simulation technologies due to the long time scales involved (on the order of tens of milliseconds to seconds) [2]. Recognizing this issue, we recently developed a knowledge-based computational approach for describing large-scale conformational transitions using a combination of several distinct enhanced sampling techniques [3–5]. In the first stage, we used non-equilibrium, driven simulations based on system-specific reaction coordinates whose usefulness in inducing the transition of interest was examined using knowledge-based, qualitative assessments along with non-equilibrium work measurements. In the second stage, we use the string method with swarms-of-trajectories (SMwST) [5–7] in a high-dimensional collective variable space to further relax the optimized non-equilibrium trajectory obtained from the first stage and found an approximate minimum free-energy pathway. Finally, we used the relaxed trajectory to initiate bias-exchange umbrella sampling (BEUS) [3–5] simulations to estimate the free energies along the transition pathway.

We have studied the structural transition between the IF and OF states in several proteins from different classes of transporters including the ATP-binding cassette (ABC) transporter P-glycoprotein (Pgp) and its bacterial homolog MsbA as well as major facilitator superfamily (MFS) transporters GlpT, GLUT1, and XylE using all-atom molecular dynamics (MD) simulations in the presence of an explicit membrane and solvent. In particular, we performed extensive non-equilibrium simulations to sample a large number of mechanistically distinct pathways for the conformational transition of MsbA [4].

The most relevant transition pathway identified using our approach was further investigated using free-energy calculations. The transition involved several distinct stages reflecting the complex nature of the structural changes associated with the function of the protein. The opening of the cytoplasmic gate during the outward- to inward-facing transition of apo (nucleotide-free) MsbA was found to be disfavored when the periplasmic gate was open and facilitated by a twisting motion of the nucleotide-binding domains that involved a dramatic change in their relative orientation. These simulations have been extended to study Pgp, a mammalian homolog of MsbA that is a multidrug resistance transporter (fig. 1).

In addition to ABC transporters, which use ATP hydrolysis as a source of energy, we have extensively studied several MFS transporters using our novel approach. Particularly, we have been able to provide a quantitative, atomic-level description of the functional thermodynamic cycle for the GlpT transporter by reconstructing the free-energy landscape governing the IF↔OF transition along a cyclic transition pathway involving both apo and substrate-bound states [5]. Our results provide a fully atomic description of the complete transport process, offering a structural model for the alternating-access mechanism and substantiating the close coupling between global structural transitions and local chemical events such as binding and unbinding of the substrate.

WHY BLUE WATERS?

Conventional MD simulations of large biomolecular systems suffer from poor conformational sampling, preventing one from achieving an accurate description of conformational ensembles and free-energy landscapes. The time scale limitation remains a great challenge, and the typical time scale of an atomistic MD simulation of biomolecular systems is much smaller than those required to describe most biologically relevant molecular phenomena. Some of the most powerful enhanced sampling techniques developed to address these issues rely on multiple-copy algorithms (MCAs) [7] that couple the dynamical evolution of a large number of copies of a system to explore the phase space and enhance the sampling. In our simulations, we employ BEUS and SMwST schemes, both of which are MCAs. Given that every copy of the simulation would require thousands of cores for its simulation, simulating a large number (typically hundreds) of interacting replicas simultaneously can only be accomplished on massive computing resources such as Blue Waters.

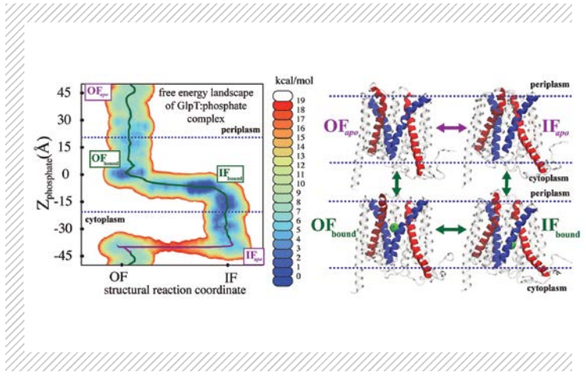

FIGURE 1: Transport cycle of MFS transporter GlpT characterized at an atomic level using enhanced sampling techniques and molecular dynamics simulations. The crystal structure of GlpT (IF_{apo} state) was used to reconstruct its entire transport cycle *(right)*. The transition pathway was then characterized energetically using 150 interacting replicas of GlpT system (each containing ~150,000 atoms) within the BEUS scheme *(left)*.

PUBLICATIONS

Li, J., P.-C. Wen, M. Moradi, and E. Tajkhorshid, Computational characterization of structural dynamics underlying function in active membrane transporters. *Curr. Opin. Struct. Biol.*, 31 (2015), 96–105, doi:10.1016/j.sbi.2015.04.001.

Moradi, M., G. Enkavi, and E. Tajkhorshid, Atomic-level characterization of transport cycle thermodynamics in a membrane transporter. *Nat. Commun.*, (accepted).

Moradi, M., and E. Tajkhorshid, Computational recipe for efficient description of large-scale conformational changes in biomolecular systems. *J. Chem.Theory Comput.*, 10 (2014), pp. 2866–2880, doi:10.1021/ct5002285.

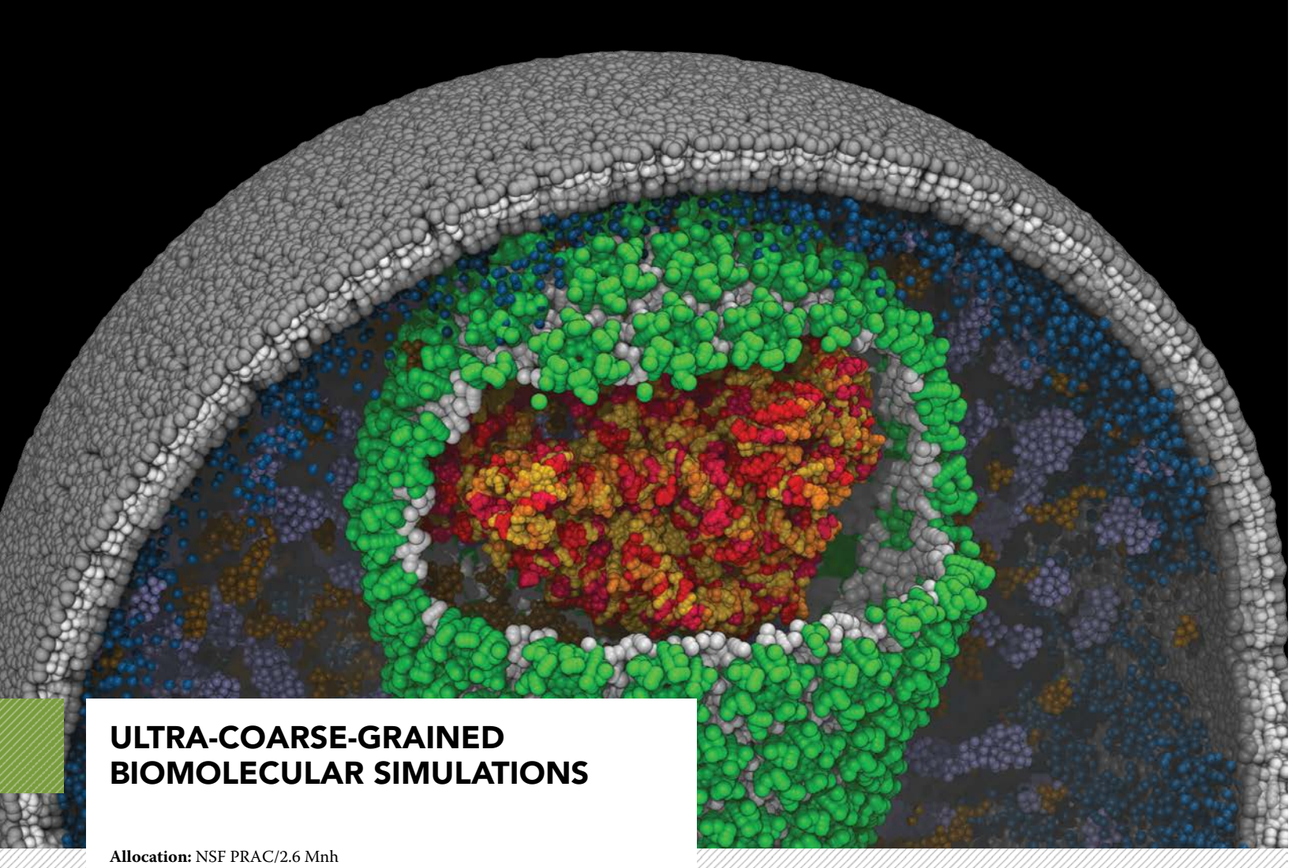

ULTRA-COARSE-GRAINED BIOMOLECULAR SIMULATIONS

Allocation: NSF PRAC/2.6 Mnh
PI: Gregory A. Voth[1]
Collaborator: John Grime[1]

[1]The University of Chicago

EXECUTIVE SUMMARY:

The spread of HIV-1 infection requires the release of a viral particle, or virion, from an infected cell. A newly released virion is non-infectious, and so a sophisticated process of "maturation" occurs in order to produce a mature and infectious virion. One critical aspect of maturation is the self-assembly of many copies of the viral capsid protein to generate a cone-shaped capsid structure that surrounds the viral RNA; where a suitable capsid is not formed, the virion remains non-infectious. In order to study the critical early stages of capsid formation, we use a simplified model of the capsid protein and examined the capsid self-assembly behaviors in a range of environments relevant to the maturing HIV-1 virion. The results of these computer simulations revealed the critical importance of protein concentration, molecular crowding, and dynamic changes in capsid protein structure in the nucleation and growth of the mature viral capsid.

INTRODUCTION

The lifecycle of human immunodeficiency virus type 1 (HIV-1) involves a complicated series of processes to successfully spread the viral infection. After the initial aggregation of many copies of the Gag polypeptide at the interior membrane of an infected cell, a viral particle (or virion) is released from the cell to potentially infect additional cells. However, the initial virion is non-infectious and must undergo a process of "maturation" before an infectious form is produced. One key process in virion maturation is the self-assembly of many copies of a single-capsid protein molecule to generate a cone-shaped viral capsid to enclose the viral RNA. The molecular nature of capsid nucleation and growth phenomena renders the experimental study of the critical early stages of capsid assembly extremely difficult with conventional experimental techniques.

Given the known natural variability across virions, it is important to consider how the spontaneous nucleation and growth of capsid protein structures is controlled in order to generate the single cone-shaped capsid expected in a prototypical mature virion. We therefore used simplified computational models of the capsid protein to examine the effects of protein concentration, molecular crowding, and protein flexibility on the spontaneous nucleation and growth of capsid structures in a variety of virion-relevant conditions. These simulations elucidated important natural contributions to both on- and off-pathway capsid assembly processes, providing insights into aspects of virion maturation that are otherwise inaccessible and furthering our understanding of the critical maturation phenomena for the HIV-1 viral lifecycle.

FIGURE 1 (BACKGROUND): Coarse-grained representation of a "mature" HIV-1 viral particle, or virion. The virion membrane *(gray)* encloses a wide variety of molecules, and the spontaneous self-assembly of a cone-shaped "capsid" structure *(green)* is necessary for the virion to infect cells. The capsid is formed by many identical copies of a specific protein molecule, and acts as a container for the viral genome *(red/orange, local capsid surface removed for clarity)*.

METHODS & RESULTS

Simplified "ultra-coarse-grained" (UCG) models [1,2] of the HIV-1 capsid protein were used in combination with specialized, highly scalable molecular dynamics software [3] to examine the self-assembly of viral capsid structures. The simulations illustrated the pronounced sensitivity of capsid protein concentration on the nucleation and growth of capsid structures, with the concentration known to vary significantly across natural HIV-1 virions. Furthermore, the effects of capsid protein structural flexibility were shown to be critical for the controlled growth of capsid structures under the highly crowded molecular conditions expected in virions. The simulations therefore elucidated the critical early stages of capsid self-assembly, providing valuable insight into a process of critical importance to HIV-1 infectivity.

WHY BLUE WATERS?

The Blue Waters HPC resource was critical for both the development of cutting-edge software and the application of this software to perform large-scale biomolecular simulations. Of particular importance was the Blue Waters project staff, providing in-depth technical information and timely advice with respect to the optimal deployment and performance tuning of our software. The next generation Track-1 supercomputing platforms will enable the application of our cutting-edge simulation techniques to still larger systems, enabling the computational study of cell-scale phenomena while retaining microscopic levels of detail.

SOCIAL SCIENCE, ECONOMICS, & HUMANITIES

178 *Policy Response to Climate Change in a Dynamic Stochastic Economy*

180 *Computational Strategies for Applying Quality Scoring and Error Modeling Strategies to Extreme-Scale Text Archives*

182 *An Extreme-Scale Computational Approach to Redistricting Optimization*

POLICY RESPONSE TO CLIMATE CHANGE IN A DYNAMIC STOCHASTIC ECONOMY

Allocation: GLCPC/0.35 Mnh
PI: Lars Hansen[1]
Senior representative: Kenneth Judd[2]
Co-PIs: William Brock[3], Thomas Hertel[4]
Collaborator: Yongyang Cai[1,2]

[1]University of Chicago
[2]Hoover Institution, Stanford University
[3]University of Wisconsin
[4]Purdue University

EXECUTIVE SUMMARY:

We extended past work on computational methods for solving dynamic programming problems by introducing an error estimation and control scheme and implementing more efficient methods for approximating multidimensional functions [1]. 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 the risks in both the climate and economic systems, and to examine how robust our results are to uncertainties about the parameters of those stochastic systems [1,2].

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. Furthermore, 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 a ~10% chance of emissions exceeding all of the IPCC emissions scenarios.

We do not have precise knowledge about many key parameters in our models of the economic and climate systems. We performed extensive uncertainty quantification and showed that the SCC highly depends on the parameter values [1,2].

INTRODUCTION

Every economy is a complex system, but economics researchers typically ignore this by analyzing simple, stylized models of pieces of the system, and use only pencil and paper analyses (supplemented by simple computational models) of these stylized models instead of computationally intensive studies of more realistic models. We are trying to change that by creating robust and general tools that use state-of-the-art numerical methods to study economics problems. In this work, we focus on applications to climate change policy analysis.

Global warming has been recognized as a growing potential threat to economic well-being. This concern has led to an increasing number of national and international discussions on how to respond to this threat. Determining which policies should be implemented will require merging quantitative assessments of the likely economic impacts of carbon emissions with models of how the economic and climate systems interact; this is the purpose of integrated assessment models (IAMs). Most IAMs are deterministic, where economic actors have perfect knowledge of future economic and climate events. Their limitations are due to economists' aversion to modern computational tools. However, all agree that uncertainty needs to be a central part of any IAM analysis.

METHODS AND RESULTS

We developed DSICE [2], a computational, dynamic, stochastic general equilibrium framework for studying global models of both the economy and the climate. We applied it to the specific issue of how the social cost of carbon depends on stochastic features of both the climate and the economy when we applied empirically plausible specifications for the willingness to pay to reduce economic risk.

The inclusion of risks and uncertainties in DSICE makes our SCC analysis among the most computationally demanding problems ever solved in economics. We solved DSICE using numerical dynamic programming—efficient multivariate methods to approximate value functions, and reliable optimization methods to solve the dynamic stochastic problems [3,4].

We showed that economic and climate risks substantially increase SCC, and the SCC is itself a stochastic process with significant variation; for example, the basic elements of risk incorporated into our model cause the SCC in 2100 to be, with 5% probability, five times what it would be without those risks [2]. A July 2014 White House report, titled "The cost of delaying action to stem climate change," incorporated our paper's conclusion that high estimates for SCC can be justified without assuming the possibility of catastrophic events.

Fig. 1 shows four emission scenarios that the IPCC adopted as the range of plausible greenhouse gas emission paths and advocates their use as inputs for climate system models. These scenarios presumably cover the range of likely emission paths. DSICE models a stochastic process, implying that the emission paths cover a range of potential values. The shaded area in fig. 1 displays the range, 1% to 99% quantiles, of business-as-usual (BAU) emission paths implied by a single run using one set of parameters in DSICE. The dotted line shows the 90% quantile of DSICE emission paths. These highlight the importance of including uncertainty in DSICE. DSICE showed that there is a 10% chance that emissions will exceed the highest IPCC scenario. Scenarios with such high emissions are the ones that could lead to the worst outcomes. We found that the representative scenarios advocated by the IPCC ignore 10% or more of the most dangerous plausible emission paths.

DSICE has demonstrated its flexibility in other applications:

- An analysis of the impact of the carbon tax under various continuous climate tipping points showed that the costs of carbon emission used to inform policy are being underestimated and that uncertain future climate damages should be discounted at a low rate [5].
- Examining the impact of ecosystem service risk on the cost–benefit assessments of climate change policies showed that the risk of a tipping point, even if it only has non-market impacts, could substantially increase the present optimal carbon tax [6].
- Examining the implications of future uncertainty in climate impacts, climate regulation, and energy prices on the long-term trajectory of global land use and associated greenhouse gas emissions produced a World Bank Policy Research working paper [7].

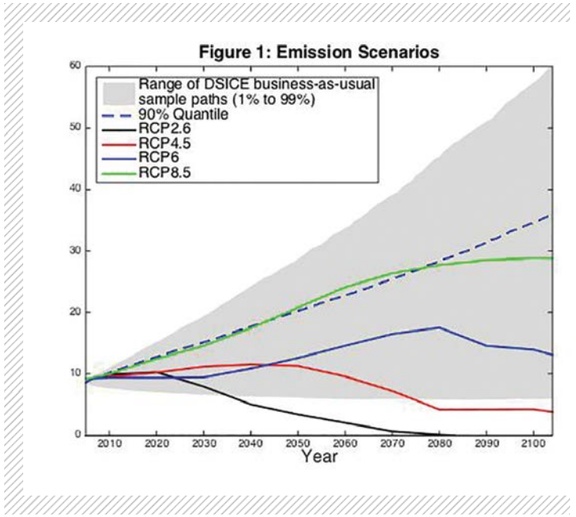

FIGURE 1: Solid lines show four emission scenarios that the IPCC adopted as the range of plausible greenhouse gas emission paths and advocates their use as inputs for climate system models. Shaded area: 1% to 99% quantiles of business-as-usual (BAU) emission paths implied by a single run using one set of parameters in DSICE. Dotted line: the 90% quantile of DSICE emission paths.

We are limited by the current state of knowledge about the critical parameters of the economics and climate system, and different decision-makers usually have different beliefs about the implications of parameter vectors. Over time, new observations will give us new information about critical parameters, such as climate sensitivity. The policy question will be whether to delay mitigation efforts until we have that new information or to proceed with aggressive interventions now.

WHY BLUE WATERS?

DSICE is a high-dimensional dynamic stochastic model that we have thus far applied to a nine-dimensional problem in [2] and a fourteen-dimensional problem in [7]. Solving these problems was very time consuming, but parallelizing the numerical dynamic programming methods allowed us to solve them efficiently [8]. For example, in the benchmark example of [2] DSICE scaled almost linearly to 84,000 cores; a serial computation would take about 77 years. Likewise, our algorithm for solving a dynamic game with multiple players and many choices [10] would take decades of serial computation. We are developing parallel computational methods to solve even higher-dimensional problems. For example, we applied a parallel algorithm to approximately solve a

400-dimensional dynamic stochastic problem in [9].

Our analyses also required extensive uncertainty quantification, which fortunately is well-suited to Blue Waters. Our applications of uncertainty quantifications used thousands of parameter specifications in [2], and hundreds in [1].

In the next year, we will study the impact of multiple interacting tipping points on the social cost of carbon in a fifteen-dimensional version of DSICE. We will also examine the effects of learning. The high-dimensional nature of both models will require Blue Waters' capabilities.

PUBLICATIONS

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. Clim. Change*, 5 (2015), pp. 441–444, doi:10.1038/nclimate2570.

Cai, Y., K. L. Judd, T. M. Lenton, T. S. Lontzek, and D. Narita, Environmental tipping points significantly affect the cost-benefit assessment of climate policies. *Proc. Nat. Acad. Sci.* U.S.A., 112:15 (2015), pp. 4606–4611, doi:10.1073/pnas.1503890112

COMPUTATIONAL STRATEGIES FOR APPLYING QUALITY SCORING AND ERROR MODELING STRATEGIES TO EXTREME-SCALE TEXT ARCHIVES

Allocation: Illinois/50 Knh
PI: Scott Althaus[1,2]
Co-PIs: Loretta Auvil[3], Boris Capitanu[3], David Tcheng[3], Ted Underwood[1]

[1]University of Illinois at Urbana–Champaign
[2]Cline Center for Democracy, University of Illinois at Urbana–Champaign
[3]Illinois Informatics Institute, University of Illinois at Urbana–Champaign

EXECUTIVE SUMMARY:

An important barrier to extreme-scale analysis of unstructured textual data digitized from printed copies using optical character recognition (OCR) techniques is the uncertain quality of the textual representations that have been made from scanned page images. We used Blue Waters to evaluate OCR errors on the HathiTrust Public Use dataset, which is the world's largest collection of digitized library volumes in the public domain, consisting of 3.2 million zipped files totaling nearly 3 TB. We also used Blue Waters to assess the impact of OCR errors on event-detection algorithms using a collection of 16 million articles from The New York Times. Our aim is to develop error quality scoring and correction strategies that can enhance the ability of data analytics researchers to work with digitized textual data at extreme scales.

INTRODUCTION

Researchers in the humanities and social sciences often analyze unstructured data in the form of images and text that have been scanned and digitized from non-digital sources. For this type of research, the most important barrier to conducting extreme-scale analysis of unstructured data is the uncertain quality of the textual representations of scanned images derived from optical character recognition (OCR) techniques. Our project is using Blue Waters to detect, score, and correct OCR errors in the HathiTrust Public Use dataset, which is the world's largest corpus of digitized library volumes in the public domain, consisting of over 1.2 billion scanned pages of OCR text. We are also using Blue Waters to assess the impact of OCR errors on natural language processing algorithms using a corpus of 16 million historical newspaper articles from The New York Times (NYT).

METHODS & RESULTS

In collaboration with the HathiTrust Research Center (HTRC), we performed two major computations with their data. Each of these implementations leveraged Akka [1] to provide a high-performance JVM-based framework featuring simple concurrency and distribution. The first HTRC-related computation on our Blue Waters allocation supported the evaluation of OCR errors in the HathiTrust public domain volumes (which at the time was 3.2 million volumes). We applied new error-detection algorithms produced by a Mellon-funded project called eMOP under Laura Mandell at Texas A&M. Our team's analysis of text-level quality problems will be compared and correlated with image-level quality analysis undertaken by Paul Conway of the University of Michigan on the page images from which the OCR text was extracted. Our aim is to identify how OCR errors are correlated with specific types of image distortion so that digital librarians can better anticipate quality problems from their digitization efforts.

The second HTRC-related computation on our Blue Waters allocation supported the creation of the HTRC Extracted Features (EF) dataset [2], where "features" are notable or informative characteristics of the text [3–5]. The dataset is derived from 4.8 million HathiTrust public domain volumes, totaling over 1.8 billion pages, 734 billion words, dozens of languages, and spanning multiple centuries. We processed a number of useful features at the page level including part-of-speech tagged token counts, header and footer identification, and a variety of line-level information. The EF dataset was used in the creation of the HathiTrust Bookworm [6], which is a tool that visualizes language usage trends in repositories of digitized texts in a simple and powerful way.

Our Blue Waters allocation is also supporting a third set of computations designed to clarify the impact of OCR error on natural language processing algorithms. The Cline Center for Democracy has access to the entire population of NYT articles from 1980 to 2005 in two different forms: "born-digital" content that contains pristine textual data, and ProQuest's digitized version of the same content that was derived from microfilm images using OCR. We used Blue Waters to deploy Phoenix civil unrest event identification software [7] produced by the Open Event Data Alliance (OEDA) over the combined 16 million articles in this NYT corpus. Phoenix requires making computationally intensive parse trees for the articles in order to identify events, and we have a prototype of this deployed, again using Akka, to distribute the work of creating parse trees, identifying events, and writing results. We also plan to compare 25 years of NYT OCR articles to the born-digital NYT articles using software developed for the eMOP project. This will help us evaluate OCR quality and determine the level and type of correction that most closely restores noisy OCR to resemble the born-digital content so that automated correction algorithms can be optimized for use with OCR data from ProQuest's historical newspaper collection.

WHY BLUE WATERS?

The computational demands of using natural language processing, machine learning, or rule-based scoring strategies on the scale of the HathiTrust Public Use corpus would severely tax the capabilities of other HPC platforms. Only Blue Waters offers the computational scale required to carry out the necessary quality scoring and error correction strategies in a timely fashion on an unstructured text corpus of this size.

AN EXTREME-SCALE COMPUTATIONAL APPROACH TO REDISTRICTING OPTIMIZATION

Allocation: Illinois/0.6 Mnh
PI: Shaowen Wang[1,2,3,4,5,6]
Co-PIs: Wendy K. Tam Cho[1,6,7,8], Yan Liu[1,2,6]

[1]CyberGIS Center for Advanced Digital and Spatial Studies, University of Illinois at Urbana-Champaign
[2]Department of Geography and Geographic Information Science, University of Illinois at Urbana-Champaign
[3]Department of Computer Science, University of Illinois at Urbana-Champaign
[4]Department of Urban and Regional Planning, University of Illinois at Urbana-Champaign
[5]Graduate School of Library and Information Science, University of Illinois at Urbana-Champaign
[6]National Center for Supercomputing Applications, University of Illinois at Urbana-Champaign
[7]Department of Political Science, University of Illinois at Urbana-Champaign
[8]Department of Statistics, University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY:

The goal of this phase is to develop a scalable political redistricting application and use the application to integrate significant computing power from Blue Waters with our computational redistricting strategies for analyzing, contextualizing, and understanding redistricting plans. Because the redistricting problem is complex and computationally challenging, existing algorithms were feasible for only very small redistricting problems. We are at the cusp of being able to draw insights from both statistical models as well as computational models for this task.

We have developed a scalable computational approach that allows us to synthesize and organize massive amounts of computation and data to evaluate redistricting schemes and tailor them to notions of "fairness" and democratic rule. Specifically, we extended a scalable parallel genetic algorithm (PGA) to a parallel evolutionary algorithm (PEA) for redistricting. In the PEA, we developed a set of spatial evolutionary algorithm (EA) operators to handle the costly spatial configuration of redistricting maps. Based on the PEA, we formulated a novel set of statistical measures for identifying partisan gerrymanders. Preliminary results of 399,562 high-quality redistricting maps were obtained on Blue Waters using up to 32,768 processor cores. These maps were used to conduct a comparative study of existing redistricting plans and draw insights that were not previously possible.

INTRODUCTION

Redistricting or drawing electoral maps amounts to arranging a finite number of indivisible geographic units into a smaller number of larger areas (i.e., districts). A computational approach is developed to harness the massive computing power provided by Blue Waters to quantitatively study redistricting, e.g., partisan gerrymandering. Partisan gerrymandering comes in two flavors. The first is partisan gerrymandering in which one party seeks to disadvantage the other party by creating districts that pack excessive numbers of minority party voters into districts. This tactic wastes some set of the minority votes while more efficiently dispersing the majority voters among districts to ensure victory. Another form of political gerrymandering is bipartisan gerrymandering where the two parties, majority and minority, join forces to create a "sweetheart deal" wherein the incumbents of both parties are protected in safe seats, thereby preserving the status quo. Both types of gerrymanders seek to minimize the role of the voters. To identify a partisan gerrymander, we must have the ability to compare large numbers of possible plans. Creating a large number of possible plans has not previously been possible, so while this idea is theoretically attractive, it remains a fundamentally new way to conceptualize gerrymanders.

METHODS AND RESULTS

In our computational model, we formulate the redistricting problem as a combinatorial optimization problem that identifies redistricting plans as it attempts to optimize objectives (e.g., competitiveness, safe districts, incumbent protection) while simultaneously satisfying legal constraints (e.g. contiguity and equipopulous districts). The redistricting problem is NP-hard. The solution space of a large redistricting problem is not characteristic of a rugged solution space. While the space landscape is hilly in the sense that it has the usual peaks and valleys, these peaks and valleys are not a rapid succession of precipices but instead a series of vast plateaus, and hence, not rugged in the traditional sense.

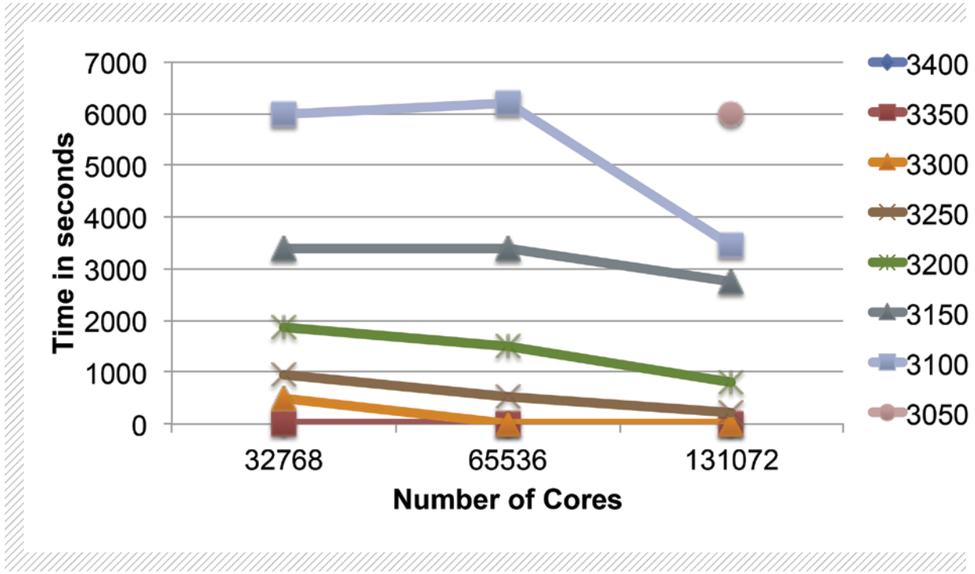

FIGURE 1: Weak scalability results. Execution time is measured at multiple solution quality thresholds. Data: North Carolina VTD.

We employ an evolutionary algorithm (EA) to efficiently find optimal or near-optimal solutions. The EA, which is a stochastic algorithm, is developed to handle spatial configuration of redistricting maps and enable the search process to avoid simple local optima. The parallelization of the EA provides a scalable computational approach for using a large number of processors to work on many plateaus simultaneously and jump from one to another through non-blocking inter-process communication for elite solution propagation and the collective but independent evolutionary searches surrounding these elite solutions.

We developed our parallel EA (PEA) by extending an existing PGA library to improve its scalability and provide explicit programming control of parallelism to overlap the computation and communication. The PEA code exhibited desirable scalability to up to 131,000 integer cores on Blue Waters with marginal communication cost (0.015% on 16,384 cores, compared to 41.38% using synchronous communication). It can conduct 22 EA iterations per second.

To demonstrate how our approach sheds light on redistricting plans, we examined North Carolina at the voting district (VTD) level using 2010 census data. On Blue Waters, 32,768 processor cores were used to identify 399,562 feasible solutions. These solutions are legally possible (contiguity requirements) and reasonable redistricting plans (with population deviation under 5%, compactness scores of 0.95 or better, and respect for political subdivisions). They form a solution pool that characterizes the underlying distribution of all solutions beyond a threshold of goodness. Based on this distribution, the 2001 redistricting plan is an outlier except on biasedness and efficiency gap, and the year 2011 plan (the current plan) is an outlier except on responsiveness.

Our computational approach has the potential to fundamentally transform our understanding of the redistricting process by providing a widely accessible objective tool that will open the redistricting process to participation by a broader and more diverse group of stakeholders. It will also provide great flexibility and enhanced capability for designing redistricting plans.

WHY BLUE WATERS

With the aid of supercomputers, no longer are we merely at the point of solely using a computer interactively to complete simple tasks, we are now contemplating the ability to employ computationally intensive models that can

FIGURE 2: An example redistricting plan found by the PEA and histograms of partisan gerrymandering measures, compared with two existing plans.

generate and evaluate alternative redistricting schemes and compare them to each other on various notions of "fairness," far exceeding human capabilities and analysis. No longer are we constrained to consider a plan in isolation. We have made considerable progress with the aid of the Blue Waters supercomputer and the superb support from the Blue Waters team, who initially created and enabled scalable computational software to support domain science research. We expect to access the next generation of Blue Waters to further develop our statistical and computational approach and achieve the capability of solving larger redistricting problems (e.g., census block-level redistricting with 0.5 million problem variables).

PUBLICATIONS

Liu, Y.Y, W.K.T Cho, S. Wang (2015). A Scalable Computational Approach to Political Redistricting Optimization. *Proceedings of the 2015 Annual Conference on Extreme Science and Engineering Discovery Environment (XSEDE'15),* July 26-30. St Louis, Missouri.

Liu, Y. Y. and S. Wang (2015), A Scalable Parallel Genetic Algorithm for the Generalized Assignment Problem, *Parallel Computing,* 46, pp 98-119.

GRADUATE FELLOWS

188 *Examining the Evolving Properties of Dwarf Galaxies using the ART Code*

190 *Energy Balance Between Thermal and Nonthermal Ions in the Solar Wind Flow*

192 *Understanding the Impact of Silent Data Corruption on HPC Applications*

194 *High Accuracy Radiative Transfer in Cloudy Atmospheres*

196 *Development of a GPU accelerated algorithm to study the absorption of light by the photoactive yellow protein*

198 *High Reynolds Number Computational Aero-Optics*

200 *Distributed State Estimation Algorithms for Electric Power Systems*

202 *Massively Parallel Graph Analytics*

204 *Reduced Communication and Improved Scaling of the Calculation of Dielectric Screening in Nanostructures, Molecules, and Solids*

206 *Sensitivity of Available Potential Energy to Changes in Surface Forcing in the Southern Ocean*

EXAMINING THE EVOLVING PROPERTIES OF DWARF GALAXIES USING THE ART CODE

Kenza Arraki, New Mexico State University
2014-2015 Graduate Fellow

RESEARCH SUMMARY

The standard cosmological model has been successful in predicting many large-scale properties of the universe. Despite these successes, it faces problems explaining the number and structure of satellite and isolated dwarf galaxies. The inclusion of hydrodynamics is critical to understanding galaxy evolution and has repeatedly been found to improve the agreement between observations and simulations. However, including gas and stars in simulations is computationally expensive, and it has only recently become possible to simultaneously resolve sub-galactic scales and include physically motivated models for star birth and feedback.

Dwarf galaxies are by definition the smallest galaxies in the universe and, as such, they provide a window into how galaxy evolution proceeds at the extreme low-mass end. The interplay of gas, stars, and dark matter is very important for these small galaxies. Evolutionary differences that may be unimportant for more massive galaxies like our Milky Way, such as the rate at which supernovae explode, determine the fate of dwarf galaxies. Understanding dwarf galaxy evolution requires a sophisticated understanding of the many physical processes that are integral to galaxy formation. Therefore, creating simulations that reproduce observational properties is a good test of the robustness of the physical processes included in our simulations. By creating realistic galaxies at an extreme end of galaxy formation, we come closer to understanding galaxy evolution more generally. This is important not just for theoretical astrophysicists but for all astrophysicists, as understanding how the universe and galaxies evolved over time remains one of the most fundamental questions in astronomy.

In order to simulate galaxy formation and evolution we use the Adaptive Refinement Tree (ART) code. ART is an adaptive mesh refinement N-body+hydrodynamics code written by Andrey Kravtsov [1] and improved upon by Daniel Ceverino [2] and Sebastian Trujillo-Gomez [3]. ART uses the adaptive refinement technique of increasing spatial and temporal resolution by refining regions of interest in order to efficiently improve the dynamic range of simulations. This allows us to simultaneously have a cosmologically representative volume as well as resolve much smaller sub-galactic scales. This is especially important for studying dwarf galaxies, which require even higher spatial and mass resolution to properly examine.

The code includes a deterministic model of star formation, metal advection, and treatment of supernovae type II and Ia, radiation pressure, and HII region photo-ionization pressure. Using these physical models of star formation and feedback we have found that the total stellar masses of massive galaxies are reduced and come into closer agreement with observations [4]. We also find that our simulated dwarf galaxies reproduce observations of rising rotation curves and increased star formation rates as late times [3]. In addition, examining the satellites of massive galaxies at redshift one has revealed a large spread in the stellar masses and total masses of satellite galaxies. These galaxies have varied star formation histories that increase in complexity with increasing stellar mass, in agreement with observations of Local Group dwarfs [5].

These results indicate that we are implementing physical processes that are able to produce galaxies with properties similar to observations across a wide range of mass scales. We are continuing this work by adding new simulations that will be run down to redshift zero in order to better compare with observations of the local universe.

WHY BLUE WATERS

Having access to Blue Waters and the staff at NCSA has been essential to the progress of this research. Examining dwarf galaxy evolution requires a large dynamic range of physical and temporal scales, which is impossible without high-performance computing. The conception and development of this project would not have been possible without the Blue Waters nodes with their large number of cores, large memory, quick I/O, and access to large amounts of short-term and long-term storage space. The Blue Waters Graduate Fellowship provided me with both the necessary node hours to run simulations of dwarf galaxies and access to knowledgeable staff members who assisted my code development projects.

FIGURE 1 (BACKGROUND): Shown is an image of a galaxy at redshift three (z=3) from an ART simulation, where the colors denote the density of the gas (high-density gas in white and low-density gas in dark blue). This is one of the simulations being used to examine isolated and satellite dwarf galaxy evolution.

Kenza Arraki is in the fifth year of her Astronomy PhD studies at New Mexico State University, working with advisor Anatoly Klypin. She expects to graduate in May 2016 and hopes to obtain a software developer position that allows her to continue creating software to improve the dissemination of scientific results.

"The Blue Waters Graduate Fellowship has provided me with many resources, including supercomputing time to complete my research project, mentors to assist with code development, and a community of peers and experts. Some the most unexpected support came from being exposed to the breadth of careers available to graduates from science PhD programs during SC14 and the Blue Waters 2014 and 2015 symposia," Arraki says. "Access to these resources has fostered my career goals of contributing to software tools that simplify the ability to obtain and share research results. This includes improving how scientists manage, share, and access data products, as well as developing open-source analysis tools and visualization software."

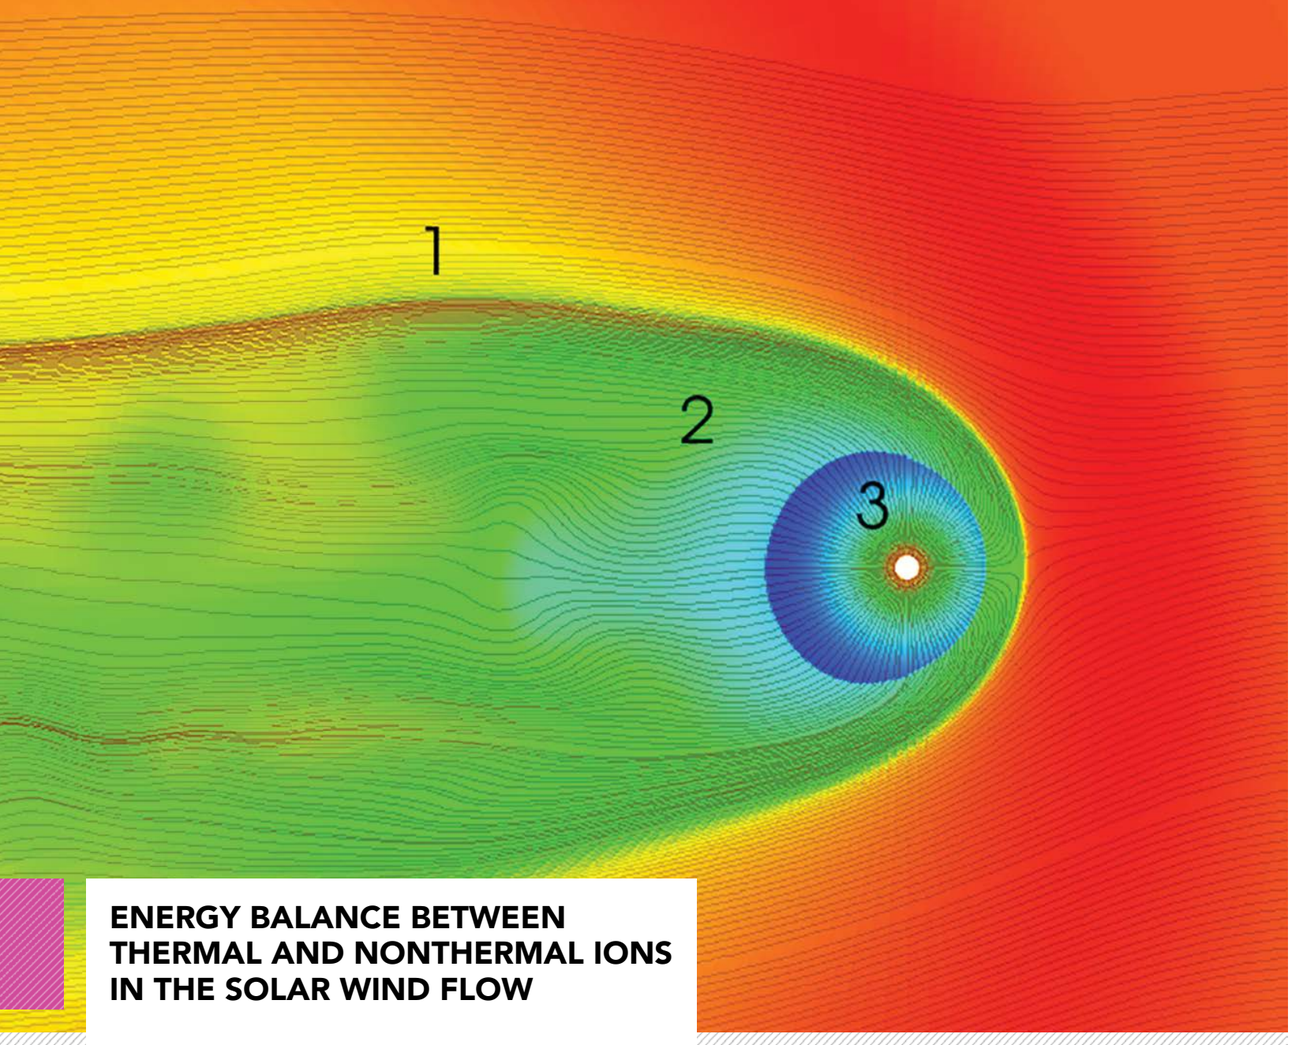

ENERGY BALANCE BETWEEN THERMAL AND NONTHERMAL IONS IN THE SOLAR WIND FLOW

Matthew Bedford, University of Alabama in Huntsville
2014-2015 Graduate Fellow

RESEARCH SUMMARY

The interaction of the solar wind (see region 3 in Figure 1) with the interstellar medium (region 1) is strongly affected by charge exchange between solar wind ions and interstellar neutral atoms, which produces hot "pickup" ions. Due to their much greater temperature, pickup ions must be modeled as a separate fluid. Observations from spacecraft such as Voyager and IBEX imply that pickup ions carry the bulk of the thermal energy in the solar wind plasma. Moreover, the sparsity of in situ measurements and large uncertainties in measured quantities make accurate numerical models essential to our understanding of the outer heliosphere (region 2). For pickup ions in particular, there are no measured quantities at all—everything we know about them is inferred from measurement of some other quantity. I use a multi-fluid magnetohydrodynamic approach with different sets of equations for the plasma mixture, pickup ions, interstellar neutral atoms, and additional neutral populations resulting from charge exchange. Physical processes at the termination shock (region 3 boundary) and charge exchange source terms govern the production of pickup ions, which in turn affects the size and shape of the heliosphere. This is important not only to space scientists but to anyone studying partially ionized plasma flows. Some of the same physical processes occur at different scales in, for example, fusion experiments.

WHY BLUE WATERS

Fluid simulations at high resolution and in 3D must be performed on a machine like Blue Waters with fast processors, lots of memory, and a file system adept at handling many large files. What makes Blue Waters different from similar machines is not just the advanced hardware but the size of the community. Fewer groups using the machine means short queue times, which is good for testing changes to the code, and more responsive, detailed support, so any problems are quickly addressed. Queue times on other machines can be quite long even for short jobs, which is very frustrating during development. The support staff also help with code optimization, profiling, etc. in addition to fixing problems.

PUBLICATIONS

M. C. Bedford, N. V. Pogorelov, A Multi-ion Model of the Heliosphere with Secondary Charge Exchange, 2014. Poster presented at American Geophysical Union, San Francisco, CA.

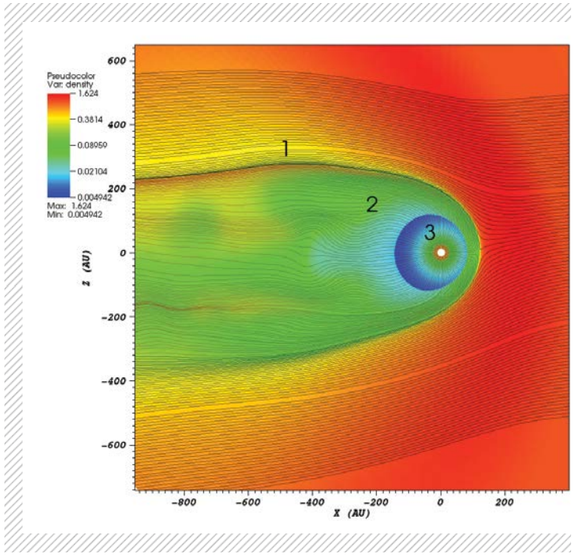

FIGURE 1: Plasma density and streamlines from a heliospheric simulation performed on Blue Waters. Interstellar medium (1) flows around the heliopause. Solar wind (3) flows from the Sun until it is compressed, slowed, and deflected at the termination shock. In the inner heliosheath (2), pickup ions are in the minority in terms of density but nonetheless carry the majority of the thermal pressure.

Matthew Bedford is in the final year of his PhD studies at the University of Alabama in Huntsville, working with advisor Nikolai Pogorelov (who is also has an allocation on Blue Waters). He plans to graduate in spring 2016.

"I plan to apply the skills I have developed in fluid modeling and high-performance computing as a result of my work on Blue Waters to a career in HPC involving space science," he says. "The Blue Waters staff has been very encouraging. I was not exactly apprehensive about job prospects before, but now I am even more focused toward working with supercomputers. Whatever career I choose after this, it will involve supercomputing—not only because that is what I have trained for, but because I genuinely enjoy it."

UNDERSTANDING THE IMPACT OF SILENT DATA CORRUPTION ON HPC APPLICATIONS

Jon Calhoun, University of Illinois at Urbana-Champaign
2014-2015 Graduate Fellow

RESEARCH SUMMARY

High-performance computing (HPC) systems are becoming larger and more complex with each generation. With machines of this size and complexity, challenging issues related to resiliency emerge. One of the most damaging resiliency-related issues of current and future HPC systems is silent data corruption (SDC). SDC refers to any change in program state that is not due to normal program operation. SDC commonly occurs as alpha particles from chip packaging and neutrons from cosmic radiation interact with transistors, causing transient disruptions in their operation that can result in incorrect signals. The incorrect signals can propagate and manifest as bit-flips in combinatorial logic or storage cells. An example of SDC from the application's point of view is when a value written to memory is not the same when it is read it back, or 2 + 2 = 5.

Operating in such a bizarre environment may seem strange, but it is commonplace as every machine experiences SDC on a daily basis. Some machines suffer bit-flips from cosmic radiation every three to four hours [1]. Protection in hardware is often the cheapest option and is the reason why memories such as DRAM have error-correcting code (ECC). Chip logic can be protected but not as efficiently as memory and is thus largely unprotected on current systems. As we scale process technologies and see mass adoption of features such as near-threshold-voltage (NTV), the rate of silent errors generated by logic is expected to increase [2]. Without protection in hardware these silent errors will propagate to program state, leading to SDC, which in turn can lead to application crashes or incorrect results.

Because SDC affects application state, understanding how much SDC applications can tolerate and how applications can detect it are open research questions. Figure 1 shows SDC can be masked by the application entirely or partially by not leading to a detectable event naturally (i.e., program crash). In this figure, as the linear solver algebraic multigrid (AMG) solves a 2D Laplacian, we flip a bit in the binary representation of a floating-point number during a residual calculation in iteration one. Flipping a bit in the mantissa has relatively no effect as we converge to the correct solution and in the same number of iterations. Only when we flip a bit in the exponent do we see effects of SDC emerge. Looking at the final converged solution, this error is masked as we converge to the correct solution, but when looking at the residual history we see that we require greater than 5x the number of iterations for some bit positions.

A multitude of SDC detection schemes have been proposed, most of which are application specific. Application- or algorithmic-based fault tolerance seeks to mitigate SDC's impact on HPC applications by various combinations of redundancy, invariant checking, and heuristics. Our approach to SDC inside AMG is threefold: 1) define recovery points inside the application at three granularities (local function, hierarchy level, and iteration) to restart from when SDC is detected; 2) create a series of detectors based on heuristics and mathematical theory; and 3) combine recovery points and detectors in various ways to yield resilient versions of AMG with varying overheads. The three configurations we test include No Detectors (default AMG), Low Cost detection, which includes abnormal residual detection/recovery and local segfault recovery, and All, which includes everything from the Low Cost configuration and monitors the energetic stability of the problem being solved. Maximal overheads for the resilient schemes are 1% and 19%, respectively, for the scales tested. We use the fault injector FlipIt, developed to support my Blue Waters Graduate Fellowship, to generate transient errors that can lead to SDC. When injecting a single fault, the probabilities of converging to the correct solution for the three configurations are 73.5%, 98.6%, and 99%, respectively. Although an unprotected AMG is not suitable for use in an environment where faults are common, paying even 1% in overhead yields a near perfect probability of convergence. In fact, our schemes have a >60% probability of converging when an average of 37 faults are injected.

Results from our AMG study show that current applications are not suitable in environments where transient errors are common. Paying even a modest amount of overhead for resiliency mitigates the effects of SDC, allowing AMG to retain a high probability of convergence even in extremely faulty scenarios. Going forward, application developers will need to address the impact of SDC on their applications not only to ensure jobs complete but also so the simulation results remain valid in spite of errors that could have occurred.

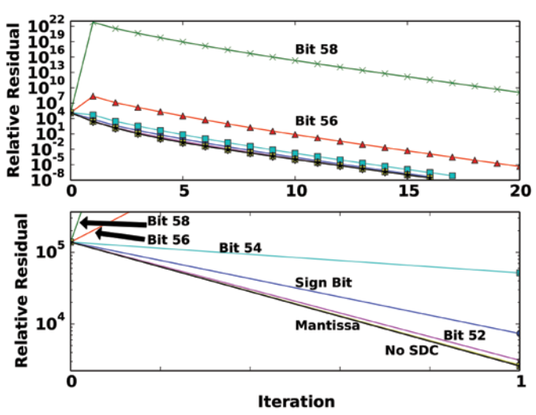

FIGURE 1:
A single bit-flip in an IEEE 754 double-precision floating-point element in the residual vector's effect on the iterations needed to converge. In this location, a flipped bit in the exponent causes an increase in the iterations until convergence, while a flipped bit in the mantissa or the sign bit, bit 63, has no effect as the application masks any SDC that emerges.

WHY BLUE WATERS

Blue Waters is a flagship of future HPC system design that allows me to test the efficacy of a myriad resiliency schemes at large scale. The project staff have been an invaluable source of information and help when working with Blue Waters. Going forward I plan to continue using HPC systems to devise and test low-overhead resiliency schemes to mitigate the effects of SDC on applications.

PUBLICATIONS

Calhoun, J., L. Olson, and M. Snir, *FlipIt: An LLVM Based Fault Injector for HPC.* In Euro-Par 2014: Parallel Processing Workshops, 8805:547-558, Springer International Publishing, 2014.

Calhoun, J., L. Olson, M. Snir, and W.D. Gropp, *Towards a More Fault Resilient Multigrid Solver.* High Performance Computing Symposium 2015 (Society for Computer Simulation International, Alexandria, VA, April 12-15, 2015.

Jon Calhoun is in the third year of his Computer Science PhD studies at the University of Illinois at Urbana-Champaign, working with advisors Luke Olson and Marc Snir. He anticipates graduating in August 2017 and then would like to teach at a university or work at a national laboratory.

"The Blue Waters Graduate Fellowship has solidified my commitment to pursuing a career in HPC," he says.

HIGH ACCURACY RADIATIVE TRANSFER IN CLOUDY ATMOSPHERES

Alexandra L. Jones, University of Illinois at Urbana-Champaign
2014-2015 Graduate Fellow

RESEARCH SUMMARY

One of the most important roles clouds play in the atmosphere is in redistributing radiative energy from the sun and that is emitted from the Earth and atmosphere. 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 satellite observed properties as well as modeled cloud properties. A Broadband 3D Monte Carlo Radiative Transfer model has been developed for massively parallel use on Blue Waters to quantify the relative importance of radiative heating and cooling in various atmospheric scenarios at higher accuracy than ever before. Such a highly accurate model has the capability to provide to the atmospheric radiative transfer community high spatial resolution radiative quantities such as volumetric heating rates, boundary fluxes, and top of atmosphere radiances that can act as standards of comparison for the less accurate—but less computationally-expensive—radiative transfer models more commonly employed.

There is no publicly available 3D broadband Monte Carlo atmospheric radiative transfer model spanning the relevant portion of the electromagnetic spectrum. Through this project one has been developed starting from a monochromatic, forward solving, solar-source-only community model [1] and will be provided back to the community so that improvements to simpler models and advances in scientific understanding can be made faster. A Monte Carlo model uses random sampling from a provided probability distribution to simulate a multitude of potential outcomes that, when aggregated, produce a mean result with a predictable standard error from the true solution that depends on the number of samples aggregated. In the case of a radiative transfer model this amounts to tracing the paths of bundles of radiative energy through the atmosphere, determining frequency of radiation, direction and distance to travel, and interaction mechanism all from provided probability distributions. As it interacts with the atmosphere its contribution to the radiative heating rate, flux, and intensity are accumulated to present an estimate of those quantities without solving the complex radiative transfer equation directly. 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 of one another.

Throughout the development process each stage of model development has been rigorously evaluated against analytical solutions, canonical values from the literature [2,3,4,5], and other models[5,6] that have previously acted as standards of comparison. Challenges in development have included finding the right balance between dynamic load management and communication latency. The ultimate goal of this work is to produce highly accurate 3D broadband benchmark simulations of radiative quantities to be used to evaluate and improve models making simplifying approximations, using the ample computational resources of Blue Waters, in the process creating a community model that can be used to address scientific questions about the interactions of clouds and realistic radiative transfer.

WHY BLUE WATERS

Radiative transfer has a reputation as one of the most computationally expensive components of modeling the dynamics of earth's atmosphere; it is generally considered too expensive to compute as frequently or accurately as other physics. Blue Waters enables me to develop and run a highly accurate radiative transfer model that returns results in a reasonable amount of real time for the scope of my project. Simulations do not wait long in the queue, and the model was easily parallelized across many nodes to reach completion faster. Access to top-of-the-line debugging and profiling tools, such as CrayPat and DDT, allowed me to streamline the development process to guide and implement updates to my model faster. Having access to a point of contact on the SEAS staff helped me think through issues that were too amorphous for the JIRA ticket system and helped me wade through myriad options to find tailored solutions for my problems. 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 plan to extend my allocation for the following academic year as I finish my degree and hope to make use of Blue Waters for the rest of its lifetime.

Alexandra L. Jones plans to complete her Atmospheric Science PhD in December 2015 and would like to then continue to conduct research in the field of atmospheric radiative transfer. Her advisor is Blue Waters Professor Larry Di Girolamo.

"Blue Waters has opened my mind to and given me the credentials for a career path that keep me at the cutting edge of both high-performance computing and atmospheric science," she says. "I strongly believe in the important role each play in driving the other forward. I want to remain in the scientific computing community and continue to promote and advocate for access to these resources to individuals and agencies who can benefit."

DEVELOPMENT OF A GPU ACCELERATED ALGORITHM TO STUDY THE ABSORPTION OF LIGHT BY THE PHOTOACTIVE YELLOW PROTEIN

Sara Kokkila Schumacher, Stanford University
2014-2015 Graduate Fellow

RESEARCH SUMMARY

The photoactive yellow protein (PYP) is a photosensor found in the bacterium Halorhodospira halophila [1,2]. When this protein absorbs blue light, it undergoes an isomerization and initiates a photocycle that involves a few key intermediates. This protein has become the model system for studying photoreceptor proteins (see Figure 1) [3]. Since the absorption of light is an inherently quantum mechanical process, the computational study of this property necessitates the use of electronic structure methods. However, there are very few electronic structure methods that are efficient enough and accurate enough to study the absorption of light in PYP or other large chemical systems.

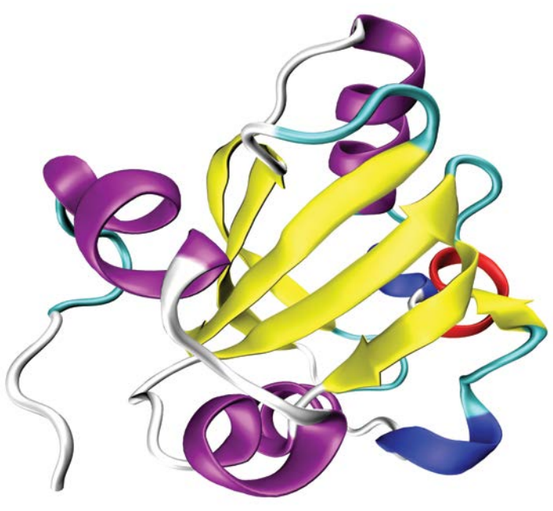

FIGURE 1: Illustration of the photoactive yellow protein. This image was created using VMD [4].

To study the absorption of light in large chemical systems, there is a need for accurate and computationally inexpensive methods. In this work, I have developed an algorithm based on the equations-of-motion approximate singles and doubles coupled-cluster (EOM-CC2) method, which is often used as a reliable benchmark method to study the absorption of light [5,6]. Unfortunately, EOM-CC2 has a computational cost of O(N5) which limits its usefulness as the size of the chemical system is increased. By incorporating the tensor hypercontraction (THC) approach into the EOM-CC2 method, a new algorithm was designed with the computational cost of O(N4) [7-9]. Further efficiency can be achieved by formulating the algorithm in a way to accelerate the method via the use of graphical processing units (GPUs). This new algorithm is not only useful for the study of PYP's spectral properties; it is a general algorithm that can be used to study the absorption of light of molecular systems.

I have developed a GPU-accelerated version of the THC-EOM-CC2 method. Timings indicate that the new algorithm does achieve a computational cost of O(N4). Comparing the timings of THC-EOM-CC2 on a single GPU to the timings achieved on a single CPU, this approach achieves an 8-15 times speedup over the CPU-based implementation of THC-EOM-CC2 with timing variations due to the size of the chemical system and the type of CPU or GPU used. Additionally, the CPU-based version of THC-EOM-CC2 was shown to accurately predict the excited states for a wide variety of molecules [9].

We are in the process of using the GPU-accelerated THC-EOM-CC2 method to study the absorption spectra of PYP and its photocycle intermediates. Initial gas phase results show that this new approach predicts the absorption maxima to be within 0.04 eV of the experimental absorption peak. The initial results indicate that the THC-EOM-CC2 approach can be used to accurately predict the absorption spectra of PYP and its photocycle intermediates. Further work includes investigating the influence of different quantum mechanical regions on the predicted absorption spectra and computations with the solvated protein.

The GPU-accelerated THC-EOM-CC2 method is a new approach that can be used to study the absorption of light in many different chemical systems. By developing a more efficient and highly parallelized electronic structure method for the study of light absorption, researchers will be able to use THC-EOM-CC2 to further research in areas such as solar cell design, photoreceptor proteins, and materials development. This work also illustrates an approach to developing software that uses GPUs in high-performance computing and a new approach to scaling coupled-cluster algorithms across nodes.

WHY BLUE WATERS

Blue Waters is one of the few supercomputers with many GPU nodes and large amounts of memory. Without this combination, some of the chemical system sizes would be inaccessible due to either limitations of CPUs or memory. Determining an absorption spectrum requires extensive sampling. The use of Blue Waters allowed us to simultaneously carry out calculations of many different configurations.

Additionally, this is the first time I have used a supercomputer. The Blue Waters staff suggested different debugging strategies, ways to improve efficiency, and data management approaches, all of which have helped me learn how to develop software for supercomputers.

This research would not have been possible without their support. I plan to continue to develop electronic structure programs designed for high-performance computing. In addition to developing software, I am interested in using supercomputing to study chemical systems that are too computationally demanding for typical computing resources.

PUBLICATIONS

Kokkila, S.I.L., E.G. Hohenstein, et al., Tensor Hypercontraction Second-Order Møller-Plesset Perturbation Theory: Grid Optimization and Reaction Energies. *JCTC*, submitted

Sara Kokkila Schumacher is in the fourth year of her Chemistry PhD program at Stanford, working with advisor Todd Martínez, and expects to graduate in spring 2016. She then plans to complete a postdoctoral position researching electronic structure methods and molecular dynamics and ultimately aims to become a professor at a liberal arts college, where she would like to develop a course that introduces high-performance computing to science majors.

"The Blue Waters Graduate Fellowship has helped me develop parallel programming skills, especially in the area of adapting programs to make use of GPU accelerators," she says. "The insight I have gained through this fellowship will help me develop parallel quantum chemistry programs for future areas of research."

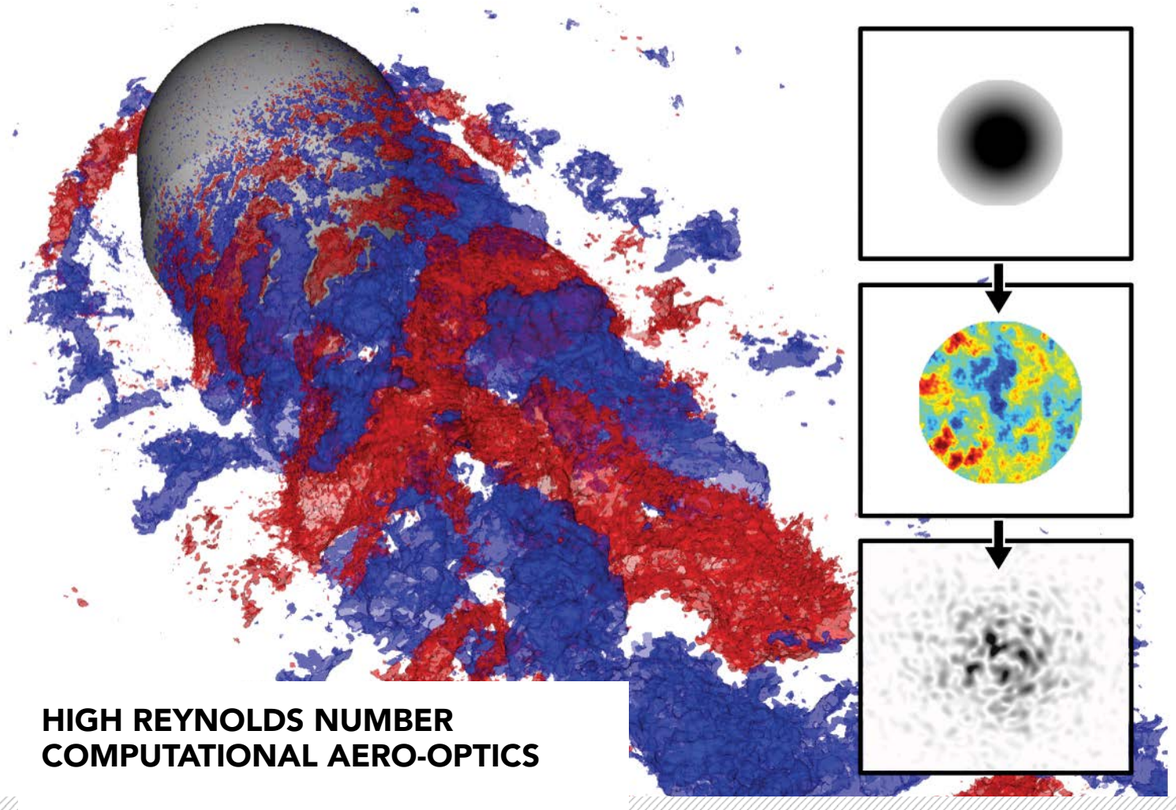

HIGH REYNOLDS NUMBER COMPUTATIONAL AERO-OPTICS

Edwin Mathews, 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 over the aperture, it is distorted due to the non-uniform speed of light resulting from density fluctuations. Even small amplitude distortions in the near field can cause severe performance degradation, as seen in Figure 1.

My research is aimed at improving the predictive capability for, and understanding of, aero-optics at realistic Reynolds numbers and Mach numbers by using wall-modeled large-eddy simulations (LES). To accomplish this, the flow over the optical turret (a hemisphere on a short cylinder mounted to a flat surface) used by the Airborne Aero-Optical Laboratory (AAOL) at Notre Dame was simulated at the actual Reynolds number of 2,300,000 and Mach number of 0.4 on a mesh of over 200 million control volumes. To date, this represents the largest aero-optic simulation of a turret and one of the highest Reynolds number wall-modeled large-eddy simulations performed.

The flow was computed using the unstructured mesh, finite volume compressible LES code CharLES, developed at Cascade Technologies, Inc. CharLES, similar to other LES solvers, computes the spatially filtered Navier-Stokes equations along with the continuity and energy equations. To account for the effect of small-scale fluid motions not resolved by the computational mesh, a subgrid-scale stress model is used. A key component in the ability to simulate the flow at the full experimental Reynolds number is the use of a wall model in conjunction with the LES. The wall model alleviates the severe grid resolution requirement imposed by the small eddies in the near-wall region by accounting for their effects with a RANS-type turbulence model.

To solve for the aero-optical distortion, an efficient parallel aero-optics code was developed and used to compute the optical path length (OPL) in almost 300 viewing directions over the whole turret surface. In all, the simulation generated over 40 TB of optical and flow data that will be analyzed to guide the design of aero-optic mitigation strategies. Beyond classical statistical approaches, we look to use data-mining techniques like Proper Orthogonal Decomposition (POD) and Dynamic Mode Decomposition (DMD) to more fully explore the dynamics of the turret aero-optics.

WHY BLUE WATERS

Beyond the processing power, the ability to generate large quantities of data and efficiently process it on Blue Waters has been vital to our research. We have been able to scale our simulations to a higher level of fidelity and have created a large database that we will use to shed light on important aero-optic phenomena. Because of Blue Waters, we have the unique capability of analyzing the entire flow field to extract important statistics that are computationally demanding due to large I/O and memory requirements. Additionally, using a system like Blue Waters effectively would not have been possible without the help of the Blue Waters and NCSA staff, especially on the use of the Lustre file system and visualization. With their help we were able to successfully accomplish our simulation goals.

FIGURE 1 (BACKGROUND): On the left, contours of density higher (red) and lower (blue) than the mean value in the wake of an airborne optical turret. On the right, a coherent optical beam shown by its intensity profile (top) is aberrated by phase distortions (center) caused by density fluctuations in turbulent compressible flow. As a result, when the beam propagates outward its integrity is severely degraded (bottom).

Edwin Mathews is a fourth-year Aerospace Engineering PhD student at the University of Notre Dame, working with advisors Meng Wang and Eric Jumper. He expects to graduate in 2016 and hopes to continue conducting research in computational fluid mechanics.

"Far beyond aero-optics, computational fluid dynamics has a vast range of applications across several fields, and LES and hybrid LES/RANS methods represent the next generation of CFD techniques for engineering applications. With techniques like wall-modeled LES, resolving unsteady turbulent flows can become practical and bring previously intractable scientific and engineering problems into the domain of possibility," Mathews says. "In the future, I believe the use of these methods coupled with data-mining methods to explore databases of large CFD simulations will empower the design of complex fluid systems. The Blue Waters Graduate Fellowship has enabled my growth in these fields and has been an invaluable experience in working effectively at the petascale level."

DISTRIBUTED STATE ESTIMATION ALGORITHMS FOR ELECTRIC POWER SYSTEMS

Ariana Minot, Harvard University
2014-2015 Graduate Fellow

RESEARCH SUMMARY

Until recently, the use of high-performance computing for advancing the electric power grid has received relatively limited attention. The electric grid is a complex interconnected system with an aging infrastructure spread over thousands of miles. For environmental and economic reasons, it is imperative to improve the operation of the electric grid, and computational science will be essential to addressing these challenges [1].

One key aspect of reliable power system operation is state estimation, which monitors the state of the grid by estimating the voltage magnitudes at each bus, or node, in the network. With the growth in penetration of renewable energy, it is becoming more demanding to estimate the system state promptly and accurately. There have been increasing research efforts in developing and integrating new sensor technology, called Phasor Measurement Units (PMUs), to create a more advanced state estimation system. Introducing PMUs leads to both an increase in the dimension of the data and in the sampling rate that will require using high-performance computing resources [2].

Our research focus is designing decentralized algorithms for coordination of power systems. In distributed approaches, information communication occurs only between neighboring processors rather than via a centralized processor, as depicted in Figure 1 (a). Advantages of distributed approaches include increased robustness as well as reduction in computation, communication, and memory requirements per area. In recent work, we proposed a new fully distributed Newton-type algorithm for state estimation of electric power systems [3]. In contrast to gradient-based methods, Newton methods are of interest due to their convergence rate, which is usually quadratic. The difficulty is that Newton methods require solving a matrix inversion at each iteration. This presents a challenge for developing a distributed method, since in general global information of the matrix entries is needed to invert a matrix. Furthermore, when the system size is large, solving a large matrix inversion may be time-consuming or even prohibitive. To overcome this challenge, we explore the use of matrix splitting techniques [4], [5]. This allows us to exploit inherent sparse structure in power systems in order to calculate the next Newton iteration in a distributed way. This distinct sparsity patterns stems from the fact that power systems follow Kirchhoff's current and voltage laws. We have shown this method to be advantageous with respect to convergence speed and memory requirements [3].

Matrix splitting allows an iterative approach for solving linear systems, Ax = b. Namely, decomposing a matrix into a difference, A = M - N, we have the following iterative scheme: xk+1 = M-1Nxk + M-1Nb. In [3], we show a construction for M and N, which guarantees convergence and allows for distributed calculation with limited communication. Different matrix-splitting schemes exhibit different trade-offs between the amount of computation done per process and the amount of communication required between processes. Characterizing how the rate of convergence and the robustness to imperfect communication depend on this trade-off is an interesting question for the field of distributed estimation algorithms.

In Figure 1 (b), we compare how the computation and communication time scales with the size of the power network. The implementation of our algorithm assigns each bus, or vertex, in the power network a separate MPI process. We see that for network sizes requiring more than one node on Blue Waters, the communication time remains constant, which is a favorable scaling. We aim to continue improving the speed of the computation time in order to make the algorithm of greater use in real-time.

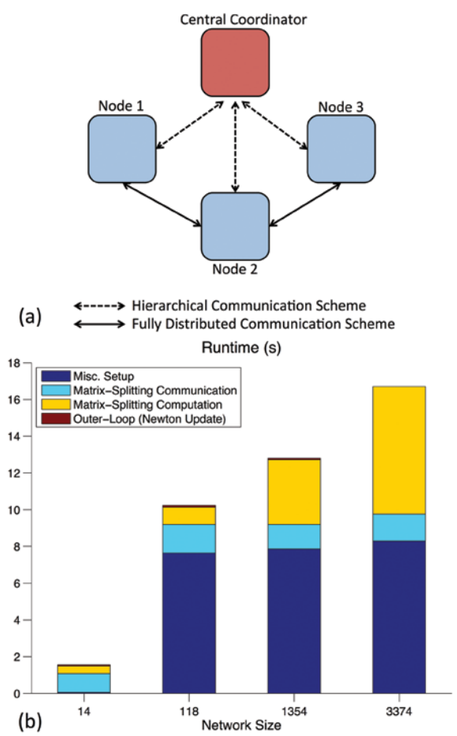

FIGURE 1: (a) Nodes 1-3 represent different control areas, or distinct groups of buses, in the power grid. (b) Using the same number of iterations, compare how computation and communication scale with power network size.

WHY BLUE WATERS

Working with Blue Waters has allowed me to develop and test our distributed estimation algorithm on large-scale realistic power systems using thousands of processes. Such simulations are infrequent in academic papers for this particular problem, and I find this to be an exciting aspect of my work. In the near future, we are interested in using supercomputing to test new parallel, iterative algorithms we are developing for 1) fault detection and 2) optimal power flow problems in power systems. I would like to acknowledge the NCSA staff for their help in sorting out a variety of issues, including building certain software on Blue Waters and overcoming hurdles encountered while scaling my algorithms up for testing on larger systems.

PUBLICATIONS

Minot, A. and N. Li, A Fully Distributed State Estimation Using Matrix Splitting Methods, American Control Conference (AACC, Chicago, Illinois, July 1-3, 2015).

Minot, A., N. Li, and Y. M. Lu, A Distributed Newton Method for Power System State Estimation, Submitted to *IEEE Trans.* on Power Systems (2015).

Ariana Minot works with advisors are Na Li and Yue Lu and expects to graduate from Harvard with a PhD in Applied Mathematics in May 2016. After completing her graduate studies, she hopes to be a researcher at a national lab.

"I have a strong interest in interdisciplinary research in applied mathematics and scientific computing with applications to electric power systems," she says. "The support of the Blue Waters Fellowship has enabled me to pursue this interest and to begin developing new algorithmic tools for transforming power systems operation."

MASSIVELY PARALLEL GRAPH ANALYTICS

George M. Slota, The Pennsylvania State University
2014-2015 Graduate Fellow

RESEARCH SUMMARY

Large-scale and highly irregular graphs, such as those arising from biological systems, social networks, and the general topology of the web, require special treatment when analyzed using parallel and distributed algorithms. This is primarily the result of two characteristics of real-world graphs. First, most real-world graphs have what is termed a highly skewed vertex degree distribution, where some vertices in the graph have orders-of-magnitude greater coincident edges than other vertices. This can create a large work imbalance across threads and tasks when considering work distribution on a per-vertex basis (i.e. a 1 million degree vertex assigned to task A will take considerably longer to process than a 100 degree vertex assigned to task B; task B will sit idle waiting for task A to finish processing). Second, many irregular real-world graphs also have what is called a small-world structure. That is, the number of edges between any two arbitrary vertices in the graph is very low relative to the size of the graph—think "Six degrees of Kevin Bacon" or how, based on Facebook and Twitter, there is only an average of about 4.5 degrees of separation (edges) between any two Internet users on the planet. This small-world nature makes scalable partitioning of graph workloads across a distributed memory environment extremely challenging.

Considering these two challenges, work during this fellowship established a methodology for parallel graph algorithm implementation at various computational levels, including per-core, per-node, and per-system level. Using the techniques and optimizations gleaned from these efforts, a complex analysis of the largest publicly available web crawl [1] at about 3.5 billion pages and 129 billion hyperlinks was performed on Blue Waters.

Initial work focused on optimizing at the task and node level, specifically for the GPUs on Blue Waters' accelerated nodes. The focus of this initial research effort sought to answer the question: "Given a highly irregular input graph and some X thousand processing elements (such as a GPU), how can one efficiently parallelize graph algorithms to ensure high performance under these conditions?" It was observed that a large number of common graph analytic algorithms follow a certain algorithmic template. Prior art considered parallelization over a certain part of this template, and while this works for a small number of processing elements, greater degrees of parallelism require more complex and direct work-balancing methods. Through explicit implementation of a known compiler optimization technique termed the "Manhattan Collapse" [2], an even work balance across all work to be performed with a minor overhead cost was achieved. Multiple common graph analytic subroutines were implemented using this collapse, along with several general GPU optimizations. Performance comparisons to highly optimized CPU code demonstrated at or better than state-of-the-art execution times despite minimal to no per-algorithm optimizations outside of the general ones following the aforementioned methodology. Up to 3.25x speedups were achieved using this methodology.

The second phase of work expanded the optimization methodology to a multi-node setting, using 16 to 64 nodes on Blue Waters. An efficient means of both work distribution, data ordering, and parallel communication for general (irregular) graph algorithms was investigated. Under consideration was the tuple (per-node parallelization and data ordering, per-dataset work distribution across nodes, per-algorithm communication methods) to determine what is most performant across several simple to complex graph analytic algorithms. Multiple vertex partitioning strategies in terms of objectives and constraints were considered. Additionally, multiple intra-task vertex ordering strategies were investigated. The effects of varying partitioning and ordering strategies were considerable. Up to 5x improvement versus naive methods with an intelligent partitioning strategy and up to 1.5x improvement with an intelligent ordering strategy was noted.

Using knowledge gained from these previous efforts, an analysis of the largest publicly available web crawl was undertaken using more than 16,000 cores of Blue Waters. The web crawl was analyzed at various levels of complexity, from simple structural characteristics, to identifying the most important web pages using various metrics, to determining intrinsic community structures. Novel insights were obtained into the structure of the web. One such insight was that community sizes followed a power law with a possible heavy tail distribution, similar to prior observations of other degree and connectivity distributions [3].

Hopefully, this work will allow several directions for future efforts. Using the optimization strategies identified through this work, more complex algorithms might be implemented for continued analysis of the aforementioned web crawl and other similar networks. Additionally, further scaling in the future to larger networks as they become available might provide additional novel insights using Blue Waters. Finally, in the forthcoming years, as both larger networks are available for analysis and larger computational systems become available for doing such analysis, some of the strategies identified might become critical to enable scalable and efficient analyses of these networks on these systems.

WHY BLUE WATERS

Blue Waters was important for this research for two reasons. Firstly, the GPUs that power Blue Waters' accelerated nodes were important for the study of many-core optimizations that composed the first parts of this research effort. Secondly, the high-performance I/O and intra-node communication were crucial to allow the efficient study of the web crawl at scale.

PUBLICATIONS

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).

Slota, G.M., K. Madduri, and S. Rajamanickam, Parallel Graph Algorithms on Modern Systems. *PhD Forum Poster at the 29th IEEE International Parallel and Distributed Processing Symposium (IPDPS15),* Hyderabad, India, May 25-29, 2015.

George M. Slota is now in the third year of his Computer Science and Engineering PhD studies at The Pennsylvania State University. His advisors are Kamesh Madduri, Penn State, and Siva Rajamanickam, Sandia National Labs. He expects to graduate in May 2016 and then plans to pursue a post-doctoral position at a national lab, in academia, or in industry.

"Blue Waters has helped me achieve these goals by giving me access to such a large-scale system necessary for me to succeed in a lot of my research efforts," he says. "The research I've accomplished during this fellowship undoubtedly strengthens me as a candidate for a position at the aforementioned places."

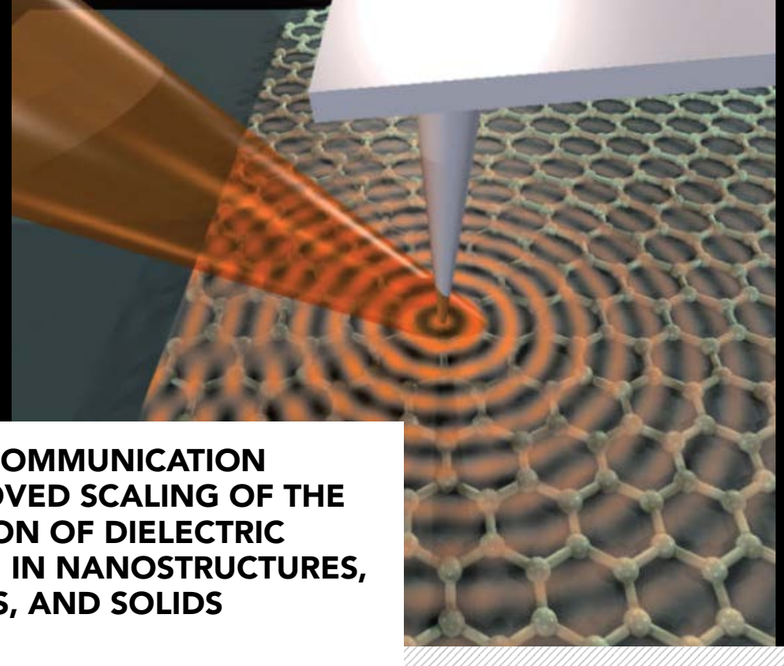

REDUCED COMMUNICATION AND IMPROVED SCALING OF THE CALCULATION OF DIELECTRIC SCREENING IN NANOSTRUCTURES, MOLECULES, AND SOLIDS

Derek Vigil-Fowler, University of California, Berkeley
2014-2015 Graduate Fellow

RESEARCH SUMMARY

In short, the dielectric response encodes all the information about how electrons in a material respond to the presence of one another in the material and to external stimulation by light or an electric field. The dielectric response is needed to compute the energy levels of quantum states, which are needed for understanding energy flow through solids, molecules, and nanostructures. Understanding the dielectric response is thus crucial for applications of semiconductors, metals, and combinations of these materials in devices like transistors or photovoltaics.

The dielectric response is encoded in a frequency-dependent matrix that must be calculated using sums over quantum states in the system, making this is a very large problem that has to be parallelized to scale to big systems. Our code, BerkeleyGW, was previously optimized when frequency dependence was calculated using a model from a single matrix and the matrix inversion step was inexpensive. With frequency dependence, this matrix inversion became a bottleneck, since the size of the matrix was not large enough to saturate the processors that were used earlier (in operations on the quantum states) and there were many frequencies done serially. My approach was to give each processor a larger piece of this matrix, but give each processor only a subset of frequencies to calculate. This allows more processors to work in parallel during the matrix inversion, reduces communication, and allows BerkeleyGW to scale to bigger system sizes.

My approach was very successful, showing speedup for all system types, but especially for more complicated systems. For the systems tested thus far, e.g. bulk semiconductors, like silicon, and molecules, my scheme showed a linear reduction in the time for communication and for the matrix inversion with the number of frequencies computed in parallel. This improvement has reduced the time to perform two components of the calculation so thoroughly that they are no longer the biggest part of the computation by an order of magnitude but instead are on par with other parts of the computation. For systems with big unit cells, like molecules and crystals with many constituent elements, the reduction of the time for inversion and communication is so dramatic as to allow the computation of whole new classes of molecules and crystals that were previously unreachable.

This work could be adopted in other contexts in high-performance computing that involve matrices that have a third index in addition to the usual two and that have operations performed in parallel on the different matrices indexed by this third index. Additionally, I think this work will open up the computation of the dielectric response to whole new classes of materials, which could be used for exciting applications in photovoltaics and next-generation transistors. I personally am excited to use this improved code to study photovoltaics made from thin-film perovskites, which are currently skyrocketing in efficiency and showing a possible way toward cheap, abundant photovoltaics. As with any code development, I am also excited to see how the BerkeleyGW user community will use this new functionality to study new and interesting systems I have not imagined, perhaps achieving breakthroughs in understanding that will enable other improvements in electronics, photovoltaics, and other applications.

WHY BLUE WATERS

The Blue Waters system is second to none in terms of the combination of the scale of jobs one can perform on the machine and the productivity it allows with its rapid queue times (itself a byproduct of the large scale and quality of the machine). I have never gotten more important work done quickly than I did on Blue Waters. Indeed, I think the Blue Waters system enables groundbreaking science to be done more than any other machine I have worked on. The Blue Waters staff is also wonderful, very responsive in requests for help on running codes optimally on the system. More than anything, the Blue Waters team seems to care the most about its users than any other system I have used, pulling out all the stops to help the science teams use the machine to do their important work. I cannot say enough about how impressed I have been with both the Blue Waters machine and staff. I will apply for an allocation on Blue Waters in the future, as I will continue to push the level of parallelism in the calculation of excited-state properties using BerkeleyGW and related codes. Access to Blue Waters will be crucial in this endeavor.

FIGURE 1 (BACKGROUND): Artist's rendition of the creation of plasmon in graphene using a laser. This is an example of dielectric response, where the electrons in a material respond to an external stimulus (the laser light). Image courtesy of Nature Publishing Group.

Derek Vigil-Fowler is in the sixth year of his Physics PhD studies at the University of California, Berkeley, working with advisor Steven G. Louie. He graduated in August 2015 and has begun a Director's Postdoctoral Fellowship at the National Renewable Energy Laboratory.

"My goal is to continue developing physical approximations and algorithmic improvements, especially those that will allow increased parallelism and scaling, to access the physics of the increasingly complicated systems relevant to technological applications, especially photovoltaics and photocatalysis," he says. "The Blue Waters Graduate Fellowship supported me in this endeavor at a critical time in my young career, helping me establish myself in this area, and for that I will always be grateful."

SENSITIVITY OF AVAILABLE POTENTIAL ENERGY TO CHANGES IN SURFACE FORCING IN THE SOUTHERN OCEAN

Varvara Zemskova, University of North Carolina at Chapel Hill
2014-2015 Graduate Fellow

RESEARCH SUMMARY

Although the global ocean is on average stably stratified, horizontal, and even more importantly vertical, mixing exists due to differential heating and cooling by the atmosphere at the ocean surface. These motions distribute heat and nutrients throughout the globe. 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, both directly in form of heat exchange and indirectly in form 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 mixing. The Antarctic Circumpolar Current, which is driven by eastward winds, goes around Antarctica unobstructed by continents, unlike currents in other ocean basins. Because of such predominantly unidirectional flow, it can be modelled in a periodic domain that decreases computational demand without sacrificing completeness of the problem. This region also experiences significant differential buoyancy forcing, such that surface waters near the Antarctic continent are cooled, making them denser and causing them to sink, and surface waters at the northern end are warmed and become less dense. Therefore, this problem can be studied in a form of horizontal convection, a model of a fluid flow forced by differential surface buoyancy forcing.

Three surface boundary conditions were investigated in this preliminary study. The first simulation was conducted only with differential buoyancy forcing with higher density near the Antarctic continent (to model cooling) and lighter fluid at the northern end of the domain (to model heating). The second simulation had the same differential buoyancy forcing as the first simulation, but unidirectional eastward wind stress was added at the surface to model the trade winds in the Southern Hemisphere. The third simulation had the same differential buoyancy forcing but divergent wind stress at the surface, with eastward trade winds in the northern part of the domain and westward polar winds in the southern part. The simulations have to be 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.

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 condition. With my continuing allocation I plan to extend this project by connecting the Southern Ocean to the rest of the global ocean and to analyze the sensitivity of vertical mixing at mid-latitudes to changing boundary conditions in the Southern Ocean.

Varvara Zemskova is in the fourth year of her Oceanography PhD studies at the University of North Carolina at Chapel Hill, working with advisor Brian White. She expects to graduate in May 2017 and then plans to continue research in physical oceanography with a focus on numerical simulations.

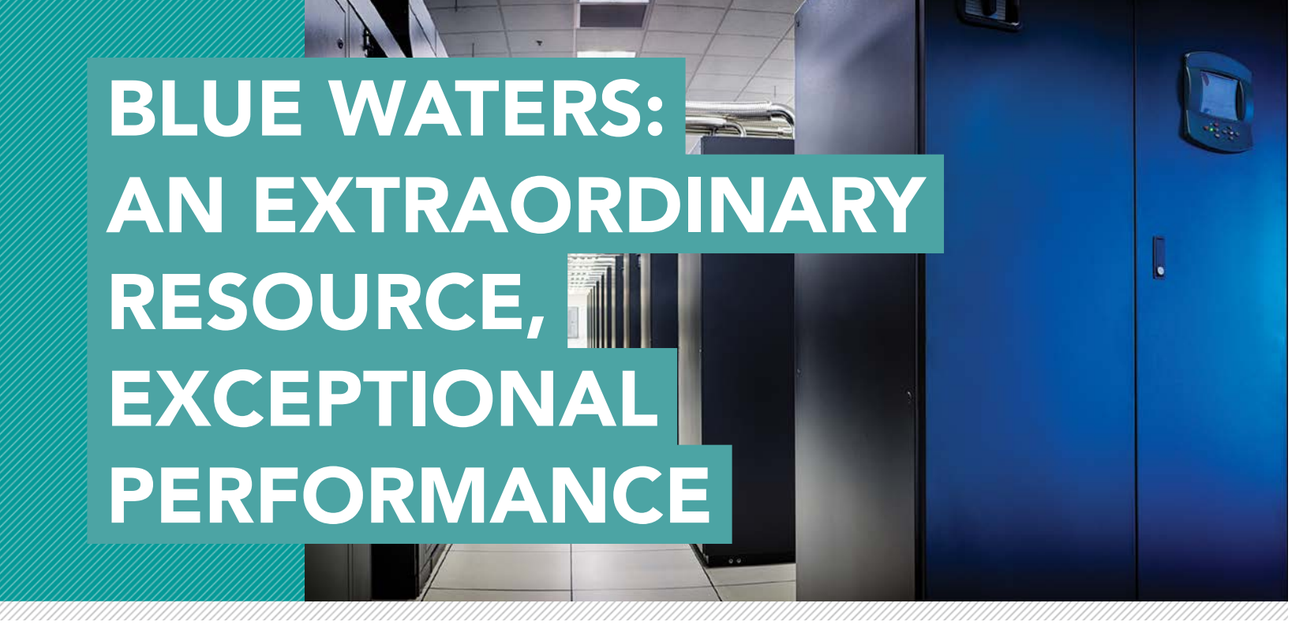

BLUE WATERS: AN EXTRAORDINARY RESOURCE, EXCEPTIONAL PERFORMANCE

Blue Waters is one of the most powerful supercomputers in the world—and given its balanced architecture and generous allocation strategy, it may be the most productive system. Many of the science and engineering projects enabled by Blue Waters would be impossible to carry out elsewhere. Researchers across the country are pushing the boundaries of a wide range of disciplines thanks to Blue Waters; many of them are graduate students who are leveraging the power of Blue Waters to open up new investigative opportunities and to launch their careers.

This year, a wide range of disciplines benefitted from use of Blue Waters. Astronomy and Astrophysics and Particle and Nuclear Physics each consumed slightly more than a quarter of the node-hours (Figure 1), followed by Biology, Biophysics & BioMedicine at 19 percent of the node hours and Atmospheric & Climate Science, Materials Science, and Physics and Chemistry, each of which use 4 to 6 percent of system time.

The Blue Waters project tracks multiple metrics—from usage to downtime to service requests—to ensure that excellent service is provided to our science team partners. In fact, Blue Waters may be the most instrumented system in the world, as we collect more than 3 billion monitoring data points every day. These metrics guide us in providing a reliable, consistent high-performance computing environment that accelerates scientific discovery, from biology to materials science to geophysics to particle physics to astronomy and across many other domains. For its second year of operations, the Blue Waters project met or exceeded the expectations for the large majority of our stringent control metrics. The metrics are reported for 11 months from July 1, 2014 to May 30, 2015.

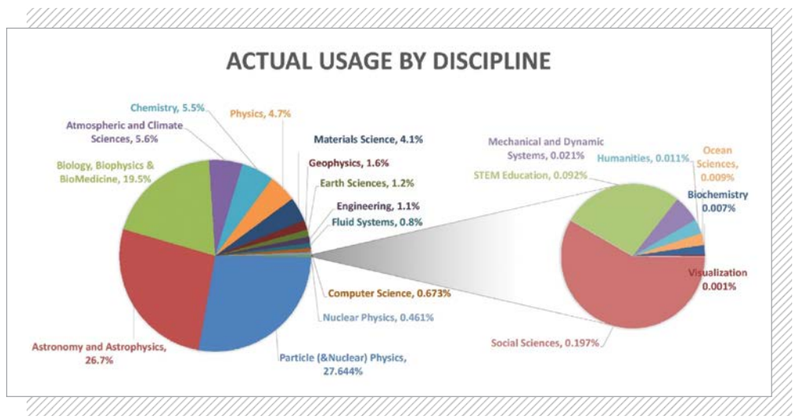

FIGURE 1: Distribution of annual actual usage by discipline

SYSTEM AVAILABILITY

System availability has significant impact on the productivity of Blue Waters' partners. For this reason, we use criteria that are more stringent than typical. For Blue Waters a service interruption is any event or failure (hardware, software, human, or environment) that disrupts the system's services level for a specified time period. Most of the outages we report are for failover events that take too long, where applications continue to process for the most part but the system is not completely serviceable. We call these events unscheduled if partners have less than 24 hours' notice of an interruption; our target is to provide at least seven days' notice. The duration of an outage is calculated as the time during which full functionality of the system is unavailable to users, from first discovery to full return to service. Partial or degraded service is counted as part of the outage, or is listed as degraded service.

Blue Waters met or exceeded all but one of its stringent system availability metrics for the year (Table 1). For example, Blue Waters significantly exceeded its target of 5 days for mean time between systemwide failure (MTBF, computed by dividing the number of hours the system was available in a month by the number of system-wide interrupts (or 1 if there are no interrupts) and then converting to days) with a MTBF of 11.2 days. The measured MTBF was above the target 10 of the 11 months in the reporting period (Figure 2), which is remarkable for a system that is 45 percent larger than any other system Cray has delivered. The system has been remarkably trouble free, with two months

METRIC	TARGET	PY2 7/1/2014-5/31/2015
Scheduled Available Computing Time	≥ 92%	91.0% BELOW EXPECTATIONS
Mean Time Between System Wide Failure	≥ 5 day	11.2 days SIGNIFICANTLY EXCEEDS EXPECTATIONS
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	96% 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% of partner service requests were resolved within three business days MEETS EXPECTATIONS
The majority of computational time goes to capability jobs	40% of all computational time will go to capability jobs	75.1% EXCEEDS EXPECTATIONS
Capability jobs are provided excellent turnaround	Expansion factor is less than 10	6.6 EXCEEDS EXPECTATIONS

TABLE 1: System availability metrics

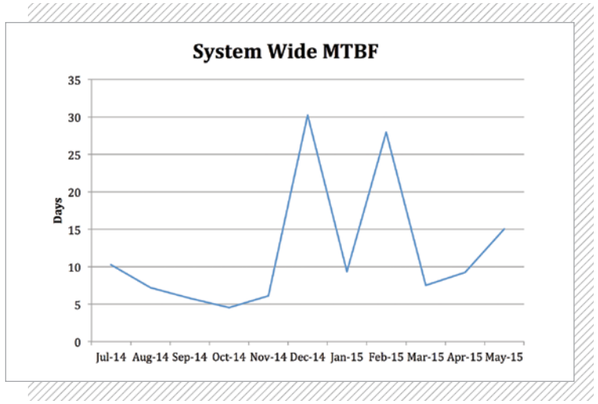

FIGURE 2: System wide mean time between failure for the entire Blue Waters system

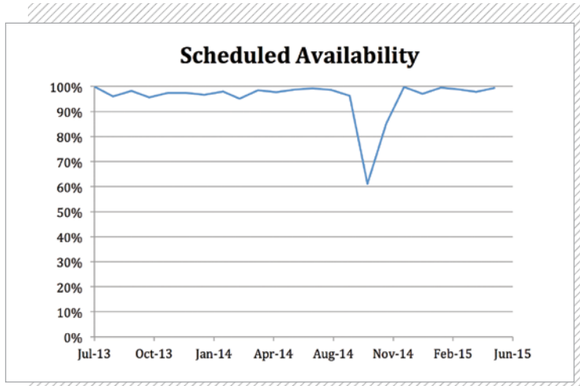

FIGURE 3: Overall Scheduled System Availability

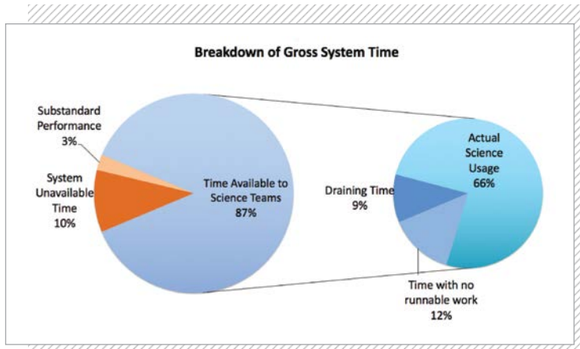

FIGURE 4: Breakdown of overall gross system time during the second year of operations. The blue shaded segments are the time available for science use.

having no systemwide outages at all (Actually at the end of writing this report, Blue Waters had been up for 98 consecutive days without a scheduled or unscheduled interruption!).

Scheduled availability (Figure 3) gives an indication of how reliable the system is for science partners during the time the system is planned to be available; here Blue Waters narrowly missed its target of 92% availability (posting 91%). Two notable outages occurred in October and November 2014 respectively; one was related to a facility wide 13.8KVA to 480V power infrastructure problem over a weekend that required shipping in a unique replacement part. The other was not an interruption per se but a severe regression in metadata performance that limited partner workloads that use the file system. The second issue was discovered while in service and all partner usage during the time of the degradation was not charged. Reducing that unscheduled outage remains a key focus of the NCSA and Cray teams.

Overall 87 percent of the total time in the year was made available to science teams (Figure 4). Sixty-six percent of the total time was actually used by science teams. The draining time refers to periods where nodes were idle, waiting on additional resources in order to start a large job. Draining time can be used for backfill jobs if workload of an appropriate size and requested duration is available. Time with no runnable work refers to periods where nodes are idle and have no future schedule jobs. There may still be queued workloads, but those workloads are not startable due to job dependencies or they require other specific resources.

PROCESSOR USE

Blue Waters is designed and allocated to enable partners to tackle the large, challenging computational jobs that: 1) cannot be performed in a feasible time period elsewhere; and 2) explore novel territory in their disciplines, forging a path for other researchers to follow. As shown in Table 1, 75 percent of the Blue Waters workload is in the capability job category, that is jobs running at a scale that is difficult or impossible to achieve on most other resources (greater than 8,000 integer cores).

Figure 5 presents another view of the usage per job size, where jobs have been categorized as Not Large (those using less than 1 percent of the system), Large (using from 2 to 20 percent of the system), and Very Large (using more than 20 percent of the system). The Large and Very Large categories make up the capability job classification. Note that even a Not Large job on Blue Waters can be larger than a full-system job on many other systems since they can be up to 16,352 integer cores. Very few systems in the world have more than 20 percent of the number of nodes on Blue Waters, and most of them have slower processors.

Looking at capability performance in another way, this year two science and engineering teams used more than 136 million node-hours on Blue Waters (more than 4.3 billion integer core equivalent hours).

The job size median corresponding to 50 percent of the actual usage is 1,250 nodes (40,000 integer cores) for the XE portion and 360 nodes for the XK portion of the Blue Waters system, marked using horizontal lines in figures 6 and 7. Note that the horizontal scale on both of these figures is logarithmic. The median is modestly smaller than last year, which is expected since some of the smaller allocation categories (e.g. Training accounts) are more numerous. Overall the XK nodes delivered 15 percent of the node-hours, which is about their relative fraction of the overall compute node count.

Blue Waters is very responsive and provides exceptional turnaround to the teams for all job sizes. Expansion factor is an indication of the responsiveness of the resource management system to work submitted by partners. Expansion factor is defined as the time an eligible job waits in the queue plus the requested wall time divided by the requested wall time of the job. On many systems, larger jobs are typically more difficult for the system to schedule. However, on Blue Waters the emphasis is on ensuring that partners can take full advantage of the size and unique capabilities of the system. For example, the scheduler has been configured to prioritize Large and Very Large jobs, thus making it easy for the partners to run their applications at scale. Not Large jobs, particularly long running jobs, may wait in the queue for about the requested wall time, on average, independent of the node type. Large jobs take about one to three times the requested wall time to start running, with jobs on XK nodes starting sooner than those on XE

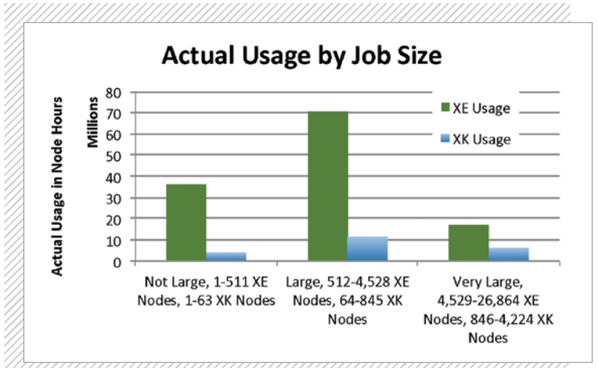

FIGURE 5: Usage per job size category in terms of absolute millions of actual node-hours

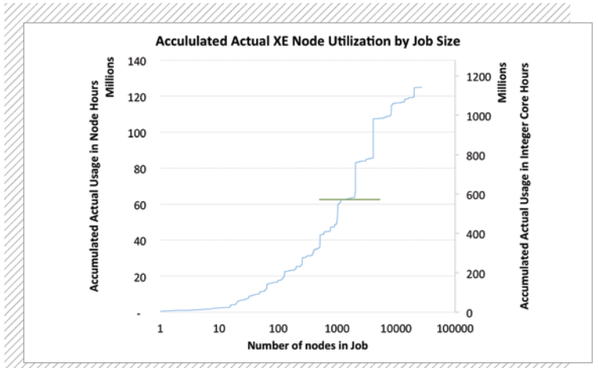

FIGURE 6: Accumulated usage for the XE nodes per job size. The key factor is half the time used on Blue Waters was for jobs larger than 40,000 integer cores.

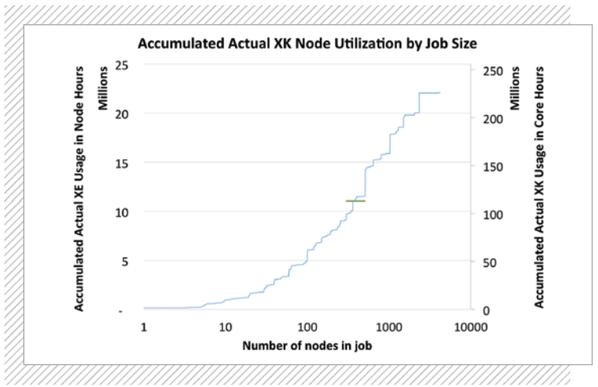

FIGURE 7: Accumulated usage for the XK nodes per job size. The key factor is half the time used on Blue Waters was for jobs larger than 360 GPUs.

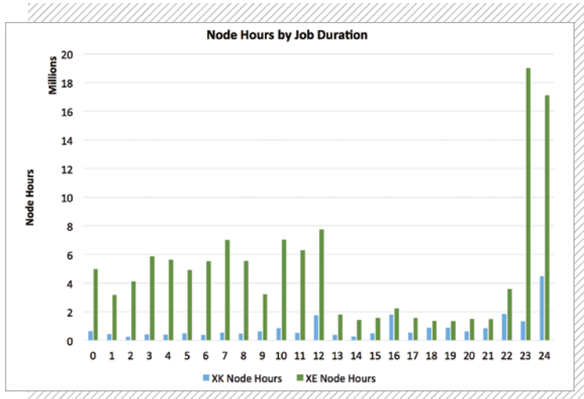

FIGURE 8: Distribution of actual usage by job duration (data from July 2014-May 2015)

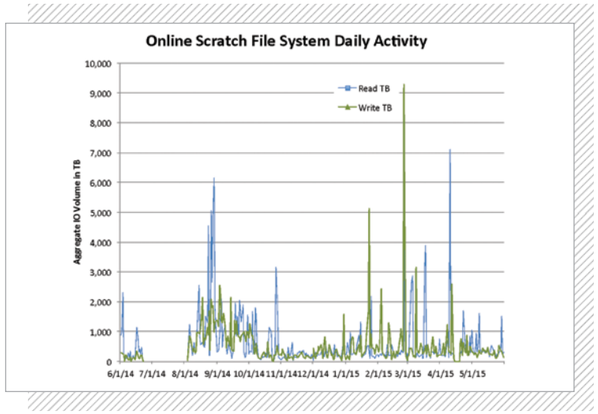

FIGURE 9: Scratch File System Daily Activity

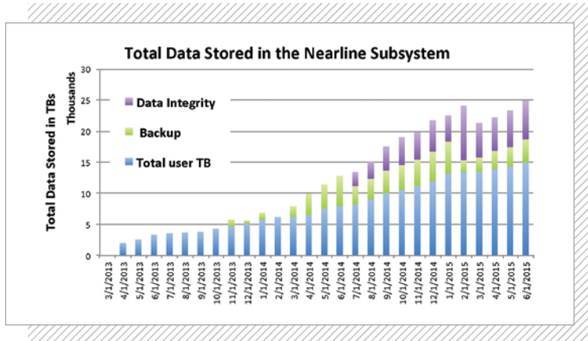

FIGURE 10: Data storage growth in the Blue Waters near-line storage. The dip in March 2015 was the deletion of older backup data that was no longer needed.

nodes. Very Large jobs wait in the queue for four to nine times the requested wall time while Blue Waters collects the massive resources required.

As might be expected, the most common run time is the current queue maximum of 24 hours (Figure 8). XE jobs have a larger distribution of run times, likely due in part to their much larger node counts. Long run times are generally beneficial to partners since this reduces the overhead cost of job startup and teardown.

STORAGE USE

The Blue Waters system has three file systems—/home, /project, and /scratch—totaling 35 raw petabytes (~26 PB usable). Home directories default to 1 TB, and project directories default to 5 TB for each project. Both are managed with user/group quotas and backed up for disaster recovery and neither is purged. Partners can request more space in their project directory as needed. Additionally the partners have access to the world's largest near-line tape storage system, with a capacity of 380 PB.

The /scratch file system consists of 21 PB of usable space. Files are purged if they have not been accessed within 30 days. This allows for very large quotas for many teams because not all teams are active all the time. The default quota for the scratch directory is 500 TB per project which is more than the total storage on many typical clusters; many teams are granted increases to the default limits for specified time periods by request.

As shown in Figure 9, there were multiple days when more than 5 PB of data was written or read to /scratch, with more than 9 PB written on the peak day for the year. Writing 9 PB of data in a day would be impossible on most other systems. The /project and /home file systems show similar variability, with much lower activity levels on the order of tens of terabytes in the home directories and hundreds of terabytes in the project directories, as expected.

At the time of this writing, there are 76 partners in 52 projects actively storing data in the Blue Waters near-line tape subsystem for a total of 25 used PB of data and more than 123 million files (Figure 10). Two very large projects have stored more than 1.2 PB each.

NETWORKING

The Blue Waters system is a very large data source and data repository. That means that a significant amount of data is transferred to and from Blue Waters each month. Figure 11 shows the aggregate monthly traffic to/from Blue Waters to external sources. This traffic has grown more slowly in this project year due in part to nearing the saturation point for some of the existing external 10Gb network links. During the next project year, those links will be upgraded to 100Gbps speeds to enable larger data transfers and higher flow rates.

A high-profile system such as Blue Waters with such excellent network connectivity is an attractive target for cyber-intruders, but the Blue Waters cybersecurity team has effectively detected and blocked these attempts and there have been no security incidents involving Blue Waters. There were approximately 150,000 scans or other malicious connection attempts per month, which were automatically detected and the attacking host blocked from further communication with Blue Waters.

SERVICE REQUESTS

Blue Waters staff focus on resolving partners' issues effectively and quickly to ensure they are able to maintain high productivity. As shown in Table 1, Blue Waters met or exceeded goals for responsiveness to user requests. Eighty percent of all service requests were resolved in less than three business days. Some requests will always take longer than three days to resolve and can be very labor intensive, such as requests for additional software or for help with code optimization and re-engineering; the average time to resolution for these more time-consuming requests was still a very reasonable eight business days.

Figure 12 shows the relative frequencies of different types of service requests, including requests for advanced support, expanded assistance, added services, and suggestions. Of the 1,435 total requests, the two largest categories were Accounts & Allocations and Applications, which each made up slightly more than a quarter of the requests.

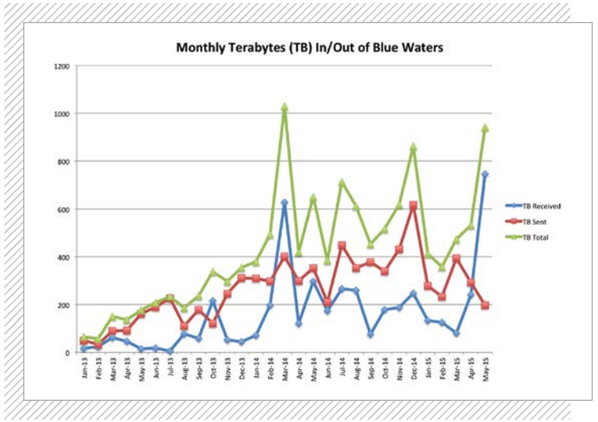

FIGURE 11: Monthly Aggregated Network Utilization.

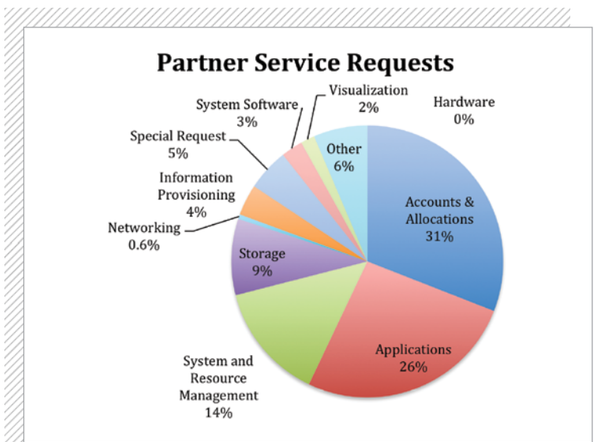

FIGURE 12: Breakdown of partner service requests by type

CONCLUSION

Thanks to the dedicated efforts of the skilled Blue Waters team, the quality of service provided by the system is exceeding expectations. By measuring our performance against stringent metrics, we are able to assess and improve our performance to ensure that the hundreds of people across the country who use Blue Waters are able to maintain high levels of productivity and to achieve breakthrough science and engineering results.

BLUE WATERS PROFESSORS

The University of Illinois at Urbana-Campaign is not just the home of Blue Waters—it's also a powerhouse of computational research in a diverse range of fields. In an effort to attract and support more top-notch researchers, the University offers Blue Waters Professorships, which provide substantial allocations on Blue Waters.

CURRENT BLUE WATERS PROFESSORS

Aleksei Aksimentiev, associate professor, Physics, has done pioneering work on nanoscale biological and synthetic motors, transmembrane transport, and DNA replication that has resulted in the development of new computational strategies for modeling nano-bio systems and technological breakthroughs in DNA sequencing. He is using Blue Waters to determine the molecular mechanism of a key step in the DNA repair process.

Daniel Bodony, associate professor, Aerospace Engineering, is a lead principal investigator for the DOE Center for Exascale Simulation of Plasma-Coupled Combustion, which will develop a multiphysics computational fluid dynamics code and software tools that use exascale-class computers. He uses Blue Waters to study the noise pollution caused by commercial and military aircraft in order to improve near-airport community conditions and reduce hearing damage and its healthcare costs.

David Ceperley, Founder Professor of Physics, is one of the leading computational condensed matter theorists in the world; he is recognized for his pioneering work on developing computational methods applicable for treating many-body fermion and boson systems. His current research is related to the Materials Genome Initiative, the federally supported program to develop computational tools to design materials.

Bryan Clark, assistant professor, Physics, focuses on the application and development of computational methods to gain greater insights into many-body and strongly correlated physics. He is the co-author of the open-source Path Integral Monte Carlo Code (PIMC++), which is widely used to study finite-temperature systems at the quantum level. He joined the University in January 2014 after post-doctoral positions at the Kavli Institute for Theoretical Physics, Microsoft Research, and the Princeton Center for Theoretical Science.

Larry Di Girolamo, professor and Daniel Shapiro Professorial Scholar in Atmospheric Sciences, has a long history of successfully using high-end computing resources for both computational modeling and processing of large satellite data sets. He uses Blue Waters and an advanced Monte Carlo radiative transfer model that he has developed and coupled to a weather prediction model to improve weather predictions and satellite retrievals of cloud properties.

Paul Fischer, professor, Computer Science and Mechanical Science and Engineering, is well known for his research on numerical methods and software for computational fluid dynamics (CFD). His contributions have improved the accuracy and speed of complex CFD simulations. At Argonne National Laboratory he developed nek5000, a widely used CFD software package that can reliably scale to over one million processes, making it the most scalable software of its kind. He joined the University in July 2014 as a professor of Computer Science and Mechanical Science and Engineering.

William Gropp, Thomas M. Siebel Chair in Computer Science, is known for his research in parallel computing, software for scientific computing, and numerical methods for partial differential equations. He uses Blue Waters for research on the effective use of extreme-scale systems, which may require new algorithms and programming systems.

Kaiyu Guan focuses on Natural Resources and Environmental Science. He is a computational environmental scientist whose current research emphasizes the impacts of hydro-climate change on vegetation dynamics and crop productivity in tropical and temperate regions. The goal of this research is to increase our society's resilience and adaptability to sustain ecosystem services, food security, and water resources under the influence of climate change.

Sharon Hammes-Schiffer, Swanlund Professor of Chemistry, focuses on the development and application of theoretical and computational methods for describing chemical reactions in condensed phases and at interfaces. The objective of her research is to elucidate the fundamental physical principles underlying charge transfer reactions. She has made important contributions to theories of proton-coupled electron transfer reactions, which play vital roles in a broad range of biological and chemical processes, including photosynthesis, respiration, enzyme reactions, and solar cells.

So Hirata, Alumni Research Scholar Professor of Chemistry, uses quantitative theories and computing technology to interpret and predict the properties and transformations of molecules, polymers, and solids. Using Blue Waters, he intends to predict a variety of properties of all molecular phases of ice and dry ice from first principles with expected impact on fields including high-pressure chemistry, geochemistry, and planetary science. He will also apply his newly developed, highly scalable stochastic algorithm of quantum chemistry to directly evaluate optoelectronic properties of molecules used in advanced organic materials, such as solar cells and light emitting diodes.

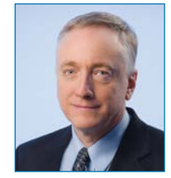

Athol Kemball, associate professor, Astronomy, uses Blue Waters to address problems in radio-interferometric image formation that are relevant to future large astronomical telescopes, such as the Square Kilometer Array. His team has developed a software framework implementing a novel solution to the problem of pixel-level uncertainty quantification and will deploy and optimize this framework for Blue Waters.

Sonia Lasher-Trapp, professor, Atmospheric Sciences. Lasher-Trapp plans to use Blue Waters to simulate entrainment, the process by which cumulus clouds bring dry air from outside the cloud inward. Entrainment initially favors precipitation formation, but eventually it dilutes the cloud and encourages its demise. A long-standing problem in meteorology has been to reproduce how quickly entrainment dilutes the cloud.

Vincent Le Chenadec, assistant professor, Aerospace Engineering. Le Chenadec develops and uses novel numerical methods that leverage computational geometry to solve a long-standing problem in simulating and modeling turbulent reacting high-density-ratio liquid-gas systems. These developments make it possible to investigate the fundamental fluid dynamics of liquid jets injected into gas environments, which could improve the design of devices for fuel-air mixing and thereby improve engine efficiency.

Andre Schleife, assistant professor, Materials Science and Engineering, uses accurate and predictive computational first-principles methods to study ground states, excited electronic states, and their dynamics in various materials, with the ultimate goal of gaining the knowledge necessary to design new and better materials. He joined the University in December 2013 after a post-doctoral position at Lawrence Livermore National Laboratory.

Klaus Schulten, Swanlund Professor of Physics, is world-renowned for developing and implementing powerful computational techniques for modeling and analyzing biomolecular structures that arise in living cells. He heads the NIH Center for Macromolecular Modeling & Bioinformatics and is co-director of the NSF Center for the Physics of Living Cells. He uses Blue Waters to simulate the physical and quantum physical properties of large-scale molecular assemblies, including recent work on the chemical structure of the HIV capsid, on the development of antiviral and antimicrobial drugs, as well as on the use of solar energy in photosynthetic organisms.

Diwakar Shukla, assistant professor, Department of Chemical & Biomolecular Engineering, focuses on developing and using novel atomistic simulation approaches for understanding complex biological processes. His work on Blue Waters focuses on plant kinases and phosphatases involved in response to these external stresses are investigated to elucidate the molecular basis of drought-resistance in plants.

Robert Jeffrey Trapp, professor, Atmospheric Sciences. Trapp uses Blue Waters to simulate severe thunderstorms and tornadoes under the weather and climate conditions predicted for the future. This will help address the impact that human-generated climate change could have on hazardous weather.

OTHER BLUE WATERS PROJECTS

Bernholc, Jerzy
Petascale Quantum Simulations of Nano Systems and Biomolecules
Allocation Type: PRAC general

Campbell, Michael
Rocstar Porting and Scalability
Allocation Type: Industrial

Chaudhuri, Santanu
Material Discovery Using Evolutionary Algorithms: Finding the Missing Ruthenium-based Ternary Phases for Breakthrough Applications in Energy Technologies
Allocation Type: Illinois general

Chew, Huck Beng
Mechanics of Deformation in High Capacity Lithium-Ion Batteries
Allocation Type: Illinois exploratory

Cox, Donna
The CADENS Blue Waters Visualization Project
Allocation Type: Illinois general

Gammie, Charles
Resolving Accretion Disk Turbulence
Allocation Type: Illinois general

Garzaran, Maria Jesus
Enabling Large Scale Irregular Computations of Sparse Data
Allocation Type: Illinois exploratory

Georgiadis, John
Real-Time Magnetic Resonance Elastography via GPU-Accelerators
Allocation Type: Illinois exploratory

Gupta, Indranil
Should Datacenters Converge and Supercomputers Converge?
Allocation Type: Illinois general

Le Chenadec, Vincent
Petascale Simulation of Turbulent Two-Phase Flows
Allocation Type: Illinois Blue Waters Professor

Leff, Laurence
Genetic Algorithms and Participatory Democracy
Allocation Type: Illinois exploratory

Nagamine, Kentaro
Peta-Cosmology: Galaxy Formation and Virtual Astronomy
Allocation Type: PRAC general

Rice, Gus
High Density Mechanism Design
Allocation Type: Industrial

Ricker, Paul
Effects of Active Galaxy Feedback on the Evolution of Galaxy Clusters
Allocation Type: Illinois exploratory

Sharma, Ashish
Very High-Resolution Numerical Modeling for Climate Extremes in Midwest U.S.
Allocation Type: GLCPC general

Snir, Marc
Parallel Programming Language and Library Research on Multicore Clusters
Allocation Type: Illinois exploratory

Sutton, Brad
Improving the Resolution of Brain Blood Flow Imaging with Advanced MRI Acquisitions and Computation
Allocation Type: Illinois exploratory

Takhar, Pawan
Multiscale Modeling of Thermomechanical Changes in Food Biopolymers During Fluid Transport Processes
Allocation Type: Illinois exploratory

Wang, Jue
BAQUS Scalability Study
Allocation Type: Industrial

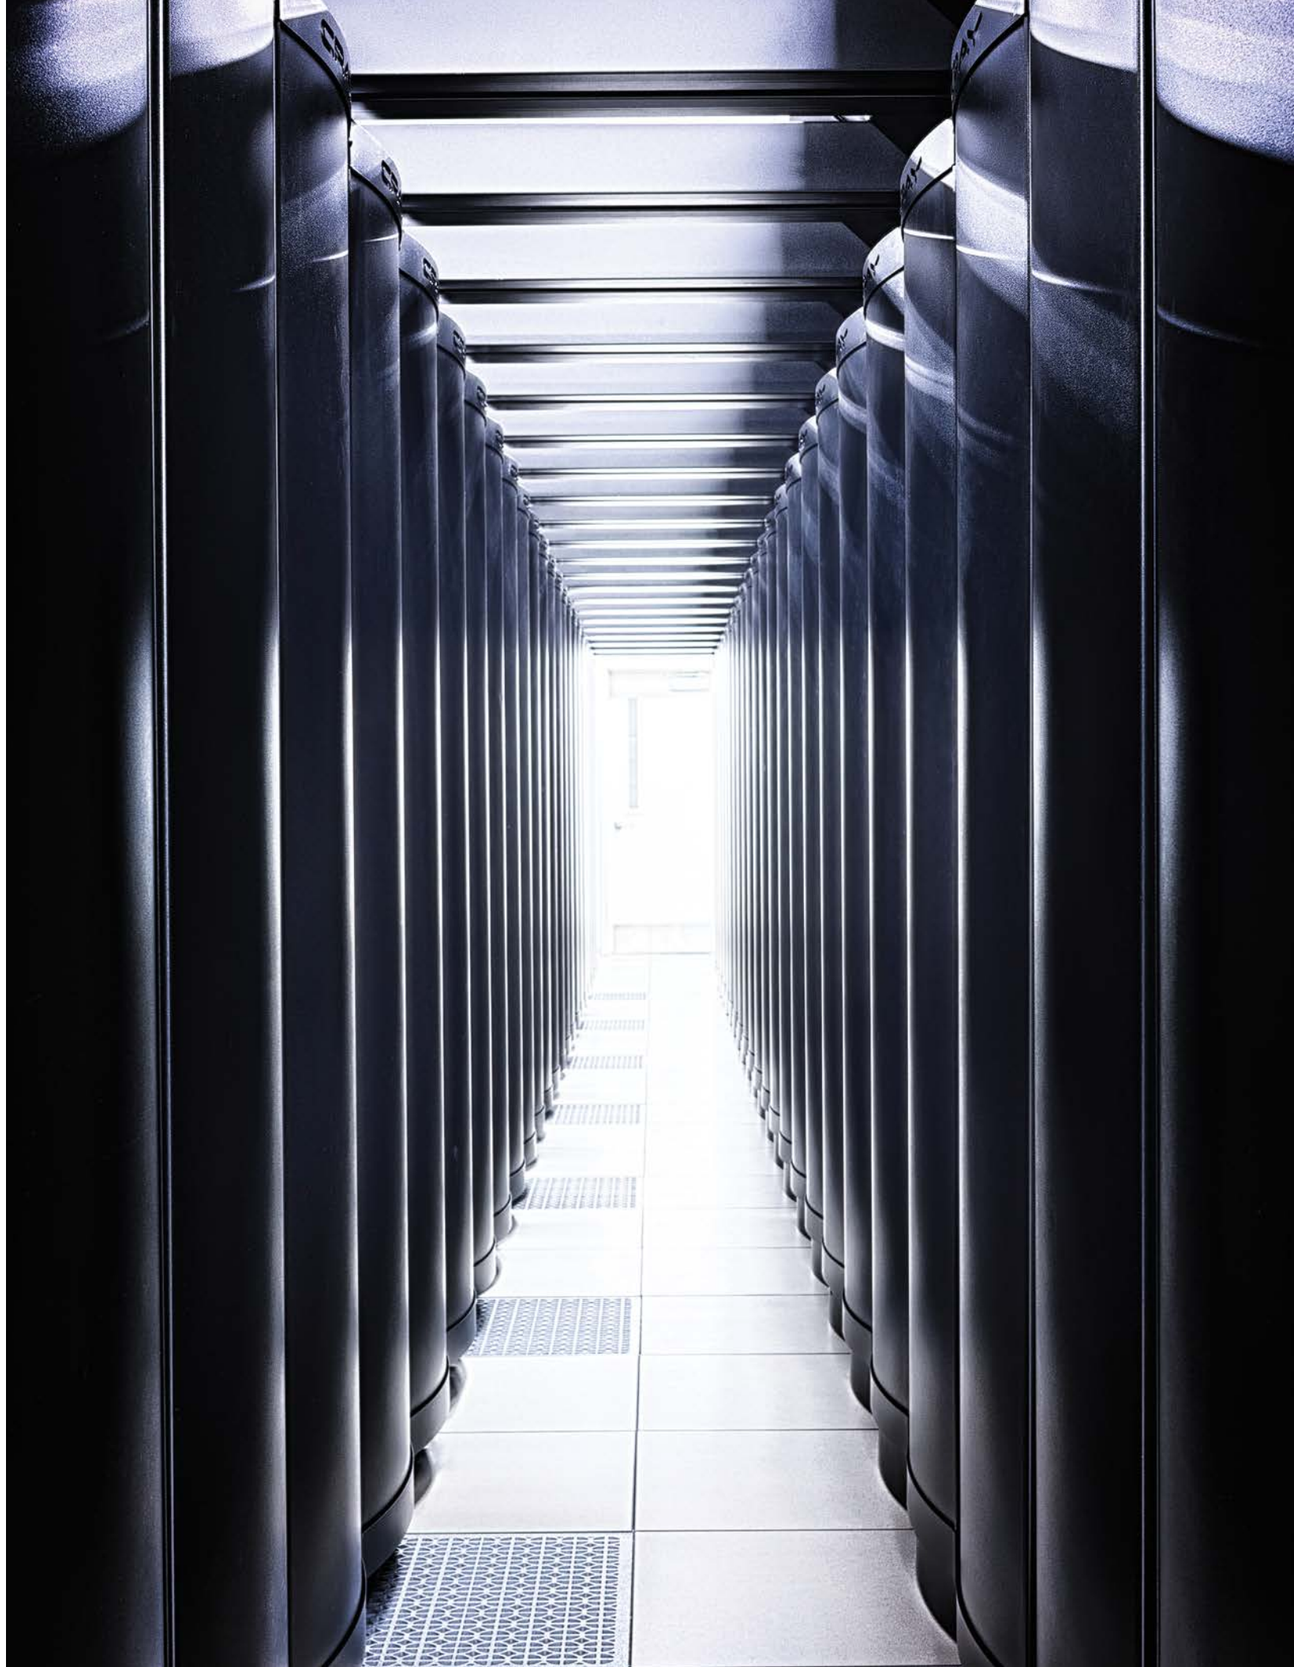

REFERENCES

SPACE SCIENCE

Di Matteo, Tiziana

[1] http://bluetides-project.org/
[2] http://github.com/mpip/pfft
[3] http://github.com/rainwoodman/MP-sort
[4] Feng, Y., M. Straka, R. Croft, and T. Di Matteo, Sorting At Scale on Blue Waters. *Cray User Group Conf. 2015*, Chicago, Ill., April 26–30, 2015. [Best Paper finalist]
[5] Feng, Y., et al., The formation of Milky Way-mass disk galaxies in the first 500 million years of a cold dark matter universe. *Astrophys. J. Lett.*, 808:1 (2015), L17, doi:10.1088/2041-8205/808/1/L17.
[6] Feng, Y., et al., BlueTides: First galaxies and reionization. *Mon. Not. Royal Astron. Soc.*, submitted.
[7] Di Matteo, T., et al., BlueTides: the next frontier of high-z quasars. *Astrophys. J. Lett.*, in prep.
[8] http://bluetides-project.org/viewer/
[9] Hopkins, P. F., A general class of Lagrangian smoothed particle hydrodynamics methods and implications for fluid mixing problems. *Mon. Not. Royal Astron. Soc.*, 428:4 (2013), pp. 2840–2856, doi:10.1093/mnras/sts210.
[10] Battaglia, N., et al., Reionization on Large Scales. I. A Parametric Model Constructed from Radiation-hydrodynamic Simulations. *Astrophys. J.*, 776:2 (2013), 81, doi:10.1088/0004-637X/776/2/81.
[11] Vogelsberger, M., et al., A model for cosmological simulations of galaxy formation physics. *Mon. Not. Royal Astron. Soc.*, 436:4 (2013), pp.3031–3067, doi:10.1093/mnras/stt1789.

Karimabadi, Homayoun

[1] Bowers, K. J., B. J. Albright, L. Yin, B. Bergen, and T. J. T. Kwan, Ultrahigh performance three-dimensional electromagnetic relativistic kinetic plasma simulation. *Phys. Plasmas*, 20 (2013), 061212, doi:10.1063/1.2840133.
[2] Wan, M., et al., Intermittent Dissipation and Heating in 3D Kinetic Plasma Turbulence. *Phys. Rev. Lett.*, 114 (2015), 175002, doi:10.1103/PhysRevLett.114.175002.
[3] Roytershteyn, V., H. Karimabadi, and A. Roberts, Generation of magnetic holes in fully kinetic simulations of collisionless turbulence. *Philos. Trans. A*, 373:2041 (2015), doi:10.1098/rsta.2014.0151.
[4] Jara-Almonte, J., W. Daughton, and H. Ji, Debye Scale Turbulence within the Electron Diffusion Layer during Magnetic Reconnection. *Phys. Plasmas*, 21 (2014), 032114, doi:10.1063/1.4867868.
[5] Guo, F., H. Li, W. Daughton, and Y.-H. Liu, Formation of Hard Power Laws in the Energetic Particle Spectra Resulting from Relativistic Magnetic Reconnection. *Phys. Rev. Lett.*, 113 (2014), 155005, doi:10.1103/PhysRevLett.113.155005.
[6] Liu, Y.-H., J. Birn, W. Daughton, M. Hesse, and K. Schindler, Onset of reconnection in the near magnetotail: PIC simulations. *J. Geophys. Res.*, 119 (2014), pp. 9773–9789, doi:10.1002/2014JA020492.
[7] Makwana, K. D., V. Zhdankin, H. Li, W. Daughton, and F. Cattaneo, Energy dynamics and current sheet structure in fluid and kinetic simulations of decaying magnetohydrodynamic turbulence. *Phys. Plasmas*, 22 (2015), 042902, doi:10.1063/1.4916492.
[8] Guo, F., Y.-H. Liu, W. Daughton, and H. Li, Particle acceleration and plasma dynamics during magnetic reconnection in the magnetically-dominated regime. *Astrophys. J.*, 806 (2015), 167, doi:10.1088/0004-637X/806/2/167.
[9] Karimabadi, H., H. X. Vu, D. Krauss-Varban, and Y. Omelchenko, Global Hybrid Simulations of the Earth's Magnetosphere. in *Numerical Modeling of Space Plasma Flows 2006*, 359 (2006), pp. 257–263.
[10] Karimabadi, H., et al., The link between shocks, turbulence, and magnetic reconnection in collisionless plasmas. *Phys. Plasmas*, 21 (2014), 062308, doi:10.1063/1.4882875.

Makwana, Kiritkumar

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. *Physical Review Letters*, 110 (2013), 091102.
[3] Mignone, A., et al., PLUTO: A Numerical Code for Computational Astrophysics. *The Astrophysical Journal 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] Bryan, G., et al., Enzo: An Adaptive Mesh Refinement Code for Astrophysics. *Astrophys. J.*, 211 (2014), 19, doi:10.1088/0067-0049/211/2/19.
[2] http://enzo-project.org
[3] Turk, M., et al., yt: A Multi-code Analysis Toolkit for Astrophysical Simulation Data. *Astrophys. J. Supp.*, 192 (2011), 9, doi:10.1088/0067-0049/192/1/9.

Pogorelov, Nikolai

[1] Borovikov, S. N., and N. V. Pogorelov, Voyager 1 near the Heliopause. *Astrophys. J.*, 783 (2014), L16, doi:10.1088/2041-8205/783/1/L16.
[2] Heerikhuisen, J., et al., The effect of new interstellar medium parameters on the heliosphere and energetic neutral atoms from the interstellar boundary. *Astrophys. J.*, 784 (2014), 73, doi:10.1088/0004-637X/784/1/73.
[3] 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, doi:10.1002/2014JA020636.
[4] McComas, D. J., et al., Global observations of the interstellar interaction from the Interstellar Boundary Explorer (IBEX). *Science*, 326 (2009),pp. 959–62, doi:10.1126/science.1180906.
[5] Pogorelov, N., S. N. Borovikov, J. Heerikhuisen, and T. K. Kim, Time-dependent processes in the sheath between the heliospheric termination shock and the heliopause in *Numerical Modeling of Space Plasma Flows: ASTRONUM-2013*, 488 (2013),pp. 167–177. [eds. N.V. Pogorelov, E. Audit, G.P. Zank]
[6] Stone, E. C., et al., Voyager 1 explores the termination shock region and the heliosheath beyond. *Science*, 309:5743 (2005), pp. 2017–2020, doi:10.1126/science.1117684.
[7] Stone, E. C., et al., An asymmetric solar wind termination shock. *Nature*, 454 (2008), pp. 71–74, doi:10.1038/nature07022.
[8] Stone, E. C., et al., Voyager 1 Observes Low-Energy Galactic Cosmic Rays in a Region Depleted of Heliospheric Ions. *Science*, 341 (2013), pp. 150–3, doi:10.1126/science.1236408.
[9] Zhang, M., P. Zuo, and N. Pogorelov, Heliospheric influence on the anisotropy of TeV galactic cosmic rays. *Astrophys. J.*, 790 (2014), 5, doi:10.1088/0004-637X/790/1/5.

Shapiro, Stuart

[1] 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, doi:10.1103/PhysRevD.82.084031.
[2] Etienne, Z. B., Y. T. Liu, V. Paschalidis, and S. L. Shapiro, General relativistic simulations of black hole–neutron star mergers: Effects of magnetic fields. *Phys. Rev. D*, 85 (2012), 064029, doi:10.1103/PhysRevD.85.064029.
[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, doi:10.1103/PhysRevD.86.084026.
[4] Beckwith, K., J. F. Hawley, and J. H. Krolik, The Influence of Magnetic Field Geometry on the Evolution of Black Hole Accretion Flows: Similar Disks, Drastically Different Jets. *Astrophys. J.*, 678:2 (2008), pp. 1180–1199, doi:10.1086/533492.
[5] Paschalidis, V., M. Ruiz, and S. L. Shapiro, Relativistic simulations of black hole–neutron star coalescence: the jet emerges. *Astrophys. J.*, 806 (2015), L14, doi:10.1088/2041-8205/806/1/L14.

GEOSCIENCE

Di Girolamo, Larry

[1] Ohring, G., B.A. Wielicki, R. Spencer, B. Emery, and R. Datla, Satellite instrument calibration for measuring global climate change: Report of a workshop, *Bull. Amer. Meteor. Soc.*, 86, (2005), pp. 1303-1313.
[2] Di Girolamo, L., L. Liang, and S. Platnick, A global view of one-dimensional solar radiative transfer through oceanic water clouds, *Geophys. Res. Lett.*, 37, L18809, (2010), doi:10.1029/2010GL044094.
[3] Marshak, A., S. Platnick, T. Várnai, G. Wen, and R. F. Cahalan, Impact of three-dimensional radiative effects on satellite retrievals of cloud droplet sizes, *J. Geophys. Res.*, 111, D09207, (2006), doi:10.1029/2005JD006686.
[4] 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.
[5] Miles, N. L., J. Verlinde, and E. Clothiaux, Cloud droplet size distributions in low-level stratiform clouds, *J. Atmos. Sci.*, 57, (2000), pp. 295–311.
[6] Wood, R., and D. L. Hartmann, Spatial variability of liquid water path in marine low cloud: The importance of mesoscale cellular convection, *J. Clim.*, 19, (2006), 1748–1764.
[7] Bennartz, R., Global assessment of marine boundary layer cloud droplet number concentration from satellite, *J. Geophys. Res.*, 112, D02201, (2007), doi:10.1029/2006JD007547.
[8] Song, X., G.-J. Zhang, and J.-L. F. Li, Evaluation of microphysics parameterization for convective clouds in the NCAR Community Atmosphere Model CAM5, *J. Clim.*, 25, (2012), pp. 8568–8590, doi:10.1175/JCLI-D-11-00563.1.
[9] Gryspeerdt, E., and P. Stier, Regime-based analysis of aerosol-cloud interactions, *Geophys. Res. Lett.*, 39, L21802, (2012), doi:10.1029/2012GL053221.

Garcia, Marcelo

[1] Bulle, H., *Untersuchungen über die Geschiebeableitung bei der Spaltung von Wasserlaufen* (V.D.I Verlag, Berlin, Germany, 1926) [Tech. Rep., in German].

[2] Giosan, L., S. Constantinescu, F. Filip, and B. Deng, Maintenance of large deltas through channelization: Nature vs. humans in the Danube delta. *J. Hydrol.*, 1 (2013), pp. 35–45, doi:10.1016/j.ancene.2013.09.001.

[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, doi:10.1002/wrcr.20139.

[4] Misra, A. K., et al., Proposed river-linking project of India: a boon or bane to nature. *Environmental Geology*, 51:8 (2007), pp. 1361–1376, doi:10.1007/s00254-006-0434-7.

[5] 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. Hydrol. Res.*, 52:6 (2014), pp. 759–774, doi:10.1080/00221686.2014.949883.

[6] Lee, S., et al., Direct numerical simulation of transitional flow in stenosed carotid bifurcation. *J. Biomech.*, 41 (2008), pp. 2551–2561, doi:10.1016/j.jbiomech.2008.03.038.

[7] Niño, Y., and M. H. Garcia, Using Lagrangian particle saltation observations for bedload sediment transport modeling. *Hydrological Processes*, 12 (1998), pp. 1197–1218, doi:10.1002/(SICI)1099-1085(19980630)12:8<1197::AID-HYP612>3.0.CO;2-U.

[8] Escauriaza, C., and F. Sotiropoulos, Lagrangian model of bed-load transport in turbulent junction flows. *J. Fluid Mech.*, 666 (2011), pp. 36–76, doi:10.1017/S0022112010004192.

[9] Sagaut, P., *Large Eddy Simulation for Incompressible Flows* (Springer–Verlag, Germany, 2010).

[10] Fischer, P. F., J. W. Lottes, and S. G. Kerkemeier, *Nek 5000 webpage* (http://nek5000.mcs.anl.gov, 2008).

[11] Deville, M., P. Fischer, and E. Mund, *High-Order Methods for Incompressible Fluid Flow* (Cambridge University Press, Cambridge, England, 2002).

Jordan, Thomas

[1] Withers, K., K. B. Olsen, Z. Shi, and S. Day, High-Complexity Deterministic Q(f) Simulation of the 1994 Northridge Mw 6.7 Earthquake. *SCEC Annual Meeting*, Palm Springs, Calif., September 6–10, 2014.

[2] Restrepo, D., and J. Bielak, Virtual topography: A fictitious domain approach for analyzing free-surface irregularities in large-scale earthquake ground motion simulation. *Int. J. Numer. Meth. Eng.*, 100:7 (2014), pp. 504–533, doi:10.1002/nme.4756.

[3] Roten, D., K. B. Olsen, S. M. Day, Y. Cui, and D. Fah, Expected seismic shaking in Los Angeles reduced by San Andreas fault zone plasticity. *Geophys. Res. Lett.*, 41:8 (2014), pp. 2769–2777, doi: 10.1002/2014GL059411.

[4] NVIDIA Article about Global Impact Award: http://www.nvidia.com/object/global-impact-award.html

[5] UC San Diego Article about NVIDIA Award: http://ucsdnews.ucsd.edu/pressrelease/sdsc_researchers_win_nvidias_2015_global_impact_award

[6] San Diego Supercomputer Center NVIDIA Award: http://www.sdsc.edu/News%20Items/PR20150317_nvidia_award.html

[7] Taborda, R. and J. Bielak, Ground-Motion Simulation and Validation of the 2008 Chino Hills, California, earthquake using different velocity models. *Bull. Seismol. Soc. Am.*, 104:4 (2014), pp. 1876–1898, doi:10.1785/0120130266.

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 Sys.*, 6:4 (2014), pp. 1065–1094, doi:10.1002/2014MS000363.

Um, Junshik

[1] Yurkin, M. A., and A. G.Hoekstra, The discrete-dipole-approximation code ADDA: Capabilities and known limitations. *J. Quant. Spectrosc. Radiat. Transfer*, 112:13 (2011), pp. 2234–2247, doi:10.1016/j.jqsrt.2011.01.031.

[2] Macke, A., J. Mueller, and E. Raschke, Single-scattering properties of atmospheric ice crystals. *J. Atmos. Sci.*, 53:19 (1996), pp. 2813–2825, doi:10.1175/JAM2501.1.

[3] Um, J., et al., Dimensions and aspect ratios of natural ice crystals. *Atmos. Chem. Phys.*, 15 (2015), pp. 3933–3956, doi:10.5194/acp-15-3933-2015.

Wentzcovitch, Renata

[1] Williams, Q., et al., Understanding the Building Blocks of the Planet: the Materials Science of Earth Processes. (2010), [report of NSF workshop, March 2–4, 2009].

[2] *Theoretical and Computational Methods in Mineral Physics: Geophysical Applications*. Wentzcovitch, R., and L. Stixrude, eds. (Mineralogical Society of America, Chantilly, Va., 2010).

[3] Giannozzi P, S. de Gironcoli, P. Pavone, and S. Baroni, *Ab initio* calculation of phonon dispersions in semiconductors. *Phys. Rev. B*, 43 (1991), 7231, doi:10.1103/PhysRevB.43.7231.

[4] Karki, B. B., R. M. Wentzcovitch, S. de Gironcoli, and S. Baroni, First-principles determination of elastic anisotropy and wave velocities of MgO at lower mantle conditions. *Science*, 286 (1999), pp. 1705–7, doi:10.1126/science.286.5445.1705.

[5] Baroni, S., S. de Gironcoli, A. Dal Corso, and P. Giannozzi, Phonons and related crystal properties from density-functional perturbation theory. *Rev. Mod. Phys.*, 73 (2001), 515, doi:10.1103/RevModPhys.73.515.

[6] Wentzcovtich, R.M., Y. Yu, and Z. Wu, Thermodynamic Properties and Phase Relations in Mantle Minerals Investigated by First Principles Quasiharmonic Theory. *Rev. Mineral Geochem.*, 71 (2010), pp. 59–98, doi:10.2138/rmg.2010.71.4.

[7] Wentzcovitch, R. M., B. B. Karki, M. Cococcioni, and S. de Gironcoli, Thermoelastic properties of $MgSiO_3$-perovskite: Insights on the nature of the Earth's lower mantle. *Phys. Rev. Lett.*, 92 (2004), 018501, doi:10.1103/PhysRevLett.92.018501.

[8] Tsuchiya, T., J. Tsuchiya, K. Umemoto, and R. M. Wentzcovitch, Phase transition in $MgSiO_3$ perovskite in the earth's lower mantle. *Earth Planet. Sci. Lett.*, 224 (2004), pp. 241–248, doi:10.1016/j.epsl.2004.05.017.

[9] Umemoto, K., R. M. Wentzcovitch, and P. B. Allen, Dissociation of $MgSiO_3$ in the cores of gas giants and terrestrial exoplanets. *Science*, 311 (2006), pp. 983–986, doi:10.1126/science.1120865.

[10] Hsu, H., P. Blaha, M. Cococcioni, and R. M. Wentzcovitch, Spin-State Crossover and Hyperfine Interactions of Ferric Iron in MgSiO3 Perovskite. *Phys. Rev. Lett.*, 106 (2011), 118501, doi:10.1103/PhysRevLett.106.118501.

[11] Oganov, A. R., and S. Ono, Theoretical and experimental evidence for a post-perovskite phase of MgSiO3 in Earth's D" layer. *Nature*, 430 (2004), pp. 445–448, doi:10.1038/nature02701.

[12] da Silveira, P. R. C., C. R. S. da Silva, and R. M. Wentzcovitch, Metadata management for distributed first principles calculations in VLab – A collaborative cyberinfrastructure for materials computation. *Comp. Phys. Commun.*, 178 (2008), pp. 186–198, doi:10.1016/j.cpc.2007.09.001.

[13] da Silva, C. R. S., et al., Virtual laboratory for planetary materials: System service architecture overview. *Phys. Earth Planet. Inter.*, 163 (2007), pp. 321–332, doi:10.1016/j.pepi.2007.04.018.

[14] da Silveira P. R. C., et al., Virtual Laboratory for Planetary Materials (VLab) – An updated overview of system service architecture. *Proc. TeraGrid '11*, (2011), 33, doi:10.1145/2016741.2016777.http://dl.acm.org/citation.cfm?doid=2016741.2016777

[15] Wentzcovitch, R. M., T. Tsuchiya, and J. Tsuchiya, $MgSiO_3$-post perovskite at D" conditions. *Proc. Nat. Acad. Sci. USA*, 103 (2006), pp. 543–6, doi:10.1073/pnas.0506879103.

[16] Badro, J., et al., Electronic transitions in perovskite: Possible nonconvecting layers in the lower mantle. *Science*, 305 (2004), pp. 383–6, doi:10.1126/science.1098840.

[17] Wu, Z., et al., Anomalous thermodynamic properties in ferropericlase throughout its spin crossover transition. *Phys. Rev. B*, 80 (2009), 014409, doi: 10.1103/PhysRevB.80.014409.

[18] Wu, Z., J. F. Justo, and R. W. Wentzcovitch, Elastic anomalies in a spin-crossover system: Ferropericlase at lower mantle conditions. *Phys. Rev. Lett.*, 110 (2013), 228501, doi:10.1103/PhysRevLett.110.228501.

[19] Xu, W. B., et al., The effect of bulk composition and temperature on mantle seismic structure. *Earth Planet. Sci. Lett.*, 275 (2008), pp. 70–79, doi:10.1016/j.epsl.2008.08.012.

[20] Murakami, M., K. Hirose, K. Kuwamura, N. Sata, and Y. Ohishi, Post-perovskite phase transition in $MgSiO_3$. *Science*, 304 (2004), pp. 855–858, doi:10.1126/science.1095932.

[21] Ni, S. D., et al., Sharp sides to the African superplume. *Science*, 296 (2002), pp. 1850–2, doi:10.1126/science.1070698.

[22] Wallace, D., *Thermodynamics of Crystals*. (John Wiley & Sons, Inc., Hoboken, N.J., 1972).

Wilhelmson, Robert

[1 Orf, L., R. Wilhelmson, and L. Wicker, A Numerical Simulation of a Long-Track EF5 Tornado Embedded Within a Supercell.

94th Am. Meteorol. Soc. Ann. Mtg., February 2–6, 2014.
[2] Bryan, G., and J. Fritsch, A Benchmark Simulation for Moist Nonhydrostatic Numerical Models. *Mon. Wea. Rev.*, 130:12 (2002), pp. 2917–28, doi:10.1175/1520-0493(2002)130<2917:ABSFMN>2.0.CO;2.
[3] The Calumet–El Reno–Piedmont–Guthrie Tornado of May 24, 2011. *Norman, OK National Weather Service*, http://www.srh.noaa.gov/oun/?n=events-20110524-tornado-b2. (Accessed April 21, 2015)
[4] Naylor, J., and M. Gilmore, Convective Initiation in an Idealized Cloud Model Using an Updraft Nudging Technique. *Mon. Wea. Rev.*, 140:11 (2012), pp. 3699–3705, doi:10.1175/MWR-D-12-00163.1.
[5] Morrison, H., G. Thompson, and V. Tatarskii, Impact of Cloud Microphysics on the Development of Trailing Stratiform Precipitation in a Simulated Squall Line: Comparison of One- and Two-Moment Schemes. *Mon. Wea. Rev.*, 137:3 (2009), pp. 991–1007, doi:10.1175/2008MWR2556.1.

Wuebbles, Donald J.
[1] Reed, K. A., et al., Impact of the dynamical core on the direct simulation of tropical cyclones in a high-resolution global model. *Geophys. Res. Lett.*, 42:9 (2015), pp. 3603–3608, doi:10.1002/2015GL063974.
[2] Liang, X.-Z., et al., Regional Climate–Weather Research and Forecasting Model (CWRF). *Bull. Amer. Meteor. Soc.*, 93 (2012), pp. 1363–1387, doi:10.1175/BAMS-D-11-00180.1.
[3] Liang, X.-Z., and F. Zhang, The Cloud–Aerosol–Radiation (CAR) ensemble modeling system. *Atmos. Chem. Phys.*, 13 (2013), pp. 8335–8364, doi:10.5194/acp-13-8335-2013.

PHYSICS & ENGINEERING

Aluru, Narayana
[1] Wells, D. B., M. Belkin, J. Comer, and A. Aksimentiev, Assessing Graphene Nanopores for Sequencing DNA. *Nano. Lett.*, 12 (2012), pp. 4117–4123, doi:10.1021/nl301655d.
[2] Branton, D., et al., The potential and challenges of nanopore sequencing. *Nat. Biotech.*, 26 (2008), pp. 1146–1153, doi:10.1038/nbt.1495.
[3] Farimani, A. B., M. Heiranian, and N. R. Aluru, Electromechanical Signatures for DNA Sequencing through a Mechanosensitive Nanopore. *J. Phys. Chem. Lett.*, 6 (2015), pp. 650–657, doi:10.1021/jz5025417.
[4] Farimani, A. B., K. Min, and N. R. Aluru, DNA Base Detection Using a Single-Layer MoS_2. *ACS Nano.*, 8 (2014), pp. 7914–7922, doi:10.1021/nn5029295.

Balsara, Dinshaw
[1] Balsara D. S., Wave Propagation in Molecular Clouds. *Astrophys. J.*, 465 (1996), 775, doi:10.1086/177462.
[2] 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.jcp.2015.05.020
[3] http://www.nd.edu/~dbalsara/Numerical-PDE-Course
[4] 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), 2197, doi:10.1093/mnras/stt1893.
[5] Burkhart, B., A. Lazarian, C. Meyer, D. S. Balsara, and J. Cho, Alvenic Turbulence Beyond the Ambipolar Diffusion Scale. *Astrophys. J.*, 805:2 (2015), 118, doi:10.1088/0004-637X/805/2/118.
[6] Li, H., and M. Houde, Probing the Turbulence Dissipation Range and Magnetic Field Strengths in Molecular Clouds. *Astrophys. J.*, 677 (2008), 1151, doi:10.1086/529581.
[7] Goldsmith, P.F., and W. D. Langer, Molecular cooling and thermal balance of dense interstellar clouds. *Astrophys. J.*, 222 (1978), 881, doi:10.1086/156206.

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; 1607-1653 (2012).
[2] Morales, M. A., C. Pierleoni, E. Schwegler, and D. M. Ceperley, Evidence for a first-order liquid-liquid transition in high-pressure hydrogen from ab initio simulations, *Proceedings of the National Academy of Sciences* 107, 12799-12803 (2010).
[3] Knudson, M. D. et al., *Science* 348, 1455 (2015).
[4] Zaghoo, M. A. Salamat, and I. Silvera (2015), arXiv:1504.00259.

Clark, Bryan
[1] Emery, V. J., and S. A. Kivelson, Superconductivity in Bad Metals. *Phys. Rev. Lett.*, 74 (1995), 3253, doi:10.1103/PhysRevLett.74.3253.
[2] Pal, A., and D. Huse, Many-body localization phase transition. *Phys. Rev. B*, 82 (2013), 17441, doi:10.1103/PhysRevB.87.134202.
[3] 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.
[4] Hetherington, J. H., and F. D. C. Willard, Two-, Three-, and Four-Atom Exchange Effects in bcc ^{3}He. *Phys. Rev. Lett.*, 35 (1975), 1442, doi:10.1103/PhysRevLett.35.1442.
[5] Motrunich, O. I., and M. P. A. Fisher, *d*-wave correlated critical Bose liquids in two dimensions. *Phys. Rev. B*. 75:23 (2007), 235116, doi:10.1103/PhysRevB.75.235116.

Draayer, Jerry
[1] Dytrych, T., et al., Collective Modes in Light Nuclei from First Principles. *Phys. Rev. Lett.*, 111 (2013), 252501, doi:10.1103/PhysRevLett.111.252501.
[2] 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.
[3] Dreyfuss, A. C., et al., Hoyle state and rotational features in Carbon-12 within a no-core shell-model framework. *Phys. Lett. B*, 727 (2013), 511, doi:10.1016/j.physletb.2013.10.048.
[4] Tobin, G. K., et al., Symplectic No-core Shell-model Approach to Intermediate-mass Nuclei. *Phys. Rev. C*, 89 (2014), 034312, doi:10.1103/PhysRevC.89.034312.
[5] Draayer, J. P., T. Dytrych, and K. D. Sviratcheva, Next-Generation *Ab Initio* Shell Model for Exotic Nuclear Structures. *Modern Phys. Lett. A*, 25 (2010), 1980, doi:10.1142/S0217732310000794.
[6] Draayer, J. P., T. Dytrych, K. D. Launey, and D. Langr, Symmetry-Adapted No-Core Shell Model Applications for Light Nuclei with QCD-Inspired Interactions. *Progr. Part. Nucl. Phys.*, 67 (2012), 516, doi:10.1016/j.ppnp.2012.01.020.
[7] Dytrych, T., et al., Symmetry-Guided *Ab Initio* Approach to Light and Medium-Mass Nuclei. *J. of Phys.: Conf. Ser.*, 403 (2012), 012015, doi:10.1088/1742-6596/403/1/012015.
[8] Draayer, J. P., et al., Symmetry-Adapted *Ab Initio* Open Core Shell Model Theory. *J. of Phys.: Conf. Ser.*, 387 (2012), 012017, doi:10.1088/1742-6596/436/1/012023.
[9] Dytrych, T., et al., First *ab initio* studies of *sd*-shell nuclei with impact on X-ray burst nucleosynthesis. to be submitted.
[10] Dytrych, T., et al., Electron-scattering form factors for ^{6}Li in the ab initio symmetry-guided framework. *Phys. Rev. C*, 91 (2015), 024326, doi:10.1103/PhysRevC.91.024326.

Fischer, Paul
[1] Fischer, P., J. Kruse, J. Mullen, H. Tufo, J. Lottes, and S. Kerkemeier, NEK5000: Open source spectral element CFD solver, http://nek5000.mcs.anl.gov/
[2] Tufo, H. M., and P. F. Fischer, Terascale Spectral Element Algorithms and Implementations. *Proc. of the 1999 ACM/IEEE Conf. on Supercomputing*, (1999), 68, doi:10.1145/331532.331599. [Gordon Bell Prize paper]
[3] Bulle, H., *Untersuchungen über die Geschiebeableitung bei der Spaltung von Wasserläufen*. (VDI Verlag, Berlin, 1926).

Glotzer, Sharon
[1] http://codeblue.umich.edu/hoomd-blue
[2] Anderson J.A. and S. C. Glotzer, The development and expansion of HOOMD-blue through six years of GPU proliferation, arXiv:1308.5587, (2013).

Kemball, Athol
[1] Kemball, A., and Martinsek, A. Bootstrap Resampling as a Tool for Radio Interferometric Imaging Fidelity Assessment, *Astronomical Journal*, 129 (2005), pp. 1760-1775.
[2] Kemball, A., Martinsek, A., Mitra, M., and Chiang, H.-F., Further Evaluation of Bootstrap Resampling as a Tool For Radio-Interferometric Imaging Fidelity Assessment, *Astronomical Journal*, 139 (2010), pp. 252-266.

Klimeck, Gerhard
[1] Steiger, S., et al., NEMO5: A Parallel Multiscale Nanoelectronics Modeling Tool. *IEEE Trans. Nanotechnol.*, 10 (2011), pp. 1464–1474.
[2] Fonseca, J. E., et al., Efficient and realistic device modeling from atomic detail to the nanoscale. *J. Comput. Electron.*, 12:4 (2013), pp. 592–600.

Koric, Seid
[1] Gupta, A., WSMP: Watson sparse matrix package, Part-I: direct solution of symmetric sparse systems. (IBM T. J. Watson Research Center, Yorktown Heights, NY) [Tech. Rep. RC, 21886, 2015]

Lele, Sanjiva
[1] Cook, A. W., Artificial fluid properties for large-eddy simulation of compressible turbulent mixing. *Phys. Fluids*, 19:5 (2007), 055103, doi:10.1063/1.2728937.
[2] Lele, S. K., Compact finite difference schemes with spectral-like resolution. *J. Comput. Phys.*, 103:1 (1992), pp. 16–42, doi:10.1016/0021-9991(92)90324-R.
[3] Kawai, S., and S. K. Lele, Localized artificial diffusivity scheme for discontinuity capturing on curvilinear meshes. *J. Comput. Phys.*, 227:22 (2008), pp. 9498–9526, doi:10.1016/j.jcp.2008.06.034.
[4] McFarland, J., et al., Experimental investigation of the inclined interface Richtmyer–Meshkov instability before and after reshock. *Experiments in Fluids*, 55:1 (2014), pp. 1–14, doi:10.1007/s00348-013-1640-1.
[5] McFarland, J. A., J. A. Greenough, and D. Ranjan, Computational parametric study of a Richtmyer–Meshkov instability for an

inclined interface. *Phys. Rev. E*, 84:2 (2011), 026303, doi:10.1103/PhysRevE.84.026303.

[6] Cook, A. W., et al., Tera-scalable algorithms for variable-density elliptic hydrodynamics with spectral accuracy. Proc. *2005 ACM/IEEE Conf. Supercomput.*, (2005), p. 60, doi:10.1109/SC.2005.70.

Levin, Deborah

[1] Korkut, B., Z. Li, and D. Levin, 3D Simulation of Ion Thruster Plumes Using Octree Adaptive Mesh Refinement. *IEEE Trans. Plasma Sci.*, 43:5 (2015), pp. 1706–1721, doi:10.1109/TPS.2015.2415458.

Ostoja-Starzewski, Martin

[1] Ostoja-Starzewski, M., *Microstructural Randomness and Scaling in Mechanics of Materials*. (Chapman & Hall/CRC, Boca Raton, Fla., 2008).

[2] Nukala, P. K. V. V., S.Zapperi, and S. Simunovic, Statistical properties of fracture in a random spring model. *Phys. Rev. E*, 71 (2005), 066106, doi:10.1103/PhysRevE.71.066106.

[3] Kale, S., M. Ostoja-Starzewski, Elastic-plastic-brittle transitions and avalanches in disordered media. *Phys. Rev. Lett.*, 112 (2014), 045503, doi:10.1103/PhysRevLett.112.045503.

[4] Kale, S., and M. Ostoja-Starzewski, Morphological study of elastic-plastic-brittle transitions in disordered media. *Phys. Rev. E*, 90 (2014), 042405, doi:10.1103/PhysRevE.90.042405.

[5] Chen, Y., and M. Ostoja-Starzewski, MRI-based finite element modeling of head trauma: spherically focusing shear waves, *Acta Mechanica*, 213:1-2 (2010), pp. 155–167, doi:10.1007/s00707-009-0274-0.

[6] Chen, Y., B. Sutton, C. Conway, S. P. Broglio, and M. Ostoja-Starzewski, Brain deformation under mild impact: Magnetic resonance imaging-based assessment and finite element study. *Int. J. Numer. Anal. Model. B*, 3:1 (2010), pp. 20–35.

Ott, Christian

[1] http://www.einsteintoolkit.org

[2] http://www.cactuscode.org

[3] Abdikamalov, E., et al., Neutrino-driven Turbulent Convection and Standing Accretion Shock Instability in Three-dimensional Core-collapse Supernovae. *Astrophys. J.*, 808:1 (2015), 70, doi:10.1088/0004-637X/808/1/70.

Schleife, Andre

[1] Varley, J. B., and A. Schleife, Bethe–Salpeter calculation of optical-absorption spectra of In2O3 and Ga2O3. *Semicond. Sci. Tech.*, 30 (2015), 024010, doi:10.1088/0268-1242/30/2/024010.

Sugar, Robert

[1] Bazavov, A., et al. (The MILC Collaboration), Leptonic decay-constant ratio $f_{K+}/f_{\pi+}$ from lattice QCD with physical light quarks. *Phys. Rev. Lett.*, 110 (2013), 172003.

[2] Bazavov, A., et al. (The Fermilab Lattice and MILC Collaborations), Determination of $|V_{us}|$ from a lattice-QCD calculation of the $K \to \pi \ell \nu$ semileptonic form factor with physical quark masses. *Phys. Rev. Lett.*, 112 (2014), 112001.

[3] Bazavov, A., et al. (The Fermilab Lattice and MILC Collaborations), Charmed and light pseudoscalar meson decay constants from four-flavor lattice QCD with physical light quarks. *Phys. Rev. D*, 90 (2014), 074509.

[4] Gamiz, A., (The Fermilab Lattice and MILC Collaborations), Charmed and strange pseudoscalar meson decay constants from HISQ simulations. *Proc. Sci. Lattice 2013*, 405 (2013).

[5] Primer, T., et al., Kaon and D meson semileptonic form factors from lattice QCD. *Proc. Sci. Lattice 2014*, 374 (2014).

[6] Bazavov, A., et al. (The Fermilab Lattice and MILC Collaborations), Charmed and light pseudoscalar meson decay constants from HISQ simulations. *Proc. Sci. Lattice 2014*, 382 (2014).

[7] Bazavov, A., et al. (The MILC Collaboration), Gradient flow and scale setting on MILC HISQ ensembles. *Phys. Rev. D*, submitted. [arXiv:1503.02769]

[8] Bazavov, A., (The MILC Collaboration), Gradient Flow Analysis on MILC HISQ Ensembles. *Proc. Sci. Lattice 2014*, 090 (2014).

[9] Winter, F. T., M. A. Clark, R. G. Edwards, and B. Joo, A Framework for Lattice QCD Calculations on GPUs. *Proc. 2014 IEEE 28th Int. Parallel Dist. Proc. Symp.*, pp. 1073–1082, doi:10.1109/IPDPS.2014.112.

Thomas, Brian

[1] *Steel Statistical Yearbook 2014*. (World Steel Association, Brussels, 2014).

[2] Chaudhary, R., B. G. Thomas, and S. P. Vanka, Effect of Electromagnetic Ruler Braking (EMBr) on Transient Turbulent Flow in Continuous Slab Casting using Large Eddy Simulations. *Metall. Mater. Trans. B*, 43B:3 (2012), pp. 532–553, doi:10.1007/s11663-012-9634-6.

[3] Wang, Z. Y., and A. Y. Tong, Deformation and oscillations of a single gas bubble rising in a narrow vertical tube. *ASME 4th Int. Conf. on Nanochannels, Microchannels, and Minichannels*, Limerick, Ireland, June 19–21, 2006.

[4] Jin K., P. Kumar, S. P. Vanka, and B. G. Thomas, Rise of an argon bubble in liquid steel in the presence of a transverse magnetic field. *Phys. Fluids*, (submitted).

[5] Jin K., S. P. Vanka, and B. G. Thomas, Three-dimensional flow in a driven cavity subjected to an external magnetic field. *J. Fluids Engineer.*, 137 (2015), 071104, doi:10.1115/1.4029731.

Woosley, Stan

[1] Nonaka, A., et al., MAESTRO: An Adaptive Low Mach Number Hydrodynamics Algorithm for Stellar Flows. *Astrophys. J. Supp.*, 188 (2010), pp. 358–83, doi:10.1088/0067-0049/188/2/358.

[2] Almgren, A. S., et al., CASTRO: A New Compressible Astrophysical Solver. I. Hydrodynamics and Self-Gravity. *Astrophys. J.*, 715 (2010), pp. 1221–38, doi:10.1088/0004-637X/715/2/1221.

Yeung, Pui-kuen

[1] Frisch, U., *Turbulence. The Legacy of A.N. Kolmogorov*. (Cambridge University Press, Cambridge, U.K., 1995).

[2] Gotoh, T., and P. K. Yeung, Passive scalar transport: a computational perspective. in *Ten Chapters in Turbulence* (2013), pp. 87–131, doi:10.1017/CBO9781139032810.004.

[3] Sawford, B. L., P. K. Yeung, and J. F. Hackl, Reynolds number dependence of relative dispersion statistics in isotropic turbulence. *Phys. Fluids*, 29 (2008), 065111, doi:10.1063/1.2946442.

[4] Yeung, P. K., D. A. Donzis, and K. R. Sreenivasan, Dissipation, enstrophy, and pressure statistics in turbulence simulations at high Reynolds numbers. *J.Fluid Mech.*, 700 (2012), pp. 5–15, doi:10.1017/jfm.2012.5.

COMPUTER SCIENCE & ENGINEERING

Taha, Ahmed

[1] Taha, A. A., Fluent Scaling Breakthrough: New Record Towards Accelerated Product Development, *Ansys Convergence Midwest Conference*, Chicago, Ill., May 23, 2014.

BIOLOGY, CHEMISTRY, & HEALTH

Amaro, Rommie

[1] Lipsitch, M., and A. P. Galvani, Ethical Alternatives to Experiments with Novel Potential Pandemic Pathogens. *PLoS Med.*, 11:5 (2014), e1001646, doi:10.1371/journal.pmed.1001646.

[2] Gamblin, S. J., and J. J. Skehel, Influenza Hemagglutinin and Neuraminidase Membrane Glycoproteins. *J. Biol. Chem.*, 285:37 (2010), pp. 28403–28409, doi:10.1074/jbc.R110.129809.

[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, doi:10.1016/S0166-3542(03)00011-1.

[4] Baz, M., Y. Abed, and G. Boivin, Characterization of drug-resistant recombinant influenza A/H1N1 viruses selected in vitro with peramivir and zanamivir. *Antiviral Res.*, 74:2 (2007), pp. 159–162, doi:10.1016/j.antiviral.2006.10.012.

[5] Ferraris, O., N. Kessler, and B. Lina, Sensitivity of influenza viruses to zanamivir and oseltamivir: a study performed on viruses circulating in France prior to the introduction of neuraminidase inhibitors in clinical practice. *Antiviral Res.*, 68:1 (2005), pp. 43–48, doi:10.1016/j.antiviral.2005.07.004.

[6] Memoli, M. J., et al., Rapid selection of oseltamivir- and peramivir-resistant pandemic H1N1 virus during therapy in 2 immunocompromised hosts. *Clin. Infect. Dis.*, 50:9 (2010), pp. 1252–1255, doi:10.1086/651605.

[7] Gubareva, L. V., et al., Evidence for zanamivir resistance in an immunocompromised child infected with influenza B virus. *J. Infect. Dis.*, 178:5 (1998), pp. 1257–1262, doi:10.1086/314440.

[8] Oakley, A. J., et al., Structural and functional basis of resistance to neuraminidase inhibitors of influenza B viruses. *J. Med. Chem.*, 53:17 (2010), pp. 6421–6431, doi:10.1021/jm100621s.

[9] de Jong, M. D., et al., Oseltamivir resistance during treatment of influenza A (H5N1) infection. *N. Engl. J. Med.*, 353:25 (2005), pp. 2667–2672, doi:10.1056/NEJMoa054512.

[10] Moscona, A., Oseltamivir resistance—disabling our influenza defenses. *N. Engl. J. Med.*, 353:25 (2005), pp. 2633–2636, doi:10.1056/NEJMp058291.

[11] Le, Q. M., et al., Avian flu: isolation of drug-resistant H5N1 virus. *Nature*, 437:7062 (2005), p. 1108, doi:10.1038/4371108a.

[12] de Jong, M. D., et al., Brief report - Oseltamivir resistance during treatment of influenza A (H5N1) infection. *N. Engl. J. Med.*, 353:25 (2005), pp. 2667–2672, doi:10.1056/NEJMoa054512.

[13] 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), e37095, doi:10.1371/journal.pone.0037095.

[14] Gubareva, L. V., Molecular mechanisms of influenza virus resistance to neuraminidase inhibitors. *Virus Res.*, 103:1-2 (2004), pp. 199–203, doi:10.1016/j.virusres.2004.02.034.

[15] Sorrell, E. M., et al., Minimal molecular constraints for respiratory droplet transmission of an avian-human H9N2 influenza A virus. *Proc. Natl. Acad. Sci. U.S.A.*, 106:18 (2009), pp. 7565–7570, doi:10.1073/pnas.0900877106.

[16] Imai, M., et al., Experimental adaptation of an influenza H5 HA confers respira-

tory droplet transmission to a reassortant H5 HA/H1N1 virus in ferrets. *Nature*, 486:7403 (2012), pp. 420–428, doi:10.1038/nature10831.

[17] Herfst, S., et al., Airborne Transmission of Influenza A/H5N1 Virus Between Ferrets. *Science*, 336:6088 (2012), pp. 1534–1541, doi:10.1126/science.1213362 .

[18] Cline, T. D., et al., Increased Pathogenicity of a Reassortant 2009 Pandemic H1N1 Influenza Virus Containing an H5N1 Hemagglutinin. *J. Virology*, 85:23 (2011), pp. 12262–12270, doi:10.1128/JVI.05582-11.

[19] Sun, Y. P., et al., High genetic compatibility and increased pathogenicity of reassortants derived from avian H9N2 and pandemic H1N1/2009 influenza viruses. *Proc. Natl. Acad. Sci. U.S.A.,* 108:10 (2011), pp. 4164–4169, doi:10.1073/pnas.1019109108.

[20] Kimble, J. B., et al., Compatibility of H9N2 avian influenza surface genes and 2009 pandemic H1N1 internal genes for transmission in the ferret model. *Proc. Natl. Acad. Sci. U.S.A.*, 108:29 (2011), pp. 12084–12088, doi:10.1073/pnas.1108058108.

[21] Qiao, C. L., et al., Pathogenicity and transmissibility of reassortant H9 influenza viruses with genes from pandemic H1N1 virus. *J. Gen. Virol.,* 93 (2012), pp. 2337–2345, doi:10.1099/vir.0.044040-0.

[22] Shelton, H., et al., Mutations in haemagglutinin that affect receptor binding and pH stability increase replication of a PR8 influenza virus with H5 HA in the upper respiratory tract of ferrets and may contribute to transmissibility. *J. Gen. Virol.*, 94 (2013), pp. 1220–1229, doi:10.1099/vir.0.050526-0.

[23] Zhang, Y., et al., H5N1 Hybrid Viruses Bearing 2009/H1N1 Virus Genes Transmit in Guinea Pigs by Respiratory Droplet. *Science*, 340:6139 (2013), pp. 1459–1463, doi:10.1126/science.1229455.

[24] Watanabe, T., et al., Circulating Avian Influenza Viruses Closely Related to the 1918 Virus Have Pandemic Potential. *Cell Host Microbe*, 15:6 (2014), pp. 692–705, doi:10.1016/j.chom.2014.05.006.

[25] Sutton, T. C., et al., Airborne Transmission of Highly Pathogenic H7N1 Influenza Virus in Ferrets. *J. Virology*, 88:12 (2014), pp. 6623–6635, doi:10.1128/JVI.02765-13.

[26] Webster, R. G., et al., Evolution and Ecology of Influenza-A Viruses. *Microbiol. Rev.*, 56:1 (1992), pp. 152–179.

[27] Ennis, F. A., Influenza-a Viruses: Shaking out Our Shibboleths. *Nature*, 274:5669 (1978), pp. 309–310, doi:10.1038/274309b0.

[28] Nakajima, K., U. Desselberger, and P. Palese, Recent Human Influenza-a (H1n1) Viruses Are Closely Related Genetically to Strains Isolated in 1950. *Nature*, 274:5669 (1978), pp. 334–339, doi:10.1038/274334a0.

[29] CDC Director Releases After-Action Report on Recent Anthrax Incident; Highlights Steps to Improve Laboratory Quality and Safety. *CDC* (2014).

[30] Harris, A., et al., Influenza virus pleiomorphy characterized by cryoelectron tomography. *Proc. Natl. Acad. Sci. U.S.A.*, 103:50 (2006), pp. 19123–19127, doi:10.1073/pnas.0607614103.

[31] 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), e1003720, doi:10.1371/journal.pcbi.1003720.

Caetano-Anolles, Gustavo

[1] Caetano-Anollés, G., M. Wang, D. Caetano-Anollés, and J. E. Mittenthal, The origin, evolution and structure of the protein world. *Biochem. J.*, 417 (2009), pp. 621–637, doi:10.1042/BJ20082063.

[2] Leszczynski, J. F., and G. D. Rose, Loops in globular proteins: a novel category of secondary structure. *Science*, 234:4778 (1986), pp. 849–855, doi:10.1126/science.3775366.

[3] Bonet, J., et al., ArchDB 2014: structural classification of loops in proteins. *Nucleic Acids Res.*, 42 (2014), pp. D315–D319, doi:10.1093/nar/gkt1189.

[4] Espadaler, J., et al., ArchDB: automated protein loop classification as a tool for structural genomics. *Nucleic Acids Res.*, 32:1 (2004), pp. D185–D188, doi:10.1093/nar/gkh002.

[5] Debès, C., M. Wang, G. Caetano-Anollés, and F. Gräter, Evolutionary optimization of protein folding. *PLoS Comput. Biol.,* 9:1 (2013), e1002861, doi:10.1371/journal.pcbi.1002861.

[6] Caetano-Anollés, G., M. Wang, and D. Caetano-Anollés, Structural phylogenomics retrodicts the origin of the genetic code and uncovers the evolutionary impact of protein flexibility. *PLoS One*, 8:8 (2013), e72225, doi:10.1371/journal.pone.0072225.

[7] Marsh, J. A., and S. A. Teichmann, Parallel dynamics and evolution: Protein conformational fluctuations and assembly reflect evolutionary changes in sequence and structure. *BioEssays*, 36:2 (2014), pp. 209–218, doi:10.1002/bies.20130013.

[8] Caetano-Anollés, G., et al., The dynamics of protein disorder and its evolution: Understanding single molecule FRET experiments of disordered proteins. *Blue Waters Annual Report* (Urbana, Illinois, 2014), pp. 120–121.

Cheatham, Thomas

[1] Galindo-Murillo, R., D. R. Roe, and T. E. Cheatham, 3rd, On the absence of intrahelical DNA dynamics on the μs to ms timescale. *Nature Commun.,* 5 (2014), 5152, doi:10.1038/ncomms6152.

[2] Galindo-Murillo, R., D. R. Roe, and T. E. Cheatham, 3rd, Convergence and reproducibility in molecular dynamics simulations of the DNA duplex d(GCACGAACGAACGAACGC). *Biochim. Biophys. Acta*, 1850:5 (2015), pp. 1041–1058, doi:10.1016/j.bbagen.2014.09.007.

[3] Bergonzo, C., et al., Multidimensional Replica Exchange Molecular Dynamics Yields a Converged Ensemble of an RNA Tetranucleotide. *J. Chem. Theory Comput.*, 2014. 10:1 (2014), pp. 492–499, doi:10.1021/ct400862k.

[4] Roe, D. R., C. Bergonzo, and T. E. Cheatham, 3rd, Evaluation of enhanced sampling provided by accelerated molecular dynamics with Hamiltonian replica exchange methods. *J. Phys. Chem. B,* 118:12 (2014), pp. 3543–52, doi:10.1021/jp4125099.

[5] Henriksen, N. M., H. S. Hayatshahi, D. R. Davis, and T. E. Cheatham, 3rd, Structural and energetic analysis of 2-aminobenzimidazole inhibitors in complex with the hepatitis C virus IRES RNA using molecular dynamics simulations. *J. Chem. Inf. Model.*, 54:6 (2014), pp. 1758–72, doi:10.1021/ci500132c.

[6] Yildirim, I., H. A. Stern, J. D. Tubbs, S. D. Kennedy, and D. H. Turner, Benchmarking AMBER force fields for RNA: comparisons to NMR spectra for single-stranded r(GACC) are improved by revised chi torsions. *J. Phys. Chem. B*, 115:29 (2011), pp. 9261–70, doi:10.1021/jp2016006.

[7] Henriksen, N. M., D. R. Roe, and T. E. Cheatham, 3rd, Reliable oligonucleotide conformational ensemble generation in explicit solvent for force field assessment using reservoir replica exchange molecular dynamics simulations. *J. Phys. Chem. B*, 117:15 (2013), pp. 4014–27, doi:10.1021/jp400530e.

[8] Steinbrecher, T., J. Latzer, and D. Case, Revised AMBER parameters for bioorganic phosphates. *J. Chem. Theory Comput.*, 8:11 (2012), pp. 4405–4412, doi:10.1021/ct300613v.

[9] Izadi, S., R. Anandakrishnan, and A. V. Onufriev, Building Water Models: A Different Approach. *J. Phys. Chem. Lett.*, 5:21 (2014), p. 3863–3871, doi:10.1021/jz501780a.

[10] Salomon-Ferrer, R., A. W. Goetz, D. Poole, S. Le Grand, and R. C. Walker, Routine microsecond molecular dynamics simulations with AMBER, Part II: Particle Mesh Ewald. *J. Chem. Theory Comput.*, 9 (2013), pp. 3878–3888, doi:10.1021/ct200909j.

[11] Hopkins, C. W., S. Le Grand, R. C. Walker, and A. E. Roitberg, Long-time-step molecular dynamics through hydrogen mass repartitioning. *J. Chem. Theory Comput.*, 11 (2015), pp. 1864–1874, doi:10.1021/ct5010406.

Ertekin-Taner, Nilufer

[1] Zou, F., et al., Brain expression genome-wide association study (eGWAS) identifies human disease-associated variants. *PLoS genetics*, 2012. 8(6): p. e1002707.

[2] Allen, M., et al., Novel late-onset Alzheimer disease loci variants associate with brain gene expression. *Neurology*, 2012. 79(3): p. 221-8.

[3] Gusareva, E.S. and K. Van Steen, Practical aspects of genome-wide association interaction analysis. *Hum Genet*, 2014. 133(11): p. 1343-58.

[4] Zou, F., et al., Brain expression genome-wide association study (eGWAS) identifies human disease-associated variants. *PLoS Genet*, 2012. 8(6): p. e1002707.

[5] Purcell, S., et al., PLINK: a tool set for whole-genome association and population-based linkage analyses. *Am J Hum Gene*t, 2007. 81(3): p. 559-75.

[6] Hemani, G., et al., EpiGPU: exhaustive pairwise epistasis scans parallelized on consumer level graphics cards. *Bioinformatics*, 2011. 27(11): p. 1462-5.

[7] Schupbach, T., et al., FastEpistasis: a high performance computing solution for quantitative trait epistasis. *Bioinformatics*, 2010. 26(11): p. 1468-9.

Gerlt, John

[1 http://efi.igb.illinois.edu/est-precompute/

[2 Gerlt, A., et al., The Enzyme Function Initiative. B*iochem.*, 50:46 (2011), pp. 9950–62, doi:10.1021/bi201312u.

[3] http://web.expasy.org/docs/relnotes/relstat.html

[4] http://www.ebi.ac.uk/uniprot/TrEMBLstats/

[5] Atkinson, H. J., J. H. Morris, T. E. Ferrin, and P. C. Babbitt, Using sequence similarity networks for visualization of relationships across diverse protein superfamilies. *PLoS One*, 4:2 (2009), e4345, doi:10.1371/journal.pone.0004345.

[6] http://weizhong-lab.ucsd.edu/cd-hit/

[7] http://efi.igb.illinois.edu/est-precompute49/

Hirata, So

[1] Hirata, S., K. Gilliard, X. He, J. Li, and O. Sode, *Ab initio* molecular crystal structure, spectra, and phase diagrams. *Acc. Chem. Res.*, 47 (2014), 2721, doi:10.1021/ar500041m.

[2] He, X., O. Sode, S. S. Xantheas, and S. Hirata, Second-order many-body perturbation study of ice Ih. *J. Chem. Phys.*, 137 (2012), 204505, doi:10.1063/1.4767898.

[3] Gilliard, K., O. Sode, and S. Hirata, Second-order many-body perturbation and coupled-cluster singles and doubles study of ice VIII. *J. Chem. Phys.*, 140 (2014), 174507, doi:10.1063/1.4873919.

[4] Sode, O., M. Keçeli, K. Yagi, and S. Hirata, Fermi resonance in solid CO_2 under pressure. *J. Chem. Phys.*, 138 (2013), 074501, doi:10.1063/1.4790537.

[5] Li, J., O. Sode, G. A. Voth, and S. Hirata, A solid–solid phase transition in carbon dioxide at high pressures and intermediate temperatures. *Nat. Commun.*, 4 (2013), 2647, doi:10.1038/ncomms3647.

[6] Li, J., O. Sode, and S. Hirata, Second-order many-body perturbation study of thermal expansion of solid carbon dioxide. *J. Chem. Theory Comput.*, 11 (2015), 224, doi:10.1021/ct500983k.

Jakobsson, Eric

[1] J. McCammon, A., B. R. Gelin, and M. Karplus, *Nature*, 267, 585 (1977)

[2] D. Branton et al., The potential and challenges of nanopore sequencing. *Nat Biotech* 26, 1146-1153 (2008).

[3] Sun, Jie, Shaoying Lu, Mingxing Ouyang, Li-Jung Lin, Yue Zhuo, Bo Liu, Shu Chien, Benjamin G. Neel, and Yingxiao Wang. "Antagonism between binding site affinity and conformational dynamics tunes alternative cis-interactions within Shp2." *Nature communications 4* (2013).

[4] S. A. Beeser, D. P. Goldenberg, and T G. Oas. *Journal of molecular biology 269*, no. 1 (1997): 154-164

Jasiuk, Iwona

[1] www.nof.org

[2] Weiner, S., and W. Traub, Bone structure: from angstroms to microns. *FASEB*, 6 (1992), pp. 879–885.

[3] Hamed, E., Y. Lee, and I. Jasiuk, Multiscale modeling of elastic properties of cortical bone. *Acta Mechanica*, 213 (2010), pp. 131–154, doi:10.1007/s00707-010-0326-5.

[4] Hamed, E, Y. Lee, T. Liszka, and I. Jasiuk, Multiscale Modeling of Elastic Moduli of Trabecular Bone. *J. Royal Soc. Interface*, 9 (2012), pp. 1654–1673, doi:10.1098/rsif.2011.0814.

[5] Hamed, E., and I. Jasiuk, Multiscale damage and strength of lamellar bone modeled by cohesive finite elements. *J. Mech. Behav. Biomed. Mater.*, 28 (2013), pp. 94–110, doi:10.1016/j.jmbbm.2013.05.025.

Johnson, Curtis

[1] Muthupillai, R., D. J. Lomas, P. J. Rossman, J. F. Greenleaf, A. Manduca, and R. L. Ehman, Magnetic Resonance Elastography by Direct Visualization of Propagating Acoustic Strain Waves. *Science*, 269:5232 (1995), pp. 1854-1857.

[2] Johnson, C. L., M. D. J. McGarry, A. A. Gharibans, J. B. Weaver, K. D. Paulsen, H. Wang, W. C. Olivero, B. P. Sutton, and J. G. Georgiadis, Local Mechanical Properties of White Matter Structures in the Human Brain. *NeuroImage*, 79 (2013), pp. 145-152.

[3] Johnson, C. L., J. L. Holtrop, M. D. J. McGarry, J. B. Weaver, K. D. Paulsen, J. G. Georgiadis, and B. P. Sutton, 3D Multislab, Multishot Acquisition for Fast, Whole-Brain MR Elastography with High Signal-to-Noise Efficiency. *Magn. Reson. Med.*, 71:2 (2014), pp. 477-485.

[4] McGarry, M. D. J., E. E. W. Van Houten, C. L. Johnson, J. G. Georgiadis, B. P. Sutton, J. B. Weaver, and K. D. Paulsen, Multiresolution MR Elastography Using Nonlinear Inversion. *Med. Phys.*, 39:10 (2012), pp. 6388-6396.

Khalili-Araghi, Fatemeh

[1] Delmar, J. A., C. C. Su, and E. W. Yu, Bacterial multidrug efflux transporters. *Annu. Rev. Biophys.*, 43 (2014), pp. 93–117, doi:10.1146/annurev-biophys-051013-022855.

[2] Du, D., et al., Structure of the AcrAB-TolC multidrug efflux pump. *Nature*, 509 (2014), pp. 512–515, doi:10.1038/nature13205.

[3] Trabuco, L. G., et al., Molecular dynamics flexible fitting: a practical guide to combine cryo-electron microscopy and X-ray crystallography. *RNA Dynamics*, 49 (2009), pp. 174–180, doi:10.1016/j.ymeth.2009.04.005.

[4] Eicher, T., et al., Transport of drugs by the multidrug transporter AcrB involves an access and a deep binding pocket that are separated by a switch-loop. *Proc. Natl Acad. Sci. U.S.A.*, 109 (2012), pp. 5687–5692, doi:10.1073/pnas.1114944109.

Klein, Michael

[1] Payandeh, J., T. Scheuer, N. Zheng, W.A. Catterall The crystal structure of a voltage-gated sodium channel *Nature*, 475 (2011), pp. 353–358

Luthey-Schulten, Zaida

[1] Hallock M. J., J. Stone, E. Roberts, C. Fry, and Z. Luthey-Schulten, Simulation of reaction diffusion processes over biologically http://dx.doi.org/10.1016/j.parco.2014.03.009relevant size and time scales using multi-GPU workstations. *Parallel Comput.*, 40:5-6 (2014), pp. 86–99, doi:10.1016/j.parco.2014.03.009.

[2] Peterson J. R., M. J. Hallock, J. A. Cole, and Z. Luthey-Schulten, A Problem Solving Environment for Stochastic Biological Simulations. *PyHPC 2013*, Denver, Colo., November 17–22, 2013.

[3] Roberts, E., J. E. Stone, and Z. Luthey-Schulten, Lattice Microbes: high-performance stochastic simulation method for the reaction-diffusion master equation. *J. Comput. Chem.*, 32:3 (2013), pp. 245–55, doi:10.1002/jcc.23130.

[4] Held W. A., B. Ballou, S. Mizushima, and M. Nomura, Assembly Mapping of 30S Ribosomal Proteins from *Escherichia coli:* Further studies. *J. Biol. Chem.*, 249:10 (1974), pp. 3103–11.

[5] Adilakshmi, T., P. Ramaswamy, and S. A. Woodson, Protein-independent folding pathway of the 16S rRNA 5' domain. *J. Mol. Biol.*, 351:3 (2005), pp. 508–19, doi:10.1016/j.jmb.2005.06.020.

[6] Adilakshmi T., D. L. Bellur, and S. A. Woodson, Concurrent Nucleation of 16S Folding and Induced fit in 30S Ribosome Assembly. *Nature*, 455:7217 (2008), pp. 1268–72, doi:10.1038/nature07298.

[7] Kim H., et al., Protein-guided RNA dynamics during early ribosome assembly. *Nature*, 506:7488 (2014), pp. 334–8, doi:10.1038/nature13039.

[8] Talkington M. W. T., G. Siuzdak, and J. R. Williamson, An assembly landscape for the 30S ribosomal subunit. *Nature*, 438:7068 (2005), pp.628–32, doi:10.1038/nature04261.

[9] Sykes, M. T., and J. R. Williamson, A complex assembly landscape for the 30S ribosomal subunit. *Ann. Rev. Biophys.*, 38 (2009), pp. 197–215, doi:10.1146/annurev.biophys.050708.133615.

[10] Bunner A. E., A. H. Beck, and J. R. Williamson, Kinetic cooperativity in Escherichia coli 30S ribosomal subunit reconstitution reveals additional complexity in the assembly landscape. *Proc. Nat. Acad. Sci. USA*, 107:12 (2010), pp. 5417–22, doi:10.1073/pnas.0912007107.

[11] Mulder A. M., et al., Visualizing ribosome biogenesis: parallel assembly pathways for the 30S subunit. *Science*, 330:6004 (2010), pp.673–7, doi:10.1126/science.1193220.

[12] Sanamrad, A., et al., Single-Particle Tracking Reveals that Free Ribosomal Subunits are not Excluded from the *Escherichia coli* Nucleoid. P*roc. Natl. Acad. Sci. USA*, 111:31 (2014), pp.11413–18, doi:10.1073/pnas.1411558111.

[13] Bakshi, S., A. Siryaporn, M. Goulian, and J. C. Weisshaar, Superresolution Imaging of Ribosomes and RNA Polymerase in Live *Escherichia coli* Cells. *Mol. Microbiol.*, 85:1 (2012), pp.21–38, doi:10.1111/j.1365-2958.2012.08081.x.

[14] Selinger D. W., R. M. Saxena, K. J. Cheung, G. M. Church, and C. Rosenow, Global RNA Half-Life Analysis in *Escherichia coli* Reveals Positional Patterns of Transcript Degradation. *Genome Res.*, 13:2 (2003), pp. 216–223, doi:10.1101/gr.912603.

[15] Cagliero C., R. S. Grand, M. B. Jones, D. J. Jin, and J. M. O'Sullivan, Genome conformation capture reveals that the Escherichia coli chromosome is organized by replication and transcription. *Nucleic Acids Res.*, 41:12 (2013), pp.6058–71, doi:10.1093/nar/gkt325.

[16] Wang W., G. W. Li, C. Chen, X. S. Xie, and X. Zhuang, Chromosome organization by a nucleoid-associated protein in live bacteria. *Science*, 33:6048 (2011), pp. 1445–9, doi:10.1126/science.1204697.

[17] Hong S.-H., et al., Caulobacter chromosome i*n vivo* configuration matches model predictions for a supercoiled polymer in a cell-like confinement. *Proc. Natl. Acad. Sci. USA*, 110:5 (2012), pp.1674–9, doi:10.1073/pnas.1220824110.

[18] O'Brien, E. J., J. A. Lerman, R. L. Chang, D. R. Hyduke, B. Ø. Palsson, Genome-scale models of metabolism and gene expression extend and refine growth phenotype prediction. *Mol. Syst. Biol.*, 9:1 (2013), 693, doi: 10.1038/msb.2013.52.

Mainzer, Liudmila

[1] Broad Institute of MIT and Harvard, "The Genome Analysis Toolkit." https://www.broadinstitute.org/gatk/ (Accessed 4/15/15)

[2] Perfsuite. http://perfsuite.ncsa.illinois.edu/ (Accessed 4/15/15)

[3] Cray Profiler. https://www.nersc.gov/users/software/debugging-and-profiling/craypat/ (Accessed 4/15/15)

[4] OVIS Wiki homepage. https://ovis.ca.sandia.gov/mediawiki/index.php/Main_Page (Accessed 4/15/15)

[5] Valgrind. http://valgrind.org/ (Accessed 4/15/15)

[6] Novocraft. www.novocraft.com (Accessed 4/15/15)

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, Quantum-classical path integral with self-consistent solvent-driven propagators. *J. Phys. Chem.*, 117 (2013), 13357, doi:10.1021/jp4043123.

[4] Feynman, R. P., Space-time approach to non-relativistic quantum mechanics. *Rev. Mod. Phys.*, 20 (1948), 367, doi:10.1103/RevModPhys.20.367.

[5] Makri, N., Exploiting classical decoherence in quantum dissipative dynamics. *Chem. Phys. Lett.*, 593 (2014), 93, doi:10.1016/j.cplett.2013.11.064.

[6] 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.

[7] 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.

Rienstra, Chad

[1] Knowles, T., M. Vendruscolo, and C. M. Dobson, The physical basis of protein misfolding disorders. *Physics Today*, 68 (2015), pp. 36–41, doi:10.1063/PT.3.2719.

[2] Gō, N., and H. Abe, Noninteracting Local-Structure Model of Folding and Unfolding Transition in Globular Proteins. I. Formulation. *Biopolymers*, 20 (1981), pp. 991–1011, doi:10.1002/bip.1981.360200511.

[3] Lai, J., K. Chen, and Z. Luthey-Schulten, Structural Intermediates and Folding Events in the Early Assembly of the Ribosomal Small Subunit. *J. Phys. Chem. B*, 117 (2013), pp. 13335–13345, doi:10.1021/jp404106r.

[4] Berendsen, H., et al., GROMACS: A message-passing parallel molecular dynamics implementation. *Comp. Phys. Comm.*, 91 (1995), pp. 43–56, doi:10.1016/0010-4655(95)00042-E.

[5] Laio, A., and M. Parrinello, Escaping free energy minima. *Proc. Natl. Acad. Sci. USA*, 99 (2002), pp. 12562–12566, doi:10.1073/pnas.202427399.

[6] Phillips, J. C., et al., Scalable molecular dynamics with NAMD. *J. Comput. Chem.*, 26 (2005), pp. 1781–1802, doi:10.1002/jcc.20289.

[7] Brooks, B. R., et al., CHARMM: The Biomolecular Simulation Program. *J. Comput. Chem.*, 30 (2009), pp. 1545–1614, doi:10.1002/jcc.21287.

[8] Piana, S., and A. Laio, A bias-exchange approach to protein folding. *J. Phys. Chem. B*, 111 (2007), pp. 4553–4559, doi:10.1021/jp067873l.

[9] Bonomi, M., et al., PLUMED: A portable plugin for free-energy calculations with molecular dynamics. *Comp. Phys. Comm.*, 180 (2009), pp. 1961–1972, doi:10.1016/j.cpc.2009.05.011.

[10] Fiorin, G., et al., Using Collective variables to drive molecular dynamics simulations. *Mol. Phys.*, 111 (2013), pp. 3345–3362, doi:10.1080/00268976.2013.813594.

Roux, Benoit

[1] Toyoshima, C., M. Nakasako, H. Nomura, and H. Ogawa, Crystal structure of the calcium pump of sarcoplasmic reticulum at 2.6 Å resolution. *Nature*, 405 (2000), pp. 647–655, doi:10.1038/35015017.

[2] Toyoshima, C., Structural aspects of ion pumping by Ca^{2+}-atpase of sarcoplasmic reticulum. *Arch. Biochem. Biophys.*, 476 (2008), pp. 3–11, doi:10.1016/S0167-7306(06)41009-7.

[3] Toyoshima, C., How Ca^{2+}-atpase pumps ions across the sarcoplasmic reticulum membrane. *Biochim. Biophys. Acta*, 1793 (2009), pp. 941–946, doi:10.1016/j.bbamcr.2008.10.008.

[4] Sørensen, T. L. M., J. V. Møller, and P. Nissen, Phosphoryl transfer and calcium ion occlusion in the calcium pump. *Science*, 304 (2004), pp. 1672–1675, doi:10.1126/science.1099366.

[5] Olesen, C., et al., The structural basis of calcium transport by the calcium pump. *Nature*, 450 (2007), pp. 1036–1042, doi:10.1038/nature06418.

[6] Toyoshima, C., et al., Crystal structures of the calcium pump and sarcolipin in the Mg^{2+}-bound E1 state. *Nature*, 495 (2013), pp. 260–265, doi:10.1038/nature11899.

[7] Allen, M. P., and D. P. Tildesley, *Computer Simulation of Liquids.* (Oxford University Press, Oxford, U.K., 1987).

[8] Pan, A. C., D. Sezer, and B. Roux, Finding transition pathways using the string method with swarms of trajectories. *J. Phys. Chem. B*, 112 (2008), pp. 3432–3440, doi:10.1021/jp0777059.

[9] Gan, W., S. Yang, and B. Roux, Atomistic view of the conformational activation of Src kinase using the string method with swarms-of-trajectories. *Biophys. J.*, 97 (2009), pp. L8–L10, doi:10.1016/j.bpj.2009.06.016.

[10] Phillips, J. C., et al., Scalable molecular dynamics with NAMD. *J. Comput. Chem.*, 26 (2005), pp. 1781–1802, doi:10.1002/jcc.20289.

Schulten, Klaus

[1] 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.

[2] Goh, B. C., et al., Atomic modeling of an immature retroviral lattice using molecular dynamics and mutagenesis *Structure*, 23:1414-1425, 2015.

[3] Ma, W. and Schulten, K., Mechanism of substrate translocation by a ring-shaped ATPase motor at millisecond resolution. *J. Am. Chem. Soc.*, 137:3031-3040, 2015

Tajkhorshid, Emad

[1] Jardetzky, O., Simple allosteric model for membrane pumps. *Nature*, 211:5052 (1966), pp. 969–970, doi:10.1038/211969a0.

[2] Li, J., P.-C. Wen, M. Moradi, and E. Tajkhorshid, Computational characterization of structural dynamics underlying function in active membrane transporters. *Curr. Opin. Struct. Biol.*, 31 (2015), pp. 96–105, doi:10.1016/j.sbi.2015.04.001.

[3] Moradi, M., and E. Tajkhorshid, Mechanistic picture for conformational transition of a membrane transporter at atomic resolution. *Proc. Nat. Acad. Sci. USA*, 110 (2013), pp. 18916–18921, doi:10.1073/pnas.1313202110.

[4] Moradi, M., and E. Tajkhorshid, Computational recipe for efficient description of large-scale conformational changes in biomolecular systems. *J. Chem. Theory Comput.*, 10 (2014), pp. 2866–2880, doi:10.1021/ct5002285.

[5] Moradi, M., G. Enkavi, and E. Tajkhorshid, Atomic-level characterization of transport cycle thermodynamics in a membrane transporter. *Nat. Commun.*, (accepted).

[6] Pan, A. C., D. Sezer, and B. Roux. Finding transition pathways using the string method with swarm of trajectories. *J. Phys. Chem. B*, 112:11 (2008), pp. 3432–3440, doi:10.1021/jp0777059.

[7] Jiang, W., et al., Generalized scalable multiple copy algorithms for molecular dynamics simulations in NAMD. *Comput Phys Commun.*, 185 (2014), pp. 908–916, doi:10.1016/j.cpc.2013.12.014.

Voth, Gregory

[1] Dama, J. F., et al., The Theory of Ultra-Coarse-Graining. 1. General Principles. *J. Chem. Theory Comput.*, 9:5 (2013), pp. 2466–2480, doi:10.1021/ct4000444.

[2] Davtyan, A., J. F. Dama, A. V. Sinitskiy, and G. A. Voth, The Theory of Ultra-Coarse-Graining. 2. Numerical Implementation. *J. Chem. Theory Comput.*, 10:12 (2014), pp. 5265–5275, doi:10.1021/ct500834t.

[3] Grime, J. M. A., and G. A. Voth, Highly Scalable and Memory Efficient Ultra-Coarse-Grained Molecular Dynamics Simulations. *J. Chem. Theory Comput.*, 10:1 (2014), pp. 423–431, doi:10.1021/ct400727q.

SOCIAL SCIENCE, ECONOMICS, & HUMANITIES

Hansen, Lars

[1] Cai, Y., K .L. Judd, and T. S. Lonzek, Numerical dynamic programming with error control: an application to climate policy. *Oper. Res.*, submitted.

[2] Cai, Y., K. L. Judd, and T. S. Lontzek, The social cost of carbon with economic and climate risks. *J. Polit. Econ.*, under review, arXiv preprint arXiv:1504.06909.

[3] Cai, Y., and K. L. Judd, Advances in numerical dynamic programming and new applications. in *Handbook of Computational Economics*, 3 (2014), pp. 479–516.

[4] Cai, Y., and K. L. Judd, Dynamic programming with Hermite approximation. *Math. Meth. Oper. Res.*, 81:3 (2015), pp. 245–267, doi:10.1007/s00186-015-0495-z.

[5] 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. Clim. Change*, 5 (2015), pp. 441–444, doi:10.1038/nclimate2570.

[6] Cai, Y., K. L. Judd, T. M. Lenton, T. S. Lontzek, and D. Narita, Environmental tipping points significantly affect the cost-benefit assessment of climate policies. *Proc. Nat. Acad. Sci. U.S.A.*, 112:15 (2015), pp. 4606–4611, doi:10.1073/pnas.1503890112.

[7] Cai, Y., J. Steinbuks, J. W. Elliott, and T. W. Hertel, The effect of climate and technological uncertainty in crop yields on the optimal path of global land use. (2014), 10.1596/1813-9450-7009. [The World Bank Policy Research Working Paper 7009]

[8] 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, doi:10.1007/s10614-014-9419-x.

[9] Cai, Y., K. L. Judd, and J. Steinbuks, A nonlinear certainty equivalent method for stochastic dynamic problems. *Quant. Econ.*, submitted.

[10] Yeltekin, S., Y. Cai, and K. L. Judd, Computing equilibria of dynamic games. *Oper. Res.*, submitted.

Althaus, Scott

[1] http://doc.akka.io/docs/akka/2.2.4/intro/what-is-akka.html

[2] http://sharc.hathitrust.org/features

[3] Capitanu, B., et al., Extracted Feature Dataset from 4.8 Million HathiTrust Digital Library Public Domain Volumes (0.2) [Dataset]. *HathiTrust Research Center* (2015), doi:10.13012/j8td9v7m.

[4] Organisciak, P., L. Auvil, S. Bhattacharyya, and J. S. Downie, The HTRC extracted features dataset. *Joint CSDH/SCHN & ACH Digital Humanities Conference*, Ottawa, Canada, June 1–3, 2015.

[5] Organisciak, P., S. Bhattacharyya, L. Auvil, and J. S. Downie. Large-scale text analysis through the HathiTrust Research Center. *Digital Humanities 2014 Conference*, Lausanne, Switzerland, July 8–12, 2014.

[6] http://bookworm.htrc.illinois.edu/

[7] http://www.phoenixdata.org/

INDEX

A

Aksimentiev, Aleksei 68
Althaus, Scott 180
Aluru, Narayana 70
Amaro, Rommie 128
Arraki, Kenza 188

B

Balsara, Dinshaw 72
Bedford, Matthew 190
Bodony, Daniel 74
Burrows, Adam 36

C

Caetano-Anolles, Gustavo 130
Calhoun, Jon 192
Ceperley, David 76
Cheatham, Thomas 132
Clark, Bryan 78

D

Di Girolamo, Larry 40
Di Matteo, Tiziana 16
Draayer, Jerry 80

E

Elghobashi, Said 82
Ertekin, Elif 84
Ertekin-Taner, Nilufer 134

F

Fischer, Paul 86

G

Garcia, Marcelo 42
Gerlt, John 136
Glotzer, Sharon 88
Gropp, Bill 124
Grosman, Claudio 138

H

Hammes-Schiffer, Sharon 140
Hansen, Lars 178
Hirata, So 142

J

Jakobsson, Eric 144
Jasiuk, Iwona 146
Johnson, Curtis 148
Jones, Alexandra 194
Jordan, Thomas 44

K

Karimabadi, Homayoun 18
Kasson, Peter 157
Kemball, Athol 90
Khalili-Araghi, Fatemeh 150
Klein, Michael 152
Klimeck, Gerhard 92
Kokkila, Sara 196
Koric, Seid 94

L

Lasher-Trapp, Sonia 46
Lele, Sanjiva 96
Lentz, Eric 22
Levin, Deborah 98
Luthey-Schulten, Zaida 154

M

Mainzer, Liudmila 158
Makri, Nancy 160
Makwana, Kiritkumar 24
Masud, Arif 162
Mathews, Edwin 198
Minot, Ariana 200
Mori, Warren 100

O

O'Shea, Brian 26
Ostoja-Starzewski, Martin 102
Ott, Christian 104

P

Pande, Vijay 164
Pogorelov, Nikolai 28

Q

Quinn, Thomas 106

R

Reed, Patrick 48
Rienstra, Chad 166
Roux, Benoit 168

S

Schleife, Andre 108
Schulten, Klaus 170
Shapiro, Stuart 30
Simpson, Jamesina 50
Slota, George 202
Sriver, Ryan 52
Stein, Robert F. 32
Sugar, Robert 115

T

Taha, Ahmed 122
Tajkhorshid, Emad 172
Thomas, Brian 112
Trapp, Robert 54

U

Um, Junshik 56

V

Virgil-Fowler, Derek 204
Voth, Gregory 174

W

Wang, Shaowen 182
Wentzcovitch, Renata 58
Wilhelmson, Robert 64
Woodward, Paul R. 34
Woosley, Stan 118
Wuebbles, Donald J. 61

Y

Yeung, Pui-kuen 110

Z

Zemskova, Varvara 206

NCSA

ILLINOIS
UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

Blue Waters is supported by the National Science Foundation

ISBN 978-0-9908385-3-1