

BLUE WATERS

SUSTAINED PETASCALE IN ACTION:
ENABLING TRANSFORMATIVE RESEARCH

2017 ANNUAL REPORT



BLUE WATERS

SUSTAINED PETASCALE IN ACTION:
ENABLING TRANSFORMATIVE RESEARCH

2017 ANNUAL REPORT

Executive Editor

Barbara Jewett

Editors

Susan Szuch
Catherine Watkins

Creative Director

Steve Duensing

Proofreaders

Susan Szuch
Catherine Watkins

Project Director

William Kramer

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

Visit <https://bluwaters.ncsa.illinois.edu/science-teams> for the latest on Blue Waters-enabled science and to watch the 2017 Blue Waters Symposium presentations.

CLASSIFICATION KEY

To provide an overview of how science teams are using Blue Waters, researchers were asked if their work fit any of the following classifications (number responding in parentheses):

DI	Data-intensive: uses large numbers of files, e.g. large disk space/bandwidth, or automated workflows/off-site transfers (39)
GA	GPU-accelerated: written to run faster on XK nodes than on XE nodes (34)
TN	Thousand node: scales to at least 1,000 nodes for production science (65)
MI	Memory intensive: uses at least 50 percent of available memory on 1,000-node runs (18)
BW	Blue Waters: Research only possible on Blue Waters (27)
MP	Multi-physics/multi-scale: job spans multiple length/time scales or physical/chemical processes (47)
ML	Machine learning: employs deep learning or other techniques, includes "big data" (9)
CI	Communication-intensive: requires high-bandwidth/low-latency interconnect for frequent, tightly coupled messaging (32)
IA	Industry applicable: Researcher has private sector collaborators or results directly applicable to industry (5)

TABLE OF CONTENTS

3	A MESSAGE FROM THE NCSA DIRECTOR
4	PROJECT DIRECTOR BILL KRAMER – SCIENCE NOT POSSIBLE ANYWHERE ELSE
6	WORKLOAD ANALYSIS REVEALS HOW BLUE WATERS IS ACTUALLY USED
10	PETASCALE APPLICATION IMPROVEMENT DISCOVERY PROGRAM
14	BLUE WATER'S EDUCATION, OUTREACH AND WORKFORCE DEVELOPMENT ACTIVITIES
16	EXTENDED ABSTRACTS
16	<i>space science</i>
62	<i>geoscience</i>
100	<i>physics & engineering</i>
168	<i>computer science & engineering</i>
190	<i>biology, chemistry & health</i>
250	<i>social science, economics, & humanities</i>
256	<i>graduate fellows</i>
290	TRANSFORMATIVE SCIENCE (MEASURING BLUE WATERS)
294	SCIENCE AND ENGINEERING ADVISORY TEAM COMMITTEE
295	OTHER BLUE WATERS ALLOCATIONS (NO REPORT SUBMITTED)
296	REFERENCES
316	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 THE NCSA DIRECTOR



I continue to be amazed by the vast range of creative, limit-pushing research that scientists submit to this publication year after year. With the support of the National Science Foundation (NSF) and the University of Illinois, the National Center for Supercomputing Applications' (NCSA) Blue Waters Project continues to empower scientists to make

discoveries that have immense impact in a diverse range of fields, spark new understanding of our world, and open new avenues for future research.

An example of a once in a generation discovery is gravitational wave detection. Even prior to NCSA's status as an official member of the Laser Interferometer Gravitational-wave Observatory (LIGO) consortium, the Blue Waters supercomputer was used by NCSA's astronomy and gravity groups, as well as many other researchers, to simulate gravitational wave sources. But verifying Einstein's Theory of Relativity is not the only cosmological work Blue Waters is performing. For example, Tiziana Di Matteo of Carnegie Mellon University successfully carried out a cosmological simulation that used all of Blue Waters that yielded information "crucial to understanding the formation of the first quasars and galaxies from the smallest to the rarest and most luminous."

There are also the discoveries that directly impact people's lives. Discoveries like high-resolution, high-precision elevation maps of the Arctic and Antarctic. Last year, Paul Morin of the Polar Geospatial Center of the University of Minnesota led a project to use Blue Waters to map the Arctic to a precision never before possible. This year, Ian Howat of The Ohio State University collaborated with Morin's team to use Blue

Waters to create the Reference Elevation Model of Antarctica (REMA). REMA requires processing over 500,000 files, totaling over 500 terabytes! This would not be possible on any other academic supercomputer. Or the work of Leigh Orf, University of Wisconsin-Madison, who leads a collaborative effort to learn more about tornadoes spawned from supercell thunderstorms. These are often the strongest, most deadly tornadoes, yet meteorologists are at a loss to forecast them. The researchers used Blue Waters to carry out 15-meter resolution simulations that utilized over 360,000 cores—approximately half of the machine—and created over 300 terabytes of compressed data. The resulting discovery? New insight into the movement of air currents and the resulting effect on tornado formation.

These are but a few of the remarkable discoveries you'll find in the pages that follow. In the past few years, the Blue Waters Project has enabled remarkable work in biology, chemistry, physics, geosciences, cosmology and astrophysics, atmospheric science, and many other fields such as economics and social sciences. Even more remarkable breakthroughs will be forthcoming as NCSA continues to partner with scientists around the nation to change the world as we know it.

Dr. William "Bill" Gropp, Director
National Center for Supercomputing
Applications

SCIENCE NOT POSSIBLE ANYWHERE ELSE



Every day, the Blue Waters Project makes previously impossible investigations and insights possible for researchers across the United States. Every day, the combination of massive computing power and the intellectual might of pioneering scientists and engineers creates opportunities for us to better understand and shape our

world. While we are more than four years into full-service operations for Blue Waters and its associated support, training, and education efforts, time hasn't dulled the sense of wonder and delight I feel when I learn of another door that we opened to discovery.

This past year, a sample of 31 science teams that have used Blue Waters were surveyed and interviewed as part of a report meant to judge the effectiveness and productivity of Blue Waters. Using information gathered in the surveys, the report's authors at International Data Corporation's HPC Division, now known as Hyperion Research, ranked the impact of each team's findings into an "innovation index"—using a methodology they developed to analyze the effectiveness of 700-plus scientific projects, including international HPC projects. The IDC/Hyperion Research analysts noted in the report that "NCSA did an unusually thorough job of preparing [science teams] for Blue Waters." In fact, a notably higher percentage of scientific innovations enabled by using Blue Waters qualified for IDC's highest impact levels than was true for IDC's global database of supercomputer-enabled scientific innovations. You can find the entire IDC report on the Blue Waters portal (bluewaters.ncsa.illinois.edu).

The Blue Waters Project has always been about much more than just providing computing cycles, memory capacity (Blue Waters still has more memory than any other system in the open research space) and storage bytes. While such high performance resources are necessary for success, they are not sufficient for effectiveness. In other sections of this report, you will read about our expanded Petascale Application Improvement Discovery (PAID) program, where the Project provided millions of dollars to science teams and computational and data experts to improve the performance of applications (in a measurable manner). Likewise, the Project is proud to play a role educating and developing the next generation extreme-scale workforce, through our workshops, symposium, graduate fellowships, undergraduate internships, the Virtual School for Computational Science and Engineering, our funding for the HPC University and our workshops, and training allocations.

As you read through the reports on the following pages, you will notice badges on reports to signify use of some of the system's special features that make Blue Waters the National Science Foundation's most powerful system for computation and data analysis. These capabilities include: data-intensive problems that use large numbers of files, large disk space/bandwidth, or automated workflows/off-site transfers; GPU-accelerated (runs faster on XK nodes than on XE nodes); thousand-node (or "highly scalable"), i.e. scales to at least 1,000 nodes for production science input; memory intensive (used at least 50 percent of available memory on 1,000-node runs); multi-physics/multi-scale - spans multiple length/time scales or physical/chemical processes; machine learning (employs deep learning or other techniques, including "big data"); communication-intensive (requires high-bandwidth/low-latency interconnect for frequent tightly coupled messaging); and industrial application (private sector collaborators or results directly applicable to industry).

And of course, many reports carry the badge indicating "Only on Blue Waters" to signify that this research would not have been possible on other currently deployed U.S. open-science resources.

You will also notice this report, our fourth, contains 18 percent more high-impact result summaries than the 2016 report, which itself was a 40 percent increase over the 2015 report, which was a 25 percent increase over the 2014 report. The fact that in four years of providing one of most effective leadership-class systems in the world, Blue Waters is making such a remarkable difference to fields from physics, biomedicine, geo and space science, biology, economics, social and political science, and big data analysis continues to be a tribute to the vision of NSF, to our exceptional science and research teams, and to the dedicated Blue Waters Project staff.

While the scientific insights gained with support from Blue Waters are perhaps the most exciting way to judge the Project's impact, there are ways to consider its unduplicated role in the open-science community. To date, the Project has provided over 20 billion core-hour equivalents. Looking through another lens, Blue Waters is supporting approximately \$500 billion of funded research—projects supported not just by the National Science Foundation but also by NIH, NASA, DOE, NOAA and other agencies. As previously noted, these funded investigations might not even be possible (and certainly would take much longer) without Blue Waters.

One of those pioneering scientists reliant upon Blue Waters' massive compute power was Klaus Schulten, who passed away very unexpectedly October 31, 2016. Dr. Schulten blazed the trail in using what he called the "computational microscope," devoting over 40 years to "establishing the physical mechanisms underlying processes and organization in living systems from the atomic to the organism scale."

Not only was Dr. Schulten a leader in computational biophysics, he was a passionate supporter of the National Center for Supercomputing Applications (NCSA), Blue Waters, and the need for a robust national program of leadership class computing for academia; even with Blue Waters he did not have enough computing power to realize his research dreams. While Dr. Schulten is dearly missed by all of the Blue Waters Project staff, his computational microscope lives on as his work transitions to other principal investigators.

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

With great respect and thanks to be working with everyone,

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

WORKLOAD ANALYSIS REVEALS HOW BLUE WATERS IS ACTUALLY USED

As the National Science Foundation's flagship "leadership-class" supercomputer, Blue Waters was designed to enable the national scientific and research community to solve "grand challenge" problems that are orders of magnitude more complex than can be done on other systems. The great majority of Blue Waters machine time is awarded directly by the NSF in large allocations of a million node-hours or more, but the machine serves all disciplines through a diverse community of 800 to 1,000 users and 120-130 projects from 47 states. Given the important and unique role that Blue Waters plays in the U.S. research portfolio, it is important to have a detailed understanding of its actual usage by the scientific community as revealed by objective measurement and analysis. NCSA staff recently collaborated with the Center for Computational Research of the University at Buffalo, SUNY to produce a 106-page workload analysis of Blue Waters, available at <https://arxiv.org/ftp/arxiv/papers/1703/1703.00924.pdf>

BLUE WATERS IS INSTRUMENTED ON MANY LEVELS

Analysis requires data, the collection of which has been designed into all of the components of Blue Waters from the beginning. Great care has been taken to avoid any impact of performance and utilization monitoring on the performance or stability of the science applications running on the machine. Also, the design of the system monitoring infrastructure on Blue Waters focuses on acquiring a great deal of application performance data in ways that do not require the application teams to do any work or add any instrumentation or performance profiling tools to the codes.

An important data collection and management tool used to do this analysis is Open XDMoD (xdmod.ncsa.illinois.edu) [1], which provided comprehensive statistics on the number and type of computational jobs run, resources (computation, memory, disk, network, etc.) consumed, job wait times, quality of service, and, of particular importance for this work, detailed job-level performance data aggregated from a variety of sources.

Job information is provided by the Torque resource manager, which operates in coordination with the Cray ALPS resource manager and the MOAB job scheduler. Log entries are created when a job passes through various states of execution, including submission, changes of scheduler status (queued, held, deleted, eligible to run, start/running, etc.), and completion. Job completion is accompanied by exit code status that indicates whether the job completed normally.

The Lightweight Distributed Metric Service (LDMS, <https://ovis.ca.sandia.gov>) [2] is used to collect metrics on load average, memory usage, limited filesystem data transfers, network utilization with one-minute granularity at the compute node level and Gemini NIC/router level. For the XK nodes, GPU utilization and GPU memory utilization are also collected. LDMS also collects machine-specific register (MSR) data from the CPUs such as the number clock ticks, the number of instructions retired and the number of floating point operations performed.

Darshan (<http://www.mcs.anl.gov/research/projects/darshan/>) [3] captures application MPI I/O behavior, including properties such as patterns of access within files, with minimum overhead. Darshan data summarized for each job includes total bytes written and read, the number of opened files, the time spent in file system I/O operations, and the number of opened files in HDF5 and NetCDF formats.

APRUN logs provide the primary means to determine what application(s) was executed during a job, including the executable path, the number of nodes requested, and the layout of the tasks on the cores of the nodes. This information is used to map the job to an application algorithm, and, in conjunction with the MSR data, to determine concurrency/parallelism.

Finally, the ALTD and XALT [4] tools track which statically and dynamically linked libraries are used by each job.

All-in-all about eight billion data points and event records are collected each day about how Blue Waters is used, making not only the most productive system in the world, but also possibly the most measured system.

WORKLOAD ANALYSIS REQUIRES A SUPERCOMPUTER

The workload analysis itself was a challenging computational problem – requiring more than 35,000 node hours (over 1.1 million core hours) on Blue Waters to analyze roughly 95 TB of input data from over 4.5M jobs that ran on Blue Waters during the period of our analysis (April 1, 2013 – September 30, 2016) that spans the beginning to Full Service Operations for Blue Waters to the recent past. In the process, approximately 250 TB of data across 100M files was generated. This data was subsequently entered into MongoDB and a MySQL data warehouse to allow rapid searching, analysis and display in Open XDMoD. A workflow pipeline was established so that data from all future Blue Waters jobs will be automatically ingested into the Open XDMoD datawarehouse, making future analyses much easier.

PHYSICS AND BIOLOGY DOMINATE, BUT LONG TAIL GROWS

The Mathematical and Physical Sciences and Biological Sciences NSF directorates are the leading consumers of cycles on Blue Waters, typically accounting for more than 2/3 of all node hours used. However, the number of fields of science represented in the Blue Waters portfolio has increased in each year of its operation – more than doubling since its first year of operation – and the historical trend for all NSF directorates is toward increasing use of Blue Waters, albeit at much different overall scales of utilization.

The applications run on Blue Waters represent an increasingly diverse mix of disciplines, ranging from broadly-used community codes to specialized codes for specific scientific sub-disciplines. Common algorithms, as characterized by Phil Colella's original "seven dwarfs" algorithm classification, are roughly equally represented within the applications run on Blue Waters aside

from unstructured grids and Monte Carlo methods, which exhibit a smaller fraction of use. Optimized numerical libraries, essential for high performance, are used throughout areas in which such libraries are applicable. The top ten applications consume about 2/3 of all node hours, with the top five applications (NAMD, CHROMA, MILC, AMBER, and CACTUS) consuming about 1/2.

THROUGHPUT AND CAPABILITY COMPUTING CO-EXIST

Blue Waters supports a diverse mix of job sizes from single node jobs to jobs that use in excess of 20,000 nodes in a single application. The patterns of usage differ between the XE (CPU-only) and XK (GPU-accelerated) nodes. Single node jobs, some of which may be attributable to high throughput computing, represent less than 2% of the total node hours consumed on Blue Waters.

For XE node jobs, all of the major science areas using more than 1 million node hours run a mix of job sizes and all have very large jobs (> 4,096 nodes/>131,072 core equivalents). The relative proportions of job size vary between different parent science areas. The job size distribution weighted by node hours consumed peaks at 1,025 – 2,048 nodes for XE jobs. Impressively, the largest 3% of the jobs (by node hours used) account for 90% of the total node-hours consumed.

The majority of XE node hours on the machine are spent running parallel jobs that use some form of message passing for inter-process communication. At least 25% of the workload uses some form of threading, however the larger jobs (> 4,096 nodes) mostly use message passing with no threading. There is no obvious trend in the variation of thread usage over time the data was available.

USE OF GPU ACCELERATION INCREASES OVER TIME

The XK (GPU) nodes are as heavily utilized as the XE nodes, and the number of science teams using them has steadily increased, but remain a relatively small number of applications can use them well. Because the XE and XK usage is charged the same rate, and we now can measure GPU usage, we know the teams that are running on the XK nodes used the GPUs and also decide they have a performance improvement over the XE nodes. The parent sciences Molecular Biosciences, Chemistry, and Physics are the largest users with NAMD and AMBER the two most prevalent applications.

The job size distribution weighted by XK node hours consumed peaks at 65 – 128 nodes for the XK jobs. Similar to the XE nodes, the largest 7% of the jobs (by node-hour) account for 90% of the node-hours consumed on the XK nodes.

The aggregate GPU utilization (efficiency) varies significantly by application, with MELD achieving over 90% utilization of the GPU and GROMACS, NAMD, and MILC averaging less than 30% GPU utilization. However, for each of the applications, the

GPU utilization can vary significantly from job to job and within code steps. It is important to note the applications with lower GPU utilization are still getting performance and time to solution improvements and the application design and problems do not expect higher utilization.

Despite the increases in the number of applications using GPUs, it is important to understand that there are still many applications that do not, and maybe should not, use GPU technology without completely redesigning their methods – which could be a tremendous amount of work.

CPU-ONLY NODES ENABLE PETABYTE-MEMORY CALCULATIONS

Most jobs that run on the XE nodes use less than 50% of the memory available on the node. However, the distribution of memory use has a substantial tail to higher memory usage, as shown in Figure 1, representing capability-class jobs exploiting the 1.38 PB of aggregate memory available on the XE nodes.

Most jobs that run on the XK nodes use less than 25% of the available memory in the node with a short tail to higher usage. GPU memory usage is very small with few jobs using more than 1GB per GPU. The XE and XK nodes show no historical differences in memory use from year to year. For almost all applications and parent fields of science, memory usage has not changed over time under study.




Figure 1: 2D binned scatter plot showing the total peak memory usage for XE node jobs. The color coding shows the total wall hours for the jobs in each bin using a log scale. The diagonal line indicates the points corresponding to 64 GB per node. The points that are to the right of 11,300 nodes and above 800 TB indicate jobs that could not be executed on any other open system in the U.S or world.

STORAGE SYSTEM HANDLES DIVERSE LOAD

On average, Blue Waters’ three filesystems (home, projects, and scratch) have a balanced reads/writes ratio with large fluctuations. The volume of traffic on the largest filesystem (scratch) peaks at 10PB per month.

Partners’ jobs exhibit a wide range of I/O patterns. Overall there is a tendency to use a very large number of small files, which is a challenge for many parallel file systems and different from the projections that file size increases and system computational capability grows. Read and write rates stay significantly below possible peak filesystem performance, which actually is good and desired.

Partner jobs spend a very small fraction of time in filesystem I/O operations (0.04% of runtime for 90% of jobs). Many jobs utilize specialized libraries for their I/O operations (about 20% use MPI-IO, HDF5, or NetCDF).

HARD DATA GUIDES OPERATIONAL DECISIONS

Beyond enabling a comprehensive and evidence-based report on the workload and utilization of the NSF’s largest supercomputing resource, the greater legacy of the Blue Waters monitoring and data analysis effort will be improved scientific productivity for users of current and future leadership-class supercomputers. The Blue Waters team at NCSA has real-time and historical data available at their fingertips, allowing system operators and SEAS support staff to identify and adapt to unanticipated new usage patterns and to identify opportunities to work with science teams to improve the performance and throughput of their jobs. Lessons learned from the workload study, such as demand for aggregate and per-node memory, filesystem bandwidth, and GPU accelerators, will inform the design of the Blue Waters follow-on systems called for in the 2017 NSF solicitation “Towards a Leadership-Class Computing Facility – Phase 1”.



PETASCALE APPLICATION IMPROVEMENT DISCOVERY PROGRAM

EXECUTIVE SUMMARY

Blue Waters has always recognized the need for a comprehensive view of high-performance computing (HPC). While hardware is an important component of doing more, it alone cannot enable frontier science. The rapid rate of change of HPC hardware technology over the last decade has made it difficult for science domain specialists to keep their applications at the cutting edge of performance on today's HPC systems. Multicore processors and manycore accelerators (aka GPUs), the ever-present memory wall, the plethora of network topologies, and idiosyncrasies in parallel filesystems are some of the areas where experts in programming models, task mapping and load balancing at scale, scalable IO, and data movement can come to the aid of domain specialists interested in improving their applications' performance on Blue Waters and other large-scale HPC systems.

The Petascale Application Improvement and Discovery (PAID) program is one effort undertaken by the Blue Waters

project to help the domain specialists reach Petascale sustained performance. This effort was accomplished through partnerships with relevant computational experts. PAID was created in response to expressed interest by National Science Foundation (NSF) Petascale Computing Resource Allocations (PRAC) teams with allocations on Blue Waters, who worked to make improvements in application performance in partnership with technology experts in areas mentioned above. Each application team and technology team cooperatively developed a work plan containing period milestones and deliverables that were tracked over the multiyear life of the program. The requirement of a performance baseline for each participating application as the first milestone was the key to understanding and eventually quantifying application improvement upon completion of the work. The successes of the PAID program, the lessons learned, and best practices uncovered as well as recommendations for future similar efforts are discussed below.

INTRODUCTION

The PAID program began by first establishing a team of experts in various technology solutions, who were referred to as Improvement Method Enablers (IMEs). The technology areas were established by reviewing the results of a less structured and more self-directed program that funded work on improvements to existing applications by NSF PRAC teams. The five general categories selected from that first-year opportunity were: task mapping and load balancing, scalable IO and hierarchical data formats (HDF), Fourier transforms (FFTs), programming models best practices, and exploitation of graphics processing unit (GPU) accelerators. Eight expert teams were formed and led by researchers from five different universities. NVIDIA supported participation in PAID by providing one of their lead application analysts to work with one of the GPU improvement teams. Each IME team developed an individual statement of work that documented the expected number of PRAC teams they would work with, participation in Blue Waters workshops and other training events, and any resulting products from the program such as libraries or best practices. Each IME was assigned a Blue Waters' staff member as their Point of Contact (PoC) to facilitate tracking progress, reviewing milestones and deliverables, and addressing any support-related issues.

The participating application teams were NSF-selected PRAC teams with an award end date far enough in the future to allow for adequate time left in their allocation period so they could reap the benefits from any improvements to their application while still running their production science on Blue Waters. In total, 15 PRAC teams from seven fields of science engaged with the eight IME teams; some application teams worked with up to three different IME teams. Both the IME's and PRAC teams received funding from the Blue Waters Project as well as assistance from the Blue Waters Science and Engineering Applications Support team. Each PRAC team-IME team pairing developed individual work plans or statements of work with the assistance of a Blue Waters PoC. All statements of work required baseline performance as a first deliverable to enable quantitatively assessing the impact of the PAID program. This baseline was used to ascertain and quantify improvements in subsequent progress milestones and associated deliverables (usually reports). In order to assure progress, the Blue Waters Project hosted monthly progress review meetings for all PRACs and IMEs. The PRAC teams also presented a public summary their PAID findings at the annual Blue Waters symposium in order to inform other Blue Waters partners of potential approaches they could adopt in their own applications. The PoC for each science team was responsible for tracking progress, reviewing milestones and deliverables, and addressing any support-related issues. The PAID report in the 2016 Blue Waters Annual Report documented in more detail work done, work in progress, or to be done by the IMEs and the science teams.

The entire PAID program was continuously assessed by an independent team lead by Dr. Lizanne DeStefano. These assessments were used to make adjustments in the program as it progress and also to do summative assessments of the impact of the program.

RESULTS

The ultimate results of each partnership of a PRAC team with an IME team is a documented level of time to solution performance improvement that clearly translates to improved science throughput for the PRAC team. In most cases, the baseline performance indicates the performance of a code that was being used by the science teams for their major science projects and the PAID program allowed for potentially substantial gains in performance. In some situations, such as for established community codes, baseline performance designates an application that is already sufficiently optimized, such as a Compute Unified Device Architecture (CUDA)-accelerated application already tuned for Fermi GPUs (the version of GPU prior to Kepler) that runs on the newer K20x GPUs on Blue Waters. Percentage performance improvements should be expected to be lower compared to the upside potential for a code being ported to the GPU for the first time. Additionally, performance may be the final result but enabling new functionality is as important as traditional performance for existing methods for the teams benefiting from the new methods enabled in the code.

"The PAID program gave us access to computing specialists in GPU and HDF5 optimizations . . . which . . . provided us with rapid insights into our code performance. The insights provided by the PAID technical experts led to faster resolutions and improved outcomes for important code improvements that are needed to support the high-frequency simulations by our science community."

- Thomas Jordan University of Southern California

PERFORMANCE IMPROVEMENTS

Blue Waters provides the ideal platform for porting production codes to the NVIDIA GPUs while continuing to conduct scientific campaigns with its balanced mix of CPU nodes and GPU nodes. Several of the applications were ported to OpenACC or CUDA, or had already-existing OpenACC or CUDA code tuned to make better use of the NVIDIA K20x GPUs. Many projects teamed up with Wen-mei Hwu's IME to enable or improve the use of the NVIDIA GPUs on Blue Waters. The IME team worked on the CyberShake code AWP-ODC of the Thomas Jordan PRAC team to get a 1.3x speedup from a message passing interface (MPI) code that was already ported to CUDA with some additional porting of code not running on the GPU. The Hwu IME worked on the ChaNGa astrophysics code with Tom Quinn's PRAC team to get an impressive time to solution 1.6x speedup from Charm++-

enabled code tuned for the Fermi generation of NVIDIA GPUs at scales of several thousand nodes.

An OpenACC example of the success of this work is the 2.8x improvement in time-to-solution for the MS-FLUKSS code of the Nikolai Pogorelov PRAC team obtained by working with the Programming Model team lead by Bill Tang for a single Blue Waters GPU node (XK) over a single CPU node (XE). Another OpenACC example is the 3.9x speedup obtained with the 3D-FDTD code of Jamesina Simpson for a single Blue Waters GPU node (XK) over a single CPU node (XE). Both codes use MPI and the work to implement the new programming model into the distributed code remains to be completed. The results obtained from the OpenACC work is in line with earlier results showing that MPI codes using XK (GPU) nodes have typically 2x to 3x faster time-to-solution than their XE (CPU) counterparts, node-for-node, for complete applications.

The Blue Waters high-speed network uses a highly scalable communication, 3D torus topology interconnect subsystem, but some applications needed improvements to ensure their communication pattern does not cause congestion due to poor MPI task placement. The IME team led by Sanjay Kale worked with PRAC several teams to understand application communication patterns and develop a new utility to automatically provide a better MPI task placement strategy using a library the team developed that is discussed below. The library collects the necessary information during a shorter instrumented run and then produces an MPI rank reorder file that works with the Cray MPT MPI library to place the tasks optimally for on-node communication and for reduced off-node congestion. A 2.2x and a 1.2x speedup in time-to-solution improvements were obtained for 65K ranks of the MILC su3_rhmc code of the Robert Sugar PRAC team and the 65K ranks of the PSDNS-CCD code of the P.K. Yeung PRAC team were obtained, respectively, simply by linking to the library. Similar improvements were obtained with the QBOX community code. At the moment, the libtopomapping library is available on the Blue Waters system, but it is being generalized to work on other HPC systems.

File IO is often the last aspect of application performance tuning to be addressed as it often is not the primary bottleneck until the scale of the code is increased. The PRAC teams were able to partner with two IMEs to improve application IO: the Bill Gropp IME team that developed a new scalable IO library (meshio) and the HDFGroup IME team led by Gerd Heber. The P.K. Yeung PRAC team worked with Gerd Heber of the HDFGroup to improve their existing HDF checkpoint code in their PSDNS-CCD code and reduce the IO time step from being 3.3x slower than a compute time step to being 3.75x faster than a compute time step at 1,024 XE nodes. A similarly impressive reduction in IO timewas obtained by the Ivan Soltesz PRAC team, working with the HDFGroup on the Neuron code. The Gropp IME team improved IO performance by more than 20x for the MILC IO pattern using their meshio

library, discussed below, turning IO from the major to a minor contributor of runtime for the application.

"In our case, the PAID program helped us to understand the issues related to IO in massively parallel cosmology simulations. In the end, we codified these understandings into a reusable computer software. The PAID program is fruitful in our case."

- Tiziana Di Mateo Carnegie Mellon University

Only a fraction of the successes are documented above. A complete report is in preparation for the NSF with details of the work done to enable these improvements. Final reports from the participating PRAC and IME teams will be provided in the report to NSF.

PRODUCTS

Several new libraries and new methods in an existing library are the direct results of the PAID program. The rank placement and topology IME team led by Sanjay Kale produced a new application communication topology library called TopoMapping that enables users of the library run their application nearly optimally by reducing contention and congestion for the shared links on the Blue Waters 3D torus network. As mentioned above, the speedups in time-to-solution for several applications that used with the library range from 1.2x for P.K. Yeung's 3D pseudo-spectral DNS code to a 2x speedup for Bob Sugar's MILC-based su3_rhmc application without having to know detailed knowledge of the communication graphs. TopoMapping is available on Blue Waters and is documented on the Blue Waters portal at <https://bluewaters.nsa.illinois.edu/topology-considerations#topomapping>.

In addition, a new IO library called meshio (available at the GitHub site— <https://github.com/oshkosh/meshio>) was developed by the IO IME team led by Bill Gropp. The library improved IO performance by more than 20x with a MILC IO pattern and by 19x for the Plascomcm application, reducing IO time from 95% of the runtime to 5%. Finally, the SpiralGen FFT team led by Franz Franchetti implemented an improved recursive kway Alltoall algorithm for 3D FFTs in their ACCFFT library, with speedups of 2x to 4x, depending on node count. The library is available at the SpiralGen GitHub site (<https://github.com/spiralgen/>). All libraries developed as part of the PAID program are available either from the developers or on the Blue Waters system.

SELF-DIRECTED EFFORTS

Two teams took advantage of the opportunity to do self-directed work or to work with an unfunded IME not part of the eight dedicated IMEs. One team enabled a new real-space method in their pseudo-spectral turbulence code by working with a collaborator who specialized in real-space methods. The resulting hybrid code, developed while working with other IMEs on other improvements, performed 2.5x faster than the initial pseudo-spectral code over a wide range of processors. The second team

initially worked with two IMEs, but after an in-depth analysis by the IMEs, the second team decided the methods would not work with data constructs already designed in the GPU-accelerated code. The investigations prompted a redesign of the existing application, resulting in a 2x faster code compared to the original code.

TRAINING AND OUTREACH

From the beginning of the PAID program, the IMEs were expected to contribute to training events by providing instructional material or other content, by presenting material, and by making the material available for self-paced use. Over the course of the program, IMEs participated in two of the annual Blue Waters symposia by providing helpful nuggets from their respective areas of expertise. IMEs also participated in Blue Waters advanced workshops and in Blue Waters webinars. Content from all these training events is available on the Blue Waters portal for use by new Blue Waters teams or by the broader community.

FINDINGS

There are many ways to look at the success of the PAID program. The primary finding that best describes the success of the program is the new or additional science made possible by the improvements to the applications as briefly discussed above. An equally important finding is the return-on-investment (ROI) by comparing the monetary investment in expertise and application development to the return of new or additional science in terms of use of HPC resources. A final, in-depth report to the NSF and others is being prepared that will contain estimates for both the additional science made possible by improved application performance and time to solution and the ROI from the cost of the node-hours that would be used by the less efficient implementations. The ROI based solely on historical use on Blue Waters should be viewed as providing a lower bound, which initial analysis indicate the ROI will be more than 2.5x. Because the application improvements done were not specific to Blue Waters, and all the PAID teams use other HPC systems, the ROI and impact become significantly greater when future science on Blue Waters as other HPC systems are taken into account.

Additional findings involve the practices needed to make sure a program like PAID is successful. There have been other programs that supported collaborative efforts among domain scientists or domain experts and computer or computation scientists that have been assessed with mixed results. However, the PAID program's implementation of required baseline performance recording, progress tracking through milestones and deliverables, and periodic program-wide reporting of progress differentiate what is needed to make a program like this truly successful.

The PAID program has also served as a workforce development program. The IMEs' postdocs, graduate and undergraduate students worked with the PRAC teams' postdocs and graduate students, enabling substantial exchange of knowledge and expertise. The science team participants had to review their code

in conjunction with the method experts, with both learning about the requirements of the other.

"... we also integrated our Blue Waters PAID activity into the graduate curriculum at University of Utah. A Fall 2016 graduate seminar on Performance Optimization required the students to work in small groups to improve the performance of existing HPC codes."

- Mary Hall University of Utah

CONCLUSION

The PAID program should serve as a model for future sponsored programs that enable deep collaborative application development by domain experts and computer and computational scientists with a quantifiable baseline, statements of work with milestones and deliverables, work plans, tracking, and reviews. The program has shown it is possible, with proper management and oversight, to quantifiably improve a broad set of applications over a wide range of technologies in a timely manner so that applications can keep up with the rapid change in architectures, programming models, and software frameworks.

"The PAID program turned [out] to be extremely beneficial for our research effort. The development of efficient software for modern and emerging high-performance computing architectures is a highly complex process that requires close collaboration of nuclear physicists, applied mathematicians, and computer scientists. We believe programs like PAID would be extremely helpful in [the] transition of computational nuclear physics toward exascale architectures that will emerge in the next several years."

- Jerry Draayer Louisiana State University

BLUE WATERS EDUCATION, OUTREACH AND TRAINING ALLOCATIONS

Blue Waters education allocations are available to the national community to support the development of a national workforce with expertise in petascale computing. Blue Waters has committed one percent (the equivalent of 60 million core-hours) of the system's computational capacity each year for educational projects, including seminars, course-ware development, courses, workshops, institutes, internships, and fellowships.

Education and training projects utilizing these allocations have ranged from one-day workshops, to week long institutes, to semester courses, as well as support for internships and fellowships. Applications for education allocations are accepted throughout the year. Applicants are asked to apply at least a month in advance of when the allocations are needed, to allow time to process the request including setting up accounts on the system.

But the Blue Waters Educational, Outreach and Training allocations are more than just computational time and storage. A project can also get special scheduling and reservations so they can do in-class, interactive exercises and assignments. Blue Waters also provides technical support to instructors, and Blue Waters staff that serve as "Points of Contact" for these projects.

To date, there have been more than 200 approved education, outreach, and training projects from organizations across the country. These allocations have directly benefitted over 3,700 individuals learning about different aspects of computational and data-enabled science and engineering at more than 160 institutions, including 41 institutions in EPSCoR jurisdictions and at 14 Minority Serving Institutions.

A sampling of the projects and programs that have benefitted include: Blue Waters Graduate Fellows; Blue Waters Interns; Parallel Computing Summer Research Internships sponsored by Los Alamos National Laboratory; Southern California Earthquake Center (SCEC) undergraduate studies for understanding seismic hazards; NASA's Oklahoma EPSCoR Summer Research Program; University of Illinois CS420/CSE402/ECE492 Parallel Programming course for advanced undergraduates and early graduate students; and the recent Scaling to Petascale Institute that engaged over 500 people participating from over 23 locations and via YouTube Live.

Each project is requested to complete a final report documenting the impact of the education allocation. Below are examples of the benefits of these allocations to the participants as reported to us by the project PIs:

- *"The Blue Waters system gave the students access to a system with much larger scaling capabilities than they typically have access to. In particular, one team was able to scale their code to hundreds of nodes, which would be impossible at their current institutions."*

- *"Blue Waters in particular supported our efforts to replicate at-scale production environments and the opportunities and challenges these present to users."*

- *"The sheer scale of the resource is making it possible for us to do combined analysis on more data sets than has been possible on any other resource. This makes us especially anxious to implement parallelism to see how much speed-up can be accomplished with improved approaches."*

- *"Students had a context where they could learn about the significance of parallelism and that there are problems in nature such as coupled n-body problems that are ideally suited for parallel programming."*

- *"This project introduced many of the students to high performance computing for the first time—setting up batch scripts, requesting interactive nodes, running remote simulations, and using gpu-accelerated programs."*

- *"The computing resources provided by Blue Waters enabled many of the students to carry out their projects for the course in which they used simulations to elucidate a biophysical problem related to their own research."*

- *"The use of Blue Waters enabled students to understand issues of programming at scale (up to 128 nodes [4,096 core equivalents]) and to tune their codes in a dedicated environment where performance is roughly repeatable."*

If you have undergraduates or professionals who are learning about high performance computing, or experienced graduate students and researchers who want to expand their petascale knowledge and skills, the Blue Waters project welcomes requests for educational allocations for preparing current and future generations to advance discovery in all fields of study.

Additional information about education allocations is available at <https://bluwaters.ncsa.illinois.edu/education-allocations>.

If you have questions please feel free to contact the Blue Waters staff at bw-eot@ncsa.uiuc.edu.

SPACE SCIENCE

AEROSPACE ENGINEERING

ASTRONOMY

ASTROPHYSICS

COSMOLOGY

HELIOPHYSICS

PLASMA PHYSICS

SATELLITE ELECTRIC PROPULSION

- 18 *Simulating Two-Fluid MHD Turbulence in Star-Forming Molecular Clouds on the Blue Waters System*
- 20 *Supermassive Black Holes at the Cosmic Frontier*
- 22 *Theoretical Astrophysics and Data Analysis*
- 24 *Detection of Gravitational Wave Sources in Dense Stellar Environments*
- 26 *Deep Neural Networks to Enable Real-Time Multimessenger Astrophysics*
- 28 *Toward Robust Magnetohydrodynamic Simulations of Galaxy Cluster Formation*
- 30 *Exploring the Nature of Exploding Massive Stars with High Resolution*
- 32 *Modeling Plasma Flows with Kinetic Approaches using Hybrid CPU-GPU Computing*
- 34 *Three-Dimensional Nature of Collisionless Magnetic Reconnection at Earth's Magnetopause*
- 36 *Transformative Petascale Particle-in-Cell Simulations*
- 38 *Mini-disk Dynamics about Supermassive Black Holes Binaries*
- 40 *Realistic Simulations of the Intergalactic Medium: The Search for Missing Physics*
- 42 *Simulating Galaxy Formation Across Cosmic Time*
- 44 *3D General-Relativistic Radiation-Hydrodynamic Simulations of Core-Collapse Supernovae*
- 46 *Modeling Physical Processes in the Solar Wind and Local Interstellar Medium with a Multiscale Fluid-Kinetic Simulation Suite*
- 48 *Unified Modeling of Galaxy Populations in Clusters*
- 50 *Kinetic Simulations of Large-Scale Plasma Turbulence*
- 52 *GPU-Accelerated Adaptive Mesh Refinement*
- 54 *Magnetorotational Collapse of Supermassive Stars: Black Hole Formation, Gravitational Waves, and Jets*
- 56 *GPU-Accelerated Simulations: Black Holes, Spaghettified Stars, and Tilted Disks*
- 58 *Advanced Space Weather Modeling*
- 60 *3D Simulations of *i*-Process Nucleosynthesis*

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

Allocation: GLCPC/275 Knh
PI: Dinshaw S. Balsara¹
Co-PI: Alexandre Lazarian²
Collaborator: Blakesley Burkhart³

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

EXECUTIVE SUMMARY

We are at the threshold of a new data-rich and simulation-rich era in star-formation studies. The question of how stars form is fascinating in itself and has a great impact on several other areas of astrophysics. There is a general consensus that a predominant amount of star formation in our galaxy takes place in molecular clouds, and specifically in giant molecular clouds (GMC). The molecular cloud material consists of a highly turbulent, very weakly ionized, strongly magnetized plasma. It is within such a plasma that we observe protostellar cores—the progenitors of the stars that will eventually form. The present project is designed to simulate the two-fluid turbulence that regulates star formation. This turbulence, and the nonideal physics it gives rise to, is fundamentally different from single-fluid magnetohydrodynamic (MHD) turbulence. The project aims to study these differences.

RESEARCH CHALLENGE

Our understanding of the star formation process requires advanced observational capabilities. Consequently, NASA has made multi-million-dollar investments in the HAWC+ instrument aboard the SOFIA airborne observatory with the specific goal of understanding the turbulent nature of star-forming clouds. At the same time, high-resolution simulations that include the

appropriate physics of GMCs are also of critical importance. The PIs of this proposal are theorists who are participating in a multi-year funded NASA observational proposal to obtain observational data associated with turbulence in the Perseus GMC.

The PIs have also done the leading simulations of two-fluid MHD simulations on a range of XSEDE resources [1–3]. The partially ionized fluid in a GMC is best modelled as a neutral fluid and as an ionic fluid that interacts with the magnetic field. Such simulations are extremely CPU-intensive, and only simulations with very limited resolution (512^3 zones) are in hand. At the present resolution, we will be unable to match the observations from HAWC+. With current simulations, NASA’s investment in HAWC+ will be in vain because the detailed match between simulations and observations will not be possible. The work on this newly funded grant will rectify this situation.

The most compelling motivation for understanding two-fluid, ambipolar diffusion-mediated turbulence comes from recent observations. Observations of differences in the linewidths between neutral and ionized tracers suggest that the dissipation of turbulence from ambipolar diffusion sets occurs on scales smaller than 0.0018 parsecs in M17 [4]. Fig. 1a from [4] shows observed velocity dispersions as a function of length for the HCN molecule (black) and the HCO⁺ ion (red). The dashed lines trace

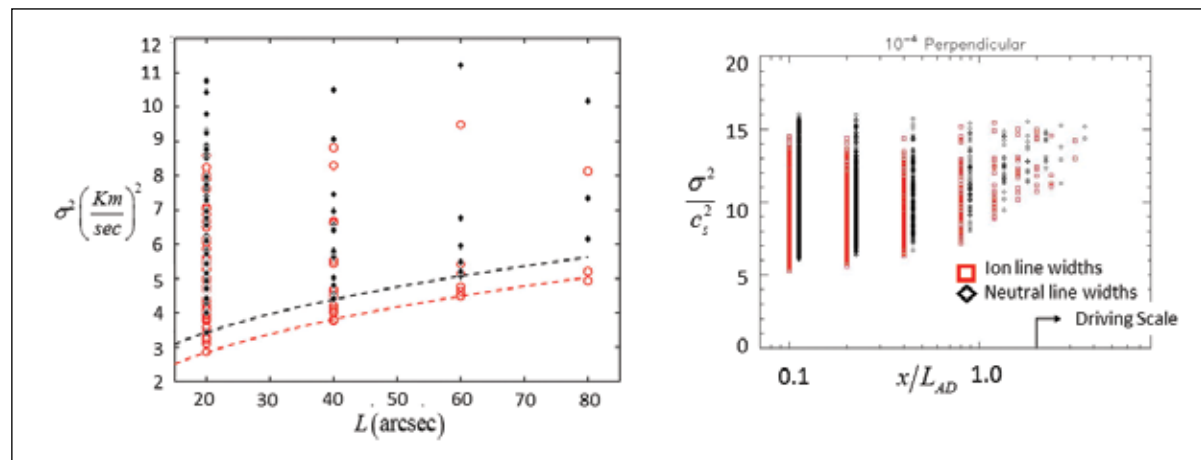


Figure 1a [4]: Observed velocity dispersions as a function of length for the HCN molecule (black) and the HCO⁺ ion (red). The velocity dispersion in the ions is lower than that in the neutrals. Figure 1b [2]: Shows simulated linewidth-size relationship from simulations.

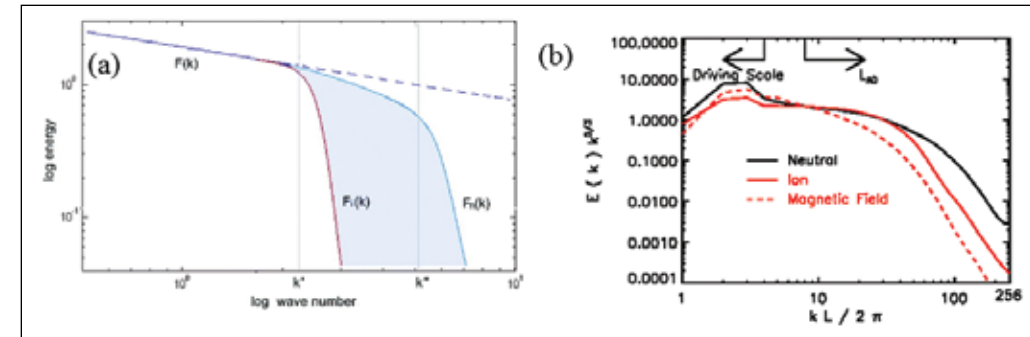


Figure 2a [4]: Shows the scaled power spectrum in the ions (subscript i) and neutrals (subscript n) needed for understanding Fig. 1. Figure 2b: Shows the velocity spectrum of the neutrals (solid black line), velocity spectrum in the ions (solid red line), and magnetic spectrum (dashed red line) from our simulations.

turbulent spectra. The velocity dispersion in the ions is smaller than that of the neutrals (especially on smaller scales), providing evidence for dissipation via ambipolar diffusion on small scales. The dissipation scales found by [4] are consistent with theoretical estimates of Alfvén wave dissipation by ion-neutral drag [5]. Fig. 1b from [2] shows the simulated linewidth-size relationship for ions (red) and neutrals (black) from one of our simulations. We see that the lower envelope of the ion linewidths is indeed lower than that of the neutrals. This is very much in keeping with the observations of [4].

Li and Houde [4] have already made the connection between their data and our theory. The theory posits that on length scales that are smaller than the ambipolar diffusion length scale (shown here as L_{AD}), the MHD waves in the ions differ strongly from hydrodynamic waves in the neutrals because of the dominance of ion-neutral drag. Such MHD waves have a strongly preferred velocity fluctuation that is orthogonal to the mean field. Consequently, we see strong differences between ion and neutral linewidths when the line of sight is orthogonal to the magnetic field.

Fig. 2a from [4] sketches their scenario. It realizes that the velocity spectrum in the ions is damped on the ambipolar diffusion dissipation scales, while the velocity spectrum of the neutrals continues undamped. As a result, the linewidths in the neutrals are systematically wider than that in the ions on a range of length scales shown in Fig. 1. Other researchers [1–3] carried out two-fluid simulations and evaluated the velocity spectra in the ions and neutrals, shown in Fig. 2b. The scaled velocity spectrum of the neutrals is shown by the black solid line. This clearly lies above the scaled velocity spectrum of the ions on length scales that are dominated by ambipolar diffusion. The linewidth-size relations and simulated linewidths in the ions and neutrals were also found to be consistent with the conjecture of Li and Houde [4].

Such simulations are extremely CPU-intensive, and only simulations with very limited resolution are in hand. At the present resolution we will certainly be unable to match the observations from HAWC+. The goal of this newly funded project is to use the petascale computing power of Blue Waters to push the resolution, accuracy, and fidelity of the simulations much higher so that we can match the theory with the observations coming from NASA-funded instruments.

METHODS & CODES

The core MHD algorithms in our RIEMANN code are based on higher-order Godunov schemes. Balsara and his collaborators have been on the forefront of the effort to develop high-accuracy schemes for computational astrophysics in general and computational MHD in particular. Two-fluid methods have been described in [6–8].

RESULTS & IMPACT

This project is newly funded and is in its initial stages where large-scale simulations have been planned and are ongoing on Blue Waters. Several papers have been published by our group using lower-resolution simulations [1–3]. The new work will be a substantial improvement on our previous work in terms of resolution as well as in the details of input physics and accuracy of simulation code.

WHY BLUE WATERS

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

PUBLICATIONS AND DATA SETS

Meyer, C., D.S. Balsara, B. Burkhart, and A. Lazarian, Observational Diagnostics for Two-Fluid Turbulence in Molecular Clouds as Suggested by Simulations. *Monthly Notices of the Royal Astronomical Society*, 439 (2014), pp. 2197–2210.

Burkhart, B., et al., Alfvénic Turbulence Beyond the Ambipolar Diffusion Scale, *The Astrophysical Journal*, 805 (2015), pp. 118–126.

SUPERMASSIVE BLACK HOLES AT THE COSMIC FRONTIER

Allocation: NSF PRAC/1,750 Knh

PI: Tiziana Di Matteo¹

Co-PIs: Yu Feng², Rupert Croft¹, Ananth Tenneti¹, Aklant Bhomwick¹

Collaborator: Steven Wilkins³

¹Carnegie Mellon University

²University of California, Berkeley

³University of Sussex

EXECUTIVE SUMMARY

The BlueTides (BT) cosmological simulation is unique because it directly probes the range of scales (masses and epochs) of galaxies that are likely to be discovered in the near future (for example, by the James Webb Space Telescope (JWST), the successor to the Hubble).

In this project, we have extended the BT simulation to cover the evolution of the first billion years of cosmic history (BlueTidesII). At this epoch, there are additional, already-existing observations that tell us that the first billion-solar-mass black holes (also known as quasars) had already formed. It is important to make contact with observations of quasars, which have not been discovered at these early stages of cosmic history. The extreme early growth depends on the early interplay of high cosmic gas densities and low tidal fields that shape the mode of accretion and allow this first generation of massive black holes to form at the most rapid pace. We have tracked the descendants of the first supermassive black holes in today's universe by running a new BT MassTracer pathfinder simulation that follows the same volume of BlueTides simulation all the way to $z=0$ (current universe) at a decreased resolution.

RESEARCH CHALLENGE

Computational cosmology, simulating the entire universe, represents one of most challenging applications for the era of petascale and beyond computing resources. We have successfully carried out a full-machine run on Blue Waters, the BlueTides cosmological simulation. The run was made possible through our new code, Massively Parallel (MP)-Gadget. The simulation has proven to be crucial to understanding the formation of the first quasars and galaxies from the smallest to the rarest and most luminous, and their effect to reionization. The simulation is being used to make predictions about what the WFIRST and JWST (expected to launch in 2018) will see. The largest next-generation telescopes (ALMA, JWST, WFIRST) currently planned aim to study the “end of the dark ages,” an epoch in the early universe when the first galaxies and quasars form and reionization of the universe takes place. The fundamental challenge to understanding this epoch is that extremely large volumes need to be simulated while at the same time extremely high resolution is required. Our BT run on Blue Waters is the first and only cosmological simulation of structure formation that has run on the full Blue Waters machine.




Figure 1: The environment of the most massive disk galaxy (top row) compared to the most massive black hole and host galaxy (bottom row) at different scales. The orange ellipses and arrow illustrate the tidal field components in the two cases. Tidal fields stretch material along one direction and compress material along the perpendicular.

In the last year, we have extended the BT simulation with an unprecedented volume and resolution to cover the evolution of the first billion years of cosmic history (BlueTidesII). The goal is to increase significantly the scientific impact to the research community of this calculation. At this later epoch, there is a wealth of additional existing observations that we would be able to compare to directly validate our models.

METHODS & CODES

The run was made possible through our new cosmological hydrodynamic simulation code Massively Parallel (MP)-Gadget. Recent radical updates to the code efficiency, and also to the smoothed-particle hydrodynamics formulation and star formation modeling, mean that we are able to meet the challenge of simulating the next-generation space telescope fields and effectively use the full Blue Waters machine. We have extended the BT run, which has an unprecedented volume and resolution, to cover the evolution of the first billion years of cosmic history.

RESULTS & IMPACT

We have run BT forward toward the epoch of observed quasars. The growth of the most massive black holes in the early universe, consistent with the detection of highly luminous quasars when the universe is less than a billion years old, implies sustained, critical accretion of material to grow and power them. It is still uncertain which conditions in the early universe allow the fastest black hole growth. Large-scale hydrodynamical cosmological simulations of structure formation allow us to explore the conditions conducive to the growth of the earliest supermassive black holes.

We used BlueTides to follow the earliest phases of black hole critical growth. At a few hundred thousand years after the Big Bang, we have shown that the most massive black holes approach 400 million solar masses. Examining the large-scale environment of hosts, we find that the initial tidal field is more important than overdensity in setting the conditions for early black hole growth. In regions of low tidal fields gas accretes “cold” onto the black hole and falls along thin, radial filaments straight into the center, forming the most compact galaxies and most massive black holes at earliest times. We have pushed the evolution of this large volume some 50 million years further and we have seen that this initial exponential growth of the first supermassive black holes is quenched even at redshift 7. This is exciting; we are starting to see effects of so-called feedback, as there are already hints in observations that this process is in action even in the first quasars.

To explore the question of where the most massive early-forming quasars are today, we have run a dark matter-only cosmological simulation that has the same volume and initial conditions of the BTMassTracer. The lower resolution allowed us to run all the way to the present universe. The results are extremely interesting: The most massive black holes that form in the early universe are not the most massive ones today and are not found in the most massive galaxy clusters. With BTMassTracer, we have traced the descendants of this first generation of supermassive black holes

and found that they are in moderately massive galaxies today. In some way, the first black holes have a very fast initial growth phase, but other black holes that assemble later are likely to grow in higher-density regions. The first massive black holes are not in very special places in the overall density field of the universe.

WHY BLUE WATERS

From the simulation/theoretical perspective, large-scale uniform volume hydrodynamic simulation of the high-redshift universe is a problem that is perfectly suited to the largest modern petascale facilities such as Blue Waters. This is because it is now feasible to run memory-limited computations with the resources that computer time allocation panels are able to grant, and so reach unprecedented volumes and resolutions.

PUBLICATIONS AND DATA SETS

Feng, Y. and C. Modi, A fast algorithm for identifying Friends-of-Friends halos. *Astronomy and Computing*, 20 (2017), pp. 44–51.

Di Matteo, T. and Croft, R.A.C., Feng, Y. and Waters, D. and Wilkins, S. (2017). The origin of most massive black holes at high- z : BLUETIDES and the next quasar frontier.

Waters, D., et al., Forecasts for the WFIRST High Latitude Survey using the BlueTides Simulation. *Mon Not R Astron Soc*, 463:4 (2016), pp. 3520–3530.

Waters, D., et al., Monsters in the dark: predictions for luminous galaxies in the early Universe from the BLUETIDES simulation. *Mon Not R Astron Soc Lett*, 461:1 (2016), pp. L51–L55.




Figure 2: The descendant in today's universe of the massive black hole in BlueTidesII (shown in Fig.1). The images show the dark matter density in a slice through the entire BTMassTracer simulation. The intensity and color of the pixels is representative of the density.

THEORETICAL ASTROPHYSICS AND DATA ANALYSIS

Allocation: Director Discretionary/50 Knh

PI: Gilbert Holder¹

Collaborators: Yuuki Omori², Yashar Hezavehe³, Warren Morningstar³, Neal Dalal¹, Joaquin Vieira¹, South Pole Telescope Collaboration

¹University of Illinois at Urbana-Champaign

²McGill University

³Stanford University

EXECUTIVE SUMMARY

We have been investigating different avenues for mapping dark matter in the universe, on both the very largest scales (extending to the cosmological horizon) and the smallest scales (>1,000x smaller than the typical size of galaxies).

We have produced the highest-fidelity map of dark matter on cosmological scales using data from the South Pole Telescope, a 10-meter (m) telescope at the Geographic South Pole that is mapping the cosmic background microwave temperature and polarization fluctuations.

We have also started a search for clumps of dark matter on small scales in galaxies that are acting as gravitational lenses of background star-forming galaxies. We have previously detected one such clump in a galaxy and we are now extending this search to a larger sample of galaxies.

RESEARCH CHALLENGE

Roughly 80% of the mass in the universe is in the form of dark matter, which appears to be a particle beyond the standard model of particle physics. It barely (if at all) interacts with particles that we have measured in the lab, other than through gravity, which is how we infer the existence of dark matter. The problem is well-studied on scales of galaxies and clusters of galaxies; our work is searching for clues about its nature by looking on much larger and smaller scales.

If dark matter is not actually “matter” at all but is instead a first crack in the theory of general relativity, it could be expected that the very largest scales may show anomalous behavior, while interesting self-interactions or other possible properties of dark matter could be expected to affect how it clumps together on the smallest scales.




Figure 1: Strongly lensed distant star-forming galaxy SDP.81 as imaged by ALMA (red), with Hubble Space Telescope image of foreground lens galaxy (blue) in which dark matter is being studied at high resolution [2]. The white dot shows the size and location of the clump of dark matter that has been inferred using a maximum likelihood fit to the ALMA emission.




Figure 2: Map showing the integrated mass on the sky over 2,500 square degrees, obtained using intensity maps from the South Pole Telescope [1].

METHODS & CODES

We used the gravitational deflection of light to map out where the mass is in two different regimes, using either the cosmic microwave background fluctuations as the light source or star-forming galaxies observed at millimeter (mm)-wavelengths.

Using the cosmic microwave background, which is the most distant light we can measure in the universe, we make mass maps on extremely large scales by looking for subtle systematic correlations in the temperature and polarization fluctuations [1]. The South Pole Telescope, a 10-m mm-wave experiment at the Amundsen–Scott South Pole Station is mapping the cosmic microwave background at high resolution and sensitivity providing the experimental data for this work. Gravitational lensing changes the local statistical properties of the maps in a way that is sensitive to the local mass density, so we use custom-made second-order statistical estimators to make a map of the mass density.

Some distant star-forming galaxies are well-enough aligned with foreground galaxies that gravitational lensing can produce multiple images, which travel through different parts of the foreground galaxy. Small-scale structure from dark matter will be different for each of the images, allowing a search for small-scale features in the dark matter. We have developed a code to perform such a search, and successfully discovered substructure in a single galaxy [2] using data from the Atacama Large Millimeter Array, and we are now applying a similar analysis to new data on five more galaxies. The image used to discover the dark matter substructure in the first galaxy studied is shown in Fig. 1.

RESULTS & IMPACT

Using Blue Waters, we were able to run Monte Carlo simulations of the South Pole Telescope data analysis pipeline at a level that allowed us to make the map shown in Fig. 2: Over 2,500 square degrees, we have a map of how much mass there is at every point on the sky. “Lensing convergence” is a weighted average of the mass density along the line of sight to the source, and since the source (the cosmic microwave background) is at the edge of the observable universe, we get a complete census of the mass in the universe. This map is being used to compare with galaxy surveys to better understand the relationship between visible mass and dark matter. Going forward, the South Pole Telescope has a new camera that is roughly 10x faster at mapping the sky, which will enable both higher precision and higher resolution in this mass map within a year or two.

The search for small-scale substructure in the dark matter is ongoing. With the data currently being analyzed, we expect to be able to confirm or rule out dark matter candidates that have either a large thermal velocity (“warm dark matter”) or dark matter composed of ultra-light particles that have a quantum mechanical fuzziness that extends to macroscopic scale (“fuzzy dark matter”).

WHY BLUE WATERS

The combination of computing power and support staff for Blue Waters made this the premier option for performing this research.

DETECTION OF GRAVITATIONAL WAVE SOURCES IN DENSE STELLAR ENVIRONMENTS

Allocation: Director Discretionary/500 Knh

PI: Eliu Huerta¹

Collaborators: Bhanu Agarwal¹, Daniel George¹, Roland Haas¹, Wei Ren¹, Hsi-Yu Schive¹, Prayush Kumar², Harald Pfeiffer², Michael Boyle³, Lawrence Kidder³, Daniel Hemberger⁴, Mark Scheel⁴, Bela Szilagyi⁴, Tony Chu⁵

¹National Center for Supercomputing Applications

²Canadian Institute for Theoretical Astrophysics

³Cornell University

⁴California Institute of Technology

⁵Princeton University

EXECUTIVE SUMMARY

The direct detection by the twin LIGO (Laser Interferometer Gravitational-Wave Observatory) detectors of gravitational waves from merging black holes has ushered in a revolution in astrophysics [1,2]. LIGO's detection campaigns are enabling the construction of catalogs of gravitational wave sources to enable accurate censuses of the mass and angular momentum distribution of black holes and neutron stars. These studies will provide new and detailed information about the formation, evolution, and environments in which compact objects reside. In particular, the detection of gravitational waves from eccentric compact binaries will provide the cleanest signature of compact object populations in dense stellar environments [3]. To detect these events, we introduce the first waveform model in the literature that reproduces the features of eccentric compact binary coalescence throughout merger. To validate this model, we used numerical relativity simulations performed with the Einstein Toolkit on the Blue Waters supercomputer.




Figure 1: Recovery of eccentric, nonspinning signals (represented by points in the figure) using a template bank of 1.5 million quasi-circular, spin-aligned waveforms. The green and black stars represent GW150914 and GW151226. The Fitting Factor is computed by maximizing the overlap, or inner product, between every single eccentric signal and the 1.5M templates in the bank.

RESEARCH CHALLENGE

Gravitational wave (GW) observations provide insights into the astrophysical sources that generate them and can be used to map the structure of spacetime in the vicinity of compact binary systems (BHs) in the galactic cluster M22 led to the development of more accurate N-body algorithms to explore the formation and detectability of binary black holes (BBHs) formed in globular clusters [4]. These studies indicate that about 20% of BBH mergers in globular clusters will have eccentricities $e_0 \sim 0.1$ when they first enter the aLIGO band at 10Hz, and that $\sim 10\%$ may have eccentricities $e_0 \sim 1$ [5]. BBHs formed in the vicinity of supermassive BHs may also merge with significant residual eccentricities. aLIGO is uniquely positioned to enhance the reach of GW astronomy by targeting these events. At the time of the detection of the first GW transient, GW150914, there was no waveform model available to describe the evolution of eccentric compact binary coalescence (eCBC) from early inspiral through merger and ringdown [6]. Furthermore, we require new data analysis techniques to capture the imprint of eccentricity on GW signals [7].

In order to detect and characterize eCBC with aLIGO, we introduce an inspiral–merger–ringdown (IMR) waveform model that reproduces the dynamics of state-of-the-art, nonspinning, quasi-circular waveform models, and the dynamics of eccentric numerical relativity simulations.

METHODS & CODES

To construct our model, we derived higher-order post-Newtonian results for eCBC and then combined these results with state-of-the-art results from the self-force formalism and BHPT. This hybrid formalism is the first of its kind to reproduce CBC accurately for compact binaries with mass-ratios between 1 and 10.

To construct the merger phase of our model, we assume that the system circularizes prior to the merger event. We have found that this assumption covers the class of eCBC that aLIGO can detect using continuous waveform signals. Highly eccentric signals that do not fall into this category can be searched for in aLIGO data using burst search algorithms, which capture very well the burst-like nature of highly eccentric eCBC.




Figure 2: For an equal mass binary black hole system with initial eccentricity $e_0=0.076$ and mean anomaly $l=3.09$ at a gauge-invariant frequency value $x=0.074$, we present a direct comparison of the dynamics predicted by our model and an eccentric numerical relativity simulation.

RESULTS & IMPACT

We carried out the first systemic analysis of the effect of eccentricity on the first two GW transients detected by aLIGO. Fig. 1 indicates that the first two GW transients could have non-zero eccentricity at a GW frequency of 15Hz and still be misclassified as quasi-circular systems due to the restricted sensitivity of aLIGO below 25Hz. Thus, once aLIGO attains design sensitivity and is able to probe lower frequencies, we will be in a better position to accurately extract the signatures of eccentricity in GW transients.

This study is the first of its kind to show that the effect of eccentricity in waveform signals cannot be mimicked by spin corrections. Indeed, the signal manifold described by GWs with eccentricity $e_0 > 0.1$ at 15Hz is orthogonal to spin-aligned GW signal manifolds.

Another important result of this study is the validation of our waveform model using eccentric numerical relativity simulations. In Fig. 2 we show that our waveform model can accurately reproduce the true, accurate dynamics of these systems throughout the merger of the binary system. This is a noteworthy result because we did not use eccentric numerical simulations to calibrate our model. The exploitation of this model for upcoming

GW detection campaigns with the aLIGO detectors will enable us to confirm or rule out the existence of compact binary populations that exist in dense stellar environments.

WHY BLUE WATERS

Blue Waters enabled us to create a large catalog of numerical relativity simulations, which required thousands of node hours that we ran in parallel to sample a deep region of parameter space. No other resource but Blue Waters can provide the required computational power to obtain a numerical relativity catalog of this nature in a timely manner. Furthermore, the Einstein Toolkit has been extensively used on Blue Waters since it began operations. NCSA Gravity Group members are key developers and maintainers of this software.

PUBLICATIONS AND DATA SETS

Huerta, E.A., et al., Complete waveform model for compact binaries on eccentric orbits. *Physical Review D*, 95 (2017), DOI: 10.1103/PhysRevD.95.024038.

DEEP NEURAL NETWORKS TO ENABLE REAL-TIME MULTIMESSENGER ASTROPHYSICS

Allocation: Director Discretionary/1,000 Knh

PI: Eliu Huerta¹

Collaborator: Daniel George¹

¹National Center for Supercomputing Applications

EXECUTIVE SUMMARY

The detection of gravitational waves has opened up a new spectrum of observation into the Universe. The LIGO flagship detection pipelines target a specific class of binary black holes that generate burst-like gravitational wave signals. In order to capitalize on the unique opportunities that gravitational wave astrophysics presents for new discoveries, it is necessary to extend the depth of gravitational wave searches to extract signals that currently go unnoticed with these pipelines. To address this issue, we introduce Deep Filtering: a new method that combines two deep convolutional neural networks for classification and regression to detect and characterize signals much weaker than background noise. We show that Deep Filtering significantly outperforms conventional machine learning techniques and enables the detection of new classes of gravitational wave signals that go unnoticed with existing detection algorithms.

RESEARCH CHALLENGE

Advanced LIGO (aLIGO) detection algorithms have confirmed the existence of a particular class of gravitational waves (GWs) using a 3D search: quasi-circular, spin-aligned binary black holes

(BBHs). Extending these searches to target the full 8D parameter space of astrophysically motivated sources presents outstanding computational challenges [1, 2, 3].

Multimessenger searches of electromagnetic (EM) and astroparticle counterparts of GW transients rely on accurate and low-latency GW analyses, which at present take from days to months to finish. To overcome these limitations, we introduce Deep Filtering, a deep learning algorithm to directly process aLIGO data, which outperforms other machine-learning methods, and is many orders of magnitude more computationally efficient than matched filtering for both detection and parameter estimation.

METHODS & CODES

We consider a 2D parameter space that describes non-spinning BHs on quasi-circular orbits, with masses between 5 and 75 solar masses, and mass-ratios of 1 to 10. We generate our data sets using the surrogate waveform family introduced in [4]. The mass-ratio values of the BBH signals are between 1 and 10 in steps of 0.1 for training, and intermediate values for testing.

We superimposed different realizations of Gaussian white noise on top of the signals over multiple iterations, thus amplifying the

size of the data sets. We then standardized the inputs to have zero mean and unit variance. The final training sets at each signal-to-noise ratio (SNR) contained about 100,000 time-series vectors produced from 4,000 templates of BBH signals by adding multiple batches of noise and shifting in time. The validation and testing sets at each SNR contained about 25,000 elements, produced from 586 clean templates and different noise realizations.

We designed simple deep neural networks (DNNs) from the ground up, since deep learning alternatives to matched filtering had not been proposed before. We tested around 80 configurations of DNNs, and found that a design with three convolutional layers followed by two fully connected layers yielded the best results. We utilized the neural network functionality in the Wolfram Language, based on the MXNet framework, which utilizes the CUDA deep learning library for acceleration with NVIDIA GPUs. We used the ADAM method as our learning algorithm.

We developed a new strategy to improve the performance and reduce the training times of the DNNs. By starting with training inputs having a high SNR of less than 16, and then gradually increasing the noise in each subsequent training session until a final SNR of 0.06, we observed that the performance of prediction can be quickly maximized for low SNRs. Our algorithm can be applied to a continuous data stream using a 1-second sliding window with offsets of 0.2 seconds.

RESULTS & IMPACT

We trained our classifier to achieve 100% accuracy with zero false positives for signals with a SNR > 0.36. For comparison, we trained standard implementations of all commonly used machine learning classifiers along with the DNNs on a training set of 8,000 elements. Unlike DNNs, none of these algorithms was able to directly handle raw noisy data, as shown in Fig. 1.

Our predictor successfully measured binary component masses given noisy GW signals that were not used for training, with an error of the same order as the spacing between templates for a SNR > 1. Although our initial goal was to create a pipeline for only non-spinning, quasi-circular BBH signals, we tested our DNNs using moderately eccentric simulations that we obtained using the Einstein Toolkit on the Blue Waters supercomputer. The classifier detected all these signals with nearly the same rate as the original test set (with over 99.7% mean accuracy for SNR > 0.36 and 100% accuracy at SNR > 0.5). The predictor was able to estimate the component masses of our eccentric simulations for SNR > 0.25 with a mean relative error less than 13%, 19%, 32%, and 34% for mass-ratio 1, 2, 3, and 4, respectively. This result is very encouraging, since these types of signals go unnoticed with aLIGO detection [5].

Our DNNs are only 2MB in size each. The average time taken for evaluating them per input of 1 second duration is approximately 6.7 milliseconds and 135 microseconds using a single CPU or GPU, respectively. For comparison, we estimated the evaluation time for time-domain matched filtering with the template bank of clean signals used for training; the results are shown in Fig. 2. Our

extremely fast inference rate indicates that real-time analysis can be carried out with a single computer even with DNNs that are significantly larger, and can be trained over much bigger template banks of signals.

WHY BLUE WATERS

Blue Waters enabled us to create a large catalog of eccentric numerical relativity simulations, which required thousands of node hours that we ran in parallel to sample a deep region of parameter space. No other resource but Blue Waters can provide the required computational power to obtain a numerical relativity catalog of this nature in a timely manner.

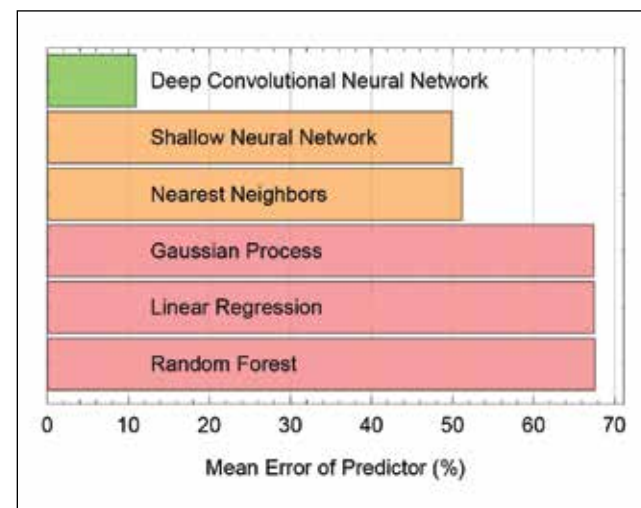


Figure 1: Mean relative error obtained by various machine learning algorithms for predicting a single parameter, i.e., mass-ratio, using a training set containing about 8,000 elements at SNR (Signal-to-Noise Ratio) = 0.36. Scaling these methods to predict multiple parameters is often difficult, whereas it simply involves adding more neurons to the final layer of neural networks.




Figure 2: For a given template bank, Deep Filtering is many orders of magnitude faster than matched filtering. The evaluation time of a Deep Neural Network is constant regardless of the size of training data, whereas for matched filtering it is proportional to the size of the template bank.

TOWARD ROBUST MAGNETOHYDRODYNAMIC SIMULATIONS OF GALAXY CLUSTER FORMATION

Allocation: GLCPC/450 Knh

PI: Thomas W. Jones¹

Co-PI: Peter J. Mendygral²

Collaborators: Krishna Kandala², Nick Radcliffe², Julius Donnert¹, Brian O'Neill¹, Chris Nolting¹

¹University of Minnesota

²Cray Inc.

EXECUTIVE SUMMARY

Our goal is to understand the laws of physics that control the dynamics of ordinary, baryonic matter during the formation of galaxy clusters. Most of that matter is very hot, very dilute plasma. A key dynamical component that results from cluster formation is turbulence in the plasma that generates magnetic fields that control the microphysical processes essential to the physical state of the plasma. Much of the important action takes place on very small physical scales, which are quite challenging to capture in simulations, especially magnetic field behaviors. The work reported here was a vital component in developing a novel, exceedingly high-performance and exceptionally scalable MHD (magnetohydrodynamics) cosmology code named “WOMBAT” that we will apply to the problem detailed above on coming exascale systems. Here, we demonstrated that WOMBAT meets our design objectives when scaled to more than 16,000 Blue Waters nodes. Initial test simulations examined behaviors of astrophysical, hypersonic plasma jets in cluster settings.

RESEARCH CHALLENGE

Galaxy clusters are knots in the cosmic web with masses that can exceed $10^{15} M_{\text{sun}}$, and sizes of several million light years. They are the largest and last structures to form by gravitational collapse from fluctuations in the Big Bang. Their formation is a critical diagnostic of cosmological theory. Galaxies are actually minor constituents of these clusters. Most of the matter is “dark matter,” whose nature remains unclear. Most of the “ordinary” baryonic matter is very hot, very dilute plasma filling the cluster. Its properties trace the current cluster dynamical state and its history. A full understanding of the dynamics of this intracluster medium plasma, or ICM, is absolutely necessary to comprehend how clusters form. Much of the important ICM action, and especially action responsible for generation of the magnetic fields that control ICM physical properties, takes place on scales that are multiple orders of magnitude smaller than the size of a cluster. Until now, simulations that could capture those small-scale actions in cosmological-scale computations, especially including magnetic field generation (which depends on activating the small-scale,

MHD dynamo), were beyond reach with existing software and high-performance computer systems. Even with petascale systems such as Blue Waters, no simulation codes existed until now with sufficient performance and scaling behavior to begin to address this research challenge. Our objective is to develop tools to resolve this problem on the coming generation of exascale systems, where the properties designed into WOMBAT will truly shine.

METHODS & CODES

To address this important astrophysics problem we have built from the ground up a novel, exceedingly high-performance and highly scalable MHD cosmology code named “WOMBAT.” The code optimizes local memory and vector performance and utilizes hybrid parallelization methods leveraging techniques pioneered at Cray that gain maximum “threading” performance within many-core nodes and MPI-RMA (message passing interface–remote-memory access) performance between nodes. The fast interconnect technology on Blue Waters is very important to our ability to test these designs. In addition, a novel strategy for decomposition of the computational domain minimizes inter-node communication and allows it to be highly asynchronous. Such features provide tremendous performance benefits that will allow WOMBAT to address exceedingly challenging problems, including cluster formation ICM dynamics down to the required physical scales.

RESULTS & IMPACT

Our primary objective on Blue Waters for this allocation was to test and tune the technologies built into WOMBAT in order to prepare to address the cosmological problems outlined above. We have succeeded very well in that. We demonstrated, for example, up to 75% weak scaling efficiency going from one Blue Waters XE6 node to 16,224 XE6 nodes with 16 floating point threads per node. As part of our team’s collaboration with Cray, the Cray Programming Environments group developed a new MPICH release, based in part on our Blue Waters results, with enhanced performance for hybrid parallel applications. This will be the default MPICH model going forward, and will considerably benefit both the Blue Waters user community and the broader high-performance community. The N-body dark matter routines in WOMBAT were still in testing at the end of the allocation period. Those tests are very promising, but we were unable to carry out planned cosmological simulations before the allocation expired. We did, however, carry out very high-resolution WOMBAT MHD simulations of hypersonic plasma jets propagating in ICM environments in order to explore their propagation behavior. The distribution of jet plasma from one such simulation is volume rendered in Fig. 1. Instabilities forming along the boundary of the light plasma cocoon are evident.

WHY BLUE WATERS

Blue Waters was essential for this project because of the importance of using a system capable of sustained petaflop

performance with very fast inter-node communications to allow proper WOMBAT performance testing. Blue Waters offers the unique opportunity to both test our strategy and to begin to carry out simulations that approach needed levels of physical fidelity.

PUBLICATIONS AND DATA SETS

Mendygral, P. J., et al., WOMBAT: A Scalable and High-performance Astrophysical Magnetohydrodynamic Code. *Astrophysical Journal Supplements*, 228:2 (2017), DOI: 10.3847/1538-4365/aa5b9c.




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

EXPLORING THE NATURE OF EXPLODING MASSIVE STARS WITH HIGH RESOLUTION

Allocation: NSF PRAC/6,000 Knh

PI: Eric J. Lentz¹

Co-PIs: Stephen W. Bruenn², W. Raphael Hix^{1,3}, O. E. Bronson Messer³, Anthony Mezzacappa^{1,3}

Collaborators: Jordi Cassanova³, Eirik Endeve³, J. Austin Harris³

¹University of Tennessee, Knoxville

²Florida Atlantic University

³Oak Ridge National Laboratory

EXECUTIVE SUMMARY

Core-collapse supernovae (CCSNe) are the violent explosions of massive stars at the end of their evolutionary lives. CCSNe play an important role in the dynamics of galaxies—injecting energy, producing and expelling heavy elements, and triggering formation of new star systems. CCSN simulation is a nexus of exotic basic physics—neutrinos, extreme gravitation, and nuclear physics. We are using 3D simulations to explore the variation in explosions and ejecta that result from the known variations in properties of massive stars (initial mass, composition, rotation, etc.) and to understand the impact of resource limitations on those outcomes. Because CCSNe are major contributors to the elemental composition of planets (including Earth), we have computed a 3D model that captures that element production in detail. Another model, limited in time and geometry, explores extremely high resolution to gauge the impact of resolution limits imposed by current supercomputers.

RESEARCH CHALLENGE

In CCSNe, the explosion is achieved by heating material in the cavity between (~200 km radius) accretion shock and the neutron star (NS) newly formed by the collapse of the star's iron core. The heating comes from neutrinos emitted from the NS, which release the gravitational binding energy of the NS and are absorbed in the cavity. This requires energy resolved (spectral) neutrino transport coupling the NS and the heating cavity. After an explosion begins (the last 80–90% of the simulation), the evolution of the nuclear isotopes in the ejecta require a reaction network. Larger networks can directly capture the evolution of more species and more accurately account for the input of energy into the expanding ejecta. In addition to the strength and composition of the explosion, CCSNe are important targets for observation by neutrino and gravitational wave detectors. The pre-collapse progenitors come in wide variety, driven largely by variations in the initial mass, rotation, and composition. To

account for this variation in the CCSN outcomes and impacts on the evolution of heavy element abundances and compare to observed CCSN properties, wide ranging sets of models must be built. We are building a modest set of 3D models that span initial stellar parameters and exploring the impact of resource-limited resolution on current and prior CCSN simulations by all groups

METHODS & CODES

We have developed the custom multiphysics code Chimera [1,2] for CCSN. It couples the hydrodynamics of a self-gravitating fluid with active nuclear reactions (burning) taking place in the ejecta, and neutrino transport to carry energy from the inner core to the heating. The hydrodynamics is solved with a dimensionally split piecewise parabolic finite volume scheme. Gravity is solved by global multipole expansion of the Poisson equation. Nuclear burning is solved by the backward Euler method for a finite sized reaction network (14 or 160 species). Neutrino transport is solved by a finite differenced implementation of energy resolved flux limited diffusion that couples all neutrino species to a comprehensive set of neutrino-matter interactions.

RESULTS & IMPACT

One simulation computes a high-resolution (1/4°) model with limited geometry (90°x90° wedge along equator) during the pre-explosive convection phase, which is compared with lower-resolution models (1/2° to 2°) computed elsewhere to understand the impact of the lower resolution on other simulations [3]. We were able to measure the development of the cascade of turbulent energy to small scales and show that it is similar to previous work [4] done using parameterized neutrino transport. We also see (Fig. 1) that the large plumes with sinking streams in the low-resolution models (resolutions similar to our pre-Blue Waters work [5] and other published results [6]) transform in the high resolution (1/4°) Blue Waters simulation to rising and sinking “clouds” of smaller-scale features. These models indicate that the 1° models of our full-geometry 3D Blue Waters models should be sufficiently resolved and allay concerns that future gains in available capability must be thrown primarily into achieving better resolution.

Another simulation follows the collapse and explosion of a lower-mass star (9.6 solar masses) with the primordial initial composition (only Big Bang hydrogen and helium) [7]. Stars at the low-end of the massive star scale have lower densities in the material around the core, so that when it is accreting onto the stalled shock, it provides less ram pressure for the neutrino-heated material in the cavity to overcome. Because it explodes more quickly with less neutrino heating, some of the outermost ejecta (Fig. 2) is neutron rich and can form rare isotopes such as calcium-48. In this 3D model, we used a larger nuclear reaction network with 160 nuclear species that allows the direct tracking of calcium-48 formation *in situ*. Previous identification of lower-mass CCSN as sites for production of calcium-48 [8] and other neutron-rich isotopes relied on post-processing with passive tracers.




Figure 1: Slice in entropy through the 90° wedge models of various angular resolutions (2° to ¼°) in pre-supernova convective phase. As resolution increases, the number of fine structures increases, but the character of the models with regions of lower entropy inflow (green) and upwelling heated material (orange) remains.




Figure 2: Slice through expanding ejecta of 9.6 solar mass model. Low-entropy (blue) material behind shock (reddish) is neutron-rich (low, orange, electron fraction “Ye” in cut-out) permitting the formation of neutron-rich calcium-48.

We are also computing 1°-resolution models of several other stellar progenitors to examine the variation in initial stellar mass. All of these models are being compared to less costly 2D models of the same and similar progenitors to better leverage the 3D results to broader applicability. For all our models, we will compute post-processed isotope production with very large (8,000 species) networks and compute neutrino and gravitational wave signals for comparison to future observations.

WHY BLUE WATERS

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

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

Allocation: Illinois/200 Knh
 PI: Deborah Levin¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Our main objective is to characterize the backflow contamination environment due to plasma created by electric-propulsion plumes, their interaction with the spacecraft environment, and neutralizer sources, using state-of-the-art high-performance petascale computations. In terms of modeling and simulation, we build on our earlier work where we have developed an object-oriented C++ Direct Simulation Monte Carlo (DSMC) code that uses adaptive mesh refinement (AMR)/octree grids to capture the vast length scales inherent in supersonic expansions to vacuum for neutral-neutral and neutral-ion collisions. A key aspect of our computational work is to take advantage of our unique, recent advances in GPU (graphics processing unit) multi-thread parallelization applied to tree-based computational strategies. Blue Waters is especially suited to this modeling since we expect to use 128 to 256 GPUs per run for the plasma plume simulations on the XK nodes. Initial results demonstrate that our approach has been successfully extended to couple DSMC and particle-in-cell (PIC) simulations on an AMR/octree grid using GPUs.

RESEARCH CHALLENGE

With the space environment becoming a home to constellations of small satellites and cube satellites, improved predictability of key surfaces of solar cell arrays and spacecraft charging in the backflow environment of chemical and electric-propulsion (EP) thrusters is crucial. Fig. 1 shows a schematic of the multiple sources and processes that must be considered. Because there are both neutral and charged species in the backflow region of an onboard EP thruster, the modeling of these highly reactive ions with thruster and ambient neutral species involves multiple time and length scales. Indirect environmental exposure of spacecraft material such as the micron-sized coatings of solar array cover glass and aluminized Mylar can cause appreciable sputtering and erosion. This is often hard to quantify and predict over the lifetime of the mission because backflow ion fluxes are about five orders of magnitude less than those due to main ion beam impingement. The well-known charge-exchange (CEX) process between xenon thruster neutrals and beam ions also occurs among plume/beam ions and ambient species such as Xe⁺ and O, although these processes are not presently modeled in spacecraft environment effect models. In addition, external hollow cathodes are a source of electrons and also emit xenon ions and neutrals. Inclusion of their plumes is essential as the electrons play an important role in the charge density distribution, which influences the generation of the slow CEX ions, and, in turn, affects the erosion of solar cell panels.

METHODS & CODES

Our new plasma modeling is an outgrowth of our DSMC code, CHAOS (Cuda-based Hybrid Approach for Octree Simulations) that was developed to study neutral flows through porous media [1]. We have adapted this approach now to include both neutral and ion species in an external electric field. Because the local mean free path for collisions is about three orders of magnitude larger than the local Debye length ($\sim 10^{-6}$ m), two linearized Morton-ordered forests of octrees (FOTs) have been implemented so that these grids can be adapted to meet these two diverse numerical criteria. To accurately model the electric field, using a Boltzmann distribution at a fixed electron temperature, or to solve Poisson's equation, a leaf node should be only one level larger than its smallest face neighbor, i.e., a "2:1 criteria." When this criterion is adhered to, any numerical discontinuities in the

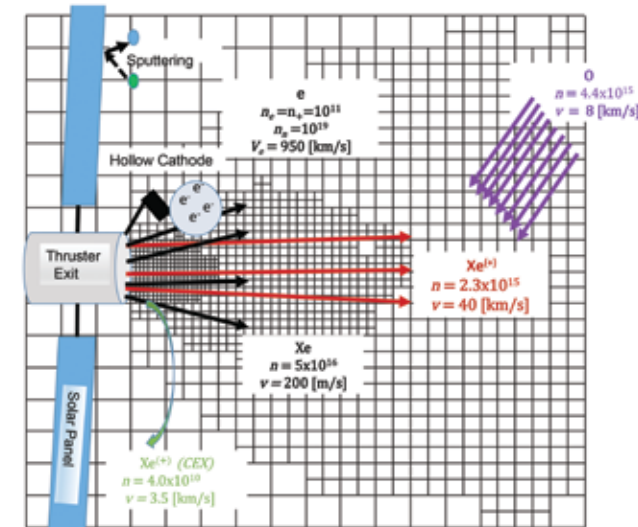


Figure 1: Key species, velocities, and concentrations (m^3) in the backflow and beam region for a spacecraft in a low- to mid-earth orbit space environment. Approximate scale is 1 m.

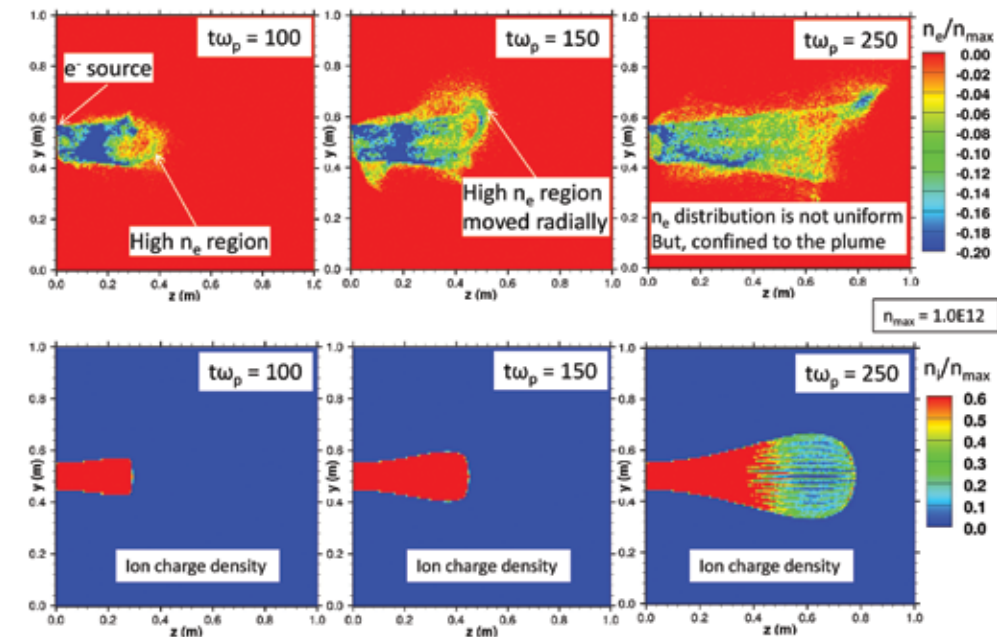


Figure 2: Unsteady behavior of ions and electrons emitted by two spatially separate sources with an electron thermal velocity, plasma frequency, ω_p , and electron to ion thermal-to-bulk velocity ratio of 5.9×10^5 m/s, 1.78×10^8 rad/s, and 1.25, respectively.

E-octree are smoothed out and we are able to maintain first-order accuracy in the gradient calculations. In our AMR/octree approach, the DSMC cells of variable size satisfy the mean-free-path and Debye length criteria but do not automatically satisfy the 2:1 ratio. Implementation of this on an AMR/octree grid is nontrivial and has been accomplished through the use of local (on a single processor) and local-global (across processors) stages.

RESULTS & IMPACT

We have recently run simulations of mesothermal, collisionless plasmas for shifted electron and ion sources [2]. We discovered very interesting unsteady plume dynamics by modeling a *shifted* electron source (in contrast to co-located electron and ion sources) as summarized in Fig. 2. The electron source, with a radius of $R_e = 0.01$ m, was placed above the ion source, with an electron thermal to ion drift velocity of 1.25. The Xe⁺ and electron number densities were 5×10^{12} and 1.0×10^{13} , ($\#/m^3$), respectively. Temperatures and bulk velocities were assumed to be 0.04 and 2 eV and 472,000 and 0 m/s, for ions and electrons, respectively. For this simulation, both ions and neutrals had the same time step of 2.8×10^{-10} s. Fig. 2 shows a summary of the interesting time-varying behavior, which is very different for the ions versus the electrons. The electron dynamics were found to be complex, and as soon as electrons are released, they accelerate toward the beam and travel *beyond* the beam front. The positively charged beam front then decelerates the electrons that have escaped and reflects them back in the reverse direction, resulting in negative electron velocities. The continual exchange between electron kinetic and potential energy results in a meandering/bouncing movement, as can be seen in the top row of the figure, as time (normalized by plasma frequency) increases. The ion beam, however, more or less propagates away

from its emitting source as time increases. Nevertheless, some diffusion of ions can be seen, particularly in the beam front at later times. In comparison to the electrons, the large ion mass keeps them moving primarily away from the source. In addition, as time progresses, the electron bouncing region is elongated with the multiple bouncing motions increasing the electron temperature. The increase in temperature is contrary to what is expected in an expanding plume and, not surprisingly, the actual electron velocity distribution obtained in the beam front region is very non-Maxwellian.

WHY BLUE WATERS

Blue Waters has allowed us to test on a large number of GPUs the potential computational savings to enable three-dimensional, fully kinetic plasma simulations. Compared to the present state-of-the-art plasma simulations, a uniform grid in 3D would require about a factor of seven more cells than our use of AMR/octree. The use of a single K20x GPU decreases the runtime by a factor of five compared to a single Interlagos processor. Very conservatively, we estimate that the octrees in combination with GPUs, decrease the total runtime by at least a factor of 10, compared to uniform grid solvers on multi-core CPUs.

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

Allocation: NSF PRAC/1,800 Knh
PI: Yi-Hsin Liu^{1,2}
Co-PIs: Michael Hesse³, William Daughton⁴
Collaborators: Ari Le⁴, Shan Wang²

¹NASA-GFSC
²University of Maryland, College Park
³University of Bergen, Norway
⁴Los Alamos National Laboratory

EXECUTIVE SUMMARY

Earth's magnetosphere shields the planet from constant bombardment by supersonic solar winds. However, this magnetic shield, called the magnetopause, can be eroded by various plasma mechanisms. Among them, magnetic reconnection is arguably the most important process. Reconnection not only allows the transport of solar wind plasmas into Earth's magnetosphere but also releases the magnetic energy and changes the magnetic topology. At Earth's magnetopause, magnetic reconnection proceeds between the shocked solar wind plasmas and the magnetosphere plasmas. The magnetosheath plasma has a typical magnetic field strength of ~20 nanoteslas and density of ~5 per cc. The magnetosphere plasma has a magnetic field strength of ~60 nanoteslas and density of ~0.5 per cc. The magnetic fields on these two sides can shear at any angle. Many three-dimensional properties of magnetic reconnection in such asymmetric geometry remain unclear. We use first-principle simulations to explore the 3D kinetic physics that control this critical energy conversion process.

RESEARCH CHALLENGE

Massive solar eruptions drive magnetic storms that impact Earth's magnetosphere and produce space weather. The consequential electromagnetic waves, electric currents, and energetic particles can do harm to satellites, astronauts, GPS systems, radio communication, and power grids on the ground. Magnetic reconnection is the key player in such solar wind-magnetospheric coupling, and space weather in general. One fundamental question in reconnection study is: Is there a simple principle that determines the orientation of the reconnection x-line in such an asymmetric current sheet? The solution of this problem remains unclear with our current understanding of magnetic reconnection, and we aim to resolve this issue. Ultimately, we hope to develop an adequate understanding of the 3D nature of asymmetric magnetic reconnection itself, which is a crucial step in the quest for predicting the location and rate of flux transfer at Earth's magnetopause. A better understanding of the nature of magnetic reconnection will advance the modeling of space weather.

METHODS & CODES

This project employs the particle-in-cell code VPIC [1]. VPIC solves the relativistic Vlasov–Maxwell system of equations using an explicit charge-conserving approach. Charged particles are advanced using Leapfrog with 6th-order Boris rotation, then the current and charge density are accumulated on grid points to update electromagnetic fields. Marder divergence cleaning frequently is employed to ensure the divergent free of the magnetic field. The level of error is bounded by the numerical round-off effect. These 3D kinetic simulations are now generating large amounts of data [~O(100) TB] for each run. We have met this challenge with help from visualization experts at LANL and LBNL to develop parallel readers that interface between these large VPIC data sets and the open-source package ParaView.

RESULTS & IMPACT

During the first award year, we have conducted several petascale simulations on Blue Waters to study the three-dimensional




Figure 2: The current density on the reconnection plane in a sequence of time shows the self-generated turbulence during magnetic reconnection. The arrows indicate the reconnection outflows.

nature of the reconnection x-line. We identify features in these simulations that can potentially bring the most scientific merit. Our approach of inducing a solitary x-line is working well in a large-scale 3D simulation. We initialize the x-line at the center of the box and then let the x-line expand freely and form a well-defined orientation. Here, the orientation is measured as ~13°, close to the angle that bisects the total magnetic shear angle made by magnetic fields on two sides of the current sheet. In comparison with other existing work, where the interaction between multiple x-lines often complicates the orientation measurement, this approach provides an ideal setup for the x-line to grow in a certain “preferential” orientation. Also, we measured the expanding/spreading velocity of the reconnection x-line, decided whether there is an intrinsic reconnection x-line extent in a fully three-dimensional system, and studied the turbulent nature of reconnection.

The same principle that determines the x-line orientation could interplay with global geometrical effects to determine the location and orientation of magnetic reconnection at Earth's magnetopause. Knowledge of the extent of x-line will enable space scientists to more accurately estimate the efficiency of flux transfer from solar wind to the Earth's magnetosphere. The work proposed here is

relevant to the study of dayside reconnection during the first phase of NASA's Magnetospheric Multiscale Mission (MMS). The close tetrahedron deployment of the MMS spacecraft cluster and the equipped capability of high time-resolution provide an invaluable chance to reconstruct the three-dimensional geometry of the reconnection diffusion region.

WHY BLUE WATERS

Because the x-line has a dimension down to electron scale, a fully kinetic description is necessary. Given the available computational capability, it has become possible to use a first-principle kinetic simulation to investigate the dynamics of the x-line in a reasonably large 3D system, which spans from electron kinetic scale to magnetohydrodynamics scale. A representative 3D run in this project traces the motion of 2 trillion charged particles under the interaction of self-generated electromagnetic fields, which are evaluated on 6 billion grids. The output data easily have a size of hundreds of TBs for each run. Blue Waters not only provides the computational resource for the calculation but also the online storage for the output and restart files.

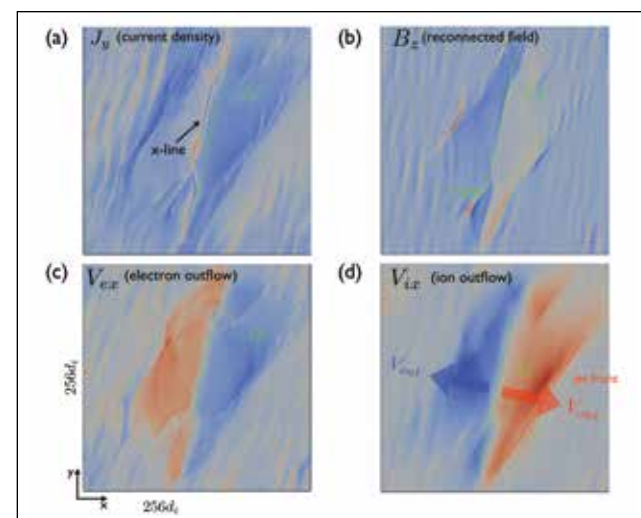


Figure 1: Quantities shown on the 2D interface between the magnetosphere plasma and magnetosheath plasma. The current density, reconnected field, electron outflow, and ion outflow all illustrate the reconnection x-line with a well-defined orientation and extent.

TRANSFORMATIVE PETASCALE PARTICLE-IN-CELL SIMULATIONS

Allocation: NSF PRAC/3,200 Knh

PI: Warren B. Mori¹

Co-PI: Frank S. Tsung¹

Collaborators: L. O. Silva², R. A. Fonseca³

¹University of California, Los Angeles

²Instituto Superior Técnico

³University Institute of Lisbon

EXECUTIVE SUMMARY

The NSF Leadership-class system at NCSA, Blue Waters, was used to study highly linear and kinetic processes that occurred in high-energy density plasmas, including plasma-based acceleration and laser fusion. These simulations have resulted in many high-impact publications in *Physical Review Letters*. In summary, Blue Waters resources allow the UCLA simulation group to perform high-impact research and train a new generation of plasma physicists capable of performing simulations on current and upcoming world-class supercomputers.

RESEARCH CHALLENGE

The research focused on three key areas with clearly developed science questions identified by the community [1–3]. This research, funded by both the National Science Foundation and the Department of Energy, addressed the following key questions:

- 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?
- Can laser plasma instabilities be controlled or even harnessed in inertial fusion plasmas?
- What are the collective processes responsible for the formation of shocks in collisionless plasmas? Are collisionless shocks in plasmas responsible for the most energetic particles in the universe?

METHODS & CODES

Based on the highly nonlinear and kinetic processes occurring in high-intensity laser and beam-plasma interactions and in plasma based acceleration, we use PIC (particle-in-cell) codes [4,5], where Maxwell's equations are solved on a grid using currents and charge densities calculated by weighting particles onto the grid. For this project, we employ the PIC codes OSIRIS, QuickPIC, and UPIC. These codes are all developed locally by the UCLA simulation group (and in collaboration with Instituto Superior Técnico), share many of the same algorithms and data structures, and have been optimized for heterogeneous leadership class supercomputers such as Blue Waters. These codes are freely available and QuickPIC and UPIC are open source (on GitHub).

WHY BLUE WATERS

Blue Waters provides the largest, time-tested, and stable supercomputing platform in the world. It has a CPU side that is time tested and very stable, and has allowed us to perform large-scale simulations in a timely manner, and a GPU side that provides a testbed for code development. This combination suits almost all research supercomputing needs, and has provided a productive computational environment and a satisfactory experience since the very beginning of Blue Waters.

RESULTS & IMPACT

Large-scale particle-in-cell simulations have been performed that will impact the design of future experiments in plasma-based accelerators and inertial confinement fusion.

Blue Waters was used to study the generation of high-quality electron beams (with high brightness and low energy spread) using the density down-ramp injection scheme. By carefully choosing the parameters of the plasma and the driver, Blue Waters simulations show that electrons with 1GeV (gigaelectronvolt) energy, 0.2% relative projected energy spread, and >10kA (kilampere) current can be generated in nonlinear plasma wakes. This is suitable for X-ray FEL (free-electron laser) applications, which can provide a compact radiation source for nuclear science. Simulations were also run to study [6] the effects of ion motions on very-high-brightness beams such as those required for future collider design. In future colliders, which confine electrons within a radius of a few nanometers, the space charge forces around the beam can pull the plasma ions inward. This generates nonlinear focus force inside the wake, which can potentially lead to beam-emittance growth. The self-consistent simulation for this problem remains a big challenge because the simulation box has a transverse size of hundreds of microns. However, the cell size needs to be a few nanometers in order to resolve the electron beam, which results in a simulation box with 10^{11} cells. Blue Waters (because of its large memory size) is ideally suited to study this particularly demanding problem, and simulations performed on it showed that ion motion does not necessarily lead to catastrophic emittance growth.

In collaboration with colleagues at the Naval Research Laboratory, Blue Waters has been used to study the effects of temporal bandwidth (laser smoothing) on laser plasma interactions in laser fusion. In current inertial confinement fusion experiments, lasers can lose a large fraction of their energy to laser-plasma interactions where the laser decays into a backscattered light and a plasma wave. Large-scale 2D OSIRIS simulations showed that, given enough bandwidth (in the order of several terahertz), laser plasma instabilities can be suppressed. Blue Waters resources allow our group to simulate, for the first time, the interaction of many speckles in 2D over several picoseconds with beam smoothing. The additional realism provided by these simulations will lead to a better understanding of laser-plasma interactions in current and future experiments in inertial confinement fusion.

MINI-DISK DYNAMICS ABOUT SUPERMASSIVE BLACK HOLES BINARIES

Allocation: NSF PRAC/1,500 Knh

PI: Scott C. Noble¹

Co-PIs: Manuela Campanelli², Julian H. Krolik³

Collaborators: Dennis Bowen², Stephane d'Ascoli², Vassilios Mewes², Hotaka Shiokawa⁴

¹The University of Tulsa

²Rochester Institute of Technology

³Johns Hopkins University

⁴Harvard-Smithsonian Center for Astrophysics

EXECUTIVE SUMMARY

The overall goal of our project is to provide the community with realistic predictions of electromagnetic (EM) signatures from supermassive black hole binaries (SMBHBs). These events are of much theoretical interest now since little is known about them, but they are ripe for discovery with current and future high-cadence large-sky observational campaigns (e.g., PAN-STARRS, Catalina Real-time Transient Survey, LSST), as well as future multi-messenger endeavors with spaced-based gravitational wave detectors. We paved the way toward this goal by performing a set of hydrodynamic and magnetohydrodynamic (MHD) simulations of these systems using the HARM3d code on the Blue Waters system. The simulations are frontier calculations and require incredible resources and significant effort to cover the large range in time and length scales inherent to their physical situation.




Figure 1: Fourier power spectrum of the fluctuating mass within the sloshing region from the 2D nonmagnetized simulation. The amount of light emitted from this region is expected to correlate with the mass, so the plot serves as a proxy for the predicted variability of the mini-disks' EM emissions. Such a signal could be used to learn properties of an observed system by comparing data to our predictions. Image credit: D. Bowen².

RESEARCH CHALLENGE

In previous years, we have performed 3D general relativistic MHD (GRMHD) simulations of the circumbinary disk that resides outside the orbit of the binary [1, 2]. These long timescale runs were affordable because we excised the binary from the evolution, which allowed us to take larger time steps yet maintain numerical stability. Unfortunately, the excision technique prevented us from learning about the gas dynamics that occurs near each black hole (BH) and the processes that develop between the "mini-disks" of gas that form there. These mini-disk simulations were the focus of our work on Blue Waters this past year, and provided us with first-of-a-kind comprehensive GRMHD simulations of accreting SMBHBs. Including the mini-disks is critically important to understanding the EM signatures because the mini-disks make up a large fraction of the total luminosity from the system, and they give rise to the most variable emission, which is a key means by which astronomers hope to characterize and identify SMBHBs.

In order to start from initial conditions closer to a steady-state and to save resources, Rochester Institute of Technology graduate student Dennis Bowen wrote a method to generate mini-disks around each BH. The initial data method was used first to explore how general relativistic terms affect the dynamics of 2D nonmagnetized disks [3].

METHODS & CODES

All the results reported here used the flux-conservative, high-resolution, shock-capturing GRMHD code called HARM3d [4]. Besides using modern computational fluid dynamics techniques and a solenoidal constraint transport scheme for the magnetic field's induction equation, GRMHD is written in a general covariant way so that any metric or coordinate system may be adopted. This last property has accommodated our implementation of a novel, time-dependent, nonuniform gridding scheme for resolving the small spatial scale features near each BH, and the larger spatial scale flow orbiting circularly that extends much further out beyond the SMBHB's orbit [5]. This coordinate system enabled us to perform the first 3D GRMHD simulations of SMBHB mini-disks already mentioned.

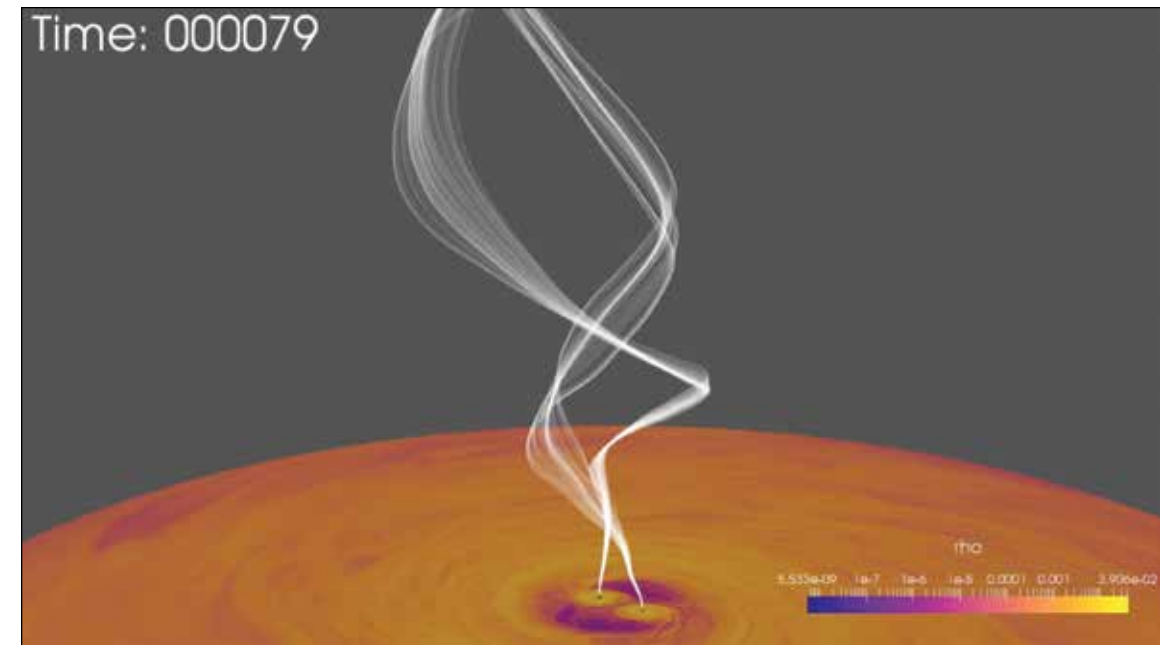


Figure 2: Magnetic field lines (white curves) emanating from two magnetized black holes from the 3D magnetized simulation. An equatorial slice of density of the accreting gas is also shown (background). Image credit: M. Van Moer (NCSA).

We are extending this idea of mesh refinement to one involving multiple coordinate patches, which we call "Patchwork." The scheme is able to solve the MHD equations of motion using multiple, overlapping coordinate system patches. We validated the method on various machines using a variety of tests, which are described in detail in [6]. After passing these tests, we performed a production-level science run of a star's disruption by a black hole from approach all the way to disruption and fallback.

RESULTS & IMPACT

In the 2D hydrodynamic simulations, we found that significant mass exchange, or "sloshing," of material occurred between the two mini-disks. The fluctuations of the mass in the sloshing region resulted in a high-quality and characteristic signal, which we would expect to produce a similarly characteristic periodic EM signal because of the high rate of dissipation there. We further found that the amount of material in the sloshing region increases as the binary shrinks, suggesting that the sloshing emission will become brighter as the BHs grow closer and proceed to the relativistic limit. Our findings suggest that binaries near merger will be bright and periodic at a time scale associated with the SMBHB's orbital period—a key to extracting information about an observed binary's orbit.

Continuing our use of the new mini-disk initial data, we have performed the first full 3D MHD evolution of mini-disks about black hole binaries in the relativistic regime, including the accretion from the circumbinary disk. This full run will begin to address the importance of spiral shocks in the overall angular momentum budget relative to internal MHD stresses. Additionally,

we will be able to explore the effects of accretion streams from the circumbinary on the overall structure of the mini-disks. Our results from this simulation will be submitted soon.

WHY BLUE WATERS

The 3D GRMHD mini-disk simulation ran for two orbital periods and used 12.9 million floating-point-core-hours, or 0.806 M node-hours, on Blue Waters. The simulation used 600x160x640, or approximately 60 million cells, on about 2 million time steps using 600 nodes or 19,200 Blue Waters cores. The simulation is challenging because of the large dynamic range of time scales between the fast behavior near the black holes and the relatively slow orbital velocity of the binary, resulting in a month-long run time. Members of the NCSA Blue Waters team, e.g., David King and Jing Li, were very helpful in arranging a reservation for our run during a busy time on the cluster. The reservation allowed us to finish the simulation before the end of the allocation period.

PUBLICATIONS AND DATA SETS

Zilhao, M., S. Noble, M. Campanelli, and Y. Zlochower, Resolving the relative influence of strong field spacetime dynamics and MHD on circumbinary disk physics. *Physical Review D*, 91:2 (2015), 024034.

Bowen, D., et al., Relativistic Dynamics and Mass Exchange in Inspiring Binary Black Hole Mini-Disks. *Astrophysical Journal*, 838 (2017), p. 42.

Shiokawa, H., R. Cheng, S. Noble, and J. Krolik, PATCHWORK: A Multipatch Infrastructure for Multiphysics/Multiscale/Multiframe Fluid Simulations. *Astrophysical Journal*, arXiv:170105610, under review (2017).

DI

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

Allocation: NSF PRAC/2,500 Knh
PI: Michael L. Norman^{1,2,3}
Co-PI: James Bordner^{1,3}
Collaborators: David Tytler^{1,2}, Pengfei Chen^{1,2}, Dan Reynolds⁴

¹University of California, San Diego
²Center for Astrophysics and Space Sciences
³San Diego Supercomputer Center
⁴Southern Methodist University

BW

MP

EXECUTIVE SUMMARY

We have performed state-of-the-art simulations of the intergalactic medium (IGM) during the epoch of helium reionization. UV radiation from quasars ionizes the intergalactic helium over a period of several billion years beginning about 1 billion years after the Big Bang. We have carried out a suite of the first fully coupled radiation hydrodynamic cosmological simulations that treat the quasars as a time-varying population of point sources. We have performed multigroup radiative transfer self-consistently coupled to the cosmological hydrodynamics of the IGM at sufficient resolution and domain size to examine the photoionization and photoheating processes in detail. We

have discovered that helium reionization completes significantly later compared to models that treat the quasar radiation as a homogeneous background. This modifies the heating history of the IGM substantially, with a maximum mean temperature of 14,000°K being achieved at a redshift of 3, consistent with observations.

RESEARCH CHALLENGE

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

METHODS & CODES

We have used an enhanced version of the hydrodynamic cosmology code ENZO to examine the possibility that inhomogeneous photoheating of the IGM by quasars is the missing physics. In the standard model of the Lyman alpha forest, quasar ionization is modeled as a homogeneous but time-varying radiation background. Photoheating is treated in the optically thin limit, which underestimates the heating behind optically thick ionization fronts. In reality, quasars are radiating point sources that ionize the helium in the IGM in their vicinity; it is the growth and eventual

overlap of these growing spheres of ionization, each centered on a luminous quasar, that ionize the IGM. This is the situation we have simulated on Blue Waters (Fig. 1).

We have carried out a suite of the first fully coupled radiation hydrodynamic cosmological simulations that treat the quasars as a time-varying population of point sources. We have used the multigroup flux-limited diffusion (MGFLD) branch of ENZO developed by Dan Reynolds to perform the simulations. Simulations with grids sizes and particle counts of 1,0243 and 2,0483 were performed on 64/1,024 and 256/4,096 nodes/cores, respectively. Using MGFLD, we accurately transport the hard UV radiation from quasars covering photon energies 54.4 eV to 500 eV and calculate its effects on the IGM self-consistently. Results are analyzed using the open source yt toolkit.

RESULTS & IMPACT

We find that the IGM photoheating is inhomogeneous and time-dependent due to multiple quasar point sources turning on and off over the 2 billion year interval we simulate between redshifts 5 and 2. Fig. 1 shows slices through the 80 Mpc volume for five quantities related to the ionization of helium, for four different redshift snapshots. Bubbles of doubly ionized helium grow and merge in the first row of images so that by the last column, the IGM is completely ionized. The temperature of the IGM is boosted to about 14,000°K in roughly spherical shells surrounding each quasar, as shown in the third row of images.

The time evolution of the IGM temperature at mean density is depicted in Fig. 2 for several simulations, overlaid on observational data. The data points show that IGM temperature peaks around redshift 3. The standard optically thin model shown with the blue line peaks at a redshift of 3.5—earlier than observed. However, the MGFLD simulations, shown in other colored lines, peak around redshift 3, in agreement with observations.

the MGFLD simulations, shown in other colored lines, peaks around redshift 3, which is in agreement with observations. The reason for the difference has to do with the finite time it takes for an ionization front to propagate across the vast distances of intergalactic space. This discovery helps resolve one, but not all, of the above-mentioned discrepancies among observations and our earlier models. The impact of these results is that all future models of the IGM must be revised to include this finite time propagation effect.

WHY BLUE WATERS

Blue Waters is required because the simulation is extremely computationally intensive. Large spatial volumes need to be simulated at high resolution, to simultaneously resolve the Lyman alpha forest absorbers and at the same time encompass a representative sample of quasars. The ratio of the outer and inner scales is about 2,000, meaning that simulations require 2,000³ grids. Additionally, a multifrequency treatment of radiative transfer is essential to accurately model the transport of the hard UV spectrum of quasars including the phenomena of ionization front pre-ionization and spectral hardening of the radiation field.

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

PUBLICATIONS AND DATA SETS

Chen, P., The Effect of Inhomogeneous He II Reionization on the Lyman Alpha Forest. Ph.D. Dissertation, University of California, San Diego (2016).




Figure 1: Redshift evolution of the thermal and ionization state of the intergalactic medium on a slice through the 80 Mpc simulation volume. Left to right: redshift snapshots at $z=4, 3.5, 3,$ and 2 . Top to bottom: He II fraction, He III fraction, gas temperature, He II photoionization rate, He II photoheating rate.

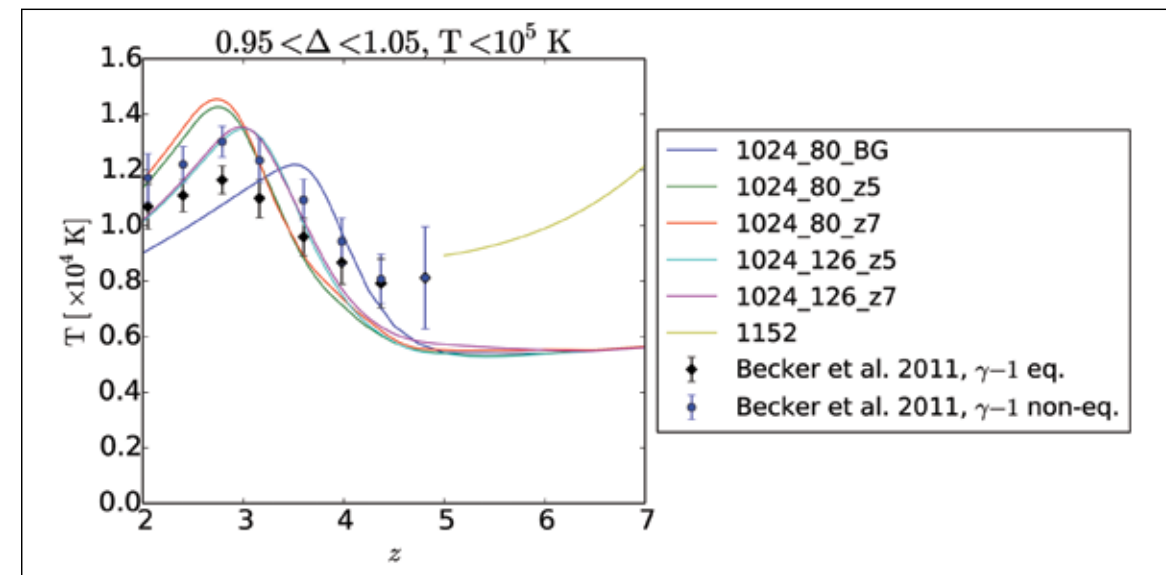


Figure 2: The time evolution of the IGM temperature for several simulations, overlaid on observational data that show that IGM temperature peaks around redshift 3. The standard model (blue line) peaks at $z=3.5$ —earlier than observed. However the MGFLD simulations, shown in other colored lines, peak around redshift 3, in agreement with observations.

SIMULATING GALAXY FORMATION ACROSS COSMIC TIME

Allocation: NSF PRAC/8,100 Knh

PI: Brian W. O'Shea¹

Co-PIs: David C. Collins², John H. Wise³

Collaborators: Devin W. Silvia¹, Cameron J. Hummels⁴, Britton D. Smith⁵

¹Michigan State University

²Florida State University

³Georgia Institute of Technology

⁴California Institute of Technology

⁵University of California at San Diego

EXECUTIVE SUMMARY

This simulation campaign addresses several pressing questions about galaxy formation and evolution by using a suite of physics-rich, high dynamic range adaptive mesh refinement simulations of cosmological structure formation. The two main thrusts of these simulations are: (1) to understand the connection of the first generations of galaxy formation with the Milky Way and its satellites, and (2) to understand the cycling of metal-enriched, magnetized plasma into and out of galaxies like the Milky Way and the regulation of star formation in these galaxies. All of these problems require simulations with extremely high dynamic range in space and time, complex physics (including radiation transport and nonequilibrium gas chemistry), and large simulation volumes. We use the Enzo code (enzo-project.org), which has been modified to scale to large core counts on Blue Waters—the only machine available where the heavy data and communication needs can be satisfied.

RESEARCH CHALLENGE

Our goals are to understand two critical issues in galaxy formation: the formation of the earliest generations of galaxies and their connections to the Milky Way through hierarchical structure formation, and the “baryon cycle” in galaxies like the Milky Way—in other words, how gas gets into and out of galaxies, and what it does while it is there. Both of these questions are important to understanding observations of galaxies over the age of the universe using telescopes such as the 10-meter Keck telescope on Mauna Kea and the Hubble Space Telescope, which are both used to observe light from very distant galaxies as well as the absorption of light by the intergalactic and circumgalactic medium. All of the calculations needed to study these problems require simulations with an extremely high dynamic range in space and time, complex physics (including radiation transport and nonequilibrium gas chemistry), and large simulation volumes.

METHODS & CODES

Our simulation tool of choice is the Enzo code [1; also see <http://enzo-project.org>], an open-source and community-developed software platform for studying cosmological structure formation. Enzo allows us to include all the critical physical components

needed to study galaxy formation—gravity, dark matter dynamics, fluid dynamics, the microphysics of plasmas, and prescriptions for star formation and feedback—and to do so using a tool that can scale to large numbers of CPUs. All analysis was done with the yt code [2; also see <http://yt-project.org>].

RESULTS & IMPACT

The analysis of the simulation performed as part of this campaign has only recently begun. However, our most important results thus far involve the cycling of gas into and out of galaxies. Both cosmological simulations of galaxy formation as well as more idealized calculations demonstrate that massively increased physical resolution in the “circumgalactic medium”—the gas outside the stellar disk of a galaxy but which is bound to the galaxy by gravity, and composes almost half of the mass of the baryons in the galaxy—is incredibly important. In particular, we find that increasing the resolution by more than an order of magnitude beyond previous state-of-the-art calculations results in the appearance of both spatial and chemical features that are seen in observations but not in previous models. Similarly, we find in our idealized simulations that galaxies can attain a dynamic equilibrium between cold gas condensing and falling into the galaxy (and thus fueling star formation), and the ejection of hot, metal-enriched gas into the circumgalactic medium. This work is changing our understanding of the interface between the stellar component of galaxies and the diffuse plasma component.

WHY BLUE WATERS

The simulations used to properly model galaxies in both the early universe and the present day require extremely high spatial and temporal dynamic range, and also require complex physics—most importantly, radiation transport, magnetohydrodynamics, and nonequilibrium gas chemistry. Furthermore, large simulation volumes (and thus many resolution elements) are needed to model the many early galaxies that will merge together to create a Milky Way-like galaxy at the present day. Additionally, in our present-day galaxy simulations, huge numbers of cells are required to accurately resolve the circumgalactic gas. Taken together, this requires the use of a supercomputer with large memory and disk space (to accommodate the tremendous dataset sizes), large




Figure 1: Composite volume rendering of several galaxies at $z=8.6$, approximately 590 million years after the Big Bang. The field of view is 250 kpc (approximately 800,000 light years) across. The blue color table shows density, red shows temperature, and green shows ionizing radiation.

computational resources, and an extremely high bandwidth, low-latency communication network to enable significant scaling of the radiation transport code. Blue Waters is the only machine available to the academic community that fits all of these requirements.

PUBLICATIONS AND DATA SETS

No publications, reports, or datasets have resulted from this simulation campaign yet. Several improvements to the open-source Enzo code have been made, and can be found at <http://enzo-project.org>.

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

Allocation: NSF PRAC/3,730 Knh

PI: Christian D. Ott¹

Co-PIs: Mark Scheel¹, Peter Diener²

Collaborators: Luke Roberts³, Philipp Mösta⁴, David Radice⁵, Roland Haas⁶, Erik Schnetter⁷

¹California Institute of Technology

²Louisiana State University

³Michigan State University

⁴University of California, Berkeley

⁵Princeton University

⁶National Center for Supercomputing Applications

⁷Perimeter Institute for Theoretical Physics

EXECUTIVE SUMMARY

Core-collapse supernovae (CCSNe) are the magnificent explosions of massive stars. They are the birth sites of black holes and neutron stars, and they enrich the interstellar medium with the chemical elements produced by thermonuclear fusion. From these elements, planets form and life develops.

Using Blue Waters, we carried out the very first *ab initio* full 3D general-relativistic radiation-hydrodynamics CCSN simulations, focusing on the phase between initial collapse and the onset of explosion. We investigated the CCSN evolution of a 27-solar-mass progenitor star and followed the supernova engine for approximately 400 milliseconds in full 3D. We found the onset of an explosion driven by a combination of neutrino energy deposition and turbulent convection. The explosion develops in a large-scale asymmetric way, which is consistent with recent astronomical observations of supernovae remnants.

RESEARCH CHALLENGE

CCSNe are cosmic engines where neutron stars and black holes are born. They expel the nuclear products of stellar evolution into the interstellar medium, driving cosmic chemical evolution and providing the foundations for planetary systems and life itself. CCSN shock waves sweep up the interstellar gas and can trigger or shut off the formation of new stars in galaxies.

After the initial collapse of the inner core, a shock wave is launched into the outer core. The shock, however, soon stalls. It requires a reenergizing CCSN mechanism to explode the star. This CCSN mechanism has evaded understanding for almost six decades of theory and computation. The reasons for this are rooted in the complexity of the problem: Understanding CCSNe requires the solution of large coupled systems of partial differential equations describing gravity, magnetohydrodynamics, and neutrino transport (i.e., Boltzmann transport, which is a 6+1D problem). Moreover, astronomical observations show that CCSNe are fundamentally 3D, and turbulent flow is a key ingredient, necessitating a 3D treatment in simulations. Finally,

CCSN simulations must be multiscale, requiring spatial resolution as small as ~50 meters to capture the inertial range of turbulent flow while at the same time encompassing at the very least the innermost ~10,000 km of the star. Ultimately, the entire star will need to be simulated (with a radius of ~10⁹ km for a red supergiant like Betelgeuse).

METHODS & CODES

We employ our 3D general-relativistic radiation-hydrodynamics code, Zelmani. This is an open-source code based on the Cactus framework that uses components of the open-source Einstein Toolkit. Zelmani is the only U.S.-based code that is fully general-relativistic and based on numerical relativity. The key new development enabling our simulations is the ZelmaniM1 component, which implements full 3D multispecies, multienergy radiation-hydrodynamics in the two-moment “M1” approximation. In M1, the Boltzmann equation is expanded in angular moments. The equations for radiation energy (a scalar) and flux (a 3-vector) are evolved and the system is closed with a local algebraic closure relation. Extensive tests show that this approach yields accurate results for CCSN neutrino radiation fields.

We employ a single 27-solar-mass progenitor star and carry out the initial collapse phase in 1D using our open-source code GR1D. We map to 3D at 20 milliseconds after core bounce (the time of protoneutron star formation). In order to study the effects of numerical resolution on shock revival, we carry out simulations with spatial resolution of 1.5 km and 3 km in the region behind the stalled shock. We study the impact of imposed symmetries by simulating in full unconstrained 3D and in an octant of the 3D cube.

RESULTS & IMPACT

In this first set of full 3D CCSN simulations, we confirm that the well-studied “neutrino mechanism,” in combination with turbulent convection, is indeed able to drive 3D CCSNe. The neutrino mechanism relies on the deposition below the shock of a small fraction (5 to 10%) of the neutrinos emitted from the edge




Figure 1: Volume rendering of the specific entropy distribution in a 3D core-collapse supernova at the onset of explosion. The region shown is 600 km³. Red and yellow colors correspond to higher entropy regions while green, blue, and dark colors correspond to regions of lower entropy.

of the protoneutron star. The neutrino energy is deposited in the gas, driving up temperature, thermal pressure, and establishing an entropy gradient that is unstable to buoyant convection. This convection is highly turbulent (theoretical Reynolds number of order 10¹³) since viscosity is extremely small. This turbulent neutrino-driven convection exerts effective turbulent pressure that jointly with the thermal pressure helps push out the shock and explode the star. Fig. 1 shows a volume rendering of the specific entropy distribution within the CCSN shock as the explosion is developing. There are large high-entropy bubbles that develop over time (“volume filling convection”) and push out the shock. The sphere-like object in the center is the low-entropy region inside the protoneutron star. The shock is strongly deformed and shows many small-to-large-scale protrusions. These are created by convective bubbles that impinge on the shock from below. The explosion is globally asymmetric and elongated along one axis. Similar geometries are found in deep astronomical observations of CCSN remnants.

Furthermore, we find that low spatial resolution artificially favors explosion, because it traps turbulent kinetic energy at large scale where it can effectively help shock expansion. This result emphasizes the need for high-resolution simulations in addition to including the full set of physics ingredients. Comparing our octant 3D simulations with full 3D, we find that in the latter, explosions develop more easily. This is because the expanding large high-entropy bubbles that form in full 3D tend to have complex geometry that cannot be captured by a simulation that is constrained to an octant.

WHY BLUE WATERS

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

PUBLICATIONS AND DATA SETS

Roberts, L. F., et al., General Relativistic Three-Dimensional Multi-Group Neutrino Radiation-Hydrodynamics Simulations of Core-Collapse Supernovae. *Astrophys. J.*, 831:98 (2016), DOI: 10.3847/0004-637X/831/1/98.

Blackman, J., et al., A Numerical Relativity Waveform Surrogate Model for Generically Precessing Binary Black Hole Mergers. *Phys. Rev. D.*, 96, 024058 (2017), DOI: 10.1103/PhysRevD.96.024058.

Radice, D., et al., Probing Extreme-density Matter with Gravitational-wave Observations of Binary Neutron Star Merger Remnants. *Astrophys. J. Letters*, 842:L10 (2017), DOI: 10.3847/2041-8213/aa775f.

Morozova, V., et al., Numerical Modeling of the Early Light Curves of Type IIP Supernovae. *Astrophys. J.*, 829:109 (2016), DOI: 10.3847/0004-637X/829/2/109.

Ott, C. D., Massive Computation for Understanding Core-Collapse Supernova Explosions. *Computing in Science & Engineering*, 18:78 (2016), DOI: 10.1109/MCSE.2016.81.

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

Allocation: NSF PRAC/2,140 Knh

PI: Nikolai Pogorelov¹

Co-PI: Jacob Heerikhuisen¹

Collaborators: Tae Kim¹, Mehmet Sarp Yalim¹

¹University of Alabama in Huntsville

EXECUTIVE SUMMARY

Blue Waters time allocation through the NSF PRAC award was used to investigate physical phenomena occurring when the solar wind (SW) interacts with the local interstellar medium (LISM): (1) the origin of the SW on the solar surface and its further acceleration to supersonic velocities; (2) the effect of transient phenomena on space weather on Earth; (3) the SW propagation throughout the heliosphere and perturbations it creates in the LISM; (4) the effect of the heliospheric boundary layer on the plasma oscillation events observed by Voyager 1 in the LISM; (5) MHD (magnetohydrodynamics) instabilities and magnetic reconnection; (6) the influence of the heliosphere on the observed anisotropy of TeV galactic cosmic rays; and (7) using observations from multiple spacecraft to reconstruct otherwise missing properties of the SW and LISM. Our simulations are important for the explanation of IBEX (Interstellar Boundary Explorer), New Horizons, Ulysses, and Voyager measurements, as well as air shower observations.

RESEARCH CHALLENGE

The grand challenge of this research is to investigate physical phenomena that start on the solar surface and result in the SW acceleration and propagation through interplanetary space toward the boundary of the heliosphere, where the SW interacts with the LISM. Our simulations are data-driven and help interpret observations from such space missions as IBEX, New Horizons, Ulysses, Voyager, and a fleet of near-Earth spacecraft. We use vector magnetogram data and STEREO observations to study the propagation of coronal mass ejections toward Earth, where they affect space weather. Voyager 1 and 2 (V1 and V2) spacecraft crossed the heliospheric termination shock (TS) in December 2004 and in August 2007, respectively, and after 40 years of historic discoveries, V1 is sampling the local interstellar medium (LISM) [1], while V2 is approaching the heliopause (HP)—a tangential discontinuity separating the 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 [2], while their observations at the same distance from the Sun are markedly different. V1 data related to the LISM properties gives the heliospheric community a unique

opportunity to study physical processes beyond the HP. IBEX is measuring line-of-sight integrated fluxes of energetic neutral atoms (ENA) in different energy bands [3]. Since most ENA are created during charge exchange between hot PUI (pickup ions) and LISM neutral atoms, they bear the plasma properties of the region in which they are created. The LISM-related objectives of the proposal are to use observational data for the analysis of the SW–LISM interaction, including the heliospheric boundary layer on the LISM side of the HP and the effect of charge exchange on the bow shock, instabilities, and magnetic reconnection near the HP, modifications to the LISM properties due to the presence of the heliosphere, and magnetic field dissipation in the heliosheath between the TS and HP.

METHODS & CODES

We solve the ideal magnetohydrodynamics (MHD) equations coupled with the kinetic Boltzmann equation describing the transport of neutral atoms. In a less strict approach, the flow of atoms is modeled with a few systems of the Euler gas dynamic equations describing different atom populations dependent on the domains of their origin. We have developed both fluid dynamics and kinetic models for PUI and turbulence generated by kinetic instabilities of their distribution function. All these are components of a Multi-Scale Fluid-Kinetic Simulation Suite (MS-FLUKSS)—an adaptive mesh refinement code we have built on the Chombo framework from LBNL.

RESULTS & IMPACT

As a result of the work supported by our Blue Waters allocation through the NSF PRAC award OCI-1615206 we have: (1) performed data-driven simulations of coronal mass ejections starting from the solar surface; (2) simulated the SW propagation along the Ulysses, Voyager, and New Horizons trajectories, and also at Uranus (Fig. 1); (3) modeled shocks propagating through the LISM and demonstrated good agreement with observational data; (4) explained the increase in the frequency of plasma waves observed by V1 in the LISM by the presence of a heliospheric boundary layer of depressed plasma density at the surface of the heliopause (Fig. 2); (5) performed high-resolution simulations of the heliopause instability, identified the areas of possible magnetic reconnection near the heliopause, and predicted that




Figure 1: Simulated interplanetary magnetic field direction, which exhibits the Sun's rotation effects, solar wind radial velocity, and proton number density (scaled to 1 AU values) at the time of the New Horizons spacecraft's closest approach to Pluto. Two cross-sections are shown: by the solar equatorial plane and the plane containing the solar rotation axis and Pluto.

V2 is likely to observe more reconnection than V1; (6) shown that the heliospheric magnetic field may be dissipating in the heliosheath producing turbulence in the SW flow; (7) analyzed the heliotail flow and quantity distribution in the heliospheric bow wave for different LISM conditions, and showed that the observed multi-TeV cosmic ray anisotropy may be explained by the LISM magnetic field distortion by the heliosphere; and (8) investigated the heliotail structure from IBEX observations and the implications for the 11-year solar cycle in the heliosheath. The results are published in 10 papers and reported at over 20 (nine invited) scientific meetings. By addressing the basic physical phenomena occurring at the interface of the heliosphere and LISM, our research is of importance for solar and heliospheric physics, physics of the interstellar medium, and plasma physics in general. Our collaboration with the Blue Waters team further promotes the application of adaptive technologies to contemporary plasma physics problems through the development of publicly available packages suitable for multiple applications.

WHY BLUE WATERS

First, neutral atoms are modeled kinetically, and we need on the order of 10^{12} particles in our Monte Carlo simulations. These simulations require particle splitting, multiple grids, and coupling with the MHD module. They also produce multiple data sets sometimes exceeding 1 terabyte, which require hybrid parallelization. Computational region sizes are very large, as in the case of long-tail simulations to about 10,000 AU (astronomical

units). (3) Very deep adaptive mesh refinement is necessary near magnetic reconnection sites.

PUBLICATIONS AND DATA SETS

Heerikhuisen, J., K.V. Gamayunov, E.J. Zirnstein, and N.V. Pogorelov, Neutral Atom Properties in the Direction of the IBEX Ribbon. *Astrophysical Journal*, 831:137 (2016), DOI: 10.3847/0004-637X/831/2/137.

Kim, T. K., et al., Modeling the Solar Wind at the Ulysses, Voyager, and New Horizons Spacecraft. *Astrophysical Journal*, 832:72 (2016), DOI: 10.3847/0004-637X/832/1/72.

Kim, T. K., N.V. Pogorelov, and L.F. Burlaga, Modeling Shocks Detected by Voyager 1 in the Local Interstellar Medium. *Astrophysical Journal Letters*, 843 (2017), DOI: 10.3847/2041-8213/aa7b2b.

López-Barquero, et al., TeV Cosmic-Ray Anisotropy from the Magnetic Field at the Heliospheric Boundary. *Astrophysical Journal*, 842:54 (2017), DOI: 10.3847/1538-4357/aa74d1.

Pogorelov, N. V., M.C. Bedford, I.A. Kryukov, and G.P. Zank, Pickup Ion Effect of the Solar Wind Interaction with the Local Interstellar Medium. *J. Phys. Conf. Ser.*, 767:012020 (2016), DOI: 10.1088/1742-6596/767/1/012020.

Pogorelov, N. V., et al., Modeling Coronal Mass Ejections with the Multi-Scale Fluid-Kinetic Simulation Suite. *J. Phys. Conf. Ser.*, 837 (2017), DOI: 0.1088/1742-6596/837/1/012014.

Pogorelov, N. V., et al., Heliopause Processes and the Structure of the Heliopause: Modeling Energetic Particles, Cosmic Rays, and Magnetic Fields. *Space Science Reviews*, DOI: 10.1007/s11214-017-0354-8 (2017).

Yalim, M. S., N.V. Pogorelov, and Y. Liu, A data-driven MHD model of the global solar corona within Multi-Scale Fluid-Kinetic Simulation Suite (MS-FLUKSS). *J. Phys. Conf. Ser.*, 837 (2017), DOI: 10.1088/1742-6596/837/1/012015.

Zirnstein, E. J., J. Heerikhuisen, and G.P. Zank, et al., Structure of the Heliotail from Interstellar Boundary Explorer Observations: Implications for the 11-year Solar Cycle and Pickup Ions in the Heliosheath. *Astrophysical Journal*, 836 (2017), DOI: 10.3847/1538-4357/aa5cb2.




Figure 2: The distribution of plasma density (left panel) along the V1 trajectory and its comparison with the plasma wave events detected by the spacecraft beyond the heliopause (right panel).

UNIFIED MODELING OF GALAXY POPULATIONS IN CLUSTERS

Allocation: NSF PRAC/3,000 Knh

PI: Thomas Quinn¹

Co-PI: Fabio Governato¹

Collaborator: Michael Tremmel¹

¹University of Washington

EXECUTIVE SUMMARY

Clusters of galaxies are both a useful probe of cosmology and a laboratory for understanding galactic feedback processes. However, modeling galactic-scale feedback processes in the context of a cluster presents a computational challenge because of the large dynamic range involved. Through the use of a highly scalable N-body/Smooth Particle Hydrodynamics code running on Blue Waters, our project is beginning to tackle this challenging problem. Preliminary results show that models that have successfully reproduced the morphology and number densities of field galaxies can also produce realistic models of cluster galaxies. Large computational resources with high-performance networks are necessary for these calculations.

RESEARCH CHALLENGE

Groups and clusters of galaxies are the largest bound objects in the Universe, containing more than a third of the warm-hot diffuse gas and a significant fraction of the galaxies in the universe. Consequently, understanding the physical processes that occur in group and cluster environments, including the interactions among the dark matter, hot diffuse gas, stars, and active galactic nuclei (AGN), is key to gaining insights into the evolution of baryons and galaxies across the age of the universe. Furthermore, galaxy clusters are one of the few places where the majority of baryons are visible via X-ray and microwave. In contrast to field galaxies, where feedback from supernovae and AGN puts gas into a mostly invisible circumgalactic medium, feedback from cluster galaxies will impact the state of the intracluster medium (ICM). Hence, clusters will provide very tight constraints on our understanding of galactic feedback processes. Clusters of galaxies are also key probes of cosmology and large-scale structure. Their size makes them visible across a wide range of redshifts, and their population statistics are sensitive to cosmological parameters such as the amplitude of the initial power spectrum and the evolution of the cosmic expansion rate. However, using clusters as cosmological probes requires understanding of the relationship between observables and the total mass of the cluster, which in turn requires the detailed modeling of the gravitational/hydrodynamic processes using large simulations.

METHODS & CODES

This project uses the highly scalable N-body/hydrodynamics code ChaNGa to model the formation and evolution of a population of galaxies in a Coma-sized galaxy cluster, including their contribution to and interaction with the ICM. This code is built on the Charm++ [1] parallel programming infrastructure. It leverages the object-based virtualization and data-driven style of computation inherent in Charm++ to adaptively overlap communication and computation and achieve high levels of resource utilization on large systems. The code has been shown to scale well to 0.5 million cores on Blue Waters [2].

The ChaNGa code includes a well-constrained model for star formation and feedback, and improved implementation of supermassive black hole formation, growth, mergers, and feedback [3,4]. In our previous Blue Waters work we have demonstrated that these models can reproduce populations of field galaxies at intermediate-to-high redshift [5], and can reproduce the observed stellar mass-halo mass relationship of galaxies from dwarfs up to galaxy groups [4].

Our simulations will be compared to observations of cluster galaxies to understand the physical and temporal origin of their morphologies. The model ICM will be compared to X-ray and microwave data (via the Sunyaev–Zeldovich effect) to understand the relation between these observations and the underlying gas properties. Finally, the overall mass distribution will be used to better understand how these clusters act as gravitational lenses for background galaxies.

RESULTS & IMPACT

We have completed a prototype simulation of a cluster one-tenth the mass of our target cluster, but at the same mass resolution. Even this simulation is pushing the state of the art in modeling galaxy clusters. Preliminary analysis of this cluster indicates that we continue to produce the observed stellar mass to halo mass even at 10^{14} solar masses. The model central galaxy also displays a morphology typical of observed brightest cluster galaxies: an old, extended stellar distribution and a complete lack of cold gas.

WHY BLUE WATERS

Our scientific goals require modeling over a large dynamic range in mass and space. We have demonstrated that we need mass resolutions on the order of 10^5 solar masses to accurately follow star formation and galaxy morphology. In addition, we




Figure 1: The gas density is shown for the central 3 megaparsecs in a simulation of a Virgo-size cluster (10^{14} solar masses). Note the lack of dense gas in the very center and the stripping of the dense gas out of the smaller galaxies.

need to model a galaxy cluster on the order of 10^{15} solar masses that is comparable to those observed over a range of redshifts. Hence, 10 billion particles are needed. Such a simulation can only be run on the largest computers available. Furthermore, the long-range nature of gravity requires a high-performance, low-latency network to perform the calculation.

KINETIC SIMULATIONS OF LARGE-SCALE PLASMA TURBULENCE

Allocation: NSF PRAC/3,850 Knh

PI: Vadim Roytershteyn¹

Co-PIs: Yuri Omelchenko¹, Kai Germaschewski²

Collaborators: John Podesta¹, William Matthaeus³

¹Space Science Institute

²University of New Hampshire

³University of Delaware

EXECUTIVE SUMMARY

This project seeks to gain a better understanding of turbulence in hot rarified plasmas typical of space and astrophysical environments. Specifically, we use large-scale kinetic simulations tracking upwards of a trillion particles at a time to analyze how the energy of turbulent motions is dissipated in such plasmas. The insights generated by the project may explain such long-standing puzzles as the properties and origins of the solar wind or the anomalously high temperature of the solar corona.

RESEARCH CHALLENGE

Plasma turbulence plays a significant role in the dynamics of many systems in the universe, from laboratory fusion experiments, to the Sun, and to astrophysical objects such as accretion disks. While parameters, geometry, and some aspects of the physics may differ among these systems, there is also a large degree of universality, which makes understanding plasma turbulence a grand challenge problem relevant to many fields of study. This project seeks to conduct simulations of plasma turbulence using codes that are capable of faithfully describing microscopic physical effects. This is important since plasma turbulence is a truly multiscale phenomenon, where the very nature of physical processes governing dynamics changes with scales. For example, in plasmas that are sufficiently hot and not too dense (as is typical of space and astrophysical systems), processes operating at small scales determine the ultimate fate of the turbulent energy. Depending on which of the many possible processes dominates, the energy could be transferred to different species (e.g., electrons, protons, or heavier ions), or to distinct populations of the same species. Understanding these issues may help advance our knowledge of the systems where turbulence operates. For example, both the solar wind and the solar corona are significantly hotter than can be explained with simple models, with local heating by turbulence often proposed as one of the likeliest explanations.

METHODS & CODES

The most complete description of the plasmas of interest is provided by the Vlasov–Maxwell equations, a six-dimensional system of partial differential equations. In order to solve these equations, we use the particle-in-cell technique, where plasma is represented as a collection of particles while equations describing electromagnetic fields are solved on a grid. A typical large-

scale simulation can simultaneously track upwards of a trillion particles in order to obtain reliable statistics. This requires petascale computational resources, such as Blue Waters. Some of the simulations performed in this project were guided by and directly compared against observations conducted by the WIND spacecraft in the solar wind.

RESULTS & IMPACT


During the first year of the project, we have mainly targeted two problems: study of current sheet formation in large-scale simulations with conditions modeling solar wind and the so-called selective decay problem.

Solar wind is a stream of plasma, emitted by the Sun, which fills interplanetary space. Interaction of the solar wind with the Earth’s magnetic field gives rise to the Earth’s magnetosphere and its complex dynamical behavior, often referred to as “space weather.” Space weather has a significant impact on us: For example, extreme solar events may affect or damage the power grid and communication satellites. Understanding of the solar wind in general and its turbulence in particular may help develop better models that improve our ability to understand and forecast space weather. Moreover, thanks to the measurements performed by multiple spacecraft operating in the solar wind, it is the best-studied example of large-scale plasma turbulence, making it also very interesting from the perspective of basic research.

While considerable progress has been made, many puzzles remain. One of them concerns identifying specific mechanisms responsible for conversion of turbulent energy into thermal energy (heat) and energetic particles. In typical space plasmas, collisions between particles are relatively rare and cannot provide sufficient randomization, as would be the case in fluid or regular gases. The relevant theoretical ideas focus on collective processes that invoke interaction of charged particles with electromagnetic fields. These processes are thought to operate most efficiently in the regions where the magnetic field is highly sheared. Geometrically, these regions typically resemble sheets or strings embedded in the plasma and are referred to as the current sheets. Due to their small transverse scale, which often approaches internal plasma scales, the dynamics of current sheets inherently depend on microscopic (kinetic) physics.

On the other hand, the large-scale motions in plasma form current sheets. Thus, in order to satisfactorily describe the

Figure 1: This figure illustrates the velocity field in a large-scale hybrid-kinetic simulation of decaying turbulence.



dynamics of current sheets, simulations need to resolve a large range of scales and include the relevant microscopic physics. We have conducted simulations using a hybrid approximation that treats heavier species in plasma (e.g., protons) using a first-principle kinetic formulation, while representing lighter electrons using an approximate fluid model. The simulations were initialized with conditions modeling solar wind plasma and were directly compared against spacecraft observations. The results were in remarkable agreement with observations [1], indicating that the hybrid model can properly describe dynamics of current sheets at scales comparable to proton kinetic scales. A detailed characterization of the current sheets was performed as the first step necessary to identify the physical processes responsible for their formation and evolution. Ultimately, this will help advance understanding of energy dissipation mechanisms in the solar wind.

The second problem considered in the project concerns the problem of selective decay in kinetic plasma. Broadly speaking, the concept of selective decay is the notion that a turbulent system decays toward special states that, somewhat counterintuitively, are characterized by a high degree of correlation among various quantities. A large body of existing work investigated these issues within the framework of a macroscopic approximation known as the magnetohydrodynamics (MHD) model. In the more complete

kinetic model used in the present study, some of the inherent symmetries of MHD are broken, which may lead to significant differences between predictions of MHD and the kinetic models for decay toward certain states. Fig. 1 demonstrates a turbulent field in one of the simulations conducted for the study. While a full analysis of the results is ongoing, preliminary results point to substantial differences with existing predictions, possibly signaling previously unappreciated significance of microscopic physics for global evolution.

WHY BLUE WATERS

The simulations needed to address the scientific questions of this project solve a global problem that cannot be split into a series of smaller simulations. As such, they require large memory, fast on-node computation, and fast internode communications. For this reason, they require an HPC resource like Blue Waters and cannot be conducted on cloud resources.

PUBLICATIONS AND DATA SETS

Podesta, J. J., and V. Roytershteyn, The most intense electrical currents in the solar wind: Comparisons between single spacecraft measurements and plasma turbulence simulations. *J. Geophys. Res. Sp. Phys.* (2017), DOI: 10.1002/2017JA024074.

GPU-ACCELERATED ADAPTIVE MESH REFINEMENT

Allocation: Director Discretionary/500 Knh

PI: Hsi-Yu Schive¹

Collaborators: Matthew Turk², John ZuHone³, Nathan Goldbaum¹

¹National Center for Supercomputing Applications

²University of Illinois at Urbana-Champaign

³Harvard-Smithsonian Center for Astrophysics

EXECUTIVE SUMMARY

GAMER is a highly scalable and multi-GPU-accelerated adaptive mesh refinement (AMR) code for astrophysics. It adopts a hybrid OpenMP/MPI (message passing interface)/GPU (graphics processing unit) parallelism model to utilize both CPU (central processing unit) and GPU computing power and to minimize MPI communication. Further, it overlaps CPU computation, GPU computation, and CPU-GPU communication to maximize computational efficiency. In this project, we conduct a series of performance benchmarks on Blue Waters and demonstrate high parallel efficiency for both weak and strong scaling using up to 4,096 XK nodes. The code thus provides a unique numerical tool to study various astrophysical phenomena requiring resolutions that are not realistically attainable by other CPU-based AMR codes. For example, based on the benchmark simulations of merging galaxy clusters, the performance of GAMER using 256 XK nodes is found to be 42 times faster than FLASH, a widely adopted CPU-based AMR code, using 256 XE nodes.

RESEARCH CHALLENGE

The AMR technique has played an indispensable role in computational astrophysics due to the large dynamical range demanded. However, compared to the uniform-resolution approaches, it remains extremely challenging for AMR codes to fully exploit the petascale computing power in heterogeneous CPU/GPU supercomputers like Blue Waters. This is mainly due to the complicated AMR data structure, load imbalance, expensive MPI communication, and the great amount of work required to convert existing time-consuming physical modules to run with high efficiency on GPUs.

Most previous GPU-AMR codes are based on considerably simplified test problems or run on a much smaller number of nodes. In comparison, here we measure performance directly from simulations of binary cluster mergers with large dynamic range capable of capturing the large-scale effects of the cluster merger as well as resolving the properties of turbulence down to the kiloparsec scale, almost an order of magnitude finer than previous work [1]. Moreover, we compare the overall performance directly with FLASH [2], a widely adopted CPU-based AMR code, and further demonstrate the parallel scalability of GAMER on hundreds to thousands of XK nodes, for which achieving good load balance becomes highly nontrivial.

METHODS & CODES

The GAMER code [3] has the following important features.

- **Hybrid OpenMP/MPI/GPUs.** GAMER uses GPUs as PDE solvers and CPUs to manipulate the AMR structure. It uses OpenMP for intra-node parallelization in CPUs, MPI for inter-node communication, and CUDA as the GPU programming interface.
- **Overlapping computation.** CPU computation, GPU computation, and CPU-GPU communication are allowed to overlap, greatly improving the overall throughput when CPUs and GPUs take a similar time to complete their own tasks.
- **Hilbert space-filling curve** for load balance.
- **Efficient usage of memory.** GAMER stores all data in the CPU memory and only temporarily transfers data to GPU, which allows exploitation of the larger CPU memory. Communication overhead between CPU and GPU is usually negligible since it can be overlapped by both CPU and GPU computations efficiently using CUDA streams.
- **Bitwise reproducibility.** The order of all floating-point operations in GAMER is carefully designed to be deterministic. It thus supports bitwise reproducibility when (1) running simulations with different numbers of MPI ranks and OpenMP threads, and (2) restarting simulations from checkpoint files. This feature is essential for scientific reproducibility.
- **Inline analysis with yt.** GAMER supports the ability to pass in-memory data structures to *yt* [4], a powerful Python-based package for analyzing and visualizing volumetric data. Moreover, using *yt* allows one to share the data analysis scripts to the community straightforwardly, which greatly improves the scientific reproducibility.

RESULTS & IMPACT

We have compared both the performance and accuracy between GAMER and FLASH based on the merging galaxy clusters simulation setup of [1]. These simulations include hydrodynamics, self-gravity, dark matter particles, and AMR, and thus provide a comprehensive test for the simulation codes. Fig. 1 shows the slices of gas temperature evolution visualized with *yt*, demonstrating that the results of the two codes are remarkably consistent.

Fig. 2 shows the strong scaling. Blue and green lines indicate the comparison with a relatively lower resolution of 7 kiloparsecs.




Figure 1: Comparison of the gas temperature evolution in the binary cluster merger simulations between GAMER (upper panels) and FLASH (lower panels). The results are remarkably consistent.

GAMER (on XK nodes) is found to be ~27 to 42 times faster than FLASH (on XE nodes) when using the same number of nodes. Most importantly, this speedup remains reasonably flat when increasing the number of nodes. The red line shows the performance of GAMER with a resolution eight times higher (~0.9 kiloparsecs), which achieves much higher parallel efficiencies of ~50% on 512 nodes and ~30% on 2,048 nodes.

The extremely high resolution and performance demonstrated in this project will allow astrophysicists, for the first time, to resolve the turbulence cascade in the intracluster medium down to scales lower than the particle mean free path, reducing the numerical viscosity down to a scale below where the effects of a physical viscosity are expected to become relevant. It will greatly help calibrate the mass estimation of galaxy clusters, refine the active galactic nucleus feedback model, and improve our understanding of the source of cosmic rays.

WHY BLUE WATERS

The benchmark conducted in this project would be impossible without Blue Waters. These simulations require a great amount of CPU and GPU memory, and also require the balance among CPU performance, GPU performance, and MPI communication to achieve an optimum throughput. Most importantly, running thousands of GPUs in parallel is extremely demanding on system stability. Blue Waters is the supercomputer that best fits these requirements.




Figure 2: Strong scaling of GAMER in the binary cluster merger simulations.

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

Allocation: Illinois/990 Knh

PI: Stuart L. Shapiro¹

Co-PIs: Vasileios Paschalidis², Milton Ruiz¹, and Antonios Tsokaros¹

Collaborator: Lunan Sun¹

¹University of Illinois at Urbana-Champaign

²Princeton University

EXECUTIVE SUMMARY

We perform magnetohydrodynamic (MHD) simulations in full general relativity (GR) of the collapse of radially unstable, uniformly rotating, massive stars to black holes (BHs). The stars spin at the mass-shedding limit, account for magnetic fields and obey a polytropic $\Gamma = 4/3$ equation of state (EOS). The calculations lift the restriction of axisymmetry imposed in previous simulations. Our simulations model the direct collapse of supermassive stars (SMSs) to supermassive BHs (SMBHs), with masses larger than $10^4 M_{\odot}$, at high cosmological redshifts z , which may explain the appearance of SMBHs and quasars by $z \sim 7$. They also crudely model the collapse of massive Population III stars to massive BHs, which could power some of the long gamma-ray bursts observed by FERMI and SWIFT at $z \sim 6-8$. We analyze the properties of the electromagnetic and gravitational wave signatures of these events and discuss the detectability of such multimessenger sources.

RESEARCH CHALLENGE

SMBHs with masses between $10^6 M_{\odot}$ and $10^9 M_{\odot}$ reside at the center of most galaxies, including our own galaxy [1]. It is not clear however what are their growth mechanisms since in the early universe (high redshift) there is comparatively a small amount of matter to feed them and less than a billion years for them to consume it. It has been suggested that first generation stars (Population III stars) could collapse and form seed BHs, which later could grow through accretion to become SMBHs [2,3]. An alternative scenario explaining the origin of SMBHs is provided by the direct collapse of stars with masses $\geq 10^4 M_{\odot}$. These so-called SMSs could form in metal-, dust-, and H_2 -poor halos, where fragmentation and formation of smaller stars could be suppressed [4]. We study the evolution of radially unstable, uniformly rotating magnetized SMSs modeled by polytropes with a radiation-dominated ($\Gamma = 4/3$) equation of state that undergo collapse to BHs with masses larger than $10^4 M_{\odot}$ in full

GR. These simulations also crudely model the collapse of massive Population III stars to massive BHs, which could power some of the long gamma-ray bursts observed by the FERMI and the SWIFT satellites at redshift $z \sim 6-8$. These calculations lift the restriction of axisymmetry imposed in previous studies and extend the initial magnetic field to full-space geometry [5].

METHODS & CODES

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

We solve the above equations through our completely independent “Illinois GRMHD code”, which has been built over many years on the Cactus infrastructure and uses the Carpet code for adaptive mesh refinement, but employs our own algorithms and coding (see e.g. [6]). This code utilizes state-of-the-art high resolution shock capturing methods to evolve scenarios involving either vacuum or matter spacetimes, with or without magnetic fields. It utilizes the Baumgarte-Shapiro-Shibata-Nakamura (BSSN) formulation of the Einstein field equations with puncture gauge conditions. It solves the magnetic induction equation by introducing a vector potential and employs a generalized Lorentz gauge condition to reduce the spurious appearance of strong B-fields on refinement level boundaries [6].

Our undergraduate research team also creates visualizations (e.g., Fig. 1) and movies of our simulations with the VisIT software on Blue Waters.

RESULTS & IMPACT

Using our latest adaptive-mesh refinement GRMHD code we performed simulations of massive stars that undergo collapse [5]. We consider a purely hydrodynamic case, and two magnetized cases, one seeded with a poloidal magnetic field only in the stellar interior and the other extended from the stellar interior to the exterior (Fig. 1, upper left). In both cases, the adopted magnetic fields are dynamically unimportant initially. To ensure reliable evolution of the exterior field, we initially impose a low-density atmosphere such that the stellar exterior is described by a constant plasma parameter, defined as the ratio of the gas pressure to the magnetic pressure. To accelerate the collapse, the pressure is initially depleted by 1% in all our cases. We then evolve the stars through the collapse and subsequent BH formation immersed in a magnetized accretion disk (Fig. 1, top right and bottom panels).

We find that, in terms of the BH mass, its spin and torus mass, the results from our hydrodynamic simulations are consistent with previous semi-analytic estimates and axisymmetric simulations in GR reported in [7-8]. We also find that the magnetic field do not affect these global quantities.

In the magnetized cases, following BH formation we observe the formation of magnetically dominated regions above the BH

poles where the magnetic field lines have been wound into a collimated helical funnel, within which the plasma flows outwards. This collimated outflow is mildly relativistic, and constitutes an incipient jet. Our analysis suggest that the Blandford-Znajek effect is likely operating in our simulations and could be the process powering these jets. The magnetization in the funnel reaches values ≥ 200 , and since for steady-state, axisymmetric jets the magnetization approximately equals the jet terminal Lorentz factor, the jets found in our simulations may reach Lorentz factors of 200, and hence, explain gamma-ray burst (GRB) phenomena. The accretion torus lifetime is $\Delta t \sim 10^5 (1+z)(M/10^6 M_{\odot})$ s. Thus, collapsing SMSs with masses $\sim 10^3 M_{\odot} - 10^4 M_{\odot}$ at $z \sim 10-20$ are candidates for ultra-long GRBs, while collapsing Population III stars at $z \sim 5-8$ are candidates for long GRBs. We estimated that for observation times of $\sim 10^4$ s, FERMI and SWIFT could detect such ultra-long GRB events from these stars [9].

WHY BLUE WATERS

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

PUBLICATIONS AND DATA SETS

Sun, L., V. Paschalidis, M. Ruiz, and S. Shapiro, Magnetorotational Collapse of Supermassive Stars: Black Hole Formation, Gravitational Waves and Jets. ArXiv:1704.04502.

Ruiz, M., R. Lang, V. Paschalidis, and S. Shapiro, Binary neutron star mergers: a jet engine for short gamma-ray burst. *Astrophys. J.*, 824 (2016), pp. L6–L11.

Shapiro, S.L., Black holes, disks and jets following binary mergers and stellar collapse: The narrow range of EM luminosities and accretion rates. *Phys. Rev. D*, 95 (2017), pp. 101303–124071.

Tsokaros, A., et al., Gravitational wave content and stability of uniformly, rotating, triaxial neutron stars in general relativity. *Phys. Rev. D*, 95 (2017), pp. 124057–124071.




Figure 1: Volume rendering of the rest-mass density normalized to its initial maximum value (log scale) at selected times. Arrows show plasma velocities and white lines show the magnetic field lines. Bottom panels highlight the system during and after an incipient jet is launched.

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

Allocation: NSF PRAC, 5,220 Knh

PI: Alexander Tchekhovskoy^{1,2,3}

Co-PI: Nicholas Stone⁴

Collaborators: Matthew Liska⁵, Casper Hesp⁵, Adam Ingram⁵, Michiel van der Klis⁵, Sera Markoff⁵, Brian Metzger⁴

¹Northwestern University

²University of California, Berkeley

³Kavli Institute for Theoretical Physics, Santa Barbara

⁴Columbia University

⁵University of Amsterdam

EXECUTIVE SUMMARY

Stars approaching supermassive black holes (BHs) can be tidally disrupted, or spaghettified, by BH tidal gravity. A tidal disruption event (TDE) generally leads to a tilted accretion disk, with an axis misaligned relative to BH spin. However, how such a disk forms and how it evolves thereafter is not understood. Using our new GPU (graphics processing unit)-accelerated 3D general relativistic (GR) magnetohydrodynamic code H-AMR, we investigated both of these questions. We carried out the first GR hydrodynamic

simulations of typical TDEs, with supermassive BHs disrupting Sun-like stars infalling from large distances. We discovered that as GR-induced precession causes the tidal stream to self-intersect, this can completely disrupt the stream and efficiently form the disk. We showed that our simulated tilted thick accretion disks launched twin magnetized relativistic jets that underwent Lense–Thirring precession together with the disk, demonstrating from first principles for the first time that jet precession can be used to probe strong-field gravity.

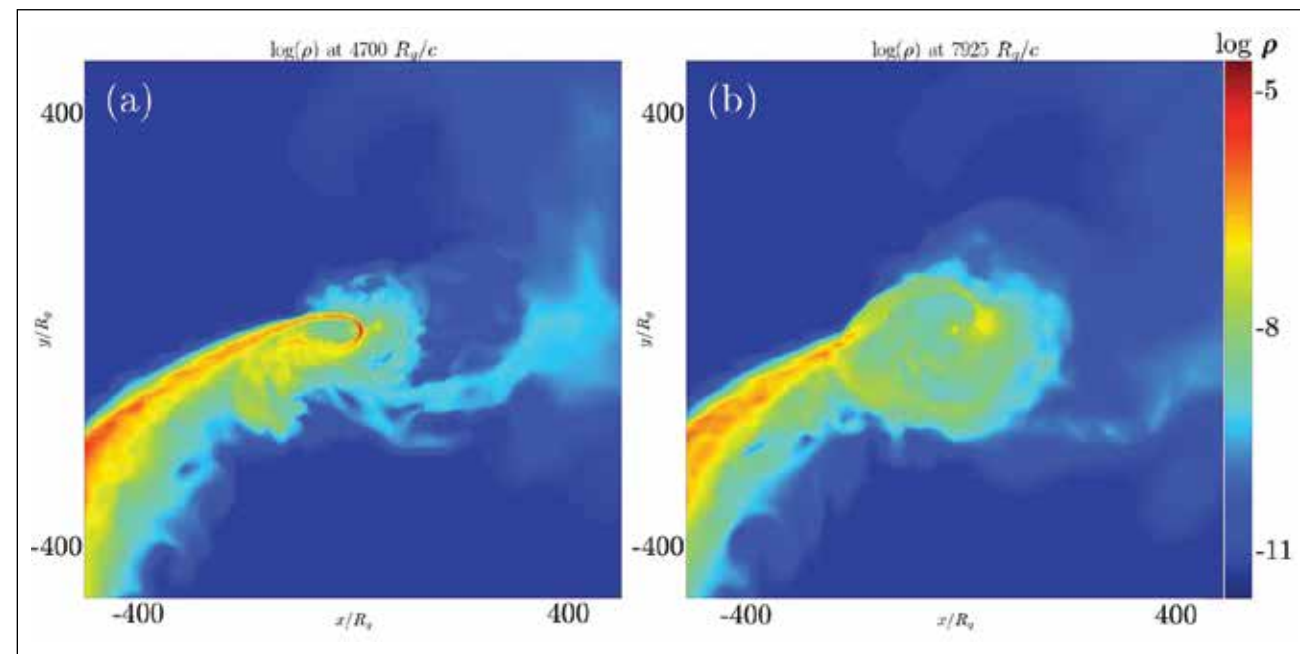


Figure 1: First general relativistic fluid dynamics simulation of typical stellar tidal disruption by supermassive black hole. It demonstrates that debris stream self-crossing can lead to its demise: complete disruption and efficient circularization into disk.

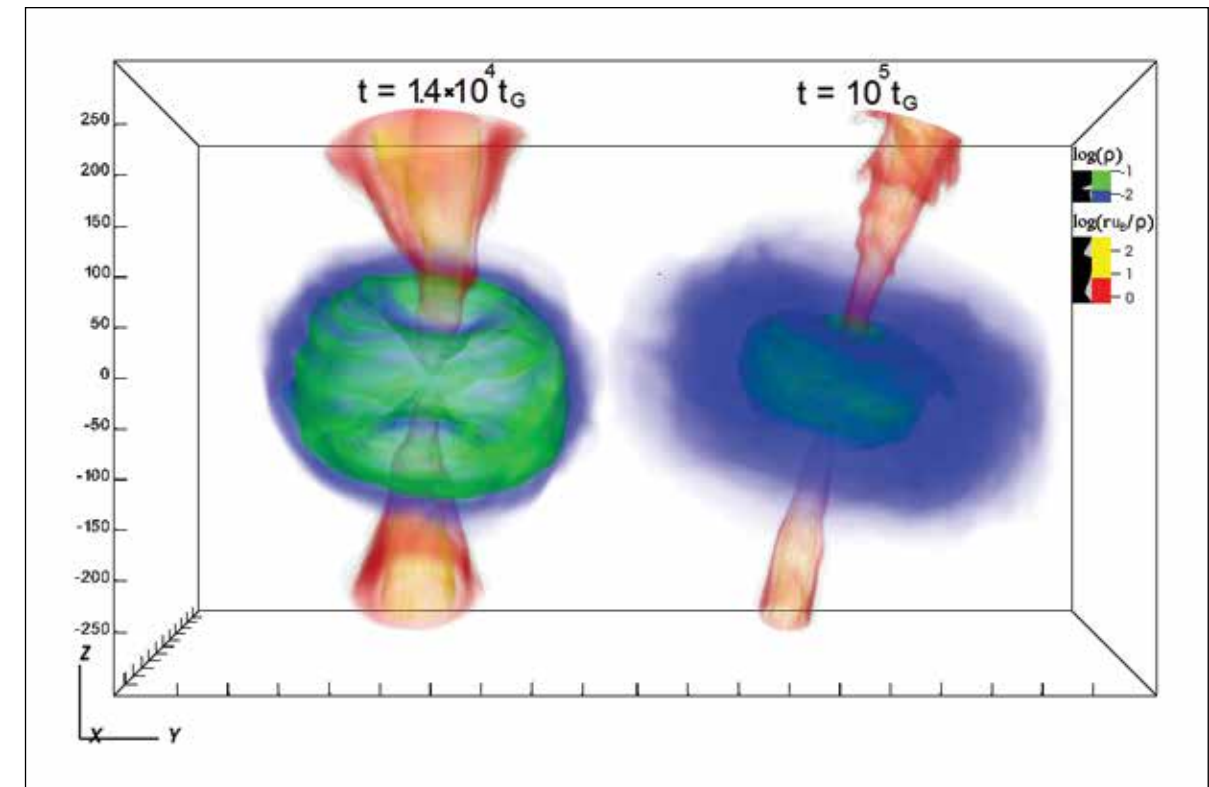


Figure 2: Our PRAC allocation allowed us to demonstrate for the first time that tilted precessing disks around BHs produce relativistic jets and that these jets precess together with the disk.

RESEARCH CHALLENGE

Numerical simulations of TDEs of Sun-like stars by supermassive BHs are extremely challenging because it is difficult to numerically resolve the debris stream because it is very thin relative to the black hole. Tilted disk simulations require high resolution to properly resolve nonaxisymmetric turbulence in the tilted disks. Both are huge numerical challenges.

METHODS & CODES

Using our new code H-AMR (pronounced “hammer”), which includes adaptive mesh refinement and efficiently runs on GPUs, we were able to overcome these challenges.

RESULTS & IMPACT

Until now, many different simplifications have been adopted when simulating TDEs. For instance, to save computational time, (1) stars were sent in on closed, elliptic orbits (instead of parabolic ones) or (2) the BH-to-stellar ratio order of magnitude was assumed to be lower than in reality (e.g., 1,000 instead of 10^6) [1,2]. The adaptive mesh capabilities of H-AMR and the high efficiency with which it runs on GPUs allowed us to properly resolve the thin debris stream for a typical encounter featuring both a parabolic orbit and a mass ratio of 10^6 .

The first simulations of tilted accretion disks were carried out a decade ago [3]. They demonstrated that tilted disks undergo Lense–Thirring precession. However, whether such disks are capable of producing relativistic outflows, or jets, and whether those jets point along the disk axis or the BH spin axis has remained a mystery. Our simulations for the first time established that such disks are indeed capable of producing jets and that the jets undergo precession together with the accretion disk. This is the first demonstration that jets can be used as probes of disk precession.

Blue Waters was instrumental in enabling our simulations, which are the highest resolution GRMHD simulations to date and set the new state of the art in the field. On a CPU (central processing unit) cluster, the effective simulation cost would have been 400-million CPU core-hours, accessible only on Blue Waters.

WHY BLUE WATERS

Our simulations require a high degree of parallelism as they run on hundreds to thousands of GPUs in parallel.

ADVANCED SPACE WEATHER MODELING

PI: Gabor Toth¹

Allocation: NSF PRAC/3,000 Knh

Co-PI: Ward Manchester¹

Collaborators: B. van der Holst¹, Y. Chen¹

¹University of Michigan

EXECUTIVE SUMMARY

Our effort aims to achieve a breakthrough advance in the understanding of space weather. The most destructive forms of space weather are due to major solar eruptions: fast coronal mass ejections (CMEs) and eruptive X-class flares. These destructive events originate with magnetic fields emerging from the solar interior, forming the active regions from where CMEs erupt into the heliosphere. Upon impacting the Earth, interplanetary CMEs impact the magnetosphere and produce geomagnetic storms. This process is controlled by the microphysics of magnetic reconnection. Our goal is to answer the most salient questions of space weather: How the buildup of magnetic energy results in solar eruptions and how magnetic reconnection results in geomagnetic storms.

RESEARCH CHALLENGE

Major space weather events are caused by large-scale expulsions of magnetized plasma from the Sun, which are known as coronal mass ejections (CMEs) and that typically travel to Earth in one to three days. These eruptions occur frequently, as often as several times per day during solar maximum, and cause geomagnetic storms by triggering sudden reconfigurations of the magnetosphere by magnetic reconnection. Extreme space weather events are caused by the most energetic CMEs, which drive sudden and extensive changes in the Earth's magnetic field producing among other effects large-scale electric impulses that can melt transformers and cause cascading blackouts. Repair times for replacing the high-voltage transformers is estimated to be several months.




Figure 1: MHD-EPIC simulation of the Earth's magnetosphere. Left panel: 2D cut; color represents pressure with projected field lines. The MHD simulation extends to 256RE; red box indicates the embedded PIC model (4REx6REx6RE). Right panel: 3D structure of a magnetic flux rope formed by magnetic reconnection. Field line colors represent plasma pressure.

Being able to predict extreme space weather is a challenging task that requires both accurate simulations of CME structures when they reach Earth and the response of the magnetosphere. The magnetic reconnection process that lies at the heart of space weather events depends on the magnetic field carried by the coronal mass ejection as well as on the plasma processes happening at kinetic scales. Accurate modeling of magnetic storms, therefore, requires prediction of the interplanetary magnetic field of CMEs and an accurate model for the reconnection process. Our proposed research addresses both of these crucial issues.

METHODS & CODES

We combine the efficiency of global fluid-type models with the physical capabilities of expensive but physically accurate local kinetic models. The resulting magnetohydrodynamic with embedded particle-in-cell (MHD-EPIC) model is 100 to 10,000 times more efficient than a global kinetic model. In addition, we found that the kinetic scales can be artificially increased, which can dramatically—by many orders of magnitude—reduce the computational cost of the embedded PIC model. Fig. 1 shows an MHD-EPIC simulation of the magnetosphere of Earth.

The flux emergence and CME initiation simulations are carried out with our high-resolution MHD code BATS-R-US in a variation called the Spherical Wedge Active Region Model (SWARM). The simulation domain extends from the convection zone into the corona with spherical wedge grid geometry with a domain the size of an active region. SWARM models the upper convection zone extending from a depth of 0.95 Rs to a height of 1.25 Rs, and extending 12x24 degrees, large enough to encompass a solar active region. Spherically adaptive grids allow us to greatly reduce the number of computational cells while also resolving the photosphere. Using SWARM, we have performed rigorous flux emergence calculations and the formation of active regions with no ad hoc assumptions on coronal or photospheric conditions.

RESULTS & IMPACT

We have used this unique opportunity to simulate space weather events with the MHD-EPIC model, where the reconnection is handled by a kinetic PIC code. With this approach, we focused on modeling of the fundamental process of reconnection and its impact on the global dynamics. Currently, the MHD-EPIC model is the first three-dimensional global study of the complex process reconnection process using a high-fidelity kinetic model for the magnetic reconnection. We also made breakthrough advances in simulating flux emergence at active-region scale in spherical geometry. This work addresses the most salient questions of space weather.

WHY BLUE WATERS

Our project uses the Blue Waters petascale computing facility to perform unprecedented space weather simulations. This capability allows us to simulate magnetic flux emergence from the convection zone into the corona to form active regions that may result in

coronal mass ejections. Using Blue Waters allows us to model a whole active region with sufficient grid resolution to capture magnetic energy buildup. Blue Waters also allows us to model the reconnection process in the magnetosphere with the MHD with embedded PIC model and gain a better understanding of the intricate interaction between the small kinetic scales and the global scales that result in magnetospheric storms.

PUBLICATIONS AND DATA SETS

Chen, Y., et al., Extended Magnetohydrodynamics with Embedded ParticleinCell (MHDEPIC) Simulation of Earth's Magnetopause reconnection. GEM 2016, Santa Fe, N.M., June 19–24, 2016.

Toth, G., et al., Scaling the ion inertial length and its implications modeling reconnection in global simulations. MMS Workshop, Boulder, Colo., June 6–8, 2016.

Manchester, W., B. van der Holst, and G. Toth, Simulating Solar Active Regions. LWS Technical Interchange Meeting, Mountain View, Calif., May 30–June 2, 2017.

3D SIMULATIONS OF I-PROCESS NUCLEOSYNTHESIS

Allocation: NSF PRAC/3,000 Knh
PI: Paul R. Woodward¹
Collaborator: Falk Herwig²

¹University of Minnesota
²University of Victoria

EXECUTIVE SUMMARY

Our team is simulating brief events in the interiors of evolved stars that result in ingestion of unprocessed new fuel into convection zones above nuclear burning shells. The new fuel can burn very violently under the much hotter conditions in the convection zone after reaching a sufficient depth within it. This burning sets off a series of reactions that dramatically affects the nucleosynthesis of heavier elements and, hence, the ultimate expulsion of heavier elements into the surrounding interstellar gas. Expulsion can be either a relatively slow expulsion of the outer envelope of the star or by an explosion of the star itself, if it is sufficiently massive. This work involves large and very detailed 3D simulations of the entirety of the stellar interior, for which Blue Waters is ideally suited.

RESEARCH CHALLENGE

Our simulations involve brief but important events in the lives of stars that can greatly affect the heavier nuclei that the stars produce. We have been concentrating so far on hydrogen ingestion

flashes in which unburned hydrogen-rich fuel is brought into the convection zone above the helium-burning shell. The ingested hydrogen reacts with carbon in the convection zone to set off a sequence of nucleosynthesis reactions that is called the i-process, since the neutron fluxes that result are intermediate between the slow and rapid s- and r-process nucleosynthesis. In the beginning of our project on Blue Waters, we simulated hydrogen ingestion events in evolved stars, and during the last year we have been gearing up to attack the more challenging problem of the potential merger of multiple nuclear burning shells in massive stars. This has involved detailed studies of the ingestion process, particularly in massive star contexts, as well as the aggressive development of a new simulation code. Our results on i-process nucleosynthesis are important as inputs for the study of the chemical evolution of galaxies. Our new work with massive stars, just begun, could have a large impact on the conditions just before those stars explode, and also on the injection of heavier elements from these explosions into the surrounding interstellar medium.




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

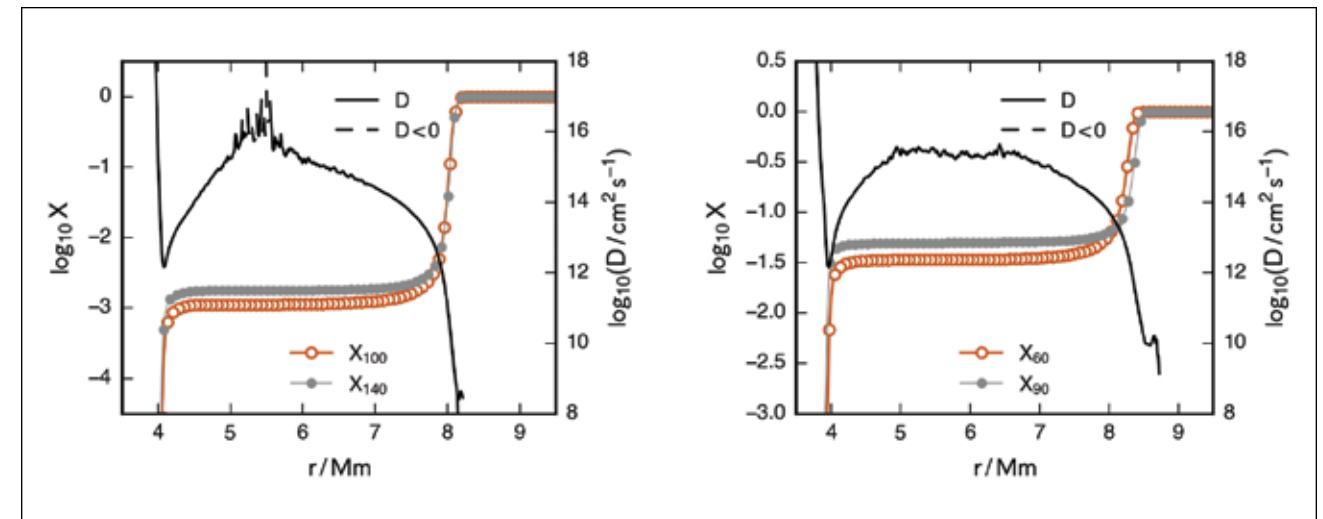


Figure 2: Effective diffusion coefficients in 1D descriptions of 3D simulations of gas entrainment from above the convection zone generated by oxygen burning in model 25-solar-mass star. Left: 2.25x luminosity of 1D model. Right: 56.25x this luminosity. Specialized simulations matching these parameters pinned down very small effective diffusion values in stably stratified gas regions.

METHODS & CODES

Our work to date simulating hydrogen ingestion flashes exploits the piecewise parabolic method (PPM) coupled with the piecewise parabolic Boltzmann (PPB) moment-conserving advection scheme for the multifluid volume fraction. PPB delivers more than double the resolving power of the PPM scheme for the single, very important variable representing the volume fraction of entrained fluid. Together with the already high resolving power of PPM, we are able to obtain very accurate results on a uniform grid. We must simulate a great many large-eddy overturning times in the convection zone above a nuclear burning shell in order to accurately approach a nonlinear, global oscillation of the burning of ingested hydrogen that increases the hydrogen ingestion rate by as much as two orders of magnitude. We are able to cover this long approach to the violent ingestion event because our PPM code scales to nearly 14,000 nodes on Blue Waters, at which scale it advances the simulation by roughly 20 time steps per second. Thus, the millions of time steps we need to simulate an ingestion event accurately are practical on Blue Waters with our code.

We have been turning our attention this year to massive stars, where the ingestion of material from above a burning shell of, for example, oxygen can allow the convection zone above that burning shell to eat its way outward in radius until it reaches the carbon-burning shell above it. A merger of these two burning shells can then result. Simulating this process is very challenging. We have been studying the relevant ingestion process in considerable detail over this last year with the goal of alternating between 1D and 3D simulation in order to span the time necessary in leading up to a shell merger. Our 3D simulations would keep models used in the 1D intervals that stitch one 3D run to the next validated as good descriptions of the full 3D results. This work is illustrated in the Figs. 1 and 2.

We have devoted an enormous effort during the last year and a half to the development of a completely new code. This code adds a Level 3 adaptive mesh refinement (AMR) grid that will enable us to contain multiple nuclear burning shells and their respective convection zones in a single simulation. It is designed to scale to 14,000 nodes while running roughly twice as fast as our older code per node by exploiting 32-bit precision and GPU (graphics processing unit) acceleration [1–3].

Results & Impact

We are producing a database of detailed simulations that investigates the phenomenon of convective boundary mixing at unprecedented accuracy for convection zones that extend over ranges in radius of more than a factor of two (see www.lcse.umn.edu). Global convection modes play an important role in these situations, making simulation difficult and costly [4–6]. Convective boundary mixing plays an important role in stellar evolution. In particular, in ingestion events that we study, it can have a dramatic impact on nucleosynthesis, which in turn affects galactic chemical evolution [7].

WHY BLUE WATERS

We have carried out our simulations on Blue Waters because of its special ability to enable our simulation code to be run at a sufficiently large scale that our large computations can each be completed in less than one week [8]. This allows our team to pose questions and get answers on a timescale that is conducive to productive thought and dynamic adjustment of our research direction.

PUBLICATIONS AND DATA SETS

For publications, see reference list in back of book.
 For shared data sets, see www.lcse.umn.edu.

GEOSCIENCE

CLIMATE

ENVIRONMENT

GEOLOGY

NATURAL HAZARDS

SPACE WEATHER

WEATHER

- 64** *High-Resolution Digital Surface Models of the 2016 Mw7.8 Kaikoura Earthquake, New Zealand*
- 66** *Image Processing to Build a Multi-Temporal Vegetation Elevation Model (MTVEM) of the Great Lakes Basin (GLB)*
- 68** *The Terra Data Fusion Project*
- 70** *Large-Eddy Simulation of Sediment Transport and Hydrodynamics at River Bifurcations: Using a Highly Scalable Spectral Element-Based CFD Solver*
- 72** *The Reference Elevation Model of Antarctica*
- 74** *Untangling Entrainment and Precipitation in Convective Clouds*
- 76** *Understanding the 4-D Evolution of the Solid Earth Using Geodynamic Models with Data Assimilation*
- 78** *Physics-Based Modeling of High-Frequency Ground Motions and Probabilistic Seismic Hazard Analysis*
- 80** *Enhanced Digital Elevation Model for the Arctic*
- 82** *Simulating the most devastating tornadoes embedded within supercell thunderstorms*
- 84** *High-Resolution Earth System Modeling Using Blue Waters' Capabilities*
- 86** *Location-Specific Space Weather Hazards to Electric Power Grids Calculated on a Global Scale*
- 88** *Impact of Ocean Coupling on Simulated Tropical Cyclone Activity in the High-Resolution Community Earth System Mode*
- 90** *Petascale Modeling of Convective Storms Under Climate Change and Variability*
- 92** *Impacts of Orientation and Morphology of Small Atmospheric Ice Crystals on in-situ Aircraft Measurements: Scattering Calculations*
- 94** *Pore-Scale Simulation of Multiphase Flow in Porous Media with Applications to Geological Sequestration of Carbon Dioxide*
- 96** *3D Particle-Resolved Aerosol Model to Quantify and Reduce Uncertainties in Aerosol-Atmosphere Interactions*
- 98** *Particulate Matter Prediction and Source Attribution for U.S. Air Quality Management in a Changing World*

DI

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

Allocation: NSF PRAC/240 Knh

PI: Marin Clark¹

Co-PIs: Michael Willis², Dimitrios Zekkos¹, Adda Athanasopoulos-Zekkos¹

Collaborator: Chris Massey³

¹University of Michigan

²University of Colorado Boulder

³GNS Science, New Zealand

TN

EXECUTIVE SUMMARY

The 4D evolution of topography is critical for understanding faulting and landslide processes and the linkages between the two. For major events that deform and erode landscapes, such as earthquakes, we lack data at an event scale for understanding

such processes because of the prohibitive cost of repeat land and airborne surveys, and because pre-event data is usually lacking. However, recent advances in high-performance computing coupled with stereo-satellite imagery collection allow the opportunity to provide regional-scale, high-resolution topographic time-series.

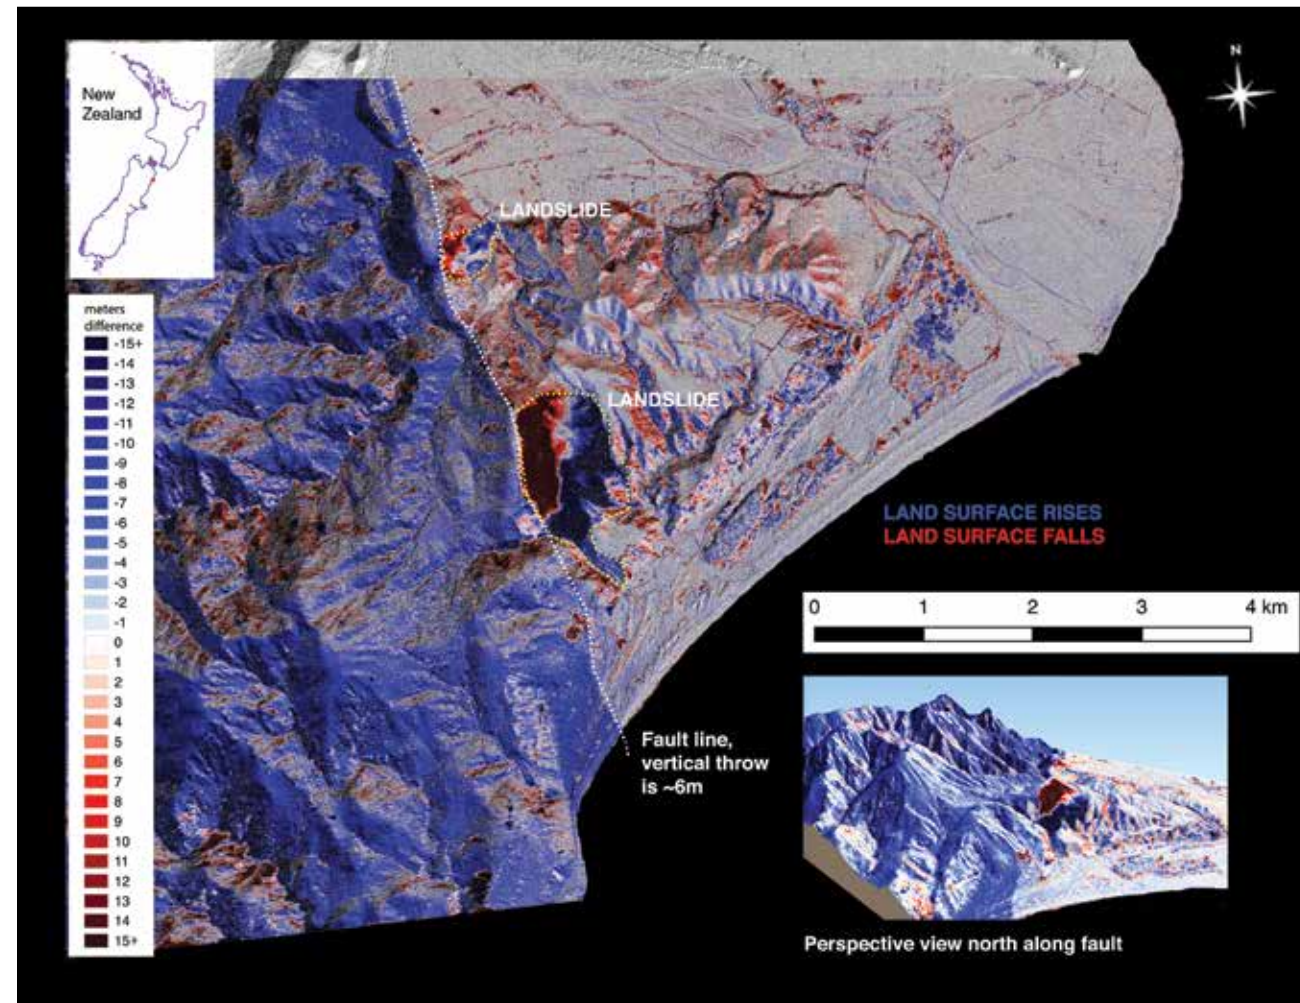


Figure 1: Elevation change from differenced digital surface models of the area near Waipapa Bay, northeastern South Island, New Zealand. Red regions—elevation drop in response to the earthquake. Blue regions—elevation gain. Landslides (yellow dots) are prominent. Satellite measurements of landslide volumes are identical to those made using UAV (unmanned aerial vehicle).

The November 14, 2016, Mw7.8 Kaikoura, New Zealand, event represents the first opportunity to execute such efforts for a regional scale, large-magnitude earthquake because a data for a large area of New Zealand were collected with stereo satellite imagery following the 2011 Christchurch earthquakes.

RESEARCH CHALLENGE

Regional assessment of landsliding and infrastructure damage in the aftermath of large earthquakes is a broad societal problem involving loss of life and property. Because hazard assessment is limited by access and resources in the immediate aftermath of an earthquake, satellite-based photogrammetry methods using high-resolution imagery and high-performance computing potentially provide an avenue to rapidly assess land surface changes and infrastructure damage in a cost-effective and time-sensitive manner.

Whereas commonly used radar-based methods rely on coherent ground deformation for constraints on finite fault slip models, these techniques break down where there is surface disruption, such as surface faulting, landsliding, and infrastructure damage. New methods using photogrammetry techniques take advantage of in-line stereo imagery capture of high-resolution (0.5 m Ground Sampling Distance) Worldview™ satellite data. Digital surface models derived from these data provide valuable information for quantitative assessment of surface change. Immediate collection preserved these data in the most pristine form, and regional assessment provides a synoptic view not provided by typical post-event land- or airborne-based surveys. The third and fourth dimension provided by satellite-derived surface models will significantly advance auto-recognition algorithms by combining image analysis with 3D surface data.

METHODS & CODES

We want to understand how the landscape evolves in time after an earthquake, examining fine-scale land-surface rupture and landslide processes. We do this by making repeat models of the terrain in northeast South Island New Zealand from a series of stereo satellite images collected both before and after the Kaikoura Earthquake. We use the Surface Extraction with TIN-based Search-space Minimization (SETSM) algorithm, developed at The Ohio State University, in a high-performance computing (HPC) workflow that bulk corrects hundreds of raw stereo satellite imagery pairs for image distortions and then performs a photogrammetric analysis. We order the surface model output in time and then use co-registering routines to make sure that the models are then further oriented in the right place in space. With this time series of landscape change we can see the immediate effects of the earthquake, how faults move in three dimensions, and how the landscape adapts to a huge number of mass flow events.

RESULTS & IMPACT

Our project has the potential to transform the way that first responders are provided information about a disaster zone. With

rapid, reactive HPC access and processing and an initial pre-event data set for areas of the globe that are likely, or have already been exposed to natural disasters, we have the ability to provide damage maps and cascading hazard monitoring in a space of hours to days. This overview information is critical in the chaotic aftermath of an event. Scientifically, we are able to study the immediate evolution of surface ruptures and landslides and provide information to geo-mechanical modelers whose work informs safety regulations and building codes.

WHY BLUE WATERS

Blue Waters-scale computing was necessary to complete the project in a timeframe useful to coordinate perishable field data collection and provide rapid feedback on the event history. Maps and volume estimates of current slope failures estimated from the surface model will immediately assist New Zealand partners in monitoring and assessing dam hazards and slope reactivation, which will continue to affect the area for months to come. Our code and workflow is optimized for Blue Waters and allowed us to instantly activate the processing chain, which would likely have taken weeks to months on cloud resources.

IMAGE PROCESSING TO BUILD A MULTI-TEMPORAL VEGETATION ELEVATION MODEL (MTVEM) OF THE GREAT LAKES BASIN (GLB)

Allocation: GLCPC/540 Knh

PI: Jennifer Corcoran¹

Co-PIs: Brian Huberty², James Klassen³, Keith Pelletier¹

Collaborators: Paul Morin¹, Joe Knight¹

¹University of Minnesota

²U.S. Fish & Wildlife Service

³SharedGeo

EXECUTIVE SUMMARY

Deriving ecosystem characteristics from stereo imagery has been used for decades, yet frequent 3D digital mapping and monitoring has not been feasible until very recently. By more frequently monitoring surface changes in the vegetation throughout the growing season in the Great Lakes Basin with satellite remote sensing techniques, land managers will be supported with new, enhanced information to address emerging stand-to-landscape scale changes in ecosystem habitats. The amount of data in a study area of this size needed to be processed and analyzed is well beyond those available from academic, private, and government systems. This computational reality is precisely why we need a leading-edge petascale resource such as Blue Waters. The results from this project will enable scientists to understand their land with greater detail in three dimensions, thereby making management, conservation, and protection of important ecosystems as modern and effective as possible, which will enhance the nation's natural resources for future generations.

RESEARCH CHALLENGE

Ecosystem management requires knowing the type, size, structure, and density of vegetation over time. These important features need to be repeatedly mapped. Stereo submeter, optical satellite imagery, and derived surface vegetation models can be used to better characterize these features, and their changes over time, with the added dimension of height. High-resolution vegetation surface canopy mapping over large geographic regions, such as the Great Lakes Basin (GLB), has never been obtained from either aerial or satellite surveys. Additionally, the binational management (Canada and U.S.) of the GLB limits consistent, repeatable coverage by either country working independently. While a few scattered vegetation surface models exist from expensive airborne-active laser sensors (LiDAR) within the GLB, these datasets represent a single date and were not planned as continuous, basinwide acquisitions. Not having this information limits the kind of science that can be done to address the multitude of questions that surround the ecosystems of the GLB.

The question remains: How are the ecosystems of the GLB changing and what can we, as a society, do about it? Continuous monitoring of surface elevation will detect both natural changes

(such as flooding, forest blowdown, fire, insects, and disease outbreaks) and anthropogenic changes (such as harvest and land cover change). Further, MTVEM will improve habitat and biological modeling. Finally, MTVEM will be used binationally to better visualize canopy change in forest habitats and freshwater wetland resources within the Great Lakes Basin.

METHODS & CODES

We currently have over 8,000 satellite image stereo pairs available and we estimate there will be an additional 5,300 pairs to be collected over the next 18 months, totaling about 13,300 image stereo pairs. Each stereo pair task is run on a single node, submitted in batches of 2 to 8 tasks per job. Complete processing of one stereo pair to 2m takes on average about 36 node hours, totaling about 483,590 node hours. We are adding ~10% time in addition for unforeseen events. Thus we asked for an allocation of 540,000 node hours. We will prioritize images based on science value to fit into our allotment, where the primary concentration is on the Great Lakes Basin but will extend temporal and geographical footprint as efficiencies and capacities improve.

RESULTS & IMPACT

As the data are processed, the resulting surface canopy models will become openly available through partner's online distribution systems, such as NOAA's Digital Coast and GEOSS Portal (www.geoportal.org). The final product, a seamless and registered surface elevation model (MTVEM) of the GLB, will enable a large range of science activities at substantially higher resolution than currently available. These canopy maps and change-detection products will provide positional accuracies of less than a couple meters with added ground control points. We have processed our entire archive of satellite image stereo pairs once. We are assessing change in priority GLB areas where LiDAR digital surface models from 2011 (DSM) are available (Figs. 1 and 2). Both gains and losses in vegetative cover over the five-year difference in the acquisition were detected in these examples. In addition, we are able to begin looking at seasonal differences by processing surface models from satellite stereo image pairs from within a single growing season. These preliminary results show great promise for providing valuable data to a myriad of coastal and terrestrial ecosystem

science research questions that need to be addressed across the entire GLB.

WHY BLUE WATERS

Stereo satellite imagery allows for the generation of highly accurate surface elevation models and we have already tasked stereo-mode acquisition through Digital Globe over the entire GLB. Processing of this dataset will require an allocation well beyond those available from most academic, private, and government

HPC systems, including the standard XSEDE allocations, which is precisely why the leading-edge petascale resource Blue Waters is necessary to address this research.

PUBLICATIONS AND DATA SETS

DigitalGlobe (2016), WorldView-3 scene 104001001C179100, Level Standard 2A, DigitalGlobe, Longmont, Colorado, 04/30/2016.

DigitalGlobe (2016), WorldView-3 scene 10400100227DC800, Level Standard 2A, DigitalGlobe, Longmont, Colorado, 10/03/2016.




Figure 1: Duluth, Minn. On the left is a digital surface model difference map, using a LiDAR-derived surface model acquired in Spring 2011 and submeter stereo imagery-derived surface model from Spring 2016. The optical image on the right is from the same acquisition as the digital surface model in Spring 2016. Imagery courtesy of DigitalGlobe and the NextView program.



Figure 2: Duluth, Minn. On the left is a digital surface model difference map, using a LiDAR-derived surface model acquired in Spring 2011 and submeter stereo imagery-derived surface model from Fall 2016. The optical image on the right is from the same acquisition as the digital surface model in Fall 2016. Imagery courtesy of DigitalGlobe and the NextView program.

THE TERRA DATA FUSION PROJECT

Allocation: Blue Waters Professor/180 Knh

PI: Larry Di Girolamo¹

Collaborators: Guangyu Zhao¹, John Towns², Shaowen Wang², Yan Liu², Kent Yang³

¹University of Illinois at Urbana-Champaign

²National Center for Supercomputing Applications

³The HDF Group

EXECUTIVE SUMMARY

The Terra Data Fusion Project that we initiated in 2015 continues to make significant progress through collaborative efforts among NASA, HDF Group, and NCSA. Large Terra data transfers (~1PB) involving the Blue Waters team and NASA are now complete. The software tool that fuses all the Terra radiance granules was developed and has produced one year of basic fusion data on Blue Waters. Besides the scientific usage demonstrated in the previous reports, the dataset has been further used to (1) characterize ice crystal roughness of cirrus clouds, resulting in a better understanding of ice cloud optical properties, and (2) gauge the effects of cloud heterogeneity on microphysical retrievals, concluding that cloud spatial heterogeneity is insufficient to explain the large biases in the MODIS *Re* standard product.

RESEARCH CHALLENGE

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

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

Solutions to these questions will: 1) facilitate greater ease in creating new geophysical retrieval algorithms that provide greater accuracy than the current single-instrument algorithms, 2) provide an easy mechanism for users to access and process the

entire Terra record, (3) greatly reduce error and redundancy in the science community among those researchers using multiple instrument datasets, (4) provide greater insight into geophysical processes through synergistic use of fusion products, and (5) provide a framework for fusion that could extend to other NASA missions and constellations. The end result will facilitate discovery and accelerate progress in Earth Science research. Use cases are presented below.

METHODS & CODES

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

Thus far, key steps (1) and (2) are accomplished. We successfully transferred mission-scale radiance data from all the five Terra instruments to Blue Waters. We built a software tool to merge all the Terra radiance granules into one basic fusion granule, which contains not only radiance measurements but also their uncertainties, geolocation, sun-view geometry, and observational time. The tool has produced one year of basic fusion data on the Blue Waters. We are currently working on key step (3) to prepare for mission-scale processing with help from the Blue Waters team. The overview of this project and progress report was given at the MISR science team meeting and data users' science symposium [3].

Initial science investigation using this Terra fusion dataset was carried out in two studies. Fusing MISR and MODIS products, we characterized the effective radius (*Re*) of the liquid cloud drop size distribution over oceans. Details were given in *Blue Waters 2015 Annual Report* and in [4]. Our results paint a radically different picture of the distributions of cloud drop sizes in our atmosphere compared to what was previously determined from the original MODIS data; e.g., 2 to 11 μm differences in the zonal means. We also processed the 15-year MISR data to examine the decadal trends in the Earth's color and texture [5]. Globally, we show that the Earth has been appearing relatively more blue (up to 1.6% per decade from both nadir and oblique views) and smoother (up to 1.5% per decade, only from oblique views) over the past 15 years.

RESULTS & IMPACT

We continue to explore the scientific usage of this dataset in two new studies. The first was that we used the dataset to generate one full year (2013) of MISR+MODIS cloud-element fusion data (~50 terabytes) and share the data with Prof. Ping Yang's group at the Texas A&M University via GridFTP. The fusion data were used to build ice crystal roughness parameters that are important to our meteorological understanding of ice clouds. Our intermediate results were presented at recent meetings and conferences [6–7]. Further scientific analyses are being conducted.

The second was that we examined the effect of cloud spatial heterogeneity on satellite-retrieved liquid cloud *Re* using 13 years of MISR and MODIS fusion for the month of January. We did find some dependency of *Re* on cloud spatial heterogeneity. Further stratification of the biases by both cloud spatial heterogeneity and cloud optical depth for various solar zenith angles and geolocations reveals that more heterogeneous and optically thick clouds tend to have larger *Re* bias, indicating cloud spatial heterogeneity is insufficient to explain the large biases in the MODIS *Re* standard product. Our results were presented at the MISR data users' science symposium [8].

WHY BLUE WATERS

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

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

PUBLICATIONS AND DATA SETS

Liang, L., L. Di Girolamo, and W. Sun, Bias in MODIS cloud drop effective radius for oceanic water clouds as deduced from optical thickness variability across scattering angles. *J. Geophys. Res. Atmos.*, 120:15 (2015), DOI: 10.1002/2015JD023256. [The main data set applied to this study is the fusion product generated by fusing the MISR Level1B radiance product (Version F03_0024) and the MODIS level2 cloud product (Collection 6).]

Zhao, G., et al., Regional changes in Earth's color and texture as observed from space over a 15-year period. *IEEE Trans. Geosci. Remote Sens.*, 54:7 (2016), DOI: 10.1109/TGRS.2016.2538723. [The main data set applied to this study is the MISR level1B radiance product (Version F03_0024).]

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

Allocation: Illinois/200 Knh
PI: Marcelo H. Garcia¹
Co-PI: Paul Fischer¹
Collaborator: Som Dutta¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Bifurcations are fundamental features of all river systems. This study focuses on a specific class of bifurcations called diversions. Experiments have shown that the distribution of near-bed sediment between the downstream channels at a diversion is not proportional to the water flow distribution, with a disproportionately higher amount of sediment going into the lateral channel. A better understanding of the aforementioned non-linear phenomena will help in efficient design of river diversions. Diversions are used for navigation and flood mitigation, and have also been put forth as a solution for reclaiming deltas going under the sea due to increasing sea level. The current study investigates the mechanisms behind this phenomenon through Large Eddy Simulations (LES) and Direct Numerical Simulations (DNS) of the flow in idealized diversions of different configurations, with the sediment being modeled as Lagrangian




Figure 1: Instantaneous velocity magnitude at the bottom of the channel for diversion angle of (a) 30-degrees and (b) 150-degrees. For both cases, the Reynolds number of the flow is 25,000, and the flow is divided equally between the two channels. One can clearly observe that most of the flow near the bottom enters the lateral channel, even though the total flow is equally divided between the two channels.

particles. The simulations have been conducted using a highly scalable spectral element based incompressible Navier-Stokes solver, Nek5000. The simulation conditions are comparable to the laboratory experiments, which make these simulations large enough to require the use of Blue Waters. The simulation results provide new insight into the mechanism behind the phenomenon, and has identified yet unexplored parameters that influences the distribution of sediment at a diversion.

RESEARCH CHALLENGE

Bifurcations are fundamental features found in river systems. This current study focuses on a specific class of bifurcations called diversions, where one source channel splits into two: a main channel and a lateral branch. Laboratory experiments conducted by Bulle in 1926 [1,2] and later by other investigators, have shown that the distribution of near-bed sediment between the downstream channels is not proportional to the flow distribution, with a disproportionately higher amount of sediment going into the lateral channel. This non-linear phenomenon is often referred to as the Bulle Effect. The current study investigates the mechanisms behind this phenomenon through high-resolution numerical simulations of the flow and sediment transport at the scale, and for the configurations similar to Bulle's experiments.

A better understanding of the aforementioned phenomenon will help in efficient design of river diversions, which among uses like navigation and flood-mitigation have also been put forth as a solution for reclaiming deltas sinking under the sea due to rise in sea level, which is a consequence of anthropogenic climate change [3]. A prime example is the Mississippi River delta, for which different potential diversion designs are being studied for diverting water and sediment from the Mississippi River [4]. Better understanding of the phenomenon will eventually help in more accurate prediction of the short and long-term geomorphological evolution of river bifurcations, thus furthering the state of art in the field of river mechanics. It will also provide insights that will help improve reduced-order models of flow and sediment transport at bifurcations. Better understanding of the fundamental mechanism behind Bulle-Effect will also help shed light on vorticity-driven transport at bifurcations in human body (e.g. the carotid bifurcation [6]) and manmade systems (e.g. grit chambers in water reclamation plants [5]).




Figure 2: The figure illustrates the division of uniformly distributed (in the water column) neutrally buoyant fine particles between the two channels, for Reynolds number of 25,000 and 50:50 flow division. The particles in the panels (a, b) are colored according to their position in the water column. Particles near the bottom are blue and those near the top are yellow. (c) shows that most of the blue particles have moved into the lateral channel and most of the yellow particles have continued into the main channel.

RESULTS & IMPACT

High-resolution Large Eddy Simulations (LES) and Direct Numerical Simulations (DNS) of the flow at different configurations of the idealized diversions have been conducted, with sediment being modeled as Lagrangian particles. A simulation being LES or DNS depended on the Reynolds number of the flow, as simulations were conducted for a range of Reynolds numbers 10–25,000. For the cases with Reynolds number of 25,000 the conditions are comparable to the laboratory experiments. The simulations have been conducted using the open-source, spectral element based higher-order incompressible Navier-Stokes solver Nek5000 [7]. The Spectral Element Method (SEM) combines the accuracy of spectral methods and the flexibility of Finite Elements Method [8]. Sediment transport in the flow has been modeled using Lagrangian particle tracking. For simulating transport of poly-disperse sediment particles efficiently, a semi-implicit Lagrangian particle algorithm was developed for the current study [1].

Initial results had shown that for a 90-degree diversion, and the flow being equally divided between the channels, the simulation was able to capture the preferential movement of bedload sediment into the lateral-channel [2]. Hence, simulations were conducted for a range of Reynolds number and flow divisions, further confirming the primary driver of the phenomenon that is preferential movement of the flow near the bottom of the channel into the lateral-channel [3,4]. This characteristic of the flow was also observed for different diversion angles (e.g. for 30,150 degrees see fig. 1).

Simulations were also conducted for sediment with different diameters, which made them travel at different depths in the water column. It was found that the phenomenon was not just valid for sediment traveling as bedload, but for any sediment traveling at the lower 30 percent of the water column. To illustrate that, thousands of neutrally buoyant fine particles were released and their paths monitored. These particles by design will follow the flow, thus providing a sense of how the flow divides between the two channels. Figure 2 illustrates one of the above-mentioned

cases, where most of the particles starting near the bottom of the channel can be observed to primarily move into the lateral-channel. This study not only provides new insights into the hydrodynamics and sediment transport at bifurcations, it also shows that high-resolution LES can be used to study complex river-mechanics problems.

WHY BLUE WATERS

The current study pushes the limit of the scale at which high-resolution LES have been used to study complex multi-phase river mechanics problems, warranting the use of a computational resource that can provide sustained computing power at an unprecedented scale, thus the need to use Blue Waters. For the current study simulations have been conducted for up to 243.648 million computational points, with the code scaling strongly up to 32,768 MPI ranks. Without access to supercomputers like Blue-Waters, completing the study within a realistic timeframe would be impossible. Visualization of a phenomenon is an effective way to understand (and explain) it's mechanics, thus we are currently working with Blue Waters project staff to create an animation of the phenomenon, using data from one of the simulations.

PUBLICATIONS AND DATA SETS

S. Dutta, P. Fischer, M.H. Garcia, A novel semi-implicit Lagrangian particle tracking model for efficient computation of particle transport at low Stokes number, under preparation for Journal of Scientific Computing, Springer.

S. Dutta, P. Fischer, M.H. Garcia, Large Eddy Simulation (LES) of flow and bedload transport at an idealized 90-degree diversion: insight into Bulle-Effect. In River Flow 2016, CRC Press (2016).

S. Dutta, P. Fischer, M.H. Garcia, Direct Numerical Simulation of flow at a 90-degree asymmetric bifurcation for different flow splits, under preparation for Journal of Fluid Mechanics.

S. Dutta, P. Fischer, M.H. Garcia, Flow dynamics and fine particle transport at an asymmetric 90-degree bifurcation: insights into the Bulle-Effect under laminar inflow, under preparation for Physical Review Fluids.

THE REFERENCE ELEVATION MODEL OF ANTARCTICA

Allocation: Innovation and Exploration/1,200 Knh

PI: Ian Howat¹

Collaborators: Paul Morin², Claire Porter²

¹The Ohio State University

²Polar Geospatial Center of the University of Minnesota

EXECUTIVE SUMMARY

The goal of this project is to create the Reference Elevation Model of Antarctica (REMA), a continuous, high-resolution (8 m), high-precision (accuracy better than 1 m) reference surface for a wide range of glaciological and geodetic applications. REMA will be constructed from stereo-photogrammetric Digital Elevation Models (DEM) extracted from pairs of sub-meter resolution Worldview satellite imagery and vertically registered using ground control from both GPS (Global Positioning System) surveyed points and coordinated airborne LiDAR (Light Detection and Ranging) surveys by Operation IceBridge of the U.S. National Aeronautics and Space Administration. The Worldview imagery is archived and provided at no cost from the Polar Geospatial Center (PGC), which will also openly distribute REMA through its web data portal. REMA processing will be done using fully automated DEM extraction and coregistration software on the Blue Waters supercomputing system.

RESEARCH CHALLENGE

Accurate surface elevation is an essential dataset for glaciology, required for mapping bed topography from ice thickness; measuring ice thickness changes; constraining ice flow and geodynamic models; mapping glacial geomorphology, terrain corrections, and filtering of remote sensing observations; and many other science tasks. It is also critical for mapping ice traverse routes, landing sites, and other field logistics planning. Continuous DEMs of the continent, however, have spatial resolutions of hundreds of meters or more, accuracies of tens of meters in coastal and mountainous areas, and no definitive time stamping for change detection, limiting their utility. The growing archive of sub-meter stereo imagery held by the Polar Geospatial Center is arguably among the most underutilized collections of satellite data in existence, with nearly complete coverage of the polar regions below 85° of latitude. The project will be the first continental-scale application of these data for terrain mapping, upscaling an established processing and distribution system that has been used to map large regions of the Arctic successfully, and will be the next step toward eventual global mapping.

METHODS & CODES

In collaboration with PGC, our team has spent three years developing an efficient algorithm for constructing photogrammetric DEMs from satellite imagery with the objective of creating a fully

automated system capable of handling large amounts of data. Development of the Surface Extraction from TIN-based Search-space Minimization (SETSM) algorithm was begun to facilitate an automated processing pipeline for the PGC operations. SETSM DEMs have been extensively validated; SETSM is node parallelized using OpenMP and has been applied to processing large-area DEM mosaics in proof-of-concept studies. SETSM's structure uniquely eliminates the need for an existing (i.e., "seed") DEM for *a priori* constraint or any data-specific, user-defined search parameters, making it a truly automated algorithm. SETSM is called from a single command line with the only required inputs being the filenames of the two stereo images and the RPC (Rich Photorealistic Content) file, typically provided in XML (eXtensible Markup Language) format. SETSM is written in stand-alone C code with no external dependencies and requires no libraries, ensuring simple, multi-platform installation, support, and optimization. SETSM is available as open source on Github.

RESULTS & IMPACT

Antarctica's ongoing rapid changes—and potential for near-future change—are of global concern due to increased rates of sea level rise. REMA will provide a benchmark for detecting these changes. It will be useful for a wide range of applications beyond glaciology, ranging from geodynamics to logistics planning.

As we near the end of the first year of the project, we are well ahead of schedule in our activities thanks to receiving a 1.2-million node-hour "innovation award" allocation for the Blue Waters supercomputer at the National Center for Supercomputer Applications (NCSA). This allocation has allowed us to complete a full, double-coverage processing of the entire Antarctic continent and subantarctic islands to 8-m resolution, totaling nearly 0.5 million individual DEMs, and resulting in 98% coverage prior to filtering. Using the remaining allocation, we are now selecting gaps in coverage due to both prior data unavailability or cloud cover, and we are processing additional coverage to fill gaps.

Processed data are undergoing manual quality control by our student assistants. In addition, we are implementing automated filtering and processing methods to improve the efficiency and consistency of the filtering. We have also finalized our algorithms for mosaicking data into seamless elevation model tiles for distribution. As a test, we produced a prototype set of tiles for the Thwaites Glacier area.




Figure 1: Hill shade representation of a preliminary Reference Elevation Model of Antarctica (REMA) consisting of over 500,000 individual 8-m resolution elevation models processed from satellite imagery on Blue Waters.

In collaboration with colleagues at the University of Washington, we have successfully implemented a method for registering the DEMs to Cryosat-2 radar altimetry, providing an elevation accuracy within the 1-m specification for REMA. Cryosat-2-derived registration has been tested against airborne LiDAR elevations. The spatial and temporal coverage of Cryosat-2 makes it an ideal source of DEM registration.

Finally, we have worked with the PGC to develop a web distribution system that will be implemented soon.

WHY BLUE WATERS

Our processing requires scheduling and rapid throughput of thousands of individual jobs, each with highly variable and unpredictable wall times. Only Blue Waters has the capacity to provide efficient throughput of this job volume. Further, our processing was able to take advantage of Blue Water's backfill scheduling capabilities, further increasing efficiency and system utilization.

UNTANGLING ENTRAINMENT AND PRECIPITATION IN CONVECTIVE CLOUDS

Allocation: Blue Waters Professor/250 Knh
PI: Sonia Lasher-Trapp¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Internal cloud circulation patterns introduce dry air inward from outside the cloud, which is called entrainment. Its effects can limit storm development, longevity, and various interdependent microphysical processes that may ultimately produce precipitation. Our understanding of entrainment and precipitation links has been limited in the past by inadequate model resolution. We are using high-resolution 3D simulations of convective clouds and storms, along with our diagnostic entrainment algorithm, to untangle the intricate web of connections between entrainment and its effects upon the generation of precipitation. Our latest results show that closer spacing between storms along a storm line can initially delay and decrease precipitation by competing for air

flowing into the cloud bases; decreases in entrainment may not appear until the spacing is so small that the storms' edges are no longer distinct. If ultimately the closer-spaced storms can generate even weak precipitation outflows, they can combine due to their proximity to produce a new generation of stronger storms that precipitate much more.

RESEARCH CHALLENGE

Deep convective clouds produce the majority of the earth's precipitation, and yet it is difficult to predict whether developing cumulus clouds will attain the depth and longevity required to produce heavy rainfall and/or become severe thunderstorms. Entrainment is the term for the process by which the circulations within clouds bring dry air from outside the cloud inward. In

time, entrainment not only reduces the cloud buoyancy, limiting its vertical development, but also depletes its liquid water by evaporation, limiting precipitation formation. A long-standing problem in meteorological models has been to understand why they tend to predict rain formation too early and in excessive amounts. While others are researching what details of precipitation processes might be incorrectly represented in models, our approach is to investigate if the under-prediction of entrainment could explain the over-eagerness of the models to produce convective rainfall. This problem affects a broad range of atmospheric science problems, ranging from short-term weather forecasts from numerical weather prediction models to climate forecasts from regional and global climate models.

METHODS & CODES

We are using the National Center for Atmospheric Research's CM1 model [1] to simulate convective clouds and storms at high resolution by employing its MPI capabilities on the many nodes available on Blue Waters. We make use of the National Severe Storms Laboratory microphysical scheme [2] within CM1. The simulations are conducted in both idealized and realistic environments. We evaluate entrainment with our own code [3] that calculates mass fluxes into the user-defined core of the cloud as the clouds and storms evolve and relate this entrainment to the ability of the clouds to produce precipitation and hail. We relate the calculated entrainment to the storm longevity and the amount of precipitation it produces.

RESULTS & IMPACT

Many atmospheric scientists have worked to try to produce parameterizations (larger-scale approximations) for cumulus entrainment, when the models they use for daily weather prediction as well as regional and global climate employ resolutions that are not high enough to represent all the cloud motions explicitly. However, these efforts are greatly hampered when we do not fully understand the process of cumulus entrainment. If more details about entrainment and its effects on precipitation can be discovered, then we will also know how best to represent its effects in larger-scale models.

We are tackling this problem using multiple approaches, including studying the process of entrainment in an individual thunderstorm at its earliest stages (Fig. 1), as well as in a line of thunderstorms.

Our latest results show that closer spacing among storms along a line can initially delay and decrease precipitation by competing for air flowing into the cloud bases; the expected decreases in entrainment for closely spaced clouds may not appear until the spacing is so small that the storms' edges are no longer distinct (and thus the storms engulf cloudy mixtures of air from surrounding storms). If ultimately the closer-spaced storms can generate even weak precipitation outflows, they can combine due to their proximity (Fig. 2) and create focal points for generating new storms that may precipitate much more heavily.




Figure 2: Vertical cross-section through a portion of a line of convective clouds that are precipitating. Rain mass indicated by color bar below; wind motions and strength denoted by direction and magnitude of plotted vectors. Rain precipitating from adjacent clouds form downdrafts that may collide near the ground to form new strong storm updrafts in between the original clouds.

WHY BLUE WATERS

Our Blue Waters allocation is essential for achieving the high resolution required within a given simulation to properly represent the smaller cloud motions that can still be important for entrainment but over the larger domains required for thunderstorms and groups of thunderstorms. 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. The hardware needed to run these kinds of simulations quickly exceeds the limits of most computers. Blue Waters staff have helped us to learn new and practical ways to visualize the output for easier analysis.




Figure 1: Visualization of the 3D core of a developing thunderstorm. Warm colors (yellow, red) denote areas of greater amounts of precipitation mass; cool colors (green, blue) denote areas with little water mass remaining. We are studying how (likely limited) entrainment resulting from the storm internal motions allows the greater precipitation mass to exist higher in the cloud. Once it falls out (not shown), the upper portions of the cloud exhibit multiple thermal circulations (nodes along the top of the storm).

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

Allocation: NSF PRAC/880 Knh

PI: Lijun Liu¹

Collaborator: Manuele Faccenda²

¹University of Illinois at Urbana-Champaign

²Università di Padova, Padova, Italy

EXECUTIVE SUMMARY

We use numerical models with data assimilation to reproduce the complex processes of Earth dynamic evolution, which requires consideration of large amounts of data distributed over many computer nodes. With Blue Waters, our research projects during the past year include: 1) investigating mantle dynamics behind the widespread intra-plate volcanism associated with the Yellowstone hotspot, 2) simulating subduction and mantle flow beneath South America since the Mesozoic, and 3) measuring the strength of the continental lithosphere by approximating its electrical conductivity as effective viscosity based on geodynamic modeling. By now, our Yellowstone work has been published in *Geophysical Journal Research*, which received extensive media attention, including from *Scientific American*, *Science Daily*, *Science News*, etc. Work on South American subduction has led to three papers in *Earth &*

Planetary Science Letters, which was also reported in *Science Daily*, *Science News*, etc. Last but not least, our project on quantifying lithosphere viscosity was published in *Science* magazine, which also generated many news pieces.

RESEARCH CHALLENGE

Dynamic processes of the deep Earth such as heat source fueling of volcanoes far away from a subduction zone, the influence of subduction on the mantle and the surface, and earthquake- and volcano-generating deformation of continents, are complex in nature. A quantitative understanding of these processes, especially their history, is vital for explaining the various geological, geophysical, and geohazard observations. The challenge in progress is the overall inaccessible nature of the present deep mantle, and more so during the geological past.

METHODS & CODES

To quantitatively understand the current and past dynamic processes within the deep Earth, we adopt various data assimilation techniques in large-scale geodynamic modeling. For example, we employ seismic and magnetotelluric tomography images to represent the present-day mantle temperature and viscosity profiles. We assimilate the motion of tectonic plates that are geologically reconstructed as velocity boundary conditions of a geodynamic model. We combine these data into a single physics-based numerical model using either sequential (forward) or adjoint (inverse) data assimilation techniques. The software package we use is a well-benchmarked finite element code, CitcomS [1]. We have further expanded the scalability and I/O performance of this code on Blue Waters.

RESULTS & IMPACT

In our first project, we evaluated the traditional hypothesis that Yellowstone volcanism is caused by a deep-rooted mantle plume [2]. We found that this hypothesis is flawed, because a hot upwelling plume is always blocked from above by the sinking oceanic plate. Therefore, Yellowstone has to have been formed by a different mechanism. This work challenges the traditional view on this important topic of intra-plate volcanism.

Our second project is to reproduce the subduction and mantle flow history below South America since 100 mega annum (100 million years). We found that the seismically observed flat slabs




Figure 2: Lithosphere viscosity (color) and mantle flow (arrow) beneath the western United States. The viscosity is converted from electrical resistivity, and the flow is calculated from a geodynamic model.

beneath the continent are mostly due to the subduction of thick and buoyant seafloor anomalies (oceanic plateaus) [3]. We also found that the formation of these flat slabs tears the slab itself, leading to the absence of earthquakes within the flat portion and formation of special volcanisms above [4]. The slab-induced mantle flow dominates the deformation of the surrounding mantle (Fig. 1), as confirmed by the model predicting the observed seismic anisotropy in the region [5].

Our third project concerns the mechanical strength and its lateral variation of continental lithosphere [6]. We demonstrate that the electrical conductivity of rocks is an excellent proxy for the effective viscosity of the lithosphere (Fig. 2). Consequently, this finding provides a practical way to “measure” the strength of the Earth’s rigid outer shell that is otherwise difficult to estimate. This also opens a new direction of geodynamic research.

WHY BLUE WATERS

The enormous amount of data processing and computation makes Blue Waters the best platform. The code CitcomS is community-based software, and this code has been designed and tested mostly on traditional supercomputers.

PUBLICATIONS AND DATA SETS

Hu, J., M. Faccenda and L. Liu, Subduction-controlled Mantle Flow and Seismic Anisotropy in South America. *Earth Planet. Sci. Lett.*, 470 (2017), pp. 13–24.

Liu, L., and D. Hasterok, High-Resolution Lithosphere Viscosity and Dynamics Revealed by Magnetotelluric Tomography. *Science*, 353 (2016), pp. 1515–1519.

Hu, J., and L. Liu, Abnormal Seismological and Magmatic Processes Controlled by the Tearing South American Flat Slabs. *Earth Planet. Sci. Lett.*, 450 (2016), pp. 40–51.

Leonard, T., and L. Liu, The Role of a Mantle Plume in the Formation of Yellowstone Volcanism. *Geophys. Res. Lett.*, 43 (2106), DOI: 10.1002/2015GL067131.

Hu, J., L. Liu, A. Hermosillo, and Q. Zhou, Simulation of Late Cenozoic South American Flat-Slab Subduction Using Geodynamic Models with Data Assimilation. *Earth Planet. Sci. Lett.*, 438 (2016), pp. 1–13, 2016.

Heller, P., and L. Liu, Dynamic topography and vertical motion of the U.S. Rocky Mountain Region prior to and during the Laramide Orogeny. *GSA Bull.* (2015), DOI: 10.1130/B31431.1.

Liu, L., The Ups and Downs of North America: Evaluating the Role of Mantle Dynamic Topography since the Mesozoic. *Rev. Geophys.*, 53 (2015), DOI: 10.1002/2015RG000489.

Liu, L., and Q. Zhou, Deep recycling of oceanic asthenosphere material during subduction. *Geophys. Res. Lett.*, 42 (2015), DOI: 10.1002/2015GL063633.

Liu, L., Constraining Cretaceous subduction polarity in Eastern Pacific from seismic tomography and geodynamic modeling. *Geophys. Res. Lett.*, 41 (2014), DOI: 10.1002/2014GL061988.




Figure 1: Geometry of the subducting Nazca slab (3D colored surface) and mantle flow (3D arrows) beneath the southern Atlantic. The regional mantle deformation represented by flow vectors is dominantly controlled by the subduction of the slab.

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

Allocation: NSF PRAC/6,600 Knh

PI: Philip J. Maechling¹

Co-PIs: Yifeng Cui², Kim Bak Olsen³, Ricardo Taborda⁴

Collaborators: Jacobo Bielak⁵, Alexander Breuer², Scott Callaghan¹, Ewa Deelman¹, Edward Field⁶, Jacquelyn Gilchrista¹, Christine Goulet¹, Robert Graves⁶, Naeem Khoshnevis⁴, Kevin Milner¹, Doriam Restrepo⁷, Daniel Roten³, William Savran³, Bruce Shaw⁸, Karan Vahi¹, John Vidale¹, Kyle Withers³

¹University of Southern California

²San Diego Supercomputer Center

³San Diego State University

⁴University of Memphis

⁵Carnegie Mellon University

⁶United States Geological Survey

⁷Universidad EAFIT, Medellin, Colombia

⁸Columbia University

EXECUTIVE SUMMARY

A research collaboration, led by the Southern California Earthquake Center (SCEC), which includes earth scientists, engineers, and computer scientists, used Blue Waters to run physics-based earthquake simulations that improve our understanding of earthquake processes and their effects on seismic hazard. SCEC's earthquake system science research program develops detailed earth models and high-performance computing software needed to perform realistic, physics-based earthquake and ground motion simulations. This past year, the SCEC team used NCSA Blue Waters to develop more accurate and scalable computational models of earthquakes and to calculate the first physics-based probabilistic ground motion forecasts for Central California.

RESEARCH CHALLENGE

Probabilistic Seismic Hazard Analysis (PSHA) [1] is the scientific framework for many seismic and risk-related engineering and social applications, including performance-based design, seismic

retrofitting, resilience engineering, insurance-rate setting, emergency response, and public education. The U.S. Geological Survey (USGS) currently uses empirical PSHA to promote seismic safety engineering and disaster preparedness across the United States, including California. SCEC's research goal is to develop physics-based seismic hazard models for California and elsewhere that are more accurate than the empirical USGS National Seismic Hazard Map Project [2] standard models. Our long-term goal is to extend physics-based PSHA across the full bandwidth needed for seismic building codes and other purposes.

METHODS & CODES

This year, SCEC researchers added improved physics into our wave propagation software and improved our software's performance on CPUs (central processing units) and GPUs (graphics processing units). For high-frequency ground motion simulations, our codes must model frequency-dependent attenuation [3], free-surface topography [4], and nonlinear yielding effects [5]. With improved codes and support through the Blue Waters PAID program, we performed the first 4-Hz nonlinear magnitude 7.7 earthquake simulation using 4,200 GPUs on Blue Waters [5,6] using a highly optimized implementation of a nonlinear computational method developed by SCEC researchers. We continued to validate our software by simulating well-recorded historic California earthquakes and comparing our simulations against the recorded ground motions [7].

Also this year, we used Blue Waters to perform CyberShake Study 17.3. This study applied the CyberShake [8] PSHA computational method to Central California for the first time. Study 17.3 calculated two seismic hazard models for Central California: one using a traditional 1D seismic velocity model and the other using a more accurate 3D velocity model, with results shown in Fig. 1. Results using the 3D velocity model show ground motion levels in the California Central Valley that are markedly lower than the levels produced from the standard ground motion




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

prediction equations (GMPEs) currently in widespread use by earthquake engineers. This is in marked contrast to our results in Southern California, where CyberShake predicts stronger shaking than the GMPEs in the deep, low-velocity sedimentary basins. The differences are related to the lateral extents of the basins, which govern their resonance frequencies and amplitudes. These results provide new evidence that CyberShake's physics-based approach can substantially improve our estimates of strong ground shaking.

We are preparing selected research codes to run efficiently on next-generation supercomputers. We have improved the performance of our wave propagation and Strain Green Tensor codes on next-generation GPUs and Xeon Phi systems [9]. To scale up the I/O performance of our software along with our improved compute performance, we optimized I/O performance of our anelastic wave propagation (AWP) software by increasing our use of third-party HPC (high-performance computing) I/O libraries including ADIOS, HDF5, and PnetCDF.

RESULTS & IMPACT

CyberShake simulations for Southern California are under review as inputs to a new Los Angeles urban seismic hazard map are under development by the USGS. The SCEC committee for Utilization of Ground Motion Simulations (UGMS) is working within the framework of the Building Seismic Safety Council activities to develop long-period, simulation-based, spectral-response acceleration maps for the Los Angeles region. Our CyberShake hazard maps are under consideration for inclusion in the National Earthquake Hazards Reduction Program, the American Society of Civil Engineers 7–10 Seismic Provisions, and for the Los Angeles City building codes. The UGMS group is using CyberShake simulations to quantify the effects of sedimentary basins and other 3D crustal structures on seismic hazard—information that is difficult to obtain with traditional empirical methods. Prototype risk-targeted maximum considered earthquake (MCER) response spectra have been mapped using a combination of the empirical approach and the CyberShake model and are being integrated into the National Institute of Building Sciences' Project 17 recommendations for tall buildings. It is no overstatement to say that our sustained work on Blue Waters is transforming and modernizing earthquake science and engineering, and thus represents a major contribution to Strategic Goal 1 (Transform the Frontiers of Science and Engineering) of the NSF 2014–2018 Strategic Plan [10].

WHY BLUE WATERS

SCEC's earthquake system science research program needs access to Blue Waters' scale computing resources for several reasons. SCEC computational research requirements continue to expand in many ways including in terms of algorithmic sophistication, geographical range, and time resolution. New simulations require more computational, memory, and storage resources. Our computational demands continue to grow because our calculations do not yet span the full range of resolution

parameter space, not all important physics have yet been included, and because individual earthquake simulations do not “solve” a problem when run just once or twice. Great uncertainty remains in the ground motions expected in future earthquakes, and society would be remiss in delaying a better resolution of such a critical scientific and public safety challenge.

PUBLICATIONS AND DATA SETS

Breuer, A., A. Heinecke, and Y. Cui, EDGE: Extreme Scale Fused Seismic Simulations with the Discontinuous Galerkin Method. *ISC High Performance'17*, Frankfurt, Germany, June 18–22, 2017.

Breuer, A., A. Heinecke, and Y. Cui, Recent Advances of the ADER-DG Finite Element Method for Seismic Simulations. *Intel XPUIG Workshop*, Chicago, Ill., September 19–22, 2016.

Cui, Y., and SCEC Collaboration, Regional scale earthquake simulations on OLCF Titan and NCSA Blue Waters. *International Workshop on Perspectives of GPU Computing in Science*, Sapienza Università di Roma, Italy, September 26–28, 2016.

Roten, D., et al., High-frequency nonlinear earthquake simulations on petascale heterogeneous supercomputers. *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis* (IEEE Press, Piscataway, N.J., November 13–18, 2016), pp. 957–968.

Roten, D., K.B. Olsen, S.M. Day, and Y. Cui, Quantification of Fault-Zone Plasticity Effects with Spontaneous Rupture Simulations. *Pure Appl. Geophys.*, (2017), DOI: 10.1007/s00024-017-1466-5.

Taborda, R., N. Khoshnevis, S. Azizzadeh-Roodpish, and M. Huda, Influence of the source, seismic velocity, and attenuation models on the validation of ground motion simulations. Poster presentation at *World Conference on Earthquake Engineering* (Santiago, Chile, January 9–13, 2017), number 4574.

Giraldo, D., and D. Restrepo, The spectral cell method in nonlinear earthquake modeling. *Computational Mechanics*, accepted for publication (2017).

Nie, S., Y. Wang, K.B. Olsen, and S.M. Day, 4th-order Staggered-grid Finite Difference Seismic Wavefield Estimation using a Discontinuous Mesh Interface (WEDMI). *Bull. Seis. Soc. Am.*, in press (2017).

Roten, D., K.B. Olsen, and S.M. Day, Off-fault Deformation and Shallow Slip Deficit from Dynamic Rupture Simulations with Fault Zone Plasticity. *Geophys. Res. Lett.*, in press (2017).

Roten, D., K.B. Olsen, and S.M. Day, and Y. Cui, Quantification of Fault-Zone Plasticity Effects with Spontaneous Rupture Simulations. *Pure Appl. Geophys.*, (2017), DOI: 10.1007/s00024-017-1466-5

Small, P., et al., The SCEC Unified Community Velocity Model Software Framework. *Seismological Research Letters*, in press (2017).

Withers, K.B., K.B. Olsen, and S.M. Day, Memory Efficient Simulation of Frequency Dependent Q. *Bull. Seis. Soc. Am.*, 105, pp. 3129–3142.

ENHANCED DIGITAL ELEVATION MODEL FOR THE ARCTIC

Allocation: NSF PRAC/9,696,484 Knh

PI: Paul Morin¹

Collaborators: Charles Nguyen¹, Claire Porter¹, Ian Howat², MJ Noh², Michael Willis³, Brian Bates⁴, Cathleen Williamson⁴

¹University of Minnesota

²The Ohio State University

³University of Colorado

⁴National Geospatial-Intelligence Agency

EXECUTIVE SUMMARY

The Earth’s land surface topography is arguably the most fundamental single data set in the geosciences, geographical sciences, and civil engineering. It is essential to research ranging from the location of rivers and the extent of watersheds in hydrology, to permafrost collapse in built-up areas, to the change in the shape of volcanoes in volcanology. The Polar Geospatial Center (PGC) and its partners at The Ohio State University and Cornell University will adapt PGC’s digital elevation model (DEM) production capabilities from small-area, on-demand production to systematically process and mosaic the entire Arctic sub-meter

stereo commercial imagery archive. Such a DEM would not only catapult the Arctic from the worst- to among the best-mapped regions on Earth, it would also allow for precise detection of change over time, creating a benchmark for measuring this rapidly evolving landscape.

RESEARCH CHALLENGE

There is a lack of high-resolution, consistent, high-quality elevation data available for the Arctic. In 2000, the Shuttle Radar Topography Mission began by acquiring Synthetic Aperture data for the Earth that was processed into an elevation model with a 30-m posting. This mission had limited geographic coverage between the latitudes of 60°N and 56°S because of the shuttle’s limited orbital inclination. The National Geospatial-Intelligence Agency (NGA), Digital Globe (DG), and PGC have built up a near-seamless archive of polar sub-meter stereo imagery that consists of almost 70,000 stereo pair strips from the Worldview 1, 2, and 3 satellites. Using photogrammetric algorithms, we are able to construct digital elevation models (DEM) from the stereo pairs, allowing for mapping of surface features at the same scale as Airborne LiDAR (Light Detection and Ranging) without the cost or logistics constraints of LiDAR. This imagery collection increases at a rate of more than one California-area equivalent per day while there is sufficient Arctic sun. These data will be used by the Arctic research community to support activities that include transportation, national defense, land management, sustainable development, and scientific studies. Further, repeat DEMs with frequencies of months or even days, can be used for change detection, with applications ranging from studies of land use, to resource management, to environmental change.

METHODS & CODES

Our team has spent three years developing an efficient algorithm for constructing photogrammetric DEMs from satellite imagery with the objective of creating a fully automated system capable of handling large amounts of data. Development of the Surface Extraction from TIN-based Search-space Minimization (SETSM) algorithm was begun to facilitate an automated processing pipeline for the PGC operations. SETSM DEMs have been extensively validated [1], are node parallelized using OpenMP, and have been applied to processing large-area DEM mosaics in proof-of-concept studies. Uniquely, SETSM’s structure eliminates the need




Figure 1: The Vavilov Ice Cap on the island of Severnaya Zemlya in Arctic Russia was known to be stable. In 1996, the ice cap was moving at 20 meters per year. By 2016, the interior of the outlet glacier from Vavilov was moving at 25 meters per day. This image pair captures the collapse of the Vavilov Ice Cap as no other resource could, given the temporal frequency and spatial resolution.

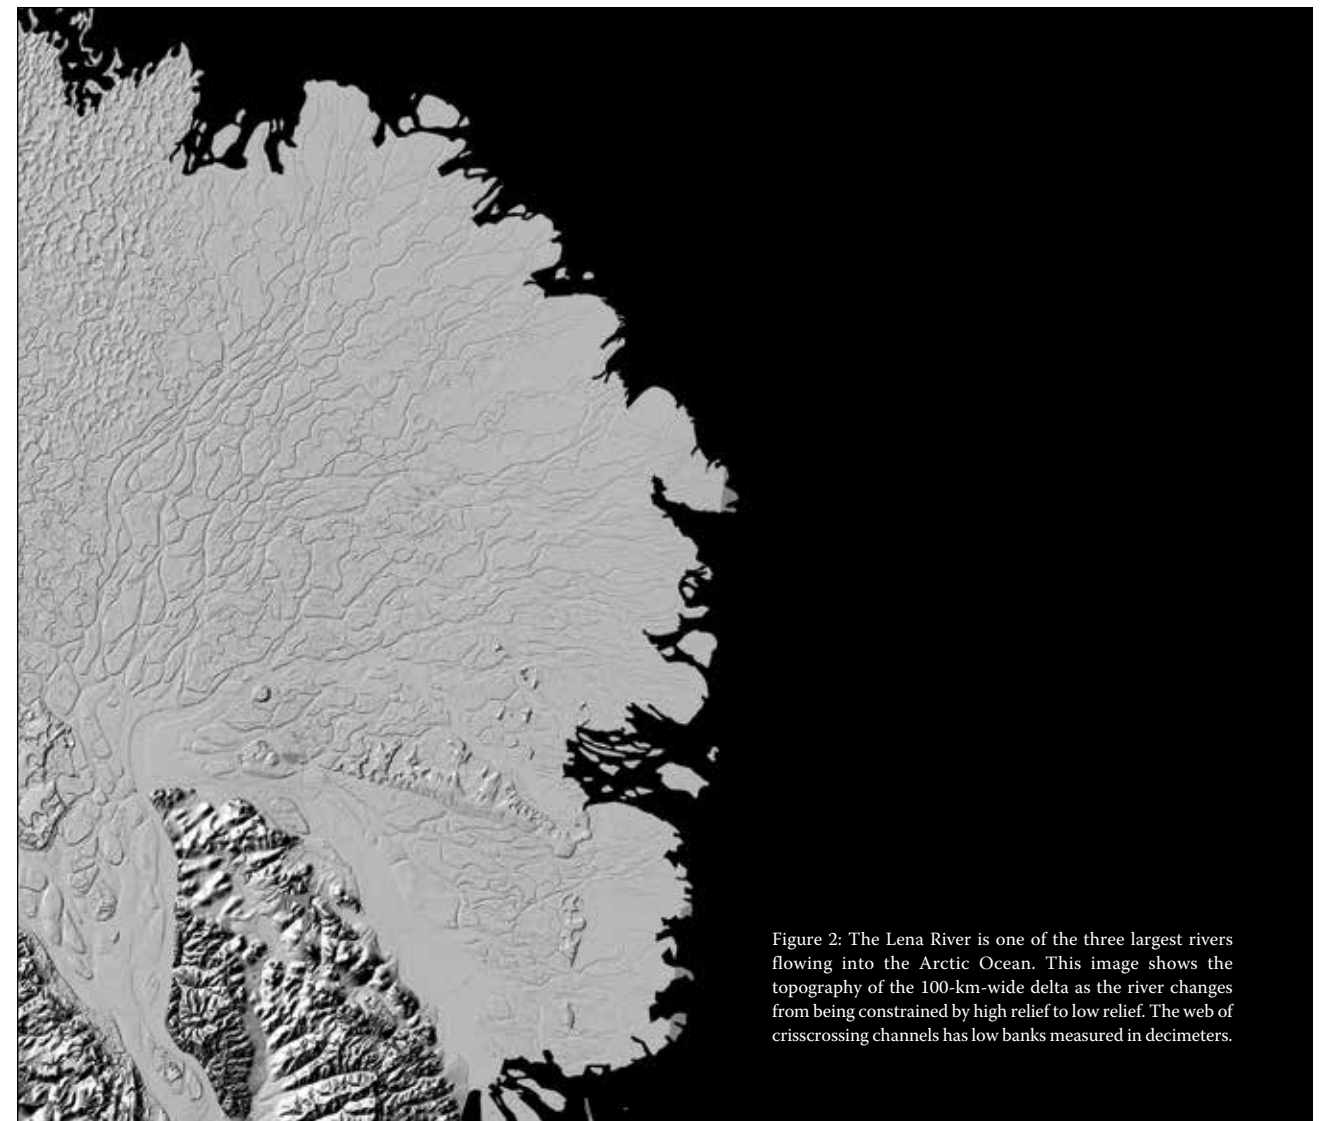


Figure 2: The Lena River is one of the three largest rivers flowing into the Arctic Ocean. This image shows the topography of the 100-km-wide delta as the river changes from being constrained by high relief to low relief. The web of crisscrossing channels has low banks measured in decimeters.

for an existing (i.e., “seed”) DEM for *a priori* constraint or any data-specific, user-defined search parameters, making it a truly automated algorithm. SETSM is called from a single command line with the only required inputs being the filenames of the two stereo images and the RPC (remote procedure call) file, typically provided in XML format.

The DEM extraction workflow starts with a preprocessing step that corrects the source imagery for sensor-specific detector alignment artifacts and outputs a GeoTIFF-formatted set of source rasters. Once the source imagery is corrected, SETSM takes the two source images and derives increasingly detailed elevation models using its pyramid-based approach.

RESULTS & IMPACT

Thus far, we have produced over 57,000 individual 2-m posting DEMs that total approximately 80,000,000 km² of the Arctic. This means that the Arctic is covered four times, on average, though

some areas are more poorly covered and some have over 100 time steps. These data were also processed into continuous mosaics for over 92% of the 20,000,000 km² Arctic. All of these data have been released to the science community and the public through ArcticDEM.org, and Esri has developed viewer and Amazon Web services. These data are now being used by scientists, national geographic surveys, and regional and local governments for a broad range of scientific, civil engineering, and mapping applications.

WHY BLUE WATERS

No other academic computer had the capacity, at the time, for this project. Blue Waters was able to execute the ArcticDEM workload without significantly impacting throughput of other projects. The project had a timeline that precluded the use of cloud services because of time required for development and porting of the SETSM code.

SIMULATING THE MOST DEVASTATING TORNADOES EMBEDDED WITHIN SUPERCELL THUNDERSTORMS

Allocation: NSF PRAC/3,630 Knh

PI: Leigh Orf¹

Co-PI: Catherine Finley²

Collaborator: Robert Wilhelmson³

¹University of Wisconsin-Madison

²St. Louis University

³University of Illinois

EXECUTIVE SUMMARY

Supercell thunderstorms produce the strongest, longest-lived tornadoes, ranked EF4 or EF5 on the Enhanced Fujita Scale. While such tornadoes are uncommon, they cause the vast majority of fatalities and damage. Forecasting such tornadoes requires a deeper understanding of supercell thunderstorms. We have simulated several EF5 tornadoes embedded within supercell thunderstorms at resolutions up to 15 meters, revealing flow features that greatly enhance our current understanding of tornado formation and maintenance. Interactions between the storm's cold pool and updraft create and reorganize vorticity (spin and shear) in such a way as to concentrate it into streams and sheets that help initiate and maintain the EF5 tornado.

RESEARCH CHALLENGE

Tornadoes embedded within supercell thunderstorms create the strongest winds found in nature at the earth's surface. Accurate forecasting of tornado behavior is of great benefit to millions of people who live in tornado-prone regions, such the Great Plains of the United States. However, the process of tornado formation, maintenance, and decay remains elusive to scientists. Our work aims to explore the factors involved in the creation within supercell thunderstorms of the strongest, most long-lived tornadoes.

METHODS & CODES

We used the CM1 model, developed at the National Center for Atmospheric Research. We modified the output driver of the model in order to enable efficient, memory-buffered, high-frequency, lossily compressed floating point HDF5 output, and wrote middleware to read, convert, and visualize raw model output utilizing techniques that included volume rendering and trajectory clouds. We explored different environmental conditions, and the most promising environments were simulated at extremely high resolution, resulting in EF5 tornadoes.

Model data are not currently published online; however, several scientific talks in video format, including model visualizations, may be found at <http://orf.media>. The code written thus far is on a private github server; we intend to make the code available to the public toward the end of this allocation.

RESULTS & IMPACT

Every year Americans die from tornadoes, despite attempts to warn the public of severe weather before it occurs. Our ability to forecast tornado behavior is limited by our current knowledge of how tornadoes form, are maintained, and decay in supercell thunderstorms. The focus of our research is on the most devastating, long-lived tornadoes embedded within supercell thunderstorms. We have simulated dozens of supercells at ultra-high resolution, and some of these supercells form long-lived EF5-strength tornadoes. These simulations have revealed newly identified flow features such as the streamwise vorticity current (SVC), a primarily horizontally oriented, helically flowing "tube" of rain-cooled air that is tilted vertically into the supercell's updraft. The SVC helps maintain the updraft's vigor near the ground, which appears to force tornado formation and assist in its maintenance.

WHY BLUE WATERS

In order to capture flow features that are critical to tornado formation, maintenance, and decay, the atmosphere must be resolved at extremely high resolution. As the resolution of the model is doubled, it requires approximately 16 times more computing resources. Our highest-resolution simulations (at 15-meter grid spacing) utilized 360,000 Blue Waters cores, or roughly half of the machine. Further, extremely high memory, communication, and I/O bandwidth were essential in these simulations, something which is not currently sufficient with XSEDE resources, much less cloud resources.

PUBLICATIONS AND DATA SETS

Orf, L., R. Wilhelmson, and L. Wicker, Visualization of a simulated long-track EF5 tornado embedded within a supercell thunderstorm. *Parallel Comput.*, 55 (2016), pp. 28–34.

Orf, L., et al., Evolution of a Long-Track Violent Tornado within a Simulated Supercell. *Bull. Am. Meteorol. Soc.*, 98 (2017), pp. 45–68.




Figure 1: A short-lived anticyclonic tornado (left) adjacent to the long-lived cyclonic EF5 tornado (right), as seen in the volume-rendered cloud field of a 15-meter resolution simulation. The cloud field also reveals a wall cloud and tail cloud, features commonly observed in the field. Rain is visible as a dark grey field and is most prominent in the rear flank of the supercell (left half of image). The cold pool is represented by the surface buoyancy field where the coldest air is dark blue.

HIGH-RESOLUTION EARTH SYSTEM MODELING USING BLUE WATERS' CAPABILITIES

Allocation: NSF PRAC/8,500 Knh
PI: Robert Rauber¹, Don Wuebbles¹
Co-PI: Gerald Meehl²
Collaborators: Susan Bates², Zach Zobel¹, Justin Small², Christine Shields²

¹University of Illinois at Urbana-Champaign
²National Center for Atmospheric Research

EXECUTIVE SUMMARY

This work contributes to a larger body of ongoing research aimed at using high-resolution climate and weather forecast models to better understand high-impact events in the present day and future warmer scenarios. Simulations completed this year include those using the Community Earth System Model (CESM) at the highest resolution currently feasible for long climate scenarios (0.25° atm/land – 1° ocn/ice) as well as the highest resolution currently possible (0.25° atm/land – 0.1° ocn/ice). For both resolutions, the simulations contribute to a set of control, climate sensitivity, twentieth-century transient, and future scenarios. In addition to ongoing studies of tropical cyclones and midlatitude storms, we also investigated atmospheric rivers and regional extreme temperatures in the northwest Atlantic Ocean. Using low-resolution global climate model (GCM) output to force the higher resolution Weather and Research Forecasting (WRF) model (~12 km) allows for dynamical downscaling over the

contiguous United States to investigate extreme temperature and precipitation.

RESEARCH CHALLENGE

Our current PRAC project employs global climate models at the highest resolution feasible for simulating multiple centuries of Earth's climate and a high-resolution regional model for informing changes in extremes over the United States, providing a leading research environment for the immediate needs in climate science, informing the upcoming national and international climate assessments, and building the avenue to scientific discovery in the next five to ten years. These analyses will be a significant contribution to the coordinated international special computational studies and model intercomparison to be done to analyze past and projected future changes in the earth's climate system. In addition, the results of these studies will be fully available to the scientific community for further analyses and resulting insights into the processes, mechanisms, and consequences of climate variability and climate change. Such high-resolution modeling studies are likely to produce important findings for other scientists, social scientists, and policymakers to achieve further understanding of climate change science, the resulting societal and ecosystem impacts, and insights for adaptation and mitigation analyses. The nature of simulating climate processes at high resolution makes this work a grand challenge. As we move to higher and higher resolution, we must rethink the need for model parameterizations for various physical processes, and from a technical standpoint we must scale the model to larger node counts.

METHODS & CODES

For long climate integrations we use the Community Earth System Model (CESM), a coupled climate model for simulating the earth's climate system. Composed of six component models that simulate Earth's atmosphere, ocean, land surface, sea ice, land ice, river transport, and one central coupler component, the CESM allows researchers to conduct fundamental research into the earth's past, present, and future climate states.

For regional downscaling over the contiguous U.S., we use output from three GCMs—Community Climate System Model (CCSM4); Geophysical Fluid Dynamics Laboratory Earth System

Model (GFDL-ESM2G); and Hadley Centre Global Environment Model, version 2-Earth System (HadGEM2-ES)—to force the Weather and Research Forecasting (WRF) version 3.3.1.

The TempestExtremes [1] and TSTORMS [2] packages were used to track midlatitude storms and tropical cyclones, respectively.

RESULTS & IMPACT

Building on our research from the previous year, we expanded storm tracking into the midlatitudes to investigate potential changes in storm number and strength over the current century. As found with tropical cyclones, the total number of midlatitude storms decreases in the future in all ocean basins (Fig. 1). The atmosphere tends to be more stable in the future in the midlatitudes, contributing to fewer storms. Atmospheric rivers (ARs), a subcategory of midlatitude storms that are able to transfer moisture from the tropics to the midlatitudes, had not previously been analyzed in the 0.25° atm/land-1° ocn/ice version of the CESM. Higher model resolution allows for better representation of counts and frequency of ARs. Results indicate that future changes in ARs depend on changes in the atmospheric jets and therefore respond differently in the different ocean basins.

Certain regional processes within the ocean may be much better represented by the high-resolution ocean (0.1°) simulations. An observed hotspot of localized warming in the Gulf of Maine is much better represented in this high-resolution simulation compared to the nominal 1° simulations. Such extremes in ocean temperatures can have a significant economic impact on fisheries.

Investigations of extreme temperature and precipitation change over the contiguous United States reveal startling realities for future climate. Two future climate scenarios were simulated: a moderate warming scenario (RCP4.5) and a strong warming scenario (RCP8.5). Fig. 2 shows the change in the number of days exceeding 95°F for the late twenty-first century under each forcing scenario. Some regions, such as the Northeast, will not experience much change due to the relative rarity of events in that region, while other areas, such as the middle part of the country, are projected to experience large changes in both scenarios. For much of country, there is an additional one to two months each year that will exceed 95°F with the RCP8.5 scenario as compared to RCP4.5. This is significant because in historical simulations, the vast majority of 95°F days take place only during June–July–August (JJA) for these regions. An additional 30–50 days per year means the “summer” months will span more than just JJA, which would have significant energy and agriculture impacts.

WHY BLUE WATERS

Multiple century-long simulations are needed in order to quantify CESM model characteristics and sensitivity and to produce a sufficiently long, stable preindustrial control simulation, followed by historical and numerous future scenarios. Furthermore, multi-member ensembles are needed to quantify and reduce uncertainty. The climate modeling community has refined horizontal resolution to 0.25° for the atmosphere and 0.1° for the

ocean, allowing for a full eddy-resolving ocean simulation within the modeling system. These simulations and analyses, at high resolution, at a minimum require petascale computing resources and cannot be completed without a computational platform like Blue Waters. Because these simulations use a modest number of nodes for long periods of time, our project requires the help of the Blue Waters staff to achieve good throughput.

PUBLICATIONS AND DATA SETS

Bacmeister, J. T., et al., Projected changes in tropical cyclone activity under future warming scenarios using a high-resolution climate model. *Climatic Change* (2016), DOI: 10.1007/s10584-016-1750-x.

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 (2016), pp. 3603–3608.

Zarzycki, C. M., et al., Impact of ocean coupling strategy on extremes in high-resolution atmospheric simulations. *Geosciences Model Development*, 9 (2016), pp. 779–788, DOI: 10.5194/gmd-9-779-2016.

Zobel, Z., J. Wang, D. J. Wuebbles, and V. R. Kotamarthi, Evaluations of high-resolution dynamically downscaled ensembles over the contiguous United States. *Climate Dynamics* (2017), pp. 1–22, DOI: 10.1007/s00382-017-3645-6.

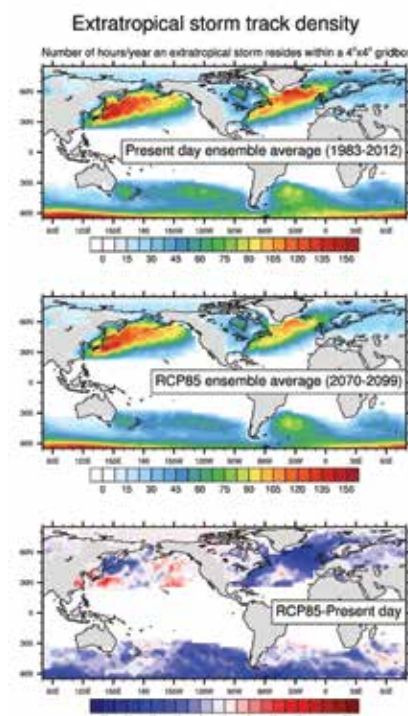


Figure 1: The total number of hours per year in which a storm is located within a 4°x4° gridbox calculated in the present day simulations (top panel), the future scenario (middle panel) and the difference between the two (future minus present day, bottom panel).

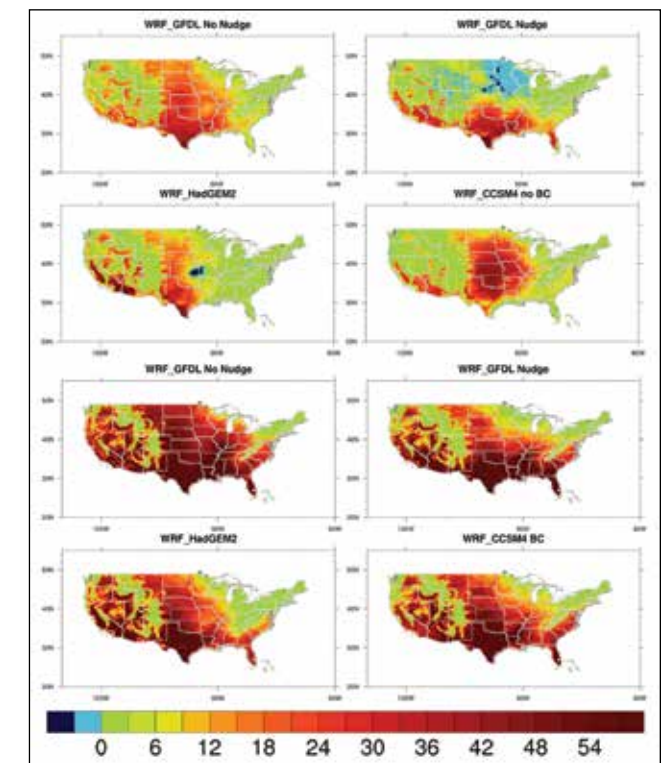


Figure 2: The change in the number of days that exceed 95°F between a moderate warming scenario and present day (top four panels) and a high warming scenario and present day (bottom four panels). The four panels show results from four different configurations of WRF using different boundary forcing (from either a global climate model or observations).

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

Allocation: NSF PRAC/5,000 Knh
PI: Jamesina J. Simpson¹

¹University of Utah

EXECUTIVE SUMMARY

The largest documented geomagnetic storm due to a coronal mass ejection on the sun occurred in 1859. This storm caused telegraph operators communicating over 100-km-long wire lines to experience electric shocks, some nearly fatal. The historical record suggests that extreme space weather is likely to impact the Earth again in the future. However, modern electrotechnologies will be affected by space weather to a much larger degree than in the past. We are utilizing Maxwell's equations models of the Earth-ionosphere waveguide to calculate location-specific space weather hazards to electric power grids in order to prevent blackouts. Blue Waters permitted us, for the first time, to study electric field behavior near ocean-continent boundaries using realistic coastal geometries. Additionally, we developed stochastic models of electromagnetic wave propagation through uncertain and variable ionosphere regions. These algorithms are allowing us to determine the confidence level that a communication or similar system will work as expected during disturbed conditions.

RESEARCH CHALLENGE

The largest documented geomagnetic storm on Earth resulting from a coronal mass ejection (CME) on the sun occurred in 1859. That storm caused telegraph operators communicating over 100-km-long wire lines to experience electric shocks, some nearly fatal. Further, business transactions requiring telegraphic exchanges were completely shut down in the world's major capitals. A 2008 National Academies report indicates that extreme space weather events, "though rare, are likely to occur again sometime in the future." However, a reoccurrence of an 1859-magnitude space weather storm could disrupt today's society to a much greater degree than in 1859 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 electric power grids are just some examples.

The focus of our research was to greatly enhance our understanding of the near-Earth electrodynamics associated with historically intense CMEs. The methodology of the proposed work was to advance and apply detailed, high-resolution Maxwell's equations models of the Earth-ionosphere waveguide developed by the PI over the past ~15 years [1]. These models are based on the robust finite-difference time-domain (FDTD) method

[2]. They currently extend from -400 km below the Earth's surface to an altitude of +100 km. They uniquely account for the Earth's complete topography, oceans, lithosphere composition, geomagnetic field, and magnetized ionospheric plasma according to altitude, position, and time of day, all while solving the full-vector Maxwell's equations. As a result, these models created by the PI's group are the most advanced electrodynamic models of the Earth-ionosphere waveguide. They can model the entire world's response to a space weather event and provide location-specific information on possible hazards to societal infrastructure. Previous studies were limited to only the sinusoidal steady-state and involved simplifying geometries and physics (e.g., infinite line currents, layered lithosphere models, nonphysical constant magnetic fields assumed within each layer, only solutions to Poisson's equation, etc.)

METHODS & CODES

The first goal of our research was to help power grid stations better understand their individual risks to different space weather impact scenarios depending on their orientation and location. Of particular focus this past year was the electromagnetic field behavior in coastal regions in order to determine whether space weather poses unique hazards to power grids along coastlines.

For this work, we used the FDTD method to calculate electromagnetic fields at the surface of the Earth near ocean-continent boundaries. As a time-domain method, FDTD permits modeling of arbitrary source time-waveforms, variable current source orientations as shown in Fig. 1, and even the finite propagation velocity of the ionospheric currents. Further, as a grid-based method, FDTD permits modeling of complex geometries, such as sloping coastlines combined with finite depth oceans (rather than a coastline having a constant, infinitely-long slope as in the previous analytical studies).

The second goal of our research was to develop the first efficient, grid-based stochastic electrodynamic models of the Earth-ionosphere waveguide. Nearly all electrodynamic solvers assume average (mean) electrical properties of materials and solve for average (mean) electric and magnetic fields. However, assuming numerically only an average state of the ionosphere yields calculated output electromagnetic field waveforms that are not as rich and complex as measured electromagnetic fields. Further, there is great uncertainty in the content of the ionosphere at any given moment. The FDTD models we developed this past year

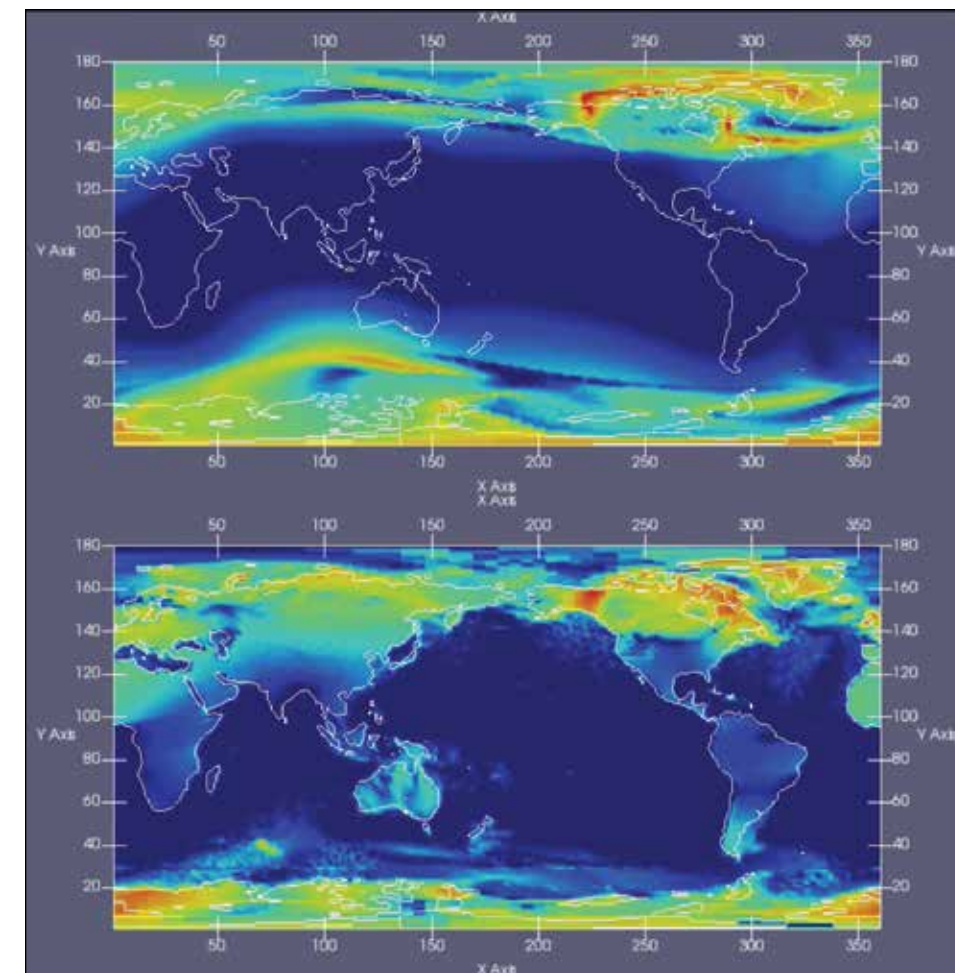


Figure 1: Snapshot of the electric field source amplitude versus position during the October 2003 Halloween geomagnetic storms as calculated by BATS-R-US model developed at the University of Michigan (top image) and the resulting surface-level electric field values calculated by the global FDTD model (bottom image).

solve for both the mean and variance of electromagnetic wave propagation through a varying/uncertain ionosphere.

RESULTS & IMPACT

In 2015, Prof. Simpson was among 20 researchers who participated in the NASA Living with a Star Working Group Institute on Geomagnetically Induced Currents. The issue of ocean-continent boundaries and whether they pose significant risks to power grids was among the list of topics that were considered important and unresolved. Blue Waters allowed us to take a closer look at this issue this past year using an established numerical technique (FDTD) that offers more flexibility than previously possible using analytical approaches, and provides more rigorous (full-vector Maxwell's equations) solutions. In order for space weather to pose a unique risk to electric power grids at ocean-continent boundaries, high electric fields must extend at least 100 km inland from the coast. Counter to the conclusions of previously published work, we were able to determine that space weather does not induce intense electric fields over a sufficiently large area to pose a risk to power grids [3].

WHY BLUE WATERS

We were able to use Blue Waters to develop higher grid resolutions than previously possible by over an order of magnitude (~1 x 1 x 1 km vs. ~40 x 40 x 5 km). This has opened up a wide variety of new applications because we can model higher frequencies of electromagnetic waves and also model smaller geometries. For example, we are now working with the Defense Advanced Research Projects Agency to examine the use of our models for developing a new electromagnetic system for geolocation. We are also now working with the Office of Naval Research to detect objects submerged in the ocean. Other possibilities due to the higher grid resolutions achieved include communications during space weather events, new remote sensing applications, and studying very low-frequency signals detected by spacecraft around the times of earthquakes.

PUBLICATIONS AND DATA SETS

S. Pokhrel, B. Nguyen, J. J. Simpson, FDTD applied to GICs at ocean-continent boundaries, *AGU Fall Meeting*, San Francisco, Calif., December 12–16, 2016.

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

Allocation: Illinois/500 Knh
PI: Ryan L. Sriver¹
Co-PI: Hui Li¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

This research seeks to advance our understanding of the relationship between tropical cyclones (TCs) and Earth’s climate system using a high-resolution state-of-the-art Earth system model, Community Earth System Model, or CESM. In this report, we highlight results from a set of 30-year simulations in which the high-resolution (25 km) atmosphere component is configured with three different levels of ocean coupling: prescribed sea surface temperature (SST), mixed-layer ocean, and full three-dimensional dynamic ocean with nominal 1-degree horizontal resolution. We find that the inclusion of ocean coupling significantly affects simulated TC characteristics, including storm frequency, geographic distribution, maximum wind, and storm intensification. Key differences in storm numbers and distributions can be attributed to variations in the modeled large-scale climate mean state and variability that arise from the combined effect of intrinsic model biases and air–sea interactions. This work addresses the importance of storm-induced ocean–atmosphere feedbacks in Earth’s coupled climate system, which can help improve our understanding of how TC activity may change in the future.

RESEARCH CHALLENGE

TCs are among the world’s deadliest natural hazards. How TC activity will vary with the changing climate is a topic of great interest. Recent research has shown that Atmosphere General Circulation Models (AGCMs) are capable of generating TC-like vortices, and the representation generally improves with increasing horizontal resolution [1]. The current generation of high-resolution (less than 50 km) AGCMs have been shown to capture realistic global TC activity and can resolve the most intense Category 5 TCs [2-4].

Correct representation of air–sea interactions under TCs is important for simulating realistic storm intensities and track durations. Local feedbacks due to ocean mixing and surface fluxes can inhibit storm develop and intensification as well as influence larger-scale ocean and atmospheric circulations. These feedbacks can influence TC characteristics and statistics in a global coupled Earth system model, and they can also contribute to the tropical mean state and variability on interseasonal to interdecadal timescales. AGCMs with prescribed surface ocean conditions do not capture these feedbacks and are thus missing important physical processes that can influence climate variability.

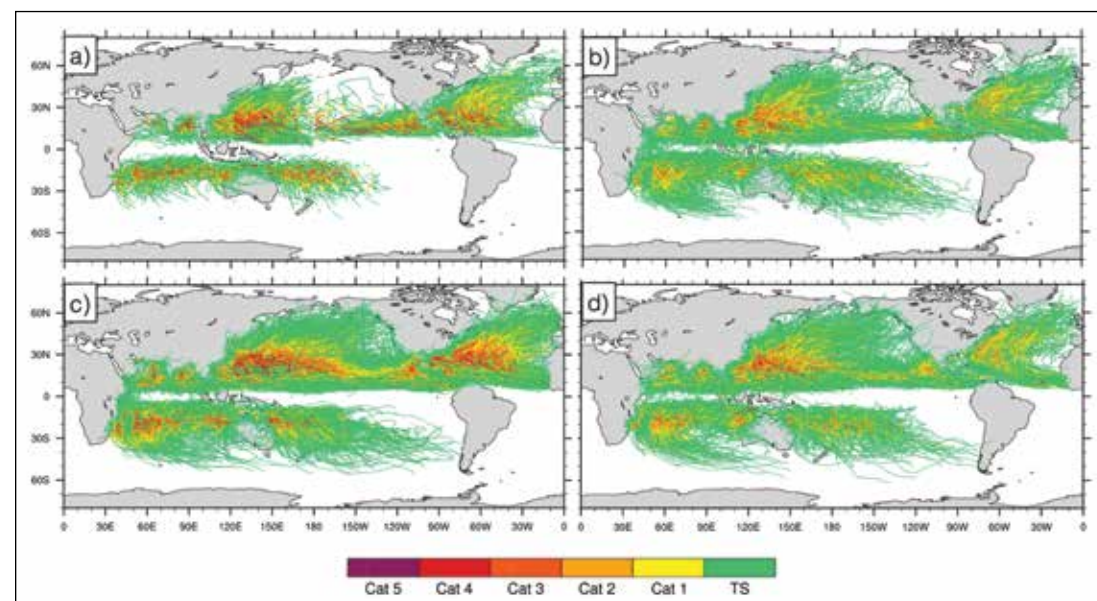


Figure 1: Global TC tracks accumulated over 30 years in the (a) observational best-track (1985–2014), (b) fully-coupled simulation, (c) atmosphere-only simulation, and (d) partially-coupled simulation.

METHODS & CODES

In this project, we assess the impact of ocean coupling on simulated TC activity using a high-resolution configuration of the CESM with a 25-km resolution atmosphere [5]. We performed three, 30-year simulations in which the atmosphere model is configured with three different levels of ocean coupling: (1) prescribed monthly-varying sea surface temperature based on the observed climatology; (2) a mixed-layer ocean model that allows thermodynamic exchanges between the atmosphere and the ocean mixed layer but does not account for ocean dynamics; and (3) a 1-degree ocean general circulation model with full dynamics and ocean–atmosphere fluxes. The models are configured with increased air–sea coupling frequency, modified surface wind drag law, and high frequency of history file output, in order to focus on ocean–atmosphere interactions associated with TCs. Each simulation is run for 30 years under the preindustrial climate conditions and with an active carbon–nitrogen cycle.

RESULTS & IMPACT

We find that TC number, geographical distributions, and intensity are sensitive to ocean coupling (Figs. 1 and 2). Differences in TC characteristics are mainly attributed to model differences in local air–sea flux exchanges and large-scale climate conditions. The fully coupled model with dynamic ocean simulates the most realistic annual TC number in the northwestern Pacific, though an expanded Indo-Pacific warm pool is likely contributing to more TCs forming in the central Pacific and Indian Ocean. The coupled model underestimates TC activity in the north Atlantic and northeastern Pacific, which is mainly due to biases in sea surface temperature and vertical wind shear, as well as the southward shift of the Inter-Tropical Convergence Zone. The partially coupled simulation with a mixed-layer ocean exhibits key features of TC distributions similar to the fully coupled simulation, including the cold surface ocean temperature bias and the underactive TCs in the north Atlantic. By design, ocean heat transport in the mixed-layer ocean model is specified based on the modeled estimates from the fully coupled model, thus biases in ocean heat transport in the coupled model are likely propagated to the partially coupled simulation. The uncoupled atmosphere-only simulation appears to perform better in the north Atlantic than the coupled model. However, the model exhibits an asymmetric bias in TC activity across the Pacific Ocean, including an overestimation of TC number in the eastern Pacific and underestimation of TC number in the northwestern Pacific. Our results show that this asymmetric pattern can be related to a weakened tropical zonal atmospheric overturning circulation (Walker Circulation) and associated biases in large-scale vertical motion.

This research addresses longstanding scientific questions about the relationship between tropical cyclones and ocean–atmosphere dynamics and variability on multiple spatial and temporal scales. It enables fundamental advancement of our mechanistic understanding about important physical processes related to TC

dynamics, ocean mixing, ocean heat storage and transport, and global ocean–atmosphere circulations. It paves the way for more comprehensive coupled climate model experiments capable of linking extreme weather events with large-scale climate. Future work includes investigating the relationship between TC activity and climate variability on intraseasonal to interannual scales, as well as characterizing model uncertainties that may affect TC projections in the future climate with enhanced CO₂ forcing.

WHY BLUE WATERS

Given the substantial computational expense of high-resolution Earth system models, it is difficult to apply these models to study tropical cyclones because of the necessary grid resolution (¼ degree), integration length (multiple decades), and high frequency output (sub-daily). Blue Waters provides the unique capabilities to handle the computational demand associated with running the model at ultra-high resolutions, including scalability to over 15,000 cores, high-frequency input and output, and post-processing and visualization of model results.

PUBLICATIONS AND DATA SETS

Li, H., and R. L. Sriver, Effects of ocean grid resolution on tropical cyclone-induced upper ocean responses using a global ocean general circulation model. *Journal of Geophysical Research-Oceans*, 121 (2016), DOI: 10.1002/2016JC011951.

Huang, A., et al., Regional variations in the ocean response to tropical cyclones: Ocean mixing versus low cloud suppression. *Geophysical Research Letters*, 44:4 (2017), DOI: 10.1002/2016GL072023.

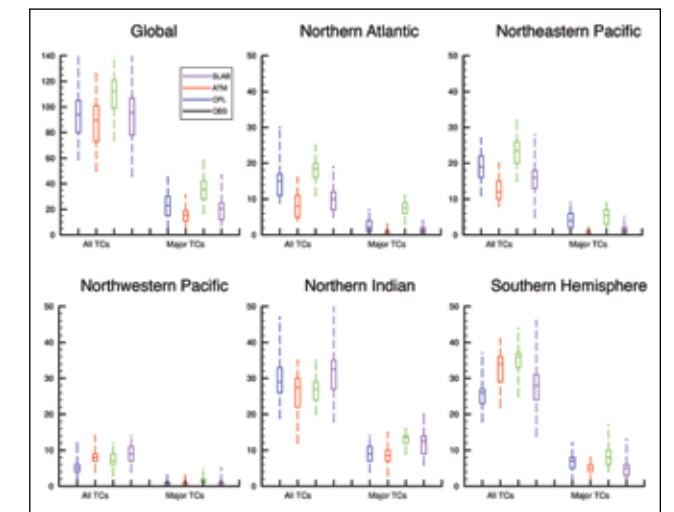


Figure 2: Global and basin-scale annual average TC number in observational best-track (blue), fully coupled simulation (red), atmosphere-only simulation (green), and partially coupled simulation (purple). In each panel, the left side accounts for the total storm number and the right side accounts for strong TCs with intensity higher than Category 3.

PETASCALE MODELING OF CONVECTIVE STORMS UNDER CLIMATE CHANGE AND VARIABILITY

Allocation: Blue Waters Professor/240 Knh
PI: Robert J. Trapp¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

This research seeks to answer the basic question of how present-day extreme storm events might be altered by human-induced climate change. The pseudo-global warming (PGW) methodology has been adapted for this purpose. Modified atmospheric states drawn from global climate model (GCM) output were used to constrain an ensemble of Weather Research and Forecasting (WRF) model simulations of select extreme events. Highly resolved grids allow for process-based analyses of the simulated

events; ensembles of such simulations facilitate quantification of uncertainty.

Our most recent work has focused on hail storms and tornadoes spawned by landfalling hurricanes. The severe hail events considered thus far are more intense but generate less-extensive and fewer hail swaths under PGW. Simulations of Hurricane Ivan (2004) are showing that the conditions under PGW promote a relatively more intense hurricane at landfall, and lead to a much higher incident of mesoscale vortices with tornadic potential.

RESEARCH CHALLENGE

A persistent uncertainty in climate-change assessments is how the frequency and intensity of local, high-impact thunderstorms, and even large thunderstorm systems including hurricanes, might be affected by human-enhanced greenhouse gas concentrations. Part of the challenge is that such storms—and especially the attendant tornadoes, hail, damaging “straight-line” winds, lightning, and localized flooding—have spatial scales that fall below the effective resolution of typical global models. Modeling approaches such as dynamical downscaling have addressed this resolution issue, but their applications thus far have generally been unconcerned with historical events, and therefore about how these events might be projected in the future.

METHODS & CODES

Drawing on the success of our previously reported work, we have further adapted the PGW methodology to investigate the impact of human-induced climate change on outbreaks of severe hail and on landfalling hurricanes. Modified atmospheric states drawn from GCM output were used to constrain WRF model simulations of these events at high resolution (inner-domain grids with lengths of 333 m). Comparison of an ensemble of these simulations with control simulations (CTRL) is facilitating the assessment of PGW effects. Experimentation with two-moment microphysical parameterization schemes adds to the hailstorm simulation ensemble.

RESULTS & IMPACT

Exemplifying the hailstorm results are the simulations of the May 19, 2013, outbreak of tornadoes, damaging wind, and hail (Fig. 1). Relative to the CTRL experiments, the PGW experiments across the three microphysical schemes (Morrison: MO, Milbrandt-Yau: MY, NSSL: NS), and the three GCM drivers (GFDL, MIROC, and NCAR), the individual hailstorms under PGW tended to be




Figure 2: Bar graph showing the domain-wide grid point hail occurrence on May 19, 2013, under PGW. The values represent fractional changes relative to the CTRL simulation. The PGW experiments are across three microphysical schemes (Morrison: MO, Milbrandt-Yau: MY, NSSL: NS) and three GCM drivers (GFDL, MIROC, and NCAR).

more intense but generated less extensive and fewer hail swaths (Fig. 2). Our analyses thus far are showing that this response is due to the PGW-induced: (1) changes in the thermodynamic environment (e.g., higher convective available potential energy, higher convective inhibition, and higher freezing level), (2) increases in precipitation loading of individual updrafts, and (3) reductions in environmental vertical wind shear. One hypothesis that we are currently exploring is that a reduction in wind shear leads to a reduction in updraft area and thus in the area that accommodates hail growth.

The hurricane simulations are of Hurricane Ivan (2004), which spawned a record 118 tornadoes and caused significant damage to inland communities. We are finding that the conditions under PGW promote a relatively more intense hurricane at landfall and lead to a much higher incidence of mesoscale vortices with tornadic potential. The relatively higher convective available potential energy under PGW is one hypothesized physical explanation for this response.

In addition to continued analyses of the hailstorm and hurricane simulations described above, we are currently developing the capability to employ the Model for Prediction Across Scales (MPAS) for studies of convective storms under climate change and variability. MPAS is one of the emerging global atmospheric models with variable-resolution grids and provides us with the

ability to isolate effects of low- and high-latitude processes (e.g., from Arctic sea ice and tropical oceans) on deep convective storms that are well resolved in middle latitudes. With the assistance of National Center for Supercomputing Applications personnel, the MPAS model codes have recently been compiled on Blue Waters, and MPAS experimentation is currently under design.

WHY BLUE WATERS

The relatively small size of thunderstorms and the ranges of relevant scales within tropical cyclones, coupled with their episodic occurrence, necessitates a research approach that can account for temporal scales from decades to minutes and spatial scales of thousands of kilometers to hundreds of meters. In other words, we require very large geospatial domains that have fine grid point spacings and long-time integrations with high rates of model output. Moreover, quantifications of uncertainty require that such realizations be repeated over multiple experiments. The Blue Waters allocation is providing us with the resources needed to achieve this unprecedented level of climate simulation.

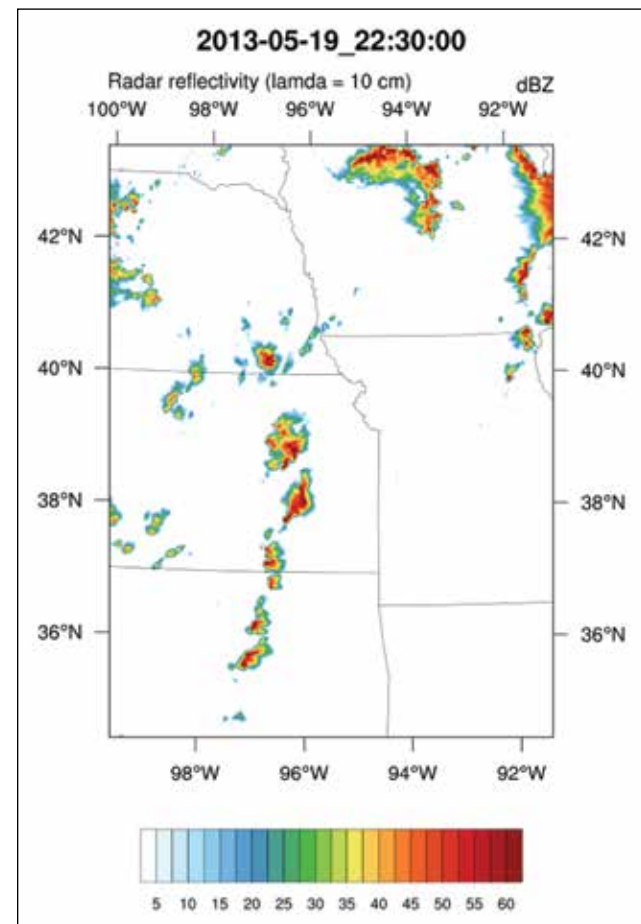


Figure 1: Simulated radar reflectivity from a high-resolution WRF model simulation of the May 19, 2013, outbreak of tornadoes, damaging wind, and hail. The radar reflectivity portrays the structure of the individual hail-producing storms in the CTRL simulation.

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

Allocation: Illinois/550 Knh
PI: Junshik Um¹
Collaborator: Greg M. McFarquhar¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

The single-scattering properties of ice crystals with maximum dimensions smaller than 50 μm were calculated at a wavelength of 0.55 μm using a numerically exact method (i.e., the discrete dipole approximation) and Mie theory. For these calculations, hexagonal columns and spheres were used to represent the shapes of the small ice crystals in natural clouds. Further, because the morphological features of nonspherical ice crystal are closely related to their single-scattering properties, varying aspect ratios were used to characterize the hexagonal column shapes in the scattering calculations. The results show the impacts of orientation, shape, and aspect ratio on the directional intensity of scattered light.

Based on these calculations, potential errors due to shape and orientation on the sizing of particles by current forward scattering probes that measure ice crystals with maximum dimensions smaller than 50 μm are quantified in an inverse problem.

RESEARCH CHALLENGE

Current *in situ* airborne probes (e.g., forward scattering spectrometer probes) that measure the sizes of ice crystals with maximum dimensions D_{max} less than 50 μm are based on the concept that the measured intensity of light scattered by a particle in the forward and sometimes backward direction can be converted to particle size. The retrieval of ice crystal size from satellites also relies on relationships between light scattering and particle size, as do parameterization schemes for numerical models. The relationship between particle size and scattered light used to process data from current forward scattering probes is based on Mie theory, which assumes the refractive index of a particle is known and that all particles are spherical. Not only are small crystals not spherical, there are also a wide variety of nonspherical shapes that have been used to represent them based on images of actual observed ice crystals.

Although it is well known that the scattering properties of nonspherical ice crystals differ from those of spherical shapes, the impacts of this nonsphericity on derived *in situ* particle size distributions (PSDs), upon which satellite retrievals and parameterizations for large-scale numerical models depend, are unknown. To improve *in situ* airborne measurements of small crystals and PSDs, precise relationships among the intensity of light scattered in multiple directions by a particle and its size and shape are required and should be based on accurate calculations of single-scattering properties. Such calculations demand large computing time and memory that rapidly increase with particle size. Large computing resources such as Blue Waters are necessary for these calculations.

METHODS & CODES

The discrete dipole approximation (DDA), is a flexible technique that can calculate the scattering properties of irregularly shaped particles. In DDA, a particle is discretized into a number N of elementary polarizable units called dipoles. Specifying the location and polarizability of these dipoles allows calculations

of the scattering and absorption of light. The number of dipoles N into which a particle is divided can be assigned to N CPUs (or cores) with the single-scattering properties then calculated within a parallel environment (e.g., message passing interface). In our study, a numerical code, the Amsterdam DDA (ADDA [1]) was used to calculate the scattering and absorption of electromagnetic waves by ice crystals.

RESULTS & IMPACT

We calculated the single-scattering properties (i.e., phase matrix, asymmetry parameter, and extinction efficiency) of hexagonal crystals with $D_{max} < 50$ μm at a nonabsorbing wavelength (i.e., $\lambda = 0.55$ μm) using Blue Waters. Because molecules in ice crystals form a hexagonal lattice structure, the most common crystal habits are hexagonal prisms [2]. To represent natural variations of hexagonal ice crystals [3], six different aspect ratios (AR=0.1, 0.25, 0.5, 1.0, 2.0, and 4.0) were used. The single-scattering properties of hexagonal columns with a width of up to 36 μm and a length of up to 48 μm were determined.

Figure 1 shows the calculated nonzero phase matrix of hexagonal crystals with AR=1.0. Based on these calculations, a new conversion table (i.e., differential scattering cross sections) was generated (Fig. 2). The differential scattering cross sections of nonspherical crystals calculated using ADDA are significantly different from those of spherical particles (the brown and black lines in Fig. 2) determined using Mie theory. Errors in the sizing of ice crystals used in current forward scattering probes due to nonsphericity of atmospheric ice crystals were quantified for the first time using the newly developed conversion table. Differences in sizing particles were larger for those with $D_{max} < 10$ μm (due to interference structures) and with $D_{max} > 10$ μm (due to nonsphericity) than for those with $D_{max} \sim 10$ μm. The differences were up to 112% (170%) in the forward (backward) direction depending on the degree of nonsphericity assumed in the orientation-averaged calculations. However, a measurement is made by a forward scattering probe within 1.0 μsec, which implies that a particle has a certain orientation. The differences became larger and were up to 515% (790%) when orientations were considered in scattering calculations.

Most research aircraft have at least one forward scattering probe to measure small cloud particles. Thus, developing a new probe that can measure sizes and shapes of small crystals and also distinguish them from liquid cloud droplets is important. However, improving the processing algorithms of current forward scattering probes that allow use of previous measurements acquired using the current model of forward scattering probes is also important.

WHY BLUE WATERS

Although numerically exact methods are typically used to calculate single-scattering properties of particles with small size parameters, approximations are often used for larger size parameters. Although exact methods can be used for particles with larger size parameters and provide more accurate results, they




Figure 2: Calculated differential scattering cross sections of hexagonal ice crystals and spheres in forward (left) and backward (right) directions.

require more computing time and memory that rapidly increases with particle size, preventing their use in many circumstances. The accuracy of radiative transfer models and satellite retrieval algorithms depends heavily on accurate calculations of single-scattering properties of ice crystals. Blue Waters is an important resource in completing these calculations by allowing more exact calculations for larger particles.

PUBLICATIONS AND DATA SETS

Baumgardner, D., et al., Chapter 9: Cloud Ice Properties - In Situ Measurement Challenges. in *Meteor. Monogr.*, 58 (2017), pp. 9.1–9.23, DOI: 10.1175/AMSMONOGRAPHS-D-16-0011.1.

Um, J., and G. M. McFarquhar, Quantifying uncertainty in forward scattering probes due to non-sphericity of atmospheric ice crystals. *17th International Conference on Clouds and Precipitations*, Manchester, UK, July 25–29, 2016.

Um, J., and G. M. McFarquhar, Accurate calculations of single-scattering properties of small ice crystals: Scattering database and application to in-situ forward scattering probes. *1st International Summer Snowfall Workshop*, Cologne, Germany, June 28–30, 2017.




Figure 1: Non-zero scattering phase matrix elements of hexagonal ice crystals with aspect ratio of 1.0 at wavelength of 0.55 μm as functions of scattering angle and volume equivalent size parameter.

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

Allocation: Illinois/250 Knh
PI: Albert J. Valocchi¹
Co-PI: Yu Chen¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Advances in noninvasive imaging of rock along with continued developments in computing power open up the exciting possibility of direct numerical simulation of pore-scale physics that are relevant to a variety of subsurface energy processes. In particular, we can investigate the fundamental pore-scale flow processes that control migration and trapping of supercritical carbon dioxide during carbon capture and storage, which is a key technology to mitigate emission of greenhouse gasses into the atmosphere. Direct numerical simulation of pore-scale multiphase flow physics is a grand computational challenge since a fine spatial grid is required to capture the complex pore geometry, while a large spatial domain must be included for a statistically representative sample. Although the lattice Boltzmann method (LBM) is generally known to be an effective numerical scheme, we did considerable work to develop a flexible code that is optimized for manycore processors. We used our code in conjunction with unique microfluidics experiments to investigate the role of inertial effects during rapid pore-scale displacements.

RESEARCH CHALLENGE

Understanding the migration of multiple fluids within pore spaces in subsurface geological formations is critical for addressing important problems such as enhanced oil recovery, groundwater

pollution from leaking tanks or pipelines, geothermal energy production, and geological sequestration of carbon dioxide (CO₂). The last application, capture and geological storage of CO₂, provides a means of reducing CO₂ emissions into the atmosphere, by capturing CO₂ from major stationary sources and injecting it into suitable deep rock formations. Fingering and fluid displacement patterns at the pore scale can have a profound impact on large-scale phenomena such as the relative permeability relationship and residual trapping of fluids. Therefore, it is of great importance to study the detailed pore-scale displacement patterns of CO₂ injection in porous media under reservoir conditions.

Laboratory rock core flooding experiments are a major tool for investigating multiphase flow in CO₂ storage and petroleum reservoir applications. However, these experiments are limited since they are difficult, time-consuming, and expensive, and it is challenging to visualize the dynamic displacement process within the rock with sufficient time and space resolution. Pore-scale direct numerical simulation (DNS) in porous media reconstructed from micro-CT (computed tomography) scanned images of real rock has become increasingly popular. In addition to visualizing fluid displacement in porous media with high resolution, DNS also provides full-flow field information and has the flexibility to independently adjust fluid parameters for parametric study. Despite the rapid development of modern processors, the computation




Figure 1: Comparison of CO₂ invasion patterns at approximately steady state between the simulations and experiments at different Ca. Red fluid represents liquid CO₂, blue fluid represents water, and solid grains are in black. E0, E1, E2, E3, and E4 denote the experimental results. S1, S2, and S3 denote the simulation results.

cost for such DNS is still very high. Fine spatial grid resolution is required to represent the pore space, yet a relatively large volume of rock must be modeled in order to have a representative volume element. Achieving this requires further development of high-performance codes to utilize GPU computing power as well as to improve parallel efficiency.

METHODS & CODES

The LBM [1] is particularly suited for numerical simulation of complex fluid flow with complex geometries. The LB color fluid multiphase model ensures a relatively sharp interface and completely immiscible fluids; therefore, it has been widely adopted for multiphase flow in porous media. We developed and optimized a portable and scalable LB code based on a variant of the multiple relaxation time (MRT) color-fluid multiphase LB model [2]. The code employs a MPI-OpenMP/OpenACC hybrid programming model so that it can run on CPUs, GPUs, and MIC processors. Preprocessors are employed to compile the code for different platforms. The code achieves two times speedup on Blue Waters' GPU node and scales almost ideally up to 512 GPU nodes thanks to the implementation of overlapped communication and computation. A typical simulation (on a 500X500X500 grid for 5 million iterations) takes 40 GPU nodes 43 hours to complete (1,736 node-hours).

RESULTS & IMPACT

During the BW project period, we further developed our existing code and successfully ported our code to GPUs and other manycore processors. We plan to make the code open source, which enables users to easily utilize computing power from manycore processors. The development of the code and its applications to pore-scale simulation of multiphase flow in porous media will be reported in a future publication.

LBM simulations and experiments of liquid CO₂ displacing water in a 2D heterogeneous micromodel at reservoir pressure conditions were performed and compared for the first time to evaluate the capability of LBM to capture detailed fluid invasion patterns. Our research collaborators at University of Notre Dame conducted the state-of-the-art experiments. The simulations reproduce qualitatively similar trends as the experiments for changes in invasion patterns as Ca (Capillary Number—ratio of viscous to capillary forces) increases, as shown in Fig. 1. The major discrepancy between experiments and simulations is the development of secondary CO₂ pathways (demarcated by ellipses in Fig. 1), which occur at a much higher Ca in the simulations. The inertial effects are ignored in most pore-scale studies due to the small bulk Re (Reynolds Number—ratio of inertial to viscous forces). However, recent studies show that local Re could be much higher than the bulk Re and hence the inertial effects may not be ignored. Due to computational cost and numerical stability issues, the LBM simulation uses proxy fluids with a much larger viscosity compared to the experiment. Our further studies of the effects of Re shows that the significantly reduced inertial effects in the

simulation alter the preferential multiphase flow pathways and are responsible for the discrepancy between the simulations and experiments. The above work has been submitted to the *Journal of Contaminant Hydrology* and is currently under review. Further investigation on the impact of inertial effects in real rocks will be performed in the future. We will demonstrate, for example, how the inertial effects alter the preferential flow pathways in a Bentheimer sandstone for different Ca, as shown in Fig. 2.

WHY BLUE WATERS

Due to the complex geometry and complex interfacial dynamics, the computational cost of DNS is very high. Blue Waters offers a large number of CPU and GPU nodes, which is essential to our research. Our code was initially designed for CPU platforms and was later ported to GPU platforms. The availability of both CPU and GPU nodes enabled us to perform required simulations while at the same time developing faster GPU code.

A single iteration step in LB simulation can be completed in a very short time, but a typical simulation requires millions of iterations. Therefore, LB simulation requires very low latency on message passing, which cloud resources cannot provide.




Figure 2: Invasion patterns of liquid CO₂ at different Ca on a Bentheimer sandstone. Brine is transparent and the solid surfaces are shown in grey.

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

Allocation: Illinois/300 Knh
PI: Matthew West¹
Co-PI: Nicole Riemer¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

This research aims at reducing key uncertainties in quantifying the impact of atmospheric aerosol particles on the Earth's climate. Aerosol particles can be brought into the atmosphere by a wide range of human activities or by natural sources. They profoundly impact the large-scale dynamics of the atmosphere because they interact with solar radiation, both directly by scattering and absorbing light and indirectly by forming cloud droplets. These impacts depend on the particles' sizes and their compositions, which continuously change in the atmosphere. The uncertainties in quantifying these impacts originate from scale interactions and the high computational cost required for modeling these. To tackle this problem, we developed the particle-resolved 3D model WRF-PartMC-MOSAIC, which has the unique ability to track size and composition information at a per-particle level. Particle-resolved simulations at the regional scale not only require efficient numerical algorithms but also a computational resource with the capabilities of Blue Waters. Together, these methods and the petascale resources allow for ultra-high-detail simulations that are needed to quantify the impact of aerosol particles on weather and climate at the regional scale.




Figure 1: Horizontal distribution of simulated number concentration of black-carbon-containing particles near the surface.

RESEARCH CHALLENGE

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

METHODS & CODES

To overcome the current limitations in representing aerosols and associated uncertainties, the particle-resolved model PartMC-MOSAIC [3] was coupled to the state-of-the-art 3D Weather Research and Forecast (WRF) model [4]. Aspects of these two models complement each other. The box model PartMC-MOSAIC is a highly detailed aerosol model that tracks the size and complex composition of individual particles in the atmosphere but is unable to resolve spatial heterogeneities of aerosol populations. The 3D regional WRF model is an advanced numerical weather model that captures the transport of chemical species in the atmosphere but assumes a crudely simplified aerosol representation. The resulting WRF-PartMC-MOSAIC model uses a 3D Eulerian grid for the atmospheric flow, while explicitly resolving the evolution of individual aerosol particles per grid cell. This next-generation model captures complex aerosol composition that current-generation models are unable to simulate.

RESULTS & IMPACT

We present results from the first-ever particle-resolved aerosol simulation for a realistic, spatially resolved three-dimensional domain (North Carolina). Aerosol and trace gas emissions were taken from the 2005 National Emission Inventory [5], and the

meteorology corresponded to January 24, 2000. On the order of 100 billion computational particles were tracked in this simulation, including their composition changes due to gas-to-particle conversion, and coagulation events.

Fig. 1 shows the modeling domain and the spatial distribution of black carbon-containing particle number concentrations near the surface after eight hours of simulation. Black carbon aerosol is of interest because of its adverse health impacts and because of its warming impact on climate. While this is a fundamental bulk quantity, common to any chemical transport model, the particle-resolved aerosol representation provides unprecedented detail of particle composition and source tracking. Fig. 2 (top) shows the originating sources of all particles within a given grid cell. This allows source attribution for any location within our domain. Fig. 2 (bottom) shows an example of the complex continuum of aerosol composition that exists within a single grid cell; particles of similar diameters can have very different chemical composition—information that is usually lost when using traditional aerosol models. The variations in particle composition are determined by their emission source characteristics, here with highway vehicles containing the largest black carbon mass fractions. During the simulation, aerosol composition evolves due to coagulation and condensation of secondary gas species, resulting in a complex continuum.

As the model tracks composition and source information of thousands of computational particles per grid cell, individual particles may also be explored. For example, a single particle with a particular size and black carbon mass fraction may be examined, marked with a red dot in Fig. 2. By tracking mass of constituent species (not shown), we can determine that the selected particle has grown considerably due to the condensation of nitrate. Additionally, by tracking source history (not shown), the contributing aerosol emissions sources can be determined where this selected particle has undergone multiple coagulation events with particles from different emission sources such as agriculture and fossil fuel combustion. These capabilities will be useful in future studies for quantifying how much individual source categories are contributing to the pollution at a certain location.

Aerosol modeling is challenging because of the multiscale nature of the problem—the macroscale aerosol impact on climate is determined by microscale processes on the particle scale. The WRF-PartMC-MOSAIC model provides a tool that represents many of these microscale processes explicitly, which allows for an improved process-level simulation of the key interactions among aerosols, clouds, and radiation. This model framework therefore serves as the first benchmark for more approximate models, and provides a basis for rigorous coarse-graining to develop physically robust parameterizations for use in larger scale models.

WHY BLUE WATERS

Access to Blue Waters allows for a cutting edge model formulation that pushes both science and computing by combining the large-scale features of state-of-the-art 3D models with the

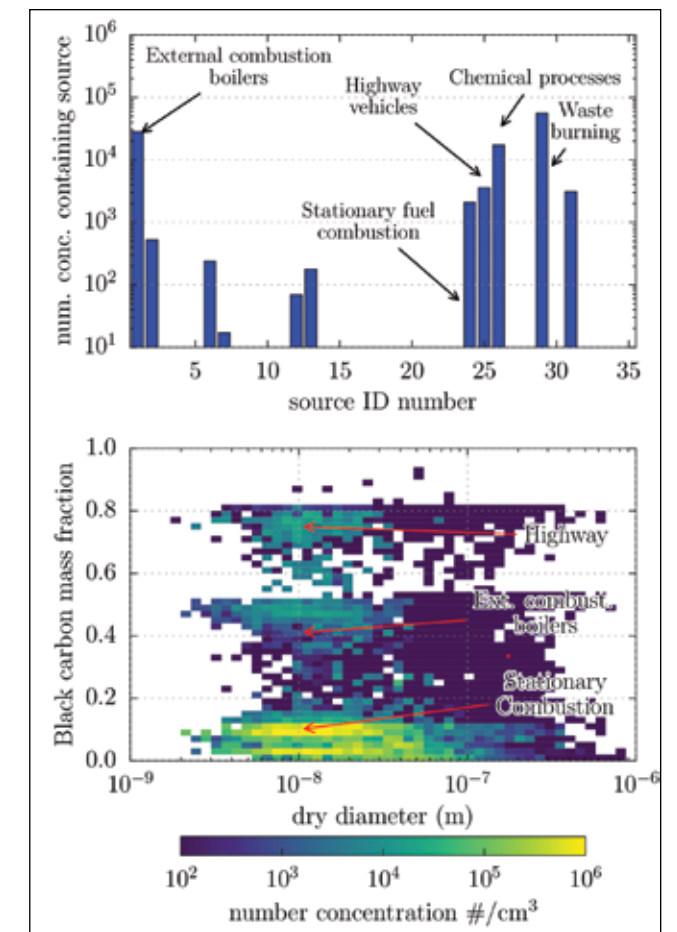


Figure 2: WRF-PartMC-MOSAIC resolves population-level details. (Top) Emission source information is tracked for each grid cell. (Bottom) Two-dimensional number distribution as a function of particle dry diameter and black carbon mass fraction indicates the amount of particles within a range of diameters and a range of fractions of black carbon mass.

process level physical representation of box models. Modeling 3D domains with on the order of 100 billion tracked particles creates many computational challenges due to computationally intensive equations per particle and memory requirements to track high-dimensional particle composition. To enable simulations of aerosols at both a high spatial and compositional resolution, there is a need for tens of thousands of cores, fast interconnections among those cores, and sufficient memory per process.

PUBLICATIONS AND DATA SETS

PartMC particle-resolved aerosol simulation code, version 2.4.0, released February 2, 2017.

PARTICULATE MATTER PREDICTION AND SOURCE ATTRIBUTION FOR U.S. AIR QUALITY MANAGEMENT IN A CHANGING WORLD

Allocation: Illinois/500 Knh
PI: Donald J. Wuebbles¹
Co-PI: Xin-Zhong Liang²
Collaborator: Swarnali Sanyal¹

¹University of Illinois at Urbana-Champaign
²University of Maryland

EXECUTIVE SUMMARY

The objectives of this study are to better understand how global changes in climate and emissions will affect pollution in the United States, focusing on particulate matter and ozone, to project their future trends, and to quantify key source attributions. We are using a state-of-the-science dynamic prediction system that couples a global climate–chemical transport model with regional climate–air quality models over North America to determine individual and combined impacts of global climate and emissions changes on air quality, with uncertainty estimate, from the present to 2050 under multiple scenarios. We are doing the long-term global climate chemistry runs using Community Earth System Model CESM1.2.2 with fully coupled chemistry using CAM5-chem at 0.9° x 1.25° horizontal resolution, then comparing the results with observational data to evaluate the model simulation.

RESEARCH CHALLENGE

Our goal is to determine the individual and combined impacts of global climate and emissions changes on U.S. air quality from the present to 2050 under multiple scenarios, quantifying pollution sources and assigning their attribution—natural vs. anthropogenic emissions, national vs. international agents, natural variations vs. climate changes—with associated probability and uncertainty. We are developing a timeline for the global change factors to become significant such that effective actions can be taken. The level of significance is defined following the cross-state air pollution rule as one percent of nonattainment areas with the goal of bringing all areas into attainment for the National Ambient Air Quality Standards. Our hypothesis is that the integration of the most advanced modeling system, most updated emissions treatment, multiscale processes representation, and a multi-climate–emission scenarios assessment will improve the predictive capability and result in more reliable projection of future changes in particulate matter, ozone, and related pollutants as well as their global and regional sources.

This research presents a state-of-the-science approach for advancing quantitative knowledge of the impacts of global changes in climate and emissions on U.S. air quality. The Global Climate Chemistry Transport model (GCCT) integrates global climate change with long-range pollutant transport that links worldwide natural and anthropogenic source emissions, while

providing lateral boundary conditions that drive the Regional Climate–Air Quality model (RCAQ) for regional climate and air quality prediction. RCAQ incorporates more complete physical representation (surface, precipitation, convection, cloud, aerosol, and radiation); comprehensive chemical mechanisms (e.g., secondary organic aerosols or SOAs); and detailed emissions treatment. Hence, it more realistically simulates interactions between surface and atmospheric processes at regional–local scales that in turn affect local air quality. This nested GCCT/RCAQ dynamic prediction system is being evaluated against observations, and subject to process-level understanding and source attribution of U.S. air quality episodes under present and future conditions.

METHODS & CODES

We are conducting three primary experiments using the dynamic prediction system: (1) historical simulations for period 1994–2013 to establish the credibility of the system and refine process-level understanding of U.S. regional air quality; (2) projections for the period 2041–2060 to quantify individual and combined impacts of global climate and emissions changes under multiple scenarios; and, (3) sensitivity analyses to determine future changes in pollution sources and their relative contributions, from anthropogenic and natural emissions, long-range pollutant transport, and climate change effects.

We are also conducting a series of 20-year runs using CESM1.2 (CAM-chem5). We have completed short-term global CAM-chem simulations driven by NASA Modern-Era Retrospective analysis for Research and Applications (MERRA) reanalysis data at 0.9° x 1.25° horizontal. We used the FSTARTMAM7 component set, which uses seven modes to model aerosols, prescribed ocean and ice, and CAM5 physics with carbon and nitrogen in the Community Land Model (CLM). CAM5-chem has 160 species with 427 reactions, and has strict enforcement of the conservation of total (organic and inorganic) chlorine and total bromine under advection. The heterogeneous chemistry module has also been upgraded to reflect the underestimation of supercooled ternary solution and surface area density, in addition to an improved aerosol treatment and inclusion of aerosol–cloud interactions, with extensive tropospheric and stratospheric chemistry. It has 30 vertical levels with model top at about 40 km. The dust emission is calibrated in the model so that global dust aerosol optical depth

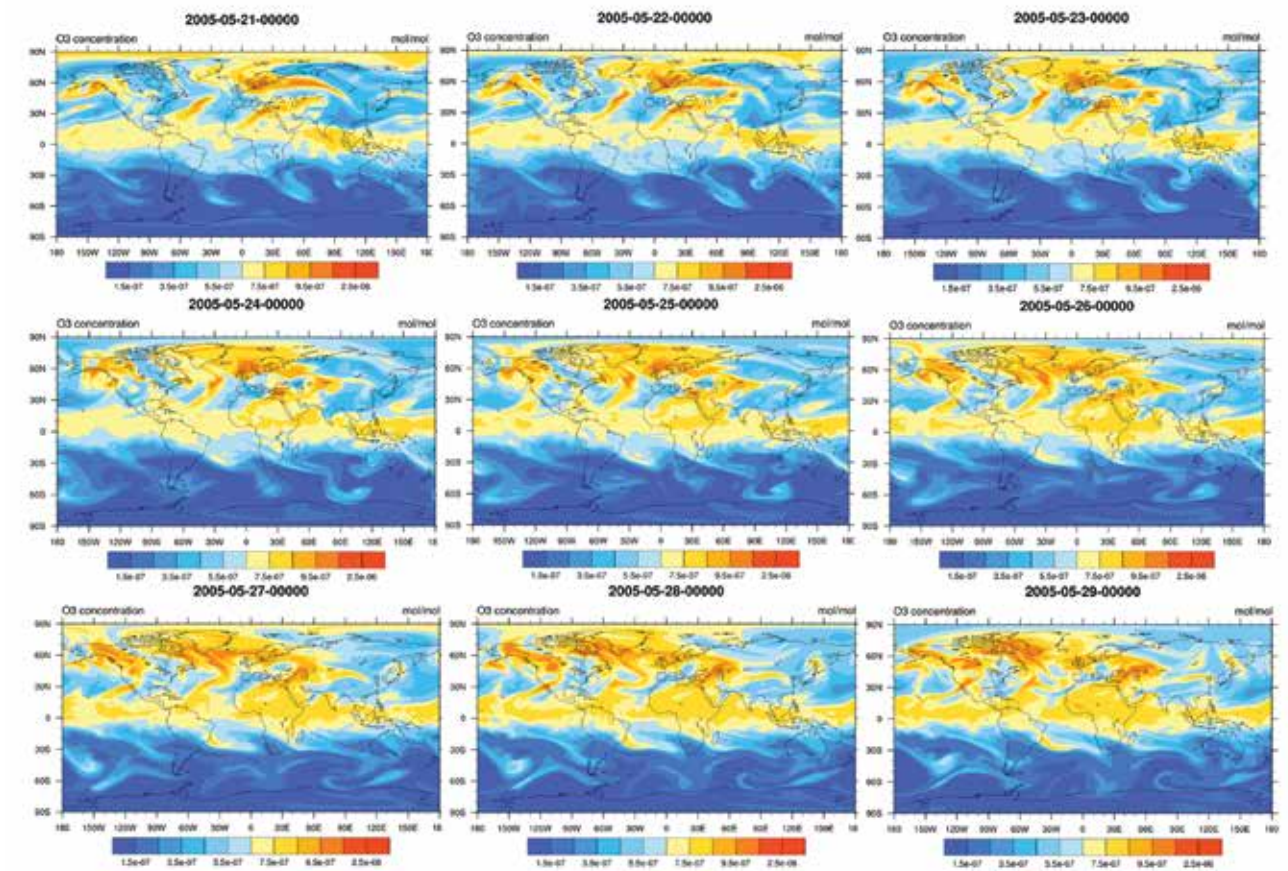


Figure 1: Evolution of daily ozone concentration for six consecutive days in May 2005 from Cam5-chem global simulation.

is between 0.025 and 0.030, and the system couples tropospheric aerosol to chemistry through heterogeneous chemistry. In this run, we used CESM1.2 default emissions, which represent surface emissions of approximately 30 species of aerosols. The surface emission of each species is composed of all possible sources of emissions, including those from biomass burning, domestic sources, transportation, waste treatment, ships, industry, fossil fuels, and biofuels, and were composed from POET, REAS, GFEDv2 and FINN emissions databases [1].

RESULTS & IMPACT

We did short-term global CAM-chem simulations driven by MERRA reanalysis data at 0.9° x 1.25° horizontal resolution and FSTRATMAM7 component set. Fig. 1 shows the global O₃ concentration for May 2005. We will be doing a variety of additional modeling simulations and comparing the results. Previous studies have shown that trans-Pacific dust and aerosols contribute significantly to North American aerosol inflow [2], while export-related Chinese pollutants contributed 3–10% of annual mean sulfate concentration, 0.5–1.5% of ozone, and one or more day of noncompliance of ozone standards over many U.S. regions in 2006 [3]. We are in the process of investigating the long-range transport of pollutants from China to the United States.

WHY BLUE WATERS

The computational demand of the high-resolution climate model used in this project is extensive. We are using the fully coupled model of the Earth’s climate system with chemistry, which is computationally expensive. Blue Waters, with its petascale computational facility, large number of nodes, storage capability for the output from the high-resolution model simulation, is essential for our project. Blue Waters staff have been critical in figuring out the various issues arising with the long-term fully coupled climate chemistry runs with CESM. The staff have helped figure out and resolve various issues with the CESM1.2.2 models. Blue Waters has given us the computational resource, data management, and support staff to perform our research.

PHYSICS & ENGINEERING

FLUIDS

MATERIALS

NANOTECHNOLOGY

NEUTRONICS AND THERMAL HYDRAULICS

NUCLEAR PHYSICS

QUANTUM

- 102** *Exploring Confinement vs. Orientation Effects in Rigid and Semi-Flexible Polymers Using a Massively Parallel Framework*
- 104** *Multi-Scale and Multi-Physics Modeling of The Strength of Geopolymer Composites*
- 106** *Harnessing Petascale Computing to Explain Fundamental Mechanisms Driving Nanopatterning of Multicomponent Surfaces by Directed Irradiation Synthesis*
- 108** *Study of DIBs with Functional Channels*
- 110** *Petaflops Simulation and Design of Nanoscale Materials and Devices*
- 112** *Reducing Jet Aircraft Noise*
- 114** *Properties of Dense Hydrogen*
- 116** *Scalable Nanopatterning of Graphene by Hydrogen-Plasma Etching*
- 118** *hPIC: A Scalable Electrostatic Particle-In-Cell for Plasma–Material Interactions*
- 120** *Innovative Ab Initio Symmetry-Adapted No-Core Shell Model for Advancing Fundamental Physics and Astrophysics*
- 122** *DNS of Pressure Fluctuations Induced by Supersonic Turbulent Boundary Layers*
- 124** *Dispersion of Fully Resolved Liquid Droplets in Isotropic Turbulent Flow*
- 126** *QMCBD: A Living Database to Accelerate Worldwide Development and Use of Quantum Monte Carlo Methods*
- 128** *Numerical Methods and Software for Computational Fluid Dynamics, NEK5000*
- 130** *Direct Numerical Simulation of Turbulence and Sediment Transport in Oscillatory Boundary Layer Flows*
- 132** *3D Probabilistic Physics-Based Seismic Hazard Maps for Regional Risk Analysis*
- 134** *Optimal Bio-Locomotion Strategies in Fluids*
- 136** *Coupled Multi-Physics of Advanced Molten Salt Nuclear Reactors*
- 138** *Understanding Hydrogen Storage in Metal Organic Frameworks using Massively-Parallel Electronic Structure Calculations*
- 140** *Atomistic Modeling of Transformations in Nanocrystals*
- 142** *Numerical Simulations of Collapsing Cavitation Bubbles on Blue Waters*
- 144** *Leading Future Electronics into the Nano Regime Using Quantum Atomistic Simulations in NEMO5*
- 146** *Kinetic Simulations of Unsteady Shock-Boundary Layer Interactions*
- 148** *High Energy Physics on Blue Waters*
- 150** *First-Principles Computer Simulations of Hydrocarbons Under Fusion Conditions*
- 152** *Direct Numerical Simulations of the Relative Motion of High-Inertia Particles in Isotropic Turbulence*
- 154** *Mapping Proton Quark Structure Using Petabytes of COMPASS Data*
- 156** *Optical Determination of Crystal Phase in Semiconductor Nanocrystals*
- 158** *Advanced Digital Technology for Materials and Manufacturing*
- 160** *Transient Multiphase Flow Phenomena and Defect Formation in Steel Continuous Casting*
- 162** *High Resolution Numerical Simulation of Oscillatory Flow and Sediment Transport through Aquatic Vegetation*
- 164** *Quantum Monte Carlo Simulations of Magnetism and Models in Condensed Matter*
- 166** *Intermittency, Resolution Effects and High Schmidt Number Mixing in Turbulence*

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

PI: David Ackerman¹
 Collaborator: Baskar Ganapathysubramanian¹

¹Iowa State University

EXECUTIVE SUMMARY

The overarching goal of this work is to develop and utilize a highly parallel computational framework to model equilibrium structures of semi-flexible polymers efficiently in complex, confined, non-periodic geometries. This will allow simulation of a wide range of technologically and biologically relevant systems such as DNA, RNA, and viral capsids; polymer chains with stiff backbones used in organic electronics (field effect transistors, photovoltaics, thermoelectrics); liquid crystals, where orientation and alignment create ordered phases; and rod-like polymers used in membranes and sensors.

To handle the range of applications, the simulation system is designed to be capable of scaling to large numbers of processors for large, complex systems of semi-flexible polymer systems. It is also designed to use high-throughput simulations for parameter sweeps of smaller or less complex systems. This work has developed tools to meet these goals.

RESEARCH CHALLENGE

An understanding of the morphology of polymers under different conditions aids in materials design and discovery. Self-consistent field theory (SCFT) [1] simulations are a tool for efficiently modeling polymer microstructures. While the

method has proven very successful, most SCFT simulations are not designed for scaling to large systems with arbitrary geometries and boundary conditions. They are also not designed to handle polymer systems with semi-flexible chains. This puts limits on the applicability of those simulations. A generic, highly scalable SCFT simulation tool capable of modeling semi-flexible polymers greatly expands the types of polymer systems that SCFT studies can address. These tools will be of great interest to polymer scientists and materials researchers. It will also enable the ultimate goal of studying confinement effects on semi-flexible chains.

METHODS & CODES.

To address the desired goals of highly scalable simulations of systems having arbitrary geometries and boundaries with complex chain models, we utilized a finite element approach with several advanced techniques. Our implementation uses a custom finite element library built upon the Portable, Extensible Toolkit for Scientific Computation (PETSc) [2].

The SCFT method utilizes an evolving chain contour, analogous to a time-dependent problem. To provide enhanced scaling, we use a coupled space-time model in the contour solver. In contrast to typical time-dependent finite element models that solve for each step sequentially, we can utilize a simultaneous solution over the




Figure 1: Example polymer microstructures generated with the framework developed in this project.

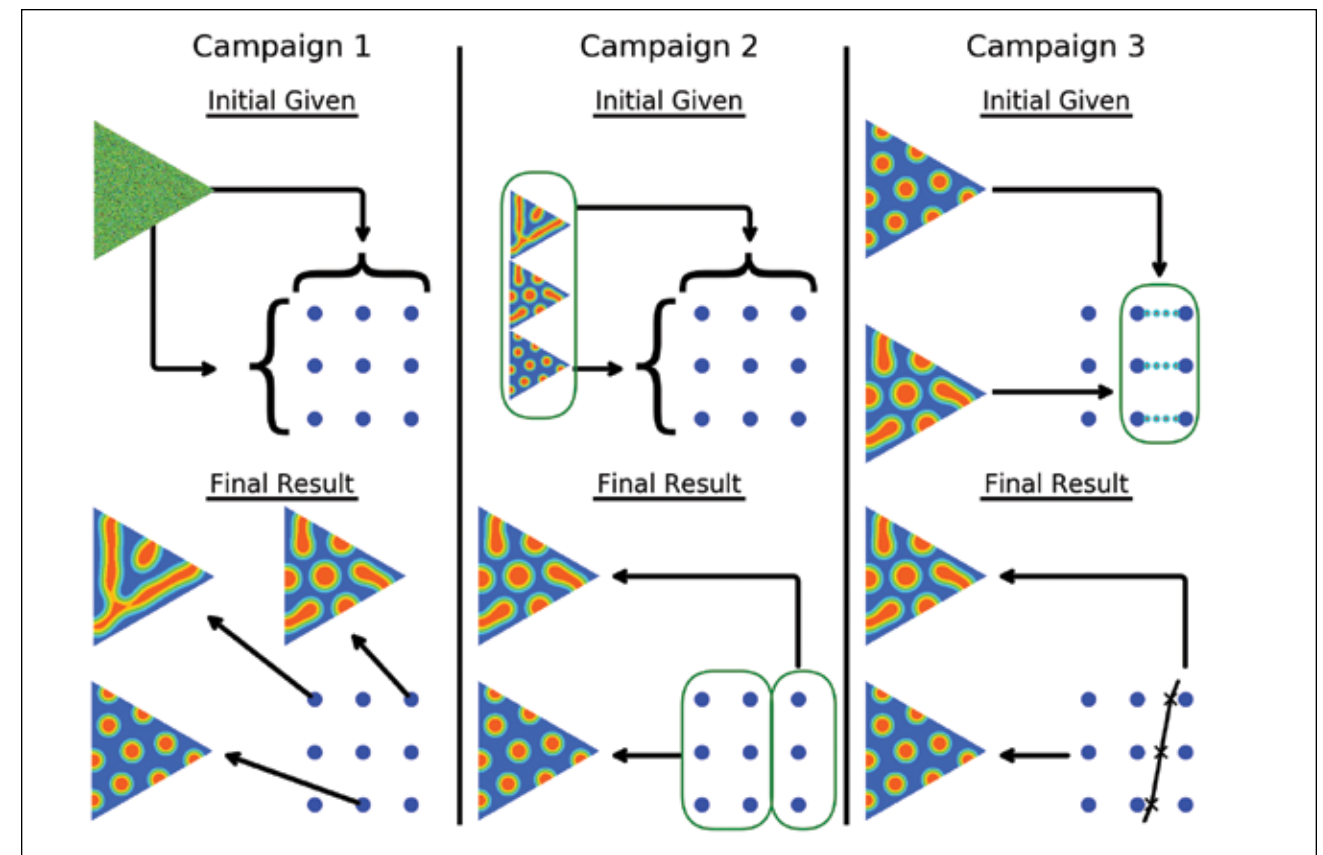


Figure 2: High throughput workflow for polymer phase diagram studies utilizing simulation “campaigns” for identifying equilibrium microstructures across a range of configuration parameters.

entire contour space. This process enables efficient adaptivity and leads to enhanced scaling as the number of processors increases.

To model semi-flexible polymers, we have developed and utilized an approach that formulates this problem in an appropriate finite element space (the product space of spatial discretization and orientation discretization). This uses a two-level finite element method [3]. In this approach, the entire spatial system is a finite element volume mesh. At each nodal point within the volume mesh, a spherical surface mesh handles the orientations. The coupling of the two simulation spaces enables orientationally dependent systems to be simulated.

RESULTS & IMPACT

This work has had three major results: First, the development, testing, and deployment of a scalable finite element SCFT code. This has been utilized in a high-throughput study of polymer confinement and will be used in several future studies. Development of this code included design of an efficient process for the construction of polymer structure phase diagrams.

Second, implementation and testing of a simultaneous space-time finite element method. Access to the Blue Waters supercomputer enabled testing of this system by providing access to numbers of CPU cores unavailable elsewhere. This ensured a

full view of the capabilities and limitations of this method. This is invaluable for a method aimed at enhancing scaling.

Third, development of a framework for modeling semi-flexible polymers using SCFT. This enables morphology studies of polymers with complex, orientation-dependent properties. The computational demands of polymer models have limited past work on these types of polymers. The resources of Blue Waters are well suited to the computational demands of this work.

WHY BLUE WATERS

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

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

Allocation: Exploratory/50 Knh

PI: Ange-Therese Akono¹

Co-PIs: Erman Guleryuz^{1,2}, Seid Koric^{1,2}, Waltraud M. Kriven¹

¹University of Illinois at Urbana-Champaign

²National Center for Supercomputing Applications

EXECUTIVE SUMMARY

Geopolymers are a class of inorganic polymeric, X-ray amorphous materials consisting of alumina, silica, and alkali metal oxides [1]. Geopolymer composites exhibit appealing properties such as high thermal stability [2], small carbon dioxide footprint [3], and high strength, which makes them suitable for many applications including alternative cementitious materials [4], passive cooling systems [5], low-level nuclear waste encapsulation [6], and potential novel biomaterials. The research goal is to understand the impact of nanoporosity on stiffness and strength via molecular dynamics and finite element modeling. Nanoporosity is governed by monovalent cation and water content. In contrast, microporosity results from the mixing procedure and from poor bonding between the geopolymer matrix and the filler phase. To increase the performance of geopolymer composites several strategies were found: reduce the microporosity, select strong inclusions, increase the volume content of stiff inclusions, or select fillers with a high aspect ratio. These findings pave a new way toward novel, high-performance, and multi-functional composites.

RESEARCH CHALLENGE

Geopolymer composites are an emerging class of composites with great potential in civil engineering, mechanical engineering,

aerospace, navy, automobile, and biomedical engineering. However, the widespread application and acceptance of geopolymer systems has been so far impeded by many roadblocks such as lack of long-term durability data [7], lack of in-service track record, and lack of standard geopolymer cements [8]. The computational approach selected based on molecular dynamics and finite elements offer a cost-effective and time-efficient means to accelerate discovery and innovation.

METHODS & CODES

At the nanoscale, atomic simulations were performed in LAMMPS, as shown in Fig. 2b. Various molecular structures were recreated using the Avogadro and Packmol software, for diverse alkali metal cations and Si:Al ratios. The amorphous structure was obtained through a melting and quenching sequence. Afterward, uniaxial tensile and compression tests along with shear tests were simulated to yield the strength behavior and failure micromechanisms. The software OVITO was utilized for visualization. At the mesoscale, the finite element package Abaqus was used (Fig. 2a). The microstructure was generated from OOF2D and MATLAB based on scanning electron microscopy observations of geopolymer composites. The constitutive behavior of the individual constituents was prescribed based on independent

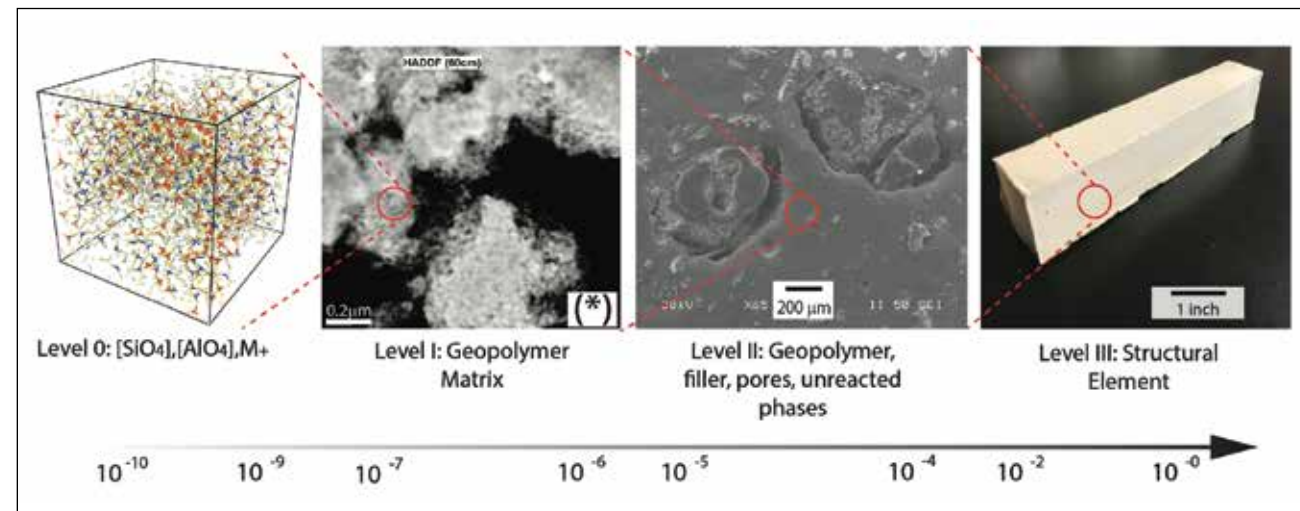


Figure 1: Multiscale nature of geopolymer composites. Our study spans levels 0 at the nanometer length-scale to level III at the macroscopic level.

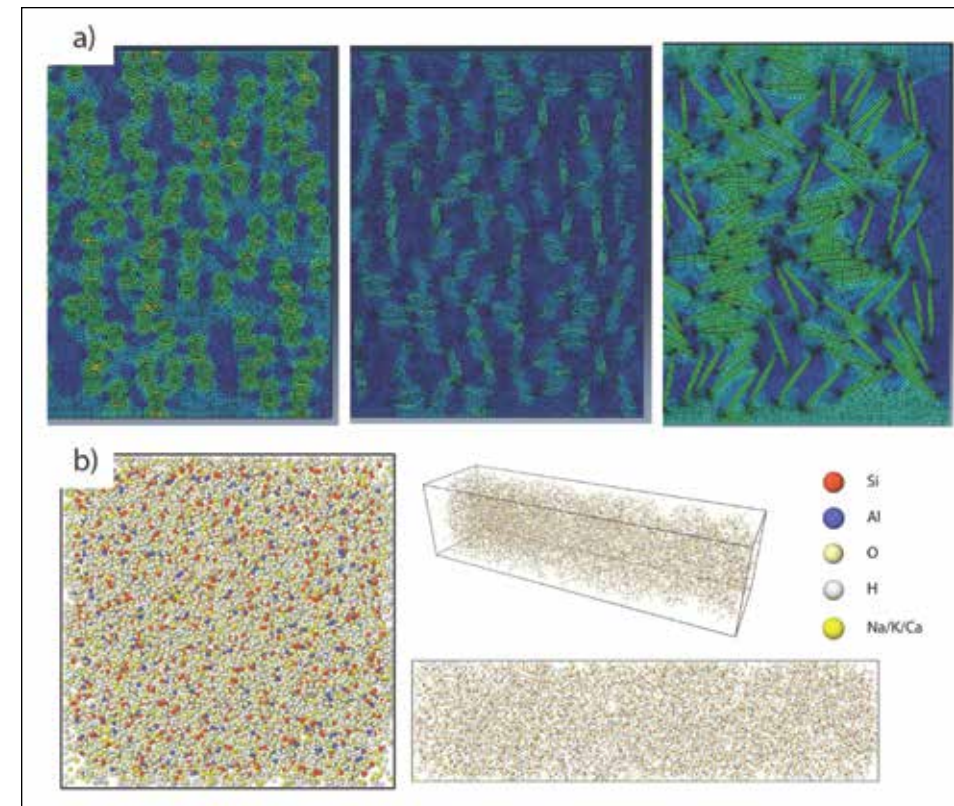


Figure 2: a) FEM model at mesoscale showing the stress distribution as a function of the filler aspect ratio. b) Molecular dynamics modelling of inorganic polysialate disiloxo.

experiments. Tensile tests with prescribed periodic boundary conditions were simulated to yield the effective response.

RESULTS & IMPACT

The strength and stiffness of amorphous geopolymers were correlated to its chemistry and density. The nanoporosity was found to be a function of the alkali and Si:Al ratio. At the mesoscale, the mechanical resistance of geopolymer composites to permanent deformation was elucidated as a function of the type, nature, size, and shape of the inclusions. Our theoretical and computational framework was validated on 31 different geopolymer-based systems, based on experiments carried out by various researchers over seven years. To our knowledge, it is the first time that a theoretical model has been proposed to upscale the constitutive behavior of geopolymers. In future efforts, we will model the evolution of the fracture toughness of pure geopolymer for various densities via molecular dynamics.

WHY BLUE WATERS

The Blue Waters platform was essential to carry out our molecular dynamics simulations as well as our finite element simulations. Standard desktop workstations do not have the required memory and computational power to resolve the complex systems involved. For molecular dynamics, it was crucial to be able to simulate large systems in a timely fashion. In the case of finite element simulations, the nonlinearity of the equations due to friction, contact, and plastic flow made it impossible to

obtain results using a desktop workstation. In both cases, the Blue Waters supercomputer has enabled us to obtain novel results that bring new insights into the origins of strength and toughness in geopolymer composites.

PUBLICATIONS AND DATA SETS

Kataruka, A., S. Koric, W. Kriven, and A.-T., Akono, Strength Properties of Particulate Potassium-Based Geopolymer Composites. *ASCE Engineering Mechanics Institute Meeting*, San Diego, Calif., June 4–7, 2017.

Cui, Y., E. Guleryuz, and A.-T. Akono, Investigation of the Mechanical Properties of geopolymer binders using molecular dynamics simulations. *ASCE Engineering Mechanics Institute Meeting*, San Diego, Calif., June 4–7, 2017.

Kataruka, A., et al., Representative Elementary Volume Modeling of Geopolymer Composites. *14th U.S. National Congress on Computational Mechanics*, Montreal, QC, Canada, July 17–20, 2017.

Cui, Y., et al., Molecular dynamics study on the mechanical and fracture properties of geopolymer binders. *8th Advances in Cement-Based Materials American Ceramic Society*, Atlanta, Ga., June 26–28, 2017.

Kataruka, A., et al., A Parametric Study of Strength Properties of Particulate Potassium-Based Geopolymer Composites. *8th Advances in Cement-Based Materials American Ceramic Society*, Atlanta, Ga., June 26–28, 2017.

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

Allocation: Illinois/250 Knh

PI: Jean Paul Allain¹

Collaborators: Michael A. Lively¹, Brandon J. Holybee¹, Michael Toriyama¹, Scott A. Norris²

¹University of Illinois at Urbana-Champaign

²Southern Methodist University

EXECUTIVE SUMMARY

Ion beam-driven synthesis of quantum dots on III–V semiconductor surfaces is a promising approach toward a device-scale fabrication process for solar cells, quantum dot lasers, and other applications. However, the physical mechanisms that determine the pattern formation and properties such as quantum dot spacing remain poorly understood. To close this knowledge gap, massive-scale molecular dynamics simulations have been performed studying ion beam bombardment of gallium antimonide (GaSb) surfaces with 500 eV Ne⁺, Ar⁺, and Kr⁺ ions. Incident energetic ions are able to induce the formation of single-element clusters within the GaSb surface, with Sb clusters being far more prevalent than Ga clusters. These clusters may provide the basis for long-timescale mechanisms beyond the scope of molecular dynamics, such as radiation-enhanced diffusion, to create the compositional depth profile observed in experimental studies.

INTRODUCTION

The ability to synthesize unique nanostructures and nanopatterns at surfaces by ion beam bombardment holds great potential as a scalable technique to create novel nanomaterials in a single process step. In particular, the ability to fabricate quantum dots on III–V semiconductor surfaces, such as Gallium Arsenide (GaAs) or Gallium Antimonide (GaSb), using ion beam irradiation could be used to efficiently functionalize those materials for use in high-efficiency solar cells or quantum dot lasers. However, since the first discovery of these nanodots [1], theoretical models have been unable to provide a complete, predictive description of the nanopattern formation on these surfaces. Consequently, the dependence of nanopattern formation on III–V semiconductor surfaces on experimental parameters such as the incident ion mass or energy remains poorly understood.




Figure 1: Snapshots of the GaSb surface after irradiation by 500 eV Kr⁺ ions to a fluence of 7.5×10^{15} cm⁻²—(a) snapshot of the entire surface after ion bombardment, showing the lack of any notable compositional depth profile; (b) snapshot showing only Ga atoms in clusters; (c) snapshot showing only Sb atoms in clusters. Sb atoms show a significantly greater tendency to form clusters within the irradiated GaSb surface.

To address this knowledge gap, atom-level computational modeling is necessary to observe and characterize the ion-induced changes in the surface as they occur, providing detailed insight into the ion-surface interactions at length and time scales far smaller than what can be experimentally observed. At the same time, any simulation efforts must also be able to connect the small-scale atomistic mechanisms to the evolution of the larger surface at experimentally observable length and time scales. Determining both the large and small scales of ion-driven nanopatterning can only be done with high-performance computational power on a massive scale, well beyond that provided by conventional computing platforms.

In previous work conducted on Blue Waters [2], massive-scale molecular dynamics simulations were conducted to connect experimental observations of the compositional depth profile [3,4] to pattern-forming surface instabilities. The results showed, for the first time, that an ion-induced compositional depth profile led to lateral phase separation, providing a pathway for the pattern-forming instability to emerge. The present simulations seek to elucidate the ion-induced mechanisms causing the formation of the compositional depth profile that drives these changes.

METHODS & RESULTS

Molecular dynamics (MD) simulations of ion bombardment are carried out with the LAMMPS package [5] for 500 eV Ne⁺, Ar⁺, and Kr⁺ ion species incident on an initially pristine GaSb surface. The simulation cells were 25×25 nm², providing sufficient space for lateral compositional variations to evolve on the same scale observed in previous simulations. For each ion species, the GaSb surface was irradiated to a fluence of 7.5×10^{15} ions/cm², an experimentally relevant fluence (the sum of the energies of the particles per unit area contained in the particles with which a material is irradiated) on the same order of magnitude as the threshold fluences for pattern formation that have been observed in experiments.

Snapshots of the resulting surfaces after ion irradiation are shown in Fig. 1 for the 500 eV Kr⁺ case. Three major observations can be made from these snapshots. First, despite irradiation to a fairly large fluence, no compositional depth profile was observed. This indicates that additional long-temporal scale mechanisms are necessary to accurately model the compositional evolution of the surface, such as radiation-enhanced diffusion or stress-driven flow. Second, while a depth profile does not emerge, the formation of smaller clusters of Ga or Sb can still be observed due to the prompt ion effects. This indicates that even at short timescales the energetic ions still induce compositional changes in the surface, which may provide “seed” structures to grow into larger-scale lateral compositional gradients. Finally, Sb atoms show a much stronger clustering tendency than Ga atoms, which may be indicative of the higher mobility of Ga compared to Sb seen in experimental studies [6].

Interestingly, different ion species were shown to have differing effects on the surface compositional evolution, as shown in Fig.




Figure 2: Fraction of Sb atoms which are members of clusters as a function of increasing fluence for each ion species. While all three ion species initially induce similar clustering behavior, for fluences greater than 2.4×10^{15} cm⁻² Kr⁺ irradiation causes a decrease in clustering compared to the lighter two ion species.

2. In the early stage of ion bombardment, the tendency for Sb clusters to form is similar for all ion species. However, at higher fluences the Sb clustering tendency decreases for Kr⁺ irradiation, while it increases asymptotically for the lighter Ne⁺ and Ar⁺ ions. While the implications of this observation are still unclear, the difference correlates well with the experimentally determined threshold fluences for nanopattern formation, which is a factor of four greater for Ne⁺ and Ar⁺ ions compared to that for Kr⁺ ions [4].

WHY BLUE WATERS

This work required carrying out three MD simulations in parallel, each of which contained around 150,000 to 200,000 atoms at any time, for over 230 million time steps each. On a conventional computing cluster, this would be practically impossible and would require several years to complete, at best. Furthermore, several terabytes of data were generated from these simulations, necessitating a high-performance parallel file I/O system to minimize the impact on the overall computation time.

PUBLICATIONS AND DATA SETS

Lively, M. A., B. Holybee, M. Toriyama, and J. P. Allain, Massive-scale molecular dynamics of ion-irradiated III–V compound semiconductors at the onset of nanopatterning, *Nucl. Instrum. Meth. Phys. Res. B* (2017), in press.

STUDY OF DIBS WITH FUNCTIONAL CHANNELS

Allocation: Illinois/975 Knh
PI: Narayana R. Aluru¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Distinguishing the different subclasses of immunoglobulin G (IgG) antibodies in blood serum can enable breakthrough advances in mapping the immune system and the health status of the human body. In this study, using petascale-based molecular simulations (containing up to ~1,000,000 atoms) and a total aggregate simulation time of 2.7 microseconds (μs), we demonstrate that an atomically thin graphene nanopore is capable of sensing and discriminating among different subclasses of IgG antibodies despite their minor and subtle variations in atomic structure. Using machine learning, we *featurized* and clustered the ionic current and the dwell times data obtained from the device during multiple antibody translocation events. In addition, the histogram of ionic current for each segment of IgG can provide high-resolution spatial detection of antibody segments. Parallel nanofluidic studies during IgG translocation reveal distinct water flux rates for IgG subclasses facilitating an additional recognition mechanism.

DESCRIPTION OF RESULTS

DNA sequencing using nanopore technology has evolved significantly during the last few years. Oxford Nanopore Technologies Ltd. currently is fabricating a USB-stick-sized device that can sequence the DNA in a few hours. In recent years, both biological and synthetic nanopores were used for “label free,” high-resolution DNA sequencing. In addition to DNA sequencing, detection of antibody proteins can lead to advances in improving human health. The challenges posed to biological molecule detection using nanopore technology are the low signal-to-noise ratio, pore degradation due to multiple uses, the identification of single bases in real time, and the high speed of translocation [1,2]. Engineering the translocation of molecules through biological/synthetic nanopores has been defined as one of the challenging problems of biotechnology. By using extensive molecular dynamics (MD) simulations, this study shows that an atomically thin graphene nanopore is capable of sensing and discriminating between different subclasses of IgG antibodies.

Protein detection via a graphene nanopore is accomplished using ionic current, dwell time, and water flux calculations. A total aggregate of 2.7 μs of simulation time has been carried out for systems containing up to ~1 million atoms. All parts of the antibody (Fig. 1) are distinguishable by ionic current measures. More specifically, the Fab, Fc, and hinge regions exhibit a unique current level when translocating through the pore. Using k-means clustering, the ionic current–dwell time and water flux–ionic

current feature plots lead to clusters with distinguishable centroids. We also compared the performance of the single-layer graphene nanopore with that of a solid-state nanopore (Si_3N_4). IgG subclasses are not distinguishable when using thick nanopores of Si_3N_4 because some of the atomic details cannot be captured.

In conclusion, we have shown that ionic current, dwell time, and water flux can detect different antibodies with high precision.

WHY BLUE WATERS

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

PUBLICATIONS AND DATA SETS

Farimani, A. B., M. Heiranian, K. Min, and N. R. Aluru, Antibody Subclass Detection Using Graphene Nanopores. *The Journal of Physical Chemistry Letters*, 8:7 (2017), pp. 1670–1676.




Figure 1: System consisting of IgG3 protein (red: chains; yellow: disulfide bonds), ions (pale green), and graphene sheet (green).

PETAFLUPS SIMULATION AND DESIGN OF NANOSCALE MATERIALS AND DEVICES

Allocation: NSF PRAC/3,120 Knh

PI: Jerzy Bernholc¹

Co-PIs: Emil L. Briggs¹, Miroslav Hodak¹, Carl T. Kelley¹, Wenchang Lu¹

Collaborators: Rui Dong¹, Yan Li¹, Zhongcan Xiao¹, Bing Zhang¹, Jiayong Zhang¹

¹North Carolina State University

EXECUTIVE SUMMARY

This project focuses on high-performance calculations for materials and devices of high current interest and on development of petascale methods for such simulations. This past year, we concentrated on three challenges. First, we investigated two paradigmatic sensor configurations for detection of biologically important molecules through *ab initio* calculations: a noncovalently functionalized nanotube for glucose detection, and a covalently functionalized nanotube for ethylene detection. Second, we analyzed high-performance dielectric materials that store and release energy electrostatically through polarization and depolarization. We showed that blending two disparate strongly polar polymers, e.g., a poly(arylene ether urea) (PEEU, $K = 4.7$) and an aromatic polythiourea (ArPTU, $K = 4.4$), the resulting mixture exhibits a very high dielectric constant, $K = 7.5$, while maintaining low dielectric loss ($< 1\%$). Third, in the pursuit of atomically precise and bottom-up fabrication of graphene-based electronics, we determined and understood the atomic mechanism responsible for controlled conversion of a polymer to a nanoribbon, stimulated by hole injection through an scanning tunneling microscope (STM) tip.




Figure 1: Optimized structures for ethylene detection in semiconducting CNT(8,0)-based systems (a, b, c) and metallic CNT(5,5)-based systems (d, e, f). The copper-carbon bond lengths (in Å) are labeled in brown; the closest distances between the ethylene molecule and the nanotube surface (in Å) are labeled in green.

RESEARCH CHALLENGE

Biomolecular detection is a rapidly growing field in biochemical and biomedical sciences. It is widely recognized as one of the key technologies for environmental monitoring and disease diagnosis. Functionalized carbon nanotube (CNT)-based nanosensors have shown great potential for detection, due to their sensitivity and selectivity. We computationally investigated the sensing mechanisms and capabilities of two paradigmatic sensor configurations based on single-walled carbon nanotubes (SWCNT): the covalently functionalized sensor for ethylene detection and the noncovalently functionalized sensor for glucose detection.

Dielectric materials store energy electrostatically through various polarization mechanisms and release it by depolarization. Dielectric capacitors are unparalleled in flexibility, adaptability, and efficiency for electrical energy storage, filtering, and power conditioning. We have discovered a practical and widely applicable mechanism for enhancing the dielectric constant through nanostructure engineering of dipolar polymers.

In the pursuit of atomically precise and bottom-up fabrication of graphene-based electronics, impressive advances have been made with the synthesis of graphene nanoribbons (GNR) by polymerizing monomers on different catalytic metal surfaces. We investigated the formation of GNRs from quasi-freestanding polymers assisted by hole injections from an STM tip. While a catalytic cyclodehydrogenation occurs in a cooperative domino-like conversion process during the thermal annealing, the hole-injection-assisted reaction happens at controlled molecular sites selected by the STM tip.

METHODS & CODES

The quantum transport and large-scale electronic structure calculations used the RMG code that we developed. For polymer simulations, the LAMMPS code was used. For calculations that include van der Waals interactions, the PWSCF code was used.

RESULTS & IMPACT

The sensing mechanism of glucose detection that we established can be generalized to detect other carbohydrate molecules, which are also known as saccharides. Carbohydrates are involved in a wide range of biological and pathological processes, such as cancer metastasis, cell signaling, protein function regulation, and

cellular communication. Detection of specific carbohydrates is thus essential for investigating their roles in numerous natural processes, as well as for obtaining insights into the mechanisms involved and diseases provoked. The glucose sensor configuration and mechanism we described could be very useful in the design of other carbohydrate sensors. Moreover, the glucose sensing mechanism and the computational framework we developed can be applied for detection of the prostate cancer biomarker, which is known as osteopontin (OPN). Experimentalists have demonstrated that nanotube devices show detectable sensitivity to OPN when covalently functionalized with a receptor called scFv. One possible explanation is that the charged sites on the scFv's surface are neutralized by opposite charges associated with the bound OPN, leading to a change of electrostatic potential at the nanotube surface. However, the precise mechanism for the observed sensing response remains to be determined. The computational approach established here could provide a quantitative understanding of the sensing mechanism.

It is a great challenge in dielectric polymers to achieve a high dielectric constant while maintaining low dielectric loss and high operating temperatures. We show that by blending two dipolar polymers in glassy state, i.e., poly(arylene ether urea) (PEEU, $K = 4.7$) and an aromatic polythiourea (ArPTU, $K = 4.4$) to form a nanomixture; the resulting blend exhibits a very high dielectric constant, $K = 7.5$, while maintaining low dielectric loss ($< 1\%$). The experimental and computer simulation results demonstrate that blending these dissimilar dipolar polymers causes a slight increase in the interchain spacing of the blend in its glassy state, thus reducing the barriers for reorientation of dipoles in the polymer chains along the applied electric field and generating a much higher dielectric response than the neat polymers.

We have established how the bottom-up synthesis of a graphene nanoribbon can be controlled by charge injections from an STM tip. From our experiments and first-principles calculations, it was found the hole injections from an STM tip can trigger a cooperative domino-like cyclodehydrogenation even when the polymers are quasi-freestanding with suppressed substrate effect. The hole injections greatly reduce the energy barrier in the key step of the carbon-to-carbon bond formation. The hydrogen atoms migrate to the edge and dissociate into the vacuum as H_2 molecules. As the STM tip treatment can be performed at selective molecular sites without involving a traditional catalytic effect from the metal substrate, the results point to a new way for bottom-up and controllable synthesis of freestanding GNRs and heterojunctions, which is critical for practical GNR-based nanodevices.

WHY BLUE WATERS

The applications described above require a very large parallel supercomputer with a high-speed interconnect among the nodes (due to the frequent exchange of substantial amounts of data among nodes). Each project required many runs to explore the various scientific issues, with a substantial amount of analysis




Figure 2: Optimized structures for glucose detection in the semiconducting CNT(8,0)-based systems (a, b, c, d) and the metallic CNT(5,5)-based systems (e, f, g, h). The closest distances between the pyrene-1-boronic acid molecule and the nanotube surface (in Å) are labeled in green.

between the runs. High availability and quick turnaround are thus also very important for timely progress in our research.

PUBLICATIONS AND DATA SETS

Li, Y., M. Hodak, W. Lu, and J. Bernholc, Selective sensing of ethylene and glucose using carbon-nanotube-based sensors: An *ab initio* investigation. *Nanoscale*, 9 (2017), pp. 1687–1698.

Ma, C., et al., Controllable conversion of quasi-freestanding polymer chains to graphene nanoribbons. *Nature Communications*, 8 (2017), p. 14815.

Thakur, Y., et al., generating high dielectric constant blends from lower dielectric constant dipolar polymers using nanostructure engineering. *Nano Energy*, 32 (2017), pp. 73–79.

Dolgova, N.V., Binding of Copper and Cisplatin to Atax1 Is Mediated by Glutathione through the Formation of Metal–Sulfur Clusters. *Biochemistry*, 56 (2017), pp. 3129–3141.

REDUCING JET AIRCRAFT NOISE

Allocation: Blue Waters Professor/100 Knh
PI: Daniel Bodony¹
Collaborator: Mohammad Mehrabadi¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Developing quieter, more fuel-efficient jet aircraft has been a goal since the 1950s. The movement of engine designs from the original turbojet to today's high-to-ultra-high-bypass-ratio turbofans has sustained noise and fuel-use reduction goals. However, future growth in aircraft utilization and more stringent requirements on acceptable noise suggest further demand on noise reduction. This requires detailed understanding of the turbulent flow in the fan stage of the engine that gives rise to the dominant broadband component of the fan noise, which cannot be provided by current reduced order models. Therefore, it is essential to utilize high-fidelity numerical tools to predict the turbulence and its radiated noise.

RESEARCH CHALLENGE

The dominant noise source of jet engines has changed as engines have evolved. Large bypass-ratio fans (Fig. 1, upper-left inset image) produce noise that predominantly originates from the fan itself, rather than from the exhaust streams from the core and fan. Fan noise is mostly due to the interaction of small unsteady flow perturbations with the fan (rotor) and stationary guide vane (stator) blade rows. Depending on the nature of these sources, they are categorized as tonal noise and broadband noise.

Broadband noise contains a wide range of frequencies and is associated with the interaction of a turbulent flow with a solid boundary (Fig. 1, main image). The prediction of the broadband component of fan noise (Fig. 1, lower-right inset image) is more

complicated than tonal noise predictions since it arises from an inhomogeneous and anisotropic turbulence interacting with fan blades. Although the methods developed to predict rotor-stator interaction tonal noise can, in practice, also be applied to the broadband component prediction, representing a collection of random incident disturbances in the time domain or a wide range of frequencies in the frequency domain—the ensemble of which represents a turbulent flow—requires a highly resolved computational domain with a significant computational cost.

METHODS & CODES

The turbulence and its radiated sound are governed by the compressible Navier-Stokes equations that are solved using a high-order, finite difference code written by the PI, capable of describing complex geometries using multiple, overset meshes. From fan airfoil geometry, including small roughness elements near the leading edge (Fig. 1, main image), a computational grid is designed which sufficiently resolves the flow around the blade as well as the radiated field far away from it (Fig. 1, lower-right inset image). The resulting mesh contains 2.5 billion degrees of freedom, which are saved in HDF5 format on the file system using a new, scalable, collective MPI-IO-based infrastructure (Fig. 2).

RESULTS & IMPACT

The direct numerical simulation of sound generated by an idealized blade shown in Fig. 1 requires 2.5 billion degrees of freedom and generates a database on the order of 1 terabyte. While useful for generating reduced-order models for broadband sound generation from isolated blades, the simulation domain shown in the main image of Fig. 1 is only a fraction of that needed to predict the noise from the entire fan shown in the upper-left inset image of Fig. 1. Thus, the main results of the simulation are being used to develop a wall-modeled large-eddy simulation (WMLES) approach to broadband noise prediction wherein the inner boundary layer dynamics (below the log layer) are modeled using a less expensive model while the outer boundary layer dynamics and external flow field are modeled using traditional LES techniques. Determining the accuracy of the WMLES simulation for on-blade and acoustic field predictions is the primary impact of the simulations. If successful, using WMLES for broadband noise prediction enables the aeroacoustic community to predict and, ultimately, reduce the environmental impact from most aircraft noise sources.

WHY BLUE WATERS

Blue Waters is critical to the research because of the system's scale, the performance of its I/O system, and the capability its software teams have in improving code performance, including I/O performance. These runs would not be feasible on any XSEDE computer.




Figure 1: (Upper-left inset image) NASA/GE Source Diagnostic Test Fan whose geometry and flow measurements provide the validation dataset. (Main image) Iso-surface plot of Q-criterion for a transitional boundary layer over a NACA0012 airfoil. (Lower-right inset image) turbulence-generated noise from the NACA0012 airfoil visualized using the divergence of the velocity field.




Figure 2: Blue Waters I/O performance using a new collective, MPI-I/O infrastructure.

PROPERTIES OF DENSE HYDROGEN

Allocation: Blue Waters Professor/100 Knh

PI: David Ceperley¹

Collaborators: Carlo Pierleoni², Marcus Holzmann³

¹University of Illinois at Urbana-Champaign

²University of L'Aquila, Italy

³CNRS, Grenoble, France

EXECUTIVE SUMMARY

The phase diagram of high-pressure hydrogen is of great interest for fundamental research, planetary physics, and energy applications. A first-order phase transition in the fluid phase between a molecular insulating fluid and a monoatomic metallic fluid has been predicted. The existence and precise location of the transition line is relevant for planetary models. Recent experiments reported contrasting results about the location of the transition. Theoretical results based on density functional theory are also very scattered. We performed highly accurate coupled electron-ion Monte Carlo calculations of this transition, finding results that lie between the two experimental results. We have also explored the transition between molecular and atomic hydrogen at lower temperatures, in the solid phase.

RESEARCH CHALLENGE

The properties of hydrogen under extreme conditions of pressure and temperature are of great interest for fundamental research, planetary physics, and energy applications [1]. Hydrogen accounts for much of the visible mass in the universe. The properties of hydrogen and helium are needed to understand 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 motions are quantum and correlated.

It has long been an open question [1] how hydrogen makes a transition from a molecular insulating state to an atomic metallic state as pressure and temperature are increased. A first-order phase transition in the fluid phase between a molecular insulating fluid and a monoatomic metallic fluid has been predicted [2].




Figure 1: Typical snapshots of molecular hydrogen approaching the transition to atomic hydrogen [7]. Shown are the positions of the protons: in red from Path Integral Molecular Dynamics within density functional theory energetics, and in blue from coupled electron ion simulations. r_s is the average inter-proton distance in units of Bohr radii.

Recent experiments [3-4] reported contrasting results about the location of the transition, different by a factor of two in pressure. Theoretical results based on density functional theory are also very scattered and hence not predictive [5]. These findings motivated us to repeat our earlier calculations on Blue Waters, to better control the convergence and utilize recent improvements in methodology. In addition, recent experiments have probed the solid phase leading up to the atomic transition [7] and the melting temperature.

METHODS & CODES

Over the past decade, we have developed new quantum Monte Carlo simulation methods to treat quantum systems at low temperature. The quantum Monte Carlo method we use (Coupled Electron Ion Monte Carlo) works with the full interaction between the electrons and protons and treats both particles fully quantum mechanically. In contrast to density functional calculations, all effects of electronic correlation are explicitly included. This is particularly important in hydrogen, because of possible self-interaction effects, difficulty in treating the hydrogen bond breaking, and the large van der Waals interactions. We model hydrogen with about 100 electrons and protons in a periodic cell. Special methods are used to extrapolate accurately to the thermodynamic limit. With our method, we simulated hydrogen for temperatures in the range of 200°K up to 5,000°K, and at relevant pressures, 100GPa (gigapascals) to 500GPa.

RESULTS & IMPACT

For temperatures below 2,000°K we observe [5] a first-order transition between an insulating molecular liquid and a denser metallic atomic liquid. Our predicted transition pressures are intermediate between the two experimental observations [3,4].

In recent unpublished work [7], we performed a detailed study of solid molecular hydrogen at temperatures between 200°K and 414°K and pressures between 180GPa and 520GPa, to understand the detailed crystal structures, the melting temperature of the solid and the transition into the metallic/atomic phase. There has been a recent controversial report of the experimental observation of this transition [6]. Shown in the figure are snapshots of a layer of one of the crystal structures (Cmca12) as it evolves from 180GPa to 420GPa.

Our calculations are needed to validate both our computational method and to resolve the different experimental measurements. 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. As a byproduct of our simulation, we are also judging the accuracy of various density functionals and force-fields for the simulations of these systems.

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.

PUBLICATIONS AND DATA SETS

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

SCALABLE NANOPATTERNING OF GRAPHENE BY HYDROGEN-PLASMA ETCHING

Allocation: Illinois/280 Knh
PI: Huck Beng Chew¹

¹University of Illinois Urbana Champaign

EXECUTIVE SUMMARY

Scalable and precise nanopatterning of graphene is an essential step for graphene-based device fabrication. Hydrogen-plasma reactions have been shown to narrow graphene from the edges, or to selectively produce circular and hexagonal holes in the basal plane of graphene, but the underlying plasma-graphene chemistry is unknown. The petascale Blue Waters supercomputing resources have enabled us to quantify the mechanisms of hydrogen-plasma etching of graphene supported on SiO₂ substrate across the range of plasma ion energies. Specifically, our molecular dynamics simulation results, based on a reactive force-field potential, have uncovered distinct etching mechanisms, operative within narrow ion energy windows, which explain the differing plasma-graphene reactions observed experimentally. These simulation results have provided rich insights into the complex plasma-graphene chemistry, opening up a means for controlled patterning of graphene.

RESEARCH CHALLENGE

The technique of plasma surface nanopatterning of materials has played an important role in the microfabrication over the past several decades of semiconductors of ever-reducing dimensions. Silicon, which has so far been the primary material for manufacturing of transistors, has reached its technological limitation with the recent 7-nm scale transistors. Use of novel materials such as graphene has been proposed in order to achieve the next milestone of the 5-nm ITRS node (International Technology Roadmap for Semiconductors). However, graphene has to be patterned to achieve its full potential in the domain of electronic applications. Several methods such as fluorination, boron nitride doping, irradiation by heavy ions, electron beam lithography and etching by hydrogen plasma have been used in the past for this purpose. Among these techniques, hydrogen plasma etching has shown the most promise in terms of scalability and cost effectiveness. While there exists ample experimental evidence for the patterning of graphene by hydrogen-plasma treatment, the reported etching reactions and the resulting graphene nanostructures have been vastly different. These reactions range from selective edge etching with no damage to its basal plane, to combined basal plane and edge etching of graphene resulting in isotropic and anisotropic hole growth in the basal plane of graphene. To date, controlling the patterned

graphene nanostructures by hydrogen-plasma treatment has not been achieved due to a lack of fundamental understanding of the complex hydrogen plasma-graphene chemistry. Furthermore, the complete parameter space of substrate temperature, ion energy, and incident flux has not been systematically studied due to the cost limitations of plasma experiments.

METHODS AND CODES

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

RESULTS & IMPACT

We quantify the extent of basal plane damage for each ion energy by defining a damage parameter D_b as the fraction of broken C–C bonds in the graphene sheet. The basal plane of graphene remains nearly undamaged at ion energies of 1 eV and 25 eV, but displays a nonmonotonous relationship with ion energies in between. Peak etching is observed at ion energy of 10 eV, as shown in Fig. 1a. The basal-plane etching initiates with the chemisorption of the H atoms on two neighboring C atoms and proceeds with the further hydrogenation of the dangling C–C bonds, as shown in Fig. 1b. To quantify the edge-etching of graphene, we define an edge-etching rate D_e and measure its variation with ion energy and for both zigzag and armchair orientations, as shown in Fig. 1c. Our results demonstrate distinct ion energy regimes for isotropic versus anisotropic etching. Between 1 and 5 eV, the average armchair and zigzag etching rates are similar with overlapping error bars. At ion energies of between 7 and 30 eV, however, armchair etching consistently proceeds faster than zigzag etching by ~15%, suggesting that the patterned holes should indeed be hexagonal with zigzag-oriented edge structures, in agreement with experimental observations. Fig. 1d shows the mechanisms of etching at the zigzag and armchair edge. H-induced etching as well as direct impact damage plays a role in zigzag edge, while only the latter is active in the more stable armchair edge. These distinctive etching mechanisms, which are operative within narrow ion energy regimes, fully explain the differing plasma-graphene reactions observed experimentally.

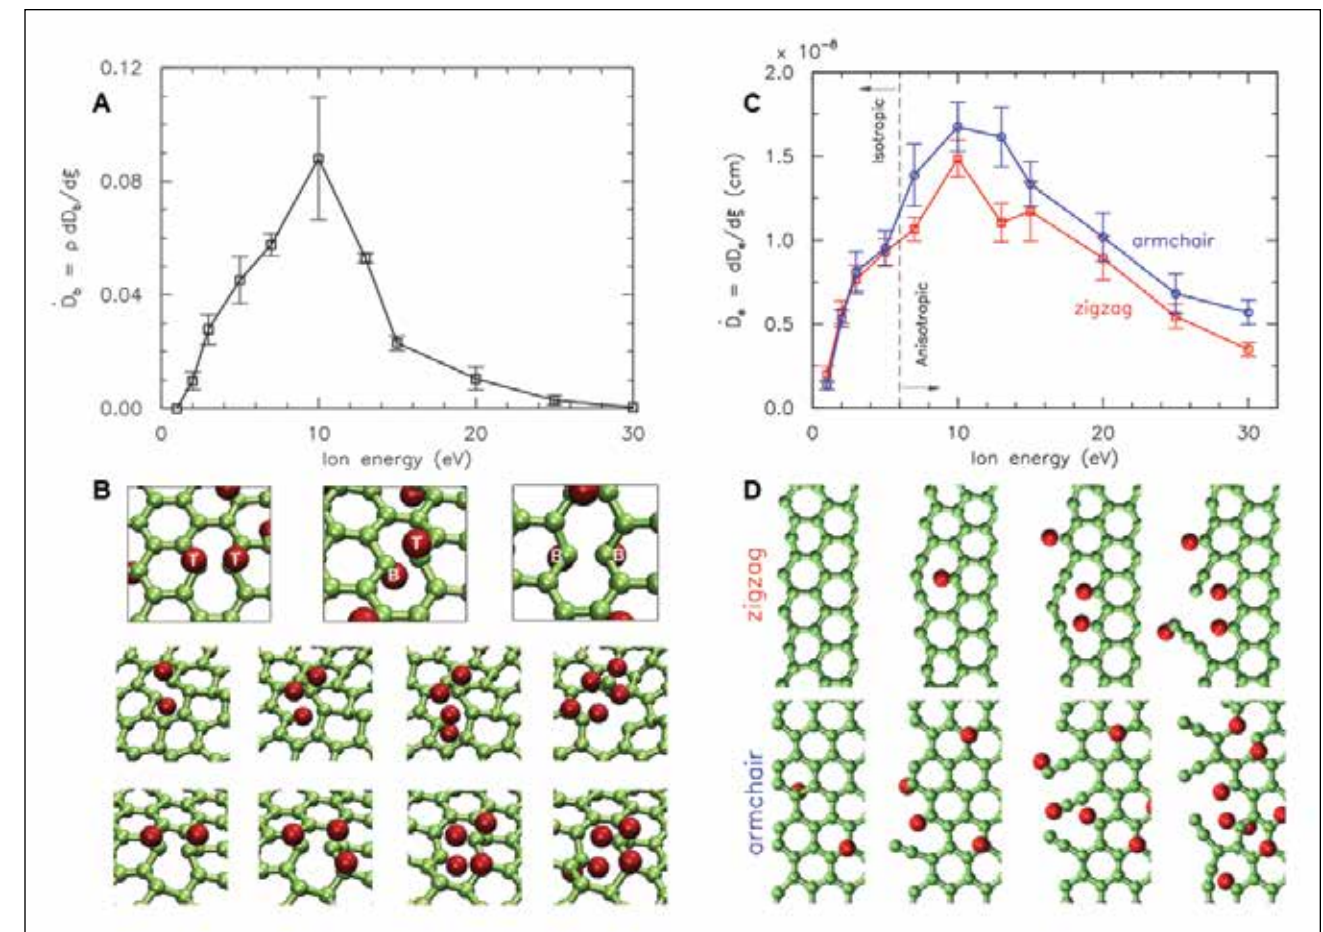


Figure 1: a) Steady state basal plane etching rate versus ion energy for monolayer graphene; b) Mechanism of etching showing the three possible configurations of damage nucleation and progression; c) Etching rates of the zigzag and armchair configurations versus ion energy; and d) Mechanism of etching for zigzag and armchair edges.

WHY BLUE WATERS

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

PUBLICATIONS AND DATA SETS

Harpale, A., and H. B. Chew, Hydrogen-plasma patterning of multilayer graphene: Mechanisms and modeling. *Carbon*, 117 (2017), DOI: 10.1016/j.carbon.2017.02.062

Harpale, A., M. Panesi, and H. B. Chew, Plasma-graphene interaction and its effects on nanoscale patterning. *Phys. Rev. B*, 93 (2016), DOI: 10.1103/PhysRevB.93.035416

hPIC: A SCALABLE ELECTROSTATIC PARTICLE-IN-CELL FOR PLASMA-MATERIAL INTERACTIONS

Allocation: Exploratory/50 Knh

PI: Davide Curreli¹

Co-PI: Rinat Khaziev¹

Collaborators: Cameron Dart¹, Alyssa Hayes¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

The hPIC code is a kinetic-kinetic electrostatic particle-in-cell (PIC) application targeted at large-scale simulations of plasma-material interactions (PMI). hPIC models the plasma sheath and presheath in strongly magnetized conditions, as normally encountered at the wall of magnetic fusion reactors. The Poisson solver is based on the PETSc conjugate gradient with BoomerAMG algebraic multigrid preconditioners. Scaling tests on the Blue Waters supercomputer have demonstrated excellent weak scaling and considerable strong scaling.

RESEARCH CHALLENGE

In the edge region of magnetically confined plasmas, plasma-material interactions can compromise the thermos-mechanical integrity of the wall, currently limiting the development of nuclear fusion reactors. When exposed to a fusion plasma environment, plasma-facing materials exhibit evidence of surface morphology

modifications, nanostructuring, and neutron damage. The boundary layer between the edge plasma and the material surface is called the plasma sheath, a region where the plasma ions accelerate up to supersonic conditions in order to maintain the quasi-neutrality of the plasma. In the strong magnetic field of a fusion machine (of the order of several Tesla), the plasma sheath is a complex three-dimensional multi-layer structure, where substantial electrostatic forces accelerate the ions toward the wall, causing damage. The hPIC code is a large-scale particle-in-cell under development at the University of Illinois and is able to perform detailed analyses of the plasma sheath in the strongly-magnetized conditions encountered in a fusion reactor. We report progress on the scaling tests (weak and strong scaling) and code optimization performed on Blue Waters. All scaling plots have been reported as a function of the number of Blue Waters' XE6 nodes; one Blue Waters node corresponds to 32 cores.

METHODS & CODES

As an electrostatic particle-in-cell, hPIC requires the solution of a large-scale Poisson problem at each time step. A well-known problem of elliptic solvers is the intrinsic nonlocality of the differential operator, hindering an efficient and scalable implementation on a large-scale computer. hPIC uses a finite-difference discretization of the Laplace operator solved with conjugate gradient and an algebraic multi-grid (BoomerAMG) as a preconditioner. The PETSc library was used as an interface for the parallel implementation. The PETSc numerical library manages the field functions of the code. The interface between hPIC and PETSc also ingests the charge density inputs required by the Poisson solver and returns electric field and electric potential.

RESULTS & IMPACT

The largest weak scaling test performed was a Poisson problem on 2,048 Blue Waters nodes (65,536 cores). A 500×500 grid ran on each core, resulting in 16.384 billion degrees of freedom. Fig. 1 shows the efficiency (w.r.t. single node) of the full PIC cycle, including both particle and field functions, for different numbers of particles-per-cell (ppc). A weak scaling efficiency greater than 78% and between 85% and 88% for all conditions relevant for practical applications has been attained.

The tests have shown that as a rule of thumb for good performance, the number of particles per cell should be kept larger

than 100, with preference to 500–600 particles-per-cell. With a small local grid of 50×50 on each MPI process, hPIC is capable of achieving up to 3,000 particles per cell, enough to achieve a satisfactory level of statistical noise. Currently, hPIC uses double precision for all physical quantities on the XE6 nodes of Blue Waters containing 64 gigabytes of random access memory per node. Thus, each node stores approximately 250 million particles.

Similarly, strong scaling tests were performed on 4 to 8,192 nodes. The problem was a Poisson equation with a spatial grid of $N \times N$ points and a corresponding Laplacian matrix of size $(NN \times NN)^2 = NN^4$. At the largest scale of 8,192 nodes, the time to solution for the two allocatable cases of $NN=16,000$ and $NN=32,000$ was 1.6 seconds. The greatest problem of $NN=100,000$ was solved on 256, 512, and 2,048 Blue Water nodes. This solution required PETSc's index to switch from a 4 byte signed integer to an 8 byte signed long long. The problem of size $NN=100,000$ has a grid size of $NN^2=10^{10}$ nodes, and a corresponding Laplacian matrix of $NN^4=10^{20}$ degrees of freedom with matrix indices running from 0 to $10^{10}-1$, which can be allocated only by using 64-bit unsigned long integers. Such a problem was solved in 13.4, 6.8, and 2.1 seconds respectively on 256, 512, and 2,048 Blue Waters nodes.

WHY BLUE WATERS

Access to Blue Waters has accelerated the development of hPIC, allowing for testing and optimization of a number of massively parallel features of the code. The code is now at the level of meeting the large computational demands of fully kinetic simulations of near-wall plasmas encountered at the edge of fusion reactors, simulations feasible only on a machine the size of Blue Waters. Resolving the plasma-material interface at scales from nano to millimeters and from picoseconds to the time scales of plasma-sheath equilibrium requires computing at the petascale and beyond. Blue Waters offers a unique opportunity to push this research forward.




Figure 1: Weak scaling efficiency of the electrostatic version of the full cycle of hPIC for a hydrogen plasma with ion and electron temperature equal to 1eV and electron density equal to 10^{17} m^{-3} ; each rank stores 100×100 grid points, with a number of particles per cell (ppc) as indicated in the figure.

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

Allocation: NSF PRAC/3,000 Knh

PI: J. P. Draayer¹

Co-PIs: K. D. Launey¹, T. Dytrych^{1,2}, J. P. Vary³, U. Catalyurek⁴, M. Sosonkina⁵

Collaborators: D. Langr⁶, T. Oberhuber⁶, W. M. Tang⁷, B. Wang⁷, R. Baker¹, G. Sargsyan¹, P. Maris³

¹Louisiana State University

²Czech Academy of Sciences

³Iowa State University

⁴The Ohio State University

⁵Old Dominion University

⁶Czech Technical University

⁷Princeton University

EXECUTIVE SUMMARY

The Blue Waters system enables us to achieve large-scale modeling of light- and medium-mass nuclei, including short-lived nuclei not yet accessible to experiment but key to understanding astrophysical processes, which are the focus of current and next-generation rare isotope experimental facilities. The scale of computational challenges inherent in the modeling of such intricate quantum many-body systems makes the Blue Waters resources essential for addressing long-lasting questions of importance to nuclear theory and experiment, as well as to astrophysics. Breakthrough theoretical advances [1,2] coupled with the Blue Waters cutting-edge computational power have opened a new region, the intermediate-mass nuclei from Fluorine to Calcium isotopes, for first investigations with *ab initio* (i.e., “from first principles”) methods. This targets nuclei far from stability, while pinpointing key features of astrophysical processes, probing fundamental symmetries in nature, as well as supporting current and upcoming experiments at radioactive beam facilities.

RESEARCH CHALLENGE

One of the quintessential open problems in contemporary physics is to design a comprehensive many-body theory for modeling and predicting nuclear structure and reactions starting from internucleon forces that are consistent with the underlying theory of Quantum Chromodynamics (QCD). The ultimate goal of *ab initio* theory is to find a solution to this problem, which is a computationally highly intensive endeavor due to dual challenges: namely, the nonperturbative nature of QCD in the low-energy

regime and the complexity of many-particle nuclei. As short-lived nuclei, currently inaccessible to experiment, are often found key to understanding processes in extreme environments from stellar explosions to the interior of nuclear reactors, first-principle nuclear models that hold predictive capabilities have been and will have a tremendous impact on advancing our knowledge at the frontiers of multiple branches of physics such as astrophysics, neutrino physics, and applied physics.

METHODS & CODES

We have developed an innovative *ab initio* nuclear structure approach, dubbed the symmetry-adapted no-core shell model (SA-NCSM) [1], with concomitant computer code “LSU3shell” [3–4], that embraces the first-principles concept and capitalizes on a new symmetry of the nucleus. The *ab initio* SA-NCSM solves the time-independent Schrödinger equation as a Hamiltonian matrix eigenvalue problem. The main computational task is to evaluate a large symmetric Hamiltonian matrix and to obtain the lowest-lying eigenvectors that correspond to the experimental regime. Accuracy is based on the degree of convergence, which is linked to the size of the model space that can be achieved. The SA-NCSM utilizes physically relevant model space of significantly reduced dimensionality compared to ultra-large model spaces encountered by standard *ab initio* approaches. These theoretical advances, coupled with the computational power of the Blue Waters system, allow us to reach medium-mass nuclei that are inaccessible experimentally and to other *ab initio* methods [2,5] (see Fig. 1).




Figure 1: Nuclear model space: (a) Explosive growth with increasing particle number and the space (“box”) in which particles reside (the largest dimension currently attainable is shown by the red horizontal line) and (b) tamed dimensions in the SA framework using symmetries known to dominate the dynamics.

RESULTS & IMPACT

The nuclei of interest represent a considerable challenge requiring computational power of nearly the entire Blue Waters (BW) machine and its system memory. Two graduate students have carried forward these studies and had the unique opportunity to work with supercomputers and massively parallel programming environments. The following list describes the results and their impact:

- We have provided the first *ab initio* description of the open-shell ^{19}Ne , ^{20}Ne , ^{24}Ne and ^{24}Si nuclei [2]. Such nuclei in the intermediate-mass region are key to further understanding the production of heavy elements, and especially X-ray burst nucleosynthesis. Of special interest are short-lived nuclei (such as ^{24}Si) that are difficult or impossible to study experimentally (see Fig. 2).
- We have studied emergent phenomena from first principle in Mg isotopes and their mirror nuclei (^{21}Mg and ^{21}F , ^{22}Mg and ^{22}Ne , with work in progress on ^{23}Mg and ^{24}Mg). While enhanced deformation and cluster substructures are difficult to describe from first principles, the BW system has allowed us to achieve *ab initio* descriptions using chiral internucleon interactions [5]. This is important for providing accurate predictions for deformed and, in the future, heavy nuclei of interest to understanding the r-process nucleosynthesis.
- Another study has focused on ^{12}C , including the most challenging Hoyle state, the resulting state of the essential stellar triple-alpha process—the study aimed at identifying important components of the internucleon interaction. A remarkable finding reveals that only a few components of the interaction can account for most of the physics, including binding energies and collectivity.
- We have performed first-principle simulations of ^{48}Ca and ^{48}Ti with the aim of studying neutrinoless double-beta decay for these heavy nuclear systems. The goal is to reduce large uncertainties in the nuclear structure matrix elements, which will, in turn, allow us to determine the neutrino type from planned experiments at the Deep Underground Neutrino Experiment.

Large investments have been made in new generations of radioactive beam facilities to enable important discoveries in nuclear science. While the above-mentioned applications focus on specific important questions, the concurrent new developments and dramatic improvements of the LSU3shell computer code, carried forward as part of the BW PAID program, may have wider impact, as multi-physics simulations in the areas of nuclear energy and national security have similar needs.

WHY BLUE WATERS

The *ab initio* nuclear structure studies represent an extremely computing-intensive endeavor. To illustrate the level of complexity, applications to medium-mass nuclei require in excess of hundreds of exabytes of memory to store the Hamiltonian matrix. The SA-NCSM drastically reduces the size of the problem and the

associated memory requirement down to hundreds of terabytes and petabytes, but this comes at the cost of a major increase in computing intensity. As a result, SA-NCSM investigations of the intermediate-mass region are beyond the scale of available academic HPC systems. Currently, only the BW system provides resources required for the *ab initio* SA-NCSM studies of medium-mass isotopes. In order to capitalize on this opportunity, we drew from the experience and expertise of the Blue Waters staff and managed to improve scalability of our code. As a result, our largest production runs efficiently utilized 717,600 concurrent threads running on 22,425 Cray XE6 compute nodes to solve the nuclear eigenvalue problem with Hamiltonian matrices that occupy up to 400 TB of memory. Clearly, the BW system represents a unique computational platform that already plays a crucial role in advancing *ab initio* nuclear theory.

PUBLICATIONS AND DATA SETS

Dreyfuss, A. C., et al., Understanding emergent collectivity and clustering in nuclei from a symmetry-based no-core shell-model perspective, *Phys. Rev. C*, 95:4 (2017), 044312, DOI: 10.1103/PhysRevC.95.044312.

Draayer, J. P., T. Dytrych, and K. D. Launey, Symmetry-adapted no-core shell model—the *ab initio* picture of nuclear collectivity. in *Emergent Phenomena in Atomic Nuclei from Large-scale Modeling: A Symmetry-guided Perspective*, in press (2017); ISBN: 978-981-3146-04-4.

Launey, K.D., et al., Symmetry-guided large-scale shell-model theory. *Prog. Part. Nucl. Phys.*, 89 (2016), p. 101, DOI: 10.1016/j.pnpnp.2016.02.001.

Dytrych, T., et al., Efficacy of the SU(3) scheme for *ab initio* large-scale calculations beyond the lightest nuclei, *Comput. Phys. Commun.*, 207 (2016), p. 202.

Draayer, J. P., et al., No-Core Symplectic Model—Exploiting Hidden Symmetry in Atomic Nuclei, accepted, *Journal of Physics: Conference Series* (2016).




Figure 2: Effect on the abundance pattern from X-ray burst (XRB) nucleosynthesis simulations (based on Hix’s Xnet) when reaction rates from the BW-enabled first-principle SA-NCSM simulations of ^{20}Ne are used (compared to current database, for fixed astrophysical conditions).

DNS OF PRESSURE FLUCTUATIONS INDUCED BY SUPERSONIC TURBULENT BOUNDARY LAYERS

Allocation: NSF PRAC/1,100 Knh

PI: Lian Duan¹

Collaborator: Meelan Choudhari²

¹Missouri University of Science and Technology

²NASA Langley Research Center, Hampton, VA 23681

EXECUTIVE SUMMARY

Pressure fluctuations are an important ingredient in wall-bounded turbulence as they are strongly correlated with turbulent vorticity dynamics and noise generation. Most existing analyses of boundary-layer-induced pressure fluctuations are based on the Poisson equation in the context of incompressible boundary layers. The pressure fluctuations induced by a supersonic turbulent boundary layer are, however, governed by the wave equation and are fundamentally more complicated than the low-speed counterpart. The objective of the research is to conduct direct numerical simulations (DNS) to advance fundamental understanding of the generic statistical and spectral features of boundary-layer-induced pressure fluctuations, including the freestream acoustic radiation at supersonic speeds and their dependence on boundary-layer parameters such as the Reynolds number. Current work with Blue Waters includes the analysis of boundary-layer data at modest Reynolds numbers, and conducting new simulations in Reynolds-number regimes difficult to reach without the large allocation provided on Blue Waters.




Figure 1: Three-dimensional representation of the spatial correlation coefficient of the pressure signal at multiple wall-normal locations for a Mach 6 cold-wall turbulent boundary layer at $Re_t \approx 450$. z_w/δ denotes the wall-normal location of the correlation origin normalized by the boundary layer thickness and three-dimensional isosurfaces are shown at 0.1 (blue) and 0.6 (green). In the horizontal planes going through the correlation origin, the contour lines shown in white range from 0.1 to 0.9.

RESEARCH CHALLENGE

Understanding the physics of the pressure fluctuations induced by turbulent boundary layers is of major theoretical and practical importance. From a practical point of view, the fluctuating pressure on aerodynamic surfaces of flight vehicles plays an important role in vibrational loading and often leads to damaging effects such as fatigue and flutter. The freestream pressure fluctuations radiated from the turbulent boundary layer on the nozzle wall are responsible for the genesis of freestream acoustic noise in supersonic wind tunnels. Therefore, the characterization of tunnel acoustic noise is critically important to experimental measurement of boundary-layer stability and transition in such wind tunnels. From a theoretical point of view, pressure is of fundamental importance to understanding the turbulent vorticity dynamics and to modeling the pressure-strain terms in the Reynolds stress closure.

One unique feature of boundary-layer-induced pressure fluctuations at supersonic speeds is that acoustic mode fluctuations emerge in the form of eddy Mach waves [1]. The pressure fluctuations thus include contributions from both vortical and acoustic modes. The characteristics of the acoustic pressure fluctuations and the relative importance of the two modes in different regions of the boundary layer are largely unknown. The current work aims to use the cutting-edge computational power of Blue Waters to provide the basis for an in-depth understanding of the global pressure field induced by supersonic turbulent boundary layers across a wide range of boundary-layer parameters. Such an understanding will advance the state-of-the-art knowledge of wall-bounded turbulence and boundary-layer-induced noise.

METHODS & CODES

DNS are conducted using HyperWENO, an in-house high-order finite-difference solver that solves the compressible Navier-Stokes equations in generalized curvilinear coordinates describing the evolution of the density, momentum, and total energy of the flow. An optimized, high-order (up to 7th order) finite-difference WENO (weighted essentially non-oscillatory) scheme [2,3] is used to compute the convective flux. The WENO scheme combines a high order of accuracy with relatively low dissipation, making it suitable for simulations of compressible turbulent flows. A 4th-order central difference scheme is used for the viscous flux terms, and a 3rd-order low-storage Runge-Kutta scheme [4]




Figure 2: Instantaneous flow visualization for a Mach 6 cold-wall turbulent boundary layer at $Re_t \approx 450$. The freestream acoustic noise is visualized using numerical schlieren and the boundary layer is colored by the vorticity magnitude. The angle θ illustrates the preferred direction of radiated acoustic wavefront.

is employed for time integration, which significantly relieves the memory requirement and is well suited for time-accurate simulations such as DNS. The turbulent inflow can be generated using either a recycling/rescaling method [5] or a digital filtering method [6]. On the wall, no-slip conditions are applied for the three velocity components, and an isothermal condition is used for the temperature. At the top and outlet boundaries, unsteady nonreflecting boundary conditions are imposed. Periodic boundary conditions are used in the spanwise direction.

RESULTS & IMPACT

We have used the computational resources of Blue Waters to investigate the characteristics of pressure fluctuations generated by nonadiabatic cold-wall turbulent boundary layers with a nominal Mach number of 6 and a Karman Reynolds number of $Re_t \approx 450$; results were published in the *Journal of Fluid Mechanics*. The DNS have overcome a number of difficulties encountered during experimental measurements of broadband pressure fluctuations and have provided access to quantities that cannot be measured easily, including multi-variate pressure statistics and large-scale structures in the pressure field (Figs. 1, 2). By comparing turbulent boundary layers with different wall-cooling rates, the study provides, for the first time, fundamental understanding of the effect of wall cooling on the global pressure field.

In the near term, we would like to perform DNS of supersonic turbulent boundary layers at a significantly higher Karman Reynolds number ($Re_t > 2,000$) to study the dependence of the global pressure field on the Reynolds number. Significant progress

has been made to re-engineer HyperWENO to overcome I/O bottlenecks and improve node-level parallelism with guidance from the Blue Waters support team. Test runs to date performed with the re-engineered DNS code have shown encouraging computation and I/O performance.

WHY BLUE WATERS

Direct numerical simulations of high-Reynolds-number turbulent boundary layers will be used to study the broadband fluctuating pressure field induced by the boundary layer, with the targeted Reynolds numbers significantly higher than the state of the art. In such simulations, extremely fine meshes are required to fully resolve all the turbulence scales in order to obtain the pressure spectra in the high-frequency/large-wave-number range. In the meantime, the simulations need large domain sizes to locate very-large-scale coherent structures in the pressure field as well as to accommodate the eddy decorrelation length and to minimize inlet transience as a result of inflow boundary conditions. A large number of time steps are also required for the study of the low-frequency behavior of the pressure spectrum. As such, the proposed computational efforts cannot be done without the world-class computing capabilities of Blue Waters.

PUBLICATIONS AND DATA SETS

Zhang, C., L. Duan, and M. M. Choudhari, Effect of wall cooling on boundary-layer-induced pressure fluctuations at Mach 6, *Journal of Fluid Mechanics*, 822 (2017), pp. 5–30.

DISPERSION OF FULLY RESOLVED LIQUID DROPLETS IN ISOTROPIC TURBULENT FLOW

Allocation: GLCPC/280 Kwh
PI: Said Elghobashi¹
Collaborator: Hengjie Wang¹

¹University of California, Irvine

EXECUTIVE SUMMARY

The objective of our research is to enhance the understanding of the four-way coupling effects in droplet-laden isotropic turbulence. (Four-way coupling means there is a two-way coupling between the droplets and turbulence, in addition to a two-way coupling between the droplets themselves.) The deformable droplets are fully resolved in 3D space and time.

The droplet-laden turbulence is simulated by directly solving the Navier–Stokes equations via a variable-density projection method. Most of the computation time is spent in solving the variable-coefficient Poisson equation. In order to improve the computational performance, we have developed an efficient Message Passing Interface (MPI)-based multigrid solver, which is further accelerated with multiple graphics processing units (GPUs).

Further, we examine the kinetic energy transfer between the droplets and turbulence. We also compare the dispersion characteristics of deformable droplets and solid particles with identical diameters and density ratios in isotropic turbulence.

RESEARCH CHALLENGE

Liquid fuel combustion devices typically atomize fuel into sprays of fine droplets. The fuel droplets disperse in the surrounding turbulent air, vaporize, and mix with it. The resulting chemical reaction transforms the chemical bonding energy into thermal energy, followed by gas expansion, which provides the desired mechanical energy. Understanding the four-way coupling effects between droplets and carrier flow is a necessary prerequisite for the efficient optimization of the energy conversion process.

The method of direct numerical simulation (DNS) of the fully resolved droplet-laden turbulence should be highly efficient in order to allow for the use of the appropriate fine mesh and timestep resolutions on Blue Waters. Our multigrid (MG) solver is essential for the efficient solution of Navier–Stokes equations.

METHODS & CODES

Our numerical procedure solves the two-phase incompressible Navier–Stokes equation. The jump condition at the interface between the carrier phase and the dispersed phase (liquid droplets) is incorporated in the Navier–Stokes equation via the Ghost Fluid Method.

The interface is accurately resolved by a conservative Level Set Method that conserves the mass of the dispersed phase. The most time-consuming part of our algorithm is the solution of

the variable-coefficient Poisson equation that is derived from the Navier–Stokes equation. In order to solve the Poisson equation efficiently, we have developed a V-cycle geometric multigrid solver that serves as a preconditioner to the Preconditioned Conjugate Gradient method. A comparison is made between the dispersion of finite-size solid particles and liquid droplets in isotropic turbulence. The dispersion of solid particles is studied using the Immersed Boundary Method under the same physical conditions as those of droplet-laden turbulence.

RESULTS & IMPACT

Using our recently developed multigrid solver with its linear strong scalability (as seen in Fig. 1), preliminary results show that the dispersion of finite-size liquid droplets in isotropic turbulence is larger than that of finite-size solid particles of the same diameter and density ratio. This is due to the reduced decay rate of turbulence kinetic energy caused by the four-way coupling between the droplets and carrier flow.

WHY BLUE WATERS

In order to accurately resolve the turbulence length- and timescales as well as the dynamical properties of the deformable droplets, the mesh size of our DNS should be $1024^3 \sim 2048^3$. Accordingly, the required number of high-performance computing cores ranges from 33×10^3 to 260×10^3 to solve the problem efficiently. Blue Waters' XE6 nodes can accommodate this requirement. No other facility available to us can provide such a large computational resource.

Our recently developed multigrid solver has been efficiently parallelized as shown in Fig. 1. It has the potential to be accelerated with GPU technology. The XK7 GPUs can serve as a good platform for us to further develop our multigrid solver. Furthermore, the Blue Waters staff have always provided us with valuable assistance in postprocessing and profiling. This assistance is essential for our project.

PUBLICATIONS AND DATA SETS

Rosso, M., H. Wang, and S. Elghobashi, Dispersion of finite size droplets and solid particles in isotropic turbulence. *ICMF-2016—9th International Conference on Multiphase Flow*, Firenze, Italy, May 22–27, 2016.




Figure 1: Scalability of our multigrid solver is indicated by the reduction of the time needed to solve a Poisson equation by increasing the number of Blue Waters processors. Our preconditioned conjugate gradient method uses a multigrid preconditioner. The mesh has 1024^3 grid points. The black line indicates the strong scalability of the solver and the red line shows a linear scalability as reference.

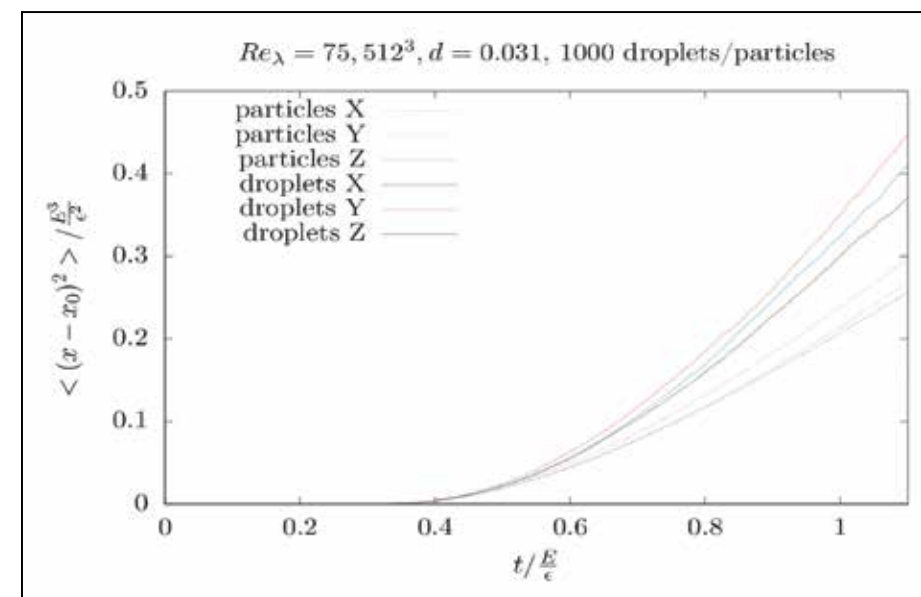


Figure 2: Comparison of the dispersion of liquid droplets and solid particles of the same diameter and density ratio in isotropic turbulence in the three coordinate directions. The dispersion is defined as the ensemble average of the square of the distance of a droplet/particle from its initial position. E is the turbulence kinetic energy. ϵ is the dissipation rate of E .

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

Allocation: Illinois/500 Knh
PI: Elif Ertekin¹
Co-PIs: Lucas Wagner¹, Ray Plante²

¹University of Illinois at Urbana-Champaign
²National Institutes of Standards and Technology

EXECUTIVE SUMMARY

Blue Waters has enabled us to carry out automated, high-throughput quantum Monte Carlo calculations of condensed matter systems, and to create a database titled QMCDB (Quantum Monte Carlo DataBase) to share these results. We have been able to simulate approximately fifteen different classical semiconductors, which represents the largest set of QMC condensed matter simulations carried out to date. The database contains formation enthalpies and optical excitation energies, calculated via the highly accurate quantum Monte Carlo technique. QMCDB will serve as a platform for easy, searchable data exchange to accelerate the knowledge base around the use of QMC for materials modeling. This work would not be possible without Blue Waters, which allows the calculation of a large class of materials ranging from classic to exotic semiconductor materials, photovoltaics, thermoelectrics, and metallic systems; we can take advantage of the near-linear scaling of our code up to several thousand nodes.

RESEARCH CHALLENGE

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

METHODS & CODES

We have developed an automated workflow that enables high-throughput simulations of solid-state materials using variational and diffusion Monte Carlo. Our open-source framework is available on our GitHub pages [1], and currently utilizes CRYSTAL

[2] to generate trial wave functions and QWalk [3] for quantum Monte Carlo simulations.

RESULTS & IMPACT

The prospect of materials design using high performance computation is one of the most exciting for future technologies. However, the properties of modern materials are rather complex, and current numerical methods can fail to describe them quantitatively. Today, QMC is a state-of-the-art suite of tools for high-accuracy ab initio modeling. It shows great promise for high-accuracy materials modeling and is already well established in the physics community, but its application to real materials with chemical identity remains fairly young. Our goal is to accelerate the development of the QMC community's collective knowledge base around the use of this method for predictive modeling of real materials. Ultimately, the impact will be to enable quantum Monte Carlo methods to emerge as a standard component of the computational materials modeling toolkit, enabling unprecedented high-accuracy simulation of complex materials, correlated systems, high-Tc superconductors, and other historically challenging materials.

Accomplishments:

- **Autogen:** The implementation of an automated workflow (Fig. 1a) for carrying out QMC calculations of solid materials on Blue Waters. The workflow was used together with Blue Waters to carry out our automated quantum Monte Carlo calculations of semiconductor materials. It is now available to the public via our GitHub pages.
- **Calculation of Thermodynamic Properties of Wide Class of Challenging Materials:** We were able to use the *Autogen* tools to carry out a large, systematic calculation of the thermodynamic properties and formation enthalpies of a large and varied class of materials (Fig. 1b). Our efforts thus far have focused on the calculation of thermodynamic properties (total energies and formation enthalpies) and the band gaps of a wide class of semiconductor materials spanning from conventional (silicon, etc.) to more exotic materials (wide band gap oxides, correlated systems, photovoltaic thin film materials).
- **The establishment of QMCDB:** Our Mongo database "QMCDB," which includes results and provenance, is described in detail

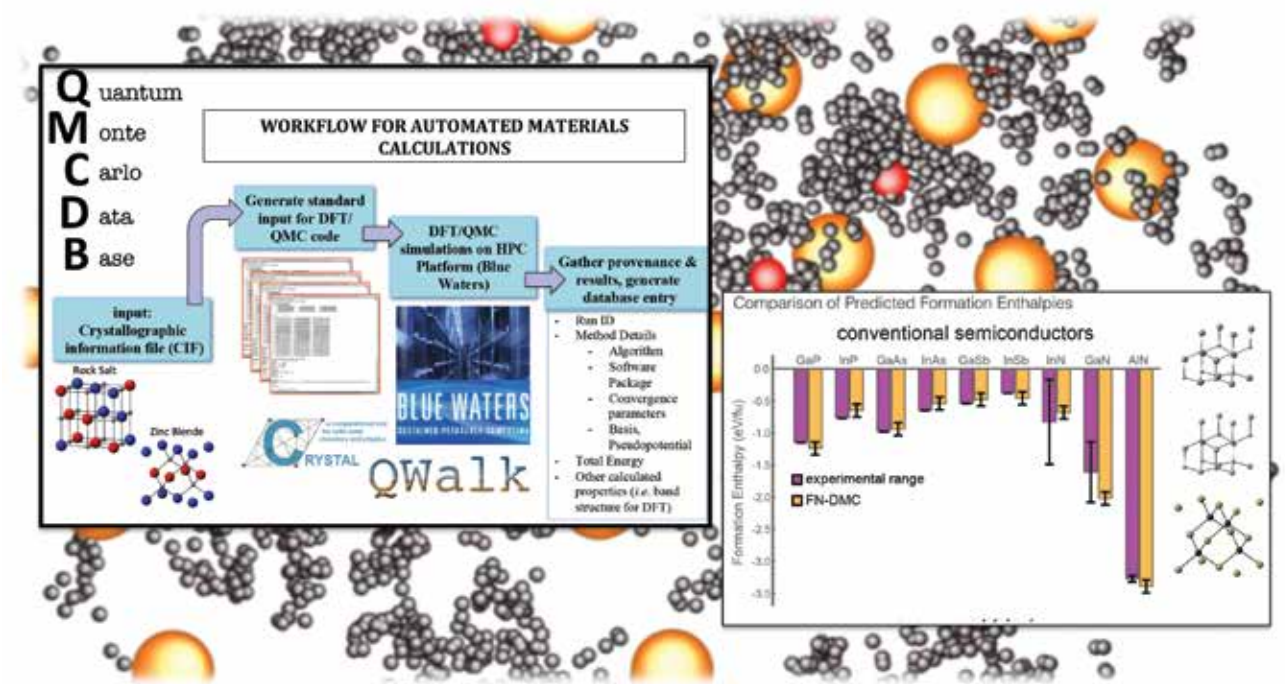


Figure 1: Background: a snapshot of a configuration of electrons around atomic nuclei in the magnesium oxide solid. Foreground: the *Autogen* workflow for automated quantum Monte Carlo calculations of solids and incorporation into quantum Monte Carlo database, and selected results for calculated formation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

on the cached National Data Service Project wiki page (<http://bit.ly/2fgoZB4>). This database will serve as a vehicle to quickly overcome the current expertise hurdle and bring the QMC methodology into the standard computational modeling toolkit. While today there are several competing materials databases for density functional theory, there is no existing QMC database. Thus, Blue Waters has provided us the opportunity to establish this tool for the worldwide QMC community. To our knowledge, this comprises the largest QMC data set that has been established to date.

WHY BLUE WATERS

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

PUBLICATIONS AND DATA SETS

Schiller, J. A., L. K. Wagner, E. Ertekin, Phase Stability and Properties of Manganese Oxide Polymorphs: Assessment and Insights from Diffusion Monte Carlo. *Phys. Rev. B*, 92:23 (2015), DOI: 10.1103/PhysRevB.92.235209.

Yu, J., L. K. Wagner, E. Ertekin, Towards a systematic assessment of errors in diffusion Monte Carlo calculations of semiconductors: case study of zinc selenide and zinc oxide. *J. Chem. Phys.*, 143:22 (2015), DOI: 10.1063/1.4937421.

Yu, J., L. K. Wagner, E. Ertekin, Fixed node diffusion Monte Carlo description of nitrogen defects in zinc oxide. *Phys. Rev. B*, 95:7 (2017), DOI: 10.1103/PhysRevB.95.075209.

Quantum Monte Carlo Database (QMCDB).

NUMERICAL METHODS AND SOFTWARE FOR COMPUTATIONAL FLUID DYNAMICS, NEK5000

Allocation: Blue Waters Professor/300 Knh

PI: Paul Fischer¹

Collaborators: Som Dutta¹, Ketan Mittal¹

¹University of Illinois at Urbana-Champaign


EXECUTIVE SUMMARY

Complex turbulent flows are one of the most important types of flows and are prevalent in both natural and built environments. We describe here simulations of two cases of turbulent flows under complex boundary conditions. The first addresses flow over biofilms and similar natural wall roughness. The second simulates sediment transport in bifurcating rivers and channels. These simulations are based on the scalable open source code Nek5000, which employs minimally dispersive spectral element discretization. Results from the biofilm simulations shed light on the interaction of the biofilm with the flow. Results from the flow at the bifurcation shed light on the dynamics of the flow near the bottom of a channel.

RESEARCH CHALLENGE

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 our research 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).

Figure 1: Instantaneous velocity magnitude at three orthogonal planes. The plane near the bottom of the channel shows the interaction of the flow with the biofilm patches, resulting in higher number of regions of low-velocity magnitude (streaks). One can also observe the effect of the biofilm on the flow, with relative increase in vortex shedding in the wall-normal direction.



The first project studies flow and fine particle transport over biofilms. Biofilms, in the form of microbial communities, serve as a key component in controlling carbon and nutrient cycling in freshwater systems. These microbial communities function as the coupling between physical and biological processes. They have a significant impact on a stream’s hydrodynamics and influence the amount and lability of carbon exported downstream. Most research efforts to date have relied on the use of experimental analysis to understand how biofilm growth affects the flow hydrodynamics and fine particle transport [1]. None of the studies, however, had the spatial and temporal resolution to unravel the interaction between the biofilm structure and the flow, and its effect on fine particle transport. Thus, we conducted Direct Numerical Simulation (DNS) of the flow over biofilms, with the structure of the biofilm provided by experiments that measured the benthic biofilm using an Optical Coherence Tomography microscope. Due to the highly irregular structure of the biofilm bathymetry, new methods for generating the computational mesh were developed.

The second project addresses 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 [2,3]. High-resolution Large Eddy Simulations were conducted for flow and sediment transport at idealized experimental scale bifurcations.

METHODS & CODES

The turbulence simulations were based on the open-source spectral element code Nek5000 [4]. 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 [5]. 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. In the simulations with sediment (or fine particle) transport, the particles have been modeled as Lagrangian point particles using a novel semi-implicit timestepping scheme developed to simulate polydisperse particles accurately [6]. For the biofilm problem, the measured biofilm bathymetry data were smoothed while keeping the general structure intact, which was then used to generate the computational mesh. For the first time, spectrally accurate DNS simulations were conducted for a channel with complex natural roughness, and this was possible due to development of sophisticated mesh smoothing algorithms for Nek5000 [7].

RESULTS & IMPACT

DNS of the flow over the biofilm, with a bulk Reynolds number of 8,000, was conducted using about 200 million computational points. Initial results show the interaction of the flow with the biofilm patches, resulting in a higher number of regions of low-velocity magnitude (streaks—see Fig. 1). One can also observe the effect of the biofilm on the flow, with a relative increase in vortex shedding in the wall-normal direction. As we are conducting the study with “real roughness,” the results will also have an impact on the study of flow over different kinds of roughness, which is important for mechanical and aerospace engineering applications.

The channel bifurcation cases have been simulated with about 250 million computational points; the Reynolds number of the flow is 25,000. One can observe that most of the flow near the bottom enters the side channel even though the total flow is equally divided between the two channels. This clearly hints at the mechanism for the sediment near the bottom to preferentially enter the side channel.

WHY BLUE WATERS

The study of flow in environmental fluid mechanics requires computational resources with sustained computing power at an unprecedented scale. Simulations have been conducted for up to 296 million computational points, with the code scaling strongly up to 32,768 MPI (message passing interface) ranks. Without access to Blue Waters, completing the study within a realistic timeframe would be impossible. In addition, since visualization of a phenomenon is an effective way to understand and explain its mechanics, we will work with Blue Waters project staff to create animations of the phenomena using data from the simulations.

PUBLICATIONS AND DATA SETS

Dutta, S., P. Fischer, M.H. Garcia, Large Eddy Simulation (LES) of flow and bedload transport at an idealized 90-degree diversion: insight into Bulle-Effect in *River Flow 2016*, CRC Press (2016).

DIRECT NUMERICAL SIMULATION OF TURBULENCE AND SEDIMENT TRANSPORT IN OSCILLATORY BOUNDARY LAYER FLOWS

Allocation: 300,000 Illinois/300 Knh
PI: Marcelo H. García¹
Co-PI: Paul Fischer¹
Collaborators: Dimitrios K. Fytanidis¹, Jose M. Mier Lopez¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Oscillatory boundary layer flows play an important role on coastal and offshore engineering and the sediment transport mechanics in coastal environments. The present work will be the first computational effort to simulate the effects of turbulent structures and bed roughness on the maximum bed shear stress phase difference with respect to the maximum free-stream velocity and the first numerical study that will examine the mixing layer and momentum exchange between the pore-scale flow and the oscillatory free-stream flow for the case of turbulent oscillatory flow over porous bed. The proposed work combines the expertise of Prof. Marcelo García's group (Ven Te Chow Hydrosystems Laboratory, Civil and Environmental Engineering) and Prof. Paul Fischer's group (Computer Science and Mechanical Science &

Engineering) with the leading-edge petascale computing resources of Blue Waters available at the University of Illinois at Urbana-Champaign and aims to become one of the most comprehensive studies on the effect of turbulent structures on the oscillatory boundary layer flows and sediment transport.

RESEARCH CHALLENGE

The growing human needs as well as global economic development have resulted in a rapid increase in marine activities. Coastal areas are usually involved in these activities in different ways, from hosting the foundation of offshore structures and breakwaters to accepting the residue of these activities in form of pollution or disturbance of the ecology and sediment transport. Numerical and theoretical models are being used by scientists, engineers, and decision-makers to design infrastructures and

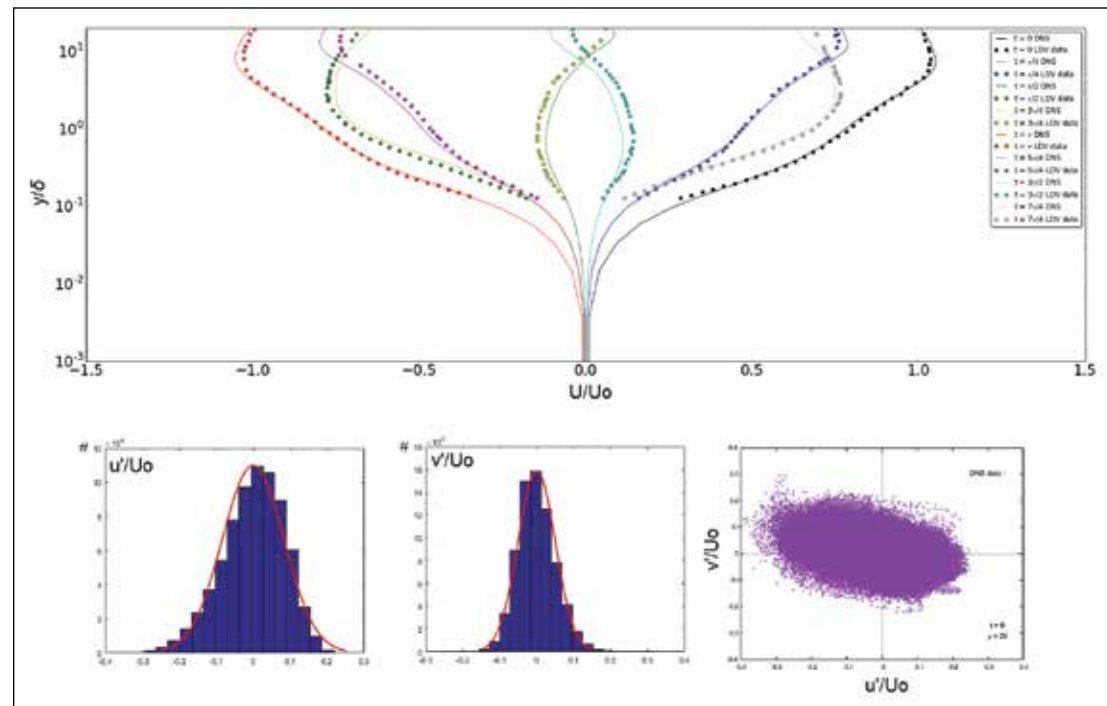


Figure 1A, B: Comparison between numerical results and experimental observations. Distribution and quadrant analysis of velocity fluctuations.

study the effect of marine activities on coastal environments and the sediment transport processes. However, most of the current state-of-the-art models fail to accurately predict the interaction of oscillatory flow with sediment transport, highlighting the existing knowledge gaps regarding the complex interactions between the oceanic flow, the coastal bottom, and sedimentation processes.

METHODS & CODES

Using advanced experimental techniques, extensive experiments have been conducted in the Large Oscillatory Water-Sediment Tunnel at the Ven Te Chow Hydrosystems Laboratory. These experiments suggest the presence of a phase-lag between the maximum bed shear stress and the maximum free-stream velocity in the case of a smooth flat bed [1]. This observation is important for the field of environmental fluid mechanics and coastal sediment transport, as this study is the first one in the literature that supports the hypothesis that the maximum shear stress is lagging instead of leading the maximum free-stream velocity over the period of each oscillation. Nevertheless, due to the limitation of the applied pointwise experimental technique used in the experiments (Laser Doppler Velocimetry), it was not possible to explicitly associate the presence of the bed shear phase lag with the development of the three-dimensional turbulence structures of the oscillatory wave flow, usually referred to as turbulence coherent structures. Thus, turbulence in complex three-dimensional flows is done mainly numerically using a combination of high-order numerical methods with backing from accurate experimental observations from the lab.

In the proposed work, we developed a Direct Numerical Simulation model capable of simulating the complex oscillatory boundary layer flow and sediment transport using the Spectral Element Method framework provided by the highly scalable open-source code *Nek5000* [2]. Except for the analysis of turbulence characteristics of OBL flow over different bed conditions representative of the coastal bottom, the present work requires use of a proper model for the simulation of the suspended sediment using an Eulerian approach and proper boundary conditions for the sediment mass exchange between the coastal bed and the free-stream flow (e.g., [3-4].)

RESULTS & IMPACT

Preliminary results show good agreement between the experimental observation [1] and numerical results for the case of the transitional/intermittent turbulent regime (Fig. 1A) [5]. Quadrant analysis of velocity fluctuations shows zones of sweeps and ejections (Fig. 1B) that dominate during the whole cycle of the oscillation. Also, velocity fluctuations seem to follow a normal distribution. These turbulence-related events (e.g., turbulent bursting and associated ejections and sweeps) are known as mechanisms that interact with sediment particles [9].

Turbulence coherent structures (Fig. 2) are also examined for the case of oscillatory flow, confirming the presence of turbulence spots in the transitional regime of oscillatory boundary layer flows




Figure 2: Turbulence flow coherent structures defined using the Lambda 2 criterion.

[6]. These lambda-shaped flow features are associated with local flow “bursts,” causing local spikes in the local wall shear stress values [6]. These local flow structures are later developed, resulting in the onset of a more intense turbulent flow field close to the wall. The above are important findings for the field of sediment transport in coastal environments as bed shear stress is considered to be the main mechanism that is the major factor in sediment entrainment into suspension (e.g., [4]).

Our research will lead to a deeper understanding of the interaction of oscillatory turbulent flow, bed shear stress, and sediment mass transport and eventually the development of new, simplified but accurate models for the analysis and design of engineering systems in coastal and oceanic environments.

WHY BLUE WATERS

The present work pushes the limit of the turbulent-resolving flow modeling of oscillatory flows. The dimensions of the computational domain are chosen based on prior knowledge of experimental observation of “turbulent spots” [6–7] to ensure that the computational domain is big enough to allow these turbulent structures to develop. A sensitivity analysis has been conducted within the first quarter of the proposed project to ensure that the computational domain size is adequate for the formation of the coherent turbulent structures. The size of the used domain, which is larger than the previous studies’ domains reported in the literature, together with the increased number of computational points (order 0.8 billion points), make the present work the first of its kind in terms of its turbulence/scale resolving results as well as its computational cost. The above, together with the requirement for the solution of several flow cycles to guarantee period-independent results, lead to increased computational requirements, making the use of a leading-edge petascale high-performance computing system like Blue Waters necessary for the success of the present work.

3D PROBABILISTIC PHYSICS-BASED SEISMIC HAZARD MAPS FOR REGIONAL RISK ANALYSIS

Allocation: Exploratory/50 Knh

PI: Paolo Gardoni¹

Co-PI: Roberto Guidotti¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Disastrous earthquakes have revealed the extreme vulnerability to seismic hazards of highly populated areas. This exploratory project contributed toward a reliable seismic risk evaluation, substantiating a future proposal for the development of site-specific seismic scenarios through physics-based 3D models of seismic wave propagation. To be reliable, those models require a high computational burden, with a significant level of parallelization and high-performance computing resources. That is why the leading-edge capability of Blue Waters is necessary. The exploratory project and the requested resources have been used to: 1) evaluate and tune the SPEED (SPectral Elements in Elastodynamics with Discontinuous Galerkin) software for Blue Waters, and 2) demonstrate its readiness for use in a future proposal submission. The next-generation work will be a proposal for a general allocation aimed at developing 3D probabilistic physics-based hazard maps.

RESEARCH CHALLENGE

A seismic event affecting a densely populated area could lead to severe damage and economic losses. The chain of disastrous earthquakes that have occurred recently, from Chile (Maule, Mw 8.8) to Japan (Tohoku Mw 9.0) to New Zealand (Darfield Mw 7.1, Christchurch Mw 6.2), revealed the extreme vulnerability of modern society to seismic hazards and the need to better estimate seismic scenarios. A reliable seismic risk evaluation, able to assess and reduce earthquake-induced damage and loss, is clearly needed, and is, therefore, a challenge of paramount significance. To improve hazard assessment through the definition of site-specific physics-based seismic scenarios will help emergency managers, planners, and the public to be prepared for future earthquakes. It will also help civil engineers to develop cost-effective mitigation measures and practices in structure design, construction, and planning. This will allow researchers and the risk-management industry to tackle the challenging task of analyzing seismic wave propagation with increased accuracy.

METHODS & CODES

The numerical simulations tested on Blue Waters within the allocated exploratory project have been carried on with SPEED [1], an existing open-source high-performance software package. The code belongs to the family of the spectral element method (SEM), a powerful, well-established, numerical technique naturally suited for three-dimensional seismic wave propagation analyses. SPEED allows seismic wave propagation modelling through viscoelastic heterogeneous three-dimensional media, both on the local and regional scale. SPEED reproduces the propagation path of the seismic wave through complex geological structures and localized superficial irregularities such as alluvial basins and civil engineering infrastructures. The code is written in Fortran90 using its pseudo-object-oriented features. It takes advantage of the hybrid parallel programming based on the Message Passing Interface (MPI) library and relies on the domain decomposition paradigm and the OpenMP library for multi-threading operations on shared memory.

RESULTS & IMPACT

The allocated resources have been used to: 1) set up the SPEED software on Blue Waters, 2) evaluate/tune the SPEED code with short tests (tutorials), and 3) evaluate/tune the SPEED code with tests at full machine scale.

Fig. 1 shows results of tests at full-machine scale. The spatial variability of peak ground east-west and north-south velocity has been estimated by a 3D numerical simulation of the February 22, 2011, Christchurch, New Zealand (Mw 6.2), earthquake, based on the available numerical model for that area [2]. The allocated resources in the exploratory project have been used to evaluate how the SPEED code utilizes the major systems elements of Blue Waters; e.g., the memory hierarchy, the communications network, the computational elements, the GPU nodes, and the I/O subsystem. We evaluated and tuned the SPEED code for Blue Waters and demonstrated the readiness of SPEED for use in a future proposal submission. In addition, the tests performed on Blue Waters confirmed the excellent scalability features of the code.

For our next work, we are developing a proposal for a general allocation aimed at developing 3D probabilistic physics-based hazard maps. These maps will help address seismic risk, contributing to a comprehensive understanding of earthquake physics and effects. The general proposal will consist of two main steps. First, we will develop a 3D physics-based seismic model of the Metro Memphis Statistical Area, which is near one of the highest seismic hazard areas in the central and eastern United States—the New Madrid Seismic Zone. This model is intended to give a comprehensive understanding of seismic wave propagation through complex media over a large area, leading to more accurate, physics-based and site-specific seismic hazard maps for use in assessing the resilience of spatially distributed large networks. Second, we will investigate the predictive power

of physics-based methods in seismic hazard analysis, introducing a probabilistic procedure based on the development of a suitable surrogate model. Based on the chosen 3D physics-based seismic scenarios, the surrogate model will allow the generation of a new set of scenarios with a significant reduction in computational burden.

WHY BLUE WATERS

Among the most appealing features of physics-based 3D models of seismic wave propagation is the possibility of capturing in its entirety the complex coupling of: 1) directivity pulses; 2) 3D basin effects (Fig. 1); 3) topographic effects; 4) wave scattering; and 5) nonlinear soil response, especially near the source of an earthquake. Accounting for all these features within a single computational model requires a high computational burden (in terms of CPU time and RAM usage), with a significant level of parallelization and high-performance computing resources. That is why the leading-edge capability of Blue Waters was necessary to address this research. To perform our 3D seismic wave propagation simulation would have been infeasible without Blue Waters' resources. Blue Waters, allowing the running of hundreds of simulations, is essential in a future work to generate the predictions needed for 3D probabilistic physics-based hazard maps.

PUBLICATIONS AND DATA SETS

Guidotti, R., A. Contento, and P. Gardoni, 3D Probabilistic Physics-based Seismic Hazard Assessment via Metamodels. *16th European Conference in Earthquake Engineering*, Thessaloniki, Greece, June 18–21, 2018.

Tian, S., R. Guidotti, and P. Gardoni, Simulation of seismic wave propagation in the Metro Memphis Statistical Area (MMSA), *11th U.S. National Conference on Earthquake Engineering*, Los Angeles, Calif., June 25–29, 2018.

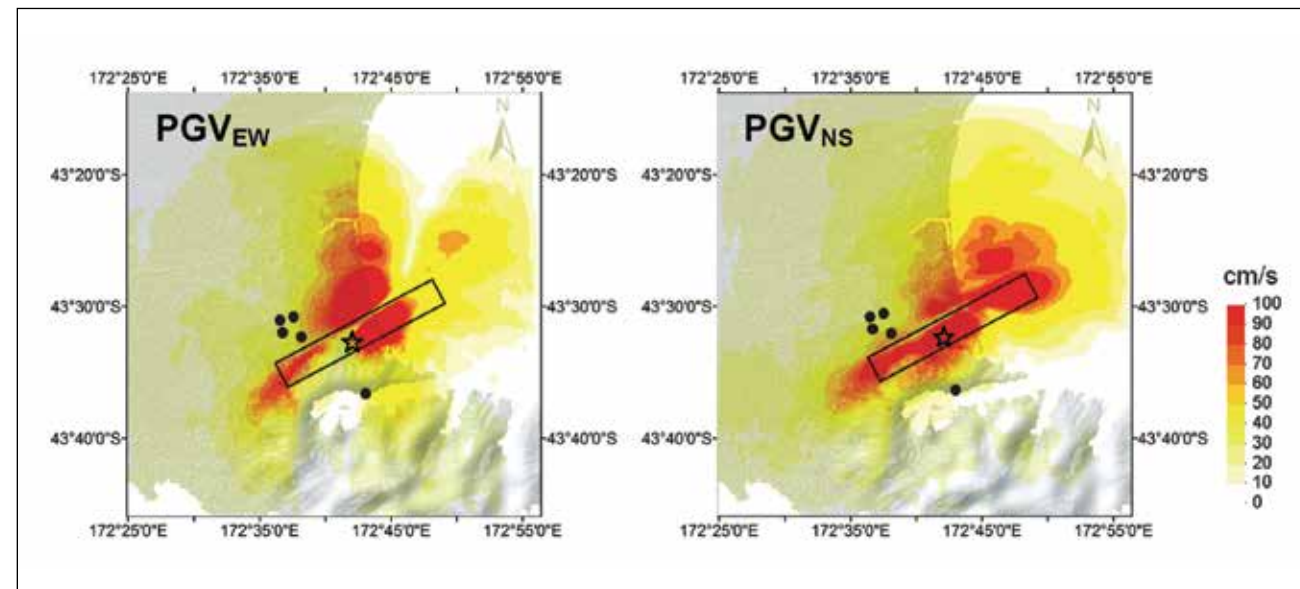


Figure 1: Spatial variability of peak ground velocity as estimated by 3D numerical simulation referring to the February 22, 2011, Christchurch, New Zealand, earthquake (Mw 6.2).

OPTIMAL BIO-LOCOMOTION STRATEGIES IN FLUIDS

Allocation: Blue Waters Professor/240 Knh
PI: Mattia Gazzola¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Is the familiar goldfish hiding a technological treasure? Seemingly simple, its wandering around a bowl involves complex interactions among its senses, its body, and the surrounding water. More generally, during millions of years of evolution animals have refined their shapes, gaits, and collective behaviors to master the complex interplay among their bodies, their senses, and the environment. How can we understand the biophysical mechanisms at play and translate them into rational design principles? To inquire into the fundamental mechanisms underlying swimming and flying we have developed an automated inverse design method based on large-scale numerical simulations and artificial intelligence techniques. This allows us to set a desired trait—for example, energy efficiency—and to reverse engineer corresponding optimal solutions. Subsequent computational analysis guides our theoretical intuition toward the identification of broader design principles.

RESEARCH CHALLENGE

All animals that swim or fly are constantly confronted with the physics of the surrounding media. Whether natural creatures can take advantage of physics to push their performance limits depends on their biological strategies, from materials, morphologies, and gaits to collective behaviors. Their struggle for survival has produced a rich array of solutions that often outperform human engineered designs and that work in ways we still do not

fully understand. They bear a great potential for technological innovation, with applications ranging from robotics to energy-harvesting devices.

In this context, bioinspired approaches rely on mimicking existing natural solutions to enhance the performance of current engineered designs. Nevertheless, one may question to what extent engineers should follow biomimicry. Indeed, natural creatures have not evolved to optimize engineering objectives, and we have limited information to determine whether a particular solution optimally serves a given function. Moreover, today's materials and components pose constraints and enable opportunities that may differ from their biological counterparts. Hence, we suggest that man-made solutions obtained through inverse design based on an automated optimization process may outperform pure biomimicry.

METHODS & CODES

The characterization of biopropulsion, from the standpoint of optimality, demands accurate, robust, fast, and flexible numerics for flow-structure interaction problems. We have been developing and implementing novel schemes for the direct numerical simulation of individual and multiple swimming bodies. Our algorithms rely on remeshed vortex methods enhanced with projection approaches to capture the effects of the fluid on the body, and with a penalization technique, to capture the effects of the body on the fluid [1]. These techniques enable wavelet-based multiresolution discretization [2], effective mapping




Figure 1: Larval fish developed an optimal escape mechanism: they bend into a “C” shape and then flee from a threat. C-starts’ optimality was demonstrated quantitatively by evolving (via evolutionary optimization) fish gaits that maximize escape distances [5]. The identified solution closely resembles real fish escape patterns. Flow analysis reveals the underlying mechanism: fish accommodate a “ball” of water (grey region) in the “C” and then push it with a backflip of the tail to gain momentum in the opposite direction. Therefore, the wake flow structures play a far lesser role than assumed thus far.

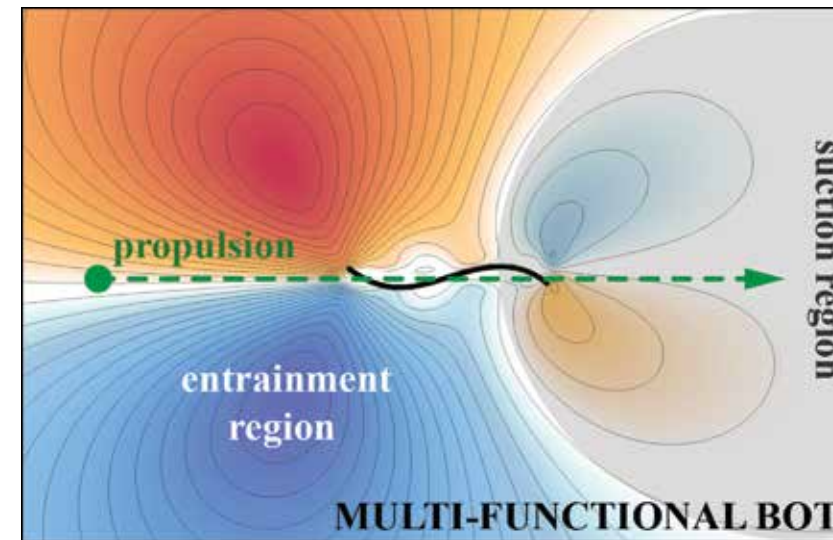


Figure 2: Flexible slender body and wave actuation produce locomotion, suction and entrainment regions.

on supercomputing architectures [3], and provide accuracy, robustness, and flexibility [1]. Simulations in 2D are performed via the Multi-Resolution Adaptive Grid (MRAG) library [2,3]. Simulations in 3D are carried out via our penalization client [1] of the Parallel Particle Mesh library [4]. These flow-structure interaction solvers are embedded in an inverse design cycle that relies on a combination of Evolutionary Optimization Strategies and Reinforcement Learning techniques for the identification of optimal biolocomotion strategies in fluids, in relation to a desired metric or task [5,6,7].

RESULTS & IMPACT

The coupling of realistic numerical simulations with artificial intelligence techniques is one of the frontiers of Computational Fluid Dynamics and is a unique aspect of this project. We have successfully demonstrated the predictive power of this approach and its ability to provide biophysical insight in the context of rational design of artificial swimmers. For example, we showed that the C-start mechanism (Fig. 1), a widespread escape response among living fish, is optimal as it maximizes the swimmers’ ability to channel water displacement into forward acceleration [5]. Furthermore, we showed that artificial swimmers have the potential to outperform biological ones [6]. We are currently focusing on the design of swimmers able to “sculpt” the surrounding flow in order to achieve multi-tasking behavior (Fig. 2).

WHY BLUE WATERS

Our inverse design process entails thousands of flow-structure interaction simulations, each one characterized by billions of computational elements. Without the sheer size and cutting-edge technology of Blue Waters, these investigations simply would not be possible.

COUPLED MULTI-PHYSICS OF ADVANCED MOLTEN SALT NUCLEAR REACTORS

Allocation: Blue Waters Professor/30 Kwh

PI: Kathryn Huff¹

Collaborator: Alexander Lindsay¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

The Advanced Reactors and Fuel Cycles Group (ARFC) conducts modeling and simulation in the context of nuclear reactors and fuel cycles and toward the improved safety and sustainability of nuclear power. In the context of high-performance computing, this work requires the coupling of multiple physics models at multiple scales to model and simulate the design, safety, and performance of advanced nuclear reactors. In particular, thermal-hydraulic phenomena, neutron transport, and fuel performance couple tightly in nuclear reactors. Detailed spatially and temporally resolved neutron flux and temperature distributions can improve designs, help characterize performance, inform reactor safety margins, and enable validation of numerical modeling techniques for those unique physics. In the work presented here, conducted on Blue Waters, ARFC has demonstrated the capability to simulate coupled, transient neutronics and thermal hydraulics in an advanced, molten-salt-fueled nuclear reactor.

RESEARCH CHALLENGE

Nuclear power provides 19 percent of the total electricity generation in the United States and is our largest source of clean energy. The current state of the art in advanced nuclear reactor simulation (e.g., the CASL DOE innovation hub) is focused primarily on more traditional light-water reactor design types. Our work extends that state of the art by enabling similar modeling and simulation capability for more advanced reactor designs that have the potential to improve the safety and sustainability of nuclear power. High-fidelity simulation of dynamic reactor performance of these designs requires development of models and tools for representing unique materials, geometries, and physical phenomena.

The current work is a finite-element-based physics application, Moltres, that extends the open source MOOSE framework to appropriately model coupled thermal-hydraulics and neutronics of molten salt flow in high-temperature liquid-fueled reactor designs. By developing the Moltres application in the open, ARFC seeks to enable both transparency and distributed collaboration on nuclear reactor concepts that promise advanced safety or sustainability.

METHODS & CODES

ARFC has developed Moltres [1], a collection of physics kernels material definitions, to extend the ecosystem of applications built upon the highly scalable, fully implicit, Multiphysics Object-

Oriented Simulation Environment (MOOSE) framework from Idaho National Laboratory [2]. These physics kernels enable Moltres to solve arbitrary-group neutron diffusion, temperature, and precursor governing equations in anywhere from one to three dimensions and can be deployed on an arbitrary number of computational processing units. Moltres is devoted to previously unmatched fidelity in coupled neutronics and thermal hydraulics MSR simulation.

MOOSE and LibMesh handle translation of Moltres-defined weak PDE forms into residual and Jacobian functions that are the inputs into Petsc Newton-Raphson solution routines. All codes use MPI for parallel communication and are easily deployed on massively-parallel cluster-computing platforms. By default, MOOSE applications use monolithic and implicit methods that are ideal for closely coupled and multiscale physics.

To solve these large systems of partial differential equations on a finite element mesh in a fully coupled, implicit way, the MOOSE framework was designed to take advantage of high-performance computing capability. Accordingly, it employs a hybrid shared and distributed memory parallel model. These simulations are memory intensive, so the exceptional memory capability of the Blue Waters resource will be essential to performant simulation times. It is also important to note that rendering visualizations of the results can be computationally intensive and that a MOOSE tool exists for taking advantage of HPC resources for conducting those rendering tasks. Explorations of the use of the *yt* visualization toolkit [5] for this purpose have also been pursued in collaboration with the Data Exploration Lab, led by Prof. Matthew Turk.

RESULTS & IMPACT

Blue Waters has enabled ARFC to develop and test a first-of-its-kind, scalable, finite element model of the transient neutronics and thermal hydraulics in a liquid-fueled molten salt reactor design. Moltres is a physics application for multiphysics modeling of fluid-fueled molten salt reactors (MSRs). It couples equations for neutron diffusion, thermal hydraulics, and delayed neutron precursor transport. Neutron diffusion and precursor transport equations are set up using an action system that allows the user to employ an arbitrary number of neutron energy and precursor groups respectively with minimal input changes. Moltres sits atop MOOSE, which gives Moltres the capability to run seamlessly in massively parallel environments. To date, Moltres has been used to simulate MSRs in 2D-axisymmetric and 3D geometric

configurations. As these simulations increase in fidelity, their results will be able to inform the safety and sustainability case for deployment of advanced commercial nuclear reactors.

Moltres solves arbitrary-group neutron diffusion, temperature, and precursor governing equations in anywhere from one to three dimensions and can be deployed on an arbitrary number of processing units. The model problem presented here has a 2D-axisymmetric geometry with heterogeneous group constants for fuel and moderator regions. Fuel volume fraction and fuel salt composition are based on the MSR experiment design. Neutron

fluxes show expected cosinusoidal shapes in radial and axial directions with visible striations between fuel and moderator regions. Additionally, the fast group neutron flux is enhanced in fuel regions while the thermal group flux is enhanced in moderator regions. Due to advection, the temperature profile increases monotonically in the direction of salt flow. The role of advection is also seen in precursor concentrations. Long-lived precursors exhibit maximum concentrations at the core outlet. As the decay constant increases across precursor groups, the maximum concentrations move toward the reactor center where the precursor production rate is maximum. Future Moltres work includes generating a high-fidelity, 3D model as well as investigating various transient accident scenarios, additional reactor configurations, and numerous design concepts.

WHY BLUE WATERS

To assess nuclear reactor performance under a variety of conditions and dynamic transients, the ARFC group must conduct myriad two-dimensional and three-dimensional finite element simulations using the MOOSE framework and our in-house developed modules. Such simulations commonly occupy tens of thousands of CPU cores at a time and vary in completion time. The MOOSE framework has been shown to scale very well up to 10,000 cores. The ARFC group has demonstrated appropriate scaling for MSR simulation above 20,000 CPU cores (600 Blue Waters nodes). Transient and multiscale simulations, which require greater capability per simulation, are on the horizon for our work. These may occupy up to 100,000 CPU cores at a time. Only a few of those larger simulations will be necessary to enable better understanding of the dynamics in these reactor systems.

PUBLICATIONS AND DATA SETS

[1] Lindsay, A., K. Huff, and A. Rykhlevskii, *ARFC/Moltres: Initial Moltres release*, Zenodo (2017), DOI: 10.5281/zenodo.801823




Figure 1: This image shows the neutron flux in a 2D cylindrical axisymmetric model of a molten-salt reactor. This flux has the anticipated magnitude and canonical cosine shape ($r = 0$ is center of core) and is undergoing validation against experimental results from the Molten-Salt Reactor Experiment.




Figure 2: This image shows the temperature in a 2D cylindrical axisymmetric model of an MSR. The reactor core temperature peaks near the reactor outlet in this model because of fuel advection ($r = 0$ is center of core).

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

Allocation: NSF PRAC/9,000 Knh

PI: Sohrab Ismail-Beigi¹

Collaborators: Glenn Martyna², Laxmikant Kale³

¹Yale University

²IBM TJ Watson

³University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Hydrogen has the potential to be a major green energy source, but, for use as a fuel, it requires efficient storage materials for retaining and releasing hydrogen in large quantities. Metal-organic frameworks (MOFs) are potential storage materials, and their hydrogen storage potential and microscopic properties need further investigation. We perform first-principles quantum-mechanical molecular dynamics calculations to understand the behavior of hydrogen inside MOFs. We include the effect of quantum nuclear motions, which are critical for the properties of hydrogen, the lightest atomic element. These challenging calculations are only possible on large-scale and tightly coupled computational platforms such as Blue Waters. The knowledge gained from our studies informs the materials research community as to the fundamental reasons underlying the properties of hydrogen in MOFs, and, potentially, how we may improve the chemical composition of MOFs for hydrogen storage.

RESEARCH CHALLENGE

We use large-scale and accurate quantum mechanical calculations on an important class of porous hydrogen storage materials: metal-organic frameworks (MOFs). Hydrogen has a strong potential to be a green energy source, but as a fuel it requires efficient storage materials that retain and release a great deal of hydrogen as desired. We study the properties of hydrogen inside MOFs to understand their physical properties and potentially how to improve MOFs to deliver improved hydrogen storage. We simulate hydrogen inside of MOFs at the atomistic scale using accurate first-principles quantum-mechanical simulations based on density functional theory. In addition to large-scale molecular dynamics simulations of the diffusion and dynamics of hydrogen inside MOFs, we describe the atomic nuclei in the entire system quantum mechanically via a path integral formalism. The effects due to the quantum fluctuations of the nuclear degrees of freedom are critical for understanding the binding and dynamics of light elements such as hydrogen.




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

METHODS & CODES

Answering the above questions requires an accurate quantum-mechanical simulation method at finite temperature. The method must work beyond the harmonic approximation as hydrogen motion can be diffusive and highly anharmonic. Due to the use of complex ligands chelated to metal ions and the metal ions themselves, standard pairwise force fields can be problematic in describing the energetics, and parameterizing a more transferable model is both time-consuming and difficult. We employ the Path Integral Car-Parrinello Molecular Dynamics simulation technique (PI-CPAAMD), which allows the nuclei to move on the Born-Oppenheimer energy surface provided by plane wave-based Density Function Theory (DFT) and includes quantum effects via Feynman Path Integrals. By including the (valence) electronic degrees of freedom explicitly, we bypass force-field difficulties; by employing path integrals, nuclear quantum effects are treated in a completely general way and are converged with a single parameter.

The use of the PI-CPAAMD method and the high-quality results it delivers requires highly scalable software, developed by us under NSF support (NSF SI2 grants ACI-1339804 and ACI-1339715), and named OpenAtom. CPAAMD, which simulates classical mechanics of the nuclei moving on the DFT-derived energy surface, has a computational workload scaling as N^3 where N is the number of atoms in the simulation. Adding Path integrals, the computational cost increases by the factor P to PN^3 where P is the number of "beads" used to discretize the Feynman path integral. Typically, P ranges between 20 and 80, which is a significant increase in computational requirements. As one can imagine, the computational cost requires petascale resources, and a supercomputer such as Blue Waters is needed to carry out this work. Work on Blue Waters is supported by NSF PRAC grant OAC-1614491.

RESULTS & IMPACT

In the first year of our PRAC project on Blue Waters, we have spent the majority of the time studying a model MOF system with hydrogen to understand the behavior of the system and its various timescales. Fig. 1 illustrates the unit cell of the MOF crystal with hydrogen molecules inside. A smaller version of the MOF was carefully studied with CPAAMD long simulations, which permitted us to obtain preliminary results. Fig. 2 shows the results of the CPAAMD simulations (without quantum effects) for the hydrogen probability distribution inside the MOF when the system is at 300°K: The hydrogen molecules spend most of their times in the "voids" (corner regions) away from the MOF backbone, signaling weak binding to the MOF backbone at these temperatures and potentially rapid diffusion. Other measures of diffusion (e.g., the hydrogen diffusion coefficient) are in good agreement with prior results, giving us confidence that the benchmarking has been successful in producing physically accurate and important results. The next step is to perform CPAAMD simulations with the full larger MOF system and then to turn on quantum nuclear effects




Figure 2: Preliminary heatmap of hydrogen diffusion inside the small MOF model system. The heatmap shows the averaged probability distribution of finding hydrogen atoms inside the MOF at 300°K. The backbone of the MOF itself is shown in white while the colormap shows hydrogen probability density with brighter colors indicating higher probabilities (linear scale).

via the path integral approach to understand how the quantum effects modify the behavior of hydrogen inside the MOF.

Successful modeling of the behavior of hydrogen inside MOFs will provide valuable information to the community of materials and energy researchers, in academia as well as at national and industrial labs, regarding the underlying reasons for the performance of these materials. Chemical modifications of the MOF structure can then be explored using such simulations to help design improved hydrogen storage performance.

WHY BLUE WATERS

Massively parallel electronic structure calculations require tightly coupled computing nodes due to intense communication loads: Electron waves are delocalized over the entire system so all parts of the system end up interacting with each other. For the MOF system of interest, the CPAAMD simulations already require a massively parallel calculation with many hundreds of nodes. The inclusion of nuclear quantum effects, however, means that only a tightly coupled petascale computer such as Blue Waters is capable of delivering results on a reasonable (one- to two-year) timescale.

ATOMISTIC MODELING OF TRANSFORMATIONS IN NANOCRYSTALS

Allocation: Exploratory/30 Kwh
PI: Prashant K. Jain¹
Co-PI: Sudhakar Pamidighantam²

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

EXECUTIVE SUMMARY

Discovery of materials drives renewable energy and optoelectronic technologies. While using an unconventional method of nanostructure synthesis in our lab, we encountered a novel form of mercury–cadmium selenide. The wurtzite structure of this crystal distinguishes it from the natural zincblende phase. The properties of the material were unknown, motivating our computational investigations of its electronic structure. These computations revealed how subtle differences in crystallographic symmetry combine with the relativistic nature of electrons to result in novel behavior. Natural mercury selenide is a zero-band-gap semimetal. However, the wurtzite form has a band gap and is a three-dimensional topological insulator, expected to exhibit spin-protected conduction on the surface, while being electrically insulating in the bulk, a property useful for next-generation logic devices and the discovery of exotic quasi-particles. The study also elucidates how crystal structure and chemical composition can be paired to tune relativistic effects and topology of electron motion.

RESEARCH CHALLENGE

Engineered nanocrystals are often utilized for making new functional electronic and optical materials such as superionic solids and battery electrodes. Our laboratory makes use of unconventional methods that enable manipulation of the chemical composition and crystal structure of nanocrystals. These techniques often produce novel compositions and crystal phases that are often

not found in the bulk phase diagram. Computational electronic structure investigations are allowing us to explore the properties of these new, unconventional materials that have been created for the first time. Alongside, we are also elucidating chemical trends in heterostructures and alloys and developing solid-state physics principles from these trends. The results from our investigations will enable the rational design of new phases and compositions with targeted applications for resolving longstanding challenges of energy storage and device efficiency.

With advances in nanotechnology and chemical synthesis, materials are becoming ever more complex. Computations can uncover chemical principles that will ultimately allow prediction of the properties of tomorrow's indispensable materials, an existing Grand Challenge. However, these studies require extensive calculation spanning a range of physicochemical parameters. As opposed to a single large and expensive calculation, our work requires a library of moderately expensive calculations. The net cost for generating such a library of data is feasible only with a resource such as Blue Waters, with considerable payoff for future scientific advances. Dissemination of solid-state physics principles, like those resulting from our project, will accelerate the discovery and design of new materials through systematic exploration, supplanting time- and energy-consuming trial-and-error pursuits.




Figure 1: In the novel wurtzite form of HgSe, uniaxial elongation of Hg–Se bonds results in the opening of a band gap between the highest energy valence (green curve) and lowest energy conduction (red curve) bands. Hg atoms are shown in purple and the electron density is shown in red.

METHODS & CODES

We used the open source Quantum Espresso software suite [1] to run our electronic structure calculations. In order to study the effects of chemical composition and structure on the alloy's electronic properties, numerous calculations were run. Each calculation is distinguished from the others in the crystal geometry, chemical formula, or both. Comparisons of energies, electron distribution, and band structure across all the calculations allowed us to make predictions about novel behavior of the wurtzite polymorph of the alloy and also understand the role of relativistic effects and symmetry breaking.

RESULTS & IMPACT

Our electronic structure calculations demonstrated that bond elongation in a novel polymorph of HgSe and $\text{Hg}_x\text{Cd}_{1-x}\text{Se}$ alloys is responsible for the opening of a band gap (Fig. 1). The presence of this band gap is of significance because it, when combined with the inverted nature of bands in HgSe, qualifies these materials as potential three-dimensional (3D) topological insulators (TIs). 3D TIs are of interest because electrons at their surface states are spin-protected from back scattering. This protection allows 3D TI materials to conduct electrons along their surface without resistance, that is, with minimal loss of energy as wasteful heat. For this reason, 3D TIs are garnering interest as components of energy-efficient logic devices that can operate at high capacity while simultaneously not contributing to the generation of heat. Such devices would reduce the need for coolants and could permit the design of more powerful supercomputers by overcoming problems caused by overheating. TI materials are also expected to harbor exotic quasi-particle states of importance in particle physics. However, only a handful of TI materials have been identified, and through our contribution we are adding this alloy system to the list of 3D TIs. In addition, our work unveils crystallographic anisotropy as a powerful synthetic handle for tuning band topology.

Our study further contributed two new principles governing band structure and topology in $\text{Hg}_x\text{Cd}_{1-x}\text{Se}$ and similar alloys of a strongly relativistic and a weakly relativistic metal. The first of these new insights involves the effect of mixing electron character (Fig. 2a–2c). HgSe and CdSe have different energy ordering of their bands; the band order is inverted in HgSe with respect to the conventional CdSe case. This inversion is due to the greater relativistic mass of the Hg valence “s” electrons resulting in the contraction and energetic downshift of their orbitals. Thus, in $\text{Hg}_x\text{Cd}_{1-x}\text{Se}$ alloys, wherein the electronic character of the two metals is mixed, the energy of the “s”-like band, the resulting band ordering, and the band gap depend on the relative abundance of Hg and Cd. At compositions of the wurtzite alloy, where the “s”-like band of spherical symmetry becomes close to the valence band edge, the light hole and heavy hole energies become less sensitive to crystallographic anisotropy. As a result, a smaller gap is exhibited between these two bands, negating the effect of the symmetry breaking caused by crystalline anisotropy. The second




Figure 2: Electron density (red) maps elucidating trends in $\text{Hg}_x\text{Cd}_{1-x}\text{Se}$ (Hg, purple; Cd: grey; Se: yellow). Unlike CdSe (a) and HgSe (c), alloys (b) show significant mixing of the states, which weakens the effect of bond elongation. (d) Creation of alternating Cd and Hg layers has a similar effect as bond elongation.

finding relates to the effect of a layered structure of the crystal (Fig. 2d). Ordering of the two metals Hg^{2+} and Cd^{2+} in alternate layers mimics the effect of crystallographic anisotropy and provides a second means by which a band gap can be opened. These principals gained from our study translate readily to existing TI systems and also enable the design of new alloys with TI behavior.

WHY BLUE WATERS

Calculations of crystal surfaces and nanocrystals may be accomplished on computational resources other than Blue Waters. However, a large number of single calculations of these structures with varying elemental composition is required to study chemical trends. The computational expense of such an effort would be prohibitive for us were it not for a Blue Waters allocation. Furthermore, the specialized hardware of Blue Waters allows the Quantum Espresso code to run even more efficiently. This is because Quantum Espresso's parallelization schemes involve sizable and frequent communication among CPUs, which rely on the speed of the Blue Waters communication hardware. Furthermore, our ability to work with Sudhakar Pamidighantam of NCSA as a co-PI has allowed us to greatly expedite code preparation and troubleshooting with use of the SeaGrid portal [2].

PUBLICATIONS AND DATA SETS

Dumett Torres, D., P. Banerjee, S. Pamidighantam, and P. K. Jain, A Non-Natural Wurtzite Polymorph of HgSe: A Potential 3D Topological Insulator, submitted to *Chem. Mater* (2017).

NUMERICAL SIMULATIONS OF COLLAPSING CAVITATION BUBBLES ON BLUE WATERS

Allocation: GLCPC/880 Kwh

PI: Eric Johnsen¹

Co-PIs: Shahaboddin Alahyari Beig¹, Mauro Rodriguez¹

¹University of Michigan

EXECUTIVE SUMMARY

The collapse of cavitation bubbles is known to cause damage, ranging from the erosion of naval propellers to soft tissue ablation. While significant attention has been dedicated to investigating this phenomenon in the case of hard materials such as metals, less is known about cavitation-induced damage to soft materials.

In order to elucidate these damage mechanisms, we perform high-fidelity numerical simulations of the collapse of individual and multiple vapor bubbles near both rigid and compliant surfaces. We have developed a computational framework to conduct massively parallel simulations of the three-dimensional compressible Navier–Stokes equations for gas–liquid flows. These simulations provide a clearer image of the detailed nonspherical bubble dynamics, pressure, and temperature fields, and stresses/deformations of the neighboring solid. This knowledge and data enable us to develop numerical models for the collapse of bubble clouds that can be used in biomedical or naval hydrodynamic applications of interest.

RESEARCH CHALLENGE

Cavitation occurs in a wide range of hydraulic applications, such as naval engineering, turbomachinery, and biomedical ultrasound. In contrast with boiling, in which liquid vaporizes as temperature is increased (i.e., a thermally driven phase transition), cavitation occurs when local pressure reductions of a liquid lead to the formation of vapor bubbles. These cavitation bubbles dramatically respond to pressure changes, growing to sizes far greater than their equilibrium radius and undergoing a violent inertially dominated collapse [1]. As a result, shock waves and high-pressure and high-temperature regions are produced, which may damage neighboring solid objects such as propellers or soft tissues [2–5].

While relatively well understood in the context of hard materials (e.g., metals), cavitation erosion to soft matters is not well known [3]. Given the prevalence of cavitating flows in a vast variety of applications, there is an essential need to investigate the detailed bubble dynamics and to predict pressures, temperatures and deformations produced. However, compressibility effects, including the propagation of shock and rarefaction waves in a multiphase medium, result in a complicated nonlinear multi-scale and multi-physics problem that is challenging to solve. Moreover, owing to the wide range of temporal and spatial scales of these flows, precise and accurate measurements are nearly impossible to obtain experimentally. Numerical simulations of effects not

available and/or feasible via experimentation have therefore emerged as a powerful tool to complement and enhance our fundamental understanding of these flows [6].

METHODS & CODES

In order to perform high-resolution simulations of the three-dimensional compressible Navier–Stokes equations for a gas–liquid system, we have developed a novel computational algorithm [7]. We used an accurate model of compressible multiphase flows capable of resolving flows around cavitating bubbles to correctly compute pressures and temperature across the material interfaces.

The algorithm employs a third-order accurate explicit strong-stability-preserving Runge–Kutta scheme to march in time [8]. For the spatial discretization, we proposed a solution-adaptive, high-order accurate, central difference/discontinuity-capturing method. This method can represent both broadband flow motions and discontinuities accurately and efficiently. The basic idea is that nondissipative methods are used where the solution is smooth, while the more dissipative and computationally expensive capturing schemes are applied near discontinuous regions. For this purpose, a discontinuity sensor discriminates between smooth and discontinuous (shocks, contacts, and interfaces) regions, which all require a different treatment. For large-scale parallel calculations, our in-house petascale production code achieves parallel efficiency over 83% on 4,096 processors using MPI; we have also been exploring alternative approaches such as GPU acceleration. Our code also uses the parallel HDF5 library to manage large and complex data collections.

RESULTS & IMPACT

The current project focuses on two specific problems: the collapse of a single bubble near solid/soft media and the collapse of multiple bubbles near solid/soft media. The first problem provides insight into the detailed dynamics of the collapse, including the nonspherical behavior of the bubble, high-velocity jet formation, propagation of shock waves, and the vortex ring convecting toward the boundary. The simulations provide the flow field pressure and temperature distribution throughout the collapse, which can be used to model cavitation erosion (Fig. 1).

Although studying single-bubble collapse is valuable specifically in exploring the flow physics, the disruptive effects of cavitation erosion are generally caused by the collapse of bubble clouds containing tens of thousands of bubbles. However, resolving




Figure 1: The microjet impacting the distal side of the bubble generates a strong outward propagating shock wave. Later on, this shock hits the wall, reflects back, and eventually interacts with the collapsed bubble which results in a second collapse (left). Thereafter, the bubble takes the form of a convoluted vortex ring, which is convected toward the adjacent wall (right). This image is produced from one of our simulations performed on a different NSF-supported supercomputer, Stampede.

every bubble is computationally prohibitive, such that a model representing the collective behavior of the bubble cloud is necessary. Currently, such models are rudimentary [9], as they only account for spherical bubble dynamics. By conducting resolved simulations of individual bubbles, we are investigating the complex interactions within the collapsing bubbles in order to establish a numerical cloud model that includes the bubble–bubble interactions and the nonspherical effects of the collapse.

This project will deepen our knowledge and understanding of the nonspherical behavior of the bubbles in cavitating flows. Developing a comprehensive model for bubble clouds will result in more precise numerical simulations of the collapse of bubble clusters. This will help to reduce the disruptive effects of cavitation erosion in naval applications and turbomachinery. Reducing cavitation erosion will significantly strengthen the structures exposed to deterioration caused by cavitating flows, extend the lifetime of machineries, and has the potential to save billions of dollars each year. Moreover, a well-known treatment in biomedicine utilizes structural damage induced by collapsing bubbles to break apart kidney stones. The same idea has been applied to destroying malignant cells and soft tissue. However, a negative side effect of these medical advancements is the unwanted damage to neighboring cells and tissues. This study delves further into these matters to provide insight into damage mechanisms to better control them in order to minimize the impairment of healthy cells and tissues.

WHY BLUE WATERS

This project utilizes two different in-house codes: our petascale production code for the large-scale simulations (based on MPI) and our next-generation code that enables larger-scale heterogeneous architectures (based on MPI and GPUs). These codes solve the compressible Navier–Stokes equations for multiple gases and liquids. The foundations of both codes are high-order accurate algorithms, explicit in time and in space, thus naturally lending itself to massive parallelization. To carry out accurate simulations of cavitating flows that effectively resolve the small-scale features, extremely high spatial resolution is essential to run even a single simulation over long compute times, which is difficult to achieve on any other NSF computing resource. Given its speed and available computation power, the Blue Waters supercomputer is capable of providing us with this opportunity.

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

Allocation: NSF PRAC/1,239 Knh

PI: Gerhard Klimeck¹

Co-PI: Tillmann Kubis¹

Collaborators: Tarek Ameen¹, Daniel Valencia¹, Pengyu Long¹, James Charles¹, Daniel Lemus¹, Harshad Sahasrabudhe¹, Kuang-Chung Wang¹, Junzhe Geng¹, Chinyi Chen¹, Xinchun Guo¹, Prasad Sarangapani¹, YuanChen Chu¹, Tillmann Kubis¹

¹Purdue University

EXECUTIVE SUMMARY

With a revenue of \$338.9 billion in 2016, and as an enabler for larger economy chains (i.e., electronic systems), the semiconductor industry influences approximately 10% of the world GDP [1]. The transistor is at the heart of this enormous industry and continuous improvements of transistors in terms of speed and power consumption are essential for the stability and growth of the semiconductor industry as well as the dependent product chains and economies. Needed improvements in transistors' performance have driven the semiconductor industry to push for smaller transistors, reaching 14 nanometers (nm) in the latest technology mode, while development is ongoing for 10 nm technology and beyond. Such aggressive downscaling into a countable number of atoms in the critical dimensions makes atomistic simulations necessary pathfinders in the quantum regime. NEMO5 is designed to comprehend the critical multiscale, multi-physics phenomena for nano-scale technology through efficient computational approaches, and enables quantitative study of new generations of nano-electronic devices even beyond transistors [2–3].

RESEARCH CHALLENGE

The U.S. has always been a world leader in the semiconductor industry with 40% of the worldwide semiconductor device-related patents originating in the U.S [4]. The U.S. semiconductor industry is one of the nation's largest and most strategic industries, and the U.S. holds one-third of the global semiconductor device market worth over \$300 billion per year. Simultaneously, a relentless downscaling is occurring, with devices expected to be about 5 nm long in their critical active region within 10 years. Further

improvements in shrinking dimensions will come only through the detailed study of device designs, materials, and of quantum effects such as tunneling, state quantization, and atomistic disorder. Fundamental questions remain about the downscaling of the CMOS (complementary metal-oxide-semiconductor) switch and its eventual replacement. What is the influence of atomistic local disorder from alloy, line-edge roughness, dopant placement, and fringe electric fields? How do lattice distortions due to strain affect carrier transport in nanometer-scale semiconductor devices such as nanowires, finFETs, quantum dots, and impurity arrays? What are the effects of interconnects' sidewall roughness, grain boundaries, electron-phonon scattering, and roughness of metal-dielectric interfaces? Can inserting new materials and device concepts reduce power consumption?

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

METHODS & CODES

iNEMO's research on Blue Waters encompasses multi-physics atomistic quantum simulations implemented in NEMO5. The needed physics vary from one device to another, but all can be simulated inside NEMO5. Several examples follow: For transistors, quantum transport simulations are performed using the self-consistent Poisson with nonequilibrium Green's function (NEGF)

approach employing semi-empirical tight binding methods. The current mechanism of tunneling field effect transistors (TFETs) is through inter-band tunneling rather than thermionic emission in typical MOSFETs (metal-oxide semiconductor field-effect transistors). For realistic performance predictions, incoherent scattering effects are also included.

For nitride devices, namely nitride-based light emitting diodes and alloy engineered Nitride TFETs, a multi-scale quantum transport model is used, which treats high-density regions as local charge reservoirs, where each reservoir serves as carrier injector/receptor to the next/previous reservoir.

RESULTS & IMPACT

For nitride devices, the I–V characteristics produced from these simulations agree quantitatively with experimental measurements. The simulations have been used to suggest improvements in the multi-quantum-well nitride-based light-emitting diode. In addition to this, a new alloy engineered Nitride TFET is proposed as a novel low-power transistor design.

Regarding transistor simulations, the tunneling from the valence to conduction band has the potential to yield significantly improved subthreshold slopes to allow lower supply voltages and much-needed lower power consumption compared to MOSFETs. Electron transport through the overall device, including the source and drain, entails significant amounts of computationally demanding scattering, which cannot be ignored in realistic device performance predictions. A new method on the rank reduction of matrices through basis transformations that retain key physical information for modeling incoherent scattering phenomena has been implemented in NEMO5. These low-rank approximations provide shorter times-to-solution and smaller memory footprints. Blue Waters was used for assessing these times-to-solution and memory improvements. Using these low-rank approximations on nano-wire with a width of 5 nm shows a speed-up of 200 and needs only 7% of the memory.

WHY BLUE WATERS

Quantum transport simulations are very computationally expensive and memory demanding due to the high degree of complexity of the equations used, especially if incoherent scattering of particles is needed. A toy quantum transport calculation of a 50 nm long wire with a 3 nm diameter requires around 1 teraflop for a single energy point where more than 1,000 energy points are needed. And this calculation must be repeated perhaps hundred times for a full current-voltage sweep. The treatment of a realistic device would require an atomistic resolution of a device with a cross section of more than (20x20) nm², which includes the core semiconductor and the surrounding gate material. Such devices of larger sizes are especially an issue due to the $O(n^3)$ scaling of matrix operation time-to-solution and $O(n^2)$ scaling of memory. Blue Waters was used for running such simulations on up to 16,384 cores per simulation. In many cases the work could not be accomplished in a reasonable amount of time without Blue Waters,

and for the larger simulations the work could not be accomplished on other available systems. Blue Waters staff provide exemplary support and user outreach to guide system usage, help with issues as they arise, and assistance with code performance and scaling.

PUBLICATIONS AND DATA SETS

Geng, J., et al., NEMO5: realistic and efficient NEGF simulations of GaN light-emitting diodes. Invited talk, *SPIE Photonics West OPTO 2017*, San Francisco, Calif., January 30–February 2, 2017.

Charles, J., et al., Non-Local Scattering with a New Recursive Nonequilibrium Green's Function Method. *IWCE 2017*, Windermere, United Kingdom, June 7–9, 2017.

Ilatikhameneh, H., T. Ameen, G. Klimeck, and R. Rahman, Universal behavior of atomistic strain in self-assembled quantum dots. *IEEE Journal of Quantum Electronics*, 52:7 (2016), pp. 1–8.

Huang, J., et al., Sb- and Al-Free Ultra-High-Current Tunnel FET Designs. *5th Berkeley Symposium on Energy Efficient Electronic Systems & Steep Transistors Workshop*, Berkeley, Calif., October 19–20 2017.

Long, P., et al., A high-current InP-channel triple heterojunction tunnel transistor design. *75th Device Research Conference*, Notre Dame, Ind., June 25–28, 2017

Valencia, D., et al., Grain boundary resistance in nanoscale copper interconnections. *SISPAD 2016*, September 6–8, 2016.

Ameen, T., et al., Alloy Engineered Nitride Tunneling Field Effect Transistor: A Promising Energy Efficient Transistor. *TECHCON*, Austin, Texas, September 10–12, 2017.

Chen, C., et al., Channel thickness optimization for TFETs. *TECHCON*, Austin, Texas, September 10–12, 2017.

Sarangapani, P., et al., Non-equilibrium Greens function method: Band tail formation in non-local polar optical phonon scattering. *International Workshop on Computational Nanotechnology*, Windermere, U.K., June 5–9, 2017.

Sarangapani, P., et al., Non-equilibrium Green's function (NEGF) Approach: Phonon and Impurity Driven Band Tail Formation in Polar Materials. *TECHCON*, Austin, Texas, Sept. 10–12, 2017.




Figure 1: A 3D visualization of a typical sub-10 nm novel transistor design made from bilayer phosphorene that is being investigated using NEMO5. (Credit: Tarek Ameen)




Figure 2: Energy-resolved electron, hole density of states (contour lines) filled with electrons and holes (color contours) of energy efficient multi-quantum well light-emitting diode simulated by NEMO5. (Credit: Junzhe Geng, Gerhard Klimeck)

KINETIC SIMULATIONS OF UNSTEADY SHOCK-BOUNDARY LAYER INTERACTIONS

Allocation: Illinois/52 Knh
 PI: Deborah A. Levin¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Hypersonic flow over configurations such as a double wedge at continuum-like free stream conditions has been a challenging problem because of the multiple shock–shock and shock–boundary layer interaction, separated flows near the hinge, shear layer, and three-dimensional effects. These conditions generate a mesh that is highly nonuniform because of very high levels of refinement near the surface due to extremely high flow gradients in temperature and pressure. The octree cells lying in this region are highly refined as compared to those in the free stream and inside the geometry leading to a high degree of load imbalance among the processors, which can cause significant increases in communication time. In addition, these flows are unsteady in nature but provide an opportunity to study flow stability mechanisms and identify near-transition behavior. However, large sizes of additional particle data must be collected in order to analyze the time-dependent signals.

RESEARCH CHALLENGE

Hypersonic compressible flows are characterized by significant gradients in fluid density, high fluid temperatures, and a large degree of nonequilibrium. The multiscale nature of these problems makes it a challenge in terms of flow physics and computational intensity. This work was motivated by the study of laminar shock wave boundary layer interaction (SWBLI) problems, which are especially difficult to accurately model due to complex shock interactions, flow separation, three-dimensional effects, and shear layers. All these effects play a significant role

on aerothermodynamic quantities such as heat transfer, skin friction, and pressure loads over different angular sections. The particle-based Direct Simulation Monte Carlo (DSMC) method is a good candidate to simulate such flows because it solves the Boltzmann transport equation exactly; hence, it is valid even in strong shock regions where the continuum assumptions inherent in the derivation of Navier–Stokes equations fail. However, the extent of its applicability to such problems has been limited by the high computational requirements arising from the need to simulate a large number of particles to satisfy DSMC requirements such as having the cell size comparable to the local mean free path, a time step that is smaller than the local mean collision time, and sufficient numbers of particles per cell to perform collisions.

METHODS & CODES

We have developed a three-dimensional MPI-parallelized DSMC code known as Scalable Unstructured Gas-dynamic Adaptive mesh Refinement (SUGAR). The code uses new techniques to simulate hypersonic, compressible flows such as an octree-based adaptive mesh refinement (AMR) implementation for capturing multiscale physics, linearized representation of an unstructured grid using Morton-Z space filling curve for efficient access of computational cells, an accurate cut-cell algorithm to compute correct volume of intersected computational cells, algorithmic improvements for efficient gas–surface interactions, and array-based data structures for optimal use of cache memory utilization [1].

A great deal of work has been done to improve the scalability of the code. In brief, the improvement in the strong scaling, as compared to the original SUGAR version presented at the Blue Waters conference approximately two years ago, is shown in Fig. 1, where a near-ideal scaling is now obtained for a 128 times increase in the number of processors for hypersonic flow over a hemisphere with 96 million computational particles (with X=32 for the blue and purple lines). The result presented two years ago (green line with X=128) achieved poor parallelization, and exhibited very long run times well. In addition, 87% weak scaling was obtained for 8,192 processors for a hemisphere flow with 24 billion particles.

RESULTS & IMPACT

We have recently simulated an even more computationally challenging flow over a double wedge using 768 nodes, with 24 billion particles and 1.8 billion computational cells. The double wedge geometry has the forward and aft angles of 30° and 55°,




Figure 2: Shock–shock interactions for hypersonic flow over a double wedge at an early time in the evolution of the unsteady flow.

with a length of 50.8 mm for the first wedge. The flow conditions corresponded to a free stream Mach number, velocity, static pressure, and unit Reynolds number of 7.14, 3,812 m/s, 0.78 kPa, and 52,200, respectively. For this simulation shown in Fig. 2, the VHS model and the majorant frequency scheme were used and gas–surface interactions were modeled using the Maxwell’s model with full momentum and energy accommodation.

Fig. 2 shows the resolution of the highly complex shock physics that occurs in this domain of hypersonic shock–boundary layer interaction flows. An oblique shock is formed at the leading edge of the lower wedge and a bow shock forms in front of the upper wedge. As the flow moves up the lower wedge, an adverse pressure gradient is formed and the boundary layer separates, which generates a separation shock. These three shocks meet at the triple point, and with their interaction a transmitted shock (also known as the reattachment shock) is formed. The reattachment shock impinges on the upper wedge and the thickness of the boundary layer reduces at the impingement point. This point is usually where the boundary layer reattaches to the surface. The zone embedded among the hinge, separation point, and the reattachment point is the separation zone where recirculation is observed. Note that the flow downstream of the bow shock is subsonic, but the flow going through the oblique, separation, and reattachment shocks remains supersonic. Therefore, a contact surface, also known as the shear layer, is formed across which the pressure is constant; however, the velocities and temperatures are different. The reattachment shock, after reflecting off the boundary layer, forms expansion

waves, which reflect at the contact surface and coalesce to form a compression wave. Even though this flow is still laminar, it is unsteady, and future calculations on Blue Waters will enable us to use global linear stability analysis to understand the different excitation and decaying eigenmodes as well as the nature of the disturbance in terms of acoustic, entropy, or vortical waves [2].

WHY BLUE WATERS

The time-accurate, large-scale DSMC simulations performed to obtain the results shown in Fig. 2 require on the order of 100,000 node-hours to reach a steady state of 1 millisecond. These are the first simulations that enable the hypersonic laminar shock boundary layer interaction community to understand the role of slip in a fully 3D simulation. In addition, we have made extensive use of the CPMAT and Perf-tools profiler on Blue Waters for testing these algorithmic improvements.

PUBLICATIONS AND DATA SETS

[1] Sawant, S., O. Tumuklu, R. Jambunathan, and D. A. Levin, Novel Use of AMR Unstructured Grids in DSMC Compressible Flow Simulations. *AIAA Aviation Forum* (AIAA, Denver, Colo., June 5–9, 2017).

[2] Tumuklu, O., D. A. Levin, and V. Theofilis, On the temporal evolution in laminar separated boundary layer shock-interaction flows using DSMC. *55th AIAA Aerospace Sciences Meeting* (AIAA, Grapevine, Texas, January 9–13, 2017).




Figure 1: Strong scaling performance of the improved SUGAR DSMC code.

HIGH ENERGY PHYSICS ON BLUE WATERS

Allocation: NSF PRAC/21,780 Knh

PI: Paul Mackenzie¹

Co-PIs: Norman Christ², Robert Sugar³

Collaborators: Alexei Bazavov⁴, Peter Boyle⁵, Carleton DeTar⁶, Steven Gottlieb⁷, Doug Toussaint⁸, Chulwoo Jung⁹, Yuzhi, Liu⁷

¹Fermi National Accelerator Laboratory

²Columbia University

³University of California, Santa Barbara

⁴Michigan State University

⁵University of Edinburgh

⁶University of Utah

⁷Indiana University

⁸University of Arizona

⁹Brookhaven National Laboratory


EXECUTIVE SUMMARY

The goal of this project is to carry out groundbreaking studies of the standard model of high-energy physics. The calculations we have in progress address basic questions in high-energy physics and are directly supportive of the very large, worldwide experimental studies in this field. This project is a joint effort of the Fermilab Lattice, MILC, and RBC Collaborations, which among them contain almost all of the high-energy physicists in the United States working on the numerical study of the standard model. The advent of petascale computers such as Blue Waters has had a transformational impact on our field. Members of our three groups have played a leading role in this transformation through the development of algorithms and community codes, and by carrying out petascale simulations. We are using Blue Waters to build upon these advances in order to make major progress in understanding the fundamental laws of physics.

RESEARCH CHALLENGE

The standard model of high-energy physics encompasses our current knowledge of the fundamental interactions of nature. It consists of two quantum field theories: the Weinberg–Salam theory of electromagnetic and weak interactions, and quantum chromodynamics (QCD), which is the theory of the strong interactions. The standard model has been enormously successful in explaining a wealth of data over the past forty years; however, high-energy physicists believe that a more general theory will be needed to explain physics at the shortest distances or highest energies. The research in which we are engaged aims at obtaining a deeper understanding of the standard model and at searching for physical phenomena that go beyond it.

Figure 1: Lattice data (squares) with a chiral-continuum-HQET fitting function. The solid curves are fit results for single lattice spacings and for the continuum limit. The quantity plotted on the x-axis is proportional to the strange quark mass, and those on the y-axis are proportional to the leptonic decay constants for the B and Bs mesons.



METHODS & CODES

QCD is formulated in the four-dimensional space–time continuum; however, in order to carry out numerical calculations, one must reformulate it on a lattice or grid. To obtain physical results, one carries out calculations for a range of small lattice spacings and then performs extrapolations to the zero lattice spacing (continuum) limit. This continuum extrapolation is one of the major sources of errors in lattice QCD calculations. Another important source of systematic errors arises because the calculations must take place in a finite box, whose physical size must be much larger than the largest lengths in the problem. Keeping both the continuum extrapolation and finite size effects under control requires working on very large lattices. The power of petascale computers is critical for enabling us to do this.

A number of different formulations of quarks, an elementary particle of matter in QCD, on the lattice are currently being used in numerical studies of QCD, all of which are expected to yield the same results in the continuum limit. We are using the two formulations most widely employed in the study of high-energy physics: domain wall fermions (DWF) and highly improved staggered quarks (HISQ). The DWF and HISQ actions each has important advantages for different aspects of our work. DWFs have nearly exact chiral symmetry at finite lattice spacings. This high degree of chiral symmetry is required for key studies of kaon decays that lead to precise tests of the standard model. On the other hand, staggered quarks are essential for studies of the decays and mixings of particles with heavy quarks for which chiral symmetry plays a less important role. However, the large lattice volumes, made accessible by the lower computational cost of staggered fermions, are necessary for accurate control of finite lattice spacing errors. The HISQ calculations are aimed at precise determination of some of the least well-known parameters of the standard model, and at making further precise tests of it.

Lattice QCD calculations proceed in two steps. In the first, one uses importance-sampling techniques to generate gauge configurations, which are representative samples from the Feynman path integrals that define QCD. These configurations are saved, and in the second step they are used to calculate a wide variety of physical quantities. Generating gauge configurations is the rate-limiting step and requires the most capable supercomputers available.

RESULTS & IMPACT

During the first year of our PRAC grant, we have used our allocation to generate DWF and HISQ gauge configurations that are among the most challenging produced to date. The initial applications of the DWF configurations will be to study two processes that are highly suppressed in the standard model and therefore favored as places where physics beyond the standard model may emerge. The first of these is the direct violation of CP (charge conjugation parity) symmetry in the decay of neutral kaons. The second is the determination of the mass difference

between the two neutral kaon-decay eigenstates, which is the smallest particle mass difference ever measured.

The first application of the HISQ gauge configurations generated under this allocation has been to enhance the determination of quark masses and leptonic decay constants of unprecedented precision.

One particularly striking result is the determination of the leptonic decay constant of the B meson to a precision of 0.4%, a factor-of-five reduction of uncertainty from the world average. This calculation is illustrated in Fig. 1. A preliminary report on this work was published in [1]. Near-final results were presented at the Lattice 2017 conference, and a journal article is in progress.

WHY BLUE WATERS

Work on lattice QCD calculations has 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 to obtain a quantitative understanding of physical phenomena controlled by the strong interactions. This progress has been enabled by the advent of petascale computers, such as Blue Waters, and could not have been made without them.

PUBLICATIONS AND DATA SETS

Komijani, J., et. al, Decay constants f_B and f_{B_s} and quark masses m_b and m_c from HISQ simulations. in *Proceedings of the 34th International Symposium on Lattice Field Theory* (Lattice, Southampton, UK, July 24–30, 2016).

FIRST-PRINCIPLES COMPUTER SIMULATIONS OF HYDROCARBONS UNDER FUSION CONDITIONS

Allocation: NSF PRAC/9,200 Knh

PI: Burkhard Militzer¹

Co-PIs: Shuai Zhang^{1,2}, Francois Soubiran¹, Kevin Driver^{1,2}

¹University of California, Berkeley

²Lawrence Livermore National Laboratory

EXECUTIVE SUMMARY

Matter in the interiors of giant planets and stars is exposed to conditions of extreme temperature and pressure. In addition to the astrophysical relevance, a rigorous and consistent theoretical description of materials properties in the warm dense matter and dense plasma regimes has been identified as a central goal in the development of key energy technologies such as advanced nuclear reactors and inertial confined fusion, shock physics, plasma science, and stockpile stewardship. The optimal design of new plasma experiments relies on computational models of the equation of state, transport, and optical properties in order to achieve desired pressures and temperatures and to make the first, key measurements. Here we use first-principles computer simulations to study the properties of hydrocarbons at extreme conditions.

RESEARCH CHALLENGE

The development of a first-principles methodology for warm dense matter (WDM) applications that treat temperature effects consistently is a key component of the stewardship of plasma science [1]. Indeed, technological progress in high-energy-density




Figure 1: Temperature–density conditions of our path integral Monte Carlo and density functional molecular dynamics simulations of CH plastic materials. The red curves show the predicted path of fusion experiments. The solid blue and dashed black lines, respectively, show isobars and isentropes that we derived from our equation of state.

physics (HEDP) applications, such as fusion energy [2], shock-wave physics [3], astrophysical processes [4], and planetary [5] and stellar interiors, relies on simulations for input and guidance. WDM is broadly described as the HEDP regime between condensed matter and ideal plasmas, where strong electron correlation, and quantum and ionization effects, are all important.

Hydrocarbons are the primary materials used for the ablator in inertial confinement fusion target capsules. It follows that the determination of the correct equation of state (EOS) of hydrocarbon ablators is important in order to optimize experimental designs to achieve desired density and temperature conditions. In this work, we perform state-of-the-art, benchmark-quality EOS calculations by combining high-temperature path integral Monte Carlo (PIMC) data and low-temperature density functional theory molecular dynamics (DFT-MD) data to construct coherent EOSs for several hydrocarbon materials over a wide range of densities and temperatures. At a commonly accessible temperature of 10^6 K, we find PIMC and DFT-MD predict consistent internal energies and pressures, validating their accuracy. While there have been several previous simulations of CH reaching into the WDM regime based on DFT-MD methods alone, our calculations provide the first set of first-principles data across the entire WDM regime that properly treats the many-body and shell-structure ionization effects. These processes affect the structure of the Hugoniot curve, which represents the density that can be achieved with various shock wave experiments. Our simulation results provide a benchmark for widely used EOS tables such as SESAME and QEOS as well as orbital-free DFT and average-atom methods. Our calculations are also timely for interpretation of ongoing spherically converging shock experiments on the Gbar platform at NIF and OMEGA.

METHODS & CODES

Since we need to cover a large temperature interval that spans five orders of magnitude (10^4 – 10^9 K), we rely on two different first-principles simulation methods (see Fig. 1). To study the high-temperature regime, we focus on the development of the PIMC method [6], which naturally incorporates finite temperature quantum effects by working within the many-body thermal density matrix formalism. The combination of Feynman’s imaginary time path integrals and efficient Monte Carlo sampling techniques makes this approach one of the most appropriate first-principles

simulation techniques for quantum systems at high temperature, T . Since the length of the path scales like $1/T$, the method becomes increasingly efficient for high temperatures. Electrons and nuclei are often treated equally as paths, but here we treat the nuclei classically because their zero-point motion is negligible for the temperatures under consideration. All PIMC simulations were performed with our own code, CUPID [7].

For the low-temperature part of the WDM regime, DFT-MD is an accurate and efficient first-principles simulation method for these conditions. The thermal occupation of electronic states is treated as a perturbation of the ground state by Fermi–Dirac smearing. The main drawback of this method is that it becomes computationally infeasible as electrons occupy a large number of bands at high temperature, which is why we switch to PIMC simulation at high temperature (see Fig. 1). All DFT-MD simulations were performed with the VASP code [8].

RESULTS & IMPACT

We performed an entirely first-principles determination of hydrocarbon mixtures in the WDM regime by including all nonideal effects. Based on PIMC and DFT-MD, we obtained coherent sets of EOS over a wide range of density and temperature conditions and derived the shock Hugoniot curves of a series of hydrocarbon materials [9]. For polystyrene, we predict a maximum shock compression ratio of 4.7, whereas earlier estimates range from 4.3–4.7. Our calculated shock Hugoniot curve (Fig. 1) agrees very well with experimental measurements and provides guidance for the interpretation of experiments on the Gbar platform at NIF. We observe a single compression maximum for hydrocarbon materials while there are two compression maxima in the Hugoniot curve of nitrogen, oxygen, and neon. We have shown that this difference is related to the properties of the L-shell ionization, which is much more gradual for carbon. We found that the linear isobaric–isothermal mixing approximation works very well, resulting in a discrepancy in the density of CH of 1% or less under stellar core conditions. This implies that it is sufficient to derive only accurate EOS tables for the end members in order to provide a thermodynamic description of deep stellar interiors. In the past, models for stellar and giant planetary interiors relied on semianalytical models and experimental results for a few key conditions. Through high-performance computation, it has become possible to provide a more rigorous theoretical description of matter at extreme conditions. While, in the past, the characterization of one material has taken a Ph.D. thesis [5], on HPC systems like Blue Waters many materials can now be characterized in parallel within a single year.

WHY BLUE WATERS

Our Blue Waters allocation is one order of magnitude larger than any other allocation that we have obtained elsewhere. This enabled us to perform simulations with more accurate fermion nodes that are more realistic but also much more expensive. Besides the hydrocarbons discussed in this report, we performed simulations




Figure 2: CH plastic at extreme pressure–temperature conditions studied with atomistic simulations. The C and H atoms are denoted as blue and pink spheres. Their interactions are governed by myriad short-lived chemical bonds that are mitigated by hot electrons. Their density is illustrated by the yellow isosurface.

for a variety of other materials, including sodium, lithium fluoride, and silicon. In each case, we studied many temperature–density conditions as illustrated in Fig. 1. Hydrocarbons, however, posed a special challenge because they are produced for a variety of chemical compositions. Because we had access to the Blue Waters system, we could meet the composition challenge rigorously. In addition to a typical 2D parameter scan in density–temperature space, we performed 3D parameter scans in density–temperature–composition space by performing simulations of C, C_2H , CH, CH_2 , CH_3 , CH_4 as well H.

PUBLICATIONS AND DATA SETS

Zhang, S., K. P. Driver, F. Soubiran, and B. Militzer, First-principles equation of state and shock compression predictions of warm dense hydrocarbons. *Phys. Rev. E*, 96:1 (2017), DOI:10.1103/PhysRevE.96.013204.

Driver, K. P., and B. Militzer, First-principles simulations of warm dense lithium fluoride. *Phys. Rev. E*, 95:4 (2017), DOI: 10.1103/PhysRevE.95.043205.

Driver, K. P., F. Soubiran, S. Zhang, and B. Militzer, Comparison of Path Integral Monte Carlo Simulations of Helium, Carbon, Nitrogen, Oxygen, Water, Neon, and Silicon Plasmas. *J. High Energy Density Physics*, 23 (2017), DOI: 10.1016/j.hedp.2017.03.003.

Zhang, S., K. P. Driver, F. Soubiran, and B. Militzer, Equation of state and shock compression of warm dense sodium—A first-principle study. *J. Chem. Phys.*, 146:7 (2017), DOI: 10.1063/1.4976559.

Hu, S. X., et al., First-Principles Prediction of the Softening of the Silicon Shock Hugoniot Curve. *Phys. Rev. B*, 94:9 (2016), DOI: 10.1103/PhysRevB.94.094109.

Zhang, S., K. P. Driver, F. Soubiran, and B. Militzer, Path Integral Monte Carlo Simulations of Warm Dense Sodium. *J. High Energy Density Physics*, 21 (2016), DOI:10.1016/j.hedp.2016.09.004.

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

Allocation: NSF PRAC/3,870 Knh

PI: Sarma L. Rani¹

Collaborator: Rohit Dhariwal¹

¹University of Alabama in Huntsville

EXECUTIVE SUMMARY

The overall objective of our research is to investigate the role of turbulence in driving the relative velocities and positions of inertial particles in isotropic turbulence. During the second year of this PRAC grant, we investigated the relative motion of high-inertia particle pairs in isotropic turbulence. We performed direct numerical simulations (DNS), as well as Langevin simulations (LS) based on a probability density function (PDF) kinetic model for pair relative motion. We developed a stochastic theory that involved deriving closures in the limit of high Stokes numbers for the diffusivity tensor in the PDF equation for particle pairs. Quantitative analysis of the stochastic theory was performed through a comparison of the particle pair statistics obtained using LS with those from DNS. The high-performance computing (HPC) resources of the Blue Waters system were invaluable and indispensable in performing the DNS and LS runs needed to validate the stochastic theory.

RESEARCH CHALLENGE

Turbulence-driven relative motion of high-inertia particles is relevant in astrophysical scenarios, such as the interstellar medium, protoplanetary disks, and the atmospheres of planets and dwarf stars. Specifically, the “sticking” of dust particles in protoplanetary disks is believed to be the mechanism for planetesimal formation. An intriguing question that astrophysicists are investigating concerns the effects of turbulence on the dispersion, sedimentation, collisional coalescence, and fragmentation of dust grains. The viscous relaxation times, τ_v , of these particles are significantly large, with estimated $St_\eta \sim 10$ -100, where $St_\eta = \tau_v/\tau_\eta$ is the Stokes number based on the Kolmogorov time scale τ_η .

The two principal quantities describing the relative motion of inertial particles in a turbulent flow are: (1) radial distribution function (RDF), which is a measure of the spatial clustering of particles, and (2) PDF of pair relative velocities, which is a key input to the particle-pair encounter rate. The RDF and the relative velocity PDF are both key inputs to the particle collision kernel, and depend sensitively on the Stokes number St_η .

Recently, we developed a stochastic theory for the relative velocities and positions of high-inertia pairs in forced isotropic turbulence [1]. The theory involved deriving a closure for the diffusivity tensor characterizing the relative-velocity-space diffusion current in the PDF kinetic equation of particle-pair separation and relative velocity. Since we had considered the

$St_\eta \gg 1$ limit, the pair PDF equation is of the Fokker–Planck form (St_η is the Stokes number based on the integral time scale of turbulence). Using the diffusivity formulation, one can perform Langevin simulations of pair relative velocities and positions, which is equivalent to simulating the Fokker–Planck equation.

In this context, the current study has two main objectives. First, we perform a quantitative analysis of the three forms of the diffusivity derived in [1]. The second objective is to compute the relative motion statistics of particle pairs using both DNS and LS, and compare the corresponding results.

METHODS & CODES

DNS of forced isotropic turbulence were performed using a discrete Fourier-expansion-based pseudospectral method. Simulations were performed over a cubic domain of length 2π discretized using N^3 grid points, with periodic boundary conditions in all three directions. The fluid velocity is advanced in time by solving the Navier–Stokes equations in rotational form, as well as the continuity equation for an incompressible fluid. Direct evaluation of the nonlinear convective terms in the Navier–Stokes equations is extremely computationally intensive. Hence, a pseudospectral approach is adopted wherein the nonlinear terms are first computed in physical space and then transformed into the spectral space. The P3DFFT library [2] is used to carry out the transforms between physical and spectral spaces.

RESULTS & IMPACT

The RDF is a well-established measure of particle clustering. In fig. 1, the RDF is presented as a function of St_η at four separations $r/\eta = 6, 12, 18,$ and 24 (η is the Kolmogorov length scale). The results from LS are compared with the data from the DNS performed in the current study, the Février, et al. [3] DNS, and also with the results from the Zaichik and Alipchenkov [4] theory. The Février, et al. [3] data were for $Re_\lambda = 69$, while the current DNS data are for $Re_\lambda = 76$. There is excellent agreement between the LS RDF and the two sets of DNS RDFs at all four separations, particularly for $St_\eta > 10$. The Zaichik and Alipchenkov [4] theory significantly overpredicts the RDFs for high Stokes numbers at all separations. The current DNS study, as well as the stochastic theory, provided the basis for the first-ever investigation of the validity of the Zaichik and Alipchenkov [4] theory in the high-Stokes-number limit.

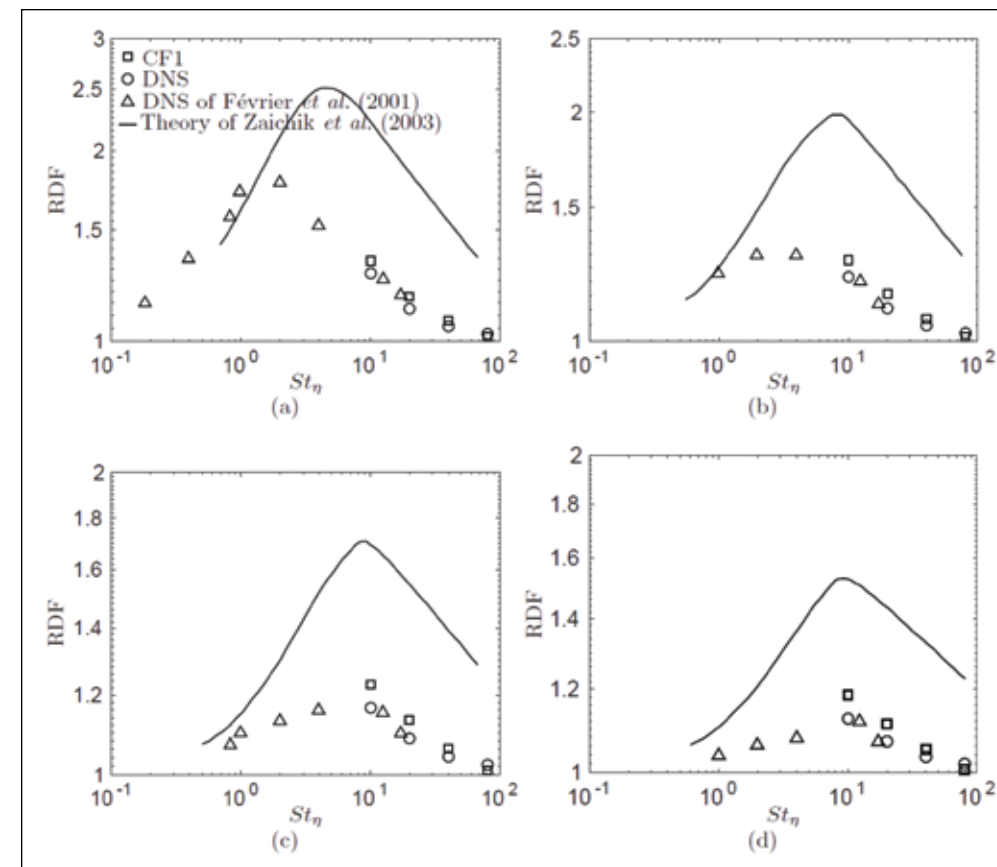


Figure 1: Radial distribution function as a function of St_η , (a) $r/\eta = 6$, (b) $r/\eta = 18$. Squares and circles represent CF1 (current stochastic theory) and current DNS at $Re_\lambda = 76$; triangles represent DNS at $Re_\lambda = 69$ from [3]. Solid line represents data from [4] for $Re_\lambda = 69$.

WHY BLUE WATERS

Direct numerical simulation is the most accurate numerical approach to resolve all the temporal and length scales in a turbulent flow. However, DNS of particle-laden turbulent flows are computationally very intensive, since in DNS the cost of a simulation scales as Re^3 , where Re is the Reynolds number. The overall computational objective of this project is to simulate particle-laden isotropic turbulence at Taylor micro-scale Reynolds number $Re_\lambda \sim 600$ with grid sizes $\sim 2048^3$. This will require running our code on tens of thousands of cores. Also, each DNS run is expected to generate several terabytes of data. Due to these central processing unit time and storage requirements, the Blue Waters supercomputer is the ideal platform to achieve our objective. It would be relevant to mention that during the last year, Blue Waters proved to be an invaluable resource in computing key inputs to our stochastic theory. For instance, the theory requires as an input the two-time Eulerian correlations of fluid relative velocities seen by particle pairs. Evaluation of the two-time correlation for nearly half a trillion pairs is a highly computationally intensive process. We were only able to compute this quantity because of the Blue Waters access, where we ran the code on 625 nodes using 20,000 cores.

PUBLICATIONS AND DATA SETS

Dhariwal, R., S.L. Rani, and D.L. Koch, Stochastic Theory and Direct Numerical Simulations of the Relative Motion of High-Inertia Particle Pairs in Isotropic Turbulence. *Journal of Fluid Mechanics*, 813 (2017), pp. 205–249.

Dhariwal, R., S.L. Rani, and D.L. Koch, Effects of Deterministic and Stochastic Forcing Schemes on Inertial Particle Statistics in DNS of Isotropic Turbulence. *69th Annual Meeting of the Division of Fluid Dynamics* (American Physical Society, Portland, Ore., November 20–22, 2016).

Rani, S.L., R. Dhariwal, and D.L. Koch, Comparison of Stochastic Theory and DNS for the Relative Motion of High-Inertia Particle Pairs in Isotropic Turbulence. *69th Annual Meeting of the Division of Fluid Dynamics* (American Physical Society, Portland, Ore., November 20–22, 2016).

MAPPING PROTON QUARK STRUCTURE USING PETABYTES OF COMPASS DATA

Allocation: Illinois/200 Knh

PI: Caroline Riedl¹

Co-PIs: Matthias Grosse Perdekamp¹, Naomi Makins¹

Collaborators: Vincent Andrieux¹, Robert Heitz¹, Marco Meyer¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

The COMPASS experiment at CERN probes proton substructure by scattering high-energy pion and muon-beams off of nuclear targets to measure the momentum and coordinate phase space of quarks inside the proton. Observing correlations between proton spin and the intrinsic transverse momentum of quarks will shed light on the quark dynamics inside the proton and will provide a critical test of fundamental predictions derived from quantum chromo dynamics, the quantum field theory describing the nuclear force. The measurements will produce 10 petabytes of experimental and simulated data. Blue Waters' balance of processing capabilities with data storage and handling is well suited for the analysis of the large COMPASS data samples as these require significant algorithmic processing per pion/muon-proton-scattering event. In addition to raw data processing and physics-level analysis, Blue Waters allows for detailed simulation of COMPASS detector properties and environmental effects.

RESEARCH CHALLENGE

Observation of the sign change of the Sivers quark distributions ("Sivers functions") in Drell–Yan scattering compared to existing measurements in semi-inclusive deep-inelastic scattering (SIDIS) is one of the few performance Nuclear Science Advisory Committee (NSAC) [1] milestones for Department of Energy- and National Science Foundation-funded research in nuclear physics. Such

measurement requires polarization-dependent Drell–Yan data. The 2015 and 2018 Drell–Yan runs of the COMPASS experiment at CERN constitute the first measurements of this kind: The negatively charged pion beam from the Super Proton Synchrotron (SPS) was impinged on a target of transversely polarized protons. Sivers functions arise from correlations between proton spin and quark transverse momentum and thus appear connected to quark orbital motion inside the proton. The experimental measurement of these correlations will provide the first tomographic images of quarks in the proton in transverse momentum space (Fig. 1).

With the 2016 and 2017 runs, COMPASS will add valuable observables in Deeply Virtual Compton Scattering (DVCS) to constrain spin-independent Generalized Parton Distributions in the so far unexplored kinematic domain between HERMES and the Jefferson Lab experiments on the one side, and the HERA collider experiments on the other side.

METHODS & CODES

For experimental data production, the raw COMPASS data collected at CERN are transferred to Blue Waters at average throughput speed of about 1 GB/s using the File Transfer System FTS3 [3], a bulk data mover created to distribute globally the Large Hadron Collider (LHC) data. For each triggered event in COMPASS, the information of the detectors is recorded by the Data Acquisition (DAQ) system. The COMPASS Reconstruction




Figure 1: Quark densities in the transverse momentum plane for a transversely polarized proton (y-direction). Deep red (blue) regions indicate large negative (positive) values. Figure from [2], projecting the impact of future measurements on the knowledge of quark densities in transverse momentum space.




Figure 2: COMPASS acceptance for pion, kaon, and anti-proton beams simulated by the Monte Carlo generator PYTHIA [6] adapted to pion and kaon beams versus the transverse momentum of the Drell–Yan particle pair. Different beam energies are indicated by different colors.

Analysis Library (CORAL) software performs the transition from raw data to physical quantities. CORAL's function is to reconstruct particle trajectories and momenta, and the position of vertices. The reconstructed information is stored in the form of Data Summary Trees, which are read and analyzed using the COMPASS PHysics Analysis Software Tools (PHAST). The production of Monte Carlo data is performed in three steps: (1) The generation of signal and background events is carried out with event-generator packages. (2) For the simulation of the detector response to the physics event, a GEANT4 [4] toolkit is used based on the description of the COMPASS apparatus. (3) Simulated hits are subjected to the same reconstruction CORAL and PHAST codes as experimental data.

RESULTS & IMPACT

COMPASS accumulates a raw experimental data set of about 0.75 petabytes per year. A first step in the data analysis is the conversion of raw data into the physical properties of the fundamental particles created in a collision event. This data production is an iterative process that requires two or three passes. Approximately 6.5 million CPU hours are needed for one data production pass, which requires usually about three calendar months given the available resources at CERN. In an exploratory mass production in 2016, 5% of a yearly COMPASS data set was processed in 15 hours [5]. Since the task is scalable, a full year can be processed in about the same time if sufficient computing resources are available. This project has recently been granted millions of node hours. As of the middle of 2017, the launching of full-scale data production on Blue Waters was in preparation.

Apart from the processed experimental data, simulated Monte Carlo data are an essential ingredient to the data analysis. Simulations of the detectors play a central role in understanding subtle detector effects and removing background events from the data sample. With the available resources at CERN and

collaborating institutions, the CPU-intensive part of the Monte Carlo—the simulation of the detector properties with GEANT—can often not be afforded for extensive studies, for example event pile-ups or detector efficiencies in a fine time binning. The production of simulated data for studies on a future fixed-target experiment at CERN using radio-frequency separated pion, kaon, and anti-proton beams has begun (Fig. 2).

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

WHY BLUE WATERS

Requiring both substantial Monte Carlo data production and up to three data campaigns in parallel, a timely analysis of COMPASS data appears difficult. A delay of several years between the end of data-taking in November 2018 and publication of all COMPASS results appears likely. Given the present computing resources at CERN and collaborating institutions, COMPASS might lose out to competing experiments scheduled to run in the future at FNAL, RHIC, and Jefferson Laboratory. The current projected schedule for data analysis also provides challenges to sustaining support from funding agencies for extended periods of times and for on-schedule completion of doctoral theses based on COMPASS data. With the petascale resources of Blue Waters, COMPASS experimental and Monte Carlo data can be processed significantly faster, and in the case of simulated data, also can be generated in greater detail.

OPTICAL DETERMINATION OF CRYSTAL PHASE IN SEMICONDUCTOR NANOCRYSTALS

Allocation: Blue Waters Professor/190 Knh; Illinois/550 Knh

PI: André Schleife¹

Collaborators: Sung Jun Lim², Andrew Smith¹

¹University of Illinois at Urbana-Champaign

²Daegu Gyeongbuk Institute of Science and Technology

EXECUTIVE SUMMARY

Semiconductor nanocrystals are highly interesting and diverse material systems with unique electronic and optical properties that are even tunable as to size, shape, composition, and internal structure. Currently, the fundamental connection between crystal structure and optical properties is not well understood, hampering wide-spread and efficient use of these nanocrystals in societally important applications such as light-emitting diodes, solar cells, bioimaging, and consumer electronics. In order to establish a connection between structure and optical properties, we use a combination of cutting-edge experimentation and first-principles theoretical spectroscopy. Our work shows that by understanding the absorption spectrum at energies well above the absorption onset, it is possible to efficiently distinguish optically between two different crystal structures of CdSe materials, wurtzite and zincblende, both in bulk crystals and nanocrystals. This may rapidly accelerate development of these materials by allowing fast and accurate structural characterization with small quantities as well as samples in liquids.




Figure 1: Visualization of wave functions of electronic energy states in bulk WZ CdSe at the kappa point (left) and the gamma point (right) of the Brillouin zone. Atoms and bonds are shown in ball-and-stick format with Cd in black and Se in grey. Isosurfaces are shown in red for the top 95% of the squared wavefunction to depict high localization and the lowest 5% in blue to depict nodes.

RESEARCH CHALLENGE

Semiconductor nanocrystals are diverse material systems currently used in light-emitting devices, solar cells, bioimaging, and consumer electronics. Their unique electronic and optical properties and the fact that these properties are tunable by changing the size, shape, composition, and internal structure of the nanocrystals, make them highly interesting material systems to a broad community of researchers and diverse commercial sectors from optical devices to medical diagnostics. Among the most fundamental questions is the *connection* between crystal structure and optical properties, as answering that question would allow control of functionality through structural properties that are accessible to synthesis. Despite highly interesting implications for materials design of semiconductor nanocrystals, this connection is not fully understood today, hampering a more widespread and more efficient use of semiconductor nanocrystals for many of the above-mentioned applications.

Establishing such a connection is extremely challenging since, in experiment, clarifying the crystal structure typically relies on X-ray crystallography. This technique is, however, not well-suited for samples in solution, for process analysis, for micro-scale reactions, or for high-throughput synthesis. In addition, X-ray diffraction patterns are less clear for very small or mixed-crystal nanocrystals. Application of fast and reliable optical techniques is highly desirable but is, so far, impossible since spectral features in the vicinity of the optical absorption onset are too similar across different crystal structures to allow reliable distinction.

METHODS & CODES

In order to extend the spectral range of interest beyond the absorption onset, we use a combination of cutting-edge experimentation and first-principles theoretical spectroscopy. This allows us to clarify important structure–property relationships for CdSe nanocrystals in two different crystal structures, wurtzite and zincblende. Here, we focus on the computational aspects of the project: We use many-body perturbation theory to explore the influence of quasiparticle effects on optical properties of CdSe. More specifically, single-quasiparticle effects are described using the HSE06 hybrid exchange–correlation functional. From the solution of the Bethe–Salpeter equation for the optical polarization function we conclude that excitonic effects are small in bulk CdSe due to the large screening of the electron–electron interaction.

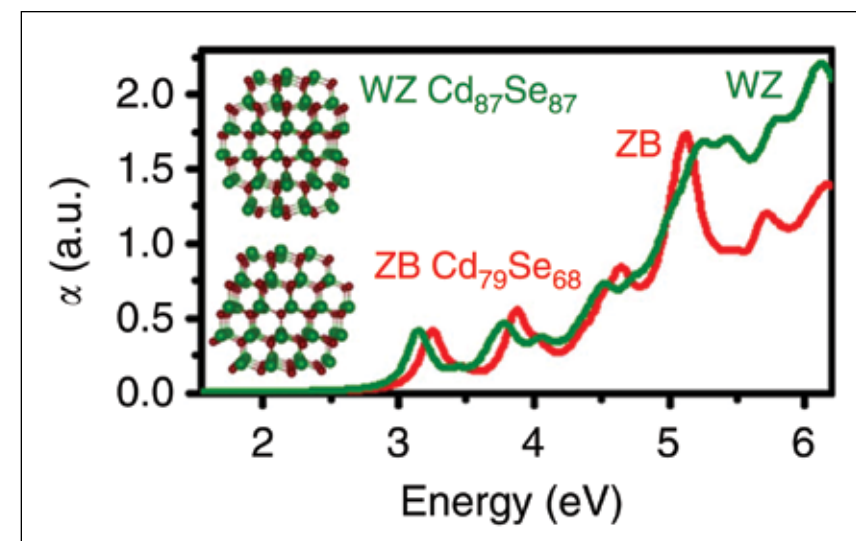


Figure 2: Absorption coefficient simulated using density functional theory for zincblende-phase (red) and wurtzite-phase (green) CdSe nanocrystals. Image credit: see [6].

This allows us to resort to a computationally more affordable density–functional theory description to extend this study toward optical properties of much more intricate hydrogen-passivated CdSe nanocrystals in both crystal structures. We use density functional theory to study both crystal structures and mixed-crystal nanocrystals where the wurtzite and zinc-blende phase are separated by a stacking fault. All calculations described here are carried out using the Vienna Ab Initio Simulation Package [1–3] as well as the Bethe–Salpeter equation implementation discussed in [4] and [5].

RESULTS & IMPACT

We report the experimental and computational identification of unambiguous optical signatures of cubic and hexagonal phases in II–VI CdSe nanocrystals: The main result of this work is that it is possible to efficiently distinguish between wurtzite and zincblende CdSe, both in bulk and nanocrystal form, *exclusively* using optical absorption spectra. Both computation and cutting-edge experimentation clearly show that it is key not only to focus on the absorption onset: While the onsets look very similar between both polymorphs, our work shows that high-energy spectral features in the 4–6 eV energy range are suitable to rapidly identify phase even in nanocrystals as small as about 2 nm. We found that this not only allows distinguishing the crystal phase of bulk CdSe but even to semi-quantitatively analyze wurtzite and zincblende contributions to nanocrystals that contain both structure types, separated by a stacking fault. Furthermore, the first-principles calculations carried out on Blue Waters explain the band-structure origin of these spectral features for both crystal structures. Our results pave the way for applying accelerated experimentation, possibly even high-throughput experiments, to nanocrystal samples in solution. Experiments carried out for this work additionally show that important ligands studied here are not affecting these conclusions.

WHY BLUE WATERS

Computing converged optical-absorption spectra of semiconductor nanocrystals is a computationally expensive task: Each nanocrystal consists of about 300 atoms, including passivating hydrogen atoms, and needs to be surrounded by a sufficiently large region of vacuum. In addition to densely sampling the Brillouin zone, a large number of valence and conduction bands needs to be included in order to achieve convergence of the computed optical spectra. These requirements lead to large matrices for the Kohn–Sham wave functions and solving the Kohn–Sham equations for such large systems can only be achieved using well-parallelized codes on fast supercomputers with fast interconnects. Such calculations are not a use case for the hardware infrastructures provided by cloud computing and, hence, are not possible on such high-throughput resources.

Instead, these calculations are exactly what Blue Waters is developed for: We benefitted from the fast CPUs, the large amount of memory per compute core, and extremely fast interconnects between nodes. Continuous access to such machines is absolutely critical for our research and cannot be accomplished by cloud services in the foreseeable future. Furthermore, Blue Water staff, e.g., Victor Anisimov, helped us identify and fix performance bottlenecks when writing and reading large wave function files, as was required especially in the early stages of this project. Only this combination of hardware and staff resources enabled the computational part of this research project.

PUBLICATIONS AND DATA SETS

Jun Lim, S., A. Schleife, and A. M. Smith, Optical determination of crystal phase in semiconductor nanocrystals. *Nat. Comm.*, 8 (2017), DOI: 10.1038/ncomms14849.

ADVANCED DIGITAL TECHNOLOGY FOR MATERIALS AND MANUFACTURING

Allocation: Exploratory/50 Knh

PI: Ahmed Taha¹

Co-PIs: Seid Koric¹, Sudhakar Parmidighantam^{1,2}, Narayan Aluru³, Gabrielle Allen^{1,3}

Collaborators: Ashraf Idkaidek³, Shantanu Shahane³, Fereshteh A. Sabet³, Ethan Shapera³

¹National Center for Supercomputing Applications

²Indiana University

³University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

This exploratory Blue Waters (BW) proposal provided computing resources to four graduate students funded by the National Center for Supercomputing Applications Materials & Manufacturing group, allowing the students to explore how their research can be furthered through the use of high-performance computing to address large-scale problem-solving. Two of the research activities that used the BW allocation, both involving analysis of bone, are reported here.

The first project, titled “Simulation of Reference Point Indentation on Cortical Bone,” was conducted by Ashraf Idkaidek. He used two different instruments—BioDent and Osteoprobe—that utilize the Reference Point Indentation (RPI) technique. The second project, “Mechanics of materials with focus on accelerated design and structure-processing-property relations of materials via high scale computations” was conducted by Fereshteh A. Sabet. This project examined and compared the performance of implicit and explicit solvers for modeling trabecular bone using Abaqus.

Project 1 - Simulation of Reference Point Indentation on Cortical Bone

RESEARCH CHALLENGE

Bone has a hierarchical architecture ranging from atomistic to macroscopic scales. At the scale of one to a few millimeters, the bone tissue is composed of cortical and trabecular bone. Osteoporosis is a bone disease characterized by low bone density, which leads to an increased risk of fractures that occur mainly in trabecular bone. Trabecular bone is also the primary site for insertion of orthopedic implant systems. Thus, the mechanical properties of trabecular bone are of great clinical and research interest for prediction of age and disease-related fractures as well as for designing improved implant systems.

METHODS & CODES

Modeling of trabecular bone entails highly nonlinear mechanical behavior along with contacts. As a result, it is of considerable interest to assess the effectiveness and efficiency of explicit solution methods. In this project, we used the implicit and explicit solvers of Abaqus to analyze micro-computed tomography (micro-CT) finite element (FE) models of trabecular bone, and compared the performance of the two solvers.

RESULTS & IMPACT

Our results show that there is a good match between micro-CT FE model results using implicit and explicit solvers (see Fig. 1). We also observed that implicit and explicit solvers scale similarly, but the explicit solver performs five times faster.

WHY BLUE WATERS

We were able to successfully scale our simulations on eight to 12 nodes on Blue Waters with the explicit solver, which significantly saved computational time. Each of our models has many millions of degrees of freedom and nonlinearities, making such models impossible to solve without the use of a supercomputer.




Figure 1: Project 1—(a) Example of von Mises stress distribution in samples under uniaxial compression, (b) implicit and explicit solver results versus experimental stress-strain curves.

Project 2 - Mechanics of materials with focus on accelerated design and structure-processing-property relations of materials via high-scale computations

RESEARCH CHALLENGE

Osteoporosis is responsible for two million broken bones at a cost of \$19 billion, annually in the United States. According to the National Osteoporosis Foundation, by 2025 this disease is expected to lead to three million fractures at a cost of \$25.3 billion per year. Assessing bone material properties in relation to its fracture resistance is important for the diagnosis and treatment of bone diseases. Using traditional materials testing approaches to measure the mechanical properties of bone, such as compression, tension, and three-point and four-point bending, are *ex vivo* and destructive.

METHODS & CODES

Cortical bone forms the outer hard shell of the whole bone. Therefore, understanding cortical bone fracture behavior is essential to evaluate fracture resistance of the complete bone. The Reference Point Indentation (RPI) technique was invented to allow *in vivo* evaluation of bone properties. Two different instruments use the RPI technique: BioDent and Osteoprobe. BioDent applies multiple indents at the same location on cortical bone, whereas Osteoprobe applies only one loading cycle at multiple neighboring locations on cortical bone. The relationship between RPI and bone properties has not been developed and is still an open topic.

In our research, we are focused on numerically relating both BioDent and Osteoprobe RPI instrument outputs to actual bone material mechanical properties. The cortical bone RPI simulation problem is highly nonlinear, where geometric nonlinearity,

material nonlinearity, and contacts must be accounted for in order to preserve the accuracy of simulation results. We use Abaqus software to simulate cortical bone RPI.

RESULTS & IMPACT

Relating each of the 10 outputs of the BioDent RPI instrument to the bone material properties has been established using the finite element method [1]. Also, simulating bone fracture using the extended finite element method on a single-osteone cortical bone sample has been evaluated and published [2]. A study to relate Osteoprobe RPI output to bone material properties and its fracture resistance is currently being developed.

WHY BLUE WATERS

Completing this study is fully dependent on the numerical finite element method. The problem is highly nonlinear, and multiple iterations are needed to be able to relate Osteoprobe device output to different bone mechanical properties. Each Osteoprobe RPI simulation iteration demands high computational power and time. Therefore, completing such study using the multi-core BW system is essential.

PUBLICATIONS AND DATA SETS

Idkaidek, A., V. Agarwal, and I. Jasiuk, Finite element simulation of Reference Point Indentation on bone. *Journal of the Mechanical Behavior of Biomedical Materials*, 65 (2017), pp. 574–583.

Idkaidek, A., and I. Jasiuk, Cortical bone fracture analysis using XFEM—case study. *International Journal for Numerical Methods in Biomedical Engineering*, 33: e2809 (2016), DOI:0.1002/cnm.2809.

TRANSIENT MULTIPHASE FLOW PHENOMENA AND DEFECT FORMATION IN STEEL CONTINUOUS CASTING

Allocation: Illinois/100 Knh

PI: Brian G. Thomas^{1,2}

Co-PIs: Seong-Mook Cho¹, Surya Pratap Vanka¹, Hyunjin Yang¹, Matthew Zappulla², Ahmed Taha³, Seid Koric^{1,3}

¹University of Illinois at Urbana-Champaign

²Colorado School of Mines

³National Center for Supercomputing Applications

EXECUTIVE SUMMARY

The objective of this project is to develop computational models to simulate transient multiphase flows and related phenomena, to apply them to gain an improved fundamental understanding of defect formation in continuous steel casting, and to find ways to further improve the process. Large-Eddy Simulations (LES) coupled with a Volume of Fluid (VOF) model were applied to track transient motion of the liquid mold flux/molten steel interface and slag entrainment into the molten steel pool during steady state continuous casting. In addition, the transport of argon bubbles in the molten steel and their capture into the solidifying steel shell were simulated using LES simulations coupled with a Lagrangian Discrete Phase Model for particle transport and particle capture criteria. Furthermore, the LES models were validated with facility measurements and applied to investigate optimal process conditions for the nozzle port angle, submergence depth of the nozzle, and Electro-Magnetic Braking (EMBr) field strength. These simulations on Blue Waters reveal deeper insights into defect formation during the continuous casting of steel and have enabled improved operation.

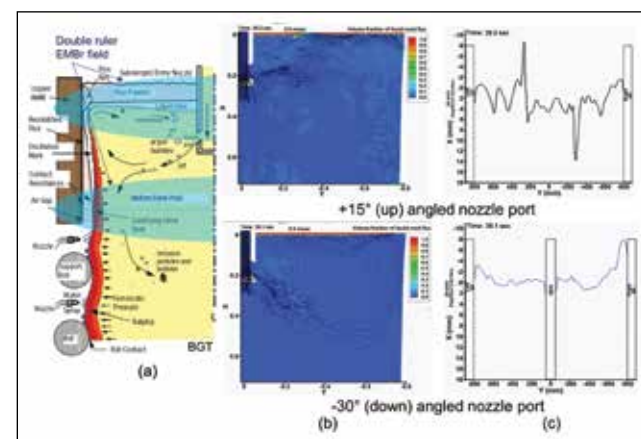


Figure 1: (a) Schematic of multiphase flow phenomena in the steel continuous-casting mold and Effect of nozzle port angle on (b) mold flow patterns, including motion of the top liquid mold flux layer, and (c) level variations at the liquid mold flux/molten steel interface in the meniscus region.

RESEARCH CHALLENGE

Continuous casting is the most widely employed solidification-process for steel manufacturing in the world [1], so even small improvements in this important process can lead to large benefits. Most defects in final products are related to transient multiphase flow phenomena in the mold region of the process (Fig. 1(a)). Severe instability at the liquid mold flux/molten steel interface can entrain some of the liquid mold flux (added on top of the molten steel pool in the mold to prevent steel oxidation) into the molten steel [2]. In addition, argon gas bubbles, injected to prevent nozzle clogging [3], can be trapped by the solidifying steel shell in the mold. To reduce these problems, transient multiphase flow phenomena should be understood, and process conditions should be optimized to reduce defect formation during continuous casting.

In this year's project, LES of several different important aspects of multiphase flow were performed to quantify the transient liquid mold flux/molten steel interface, transport of argon gas bubbles, and bubble capture into the steel shell in the mold during nominally steady continuous casting of steel slabs for different process conditions. The modeling results have been validated with plant measurements and applied to find optimal process conditions, including nozzle port angle, nozzle submergence depth, and EMBr field strength.

METHODS & CODES

LES coupled with VOF were applied to model transient molten steel flow and to track the liquid mold flux/molten steel interface. These models were implemented into the commercial package ANSYS Fluent on Blue Waters (BW) XE nodes. To calculate bubble transport and capture into solidifying steel shells with and without EMBr, LES coupled with Lagrangian particle capture (based on a force balance on each particle at the solidification front) [5] and MagnetoHydroDynamics models [8] using the GPU-based in-house code CUFLOW were employed on BW XK nodes.

RESULTS & IMPACT

Turbulent swirl flow from the upward-angled nozzle ports produces jet wobbling in the mold [4]. Sometimes, the jet impinges onto the liquid mold flux/molten steel interface and drags some of the liquid mold flux into the steel pool, resulting

in slag entrainment, as shown in Fig. 1(b). Most of the entrained slag becomes entrapped into the solidifying steel shell to form defects. Calculated interface instability reveals level variations greater than ~20 mm, especially at the meniscus region around its perimeter (Fig. 1(c)). These severe level instabilities can cause the liquid mold flux to touch the solidifying steel shell, and to be captured into the steel shell via meniscus hooks. However, the jet flow from the downward-angled nozzle ports with well-optimized casting conditions makes a classic double-roll pattern with less jet wobbling, resulting in better stability of the surface level and velocity in the mold. Velocity variations are smaller and the interface shows only ~2 mm fluctuations, so slag entrapment defects are drastically reduced.

Argon bubble motion is affected by turbulent jet flow in the mold, as shown in Fig. 2(a). Most bubbles larger than 3-mm in diameter float up toward the top surface due to their large buoyancy. On the other hand, small bubbles move along with the jet flow and easily reach the narrow face to be carried deep into the mold cavity. Many of the small bubbles move between the dendrites to be captured into the steel, especially without EMBr. However, EMBr slows and deflects the jet flow [6–7], so more bubbles float upward near the nozzle and fewer bubbles are transported to the narrow face and deep into the mold. As shown in Fig. 2(b), the bubble capture rate increases with time (until ~15-18 sec after gas injection for EMBr off and ~15 sec for EMBr on, because the jet flow path is shorter with EMBr). Then, after the flow has stabilized at nominally steady casting, EMBr is observed to reduce bubble entrapment significantly.

Parametric studies with these multiphase LES models have enabled better understanding of the complex multiphysics phenomena related to defect formation, including slag entrainment, inclusion and bubble transport, and capture. This has led to suggestions of nozzle geometry/casting condition combinations that lead to fewer defects, and, consequently, to significant savings to the steel plants.

WHY BLUE WATERS

Blue Waters enabled high-resolution multiphase flow simulations of the continuous caster needed for accurate predictions. In particular, the transient transport of small volume secondary-phases (entrained liquid mold slag and argon bubbles) in the turbulent flow require very small cells (~1 mm³) in a huge domain, and simulations of over 50-seconds flow time (with 0.001-second time steps). Furthermore, Blue Waters resources (both ANSYS Fluent HPC on XE nodes and our in-house multi-GPU code CUFLOW on XK nodes) showed speed-up breakthroughs (e.g., over 3000x with ANSYS Fluent HPC on BW) needed to provide this modeling capability for the steel continuous casting process.

PUBLICATIONS AND DATA SETS

Cho, S.-M., et al., Effect of Nozzle Port Angle on Transient Mold Flow, Surface Level, and Heat Transfer in Steel Caster. *CCC Annual Report*, Golden, Colo., August 16–18, 2016.




Figure 2: Effect of EMBr on (a) instantaneous argon bubble distributions and (b) bubble capture rate in the mold.

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. *Physics of Fluids*, 28:9 (2016), DOI: 10.1063/1.4961561.

Jin, K., S. P. Vanka, R. K. Agarwal, and B. G. Thomas, GPU Accelerated Simulations of Three-Dimensional Flow of Power-law Fluids in a Driven Cube. *Int. J. Computational Fluid Mechanics*, 31:1 (2017), pp. 1029–10257. DOI 10.1080/10618562.2016.1270449.

Jin, K., S. P. Vanka, and B. G. Thomas, Large Eddy Simulation of the Effects of EMBr and SEN Submerged Depth on Turbulent Flow in the Mold Region of a Steel Caster. *Metallurgical and Materials Transactions B*, 48:1 (2016), pp. 162–178, DOI: 10.1007/s11663-016-0801-z.

Jin, K., S. P. Vanka, and B.G. Thomas, Large Eddy Simulations of Argon Bubble Transport and Capture in Mold Region of a Continuous Steel Caster. *Proc. TFEC2017*, Las Vegas, Nev., April 2–5, 2017.

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

PI: Rafael Tinoco Lopez¹
Allocation: Illinois/200 Knh
Co-PI: Paul Fischer¹
Collaborators: Som Dutta¹, Pallav Ranjan¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Aquatic vegetation provides a wide range of services to the ecosystem: improving water quality through nutrient uptake and oxygen production, providing flood buffering and coastal protection services, and regulating erosion and deposition patterns, thus playing a paramount role in habitat creation and promotion of biodiversity. While vegetation–flow interactions have been studied extensively for unidirectional flows, much less is known about oscillatory conditions. The current study is geared toward increasing our understanding of the interactions among vegetation, flow, and sediment under oscillatory flows. Direct Numerical Simulations (DNS) and Large-Eddy Simulations (LES) through different arrays of idealized vegetation, represented as cylinders, are conducted using the higher-order spectral element-based computational fluid dynamics (CFD) solver Nek5000. Different arrangements and numbers of cylinders have been simulated in 2D and 3D, with the largest simulation having ~296 million computational points, using up to 32,768 MPI ranks.

RESEARCH CHALLENGE

Seagrasses are commonly referred to as “ecosystem engineers” due to their ability to modify and stabilize their environments [1]. They are a fundamental component of near-shore ecosystems, providing a wide range of services [2] ranging from increasing water quality through nutrient uptake and oxygen production,

creating habitats through spatial heterogeneity of the flow velocity, to dampening erosion on coastal wetlands. Past studies have focused mostly on unidirectional flows, relying strongly only on experimental approaches [3], with limited applications to oscillatory conditions. Such experiments provide ambient conditions closer to nature, although their measurements often lack the spatial and temporal resolution required to fathom the fundamental physical processes in detail. On the other hand, most numerical studies to date have primarily used CFD models based on temporal averaging of the Navier–Stokes equations, which approximate the turbulence in the system rather than accurately calculating it, and a few LES studies, which had to settle for a relatively small number of vegetation elements.

Our study is geared at bridging this gap by conducting numerical simulations at unprecedented scales, based on previous and ongoing experiments at the Ven Te Chow Hydrosystems Laboratory at the University of Illinois at Urbana-Champaign. We investigated flow through random and staggered arrays of cylinders to understand the effect of spatial heterogeneity of the vegetation on the flow. The study focuses primarily on oscillatory flow, though a few cases of unidirectional flow will be conducted for comparison purposes. Coupling the experimental and numerical study will yield further understanding of sediment dynamics under the influence of vegetation [4].




Figure 2: Results from 3D simulation of turbulent flow through a random arrangement of cylinders. About 296 million computational points are used. Instantaneous velocity magnitude at different elevations. Planes at (a) 0.5 %, (b) 1 %, (c) 10 %, (d) 50 %, (e) 75 %, and (f) 95 % water depth.

The scale of the experimental setup to be modeled is a challenge. The number of computation points required to model the whole domain is near 1.2 billion. While such simulations are still tractable on a petascale platform like Blue Waters, the computational cost is high, resulting in a reduction of the number of cases one can run, thus constraining the insights a broader range of parameters could yield. To increase variable space, a wide range of conditions are first simulated in 2D for the whole domain (~4 million computational points) to get an overview of the effect of different parameters such as Reynolds number, vegetation density, period and amplitude of the oscillatory flow, and spatial heterogeneity. Once the effects of different parameters are well understood, 3D simulations will be conducted for partial domains, big enough to accurately capture the general dynamics, but within manageable computational costs.

A better understanding of vegetation–flow dynamics will not only advance fundamental knowledge of physical processes but also will guide design efforts for scour protection and artificial wetlands. Part of the study is to characterize drag coefficients of the vegetation array, improving accuracy of reduced-order models of flow through vegetation. The study will also identify the ideal quantity and locations to place instrumentation in experiments through large arrays of cylinders.

METHODS & CODES

High-resolution LES and DNS of the flow at different configurations of the idealized vegetation were conducted using the open-source, spectral element-based higher-order incompressible Navier–Stokes solver *Nek5000* [5]. The Spectral Element Method combines the accuracy of spectral methods and the flexibility of Finite Elements Method [7]. In the planned simulations with sediment transport, sediment would be modeled as Lagrangian particles using a novel semi-implicit time-stepping scheme developed to simulate polydisperse sediment accurately.

RESULTS & IMPACT

2D simulations have been conducted for the full domain on different configurations. Two cases having the same vegetation

density and Reynolds number, but different array configurations, are presented here. The velocity magnitude, along with the pressure field, is shown in Fig. 1. For the staggered case, in contrast with the random array, a vortex being shed from a cylinder is impeded by the ones behind it. This is evident in the pressure plots, where more and larger low-pressure areas, indicating the low-pressure core of rotating vortices, appear in the random case. For comparison, the random configuration was subjected to unidirectional flow at the same Reynolds number. High-flow zones near the walls arise, resulting in stronger vortices being shed from near-wall cylinders. Compared with the oscillatory flow case, more high-speed regions are also found among cylinders in the unidirectional cases.

Results from 3D simulations of turbulent flow through a random arrangement of cylinders have also been conducted (Fig. 2). A quarter of the full domain was simulated, with ~296 million computational points. Fields of instantaneous velocity magnitude at different elevations show that the length of the vortex being shed increases with distance from the bottom. This simulation, one of the largest high-resolution eddy-resolved hydrodynamic simulations in this field, will provide as yet unseen details of the physical processes involved.

WHY BLUE WATERS

The study pushes the limit of the scale at which high-resolution simulations are used to study complex multi-phase flow in environmental fluid mechanics, requiring computational resources with sustained computing power at an unprecedented scale, such as Blue Waters. Simulations have been conducted for up to 296 million computational points, with the code scaling strongly up to 32,768 MPI ranks. Without access to petascale HPC like Blue Waters, completing the study within a realistic timeframe would be impossible. In addition, since visualization of a phenomenon is an effective way to understand and explain its mechanics, we will work with Blue Waters project staff to create animations of the phenomenon using data from the simulations.




Figure 1: Instantaneous velocity magnitude (a,c) and hydrodynamic pressure (b,d) for oscillatory flow at Reynolds number ~10000. For staggered (a,b) and random (c,d) configuration of cylinders the flow has been captured while accelerating from left to right. For the staggered case, the vortex being shed from a cylinder is impeded by the ones behind it, which is not the case for the random case.

QUANTUM MONTE CARLO SIMULATIONS OF MAGNETISM AND MODELS IN CONDENSED MATTER

Allocation: Illinois/202.8 Knh
 PI: Lucas K. Wagner¹
 Collaborators: Elif Ertekin¹, Narayan Aluru¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

This project uses quantum Monte Carlo techniques to study the behavior of electrons in materials where traditional methods of calculation are too inaccurate to be of use. These large-scale calculations will provide unprecedented detail about the correlated behavior of electrons in these materials at the subnanoscale. The simulations will then be analyzed in a systematic way to extract a compressed description of their behavior, which will help us to understand the difference at a microscopic level between a simple magnetic material and unconventional superconductors.

RESEARCH CHALLENGE

There are special materials in which quantum mechanics are noticeable even at the scale of the entire material. The simplest examples of this are semiconductors, which are by now quite well understood by band structure theory and have had a huge impact on technology. However, materials such as unconventional high-temperature superconductors, magnetic materials, topological materials, and many others have quantum effects that are very challenging to describe under the normal band structure framework. Current tools also struggle to accurately predict the properties even of semiconductors.

The Schrödinger equation, elucidated in the 1920s and carefully tested, tells us how, in principle, to predict the properties of materials. Until recently, however, directly solving the Schrödinger equation for complex materials has been out of reach because of the computational cost.

METHODS & CODES

We use the QWalk package [1], which was developed at the University of Illinois at Urbana-Champaign, to perform quantum Monte Carlo calculations. These methods sample the positions of the electrons using a stochastic process. How often the positions appear in the random sample represents the magnitude of the wave function. In this way, we can incorporate correlation effects between electrons. This method uses a single approximation, called the fixed node approximation, to become efficient. Since there is just one major approximation, it can be improved as more computational resources become available to achieve more and more accurate results.

RESULTS & IMPACT

With Blue Waters (BW), it is now possible to achieve unprecedented accuracy in solving the Schrödinger equation using Monte Carlo techniques. Our group has used BW to show that the simulations represent a step forward in the state of the art in computing many particle quantum systems. We have been able to describe the properties of water interacting with graphene, defects in semiconductors, the properties of superconducting materials, and transitions between metallic and insulating behavior.

The figure shows work performed on BW on vanadium dioxide [2], which transitions from a metal to an insulator at around 340°K (150°F). This transition can also be induced by strain [3] or electric field. For many years, it has been a mystery why this happens, mainly because electron correlation is important in the effect, which is difficult to access with inaccurate quantum calculations. By using BW to perform quantum Monte Carlo calculations of the electrons in this material, we were able to describe the transition from first principles, meaning that we used only the positions of the atoms to compute its properties. This new predictive capability opens the door to computer-based design of these extraordinarily sensitive materials that can be used in new devices.

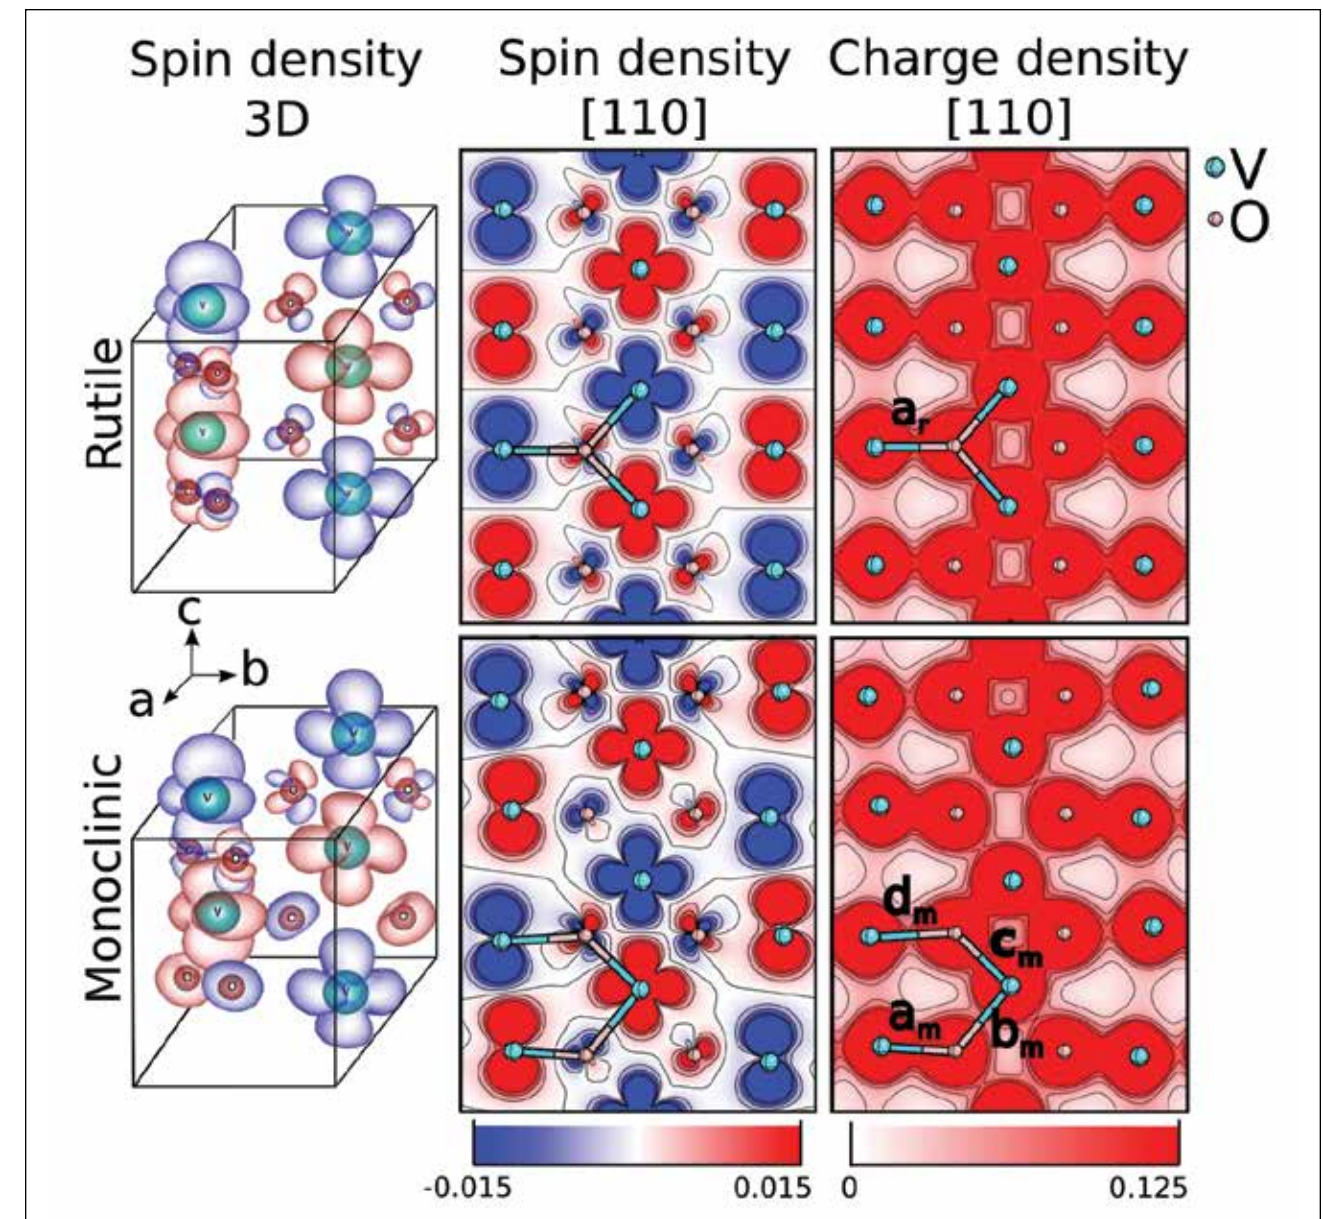
Now that we have accurate calculations available, we would like to learn as much as possible from them. We would like to be able to know how to think about a given material. In physics, this is formally given by using an object known as an effective Hamiltonian, which encodes how a system behaves under experimental conditions. Toward that objective, we are using BW to probe quantum systems and extract effective Hamiltonians by performing rigorous multiscale quantum simulations.

WHY BLUE WATERS

The BW system provides excellent technical support and access to high-performance computing with a minimal amount of effort required from scientists, so that we can focus on the science. As one of the few National Science Foundation-supported resources that provides large allocations, it has been critical to the success of the projects listed here.

PUBLICATIONS AND DATA SETS

Wagner, L.K., and P. Abbamonte, Effect of electron correlation on the electronic structure and spin-lattice coupling of high- T_c cuprates: Quantum Monte Carlo calculations. *Physical Review B*, 90 (2014), 125129.



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

Shiller, J.A., L.K. Wagner, and E. Ertekin, Phase stability and properties of manganese oxide polymorphs: Assessment and insight from diffusion Monte Carlo. *Physical Review B*, 92 (2015), 235209, DOI: 10.1103/PhysRevB.92.235209.

Wu, Y., L.K. Wagner, and N.R. Aluru, The interaction between hexagonal boron nitride and water from first principles. *The Journal of Chemical Physics*, 142 (2015), 234702.

Wu, Y., L.K. Wagner, and N.R. Aluru, Hexagonal boron nitride and water interaction parameters. *The Journal of Chemical Physics*, 144 (2016), 164118.

Yu, J., L.K. Wagner, and E. Ertekin, Towards a systematic assessment of errors in diffusion Monte Carlo calculations of semiconductors: Case study of zinc selenide and zinc oxide. *The Journal of Chemical Physics*, 143 (2015), 224707.

Yu, J., L.K. Wagner, and E. Ertekin, Fixed-node diffusion Monte Carlo description of nitrogen defects in zinc oxide. *Physical Review B*, 95 (2017), 075209.

Zheng, H., and L.K. Wagner, Computation of the correlated metal-insulator transition in vanadium dioxide from first principles. *Physical Review Letters*, 114 (2015), 176401.

INTERMITTENCY, RESOLUTION EFFECTS AND HIGH SCHMIDT NUMBER MIXING IN TURBULENCE

Allocation: NSF PRAC/9,547.165 Knh
PI: P.K Yeung¹
Collaborator: T. Gotoh²

¹Georgia Institute of Technology
²Nagoya Institute of Technology, Japan

EXECUTIVE SUMMARY

Substantial advances have been made in both science and computing for fluid turbulence using Blue Waters. Simulations on grids of up to 4 trillion points show clear differences between energy dissipation rates averaged locally in three versus one dimension (the latter being common in the literature). Resolution effects on extreme events are examined critically using a multi-resolution approach. We have also developed a new algorithm for turbulent mixing at low diffusivity, where scalar fluctuations arise at scales much smaller than in the velocity field. We use a dual communicator approach where different groups of MPI processes compute the velocity and scalar field at different resolutions and using different numerical schemes. Through careful use of inter-communicator communication, as well as multithreading via nested OpenMP constructs, the code scales well up to 524,288 cores at close to six percent of theoretical peak performance.

RESEARCH CHALLENGE

Turbulent flows with disorderly fluctuations over a wide range of scales are an important agent of efficient mixing in many fields of science and engineering. An enduring challenge in the theory of this subject is to understand the nature of intermittency [1] in terms of fluctuations of the energy dissipation rate (a measure of local straining) and enstrophy (a measure of local rotation) over a wide range of domain sizes. Direct numerical simulations of the scale enabled by Blue Waters are the best source of data for this purpose, but the complexity of the flow physics requires a

critical examination of the accuracy and reliability of the results. Accordingly, in addition to analyzing data from a 0.5-trillion-grid-points simulation [2] we have also performed a short simulation using as many as 4 trillion grid points. Massive datasets at this level allow us to investigate the effects of differences in the local averaging procedure as well as the effects of finite resolution.

A second focus area is the development and application of a new parallel algorithm that is uniquely suited to the study of turbulent mixing at high Schmidt number, where low molecular diffusivity leads to fluctuations at scales smaller than those in the velocity field. Our objective includes checking a scaling relation proposed in classical theory for which confirmation via either simulation or experiment has been limited by grid resolution and/or Schmidt number. Our new simulation uses a new dual-resolution parallel paradigm to make the calculation practical.

METHODS & CODES

Our technical approach is direct numerical simulation (DNS), which is based on exact physical laws and can be carried out efficiently in simplified geometries. We use Fourier pseudo-spectral methods for the velocity field, and perform local averaging both along a line (one dimension) and over a cube (three dimensions). To study resolution effects via filtering we apply successive truncations in wavenumber space before transforming to physical space coordinates. Results that differ greatly after truncation may indicate substantial errors.




Figure 1: Moments of orders 2 to 6 (in direction of arrows) of normalized and locally averaged dissipation (solid lines) and enstrophy (dashed lines) versus scale size normalized by the Kolmogorov scale. Frames (a) and (b) show data obtained from 1D and 3D averages, respectively.

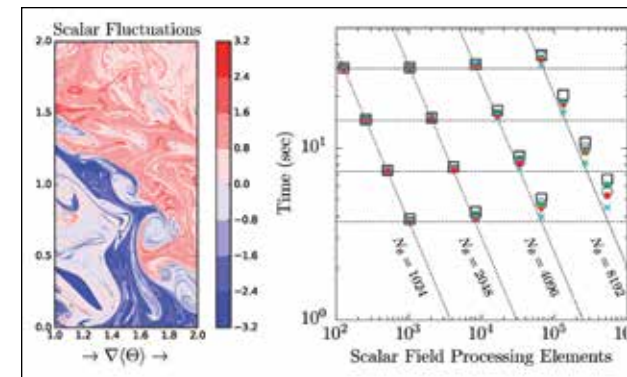


Figure 2: Left—Isocontours of normalized scalar fluctuations in a local region within a two-dimensional plane. Right—Scalability of scalar field computation for different grid resolutions using different versions of the CCD routines, from the baseline single-threaded nonoverlapped version (open squares) to the best version with dedicated communication threads (stars).

For high Schmidt number mixing we retain pseudo-spectral methods for the advecting velocity field but use combined compact finite differences (CCD) on a finer grid to satisfy stringent resolution requirements for the scalar [3]. We have devised a new dual-communicator parallel algorithm [4] where distinct groups of MPI parallel processes work on the velocity and scalar fields at different resolutions. This particular setup is less communication-intensive than full pseudo-spectral codes. Rigorous efforts have been made to improve performance by overlapping communication with computation. The best performance is obtained by dedicating one OpenMP thread to communication and dividing computational work among the other threads using nested OpenMP parallelism.

RESULTS & IMPACT

In the study of intermittency in high Reynolds number turbulence there is great interest in whether higher-order moments of the local averages of dissipation and enstrophy exhibit power law behaviors at intermediate scale sizes. Most past data in the literature were based on one-dimensional averaging, which is conceptually not ideal. Fig. 1 shows data on logarithmic scales with a clear scaling range (between the two vertical dashed lines) but only for averages taken in three dimensions, in a long simulation with 8,192³ grid points in each direction. Dissipation and enstrophy scale in the same manner, as well. A short simulation at an even higher resolution has further confirmed these results. The values of "scaling exponents" deduced from the data are helpful in allowing intermittency theories to be evaluated more definitively than in the past.

For turbulent mixing at high Schmidt number, Fig. 2 shows the delicate fine-scale structure of a scalar field at Schmidt number 512 (comparable to the salinity of the ocean), and the high scalability of the dual-communicator hybrid algorithm over a wide range of problem sizes and Cray XE6 core counts on Blue Waters. Careful observation aided by the color annotations in the figure suggests the presence of sharp interfaces where large

fluctuations of the scalar gradients exist. In the right frame of Fig. 2, different symbols represent different coding implementations of the hybrid algorithm, differing mostly in the manner and degree in which overlapping between computation and communication is implemented. Horizontal and sloping lines indicate limits of perfect weak and strong scaling, respectively. At the largest problem size (8,192³), strong scaling efficiency from 65,536 to 262,144 cores is as high as 94 percent. This highly scalable code enables studies of turbulent mixing in this previously unreachable Schmidt number regime.

WHY BLUE WATERS

The 8,192³ grid resolution of our production simulations requires access to a world-class machine such as Blue Waters, and the memory capacity of Blue Waters has enabled us to obtain data at 16,384³ resolution, the highest known worldwide in the turbulence community. Other important machine characteristics include support for Co-Array Fortran, topologically aware scheduling, and a large number of cores per node, which is conducive to nested OpenMP parallelism in our high-Schmidt-number algorithm.

PUBLICATIONS AND DATA SETS

Clay, M. P., Strained turbulence and low-diffusivity turbulent mixing. Ph.D. thesis, Georgia Institute of Technology (2017).

Clay, M. P., D. Buaria, T. Gotoh, and P. K. Yeung, A dual communicator dual grid-resolution algorithm for petascale simulations of turbulent mixing at high Schmidt number. *Computer Physics Communications*, in press (2017), DOI: 10.1016/j.cpc.2017.06.009.

Iyer, K. P., K. R. Sreenivasan, and P. K. Yeung, Reynolds number scaling of velocity increments in isotropic turbulence. *Physical Review E*, 95:2 (2017), DOI: 10.1103/PhysRevE.95.021101.

Yeung, P.K. Turbulence in fluid dynamics: the science, and the need for exascale. *First International Symposium on Research and Education for Computational Science*, Tokyo, Japan, November 29–30, 2016.

COMPUTER SCIENCE & ENGINEERING

COMPUTATIONAL BIOLOGY

NETWORK CONGESTION

OPTIMIZATION

SCALABILITY

VISUALIZATION

170 *CADENS NSF Project: Digital Literacy, Data Visualization, and the Cinematic Presentation of Science*

172 *Algorithms for Extreme-Scale Systems*

174 *Parallelization of the Multilevel Fast Multipole Algorithm (MLFMA) on Heterogeneous CPU-GPU Architectures*

176 *Predicting Performance Degradation and Failure of Applications through System Activity Monitoring*

178 *Parallel Algorithms for Solving Large Assignment Problems*

180 *Localizing Communication in Sparse Matrix Operations*

182 *Performance Evaluation of New Algebraic Algorithms and Libraries*

184 *Parallel Algorithms for Big Data Phylogenomics, Proteomics, and Metagenomics*

186 *Hardware Acceleration of Deep Learning*

188 *A Critical Evaluation of the OP2/OPS Parallel Meshing and Code Generation Software*

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

Allocation: Illinois/300 Knh

PI: Donna Cox¹

Co-PIs: Kalina Borkiewicz¹, Jeff Carpenter¹, AJ Christensen¹, Stuart Levy¹, Robert Patterson¹

Collaborators: Nathan Goldbaum¹, Michael Norman², Leigh Orf³, Brian O'Shea⁴

¹National Center for Supercomputing Applications

²University of California, San Diego

³University of Wisconsin–Madison

⁴Michigan State University

EXECUTIVE SUMMARY

The Advanced Visualization Lab (AVL) continues work on the NSF-funded CADENS project (The Centrality of Advanced Digitally ENabled Science, ACI-1445176). The AVL co-produced and rendered visualization scenes for the recently released science documentary, “Seeing the Beginning of Time.” We have used Blue Waters for processing data as well as rendering scenes in 4K monoscopic, stereoscopic, and dome formats. A few technological byproducts of this work include a new pipeline for rendering on Blue Waters called Blurend and an open-source cinematic scientific visualization tool called Ytini.

RESEARCH CHALLENGE

The AVL produces cinematic scientific visualizations for public outreach. Cinematic visualization is an especially accessible form of science communication because it leverages a powerful and familiar visual language established by Hollywood films. Creating production-quality renderings of supercomputer datasets that are both cinematic and scientifically accurate is challenging. There are many existing tools for cinematic animation (Houdini, Maya, 3DS Max, Blender, etc.) and many tools for scientific visualization (yt, ParaView, Partiview, VisIt), but few tools exist at the intersection of cinema and science.



Figure 1: A visualization of the Renaissance simulations [1] done by the Advanced Visualization Lab. The universe 400 million years after the Big Bang.

METHODS & CODES

To create visualizations for the science documentary “Seeing the Beginning of Time,” the AVL primarily used the commercial visual effects tool Houdini and the scientific analysis and visualization Python package yt. We have created middleware and a pipeline that uses both of these tools, called Ytini. We have run Ytini on Blue Waters to convert raw data into Houdini-compatible OpenVDB sparse volume data files. Further, we have created a Python tool called Blurend for preparing Houdini scene files for rendering on Blue Waters.

RESULTS & IMPACT

We have rendered the following sequences on Blue Waters:

- “First Light in the Renaissance Simulations” [1] in 4K stereoscopic format
- “First Light in the Renaissance Simulations” [1] in 4K dome format
- “El Reno 2011 Tornado” [4] in 4K dome format
- “Large Synoptic Survey Telescope Model” [6] in 4K monoscopic format

We processed data on Blue Waters for the following sequences:

- “Cosmic Bubble Bath” [2] for the 4K documentary “Seeing the Beginning of Time”
- “Formation of the Moon” [5], for our work-in-progress dome show

“Seeing the Beginning of Time” is currently available for streaming on Amazon Prime, where it is rated 4.3/5 stars.

AVL’s new tool, Ytini, has significantly optimized our pipeline by allowing us to read many scientific data formats using the Python package yt. This has cut out weeks of development time from our visualization process. Ytini is an open-source project, making it easier for both artists and scientists to make their own cinematic scientific visualizations. Files and tutorials are distributed through www.ytini.com.

WHY BLUE WATERS

We are currently using Blue Waters to prepare data for import into the commercial visual effects tool, Houdini, where we can create cinematic treatments of the data. Converting a particular dataset from particles to a Houdini-compatible Adaptive Mesh Refinement-like volumetric format took about four days on our local cluster. It would have been incredibly difficult to iterate at this time scale. Running the same script across nodes on Blue Waters took four hours.

Additionally, we have created a new pipeline for preparing Houdini scene files for rendering on Blue Waters, resulting in a new software tool called Blurend. We used this tool to render several sequences for “Seeing the Beginning of Time” as well as a tornado simulation [4] in dome format.

Several of our scientific collaborators used Blue Waters to run their simulations and create large datasets residing there. It was



Figure 2: A dome rendering of a tornado simulation [4], done by the Advanced Visualization Lab.

is more productive and save valuable time to be able to work with their data in situ on the Blue Waters filesystems rather than transfer large amounts of data to specialized visualization systems.

PUBLICATIONS AND DATA SETS

Naiman, J., K. Borkiewicz, and A. Christensen, Houdini for Astrophysical Visualization. *Publications of the Astronomical Society of the Pacific*, 129:975 (2017), DOI:10.1088/1538-3873/aa51b3.

Borkiewicz, K., A. Christensen, and J. Stone, Communicating Science Through Visualization in an Age of Alternative Facts. *SIGGRAPH Conference on Computer Graphics and Interactive Techniques*, Los Angeles, Calif., July 30–August 3, 2017.

Cox, D., et al., First Light in the Renaissance Simulation Visualization: Formation of the Very First Galaxies in the Universe. *Supercomputing, Scientific Visualization Showcase Submission*, Denver, Colo., November 12–17, 2017.

Cox, D., et al., Milky Way Analogue Isolated Disk Galaxy Visualization. *Supercomputing, Scientific Visualization Showcase Submission*, Denver, Colo., November 12–17, 2017.

“Seeing the Beginning of Time,” science documentary directed by Thomas Lucas (50 minutes). Long-form web distribution via Amazon Prime.

ALGORITHMS FOR EXTREME-SCALE SYSTEMS

Allocation: Blue Waters Professor/80 Knh

PI: William Gropp¹

Collaborator: Luke Olson¹

Students: Amanda Bienz¹, Paul Eller¹, Philipp Samfass¹

¹University of Illinois at Urbana-Champaign

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 both the increasing number of cores per processor and the use of large numbers of specialized computing elements in GPGPUs (general-purpose graphics processing units). 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. Over the last year, we have explored two ways to reduce communication costs in large-scale systems. One is the redesign of algorithms to take into account the difference between on-node and off-node communication performance. The other is to look at the effect of performance irregularity and the use of nonblocking collectives to improve performance of algorithms that use MPI (message passing interface) collectives. We also developed an improved communication model that better matches the performance of modern parallel processors.

RESEARCH CHALLENGE

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 sustained performance.

This work directly targets current barriers to effective use of extreme scale systems by applications. For example, Krylov methods such as Conjugate Gradient are used in many applications currently being run on Blue Waters (MILC is one well-known example) and other leadership-class systems. Developing and demonstrating a more scalable version of this algorithm would immediately benefit those applications. Also of importance to many computations, including Krylov methods for solving large systems of linear equations as well as methods for large-scale graph computations, are sparse-matrix vector multipliers. These involve significant communication between nodes and can lead to scalability limits; by improving methods to exploit internode and intranode communication, many applications can improve their scalability. In the longer term, the techniques that are developed in this project will provide guidance for the development of highly scalable applications.

METHODS & CODES

To address the challenges of parallelism and scale, we developed several codes that allow us to benchmark the performance of these operations, gather detailed timing results, and perform experiments with different approaches. For example, we have been developing a “noise injector” to allow us to better experiment with different amounts of performance variation in multicore nodes. We have also developed a set of benchmark codes that better measure the achievable communication performance of the communication patterns commonly used in applications.

RESULTS & IMPACT

Early results with alternative Krylov formulations have revealed several performance effects that can provide a factor of two or more improvement in performance at scale [1]. We have been using Blue Waters over the last year for an investigation into the impact of large-scale system performance variation on parallel numeric algorithms. This includes developing code for measuring and processing network performance counters, injecting network noise into nodes running another algorithm, and kernels with a variety of common communication patterns. Our initial experiments have used smaller core counts to assist in developing the code and improving the experiments in preparation for a more detailed study in the coming months involving large runs. The goal of this study is to better understand what network noise looks like on supercomputers and how we can develop parallel numeric algorithms that perform well despite noise.

We have also explored the performance models used to guide both the development of algorithms and the analysis of application performance. We discovered that the classic “postal model” is no longer effective for systems with multi-core nodes, and we developed a simple extension to the postal model that explains the observed performance of current systems [2]. This work has also informed the development of an improved approach to sparse matrix–vector multiplication that takes into account the different performance of inter- and intra-node communication [3].

PingPong Performance Between Two Nodes




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

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 environments available for large-scale experiments. 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.

PUBLICATIONS AND DATA SETS

Eller, P. R., and W. Gropp, Scalable nonblocking preconditioned conjugate gradient methods. *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis* (IEEE Press, Salt Lake City, Utah, November 13–18, 2016), pp. 1–12.

Gropp, W., et al., Modeling MPI Communication Performance on SMP Nodes: Is it Time to Retire the Ping Pong Test? *Proceedings of the 23rd European MPI Users' Group Meeting* (ACM Digital Library, Edinburgh, United Kingdom, September 25–28, 2016), pp. 41–50.

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

Allocation: Exploratory/62.5 Knh

PI: Levent Gürel¹

Co-PI: Wen-mei Hwu¹

Collaborators: Weng Cho Chew¹, Narayan Aluru¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

The aim of this allocation was to solve unprecedentedly large scattering problems requiring large amounts of memory and to utilize a multitude of GPUs (graphics processing units) available in Blue Waters. In the scope of this project, we solved a scattering problem with 1.5 billion unknowns on 512 XE nodes; this was beyond what we could achieve before using Blue Waters. Our efficient parallelization strategy [1,2] allowed us to spread problems among large numbers of nodes to access more memory, which is required by large problems. Additionally, we achieved the largest full-wave inverse-scattering solutions in near-real time on 256 XK nodes, and it was 4.34 times faster than the same number of XE nodes [3]. This was the first GPU implementation for nonlinear inverse solutions in the literature. We implemented an efficient multi-GPU MLFMA (multilevel fast multipole algorithm) to serve as the kernel of an inverse-scattering solver.

RESEARCH CHALLENGE

We have been working on solutions of extremely large problems that are derived mainly from electromagnetics, acoustics, and optics, to name a few. An electromagnetic scattering problem may become very large as either the frequency increases or the target




Figure 1: Convergence of the iterative inverse-scattering algorithm. The image details can be well seen after 40 iterations. It takes about two minutes to solve 38,400 forward-scattering problems with MLFMA on 256 XK nodes, and this is 4.34 times faster than employing the same number of XE nodes [5].

size gets larger. Simulating the scattering of high-frequency radar waves from a large aircraft, for instance, may require the solution of millions or even billions of unknowns. Similarly, medical imaging with microwaves requires solutions of thousands of large problems. As the problem size (i.e., the number of unknowns) grows, memory requirement increases so we need more nodes, providing access to more memory. MLFMA operates on a multilevel tree structure and is mathematically involved; therefore, it is difficult to implement the parallel algorithm on multiple CPUs. Moreover, it is not easier to implement MLFMA on GPUs because the algorithm is memory-bound due to its $O(N)$ computational complexity. The challenge is to develop efficient parallelization strategies to spread problems evenly among large number of nodes to obtain larger solutions, and to exploit the hierarchical memory architecture of GPUs to obtain large computational throughput.

METHODS & CODES

The memory requirement of large problems is a challenge. Prior to Blue Waters, we reached a point where both the available total memory and the single-node memory became insufficient. Therefore, we experimented with out-of-core methods to use the disk storage as additional memory, even though we had to endure a huge penalty in the time required to solve problems [4]. The huge number of computing nodes on Blue Waters makes available a much larger total memory, thus enabling the solution of such enormous problems that were impossible for us to solve before. Additionally, improving parallelization and using more nodes led to much shorter solution times. The hierarchical parallelization strategy allows us to partition the MLFMA data structures among 512 XE nodes efficiently and evenly (i.e., in a load-balanced way). A scattering problem involving dense linear systems with 1.5 billion unknowns is solved within the 32 TB (terabytes) of memory provided by 512 XE nodes.

Unfortunately, MLFMA is a memory-bound algorithm. This is mainly due to the fast nature of MLFMA, where its $O(N)$ complexity prevents data reuse on GPUs [5]. To exploit the hierarchical memory architecture of GPUs, we reformulate the MLFMA operations as matrix-matrix multiplications; this provides good data reuse and high computational throughput. To implement matrix-matrix multiplications efficiently, we use a




Figure 2: Radar cross-section of a conducting sphere of 1,200 wavelengths in diameter requires the solution of a $1,500,000,000 \times 1,500,000,000$ dense matrix equation. The computed bistatic scattering values are compared to semi-analytical Mie-series solution to demonstrate the high accuracy of the solution.

hybrid shared-memory and register tiling algorithm with thread-coarsening methods.

For multi-GPU parallelization, the MLFMA tree structure is partitioned among GPUs and the operations among the branches are categorized as intra-GPU and inter-GPU. The inter-GPU data are first transferred to CPUs (central processing units), and then MPI (message passing interface) communications are performed to swap the data among computing nodes, and finally the data are transferred back into GPUs. The MPI communications are optimized to eliminate redundancy and to minimize the communication time. To reduce idle time of CPUs and GPUs, we overlap both MPI communications and CPU-GPU transfers with GPU computations. To do that, we propose a different order than the standard one such that the data-transfer time is completely overlapped with GPU computations. This provides excellent inter-node parallelization of MLFMA [6].

All of our codes are developed in-house with Fortran, C++, and CUDA. The PETSc framework is used in the Fortran code for iterative solutions.

RESULTS & IMPACT

We have increased the number of unknowns in the problems we can solve by achieving the solution of an electromagnetic scattering problem with 1.5 billion unknowns. The ensuing $1,500,000,000 \times 1,500,000,000$ dense matrix equation is solved with MLFMA and parallelized on 512 XE nodes of Blue Waters.

We obtained an efficient multi-GPU MLFMA implementation for volumetric problems for the first time within the scope of this project. A single GPU speedup is 55.14 and 3.97 times with respect to the sequential and 16-core execution baselines. The CPU executions are obtained on XE nodes whereas the GPU executions are obtained on XK nodes. This implies 3.97 times speedup on XK nodes over XE nodes. The speedup of 16 XK nodes is 846.41 and 15.34 times over sequential and a single XK node. This implies 96% parallelization efficiency among XK nodes.

Large-scale inverse-scattering solutions on GPUs are obtained on 256 XK nodes, where each node is employed with an MLFMA solver. This parallelization scheme decreases an inverse solution from 11.5 hours (sequential execution) to 7 seconds on 128 XK nodes. This provides images in near-real time, allowing our fast

and massively-parallel algorithms to be employed for real-life applications. We plan to employ supercomputers (e.g., Blue Waters) for the benefit of complex imaging problems.

WHY BLUE WATERS

Our team is located at the University of Illinois at Urbana-Champaign and, therefore, it is very easy contact the Blue Waters staff. Occasionally, our Ph.D. students visit the NCSA building several times a day to solve their issues on compilers, libraries, and other Blue Waters utilities with the help of the staff. The vast amount of DRAM (dynamic random-access memory) and the number of GPUs available on Blue Waters enable scaling our science to much larger problems. Blue Waters (with its large number of CPU and GPU nodes) is ideally suited for our research. We can also make short-term (e.g., one-hour) reservations for a large number of nodes (e.g., 1,024 nodes) for large benchmark runs [7].

PUBLICATIONS AND DATA SETS

Hidayetoğlu, M., C. Pearson, W. C. Chew, L. Gürel, and W.-M. Hwu, Large Inverse-Scattering Solutions with DBIM on GPU-Enabled Supercomputers. *International Applied Computational Electromagnetics Society Symposium*, ACES, Florence, Italy, March 26–30, 2017.

Hidayetoğlu, M., et al., Large-Scale Inverse Scattering Solutions with Parallel Born-Type Fast Solvers. *Progress on Electromagnetics Research Symposium* (PIERS, Shanghai, China, August 8–11, 2016), pp. 916–920.

Hidayetoğlu M., et al., Incorporating Multiple Scattering in Imaging with Iterative Born Methods. *USNC-URSI National Radio Science Meeting*, Boulder, Colo., January 4–7, 2017.

Hidayetoğlu M., et al., Fast DBIM Solutions on Supercomputers with Frequency-Hopping for Imaging of Large and High-Contrast Objects. *Progress on Electromagnetics Research Symposium*, St. Petersburg, Russia, May22–26, 2017.

Hidayetoğlu M., W.-M. Hwu, and W. C. Chew, Performance Considerations on Various Iteration Schemes for the Distorted-Born Iterative Method. *2017 IEEE AP-S Symposium on Antennas and Propagation and USNC-URSI Radio Science Meeting*, San Diego, Calif., July 9–14, 2017.

PREDICTING PERFORMANCE DEGRADATION AND FAILURE OF APPLICATIONS THROUGH SYSTEM ACTIVITY MONITORING

Allocation: Exploratory/50 Knh
PI: Ravishankar K. Iyer¹
Co-PI: Zbigniew Kalbarczyk¹
Collaborator: Saurabh Jha¹, Benjamin Lim Wen Shih¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Our overarching investigation addresses complex data-driven problems associated with online system monitoring for understanding causes of application failures and performance degradations. Our contributions include the design of machine learning-based deep analytics framework to distinguish the cause of application and system performance degradation due to failure- and nonfailure-related issues. The tool leverages probabilistic graphical models to conduct machine learning at scale for runtime detection of congestion parameters and its effects on running applications. Our approach correctly identifies the cause of congestion in 74% of the cases (of 302 cases found in our dataset) and attributes it either to resource contention issues or to failures in the system. In addition, working with Sandia National Laboratories and the National Energy Research Scientific Computing Center, we are addressing the issue of congestion in the Cray Aries interconnect to help advance scheduling decisions of applications by porting our algorithms for online system monitoring and mitigation.

RESEARCH CHALLENGE

Extreme-scale high-performance computing (HPC) systems require a holistic approach to monitoring and coordinating many disparate subsystems (both hardware and software) to enable continued scaling and efficient execution of applications. HPC systems are typically used for executing tightly coupled simulation applications across hundreds of thousands to millions of processor

threads. Resource contention due to failures or design issues (of applications/systems) impacts applications in two ways: (1) by degrading application performance, which causes application runtime unpredictability and limits scaling to full system; and (2) through propagation of errors and failures in application logic and code-flow, which causes application failures or invalid/error-prone outputs. As a result, effective failure/degradation mitigation response(s) in complex systems require analysis of the propagation of faults/errors and of performance issues due to interference among applications or resource exhaustion.

Our analysis approach addresses more general scenarios that result in performance degradation in which timely and appropriate response can significantly improve both application runtimes and system throughput. The eventual product of this work will be an interoperable set of capabilities for extreme-scale systems that provides monitoring, analysis, and appropriate response to both resilience and performance issues. These capabilities will support both automated and exploratory analysis, both at runtime and in post-processing. Specifically, the Blue Waters grant allocation was used to characterize performance logs and machine-generated error logs to holistically understand the propagation of performance degradation and faults affecting applications. We use Bayesian network-based machine learning methods to distinguish between failure-based resource contention issues and nonfailure-based contention issues.

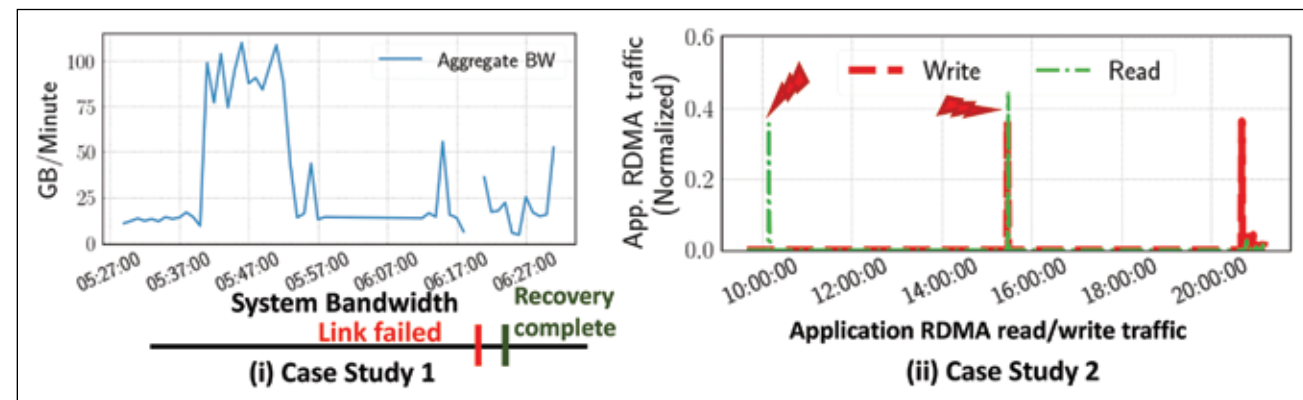


Figure 1: Case study 1 shows link-failure-related performance degradation. Case study 2 shows design issue-related performance degradation.




Figure 2: Bayesian network for diagnosing the cause of application degradation. Green-filled circle values are obtained using unsupervised clustering algorithms from raw data. Yellow-filled circle values are obtained directly from measure-derived metrics.

METHODS & CODES

Regardless of whether performance problems such as contention for shared resources are failure-related (e.g., caused by link failures) or nonfailure-related (e.g., caused by design issues in an application or system), they can quickly propagate in systems and result in severe performance degradation, rather than application failure as discussed below.

Case Study 1: Fig. 1 (i) shows the impact of network link failure and its corresponding recovery on the Gemini 3D torus interconnection network of Blue Waters. This impact occurs because a single-link failure can cause the traffic pattern of the application to change in the underlying network. Event logs corresponding to this failure are shown in the figure. It also shows the aggregate data passing through all the Gemini routers. Basic statistics such as these indicate bandwidth utilization of the system. To understand whether there is actual performance degradation and, if so, which Gemini router may be the source of the problem, we calculate a derived metric, “average packet latency,” which captures the average time taken by a router to deliver packets.

Case Study 2: Fig. 1 (ii) shows the impact on performance of the system and application due to interplay of design issues between the application (write pattern) and system (scheduling strategy). Specifically, a 32-node job caused high congestion in the system interconnect, triggering two congestion protection events, first at 10:00 a.m. (within 10 seconds of job launch) and then at 15:20 (red lightning bolts). Congestion within the torus can adversely impact the performance of the application generating the messages and of other running applications, and is a major cause of inconsistent application runtimes. This job was assigned a linear shape in the “Z” direction by the topology-aware scheduler. Hence, when it was reading large amounts of data (over RDMA, or remote direct memory access), most of the I/O calls were funneled through specific links, causing the congestion.

RESULTS & IMPACT

Our tools take a holistic approach for differentiating design-related performance issues from failure-related issues. This is done by building a Bayesian network-based diagnostic model using features obtained from systems and applications. These features are either obtained directly from measured raw metrics (yellow-filled circles in Fig. 2) or by running unsupervised clustering algorithms on raw metrics (green-filled circles). We tested our design and tools by manually verifying the cause of congestion for 302 reported cases and achieved 74% accuracy.

WHY BLUE WATERS

Blue Waters is one of the few open-science capacity systems that provides a test bed for scaling computations to tens or hundreds of thousands of cores on central processing units (CPUs) and graphics processing units (GPUs). It also enables the study of failures and degradation of applications in production petascale systems because of its unique mix of XE6 and XK7 nodes. This allows us to understand the performance–fault-tolerance continuum in HPC systems by enabling the investigation of application-level designs for mixed CPU and GPU node systems, and fault isolation in system components to mitigate failures at the application level.

PUBLICATIONS AND DATA SETS

Jha, S., et al., Resiliency of HPC Interconnects: A Case Study of Interconnect Failures and Recovery in Blue Waters. *Proceedings of the IEEE Transactions on Dependable and Secure Computing*, PP:99 (2017), DOI: 10.1109/TDSC.2017.2737537.

Formicola, V., et al., Understanding Fault Scenarios and Impacts through Fault Injection Experiments in Cielo. *CUG 2017*, Redmond, Wash., May 7–11, 2017.

Jha, S., et al., Holistic Measurement Driven System Assessment. *Workshop on Monitoring and Analysis for High Performance Computing Systems Plus Applications*, Honolulu, Hawaii, September 5, 2017.

PARALLEL ALGORITHMS FOR SOLVING LARGE ASSIGNMENT PROBLEMS

Allocation: Illinois/210 Knh
PI: Rakesh Nagi¹
Co-PI: Ketan Date¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

The objective of our project is to develop fast and scalable algorithms for obtaining strong lower bounds and exact solutions for large instances of the Quadratic Assignment Problem (QAP), using Graphics Processing Unit (GPU) clusters. The QAP is an NP-Hard problem, in the strong sense. To solve a linearized model of the QAP using branch-and-bound, lower bounds must be calculated using the Lagrangian dual technique, in which a large number of Linear Assignment Problems (LAP) are solved efficiently, using our GPU-accelerated Hungarian algorithm. Additionally, in a branch-and-bound scheme, a large number of nodes must be explored in order to find a provable optimal solution. To this end, we have used Blue Waters to develop a GPU-accelerated Lagrangian dual ascent heuristic for obtaining lower bounds on the QAP, which is used in the parallel branch-and-bound scheme to solve large QAPs to optimality.

RESEARCH CHALLENGE

The Quadratic Assignment Problem (QAP) was introduced by [1] as a mathematical model to locate indivisible economical activities (such as facilities) on a set of locations so as to minimize a quadratic cost function. Typical applications of the QAP may be found in facility layout (re)design in manufacturing, distribution systems, services (retail outlets, hospital floors, etc.), and electronic circuit design. QAP may also serve as a specialization to many “harder” optimization problems, such as the Graph Association (GA), Traveling Salesman Problem (TSP), Vehicle Routing Problem (VRP), etc., in alternative formulations. Therefore, to solve these problems efficiently, we need to solve the QAP sub-problems efficiently. As a result, a fast and scalable QAP solver can be a powerful tool for researchers working on such NP-hard problems, or problems that are exceeding difficult to

computationally solve on a finite resource in a finite time. The sequential QAP solver can become computationally intensive and, therefore, the algorithm can benefit from parallelization on an appropriate parallel architecture, such as Blue Waters.

METHODS & CODES

We chose to parallelize the Lagrangian dual ascent algorithm for Level-2 Refactorization-Linearization Technique (RLT2) proposed by [2], in which we need to solve $O(n^4)$ LAPs and adjust $O(n^6)$ Lagrange multipliers to obtain a strong lower bound on the QAP. We designed a parallel Lagrangian dual ascent heuristic for solving RLT2 using hybrid MPI+CUDA architecture. The $O(n^4)$ LAPs are split across these GPUs and solved using our GPU-accelerated Hungarian algorithm [3], while the $O(n^6)$ Lagrange multipliers are updated by multiple CUDA threads in parallel.

We used this GPU-accelerated dual ascent algorithm in a branch-and-bound scheme to solve QAP instances to optimality. For a node in the search tree, we fix a facility to a location and solve the corresponding RLT2 sub-problem, whose objective value provides a lower bound on the QAP. If this value is greater than the incumbent solution then the node is fathomed; otherwise, it is branched further. Each node is processed using a bank of GPUs. By using multiple such banks, we can process multiple nodes in parallel. The algorithms were tested on the problem instances from the QAPLIB [4].

RESULTS & IMPACT

Lower bounds: The results for lower bounding tests for the various problem instances are summarized in Table 1. With our architecture, we are able to obtain strong lower bounds on problem instances with up to 42 facilities, which is a tremendous achievement.

Problem	GPUs	Optimal Value	Lower Bound	% GAP	ltn Time (s)
Nug25	4	3744*	3610	3.58	7.99
Nug27	7	5234*	5076	3.02	8.49
Nug30	12	6124*	5846	4.54	9.08
Tai25a	4	1,167,256*	1,091,480	6.49	7.75
Tai30b	12	637,117,113*	620,444,000	2.62	10.73
Tho40	71	240,516	213,372	11.29	18.41
Sk042	95	15,812	14,741	6.77	21.54

Table 1: RLT2 dual ascent lower bounds on QAPLIB instances (selective)

Scalability study: Although there is a minimum required number of PEs for applying accelerated RLT2 dual ascent to a QAP of specific size, the number of GPUs can be increased and the LAPs can be solved in parallel on multiple GPUs. This allows us to achieve some parallel speedup. We performed strong scalability with 1 to 32 GPUs. We obtain good speedup in the initial stages. However, as we continue to increase the number of GPUs in the system, we get diminishing returns in the execution times, due to increased MPI communication.

Parallel branch-and-bound: The results for the parallel branch-and-bound tests are shown in Table 2. We can see that the number of nodes explored and the completion times increase exponentially with the problem size. The most challenging Nug30 problem instance required more than four days to solve optimally, using 300 GPU banks with 4 GPUs each, which is a significant achievement.

Problem	Optimal Value	GPU Banks	GPUs/Bank	Nodes Explored	Time (d:hh:mm:ss)
Nug25	3744*	50	2	3868	0:02:44:24
Nug27	5234*	100	2	55761	1:02:28:32
Nug30	6124*	300	4	840273	4:14:06:21
Tai25a	1,167,256*	100	2	523005	3:13:53:33
Tai30b	637,117,113*	60	4	30523	2:09:55:17

Table 2: Branch-and-bound results for QAPLIB instances (selective)

Impact: With our architecture, we are able to obtain strong lower bounds on problem instances with up to 42 facilities, and optimally solve problems with up to 30 facilities, using only a modest number of GPUs. To the best of our knowledge, this is a first-of-its-kind study that has pushed the RLT2 formulation to solving such large problems. With some additional work in memory management and CPU + GPU collaboration, our proposed algorithms may be used effectively to solve truly large QAPs with over 30 facilities, which has been impossible so far.

WHY BLUE WATERS

In a typical branch-and-bound tree, we need to explore a large number of nodes in order to find an optimal solution. As the problem size grows, the number of nodes that need to be explored grows exponentially. Therefore, we need a large number of processors that can explore the solution space in parallel. Additionally, the GPU-accelerated dual ascent procedure benefits from the large number of powerful GPU-enabled processors available at the Blue Waters facility. As seen in the results, we have used over 1,200 XK compute nodes for over 110 hours to solve Nug30 problems, which would have incurred significant costs on the proprietary systems such as the AWS. We are grateful to Blue Waters and the project staff for providing this invaluable service to the scientific community.

PUBLICATIONS AND DATA SETS

Date, K., and R. Nagi, GPU-accelerated Hungarian algorithms for the Linear Assignment Problem. *Parallel Computing*, 57 (2016), pp. 52–72, DOI: 10.1016/j.parco.2016.05.012.

LOCALIZING COMMUNICATION IN SPARSE MATRIX OPERATIONS

Allocation: Illinois/50 Knh
PI: Luke Olson¹
Collaborator: Amanda Bienz¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Parallel sparse iterative methods, such as algebraic multigrid (AMG), solve a variety of linear systems in virtually every field of science and engineering. Supercomputers such as Blue Waters provide sufficient memory and bandwidth to solve extremely large systems, enabling the simulation of more complex problems. However, the standard computational kernels for sparse matrix operations in methods such as the AMG algorithm lack the scalability required to take full advantage of current hardware. The Blue Waters allocation has exposed an opportunity to utilize the topology of the underlying communication network to reduce the communication costs associated with the parallel matrix operations that dominate the cost of each iteration of AMG.

RESEARCH CHALLENGE

Sparse matrices arise in many large-scale simulations. Sparse matrix operations such as matrix-vector multiplication and matrix-matrix multiplication are key computational kernels

and demand significant resources on the machine, largely in communication among processing elements. When multiple processors on a node of Blue Waters communicate with processors on distant nodes, communication costs are increased further. The key challenge of this project is to limit costly, internode communication through localization of the sparse matrix routines.

METHODS & CODES

Current supercomputing architectures consist of a large number of nodes, each with several multi-core processors. For example, many sparse matrix operations use 16 or 32 processors per node on Blue Waters. The standard approach to sparse matrix communication does not consider the physical location of processors on the network. Yet, the cost of communication varies greatly depending on the locations of each endpoint. For instance, communication between two processes located on the same node incurs a significantly lower cost than communicating between two different nodes.




Figure 1: Communication patterns in the standard and topology-aware algorithms.

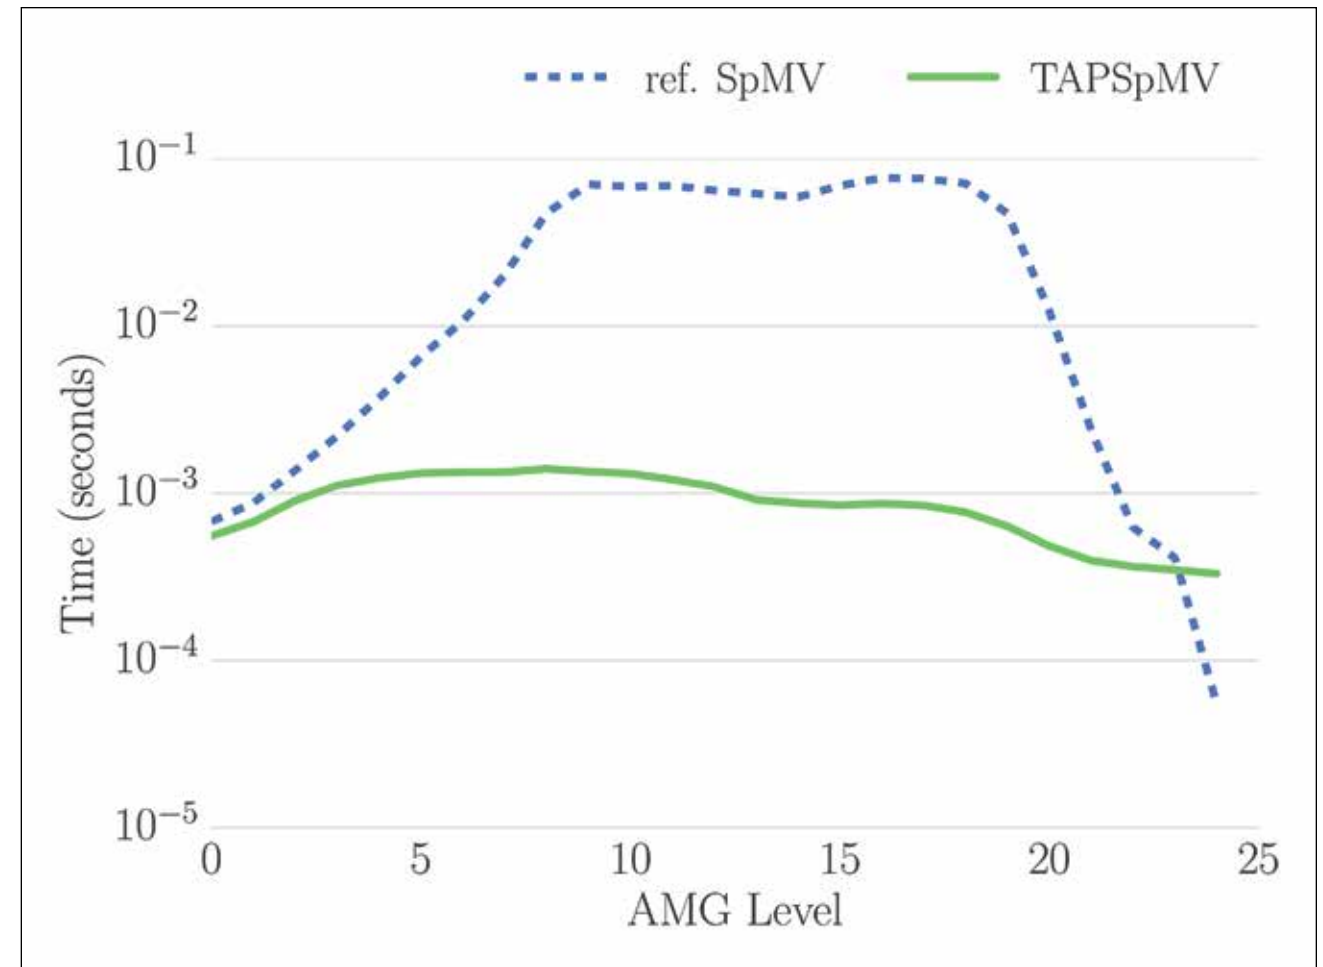


Figure 2: Time for standard and topology-aware sparse matrix-vector multiplication (SpMV).

The method developed in this work reroutes communication inside of sparse matrix operations so that MPI messages are aggregated on a node before executing internode communication on the network. This is shown in Fig. 1, where processes on nodes n and m typically send to some process q , sending directly to process q regardless of the starting location. In contrast, the method developed here considers the topology for communication. All processes on node n send everything with destination on node m to some local process p . After all data are collected, process p sends a single message across the network to process q , where it is then distributed to processes on node m .

RESULTS & IMPACT

Topology-aware methods on Blue Waters have shown the potential to greatly reduce the dominant communication costs of sparse matrix-vector multipliers (SpMVs) when many messages are sent between sets of nodes. Fig. 2 shows a large reduction in cost for SpMVs on coarse levels of an AMG hierarchy; in contrast, standard communication sends a large number of small messages.

The results highlight the value in taking advantage of processor layout and topology in irregular communication demands, such as those introduced through sparse matrix operations. As applications and data demands continue to grow in complexity and dimension, localizing communication will be critical to achieving efficiency and taking advantage of the full capacity of the network.

WHY BLUE WATERS

Blue Waters is an ideal platform for testing scalable algorithms for future machines. The method developed in this project could be extended to additional elements of the machine topology, including the socket level and also the Gemini hubs on the network. The scale and network type made Blue Waters a necessary component of the experimentation to support the algorithm development and modeling.

PERFORMANCE EVALUATION OF NEW ALGEBRAIC ALGORITHMS AND LIBRARIES

Allocation: Illinois/50 Knh
PI: Edgar Solomonik¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

High-performance parallel algorithms for numerical linear algebra play a crucial role in most large-scale computational science problems. We have made progress in two aspects of advancing these types of algorithms: (1) introduction and tuning of distributed-memory functional abstractions for tensor operations, and (2) development of new communication-avoiding algorithms for matrix factorizations. Blue Waters has allowed us to deploy and evaluate these new methods on a leadership-class computing platform. In particular, we have done performance studies on the use of distributed symmetric tensor contractions for the atomic-to-molecular orbital transformation (a key kernel in quantum chemistry computations), of parallel sparse matrix multiplication routines and their use in graph analytics (betweenness centrality), and are currently evaluating a novel communication-avoiding algorithm for QR factorization of rectangular matrices.

RESEARCH CHALLENGE

The biggest challenge facing the parallel scalability of methods in computational science is the overhead of moving data between processors. Our goal is to develop algorithms and libraries that minimize communication in the number of messages as well as in the amount of data moved. To do so in the most useful way, we leverage the ubiquity of numerical linear algebra in scientific computing, targeting the development of algebraic algorithms and libraries. Matrices and tensors (multidimensional matrices) provide high-level abstractions for data sets and transformations thereof. Our aim is to push to the limit the parallel scalability of fundamental computational kernels on matrices and tensors by formulating algorithms that have a provably minimal communication complexity, and evaluating them in large-scale execution on Blue Waters.

METHODS & CODES

Our work does not only study hypothetical algorithms but contributes directly to libraries that are available to application developers using Blue Waters or other supercomputing platforms. In particular, our research has focused on extending and tuning Cyclops Tensor Framework (CTF). This library provides distributed-memory support for sparse and dense tensors, automatically mapping contractions and other functions on these distributed data sets. CTF uses performance models to make

runtime mapping decisions, using autotuning to train the model parameters. After performing tuning at scale on a training suite of CTF driver-routines, we studied performance of two kernels from radically different domains: quantum chemistry and graph analysis.

RESULTS & IMPACT

We conducted a performance study of in-memory and out-of-core CTF versions of an atomic-to-molecular orbital (AO-MO) transformation, immediately showing scalability on problems that are comparable in scale to the largest previously executed. This transformation appears in many high-accuracy quantum chemistry methods and is the most expensive step in some newly proposed methods. This work served to provide preliminary results for a many-principal investigator interdisciplinary proposal focusing on catalysis in chemical reactions.

Additionally, we evaluated the performance of a betweenness centrality code—MFBC—that leverages sparse matrix multiplication functionality in CTF as well as its support for user-defined tensor element-types and functions. We were able to calculate centrality scores for some of the largest graphs publicly available, including the Friendster graph, which has 1.8 billion edges, leading to a paper in *Supercomputing '17*. Fig. 1 displays the parallel scalability of the CTF MFBC code in terms of millions of edge traversals per second.

Over the past year, we have also developed a new parallel algorithm for QR factorization as well as a parallel implementation thereof. The algorithm aims to realize a QR code that achieves optimal communication and synchronization complexity in theory and is efficient in practice. While the new algorithm is not asymptotically more efficient than the state of the art, it is substantially more simple and easier to implement (no algorithm with the same communication complexity has been implemented previously). The basic idea is to use the Cholesky–QR2 algorithm and leverage communication-optimal parallel Cholesky and matrix multiplication routines. The Cholesky–QR2 algorithm is numerically stable, so long as the matrix is reasonably well-conditioned. We plan to compare the performance of our implementation to the QR routine in ScaLAPACK and test its stability on very large matrices.

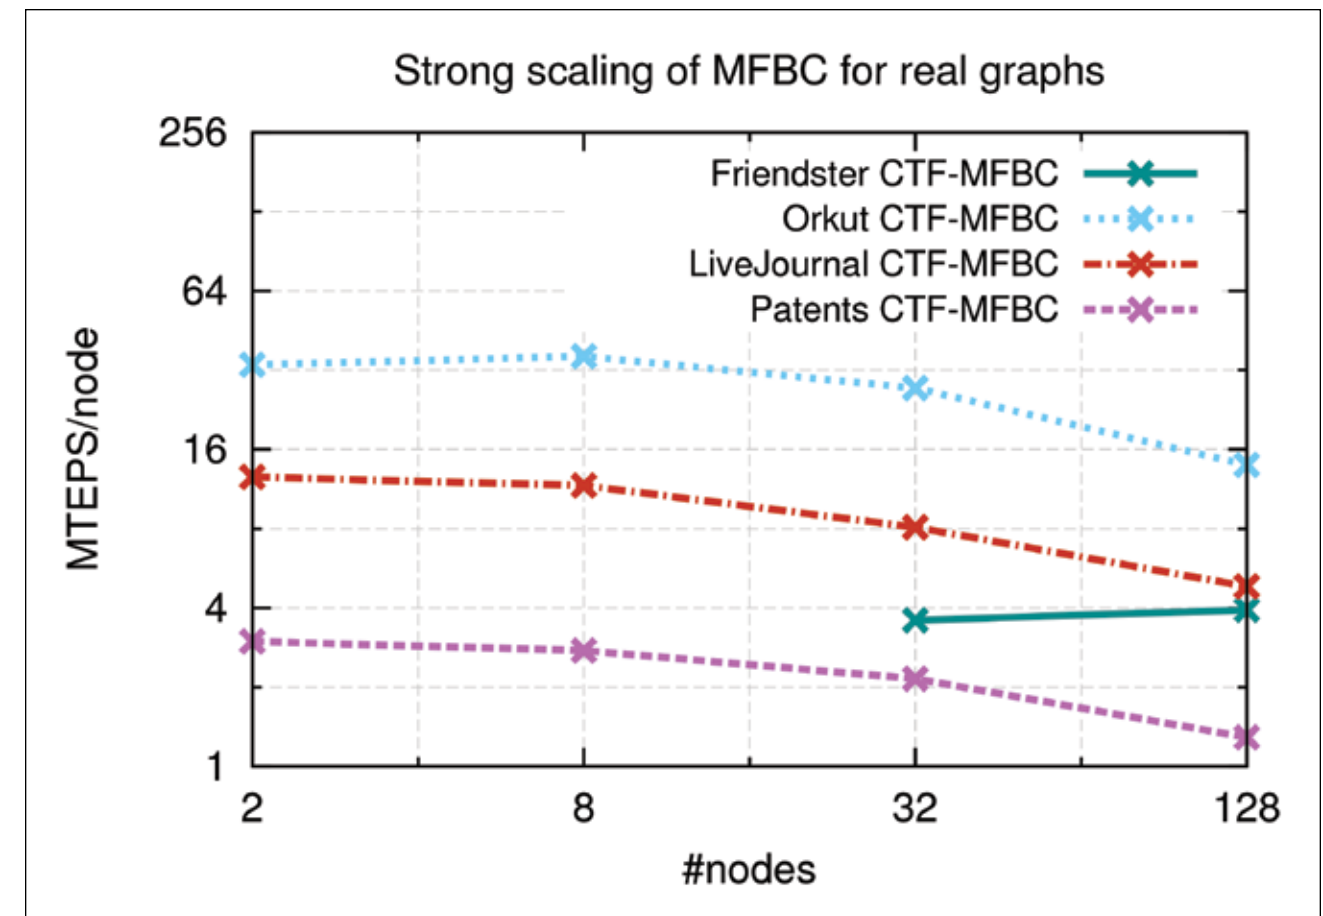


Figure 1: Parallel scalability of betweenness centrality using sparse matrix multiplication for massive graphs.

WHY BLUE WATERS

As our goal is to design software and algorithms for applications running on supercomputing systems, access to Blue Waters is essential for testing and evaluation. While all of the codes developed are designed to be portable, demonstrating performance on Blue Waters will help foster local collaborations and deployment of parallel numerical library software.

PUBLICATIONS AND DATA SETS

Solomonik, E., M. Besta, F. Vella, and T. Hoefler, Scaling betweenness centrality using communication-efficient sparse matrix multiplication. *ACM/IEEE Supercomputing Conference*, Denver, Colo., November 12–17, 2017.

Cyclops Tensor Framework: distributed-memory tensor algebra data sets, <https://github.com/solomonik/ctf>.

PARALLEL ALGORITHMS FOR BIG DATA PHYLOGENOMICS, PROTEOMICS, AND METAGENOMICS

Allocation: Illinois/125 Knh
PI: Tandy Warnow¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

This project addressed three interrelated problems in computational molecular biology, where large data sets present substantial computational and statistical challenges: phylogenomics (genome-scale phylogeny estimation), proteomics (protein structure and function prediction), and metagenomics (analysis of environmental samples from shotgun sequence data sets). Highlights of this project's activity include: (1) SVDquest, a method for species tree estimation from multilocus data sets that bypasses gene tree estimation (Tandy Warnow, with Ph.D. student Pranjal Vachaspati); (2) HIPPI: a method for protein family classification (Tandy Warnow, with Ph.D. student Mike Nute and two others); (3) an evaluation of the impact of screening genes in multilocus phylogenomic analyses (Tandy Warnow, with Ph.D. student Erin Molloy); and (4) an evaluation of statistical methods for multiple sequence alignment on protein benchmark data sets (Tandy Warnow, with Ph.D. students Ehsan Saleh and Mike Nute, and undergraduate researcher Kodi Collins). Five journal papers based on this work were published this year and another two were submitted.

RESEARCH CHALLENGE

This project aimed to develop methods for large-scale statistical estimation problems of phylogenies, multiple sequence alignments, and analyses of metagenomic data sets, where standard methods either do not run or provide poor accuracy. The need for new methods is particularly urgent as more and more studies attempt to analyze phylogenomic data sets with many thousands or tens of thousands of genes, and hence encounter massive gene tree heterogeneity, which can be due to multiple biological processes (incomplete lineage sorting, gene duplication and loss, horizontal gene transfer, etc.). The Genome 10K group is encountering these challenges in its plans to assemble phylogenies of the major groups of life on earth.

METHODS & CODES

We made progress on each problem using a combination of algorithmic approaches. In many cases we used divide-and-conquer, which allows powerful statistical off-the-shelf techniques to be applied to small subsets of a large data set, followed by innovative approaches to combine results from the small data sets.

RESULTS & IMPACT

One of the main outcomes of this project is HIPPI [1], a method for protein family classification. Family identification is a basic step in many bioinformatics pipelines, such as metagenomic taxon identification and abundance profiling (first steps in microbiome analysis) and is closely related to remote homology detection, which is a basic step in protein function and structure prediction. BLAST [2] is the most well-known method for this problem, but other approaches based on profile Hidden Markov Models (HMMs) have been used as well. In this work, we developed a novel machine-learning technique to detect membership in existing protein families, where we construct an ensemble of profile HMMs to represent each protein family, and then compare each sequence (which can be short reads or full-length sequences) to each HMM in each ensemble to find the best-fitting protein family. We provided an extensive study based on the PFAM [3] database of protein families and their associated profile HMMs from HMMER [4] to compare our method to the previous best methods. This study showed that the technique outperformed all the current methods (including BLAST, HMMER, and HHsearch [5]) in terms of both precision and recall, especially when analyzing short sequences (Fig. 1).

WHY BLUE WATERS

Blue Waters is necessary for at least two reasons. First, the development of these methods requires extensive testing, which is not feasible on other platforms. Second, the analysis of large biological data sets (and even of moderate-sized data sets) often requires years of CPU time (e.g., the avian phylogenomics project spent 450 CPU years to analyze approximately 50 whole genomes). Blue Waters makes this feasible and enables biological discovery.

PUBLICATIONS AND DATA SETS

Vachaspati, P., and T. Warnow, FastRFS: Fast and accurate Robinson-Foulds Supertrees using constrained exact optimization. *Bioinformatics*, (2016), DOI: 10.1093/bioinformatics/btw600.

Nute, M. and T. Warnow, Scaling statistical multiple sequence alignment to large datasets. *BMC Genomics*, 17 (Supplement 10):764 (2016), DOI: 10.1186/s12864-016-3101-8.

Nguyen, N., M. Nute, S. Mirarab, and T. Warnow, HIPPI: Highly accurate protein family classification with ensembles of HMMs. *BMC Genomics* 17 (Supplement 10), 765 (2016), DOI: 10.1186/s12864-016-3097-0.

Boyd, B. M., et al., Phylogenomics using Target-restricted Assembly Resolves Intra-generic Relationships of Parasitic Lice (Phthiraptera: Columbicola). *Systematic Biology*, (2017), DOI: 10.1093/sysbio/syx027.

Allen, J.M., et al., Phylogenomics from Whole Genome Sequences Using aTRAM. *Systematic Biology* (2017), DOI: 10.1093/sysbio/syw105.

HIPPI: <https://github.com/smirarab/sepp>, a github site maintained by Siavash Mirarab (former student).

FastRFS: <https://github.com/pranjalv123/FastRFS>, a github site maintained by Pranjal Vachaspati (current Ph.D. student).

PASTA+BAli-Phy: <https://github.com/MGNute/pasta>, a github site maintained by Michael Nute (current Ph.D. student).

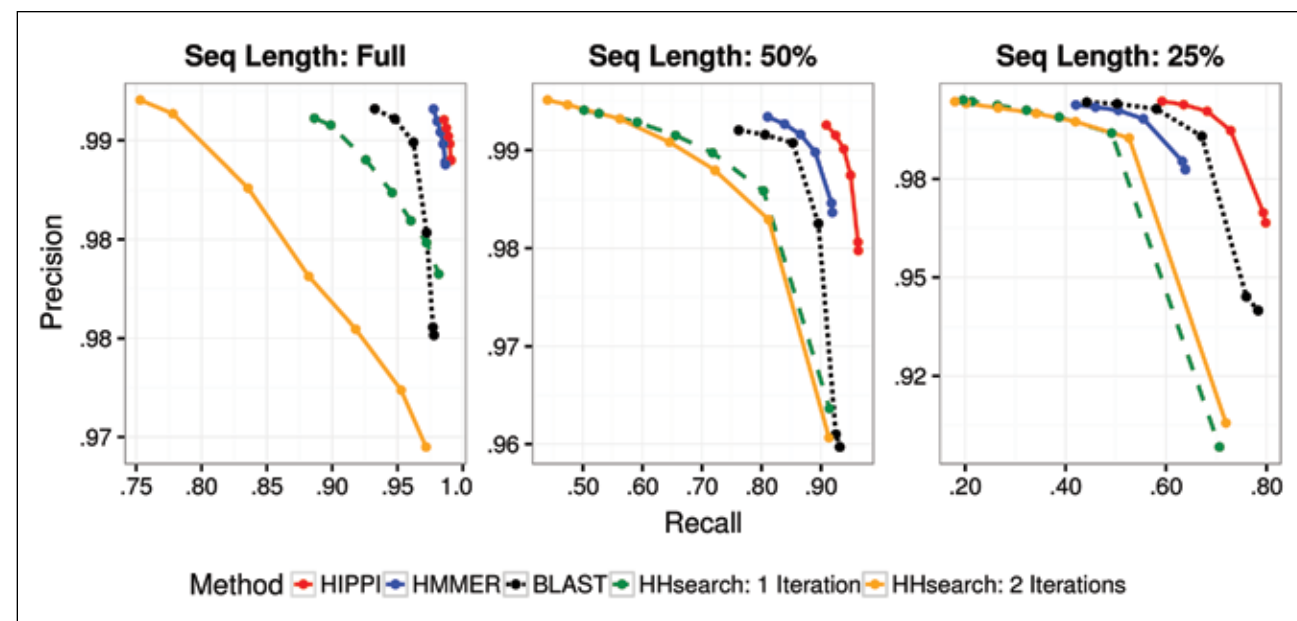


Figure 1: Precision-recall curves for methods for protein family classification, evaluated on one cross-fold subset of the PFAM seed sequence data. Our new method—HIPPI (shown in red)—strictly dominates all the other methods. Figure taken from [3].

HARDWARE ACCELERATION OF DEEP LEARNING

Allocation: Illinois/50 Knh
PI: Tao Xie¹
Co-PI: Yuan Xie²

¹University of Illinois at Urbana-Champaign
²University of California, Santa Barbara

EXECUTIVE SUMMARY

Our project aims to use the Blue Waters platform for hardware acceleration of deep learning for big data image analytics. To achieve near real-time learning, efforts must be paid to both scaling hardware out (increasing the number of compute nodes in a cluster) and scaling up (improving the throughput of a single node by adding hardware accelerators). In this work, we evaluated the performance of scaling up using the GPU-enabled node (XK7) for training convolutional neural networks. The key observation we obtained is that implicit data synchronization across different nodes severely limits the training process. We propose a data manager that explicitly overlaps the data transfer overhead with computation. In the first step, we test the proposed strategy on a single Blue Waters XK7 node. Experimental results show that this strategy achieves a speedup of 1.6X over the implicit data transfer implementation.

RESEARCH CHALLENGE

Deep learning has been widely used in applications such as image classification, speech processing, and object recognition. The huge amount of training data required by the deep neural networks requires more computing power to keep pace of the advance in the state-of-the-art accuracy of these tasks. Mainstream deep learning facilities are CPU-based clusters, which usually consist of thousands of compute nodes. Because the major computation step in deep learning is convolution and matrix multiplication, which is suitable for Graphic Processing Units (GPUs) to compute, modern deep learning facilities are often equipped with GPUs as hardware accelerators.

However, the straightforward implementation of deep neural networks on such GPU-enabled compute nodes will lead to underutilization of compute resources, especially for multi-node systems such as Blue Waters. Therefore, there is a strong

motivation to evaluate and characterize the deep learning workload on the GPU-enabled nodes.

In this work, we evaluated the performance of popular types of deep neural networks on a GPU-enabled supercomputer (the XK7 nodes on Blue Waters). From the evaluation results, we observed good scalability of neural networks. Meanwhile, among the three types of neural network layers we evaluated, that is, convolutional layers, fully connected layers, and long short-term memory (LSTM) [1] layers, the convolutional layers have the best scalability. The difference among the three types of network layers in terms of scalability comes from the variation of computation per byte in each layer. Given the same number of weights, the convolution layers have a one order magnitude larger number of multiply-accumulation (MAC) operations since the computation complexity of convolution is higher than matrix multiplications. Furthermore, the convolutional layers employ the weight sharing technique, which dramatically increases the computation per byte of the network.

Based on these observations, we continue to explore the design space of mapping different kinds of neural networks onto the GPU-enabled supercomputer. In real-world data centers, there are numerous neural network-based applications running concurrently. Since the optimal number of nodes allocated for each type of neural networks varies, we should design a scheduling method to achieve the best efficiency. In the next generation of work, we will conduct more application characterization on the multiple-type neural network workload on Blue Waters.

METHODS & CODE

To evaluate the performance of different types of neural networks, we chose a popular neural network, AlexNet [2], for reference of the convolutional layer and fully connected layer topologies. AlexNet has one convolution layer of (224, 3, 11—the numbers are the size of input image, the number of channels and the size of filter kernels); one convolution layer of (55, 96, 5); one convolution layer of (27, 256, 3); and two convolutional layers of (13, 384, 3). The sizes of the fully connected layers in AlexNet are 4,096; 4,096; and 1,000, respectively. For the topology of LSTM RNNs, we chose a character-based language model of which all recurrent layers have 128 neurons. Since all the LSTM layers are the same, we use only one LSTM layer to run the experiment.

We implemented these neural network layers based on DeepBench, which is a performance benchmark for deep learning hardware accelerators. We modified DeepBench to change the OpenAPI originally used to the platform API of Blue Waters. All nodes allocated are XK7 GPU-enabled nodes.

RESULTS & IMPACT

Fig. 1 shows the running time of each type of neural network on different numbers of nodes. In the figure, we can see that the running time of all three types of layers reduces along with the increase of the nodes used in parallel. The Blue Waters system shows good scalability, although there is communication overhead that makes the speedup sublinear. From the figure, we observe that the speedup of different types of neural networks is different since they have different computation per byte. This observation indicates that we cannot achieve the best performance or system efficiency if we use one single resource allocation scheme for all three types of neural networks. For example, communication dominates the latency for LSTM layers and fully connected layers in the case where we allocate eight nodes, while convolutional layers are still computation-bound. Based on this, we will design new resource allocation and algorithm mapping techniques to achieve better system performance given a fixed amount of workload.

WHY BLUE WATERS

Blue Waters offers us an opportunity to do research on the optimization of deep learning on computational clusters with GPUs. Blue Waters' XK7 nodes, which consist of one AMD eight-core CPU and one NVIDIA K20 GPU, allows studying of scaling up the computation per node through the addition of GPUs. As GPUs are more suitable than CPUs for convolution and matrix multiplications, which are the major computation in deep learning, state-of-the-art deep learning facilities widely employ GPUs as their hardware accelerators.




Figure 1: Performance Evaluation of Different Layer Types

A CRITICAL EVALUATION OF THE OP2/OPS PARALLEL MESHING AND CODE GENERATION SOFTWARE

Allocation: Industry/8 Knh

PI: Kevin Olson¹

Co-PIs: Reid Atcheson¹, Edward Smyth¹, Mike Dewar¹

¹Numerical Algorithms Group

EXECUTIVE SUMMARY

OPS/OP2 is a software package designed to construct algorithms that are easier to implement, maintain, and perform well on different computer architectures, including symmetric multiprocessing (SMP), clusters and clusters of graphics processing units (GPUs). This is accomplished by a programming model where the user defines stencils for performing operations, defining data structures, defining the data and how it is stored on the data structures, and constructing parallel kernels that operate on the data. Once these steps are completed, OPS and OP2 will construct underlying code for different parallel programming paradigms that include sequential, MPI, OpenMP, CUDA, MPI+OpenMP, and MPI+CUDA. Both C/C++ and Fortran programming languages are currently supported by OPS and OP2.

RESEARCH CHALLENGE

The challenge of this research was to test the OPS/OP2 code generation software to determine how easy it is to use and if it can generate code that scales well on large, parallel computer architectures.

METHODS & CODES

OPS/OP2 is software for constructing parallel computer programs [1,2]. OPS focuses on algorithms using structured meshes while OP2 is used for unstructured meshes. Using OPS/OP2 code involves writing either Fortran or C/C++ source code that embeds OPS/OP2 function calls that abstract out the parallelism in the algorithm. A code generation script is then run on the user-written code that generates source code for different parallel programming paradigms that include MPI, OpenMP, and CUDA as well as any combination of them.

RESULTS & IMPACT

OPS was used to construct a finite volume, compressible hydrodynamics code. OP2 was used to construct a Laplace equation solver using an unstructured mesh. The main result was that OPS and OP2 can produce code that runs using all the parallel programming paradigms listed above. Scaling was good in all cases, but MPI-“only” code gave the best performance. Using OPS and OP2 requires an investment in learning its syntax and use.

OPS and OP2 could be very useful tools for developing parallel algorithms since it would ease the amount of maintenance work for developing parallel algorithms. We note that OPS and OP2

are continuing to be developed and have evolved several versions ahead of those used for this study.

Whether OPS/OP2 can be used in further work at Numerical Algorithms Group is still under evaluation.

WHY BLUE WATERS

Blue Waters was chosen because Numerical Algorithms Group is one of its corporate partners and Blue Waters has the right combination of hardware available: a large number of SMP and GPU nodes that allowed us to perform the required scaling studies to evaluate OPS and OP2.

BIOLOGY, CHEMISTRY & HEALTH

BIOPHYSICS

CELLULAR

MEDICINE

MOLECULAR

NEUROSCIENCE

PHYSICAL CHEMISTRY

PUBLIC HEALTH

- 192** *DNA Origami Membrane Channels*
- 194** *Molecular Mechanism of Nuclear Transports*
- 196** *Improving NWChem Scalability Using the DataSpaces Framework*
- 198** *How Function Shapes Dynamics in Protein Evolution*
- 200** *Cellulosome Structure Determination by Atomistic Simulations Combined with Experimental Assays*
- 202** *Mechanism of Temperature Sensitivity in TRPV1 Channel*
- 204** *Exploring the Structure and Dynamics of Converged Ensembles of DNA and RNA Through Molecular Dynamics Simulations*
- 206** *Structural Basis for Extreme Cold Tolerance in the Eye Lenses of Teleost Fishes*
- 208** *Predicting Protein Structures with Physical Petascale Molecular Simulations*
- 210** *Multiscale Modeling of Biofilm Dynamics in Drinking Water Distribution Systems: Toward Predictive Modeling of Pathogen Outbreaks*
- 212** *Comprehensive in silico Mapping of DNA-Binding Protein Affinity Landscapes*
- 214** *Non-Born–Oppenheimer Effects Between Electrons and Protons*
- 216** *Brueckner–Goldstone Quantum Monte Carlo*
- 218** *How Membrane Organization Controls Influenza Infection*
- 220** *A Hybrid Stochastic-Deterministic Simulation Method Enables Fast Simulation of Cellular Processes in Eukaryotes*
- 222** *Quantum-Classical Path Integral Simulation of Charge Transfer Reactions*
- 224** *Patient-Specific HPC Models and Simulation-Based Imaging for Cardiovascular Surgical Planning*
- 226** *Quantum Effects of Proton Transfer in Biological Systems*
- 228** *Thermodynamic Characterization of Conformational Landscape in Proton-Coupled Oligopeptide Transporters*
- 230** *Machine Learning Reveals Ligand-Directed Conformational Change of μ Opioid Receptor*
- 232** *Elucidating the Molecular Mechanism of C-type Inactivation in Potassium Channels*
- 234** *Studying Cellular Processes through the Computational Microscope*
- 236** *Understanding the Protein Allostery in Kinases and GPCRs*
- 238** *Data-Driven, Biologically Constrained Computational Model of the Hippocampal Network at Full Scale*
- 240** *Stretching the Cadherin Molecular Velcro® of Cell-Cell Junctions*
- 242** *Simulation of Viral Infection Propagation through Air Travel*
- 244** *High-Resolution Magnetic Resonance Imaging of Mechanical Properties of the Brain*
- 246** *A Crystal Ball of Bacterial Behavior: from Data to Prediction using Genome-Scale Models*
- 248** *Large-Scale Coarse-Grained Molecular Simulations of the Viral Lifecycle of HIV-1*

DNA ORIGAMI MEMBRANE CHANNELS

Allocation: Blue Waters Professor/240 Knh

PI: Aleksei Aksimentiev¹

Collaborators: Chen-Yu Li¹, Jejoong Yoo¹, Kerstin Göpfrich², Iwona Mames³, Satya Prathyusha Bhamidimarri⁴, Maria Ricci², Adam Mames³, Alexander Ohmann², Mathias Winterhalter⁴, Eugen Stulz³, Bertalan Gyenes², Ulrich F. Keyser²

¹University of Illinois at Urbana-Champaign

²University of Cambridge

³University of Southampton

⁴Jacobs University Bremen

EXECUTIVE SUMMARY

DNA nanotechnology utilizes self-assembly with nanometer precision for the high-throughput construction of sub-micron-size objects. In comparison to conventional nanofabrication approaches, the DNA origami method is relatively low cost, easy to use, and has an infinite number of possible applications. Using Blue Waters, we have performed landmark molecular dynamics (MD) simulations to characterize the structure and transport properties of two biomimetic DNA origami channels—the smallest [1] and largest [2] DNA channels ever made—working in collaboration with the experimental Keyser lab (University of Cambridge). Once the technology is perfected, the DNA channels could be used to replace biological membrane channels or to deliver drugs across cellular membranes.

RESEARCH CHALLENGE

Membrane protein channels are biological sensors with high selectivity and efficiency. One important avenue of medical research is building a synthetic channel that has the same functionality as a biological channel or that performs a user-defined role. Recently, researchers demonstrated that DNA origami-based channels could mimic the ionic conductance and transport properties of membrane protein channels [3–8]. Only after characterizing their structural and electrical properties can these DNA channels be applied to biosensing and drug delivery.

METHODS & CODES

We performed explicit-solvent all-atom MD simulations with the latest version of NAMD2 [9–10] of the smallest and largest synthetic DNA channels ever designed, complementing the experimental work of our collaborators in the Keyser Lab. The smallest DNA channel was built using a single DNA helix, and the largest was a megadalton funnel-shaped DNA origami porin (a protein channel that allows passive diffusion). Chemical tags were added to stably embed the channels in a lipid bilayer. Simulations with an applied electric field were then performed to measure the ionic conductance of each channel.

RESULTS & IMPACT

Through all-atom MD simulations, we have shown that a membrane-spanning single DNA helix decorated with chemical

tags can provide a pathway for ions across the lipid membrane despite the lack of an internal physical channel. Lipid molecules were found to rearrange around the helix, forming a narrow water-filled passage at its circumference, allowing ions and water molecules to pass through the membrane. The average conductance calculated from simulation was in excellent agreement with experiment, and the simulations provided a microscopic explanation for the large variation in ionic conductance measured in experiment and simulation. Results of the single-DNA helix channel were published in *Nano Letters* [1].

Following the same all-atom MD approach, the conductance of the large funnel-shaped DNA channel was measured to be an order of magnitude larger than any previous man-made channel, and its cross-sectional area was similar to that of the nuclear pore complex. Consistent with the results of the single DNA helix, the ion current was found to flow through both the central pore of the channel and along the channel's walls. Results of the large funnel-shaped DNA channel were published in *ACS Nano* [2].

This work could lead to important applications at the frontier of medical science. Researchers could use synthetic DNA channels as a syringe for specific drug molecules by modifying the channels to recognize selective tissues and to open up the membrane. Furthermore, synthetic channels could be used in artificial tissues to give neighboring cells a new way to communicate.

WHY BLUE WATERS

Explicit-solvent all-atom MD simulation is the only computational method that can treat DNA origami objects enhanced by nonstandard functional groups and accurately characterize their structural fluctuations and transport properties [11]. Because of the size of the DNA origami structures, such MD simulations are computationally demanding. The large number of XK nodes on Blue Waters with graphics processing unit accelerators connected by the fast Gemini interconnect make it one of the best publicly available systems for performing DNA origami simulations. Over the past several years, our group has used Blue Waters to carry out a set of landmark simulations in the area of DNA nanotechnology, bringing high-performance simulations to the forefront of this research field.




Figure 1: The smallest and largest synthetic DNA channels (in blue and yellow) embedded in a lipid bilayer (in green and gray). Chemical tags were used to anchor the channels in place (shown in red).

PUBLICATIONS AND DATA SETS

Göpfrich, K., et al., Ion channels made from a single membrane-spanning DNA duplex. *Nano Letters* 16:7 (2016), pp. 4665–4669.

Göpfrich, K. and Li, C., et al., Large-conductance transmembrane porin made from DNA origami. *ACS Nano* 10:9 (2016), pp. 8207–8214.

MOLECULAR MECHANISM OF NUCLEAR TRANSPORT

Allocation: Illinois/986 Knh

PI: Aleksei Aksimentiev¹

Collaborators: David Winogradoff¹, Wei Si^{1,2}, Cees Dekker³

¹University of Illinois at Urbana-Champaign

²Jiangsu University

³Delft University of Technology

EXECUTIVE SUMMARY

The nuclear pore complex (NPC) regulates the transport of all ribonucleic acid (RNA) and proteins across the nuclear envelope of eukaryotic cells. Large molecules can only cross the nuclear pore when escorted by nuclear transport proteins known to interact with the central element of the NPC—nucleoporins, or nups. Several mutually exclusive theoretical models have been proposed to account for selective transport across NPCs; however, the microscopic mechanism remains largely unknown. Using Blue Waters, we have performed molecular dynamics (MD) and Brownian dynamics (BD) simulations to characterize the architecture and electrical properties of the central mesh of the NPC, essential to understanding nuclear transit.

RESEARCH CHALLENGE

The NPC serves as a gatekeeper, controlling the molecular traffic into and out of the nucleus. The central channel of the NPC is composed of the intrinsically disordered nups, which are known to be key to diffusion across the nuclear pore. The NPC is challenging to investigate *in vivo*. Hence, experimentalists developed synthetic pores that can faithfully mimic the selective transport of proteins through NPCs, while providing control over the pore dimensions and surrounding conditions [1–2]. Furthermore, advances in electron microscopy over the past several years have greatly increased our knowledge of the NPC scaffolding [3–4], but the structural architecture of the central channel remains elusive. The simulations performed on Blue Waters will offer new insight

into the physical mechanism of nuclear transport, with important implications for several human diseases and the development of novel gene therapies.

METHODS & CODES

We performed explicit-solvent all-atom MD simulations with the latest version of NAMD2 [5–6] of a solution of nup fragments and varied the protein volume fraction within a confined volume, analogous to varying synthetic pore diameter in experiment. Simulations with an applied electric field were then performed to calculate ionic conductivity as a function of nup density. In addition, we used atomic-resolution Brownian dynamics (ARBD) and traditional all-atom MD to model the entire NPC, including the disordered central channel.

RESULTS & IMPACT

Through an array of all-atom MD simulations, we characterized the structural fluctuations and electrical properties of a solution of nups varying in protein density. We measured the ionic current of each system in the presence of an applied electric field and observed a transition from conducting to not-conducting ions at a critical nup density, matching a key result seen in experiment [2].

Building upon an atomic model of the NPC scaffold mapped through experiment just last year [4], coarse-grained ARBD simulations were performed to determine the architecture of the central channel. We then built and performed MD simulations of the first complete all-atom model of the NPC—140 M atoms in total—including the outer scaffolding, the disordered inner channel, the nuclear envelope, and the surrounding solvent (shown in fig. 1). This system is one of largest ever modeled with NAMD on Blue Waters to date.

This work could lead to important medical applications. A number of human diseases—including cancer, viral infections, and neurodegenerative diseases [7–9]—are caused by disturbances in nuclear pore transport, and this project may offer new insights into the molecular origin of those diseases and have implications for the development of gene therapy treatments.

WHY BLUE WATERS

Explicit-solvent all-atom MD simulations are needed to characterize the structural fluctuations and electrical properties of the disordered central channel of the NPC. Because of the immense size of the entire NPC—140 Million atoms in total—such MD simulations are only possible on a supercomputer with the computational power of Blue Waters. Furthermore, the large number of XK nodes on Blue Waters with graphics processing unit accelerators connected by the fast Gemini interconnect makes it one of the best publicly available systems for performing simulations of the entire NPC.




Figure 1: An atomic model of the entire nuclear pore complex, 140 M atoms in total (water omitted for clarity). Colors highlight the outer scaffolding (in cyan), the disordered central channel of nups (in purple), and the nuclear envelope (in green).

IMPROVING NWCHEM SCALABILITY USING THE DATASPACE FRAMEWORK

Allocation: Innovation and Exploration/500 Knh

PI: Gregory Bauer¹

Collaborators: Victor Anisimov¹, Manish Parashar², Melissa Romanus²

¹National Center for Supercomputing Applications

²Rutgers University

EXECUTIVE SUMMARY

The purpose of this project is to speed up coupled cluster CCSD(T) computation in NWChem and then incorporate this highly reliable level of theory into parameter optimization procedures for the AMBER classical force field for DNA, with the potential in future of extending this approach to optimization of other popular force fields. The development approach relies on revising the most memory-demanding part of the CCSD algorithm with the possibility of offloading the largest memory arrays to a DataSpaces server. The purpose of using CCSD(T) computation in the parameter optimization procedure is to employ the wealth of available experimental data, which is typically underutilized in common parameter-refining techniques. Moving the parameter optimization engine to a supercomputing platform, and automating the entire process of data extraction and parameter refining, makes the process reproducible, extendable, and portable.

RESEARCH CHALLENGE

Molecular Dynamics (MD) simulations based on classical force fields are the major tool to breach the gap between experiment and theory in materials science, engineering, and biomedical research. However, the predictive ability of MD simulations depends heavily on the quality of the underlying parameters. The limited pool of experimental data and the extremely labor-intensive nature of parameter optimization represent the major limiting factors in improving the quality of classical force field parameters. Theoretically, the use of highly reliable electronic structure computations at the CCSD(T) level of theory could help mitigate the lack of experimental data in parameter optimization. However, the high computational cost of the CCSD(T) method precludes its routine use even on modern supercomputing platforms. Therefore, faster and more computationally efficient CCSD(T) implementations are greatly needed. Additionally, the labor-intensive nature of classical force field determination, and the manually driven optimization, makes it difficult to reproduce the work, to revise and improve the protocol, and to introduce a quality control into the optimization process. Because parameter optimization is a computationally demanding effort, the efficient use of supercomputing resources becomes an additional challenge.

METHODS & CODES

The work on performance optimization of the CCSD(T) method in NWChem, which is maintained at Pacific Northwest National Laboratory, relies on offloading large memory arrays from compute nodes to a dedicated data storage employing the DataSpaces data management framework developed by the team headed by Manish Parashar at Rutgers University. While this work is in progress, the previously optimized version of the CCSD(T) method [1] in NWChem [2] has been used to generate a database of experimental data on intermolecular interactions from X-ray crystallographic data of molecular crystals of DNA bases. An automated procedure was developed to extract unique intermolecular pairs from the crystal data, to relax hydrogen atoms, to compute CCSD(T) interaction energies, and to arrange the data in a structured form. A scalable tool to optimize Lennard-Jones parameters in the AMBER (Assisted Model Building with Energy Refinement) force field to fit the parameters to intermolecular interaction energies for experimental geometry of monomers has been developed. It has been tested to run on 1,000 nodes using 32,000 processing units on Blue Waters. The optimization tool generates 200 alternative parameter sets of comparable quality determined by the fitting criterion. The number of alternative parameter sets is easily adjustable in this method. The optimized parameter sets proceed to the final validation step running in parallel to perform MD simulations, to compare the geometry of the system to that from the X-ray data, and to reproduce experimental heat of sublimation, unit cell volume, and interatomic distances.

RESULTS & IMPACT

This project introduces a new approach to parameter optimization for classical force fields that combines a high-level electronic structure calculation method to extract additional previously inaccessible information from experimental data, to reengineer the optimization procedure, and to tailor it to maximally utilize high-performance computational resources. The developed procedure shortens the time needed for parameter optimization by roughly a factor of 10. It resolves the issue of accessibility to only a few highly capable teams, transforms the force field optimization from an empirical to a well-structured discipline, educates, and makes the results easy to reproduce by the community.

WHY BLUE WATERS

Blue Waters, with its fast interconnect and large memory per core, is unique in its ability to conduct CCSD(T) computations of molecular systems encountering a thousand basis functions, which is vital for the success of the developed parameter optimization procedure. Since the parameter optimization procedure is extremely resource demanding, the availability of large numbers of nodes is essential for the exhaustive exploration of parameter space.

HOW FUNCTION SHAPES DYNAMICS IN PROTEIN EVOLUTION

Allocation: Illinois/350 Knh
PI: Gustavo Caetano-Anolles¹
Co-PIs: Frauke Gräter²
Collaborators: Fizza Mughal¹

¹University of Illinois at Urbana-Champaign
²Heidelberg Institute for Theoretical Studies (HITS)

EXECUTIVE SUMMARY

Protein loops are found to be chiefly responsible for the wide functional diversity of proteins. This stems mostly from the property of protein flexibility, which is evolutionarily conserved. This suggests specific molecular motions may have been selected for individual functions. By harnessing the power of Blue Waters, we aim to decipher patterns and processes underlying the origin, evolution, and structure of the molecular dynamics of proteins. In our current allocation, we have completed 116 molecular dynamics simulations of loop regions of protein structural domains found in metaconsensus enzymes. In addition, we completed aminoacyl-tRNA synthetase (aaRS) simulations that were pending from a previous allocation. The collected data were subjected to a preliminary analysis of molecular trajectories. Variables were computed that described the dynamic properties in these trajectories permitted to construct a dynamics space, a “dynamosome,” that we intend to map onto a “structure–evolution” protein space.

RESEARCH CHALLENGE

The biophysical properties of protein loops may hold answers to the discovery, prediction, and annotation of protein functions [1]. Some of these biophysical properties may be governed by yet-to-be discovered evolutionary drivers that could significantly impact synthetic biology and translational medicine [2]. Incorporation of biophysics in protein structure–function studies is becoming increasingly common. However, biophysics is rarely used in evolutionary studies [3]. The main objective of our studies is to bridge disparate disciplines of biology and physics with Molecular Dynamics (MD) simulations performed at nanosecond (ns) timescales to capture evolutionary dynamics on a scale of billions of years. Here, we explore biophysical variables of the MD simulations by studying community structures of protein loop residues that describe the molecular trajectories of the loop regions. Our goal is to dissect evolutionary relationships in these data using evolutionary timelines reconstructed from robust phylogenomic methods [4].

Figure 1: A: Distribution of functional annotations across 86 of 87 loops that exhibited a community dynamics structure. A total of 72 out of the 87 loops examined possessed (single or multiple) Gene Ontology (GO) functional annotations. B: Principal Component Analysis of Protein loop 1B7Y_B_408 associated with the a.6.1.1 SCOP domain. C: Protein loop 1B7Y_B_408; residues connected to each other due to positive (red) and negative (blue) correlations of motions during MD trajectory. D: All-residue network of 1B7Y_B_408 with community structures highlighted as groups. E: Dynamical Cross Correlational Map of the protein residues of 1B7Y_B_408. Featured in *Blue Waters Annual Report* (Urbana, Illinois, 2016).


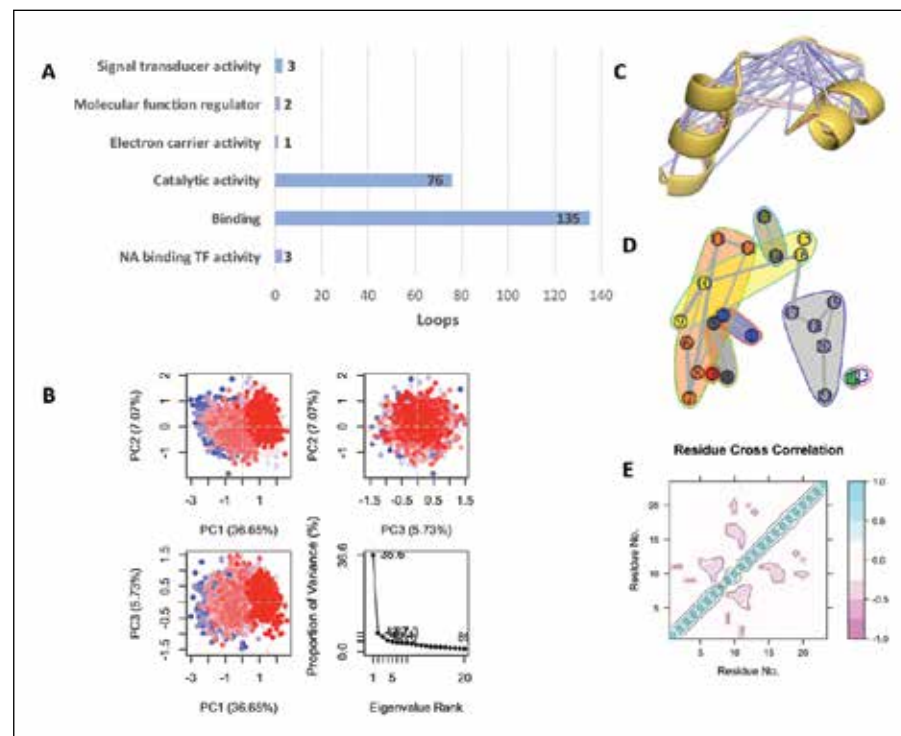


Figure 2: The community dynamics network morphospace of aaRS protein. A Pareto front defines a Zooko-like triangle of trade-off solutions among economy, robustness and flexibility for protein dynamics.

METHODS & CODES

We completed MD simulations that were left pending in a previous allocation, in which we analyzed 87 protein loops from aaRS structural domains on a timescale of 50–70 ns. In addition, we simulated 116 protein loops belonging to single-domain metaconsensus enzymes. The protein loops of aaRS domains were mostly associated with the Gene Ontology (GO) level-1 molecular function of “binding” followed by that of “catalytic activity” (Fig. 1.A). We constructed a dynamics space, a modified version of the dynamosome [5], by calculating the eigenvalues of the top five principal components from principal component analysis (Fig. 1.B) and centrality metrics from a network (Fig. 1.D) based on the dynamic cross-correlation matrix of the motions of protein residues (Fig. 1.E). In order to assess the presence or absence of a specific network topology, we also calculated maximum modularity scores, alpha values to test power law behavior, and Bartel’s test statistic for measuring the extent of modularity, scale-freeness, and randomness of the network (Fig. 2). We are currently in the process of performing unsupervised clustering of the trajectories using the dynamosome variables. We also plan to use methods that classify community structure patterns (Fig. 1.D) exhibited by loops and their correlation to specific function. Our goal is to reconstruct a “structure–evolution” space that would complement our dynamosome.

RESULTS & IMPACT

The aim of our investigation is to detect the presence or absence of patterns of motion in molecules. We focus on an analysis of dynamic network topologies defining a three-dimensional morphospace delimited by the conceptual axes of modularity, scale-freeness, and randomness [6]. Modularity, a feature persistently

observed in biological networks [7], embodies flexibility and diversity of the molecular components that make up the whole. Scale-freeness is an indicator of heterogeneity in patterns of connectivity of the network. It is a measure of “economy” (i.e., how easy it is to traverse the network structure). Randomness entails uniform connectivity of nodes throughout the network, a property that confers network fault tolerance. Fig. 2 shows network topology tendencies for 72 of the 87 protein loops that have been annotated with GO functions. This plot encapsulates trade-offs among flexibility, economy, and robustness that result in a “noisy” 2-polytope Pareto front. Interestingly, Zooko’s triangle, a concept used in the design of internet domain-name systems [8], can be thought to illustrate the design space of community networks obtained from protein dynamics. Protein loops possessing various functions tend to cluster at the center of this “triangle” (Fig. 2). They also prefer enhancing modularity in their quest to seek temporal persistence.

WHY BLUE WATERS

The petascale competencies of Blue Waters have been of great value to our project. With the help of such a state-of-the-art system, we have been able to achieve our goal of simulating a significantly large number of proteins (each at a timescale of ~70 ns) in a short time period. NAMD scales well on the Blue Waters architecture, especially when combined with GPU (graphics processing unit) nodes. This provides a significant boost in acceleration [9]. Apart from system specifications that are well-suited to our project, the domain experts/scientists in the Blue Waters support team have helped us smooth out any technical issues that have arisen during the current and previous allocations.

CELLULOSOME STRUCTURE DETERMINATION BY ATOMISTIC SIMULATIONS COMBINED WITH EXPERIMENTAL ASSAYS

Allocation: Illinois/680 Knh

PI: Isaac Cann¹

Co-PI: Rafael C. Bernardi¹

Collaborators: Klaus Schulten¹, Edward Bayer², Hermann Gaub³, Michael Nash⁴

¹University of Illinois at Urbana-Champaign

²Weizmann Institute

³Ludwig Maximilian University of Munich

⁴Universit of Basel

EXECUTIVE SUMMARY

Cellulosomes, which deconstruct cellulose and hemicellulose, are present in many anaerobic bacteria. Cellulosomal cohesion-dockerin interactions enable the incorporation of catalytic cellulases and hemicellulases onto the cellulosome scaffoldin. Their very efficient mechanism of degrading plant cell-wall biomass makes cellulosomes of interest to the second-generation biofuel industry, which aims to produce ethanol from agricultural waste. Furthermore, the recent discovery of cellulosomal bacteria in the lower gut of humans is paradigm-shifting as it has allowed demonstration of the capacity to degrade both hemicellulose and cellulose, at least, in the gut of some humans. Employing molecular dynamics simulations to complement single-molecule and biochemistry experiments, we characterized cellulosome's components, showing that even a single mutation can cause a large change in cellulosome structural stability.

RESEARCH CHALLENGE

Symbiotic bacteria greatly influence human health and play a significant role in pathogenesis, disease predisposition, physical fitness, and dietary responsiveness [1]. Moreover, bacteria play a key role in the second-generation biofuel industry where their cellulolytic enzymes are used for plant cell-wall degradation [2]. Here, we investigate key processes underlying bacterial activity, namely plant fiber metabolism. Specifically, we seek the structure and function of cellulosomes, the highly cooperative macromolecular complex that is central for this metabolic process in some bacteria [3].

Cellulosomes are multi-enzyme complexes that enable deconstruction of cellulose and hemicellulose in anaerobic cellulosome-containing bacteria. Integration of cellulosomal components occurs via highly ordered protein-protein interactions among three major components. In cellulosome assembly, a large noncatalytic polypeptide called the scaffoldin, embedded with various cohesins, anchors dockerin-containing enzymes through cohesion-dockerin interactions (see Fig. 1). Specificity of the cohesion-dockerin interaction allows incorporation of different catalytic cellulases and hemicellulases onto the scaffoldin, which may or may not be bound to another domain tethered to the

cell wall [4]. Cellulosome assembly promotes the exploitation of enzyme synergism because of spatial proximity and enzyme-substrate targeting.

Using stochastic search algorithms connected to MD tools, we are building the first comprehensive structure of a cellulosome including enzymatic domains. We expect that a complete model of a cellulosome's structure will shed light on the mechanism that allows these enzymatic complexes to be highly efficient.

METHODS & CODES

Combined with biochemical and single-molecule experiments, we employ molecular dynamics (MD), steered MD (SMD), and generalized simulated annealing (GSA) simulations on Blue Waters utilizing QwikMD [5], a graphical interface connecting Visual MD (VMD) and Nanoscale MD (NAMD). Using stochastic search algorithms coupled to NAMD we can generate thousands of different structure conformations for the cellulosome [6]. GSA [7] analysis shows that the different linkers between cohesin and dockerin (a carbohydrate-binding molecule) in cellulose-integrating protein A (CipA) scaffoldin assume a number of stable conformations. Small angle X-ray scattering analysis has previously shown that three conformations are observed for one of the cellulosomal linkers. GSAFold can predict these three conformations and all the other conformations for CipA. To perform this analysis, 20,000 conformations were obtained per linker and clustered. Combined, these linker conformations would give us 10^{43} CipA conformations. From clustering, we reduce this number to 3,888 structures that were obtained and subjected to a cluster analysis that gave rise to the five most significant structures.

RESULTS & IMPACT

Following well-established protocols for large macromolecular systems, and using one of the CipA conformations that we obtained using GSAFold, we built a first model of an entire cellulosome structure. MD simulations are now being employed to study the quaternary structure stability. Also, combining biochemical and single-molecule experiments with MD and SMD simulations, we investigated a series of cellulosomal cohesins from *Acetivibrio cellulolyticus*. We revealed that these cellulosomal components withstand different amounts of force depending on their position




Figure 1: Illustration of a complete cellulosomal structure over a bacterial cell surface. The cohesin-based scaffoldin (red) can be attached by a specific cohesion-dockerin interaction to a cell-anchoring domain (in orange). Another cohesion-dockerin interaction is responsible for attaching the enzymatic domains (blue) to the scaffoldin.

in the protein network. In this study, we combined one-step *in vitro* expression and specific covalent pulldown of protein constructs to assess the mechanical stability of highly related proteins in a parallel single molecular force spectroscopy assay. Using SMD simulations, we reproduced the experimental results and identified important amino acids. In addition, we used the simulations performed on Blue Waters to suggest mutations that were experimentally performed by site-directed mutagenesis, engineering proteins to pin down single crucial amino acids promoting force resilience.

WHY BLUE WATERS

Investigating the structure and functional processes of large enzymatic complex machineries such as the cellulosomes is only possible on petascale computing resources like Blue Waters. Structures obtained using enhanced sampling techniques such as GSA are only reliable if thousands of conformations (models) are predicted. Employing GSA for the numerous linkers of the cellulosome is a well-suited task for the large-scale parallel architecture of Blue Waters.

PUBLICATIONS AND DATA SETS

Schoeler, C., et al., Ultrastable cellulosome-adhesion complex tightens under load. *Nat. Commun.*, 5 (2014), p. 5635.

Schoeler, C., et al., Mapping mechanical force propagation through biomolecular complexes. *Nano Letters*, 15 (2015), p. 7370.

Bernardi, R.C., M.C.R. Melo, and K. Schulten, Enhanced sampling techniques in molecular dynamics simulations of biological systems. *Biochim. Biophys. Acta-Gen. Subj.*, 1850 (2015), pp. 872-877

MECHANISM OF TEMPERATURE SENSITIVITY IN TRPV1 CHANNEL

Allocation: NSF PRAC/6,510 Knh

PI: Vincenzo Carnevale¹

Collaborator: Tibor Rohacs²

¹Temple University

²Rutgers New Jersey Medical School

EXECUTIVE SUMMARY

TRPV1 is an ion channel crucially responsible for transduction of nociceptive stimuli into pain signals. Accordingly, inhibition of TRPV1 is one of the major strategies for designing next-generation antipain drugs. The polymodal nature of TRPV1 activation, i.e., the fact that a variety of stimuli can open the channel, suggests a complex molecular mechanism of activation whose details are still largely unknown. In particular, we do not currently have a satisfactory microscopic model to explain TRPV1 temperature sensitivity. In this project, we revealed the crucial role played by four nonpolar cavities whose presence and involvement in activation have not been described before. Free energy calculations show that dehydration of these cavities triggers activation of the channel. This observation is able to explain the puzzling response of TRPV1 to diverse environmental factors such as increased cytosolic hydrostatic pressure and osmolarity.

RESEARCH CHALLENGE

The nonselective channel TRPV1 is a crucial player in the human nociceptive system. TRPV1 is responsible for the detection of several harmful stimuli such as heat, low pH, and irritating chemicals that are transduced in painful signals originating from peripheral nerves. In chronic pain syndromes, this channel is hypersensitized and the threshold for activation is so low that innocuous stimuli result in constant pain. This involvement in

the pain pathway makes TRPV1 an appealing target for designing novel antipain drugs: selective modulation of this channel would inhibit the generation of the pain signal without interfering with other physiological pathways. Drug discovery campaigns aimed at this target hold promise to deliver pain killers that are virtually devoid of side effects. However, several promising molecules have failed in late stages of clinical trials. The reason for these failures is that TRPV1 is involved in body temperature regulation, and inhibitors of this channel might cause hyperthermia in patients. To make progress in identifying molecules that aptly modulate the channel without interfering with the mechanism of temperature sensitivity, a microscopic understanding of TRPV1 is needed. This project aimed at filling this gap by using long time-scale molecular dynamics (MD) simulations together with enhanced sampling techniques to estimate free energies.

METHODS & CODES

The structure of the TRPV1 capsaicin-bound (CAP-bound) state was taken from the Protein Data Bank: the PDB code is 3j5r [1]. The structure was refined and the missing residues were modeled using Rosetta software [2]. Four capsaicin molecules were docked following the protocol described in [3]. The protein with the ligands was embedded in a hydrated 1-palmitoyl-2-oleoylphosphatidylcholine (POPC) bilayer and surrounded by 150 mM NaCl solution. The overall size of the system was

$\sim 170 \times 170 \times 160 \text{ \AA}^3$; the total number of atoms was $\sim 400,000$. Two MD trajectories were generated with the peripheral cavities (PCs) either empty or hydrated. The CHARMM36 force field [4] was used to describe the protein and the POPC lipids. For capsaicin, we used the parameters derived in [3]. The TIP3P model was used to describe water [5]. An analogous setup was used to simulate the TRPV1 apo state (PDB code 3j5p [6]). The equilibration of the systems (three in total: the CAP-bound state with empty and hydrated PCs, and the apo state) was performed using NAMD 2.10 software [7] in several steps. Simulations were performed at constant temperature and pressure (1 atm) using the Langevin piston approach. For the Van der Waals interaction, we used a cutoff of 11 Å with a switching function between 8 and 11 Å. The long-range component of electrostatic interactions was calculated using the Particle Mesh Ewald approach [8] with a cutoff for the short-range component of 11 Å. The equations of motion were integrated using a multiple time-step algorithm, with a time step of 2 fs (femtoseconds) and long-range interactions calculated every other step.

We performed metadynamics simulations using the preliminary unbiased trajectories to estimate an upper bound for the free energy barrier and the diffusion constant along the biased collective variable. These were used to obtain an *a priori* estimate of the error on the reconstructed free energy profile using the expressions reported in [9], which relate the error to the width, height, and deposition rate of the hills. This estimate informed our choice of the metadynamics parameters. Metadynamics simulations were performed using the collective variable module implemented in NAMD 2.10 [10] at three temperatures: 280°K, 300°K and 340°K.

RESULTS & IMPACT

We found that the lower gate is open or closed depending on the conformation of N676, an amino acid located on S6 at the edge of a π -helix segment (Fig. 1). This segment is characterized by extreme conformational flexibility [11]: Not all the backbone hydrogen bonds can be simultaneously satisfied and therefore their pattern is dynamic. The presence of the π -helix allows N676 to easily rotate in and out the central pore. This motion is, in turn, controlled by the hydration state of the adjacent protein cavity (PC). Besides extensive accessibility experiments by Salazar, et al. [12], the presence of these cavities is supported by alanine scanning mutagenesis performed on the S6 segment [13]; the residues lining the PCs were shown to produce the greatest perturbation to the channel activation in response to several stimuli, including capsaicin and heat. Interestingly, our molecular mechanism does not entail any large conformational rearrangement of the TRPV1 central pore, whose radius profile is not dramatically altered by the closed-open transition. The hydration/dehydration of this compartment is, in fact, controlled by N676, which upon rotation changes the hydrophobic character of the molecular surface lining the central pore. This susceptibility to perturbations is not uncommon in the pores of ion channels. Wet-to-dry transitions have been reported several times [14] and

are, arguably, the result of a precise evolutionary optimization. This microscopic picture provides the basis for rational design of precise modulators of TRPV1.

WHY BLUE WATERS

We investigated a system of approximately 400,000 atoms using MD simulations. The time scales involved in the activation process of TRPV1 dictated trajectory lengths on the microsecond time scale. This was possible thanks to a massively parallel calculation enabled by the computational capabilities of Blue Waters.




Figure 1: TRPV1 pore domain with hydrated (left) and empty (right) protein cavities (PCs). Water in the PCs and in the central pore is shown in blue and cyan, respectively. In the state with hydrated PCs, there are two interruptions of water density, while with empty PCs the water density is continuous. N676 is shown in green.

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

Allocation: NSF PRAC/12,000 Knh

PI: Thomas E. Cheatham III¹

Co-PIs: Adrian Roitberg², Carlos Simmerling³, David Case⁴

Collaborators: Rodrigo Galindo-Murillo¹, Christina Bergonzo⁵, Daniel R. Roe⁶

¹University of Utah

²University of Florida

³Stonybrook University

⁴Rutgers University

⁵Institute for Bioscience and Biotechnology

⁶National Heart, Lung, and Blood Institute

EXECUTIVE SUMMARY

Over the past two years of using Blue Waters and taking full advantage of its computational power, we have made extensive progress in the realm of biomolecular simulation methodologies, specializing in nucleic acid structure, dynamics, and ligand–protein binding. The AMBER package for biomolecular simulation and its GPU (graphics processing unit) code has proven to be high-performance and reliable, taking full advantage of Blue Waters. As our group continues to develop specialized software and methodologies to analyze the vast amount of sampling information, we further increase our understanding of relevant biological DNA, RNA, and protein structures.

RESEARCH CHALLENGE

Molecular Dynamics (MD) simulations have been one of the most important tools in the computational chemist's toolbox for the last 25 years. Useful as it is, this technique has considerable limitations, mainly in two areas: force field validation and conformational sampling. These two problems are deeply related. As more sampling time is achieved, force field discrepancies are found. The force field is temporarily fixed; however, as computational power grows, more sampling time is achieved and new force field issues are found. Blue Waters has enabled our group to achieve enough sampling time to rigorously validate and assess protein, DNA, and RNA force fields in order to expose existing drawbacks.

METHODS & CODES

Blue Waters provides the necessary computing power to perform tests and benchmarks of several simulation methodologies. This allows us to increase the amount of space sampled for a particular biomolecular system and to create new enhanced sampling techniques in order to obtain a converged ensemble. Our group is focusing on two areas: multi-dimensional replica exchange (M-REMD) and ensemble simulations. In the case of replica exchange, we have explored multiple small RNA systems

(tetranucleotides, hairpins, loops, etc.) and we have achieved a converged ensemble that generated insight into successes and failures of force fields, ion models, water models, and modeling procedures. This information helps to pinpoint problem areas in the models used and to guide the next steps of research.

The other methodology we have used involves multiple independent copies of a particular biomolecular system, or ensemble simulations. This lets us explore increasing sampling space without the introduction of any biasing or enhanced mechanism. These simulations allow us to study the process of the DNA–ligand binding mechanism in fully atomistic ways, which provide further insight in order to design novel small molecules to increase biological activity. Of major interest is the study of a family of planar copper-compounds (with general formula $[\text{Cu}(\text{N}-\text{N})(\text{N}-\text{O})]\text{NO}_3$ and $[\text{Cu}(\text{N}-\text{N})(\text{O}-\text{O})]\text{NO}_3$; where the N–N ligand denotes either 2,2'-bipyridine or 1,10-phenanthroline (the aromatic ligand); N–O represents an essential amino acid or peptides; and the O–O represents a nonaromatic ligand—either acetylacetonate or salicylaldehyde) that experimentally show higher biological activity at a lower dosage with respect to the hallmark of transition-metal drugs: cisplatin. Using Blue Waters and several milliseconds of sampling time, we observed five principal binding sites (Fig. 1). Binding modes a) through e) represent the result of unbiased interactions of the ligands with the Drew Dickerson dodecamer, with the sequence GCGCAATTGCGC, using the bsc0 force field for nucleic acids. Binding mode a) represents stacking at the edges of the DNA, modes b) and d) are where the ligand binds into the minor groove of the helix. Mode c) starts in the minor groove and then the ligand moves into an AT base pair, pushing the AT bases toward the major groove as the ligand slides into the resulting cavity. Mode e) is an intercalated mode as a result of the terminal base pairs fraying, allowing the ligand to stack into the exposed bases.

RESULTS & IMPACT

Key results are described in detail in our publications (below).




Figure 1: Modes of binding between a 12-mer double-stranded DNA and a planar copper complex—a) stacking at the end of the DNA; b) and d) minor groove interaction; c) base-pair eversion mechanism; and e) intercalation at the terminal base pairs.

WHY BLUE WATERS

Because Blue Waters is optimized for GPU-centered calculations, as well as massive CPU parallelization, it has been a key tool for our group and allowed us to independently run hundreds of biomolecular simulations of myriad systems of DNA and RNA. In some cases, we have reached well over millisecond timescales of combined sampling time. In addition, the highly trained Blue Waters staff have been in close collaboration to develop custom and specialized code capable of performing the analysis of terabytes of simulation data.

PUBLICATIONS AND DATA SETS

<http://amber.utah.edu>

Galindo-Murillo, R., et al., Intercalation processes of copper complexes in DNA. *Nuc. Acids Res.*, 43 (2015), pp. 5364–5376.

Galindo-Murillo, R., et al., Assessing the current state of Amber force field modifications for DNA. *J. Chem. Theory Comp.*, 12 (2016), pp. 4114–4127, DOI: 10.1021/acs.jctc.6b00186.

Heidari, Z., et al., Using Wavelet Analysis to Assist in Identification of Significant Events in Molecular Dynamics Simulations. *J. Chem. Inf. Model.*, 56 (2016), pp. 1282–1291.

Bergonzo, C., and T. E. Cheatham, III, Mg^{2+} binding promotes SLV as a scaffold in Varkud Satellite Ribozyme SLI-SLV kissing loop junction. *Biophys. J.*, 113 (2017), pp. 313–320, DOI: 10.1016/j.bpj.2017.06.008.

STRUCTURAL BASIS FOR EXTREME COLD TOLERANCE IN THE EYE LENSES OF TELEOST FISHES

Allocation: Illinois/113.5 Knh
PI: Christina C.H. Cheng¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Eye lenses of endothermic mammals such as the cow develop cold cataracts at a mild 17°C. In contrast, ectothermic teleost fish lenses remain transparent down to -12°C. Cold-induced cataracts arise from a liquid-liquid phase-separation of lens proteins (crystallins) resulting in a protein-rich and a protein-poor phase. Crystallins are tightly packed at high concentrations to enable refraction of incident light, and teleost lenses are especially protein-dense to achieve a refractive index change in aquatic environments. Attractive forces would enable crystallins to tightly pack in the lens but risk increasing propensity for phase separation. We propose that teleost crystallins are structurally more flexible than mammalian paralogs to minimize the propensity of phase separation at the high concentrations necessary to function in aquatic environments, conferring the observed tolerance to very low temperatures as a side benefit.

RESEARCH CHALLENGE

Attractive forces that are responsible for maintaining proper density of the lens are subject to alterations by physical factors such as low temperature, resulting in the cold cataract phenomenon in

endotherms [1]. Reduction in attractive forces can increase cold resilience but would negatively impact the packing density of lens crystallins necessary for the refraction of light in ectothermic teleost fishes. Teleost lens crystallins, therefore, must have evolved adaptive mechanisms to pack at high concentrations, remain soluble, and avoid phase separation. Protein-protein interactions can be attenuated by modulation of flexibility at sites of interaction [2–4], and we propose that the abundant γ -crystallins in fish lenses evolved enhanced flexibility at interaction sites relative to mammalian paralogs.

γ -Crystallins have been identified as the mediator for phase separation [5]. Teleost fishes possess a unique γ class of crystallins, the γ M, which may confer the ability to maintain homogeneity at very high concentrations and extremely cold temperatures [6]. While mammals typically express between 6 to 7 γ -crystallin isoforms, teleost fishes express between ~20–40 unique isoforms depending on species, all except five belong to the γ M class. The large number of γ -crystallin isoforms in teleosts relative to mammals suggests inherent functional importance. Uniform flexibility across all γ -crystallin isoforms may negate the effects of attractive forces necessary for the tight packing of lens crystallins that maintains a high refractive index for teleost fish lenses. Therefore, we additionally propose that flexibility profiles across the lens crystallin landscape will be diverse.

We are currently utilizing the computational power of Blue Waters to run extensive molecular dynamics simulations to address our hypotheses regarding flexibility and extreme cold tolerance. With this resource, we are able to ascertain the potential contribution of flexibility to resist cold cataracts at cold temperatures by assessing the flexibility of a large suite of γ -crystallin isoforms among teleost fishes and mammals.

METHODS & CODES

We ran molecular dynamics (MD) simulations on eleven zebrafish and five mammalian isoforms at a cold temperature (0°C), and 3 zebrafish and 5 mammalian γ -crystallin isoforms at the normal body temperature (25°C and 37°C respectively). Three replicates of each γ -crystallin isoform were simulated for 50 nanoseconds in NAMD 2.12 using CHARMM27 force field parameters. Each of the five mammalian γ -crystallin isoforms was simulated using solved structures obtained from the Protein Data Bank. Simulation of zebrafish γ -crystallin isoforms used one known structure, the γ M7-crystallin, and the remaining 10 γ -crystallin isoforms simulated in this study were modeled onto the γ M7-

crystallin using iTasser. VMD 1.9.3 was used to quantify flexibility via root mean square fluctuations (RMSFs), which measure the average distance (Angstroms) of aligned backbone C α atoms per residue of a protein over the duration of the simulation. Average RMSF values were taken for each isoform, then formatted based on amino acid sequence alignment generated by MUSCLE 3.8.31 for comparison.

RESULTS & IMPACT

At 0°C, it is evident that zebrafish γ -crystallins are largely more flexible than the mammalian isoforms, most notably on surface loops (Fig. 1a). Fig. 1b shows simulations of mammalian γ -crystallins at 37°C and presents flexibility profiles similar in amplitude to zebrafish isoforms at 0°C. The zebrafish γ -crystallin isoforms tested at 25°C exhibit greater flexibility profiles compared to 0°C (Fig. 2a). In accordance with our hypothesis concerning variation of flexibility profiles among the γ M-crystallins, the 11 zebrafish γ M-crystallins are not identical in amplitude across all isoforms (Fig. 2b). This is unlikely to be due to functionally neutral changes along the γ M-crystallin evolutionary trajectory, but, rather, is likely to be due to functional diversity. Identical flexibility at regions with significant attractive forces may mitigate the effects the attractive forces have in maintaining high concentrations necessary for light refraction in teleost fish lenses. Diversity with regard to flexibility, spatial distribution of attractive forces, size, and shape may all be essential parameters in maintaining homogeneity in the lens at high concentrations and over a range of temperatures. Flexibility, in part, appears to explain the large standing question regarding the incredible ability of teleost fish lenses to maintain transparency at extremely cold temperatures.

WHY BLUE WATERS

Our work requires simulating three (3) trials of 49 proteins at two temperatures and over a long timecourse of 50 nanoseconds to detect meaningful molecular behavior. This work is at the core of a Ph.D. project in determining the extreme cold tolerance observed in teleost fish lenses. Only the petascale computational power and resources of Blue Waters could allow us to achieve this core portion of the project in a reasonable amount of time for downstream analyses to test our hypotheses. Without Blue Waters, we would not be able to finish this project in a reasonable time for a P.D. project.




Figure 1: a. Flexibility of five mammalian γ -crystallin isoforms (blue) and 11 zebrafish isoforms (orange) at 0°C. b. Flexibility of five mammalian γ -crystallin isoforms at 0°C (blue) and 37°C (red).




Figure 2: a. Flexibility of three zebrafish γ -crystallin isoforms at 25°C (orange) and 0°C (purple). b. Flexibility of 11 zebrafish γ -crystallin isoforms at 0°C.

PREDICTING PROTEIN STRUCTURES WITH PHYSICAL PETASCALE MOLECULAR SIMULATIONS

Allocation: NSF PRAC/2,900 Knh

PI: Ken A. Dill¹

Co-PI: Alberto Perez¹

Collaborators: Cong Liu¹, Emiliano Brini¹, James Robertson¹, Lane Votapka¹, Roy Nassar¹

¹Stony Brook University

EXECUTIVE SUMMARY

Blue Waters has enabled the first use of physics-based methods to make accurate atomistic predictions of protein structure based on sequence information alone in a double-blind international competition called CASP (Critical Assessment of protein Structure Prediction). This is a step forward for the computational biophysics community and an unprecedented result in the last 24 years of CASP, where physics has not had a role in *ab initio* modeling.

Further, we have used this methodology with great success to complete the pipeline that would lead to drug discovery: from sequence to structure, to interactions with small drugs, peptides, and other proteins. We are developing new protocols in these areas and have performed well in matching experimental results for the MDM2/p53 system, a cancer target. We continue to make headway in improving the accuracy and speed of our unique technology (MELD, or Modeling Employing Limited Data), which requires the significant GPU (graphics processing unit) resources available on Blue Waters.

RESEARCH CHALLENGE

A grand challenge in the last 60 years is to use computers to predict how proteins fold from a polymer sequence into a 3D structure. Proteins are the workhorses of the cells, made from 20 different amino acids whose arrangement in a linear chain gives rise to the 3D structure. We want to learn the physical principles that lead to the 3D structure so we can design drugs to bind to these proteins and stop them from misbehaving, or to design new protein sequences that carry out new functions (e.g., digest oil spills).

Computationally cheap knowledge databases are often used to copy and mimic structures based on sequence similarity. Computationally expensive physics-based methods are necessary to push the boundary of our predictive ability.




Figure 1: There are 465 nonthreadable proteins in the Protein Data Bank whose structure is known. When predicting the structure based on sequence using state-of-the-art homology methods (left) the predicted structure is wrong. Using MELD and physics, we get the right answer (right: experiment in red, our prediction in blue).

METHODS & CODES

Physics models are accurate but too slow to tackle problems such as folding. MELD [1,2] is our solution. In this approach, we accelerate molecular dynamics (MD) simulations while integrating ambiguous, noisy, and sparse data through a Bayesian inference approach. MELD allows us to find the best agreement between physics and a subset of the data. On one hand, this reduces the search space of the system, accelerating the convergence time of physics-based methods. On the other hand, MELD allows us to deal with imperfect sets of information, such as data derived from a generic knowledge of the system from experiments. Our MD engine is based on the OpenMM program [3] and AMBER [4] for force fields and setup.

RESULTS & IMPACT

With Blue Waters, we have taken MELD and physics to the limit. We have participated in CASP, the blind protein structure prediction competition. This worldwide event in which more than 200 groups participate has taken place every other year for the last 24 years. During the three months of competition, hundreds of targets are released and predictions have to be submitted in a timely manner. Before MELD, atomistic physics-based simulations had not been possible this quickly. Thanks to MELD and Blue Waters, not only were we able to make predictions, some of them were the best of the entire competition. This is an unprecedented result in the last 24 years of CASP and in using physics-based approaches for structure prediction.

We put MELD to work on problems where current technologies fail. MELD is particularly suited for predicting the structure of small proteins with little homology. These are at the intersection of cases where the faster database methods do not work and where the computational cost of MELD is still acceptable. We are working on microproteins (<70 amino acids) with no known experimental structure as well as a glass-ceiling set of proteins for which the experimental structures are known [5]. We are obtaining encouraging results in both cases, and we have active collaborations with experimentalists to prove the quality of our prediction (using circular dichroism and nuclear magnetic resonance spectroscopy). In the second case, we have so far had a 25% success rate (see Fig. 1).

Sometimes protein folds upon binding to other proteins. This means that a disordered protein in solution obtains a specific 3D shape when it interacts with a specific receptor. Common computational tools fail to predict the binding pose due to the rearrangement of the structure during the process. With MELD, we have been able to predict binding poses and relative binding free energies for the complex P53-MDM2, which is involved in the development of cancer. The protocols we designed for this project are unique and have significantly advanced the field. Without Blue Waters, this project would have taken us about 16 times longer to simulate.




Figure 2: Protein-protein conformations predicted with MELD (top left). Populations (lower left) are related to free energies; hence, we can pick up the right answer by clustering results. Our method so far succeeds in 15 of 20 cases (right).

Proteins bind to each other and to other molecules. With MELD, we are able to investigate the way complexes of two proteins are arranged in space (see Fig. 2) and how small ligands (i.e., drugs) bind to a protein. Ultimately, physics governs both processes, and MELD allows for a tightly focused search process. This is an exciting new line of research made possible by the computational resources of Blue Waters.

WHY BLUE WATERS

Blue Waters has been the perfect resource for us. It is the only cluster in the United States where we can get enough throughput GPU (graphics processing unit) usage, especially for the time-sensitive CASP competition. The staff is also helpful and quick to respond to solve issues and help us to maximize our Blue Waters usage for efficiency. We have a 100-GPU cluster in our lab (slow 2050s and 2070s models). Every MELD calculation requires at least 30 GPUs, so at most three calculations can run in the lab, whereas on Blue Waters we can sometimes run up to 30 calculations.

PUBLICATIONS AND DATA SETS

Perez, A., et al., Blind protein structure prediction using accelerated free-energy simulations. *Science Advances*, 2 (2016), pp. e1601274–e1601274.

Morrone, J. A., et al., Molecular Simulations Identify Binding Poses and Approximate Affinities of Stapled α -Helical Peptides to MDM2 and MDMX. *J. Chem. Theory Comput.*, 13 (2017), pp. 863–869.

Morrone, J., A. Perez, J. MacCallum, and K. Dill, Computed Binding of Peptides to Proteins with MELD-Accelerated Molecular Dynamics. *J. Chem. Theory Comput.*, 13 (2017), pp. 870–876.

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

Allocation: Exploratory/50 Knh

PI: Ahmed Elbanna¹

Co-PI: Helen Nguyen¹

Collaborator: Jaehyuk Kwack²

¹University of Illinois at Urbana-Champaign

²National Center of Super Computing Applications, UIUC

EXECUTIVE SUMMARY

Biofilms are aggregates of cells and extracellular polymeric substances. They are found ubiquitously in both natural and engineered systems, such as on the surfaces of pipes in drinking water distribution systems (DWDS). Biofilms in DWDS have been reported to be capable of attracting and harboring pathogens. In addition, the biofilm matrix may prevent disinfectants from reaching the cells located deep inside the biofilm. As a result, pathogenic microorganisms have been found in DWDS biofilms and have been linked to outbreaks and severe health problems. Understanding the mechanisms of pathogen attachment to biofilms developed in DWDS is of crucial interest to ensure the quality of drinking water and is a critical public health issue. This project contributes to establishing a comprehensive multiscale framework for studying the dynamics of growth and detachment of biofilms in DWDS under different operating conditions by integrating fluid mechanics, solid mechanics, and chemistry to predict pathogenic outbreaks and related public health hazards.

RESEARCH CHALLENGE

The quality of drinking water is a critical public health issue. The formation and detachment of biofilms in DWDS have been linked to pathogenic outbreaks. Yet, an understanding of the response of biofilms to different types of disinfectants and their susceptibility to fracture under different flow rates and growth environments is still lacking. Several challenges exist, including the absence of a comprehensive constitutive model for the nonlinear and inelastic response of the soft biofilm, the extreme computational demand of the 3D coupled fluid-film simulation, and the numerical challenges associated with meshing the complex biofilm geometry, integration of the tightly coupled partial differential equations governing the chemo-hydro-mechanical response as well as arbitrary growing biofilm fractures in 3D. Progress along these different fronts contributes not only to getting us closer to being able to make quantitative predictions for biofilm response in different operating conditions but also pushes the frontiers of computational science in material modeling and multiphysics simulation.

METHODS & CODES

Over the past allocation period, we have been able to make fundamental progress in addressing some of these challenges. We have developed a continuum nonequilibrium statistical thermodynamics framework based on the effective temperature approach to model the nonlinear elasto-plastic response of soft amorphous materials. We implemented this material model in a nonlinear finite deformation framework within the MOOSE platform, a finite element framework from Idaho National Lab for multiphysics simulations. MOOSE provides an advanced modular computational infrastructure including the libMesh finite element library and the PETSc solver library. The program requires input of the weak form for partial differential equations to be solved as well as the material model, but it provides a variety of shape functions, stabilization options, and locking control algorithms.

For the bio-film fluid interaction, we used a finite element code for Computational Fluid Dynamics and Fluid-Solid Interaction developed by Dr. JaeHyuk Kwack, who has worked in the Scientific and Engineering Applications Support group at NCSA.

The application code employs a variational multiscale approach for turbulence and nonlinear material modeling; as a result, it provides an optimal computational performance for multiphysics problems that will be investigated through this project. For parallel computation, the code is implemented with the Message-Passing Interface (MPI). As a preprocessing step, numerical meshes are decomposed by the METIS/ParMETIS library. Each MPI process accesses its own decomposed numerical domain (i.e., a series of elements and nodes) and locally computes its residual vectors and tangent matrices. After that, global communications for matrix-vector calculations, assembly process, and residual norm calculations are implemented via the PETSc library. Krylov subspace methods from PETSc and sparse direct matrix solvers (e.g., SuperLU_dist and MUMPS) are employed to solve linear equations with the globally assembled matrices and vectors. The output files are written in VTK file formats and they are post-processed with ParaView.




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

RESULTS & IMPACT

Our results thus far have been pushing the limits of the state of the art in modeling biofilm mechanics. To the best of our knowledge, prior work in the field has been limited to two-dimensional models. We have been able to run full 3D fluid-structure models of biofilm with complex surface geometry. As a result, we can realize complex 3D turbulent structures near surface irregularities as well as complex stress patterns that are not apparent from the 2D simulations. Furthermore, the 3D implementation of the amorphous material model enables us to realize—for the first time—complex strain localization patterns (which will potentially map into complex 3D fracture surfaces) that have not been reported before.

WHY BLUE WATERS

The progress accomplished so far would not have been possible without Blue Waters. Several of the problems we have tested required 10,000 or more core-hours of runtime and generated tens to hundreds of gigabytes of data per simulation. These resources are not accessible on conventional cluster scales. Furthermore, cloud computing is not suitable for the nature of our problems since any parallelizable discretization of the coupled system of equations of interest would require domain decomposition and efficient communication among processors, making cloud computing out of question.

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

Allocation: GLCPC/310 Knh

PI: Peter Freddolino¹

Co-PI: Morteza Khabiri¹

Collaborator: Arttu Jolma²

¹University of Michigan Medical School

²University of Toronto

EXECUTIVE SUMMARY

Transcription factors (TFs) and other DNA-binding proteins shape the behavior of all cells, coordinating appropriate gene expression patterns in response to internal or external cues. For any particular transcription factor, maps of the binding affinity for different DNA sequences must be obtained through laborious and expensive experiments. We used Blue Waters as our computing resource to computationally map the TFs' binding free-energy landscapes for several well-studied transcription factors with known crystal structures. Comparing our results with

experimental data set on the same systems, we observe generally poor correlations among the computational predictions and experimental data. We used a robust computational protocol for reliable *in silico* determination of TF affinity landscapes; however, there is still some element of uncertainty with nonequilibrium molecular dynamics simulations likely to play a key role in defining the binding free-energy landscape. The remaining bottleneck in high-accuracy prediction of protein-DNA binding free-energy landscapes using these methods thus remains an area of active investigation.




Figure 1: Binding free energy landscape for ELK1. Computational values (black) are compared with human (green) and mouse (red) values from experimental data. The small box in the landscape plot indicates example of work distribution achieved from mutant→wildtype (blue) and wild type→mutant (green) transition. Data from [5].

RESEARCH CHALLENGE

Transcriptional regulation is driven in large part by the action of TFs and other DNA-binding proteins that either recruit or inhibit the recruitment of RNA polymerase. The activity of TFs that interact directly with DNA depends on their potent ability to recognize and bind to specific DNA sequences. In reality, most TFs bind to a range of related sequences; however, the one that provides enough information or has the lowest energy will define a reliable TF binding site. Thus, understanding the binding affinity landscape of each transcription factor to all possible sequences of DNA will help us to predict the behavior of regulatory networks. There have been a number of efforts to predict these binding sites quantitatively based on both experimental methods (e.g., protein-binding microarrays [1] or HT-SELEX experiments [2]). However, most of these experiments carry high costs in terms of time and money. Recent computational efforts by Seeliger and co-workers [4] predicted an accurate DNA-binding affinity landscape based on atomistic molecular dynamics simulations and reproduced experimental results to near-chemical accuracy. This proved to be much more accurate than previous computational methods. By using the massive computing resources provided by Blue Waters, we attempted to apply a similar method for free-energy calculations to that used by Seeliger and coworkers to map the complete binding free-energy landscapes of four TFs with well-characterized binding affinity landscapes.

METHODS & CODES

Building on previous results that showed accurate calculation of protein-DNA binding free energies for a small number of cases [4], we applied the Crooks-Gaussian intersection (CGI) method [6] to calculate the free-energy changes for base pair substitutions in the binding site of the transcription factors of interest. This method requires calculations of very long equilibrium simulations of the protein-DNA complex and the DNA alone for each of two sequences to be compared, followed by many short simulations morphing the system between the two sequences. We performed the free-energy calculations for all possible single nucleotide perturbations of the consensus binding site for the transcription factor of interest. Our results illustrate both very poor overall correlations with the experimental results and, frequently, unphysically large magnitudes of binding free-energy changes [5]. Excluding many sources of possible errors in our simulation setup, we realized that neither extensive control simulations using simplified systems or other free-energy calculation methods, nor careful characterization of the structural features involved in the protein-DNA interface of the simulated complexes, provided an explanation for the poor correlation between calculated and experimental binding free-energy landscapes. We are currently working to resolve this difficulty so that we can realize the promise of computational predictions of DNA-binding affinity landscapes.

WHY BLUE WATERS

The computational work described here requires the capability to efficiently bring huge numbers of nodes together to run dozens of simulations of independent trajectories using GPU-accelerated molecular dynamics software, and then, for each such trajectory, to perform more than 100 short follow-up simulations using CPU-only code for the free-energy calculation. The hybrid architecture of Blue Waters has been absolutely ideal for these applications, providing us with the most efficient possible environment for each portion of our workflow, and allowing us to make progress on huge numbers of mutational calculations simultaneously.

PUBLICATIONS & DATA SETS

Khabiri, M. and P.L. Freddolino, Deficiencies in Molecular Dynamics Simulation-Based Prediction of Protein-DNA Binding Free Energy Landscapes. *J. Phys. Chem. B*, 121:20 (2017), pp. 5151–5161.

NON-BORN-OPPENHEIMER EFFECTS BETWEEN ELECTRONS AND PROTONS

Allocation: Blue Waters Professor/240 Kwh
PI: Sharon Hammes-Schiffer¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

The quantum mechanical behavior of nuclei plays an important role in a wide range of chemical and biological processes. The inclusion of nuclear quantum effects and non-Born–Oppenheimer effects between nuclei and electrons in computer simulations is challenging. Our group has developed the nuclear-electronic orbital (NEO) method for treating electrons and select nuclei quantum mechanically on the same level using an orbital-based formalism. We have used Blue Waters to perform NEO calculations on systems in which all electrons and one proton or positron are treated quantum mechanically. In addition, we have studied the active site conformations of a blue light using flavin adenine dinucleotide (BLUF) photoreceptor protein and their relation to photocycle activation. This work serves as groundwork for subsequent nonadiabatic dynamics studies of photoinduced proton-coupled electron transfer (PCET) reactions that have been proposed to be involved in biological signaling by BLUF proteins [1].

RESEARCH CHALLENGE

The inclusion of nuclear quantum effects such as zero-point energy and tunneling in electronic structure calculations is important in a variety of chemical systems, particularly those involving hydrogen transfer or hydrogen-bonding interactions. Moreover, nonadiabatic effects—also called non-Born–Oppenheimer effects—between electrons and certain nuclei are significant for many of these systems. In this case, the electrons cannot be assumed to respond instantaneously to the nuclear motions, and the concept of the nuclei moving on a single electronic potential energy surface is no longer valid. This type of nonadiabaticity has been shown to play a critical role in PCET reactions, which are essential for a wide range of chemical and biological processes, including phototaxis, 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 the design of novel materials for optogenetics to the design of more effective catalysts for solar energy devices.




Figure 1: (Left) A schematic diagram of the photocycle in BLUF photoreceptor proteins. The dark state to light state conversion is believed to involve photoinduced PCET and shifts the equilibrium between the “off” and “on” states of the effector domain. (Right) Our studies indicate that the proton relay conformation of the active site Tyr and Gln residues is important for electronic charge transfer from Tyr to the flavin.

METHODS & CODES

Typically, in the NEO approach, all electrons and one proton or positron are treated quantum mechanically, and a mixed nuclear-electronic time-independent Schrödinger equation is solved [2]. 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 [3]. In this approach, denoted NEO-RXCHF, only select electronic orbitals are explicitly correlated to the nuclear or positronic orbital, and certain exchange terms are approximated. We recently incorporated an improved integral package into the NEO code. The new and fully optimized integral code enables the use of a direct algorithm for the NEO-RXCHF method and provides a multiple order of magnitude speed-up compared to the old integral code while significantly reducing the memory overhead.

The study of photoinduced PCET in the BLUF photoreceptor proteins requires resolution of the debate over the conformations of key active site residues arising from inconsistent X-ray crystallographic and solution nuclear magnetic resonance structures. Unlike most computational studies reported in the literature, we carried out comprehensive free energy calculations to elucidate the dominant conformations of key residues on different electronic states and studied the effects of active site conformations on the photocycle activation mechanism. The free energy calculations were carried out using the adaptively biased path optimization (ABPO) [4] and umbrella sampling [5] techniques implemented in the CHARMM [6] software package, while excited electronic states were calculated using the time-dependent density functional theory [7] method implemented in Gaussian 09 [8].

RESULTS & IMPACT

The NEO-RXCHF method was used to compute the binding energies of a positron to neutral atoms. The NEO-RXCHF results agreed well with highly accurate stochastic variational method (SVM) calculations. The inclusion of explicit electron-positron correlation was found to be essential for even qualitative accuracy, as NEO methods that neglect explicit electron-positron correlation did not predict the positron to bind to the atomic systems. The NEO-RXCHF method has better computational scaling characteristics than SVM, and SVM calculations have never been performed on a system with more than five quantum particles. Because of this limitation, the SVM method cannot be used for larger molecular systems. No *ab initio* or first principles method has ever predicted that a positron will bind to an alkane, which has been observed experimentally. The accuracy of the NEO-RXCHF calculations combined with this method’s computational tractable scaling characteristics demonstrate that NEO-RXCHF is a promising method for computing the binding energy of positrons to molecules such as alkanes.

Our free energy simulations of the AppA BLUF domain have elucidated the active site conformations before and

following photoexcitation [9]. For the dominant Trp_{in}/Met_{out} conformation, the hydrogen-bonding pattern conducive to the proton relay from Tyr21 to the flavin chromophore is not thermodynamically favorable on the ground electronic state but becomes more favorable, corresponding to approximately half of the configurations sampled, on the locally excited state of the flavin. The calculated energy gaps between the locally excited state and the charge-transfer state associated with electron transfer from Tyr21 to the flavin suggest that electron transfer from Tyr21 to the flavin is more facile for configurations conducive to proton transfer. When the active site conformation is not conducive to PCET from Tyr21, Trp104 can compete directly with Tyr21 for electron transfer to the flavin through a nonproductive pathway, impeding the signaling efficiency. These insights have contributed to the resolution of a long-standing debate in the field of BLUF protein research and pave the way for nonadiabatic dynamics studies that will be crucial in understanding the signaling mechanism of the AppA BLUF domain.

WHY BLUE WATERS

The NEO-RXCHF method requires the calculation of trillions of integrals. Our in-house NEO code has been parallelized using the MPI (Message Passing Interface) protocol and requires a large number of processors. Additionally, with the implementation of a direct algorithm for the NEO-RXCHF method, the fast calculation of the integrals is even more important, as the integrals must be computed many times during the calculation instead of a single time at the start. The speed and scalability of Blue Waters is crucial in enabling large NEO-RXCHF calculations.

To investigate the conformations of the active site tryptophan and methionine residues in the AppA BLUF domain, we used the ABPO method implemented in CHARMM. The calculations employed 21 images of the system and were carried out on a total of 1344 processors. Blue Waters provided a large number of simultaneously available nodes and efficient inter-node communication important for these calculations. Assistance from the project staff was crucial in successful building and testing of the locally modified CHARMM code on Blue Waters.

PUBLICATIONS AND DATA SETS

Borsen, K. R., M. V. Pak, and S. Hammes-Schiffer, Calculation of Positron Binding Energies and Electron-Positron Annihilation Rates for Atomic Systems with the Reduced Explicitly Correlated Hartree-Fock Method in the Nuclear-Electronic Orbital Framework. *J. Phys. Chem. A*, 122:2 (2017), pp. 515–522.

Borsen, K. R., Y. Yang, M. V. Pak, and S. Hammes-Schiffer, Is the Accuracy of Density Functional Theory for Atomization Energies and Densities in Bonding Regions Correlated? *J. Phys. Chem. Lett.* 8:9 (2017), pp. 2076–2081.

Goyal, P., and S. Hammes-Schiffer, Role of active site conformational changes in photocycle activation of the AppA BLUF photoreceptor. *Proc. Natl. Acad. Sci. USA*, 114:7 (2017), pp. 1480–1485.

BRUECKNER–GOLDSTONE QUANTUM MONTE CARLO

Allocation: Blue Waters Professor/200 Knh

PI: So Hirata¹

Collaborators: Cole M. Johnson¹, Alexander E. Doran¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

We introduced and fully developed novel scalable algorithms and software for predictively accurate (*ab initio*) electronic-structure calculations for large molecules and solids, which are not easily subjected to fast calculations by fragmentation. We transform the usual, non-scalable sum-of-products expressions of many-body perturbation and Green's function theories in the complete-basis-set limit into a few high-dimensional integrals, which are then evaluated by a highly scalable Metropolis Monte Carlo algorithm. They can compute energy differences (including quasiparticle energy bands) directly without a sign problem at an operation cost whose size dependence is one or two ranks lower than their deterministic counterparts. They execute efficiently on many CPUs or many GPUs, easily achieving an unprecedented speedup (for an *ab initio* electron-correlation calculation) by a factor of 31,000 (on 256 GPUs) relative to a serial calculation.

RESEARCH CHALLENGE

Existing algorithms of predictive computational chemistry are not scalable with respect to either system size or computer size. For example, the memory and arithmetic operation costs of the simplest *ab initio* electron-correlated theory, i.e., the second-order Møller–Plesset perturbation (MP2) theory, increase as the fourth and fifth power of system size, respectively, and its computational kernel is poorly parallelizable. The situation is even worse for

higher-order MP and other *ab initio* electron-correlation methods such as coupled-cluster theory. This is compounded with the extremely slow convergence of their results with the size of an expansion basis set, further driving up the cost to achieve the complete-basis-set (CBS) limit. This project seeks to invent and fully develop fundamentally scalable algorithms for predictive computational chemistry (i.e., without sacrificing accuracy by introducing arbitrary approximations with uncontrollable errors) with respect to both system and computer sizes by combining quantum-Monte-Carlo-like stochastic algorithms and *ab initio* electron-correlated theory. We will then deploy such algorithms on Blue Waters in chemistry/solid-state-physics applications with unprecedented accuracy and problem sizes.

METHODS & CODES

We mathematically transformed the usual sum-of-products expressions of MP2, second-order Green's function (GF2) theory, and their CBS corrections by explicitly correlated (F12) ansätze into single high-dimensional integrals by a Laplace transform. These integrals are then evaluated by a Metropolis Monte Carlo method with judiciously chosen weight functions. The resulting stochastic (Brueckner–Goldstone quantum Monte Carlo) methods [1—Monte Carlo MP2 (MC-MP2) [2], Monte Carlo GF2 (MC-GF2) [3], Monte Carlo explicitly correlated MP2 (MC-MP2-F12) [4,5], and Monte Carlo explicitly correlated GF2 (MC-GF2-F12)




Figure 2: The universal size (the radius of 0.8 Ångstrom) and concave shape of a correlation hole uncovered by Monte Carlo MP2 calculations with 17 different correlation factors.

[6]—can compute energy differences (correlation energies and electron detachment/attachment energies) directly without a sign problem in a scalable manner with respect to both computer size (on thousands of CPUs or hundreds of GPUs) and system size (the operation cost is linear scaling per MC step and cubic to quartic scaling to achieve given relative accuracy and the memory cost is negligible) [7]. They can also calculate quasiparticle energy bands of a solid for the entire Brillouin zone as nearly continuous curves of a wave vector [8] and have been extended to third-order MP (MP3) [9] using an expedient interpretation of Brueckner–Goldstone diagrams as well as a convergence-acceleration scheme (redundant-walker algorithm) [10].

RESULTS & IMPACT

The MC-MP2-F12 method enabled an exact (CBS-limit) MP2 energy calculation of tetrahydrocannabinol (472 basis functions) without a local-correlation scheme (Fig. 1). Exploiting the extraordinary flexibility of this algorithm in using virtually any explicitly correlated factor, we numerically characterized the performance of 17 such factors. We observed that highly performing factors share the same short-range behavior within the radius of 1.5 Bohr, while differing greatly in the long-range behavior. This result reveals fundamental electron-correlation physics that a correlation hole of a pair of electrons has a universal size (1.5 Bohr) and concave shape (dictated by Kato's cusp condition) regardless of its molecular environment or energy (Fig. 2). We have completed the development of the MC-GF2-F12 method, which can compute electron detachment/attachment energies directly in the CBS limit. With this, we computed exact GF2 electron affinities with a statistical uncertainty of 0.03 eV for C_{60} and C_{70} , which play important roles in heterojunction solar cells as an electron acceptor but resist a local-correlation scheme for fast calculations.

These calculations were based on the redundant-walker algorithm, which propagates more walkers than minimally necessary and permutes them in all possible ways when being substituted into the integrand, thereby multiplying the sampling efficiency. We introduced [7] a two-level parallelism in which dense matrix multiplications for many walkers are fine-grained on a GPU and a Monte Carlo integration itself is coarse-grained across multiple CPU-GPUs. In this way, not only did we observe a speedup by a factor of 31,000 on 256 GPUs relative to a serial execution, but we also found that the saturation point of the acceleration is significantly delayed to a much greater number of walkers. This is a rare instance in which the parallel architecture (GPU) and algorithm (the redundant-walker algorithm) mutually enhance each other.

WHY BLUE WATERS

The stability and ease of use (OS, compilers, libraries, and NCSA expertise) as well as the balanced deployment of CPUs and GPUs are all essential for rapid coding/profiling of new scalable algorithms from scratch and their capacity testing.

PUBLICATIONS AND DATA SETS

Johnson, C. M., S. Hirata, and S. Ten-no, Explicit correlation factors. *Chem. Phys. Lett.*, published online (2017): <https://doi.org/10.1016/j.cplett.2017.02.072>.

Johnson, C. M., A. E. Doran, J. Zhang, E. F. Valeev, and S. Hirata, Monte Carlo explicitly correlated second-order many-body perturbation theory. *J. Chem. Phys.*, 145 (2016), p. 154115.

Doran, A. E. and S. Hirata, Monte Carlo MP2 on many graphical processing units. *J. Chem. Theory Comput.*, 12 (2016), pp. 4821–4832.




Figure 1: Monte Carlo explicitly correlated MP2 energies of tetrahydrocannabinol (472 basis functions).

HOW MEMBRANE ORGANIZATION CONTROLS INFLUENZA INFECTION

Allocation: GLCPC/640 Kwh

PI: Peter Kasson¹

Collaborator: Steven Boxer²

¹University of Virginia

²Stanford University

EXECUTIVE SUMMARY

Viruses such as influenza infect cells via a process of membrane fusion between the viral envelope and cellular membranes. We have shown experimentally that changing the membrane composition of either the influenza virus or its target cellular membranes can affect the rate and efficiency of infection, but a detailed molecular

explanation for these effects is not experimentally accessible. We are, therefore, using molecular dynamics simulations to study how both the viral and target membranes reorganize in response to changing composition. We analyze these via machine-learning approaches and compare to experimental data that our collaborators and we measure.




Figure 1: This image shows a model of an influenza membrane patch. Components of the membrane are varied to match experimental perturbations and compared to experimental measurements.

RESEARCH CHALLENGE

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

METHODS & CODES

We incorporate structural data from X-ray crystallography and cryo-electron microscopy along with mass spectrometry to construct initial models of influenza viral and target membranes. These membrane patches are then simulated using the molecular simulation software that we helped develop, Gromacs, and are measured under perturbations that match what we do experimentally. We use in-house machine-learning methods to analyze the simulations and connect them back to experiments. This then yields a model of the membrane response and its infectious consequences at a high level of molecular detail yet also with experimental validation.

RESULTS & IMPACT

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

WHY BLUE WATERS

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

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

Allocation: Illinois/169 Knh
PI: Zaida Luthy-Schulten¹
Collaborators: David Bianchi¹, Joseph Peterson¹, Tyler Earnest², Michael Hallock¹

¹University of Illinois Urbana-Champaign
²National Center for Supercomputing Applications

EXECUTIVE SUMMARY

Stochasticity in transcription is an important source of noise that can have a profound effect upon the fate of a living cell. In recent years, we have seen the advent of a community of researchers interested in performing stochastic simulations of large biological systems (e.g., millions of particles). The Lattice Microbes (LM) software suite and its pyLM problem-solving environment provide a convenient way to set up simulations of complex biological systems. However, simulations of large systems performed using the exact Stochastic Simulation Algorithm (SSA) to solve the Chemical Master Equation (CME) are computationally expensive. To alleviate this issue, we have implemented a hybrid CME-ODE (ordinary differential equations) method for LM similar to previous mixed methodologies. Our robust hybrid implementation gives protein distributions and species dynamics similar to the pure SSA and can run approximately 50–100X faster in a simulation of the well-studied galactose switch in *S. cerevisiae* (yeast) containing millimolar particle concentrations.

RESEARCH CHALLENGE

Many processes within living cells, especially gene expression, are characterized by low particle numbers and a high degree of randomness. CME and its spatially resolved analog the Reaction-Diffusion Master Equation (RDME) are descriptions of cellular processes where the system is considered to follow a Markov jump process on the state space of particle numbers in time, capturing the discreteness of the particles and the random nature of individual chemical reactions. Gillespie's widely used SSA method [1] provides an effective technique for obtaining unbiased realizations of these Markov processes.

However, this algorithm is limited because reaction events are accounted for explicitly by the SSA, making simulations of highly reactive systems, where the time between reactions is

small, computationally expensive. Highly reactive systems are characterized by large reaction propensities that can arise in the case of high copy numbers, such as metabolites in millimolar concentrations, and/or large rate constants (fast reactions). A challenging and typical scenario is when species participating in slow reactions interact with species involved in fast reactions, making the dynamics of the slow reactions dependent on the fast reactions. To alleviate the issues faced by the SSA for high particle number systems, many researchers have developed hybrid multiscale stochastic approaches [2,3,4] in which the highly reactive parts of the system are described by ODE and the slow reactive parts are described stochastically. Our hybrid method along with an easy-to-use interface through LM [5] and pyLM [6] provides an effective way to study stochastic behavior in highly reactive systems.




Figure 1: (top) The *S. cerevisiae* galactose switch system with ODE/CME separation. The reactions depicted (boxed area) are simulated continuously; those outside this region are simulated stochastically. (bottom) Distribution at steady state of the G3i species. As the communication timestep is decreased from 5 minutes to 1 second the distributions given by the hybrid algorithm more closely match pure SSA results.

METHODS & CODES

The galactose switch system, with its four feedback loops and millimolar galactose concentration, is separated into a regime of species whose reactions will be simulated stochastically and another whose reactions will be simulated deterministically (Fig. 1). At the beginning of each timestep, the differential equation solver (DES) is updated with the species counts obtained from the stochastic regime (transcription, translation) simulated via the SSA. The DES then takes adaptive timesteps to evolve the high particle number species through time in the deterministic regime. At the conclusion of a timestep, the stochastic rates of reactions involving low particle number species interacting with high particle number species are updated with the species counts found by the ODE solver. At this time, the hybrid algorithm also communicates updated species counts generated from reactions in the CME regime to the ODE regime. The optimal communication times between the stochastic and deterministic descriptions, as well as the timesteps for each method, need to be assessed (Fig. 1) to verify that the hybrid description accurately describes the stochastic dynamics, which often have great impact on the cell's behavior.

RESULTS & IMPACT

Such a CME-ODE partitioning works well for both bacterial and eukaryotic systems where stochastic effects are important. Partitioning typically improves the speed of the numerical simulations by a factor of 50–100, making it an indispensable tool for complex cell simulations with a large number of species types, cellular components, and high concentrations of metabolites (sugars, etc.) inside and outside the cell. Simulations enabled by this type of hybrid algorithm will allow researchers to study larger and more detailed systems, capturing the effects of reactions involving high particle count species such as metabolites, which have a crucial role in systems such as the genetic switch studied here. We have already used this hybrid approach to perform a spatially resolved RDME-ODE study (Fig. 2) of the galactose switch system, experiencing similar speedup to what is seen in the CME implementation.

WHY BLUE WATERS

Blue Waters was essential to generate over 1,000 replicate hybrid simulations over the simulation time of 750 minutes and a range of concentrations. Only then did we have sufficient data to make the results statistically reliable and to determine the optimal communication time. In the worst-case scenario, the full CME simulations take nearly two days of wall-clock time, while the hybrid CME-ODE implementation with a communication time of one second requires less than 30 minutes. The response of the switch guided the setup for much more computationally costly RDME-ODE simulations on Blue Waters, which account for the spatial heterogeneous environment (nucleus, cytoplasm, membrane, etc.) of the yeast cell.

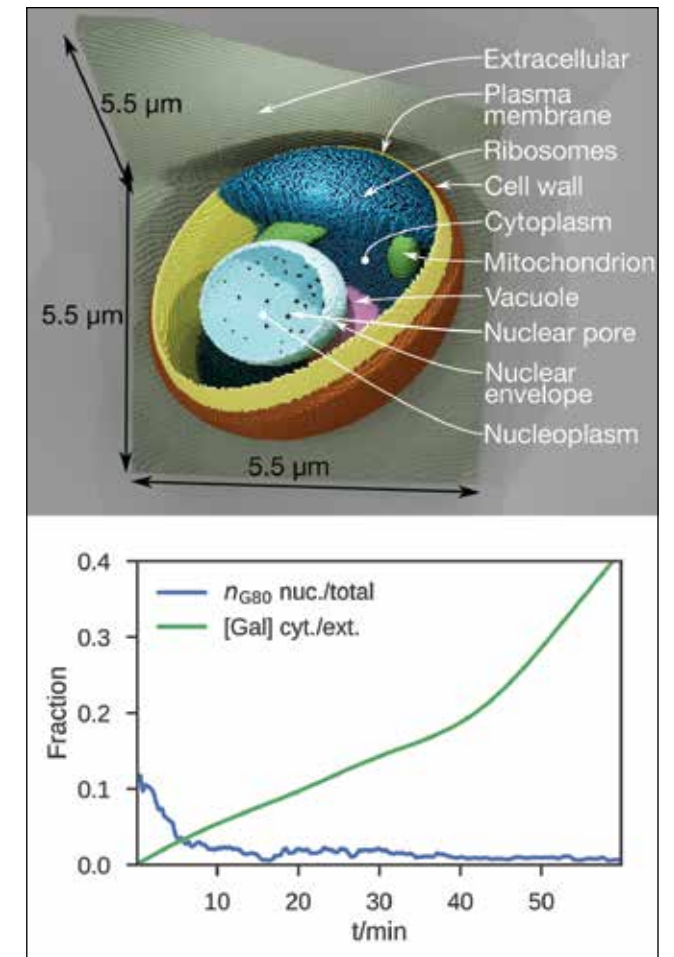


Figure 2: (top) A visualization of the cell compartmentalization that is imposed in a hybrid RDME-ODE simulation, with structural insight taken from tomograms. (bottom) The decrease in the amount of transcription factor present in the nucleus as the galactose concentration increases, matching expectations from experimental results.

QUANTUM-CLASSICAL PATH INTEGRAL SIMULATION OF CHARGE TRANSFER REACTIONS

Allocation: Illinois/100 Knh 175,000

PI: Nancy Makri¹

Co-PI: Peter L. Walters¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Quantum mechanical calculations on condensed phase or biological systems are prohibitively expensive. Mixed quantum-classical approximations are highly efficient and attractive. However, conventional quantum-classical methods introduce major assumptions in the treatment of the interaction between quantum and classical degrees of freedom. We have developed a rigorous quantum-classical path integral (QCPI) methodology that is free of assumptions and that treats all interactions in full atomistic detail. Our QCPI simulation on Blue Waters of the ultrafast ferrocene–ferrocenium electron transfer reaction in liquid hexane, a standard in electrochemistry, yielded results of unprecedented accuracy and enabled the first quantitative demonstration of Gaussian behavior for a complex molecular solvent. Additional theoretical and algorithmic advances enabled simulation of reactions that span very different time scales. In particular, simulation of this process in a series of solvents with increasing polarity exemplified the interplay between electronic and molecular time scales and energetics and its effects on the kinetics of charge transfer reactions.

RESEARCH CHALLENGE

Tunneling, quantum dispersion, and phase interference play a subtle but very important role in many chemical and biological processes. Quantum mechanical simulations of dynamical processes in the condensed phase continue to be extremely challenging because quantum mechanics is a nonlocal theory, implying an exponentially growing cost with the number of interacting particles.

For many processes of interest, quantum mechanical effects are essential only in the treatment of a small number of particles (e.g., a proton, or a small number of electronic states associated with a charge transfer reaction). The remaining particles (solvent molecules or biological medium) can be adequately described via classical dynamics. Unfortunately, the traditional Schrödinger formulation of quantum mechanics (which is based on delocalized wave functions) is incompatible with Newtonian trajectories (which are local in space). In the Schrödinger formulation, the interaction between the quantum and classical partitions cannot be treated without resorting to severe approximations.

METHODS & CODES

The focus of our work has been the development of a rigorous quantum-classical formulation based on Feynman path integral formulation of quantum mechanics. The local nature of the Feynman paths leads naturally to quantum-classical treatments that are free of *ad hoc* assumptions. The quantum-classical path integral (QCPI) methodology that we developed is a rigorous quantum-classical formulation. However, the QCPI expression appears impractical, as it contains an astronomical number of terms. Several advances in the understanding of interference and decoherence have recently made the QCPI methodology practical for the simulation of condensed phase reactive processes.

QCPI treats the small system of interest by full quantum mechanics, while the effects of the environment are captured via standard molecular dynamics (MD) procedures. Two widely used MD packages, NAMD and LAMMPS, are used to yield trajectories subject to forces obtained using the coordinates of the charged particle, which are specified by the given path. The dynamics captured along a classical trajectory augments the Hamiltonian of the quantum system through a time-dependent term, which leads to level fluctuations and eventually to decoherence. The exponential proliferation of trajectories with propagation time is avoided through a tensor decomposition that exploits the memory-quenching effects of condensed phase environments. The QCPI algorithm is characterized by classical molecular dynamics scaling and is fully parallelizable.

RESULTS & IMPACT

The QCPI methodology has enabled the simulation of charge transfer reactions in solution with unprecedented accuracy. The dynamics of the prototypical ferrocene–ferrocenium electron transfer pair was studied in a series of organic solvents of increasing polarity. Our first simulation of this electron transfer reaction in liquid hexane employed 1,320 atoms interacting via CHARMM force fields. More recent work has studied the same reaction in a series of solvents of increasing polarity.

The all-atom QCPI simulations shed light on the complex interplay among molecular/solvent time scales, electronic couplings, and reorganization energy, particularly on the way that these effects determine the rate of the reaction, the nature (exponential or nonexponential) of the dynamics, and the validity of Gaussian response. They also demonstrate how the interference among quantum mechanical phases leads to decoherence, and




Figure 1: Solvent trajectories during the first two time steps of the ferrocene–ferrocenium electron transfer simulation. Blue and red correspond to the reduced and oxidized iron atoms. The solvent molecules surrounding the ferrocene–ferrocenium complex are seen to respond to the change of electronic configuration.

reveal substantial quantum delocalization of the otherwise classical solvent as a result of its interaction with the quantum electron transfer pair.

WHY BLUE WATERS

The QCPI formulation is well suited to a decomposition based on multi-level parallelism, and Blue Waters provides the ideal platform for its implementation. Specifically, the set of system paths is distributed across nodes; one processor within each node is assigned to the quantum mechanical calculations, while the other performs supporting trajectory computations. Moreover, because the trajectories are independent and generally relatively short, it is possible to assign a single trajectory to each core within a given processor while maintaining computational efficiency. This multi-level approach has the benefit of minimizing communication time while maximizing concurrent processing, since related classical and quantum-mechanical calculations are performed within the same node, where inter-processor communication should be much faster than if the information were more widely distributed.

PUBLICATIONS AND DATA SETS

Makri, N., Blip-summed quantum-classical path integral with cumulative quantum memory. *Reaction Rate Theory: Faraday Discussion*, 195 (2016), DOI: 10.1039/C6FD00142D.

Makri, N., Quantum-Classical Path Integral. *49th Midwest Theoretical Chemistry Conference* (Michigan State University, East Lansing, Mich., June 1–3, 2017).

Bose, A., and N. Makri, Quantum Mechanical Rate Calculations in Condensed Phase Reactions. *49th Midwest Theoretical Chemistry Conference* (Michigan State University, East Lansing, Mich., June 1–3, 2017).

Makri, N., Rigorous Quantum-Classical Path Integral Formulation of Real-Time Dynamics. *Workshop on Understanding Quantum Phenomena with Path Integrals: From Chemical Systems to Quantum Fluids and Solids* (ICTP, Trieste, Italy, July 3–7, 2017).

Makri, N., Interference, Decoherence, and Quantum-Classical Path Integral Simulations with Thousands of Atoms. *American Conference on Theoretical Chemistry* (Boston University, Boston, Mass., July 17–20, 2017).

Walters, P.L., and N. Makri., Quantum-classical path integral calculations of ferrocene-ferrocenium charge transfer in solution. *American Conference on Theoretical Chemistry* (Boston University, Boston, Mass., July 17–20, 2017).

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

Allocation: Illinois/200 Knh

PI: Arif Masud¹

Collaborators: Soonpil Kang¹, Elizabeth R. Livingston¹, Daniel R. Sheehan¹, Naveed Adoni², Tor Jensen^{1,2}

¹University of Illinois at Urbana-Champaign

²Carle Foundation Hospital

MP

EXECUTIVE SUMMARY

Modern computational technology provides a platform to take Magnetic Resonance Imaging (MRI) and Computed Tomography (CT) images and construct patient-specific geometric models for high-performance computing (HPC) as well as for 3D printing. Computed data obtained from numerical simulations can be converted into graphical images and used for clinical diagnostics as well as for vascular surgical planning. Geometric models can be used for developing STereoLithography (STL) files for 3D printing and generating patient-specific prototypes. In this project we focused on HPC aspects and extended and applied our methods [1,2] for blood flow modeling to patient-specific geometries.

RESEARCH CHALLENGE

The basic scientific and medical question that this research is aimed to address is to determine if computational models can provide insights into identifying regions of significant atherosclerotic plaque, luminal narrowing, and loss of wall elasticity. An important aid in the vascular access decision-making process in Transcatheter Aortic Valve Replacement (TAVR) patients is the ability of reproducing three-dimensionally the aortic, iliac, and femoral artery bifurcations from the data obtained from CT angiograms. This research helped in developing

software to generate STL files for use in 3D printing. A third facet of investigation was to investigate if CT scan-based virtual models can serve as a “Virtual Patient” and provide an inside view of the arteries, which otherwise is available only through angiography.

METHODS & CODES

Automated generation of patient-specific models from CT scan images for cardiovascular models and blood flow analysis is still a bottleneck in the application of HPC to clinical applications. This research was focused on developing a framework and software that can take CT scan images to create high-fidelity patient models. Since patient-specific calculations involve uncertainty in the data, a series of numerical simulations were carried out based on statistical distribution of the data. Furthermore, blood artery interaction models are computationally expensive and therefore Blue Waters resources were needed to further explore the mathematical attributes of our non-Newtonian constitutive models as well as the coupling schemes for blood–artery interaction.

CT-Scan Images for Patient Specific Model Construction: Fig. 1 shows CT scan images of the femoral arteries of a TAVR patient. Students in Masud’s group developed a computer program that identifies arteries and veins in these cross-sectional abdominal scans. Once they are identified and registered in a given scan, the

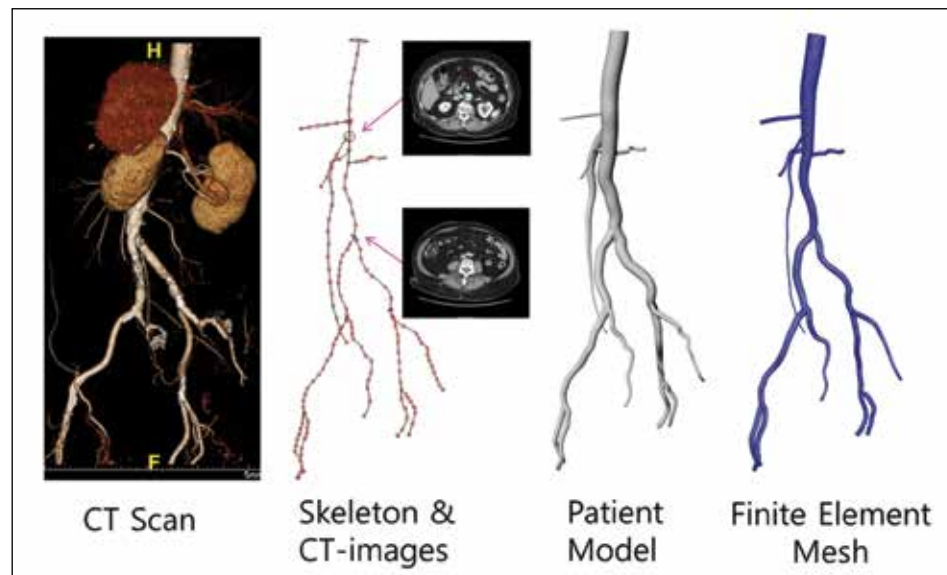


Figure 1: This is the CT Image of the branch structure created using the program developed by students, surface model of the patient, and the computational grid.

program automatically picks a continuation of an artery or a vein in the subsequent image via an algorithm that is based on close proximity of the geometric object to its corresponding image in the previous scan. Once the data is read from the CT scans, a program written by Elizabeth Livingston uses the point cloud method for developing patient-specific arterial tree models that are then used to generate 3D geometric models and Finite Element (FE) meshes. Various stages in the process of development of the model are shown in Fig. 1(c) and 1(d).

Simulation and visualization of patient-specific models: The patient-specific model was solved on the Blue Waters supercomputing platform. Spatial distribution of the pressure field for two time points during the cardiac cycle are shown in Fig. 2. A key step in assignment of outlet boundary conditions is the prescription of unique resistance values for each outlet, based on the morphometry laws. Masud and colleagues have developed dynamic resistance boundary conditions [2] that can accommodate patient-specific clinical data of flow rate distal to the region of interest, into the mathematical model. These boundary conditions help embed clinically measured patient-specific pressure variation into the computational model for a clinically relevant blood flow simulation.

Employing Data Mining techniques: Using open source software developed by Sandia National Labs, Daniel Sheehan developed a graphical visualization package that takes the large amount of computed data and converts it into time-dependent images and movies for easy visualization by doctors and clinicians.

3D Printing of Patient Geometry: Soonpil Kang developed a postprocessing capability that takes patient-specific models that are used in virtual modeling and develops STL files for 3D printing of the model. This feature is important for physical experimentation in the lab, thus providing an insight into blood flow rheology in patient specific models that was not possible before.

RESULTS AND IMPACT

Recent advances in computational fluid dynamics and image-based modeling permit determination of flow and pressure from CT scans, without the need for additional imaging, modification of acquisition protocols, or administration of medications. To analyze the massive amount of data in HPC one needs simulation-based images for clinical diagnostics. The software developed by Sheehan can help bring HPC and patient modeling for use by clinicians and surgeons.

WHY BLUE WATERS

Blue Waters was critical for both the development of cutting-edge software and the application of this software to perform large-scale biomechanics simulations. From a computational and algorithmic perspective, the newly developed coupled hierarchical multiscale methods tremendous benefit HPC by exploiting its large local resident memory on the processing nodes in favor of reducing the size of the global problem to be solved.

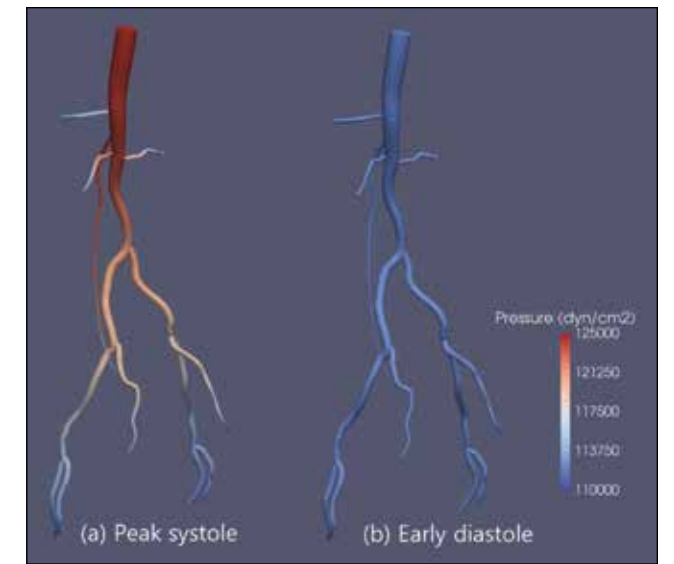


Figure 2: This is the pressure field at two instants in a typical cardiac cycle: (a) peak systole, (b) early diastole.

PUBLICATIONS AND DATA SETS

Weddell, J.C., J. Kwack, P.I. Imoukhuede, and A. Masud, Hemodynamic Analysis in an Idealized Artery Tree: Differences in Wall Shear Stress between Newtonian and Non-Newtonian Blood Models. *PLoS ONE*, 10:4 (2015), DOI: 10.1371/journal.pone.0124575.

Kwack, J., A. Masud and K.R. Rajagopal, Stabilized Mixed Three-field Formulation for a Generalized Incompressible Oldroyd-B Model, *International Journal for Numerical Methods in Fluids*, 83:9 (2017), pp. 704–734.

QUANTUM EFFECTS OF PROTON TRANSFER IN BIOLOGICAL SYSTEMS

Allocation: Innovation and Exploration/20 Knh
PI: James W. Mazzuca¹

¹Alma College

EXECUTIVE SUMMARY

Quantum effects of nuclei have been shown to affect chemical reactions in a variety of condensed-phase chemical systems, including certain biological enzymatic reactions. In this study, we investigate how the quantum effects of hydrogen and its isotopes affect the rate of charge transfer in a model biological system in which a water molecule acts as an intermediary between the donor and acceptor molecule. Our calculations show an abnormally high kinetic isotope effect, which is indicative of a quantum mechanical tunneling mechanism, and we see dramatically different behavior when these quantum effects are turned off. This study makes a strong case for examining larger, more realistic biological systems using this approach in the future.

RESEARCH CHALLENGE

In most chemical dynamics calculations, it is assumed that all nuclei are classical particles and all electrons are quantum particles represented by molecular orbitals. This approach is typically quite good, especially for high-temperature chemistry and heavy nuclei, or basically anything other than hydrogen. There are certain cases, however, where the quantum mechanical effects

of nuclei become important. Such cases include condensed-phase chemistry in which a hydrogen/proton transfer is the rate-limiting step. There is evidence to suggest that many biological chemical systems fit into this category, and we examine a model system in this study to determine how these quantum effects influence the mechanism of charge transfer.

These quantum effects can potentially impact both the efficiency and overall mechanism of charge transfer. The primary challenges in this field revolve around how exactly one should include the quantum effects of nuclei, but also how the electronic structure should be evaluated as the chemical system evolves from a reactant to a product state. In this particular study, we chose a model that we believe will capture the important physics of the proton transfer mechanism, but it is also simple enough that we can extract the important properties of this mechanism before examining larger, much more complicated systems.

Our calculations can give us a close look at exactly how this charge transfer process occurs, and we can use this information to improve our fundamental understanding of how nature has taken advantage of quantum mechanics in a way that has a profound influence on the chemistry of living organisms.

METHODS & CODES

We employed the ring polymer molecular dynamics (RPMD) method to include quantum effects of proton motion in our system. This method was chosen because it has been shown to be accurate for a wide range of applications, it scales linearly with respect to the system size, and the calculations can be done in parallel across many cores. Our code primarily simulates the physics of our chemical system as it evolves from a reactant state to a product state. To this end, many simulations containing many particles need to be performed in a highly parallel computing environment. Our group developed a Fortran-based code, which was highly scalable and well-suited for supercomputing environments such as Blue Waters.

RESULTS & IMPACT

Our calculations show an abnormally high kinetic isotope effect (KIE) for proton transfer in this model system. The kinetic isotope effect, in this case, is a measure of how much faster the reaction occurs when a proton, rather than a deuteron, is the charge carrier. In standard classical chemical dynamics calculations we observe a KIE of approximately 1.4 for all temperatures. In our quantum mechanical calculations we see a KIE of 35 at room temperature,




Figure 1: This image shows how potential energy changes as a proton is passed from a donor to acceptor molecule, with the help of an intervening water. The arrows represent two paths that the proton may travel to reach the product state, and quantum effects of this phenomenon were included using the ring polymer molecular dynamics (RPMD) method. Image from *J. Phys. Chem. A*, 121, pp. 819–826, 2017.




Figure 2: The inclusion of quantum effects in the proton transfer mechanism shows an extremely high kinetic isotope effect when the rate of proton transfer is compared to the rate of deuteron transfer, and this points to the importance of quantum effects in this process. Image from *J. Phys. Chem. A*, 121, pp. 819–826, 2017.

and an even higher KIE at lower temperatures. This indicates that quantum tunneling of nuclei plays an important role in this charge transfer mechanism, and a physically correct picture of this process depends on the inclusion of these quantum effects.

These calculations have provided us with a very close and detailed view of a physical phenomenon that plays an important role in many biological systems and would not have been possible without highly parallel supercomputing resources. The seemingly obscure phenomenon of quantum mechanical tunneling has been shown, in this study, to profoundly impact the rate of charge transfer in a simple model biological system, and this has similar implications for more complicated but fundamentally similar systems that are found in many different natural biological systems.

WHY BLUE WATERS

The RPMD method relies on taking averages among many chemical dynamics simulations with different initial conditions. Often, we need to perform hundreds of thousands of calculations in order to see converged results. Fortunately, these jobs can be run in parallel with one another, and a supercomputing environment in which we can run massively parallel jobs over thousands of cores is ideal. Additionally, the scalability of this method typically means that if enough compute cores are available, no job should take more than 30 minutes. The efficient parallelization, short job queues, and responsiveness of the system made Blue Waters a perfect platform for our code development and job submission. Additionally, the Blue Waters Student Internship Program provided a perfect gateway into using these resources while funding an undergraduate student's research experience on Blue Waters.

PUBLICATIONS AND DATA SETS

Mazzuca, J. W., and C. P. Schultz, Quantum Mechanical Enhancement of Rate Constants and Kinetic Isotope Effects for Water-Mediated Proton Transfer in a Model Biological System. *J. Phys. Chem. A*, 121 (2017), pp. 819–826.

THERMODYNAMIC CHARACTERIZATION OF CONFORMATIONAL LANDSCAPE IN PROTON-COUPLED OLIGOPEPTIDE TRANSPORTERS

Allocation: GLCPC/530 Knh

PI: Mahmoud Moradi¹

¹University of Arkansas, Fayetteville

EXECUTIVE SUMMARY

Proton-coupled oligopeptide transporters (POTs) use the inwardly directed proton flow to uptake small peptides and peptide-like molecules. The human POT transporters PepT1 and PepT2 provide the main route through which the body absorbs and retains dietary proteins. Human POTs also recognize several important families of peptide-like drug compounds such as β -lactam antibiotics. In order to function, POTs undergo large-scale conformational changes, whose characterization is the key in understanding the mechanism of transport by these proteins. The inward- (IF) to outward-facing (OF) structural transition of POTs, however, has remained elusive despite many experimental and computational efforts. We have employed all-atom molecular dynamics (MD) simulations along with novel enhanced sampling techniques to, for the first time, characterize the large-scale conformational changes of a bacterial POT transporter, namely GkPOT. By employing novel loosely coupled multiple-copy (LCMC) algorithms, our simulations provide a detailed description of the GkPOT conformational landscape, which sheds light on the structure–function relationship in POTs.

RESEARCH CHALLENGE

Membrane transporters provide the machinery to intimately couple active transport of materials to various forms of cellular energy. POT transporters couple the energy from proton flow to the transport of dipeptides, tripeptides, and their analogs [1]. A key feature of POTs is their substrate promiscuity [2], which is of great interest from a biomedical perspective. Human POT transporters PepT1 and PepT2, which play a key role in absorbing and retaining dietary proteins (in the small intestine and kidney, respectively) [3] recognize several important families of peptide-like drugs such as β -lactam antibiotics [4] and can improve the uptake of poorly absorbed/retained medications if attached to amino acids or dipeptides (prodrugs) [5]. Recent structural studies have resulted in several crystal structures of bacterial POTs [6–10], among which GkPOT, the POT transporter found in the bacterium *Geobacillus kaustophilus*, has the highest resolution (1.9 Å) [10]. These crystal structures, which are all in the IF state, provide the basis of our understanding of POTs' transport mechanism at the structural level. However, in order to function as active transporters, POTs are known to alternate between distinct IF and OF states. The conformation of the OF state and the transition pathway between the two functional states have remained elusive.




Figure 1: Top and side views of GkPOT transporter (cartoon representation) in its IF state (PDB: 4IKV [1]) and OF state (our model generated using enhanced sampling techniques), along with the water count along the pore as measured from equilibrium simulations of each model.




Figure 2: Strong and weak scaling of the Riemannian BEUS and SMwST algorithms [2] implemented in NAMD 2.10. Time-to-solution is reported for 40 ps of GkPOT simulations in explicit membrane and water (~85,000 atoms).

Previous MD simulations that relied on equilibrium simulations have failed to reliably characterize large-scale conformational changes such as those involved in POTs [10–11]. While the conventional MD can provide information on local conformational changes of a protein upon binding or unbinding of a substrate, ion, or proton, the global conformational changes observed are not often statistically significant. Functionally important conformational changes such as the IF–OF transition in membrane transporters typically occur on timescales beyond those accessible to conventional all-atom MD. The large-scale conformational changes, on the other hand, are typically studied using simplified modeling techniques such as coarse-graining, which could completely ignore or misrepresent the role of chemical events in the transport process. The main challenge in characterizing the large-scale conformational changes of proteins such as those associated with GkPOT is to reach the functionally relevant timescales without compromising the chemical details.

METHODS & CODES

We used a novel ensemble-based simulation approach [12–15] to reconstruct the entire transport cycle of GkPOT. Bias-Exchange Umbrella Sampling (BEUS) and String Method with Swarms of Trajectories (SMwST) are two independent LCMC algorithms employed. Both methods require parallel execution of hundreds of MD simulations of large protein systems with explicit representation of water and membrane, which requires hundreds of nodes. Our methodology was partly based on the techniques developed and used to investigate the thermodynamic cycle of the glycerol-3-phosphate transporter (GlpT) [14]. This is a rigorous and practical approach in characterizing large-scale conformational changes of proteins and their coupling to chemical events, which efficiently takes advantage of petascale computing. An important modification to the methodology to increase the accuracy of the results was introduced that involves a Riemannian formulation of free energy calculation and path-finding algorithms [15].

The software engine used for the simulations is NAMD: a highly scalable MD code implemented in Charm++, an object-based message-driven execution system based on C++. NAMD has been enhanced to support extremely scalable LCMC algorithms on massively parallelized computers of CPUs and GPUs. Multiple concurrent NAMD instances are launched with internal partitions of Charm++ and are located continuously within a single communication world. Messages between NAMD instances are passed by low-level point-to-point communication functions, which are accessible through NAMD's Tcl scripting interface.

RESULTS & IMPACT

The OF structure shown in Fig. 1 represents the first OF model of POT transporters, generated using our all-atom MD simulations in combination with LCMC algorithms discussed above. Our model is verifiably a stable OF structure since the subsequent equilibrium simulations show a water accessibility

consistent with an OF state (see the water profiles in Fig. 1). Our simulations also suggest that the full IF–OF transition requires the binding of both proton and substrate (data not shown). The pathways generated using parallel SMwST [14] and modified using our recently introduced Riemannian formalism [15] reveal that the proton-bound GkPOT cannot transition to the OF state (data not shown). Unlike the common simulation studies, which either rely on unbiased equilibrium simulations or simple representations (e.g., coarse-graining), our approach combines the accuracy of all-atom MD with the accessibility of long timescales provided by enhanced sampling techniques. The successful employment of these cutting-edge multiple-copy algorithms using Blue Waters resources opens a new window to structural biology of membrane transporters that bypasses the limitations of computational approaches to study structure–function relationships in these proteins.

WHY BLUE WATERS

We have explicitly shown that the unbiased all-atom MD, which is routinely used in the field, could be quite misleading in deciphering mechanistic features of membrane transporters due to the great gap in the timescales associated with the conventional simulations and the function of these proteins [11]. On the other hand, we have shown that the LCMC algorithms such as BEUS and SMwST [14–15] can be used to reconstruct unknown conformational transitions of membrane transport proteins. Unlike the conventional all-atom or coarse-grained MD that can be performed on sub-petascale machines, LCMC MD simulations of membrane transporters are well suited for large petascale computational resources such as Blue Waters as they require hundreds of nodes for a single job. We note that the “weak scaling” of these algorithms (Fig. 2) makes them particularly attractive for large petascale machines, as they can utilize hundreds of compute nodes with almost perfect efficiency.

PUBLICATIONS AND DATA SETS

Immadietty, K., J. Hettige, and M. Moradi. What can and cannot be learned from molecular dynamics simulations of bacterial proton-coupled oligopeptide transporter GkPOT? *J. Phys. Chem. B.*, 121:15 (2017), pp. 3644–3656, DOI: 10.1021/acs.jpcc.6b09733.

Moradi, M., K. Immadietty, and J. Hettige, Couplings between local and global conformational changes in proton-coupled oligopeptide transporters. *Biophysical Society Meeting*, New Orleans, La., February 11–15, 2017.

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

Allocation: NSF PRAC/2,300 Knh

PI: Vijay Pande¹

¹Stanford University

EXECUTIVE SUMMARY

The μ Opioid Receptor (μ OR) is a G-Protein Coupled Receptor (GPCR) that mediates pain and is a key target for clinically administered analgesics (i.e., pain medicines). The current generation of prescribed opiates—drugs that bind to μ OR—engender dangerous side effects such as respiratory depression and addiction due to the ligand-induced off-target conformations of the receptor. To determine both the key conformations of μ OR to atomic resolution as well as the transitions between them, long timescale molecular dynamics (MD) simulations were conducted using the Blue Waters supercomputer. These simulations predict new and potentially druggable metastable states that have not been observed by crystallography. We used statistical algorithms (e.g., tICA and Transfer Entropy) to perform our analysis and discover key conformations from simulation, presenting a transferable and systematic analysis scheme. Our approach provides a complete, predictive model of the dynamics, structure of states, and structure–ligand relationships of μ OR with broad applicability to GPCR biophysics and medicinal chemistry applications.

RESEARCH CHALLENGE

Because of its remarkable pain-reducing and induced-euphoria properties, opium has been used recreationally and for medical purposes for more than 4,000 years. Unfortunately, an epidemic of opioid abuse has increasingly bedeviled the United States [1]. An efficient opioid (medicinally perfect) would be a potent pain

reliever without side effects such as harmful respiratory effects or constipation, would show sustained efficacy in chronic treatments, and would not be addictive. The grand challenge questions are “How to design a perfect opioid?” and “How can we use the Blue Waters supercomputer to do that?”

Following the solution of the first structure, crystallography of GPCRs has both illuminated the structural biology and empowered medicinal chemistry of this class of receptors [2,3]. Recently, crystal structures of μ OR itself were solved in its “inactive” and “active” conformations [4,5]. However, other biophysical and pharmacological experiments have definitively demonstrated that μ OR traverses multiple functionally important conformational states [6]. These states are important in designing the “perfect opioid” but are not tractable by the experiments and crystallography.

Using Blue Waters, we attempted to discover states of μ OR that have different conformation compared to crystallography but are physiologically significant. In addition, we also tried to unravel how opiates of different scaffold classes tune the receptor toward distinct conformational energy landscapes.

METHODS & CODES

We performed multiple rounds of MD simulations on Blue Waters starting from the active with ligand, active APO, and inactive crystal structures. We used *MDTraj* (the package developed in the Pande lab) to convert and assemble the trajectories. Then, we used the Conformation software package written for and applied to the

featurization of this large GPCR MD dataset. To summarize, all residue–residue pairs within 6.6 Angstroms measured by closest heavy atom distance in either crystal structure were selected. Then, for each of these approximately 2,200 residue–residue pairs, both the closest heavy atom distance and Calpha distance were computed for each frame in each trajectory, leading to 4,400 “features” for each trajectory frame. Next, the Sparse tICA algorithm was applied to determine the reaction coordinates, or slowest collective degrees of freedom (up to the 10 slowest in this case), of the protein. Finally, a Markov State Model (MSM) was constructed with a lag time of 25 nanoseconds and prior counts of 1×10^{-5} . The equilibrium state probabilities from the MSMs were used individually in each condition to generate the free energy surfaces projected onto the features and tICA coordinates in Fig. 1.

RESULTS & IMPACT

Using Blue Waters, we mapped the complete free energy landscape of μ OR (Fig. 1) and discovered the novel and significant conformation states not tractable by experiments (Fig. 2). Using the state-of-the-art machine-learning algorithms developed at the Pande lab (tICA and relative entropy), we reduced the high dimensionality and sheer number of data points that render analysis so difficult. The important signal relaying residue switches we discovered (Fig. 2) shed new physical insight into the deactivation mechanism and pathway of μ OR. The newly discovered protein structures will be publically available for drug discovery projects.

We also produced a replicable framework that might enable other labs to gain actionable knowledge about their protein systems of interest through the “computational microscope” that is MD simulation. To this end, we published our simulations—both raw and featurized trajectories—as an open-source, downloadable resource, useful to both opioid researchers in particular as well as the wider structural biology and medicinal chemistry communities.

WHY BLUE WATERS

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

PUBLICATIONS AND DATA SETS

Harrigan, M. P., et al., Modeling Reveals Novel Intracellular Modulation of the Human μ 2 Selectivity Filter. *Scientific Reports*, 7 (2017), p. 632.



Figure 2: μ OR samples from at least two deactivating pathways that are gated by two distinct configurations of the DFWY motif. “Canonical” pathway (orange trace) entails (1) D^{147,32} relaxation away from Y^{326,43}, (2) W^{296,48} rotation toward Y^{326,43}, (3) F^{289,44} rotation toward binding pocket to maintain π -stacking with W^{296,48}, (4) I^{153,40} switches to take space previously occupied by F^{289,44}, (5) P^{245,50} relaxes away from I^{153,40} toward inactive position, triggering deactivation of TM6–TM5 packing and ultimately the relaxation of TM6 toward its inactive pose occluding the G Protein coupling domain.




Figure 1: Markov State Model reweighted free energy plots (kcal/mol) of μ OR projected onto tICA coordinates ω_1 and ω_4 in three different conditions: from left to right—Apo, BU72, and Sufentanil.

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

Allocation: NSF PRAC/3,710 Knh

PI: Benoit Roux¹

Collaborator: Jing Li¹

¹University of Chicago

EXECUTIVE SUMMARY

C-type inactivation of K⁺ channels is a molecular process of great physiological significance that affects the firing patterns of neurons in the central nervous system and the repolarization of cardiac cells in the heart. Despite years of studies, recent experimental results with semi-synthetic channels inserting unnatural amino acids in the structure suggest that a constricted filter conformation may not correspond to the C-type inactivated state. To resolve this issue, molecular dynamics simulations and free energy computations based on atomic models of the KcsA potassium channel were carried out. The computational results support the notion that the constricted conformation of the selectivity filter correspond to the functional C-type inactivated state of the KcsA channel. Our previous studies have put us on a strong path to execute this computational research [1–3].

RESEARCH CHALLENGE

Identifying the structural features associated with C-type inactivation in K⁺ channels is a very important goal and, in this regard, the bacterial KcsA channel has played a critical role. Although C-type inactivation has traditionally been associated with a functional behavior of voltage-gated channels such as Shaker, all the known functional markers of C-type inactivation are experimentally recapitulated by KcsA, providing a coherent view of the molecular determinants that affect inactivation. In addition, because the TTVGYGD canonical selectivity filter sequence is highly conserved, it is expected that the accessible conformational states of the selectivity filter, i.e., conductive and inactivated, ought to be fairly similar throughout most of the K⁺ channel family.

METHODS & CODES

We have carried out molecular dynamics simulations based on atomic models of the channels with the program NAMD using the CHARMM force field PARAM27.

RESULTS & IMPACT

Molecular dynamics simulations and free energy calculations were used to investigate the molecular mechanism of C-type inactivation. The impact of different inner-gate openings on the conformation and dynamics of the selectivity filter was systematically studied, and—for the first time—our 2D free energy calculation quantitatively characterized the conformational

preferences of the selectivity filter in which a fully open gate (~23 angstroms) highly favors a constricted filter, whereas a partially open gate (~16 angstroms) notably prefers to maintain a conductive filter. The spontaneous, rapid, and consistent conductive-to-constricted transition observed in simulations with a fully open gate reveals a hydrogen-bond network controlling the cooperativity for the tetramer constriction. By contrast, the partially-to-fully open transition for the inner gate is much slower, and the rearrangement of inner helices demonstrates considerable conformational heterogeneity. Analysis based on the simulations provides comprehensive details of the long-range allosteric gating coupling, and, more importantly, reveals the molecular basis for the delayed kinetics of C-type inactivation. These recent results are described in manuscripts that are currently under review.

WHY BLUE WATERS

Blue Waters offers the ability to carry out extensive umbrella sampling computations with Hamiltonian replica-exchange molecular dynamics simulations (i.e., window swapping) using multiple copies of the system.

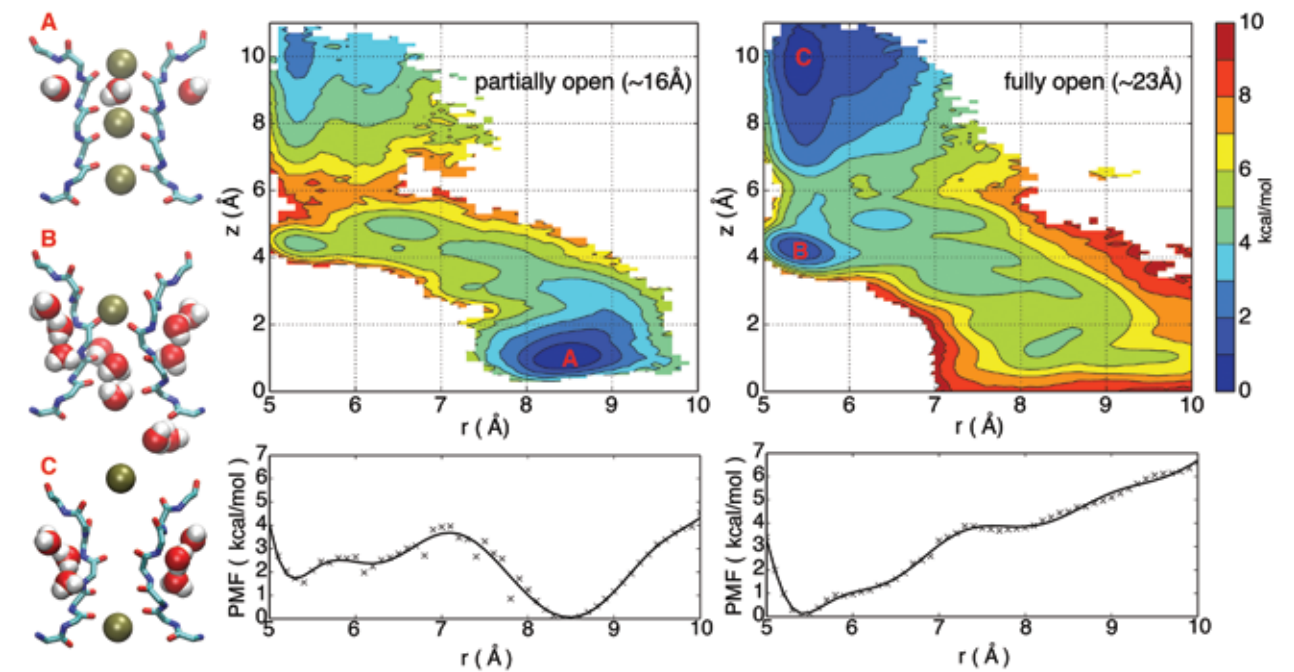


Figure 1: 2D PMF for the conductive-to-constricted transitions for the selectivity filter with different opening degrees of inner gate. The horizontal reaction coordinate r describes the width of the selectivity filter and is defined as the average cross-subunit pinching distance between the C α atoms of G77. The vertical reaction coordinate z indicates the position of the external K⁺ ion along Z axis relative to the center of the selectivity filter. The lower panel is the one dimension PMF along horizontal reaction coordinate r , with integration of the vertical reaction coordinate z . The typical conformations for three free energy basin (left) are shown in stick for protein and VDW representation for both water and K⁺ ion.

STUDYING CELLULAR PROCESSES THROUGH THE COMPUTATIONAL MICROSCOPE

Allocation: NSF PRAC/13,940 Knh

PI: Klaus Schulten¹ (deceased), Emad ¹ (successor)

Co-PIs: Juan R. Perilla², James C. Phillips¹, John E. Stone¹

Collaborators: Peijun Zhang³, Tatyana Polenova², Angela M. Gronenborn⁴, Christopher R. Aiken⁵, Adam Zlotnick⁶, Yale E. Goldman⁷, Robert J. Woods⁸

¹University of Illinois at Urbana-Champaign

²University of Delaware

³University of Oxford

⁴University of Pittsburgh School of Medicine

⁵Vanderbilt University School of Medicine

⁶Indiana University

⁷University of Pennsylvania School of Medicine

⁸University of Georgia

“In biology and in biomedicine, we have to realize that basically all organisms are large societies of molecules. We need a supercomputer to see that society for the first time.”

—Klaus Schulten (1947–2016)

EXECUTIVE SUMMARY

Blue Waters provides a powerful platform at the interface of biology, physics, and computer science that is being leveraged to shed light on cellular processes toward direct impact on public health. Utilizing Blue Waters as a “computational microscope,” researchers are revealing key new insights into viral infection and the mechanisms by which it can be disrupted with drugs. Virus capsids represent promising drug targets, and the world’s first microsecond molecular dynamics (MD) simulations of these remarkable molecular machines performed at all-atom resolution—a feat made achievable only by the petascale computing resources of Blue Waters—expose details underlying their biological functions and potential vulnerabilities to drug compounds.




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

RESEARCH CHALLENGE

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

METHODS & CODES

MD simulations provide a powerful technique to investigate the dynamical structure and chemical-physical properties of virus capsids [1]. Our work has demonstrated that, when performed at the all-atom level of detail, simulations are capable of capturing even subtle effects on capsid structure and dynamics induced by bound drug molecules [2]. We employed Nanoscale Molecular Dynamics, or NAMD [3] for our simulations, a highly scalable MD code optimized specifically for Blue Waters that boasts a long and successful track record of deployment on the machine. While all-atom simulation of virus capsids comes at great computational expense, access to NAMD on Blue Waters has enabled us to reveal critical new insights into the structure and function of capsids, as well as to suggest mechanisms by which drug molecules can disrupt them. Importantly, our discoveries were inaccessible to state-of-the-art experimental methods, and were made possible only through access to the petascale computing power of Blue Waters.

RESULTS & IMPACT

HIV-1. Infection with human immunodeficiency virus type 1 (HIV-1) is classified as a global pandemic by the World Health Organization. Due to the extremely high mutation rate of the virus, new drug treatments must be constantly developed. We previously utilized Blue Waters to solve the all-atom structure of the mature HIV-1 capsid [1], providing an essential platform




Figure 2: (a) The HIV-1 capsid simulation system contains more than 1,300 proteins immersed in a native solvent environment, totaling over 64 million atoms. (b) The HBV capsid simulation system contains 240 proteins immersed in a native solvent environment, totaling about 6 million atoms.

for study of its dynamical and structural properties, as well as its interaction with host-factors [3] and drug molecules [2]. At present, we have leveraged Blue Waters to perform the most monumental all-atom simulation achieved to date, characterizing the dynamical behavior of the HIV-1 capsid system (64 million atoms, Fig. 2a) over the timescale of one microsecond.

Our simulation reveals charge-specific channels in the surface of the capsid through which solvent ions translocate; these channels are likely capable of regulating translocation of DNA nucleotides, which must travel from the exterior to the interior of the capsid during reverse transcription. Further, analysis of the capsid’s collective motions reveals a belt dividing the capsid into two hemispheres, suggesting a potential mechanism by which the capsid may break apart to deliver its genome. Our results, which additionally included characterization of the electrostatic and acoustic properties of the capsid, indicate new avenues for the development of drugs that seek to disrupt the capsid by altering its complex biophysical properties.

HBV. Hepatitis B virus (HBV) is a leading cause of liver disease worldwide, including cancer, and the World Health Organization estimates that 240 million people suffer from chronic infection. The capsid of HBV is icosahedral, and most prior structural studies imposed assumptions of icosahedral symmetry to enhance experimental resolution and reduce computational complexity. We previously utilized Blue Waters to perform the first simulation of the HBV capsid undertaken without symmetry bias and characterized drug-induced structural changes likely related to the drug’s mechanism of capsid disruption [2]. At present, we have leveraged Blue Waters to perform the most extensive unbiased simulations achieved for an icosahedral virus capsid to date, characterizing the dynamical behavior of the HBV capsid system (6 million atoms, Fig. 2b) in the presence and absence of three distinct drug compounds over timescales of one microsecond.

Our simulations reveal remarkable asymmetry in capsid motions, supporting hypotheses that the capsid can distort asymmetrically to accommodate unevenly-distributed internal strain resulting from

conversion of pgRNA (pre-genomic ribonucleic acid) to DNA during reverse transcription. Further, the localization of ions during simulation provides a structural explanation for experimentally observed enhancement of capsid assembly under high salt concentrations. Finally, our results capture a variety of morphological and allosteric changes induced by bound drug molecules, providing insight into their complex mechanisms of action.

Viral host factors. Integrating the results of additional simulations with experiments performed by our collaborators, we also reveal key details underlying the structures and mechanisms of large host-cell molecules that are implicated in viral infection, including the myxovirus resistance protein B (HIV-1) and cytoplasmic dynein (HIV-1 and HBV).

WHY BLUE WATERS

Due to their formidable computational expense, microsecond simulations of virus capsids are only possible on a petascale machine like Blue Waters. Capsid systems encompass millions of atoms, and computing the interactions among such large numbers of particles over such long timescales can take months, even on tens of thousands of processors. Further, analysis of the colossal data sets generated by our capsid simulations was made feasible only through access to the massively parallel computing power and high-performance Lustre filesystem provided by Blue Waters. The exciting discoveries revealed by our research underscore an essential role for petascale resources like Blue Waters in the development of anti-viral treatments, and demonstrate that access to leadership-class computing facilities holds the potential for significant impact on overall public health.

PUBLICATIONS AND DATA SETS

Perilla, J.R., and K. Schulten, Physical properties of the HIV-1 capsid from all-atom molecular dynamics simulations. *Nat. Commun.*, 8 (2017), p. 15959, DOI: 10.1038/ncomms15959.

Lippert, L.G., et. al., Angular measurements of the dynein ring reveal a stepping mechanism dependent on a flexible stalk. *Proc. Natl. Acad. Sci. USA*, 114:23 (2017), pp. E4564–E4573, DOI: 10.1073/pnas.1620149114.

Perilla, J.R., and G. Zhao, et al., CryoEM structure refinement by integrating NMR chemical shifts with molecular dynamics simulations. *J. Phys. Chem. B.*, 121:15 (2017), pp. 3853–3863, DOI: 10.1021/acs.jpcc.6b13105.

Stone, J.E., J.R. Perilla, C.K. Cassidy, K. Schulten, Chapter 11: GPU-accelerated molecular dynamics clustering analysis with openACC. in *Parallel Programming with OpenACC*, Elsevier (2016), pp. 215–240.

Zhang, H., et al., HIV-1 capsid function is regulated by dynamics: Quantitative atomic-resolution insights by integrating magic-angle-spinning NMR, QM/MM, and MD. *J. Am. Chem. Soc.*, 138:42 (2016), pp. 14066–14075, DOI: 10.1021/jacs.6b08744.

Thieker, D.F., J.A. Hadden, K. Schulten, R.J. Woods, 3D implementation of the symbol nomenclature for graphical representation of glycans. *Glycobiology*, 26:8 (2016), pp. 786–787, DOI: 10.1093/glycob/cww076.

UNDERSTANDING THE PROTEIN ALLOSTERY IN KINASES AND GPCRS

Allocation: Blue Waters Professor/305 Knh; Illinois/300 Knh
PI: Diwakar Shukla¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Kinases and GPCRs (G-protein-coupled receptors) are key cellular signaling proteins involved in various pathophysiological functions. These proteins are coupled to another protein or effector molecule and allosterically modulate the biological activity of the downstream signaling proteins. The allosteric-mediated effects are poorly understood and the understanding of the molecular basis of allostery remains elusive. We investigated the cyclin-mediated effects on cyclin-dependent kinase (CDK) and endogenous sodium ions on GPCRs.

To understand the allosteric effects on CDK family kinases and sodium ions on GPCRs, we performed extensive simulations of computationally reconstructed ancestral (protein of CDK) CMGC family kinase, CMGI, and sodium ion-binding mechanism in various GPCRs. Using Blue Waters, we determined the long-range coupling of protein domains and ions that affect the pharmacology. Our results show that the helix at the beginning of the A-loop (activation loop) locks the modern CDK proteins in inactive conformation. In GPCRs, the sodium ions bind with different specificity in various GPCRs. This is the first study reported to date that estimates the free energy profiles show specificity of ion binding. Our finding sheds light on the allosteric interaction of biological molecules at an atomic level.

RESEARCH CHALLENGE

G-protein-coupled receptors and protein tyrosine kinases are the two large protein families that represent two prominent pathways for cellular signaling. Allosteric coupling of these proteins with effector proteins or modulators restricts the proteins in specific conformational state, which triggers the signals to the downstream proteins. Kinases are cellular signaling proteins involved in a variety of cellular pathways that control cell growth. They coordinate the cell cycle by switching between active and inactive states, which are considered as on/off states. In CDK, the association of another protein (cyclin) is required for the kinase activation. Due to the dependence of CDK's activity on cyclin, its activity can be further regulated. CDKs can be activated to stimulate different signals in different cell phases via cyclin-CDK intermolecular regulation. GPCRs are flexible molecules that shift the equilibrium from inactive to active states to transduce signals to the downstream proteins. In GPCRs, sodium ions bind at the intracellular site and act as an allosteric modulator, thereby restricting the receptor-mediated signaling. Our study aims to discover the long-range coupling effects on kinase and GPCR and its effect on signaling mechanisms.

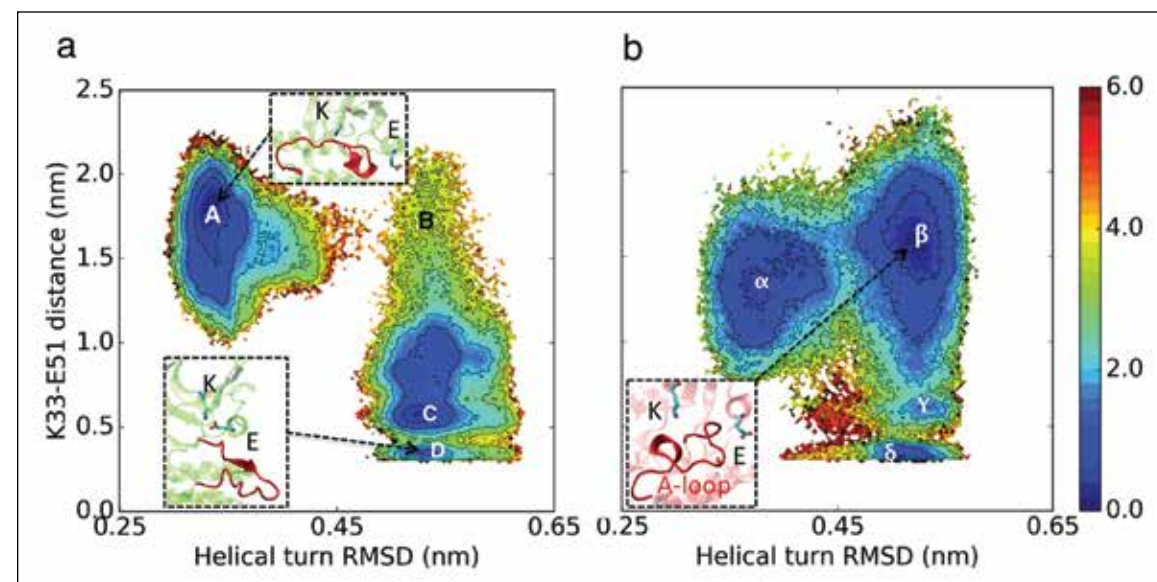


Figure 1: Free energy landscape of CDK and CMGI kinase indicating the key differences in their free energy landscapes.




Figure 2: Free energy profiles of sodium ion binding to various GPCRs. The various stages of ion binding have been shown and the ion recognition and binding residues are shown respectively.

A-loop to unfold with a lower barrier. These free energy landscapes support the hypothesis that the differences in the stability of the helical turn between CDK2 and CMGI are one of the factors causing differences in their activation process.

Our MSMs built from GPCR simulation data provide kinetic data unobtainable from experiments describing ion-binding pathways in addition to the crucial residues that drive the ion-mediated interaction. The sodium ion is recognized by residues at the extracellular loop in all other GPCRs except S₁P and LA₁R as the ion interacts with extracellular residues involving different helices. The residues slowly diffuse to the primary ligand binding site and form stable interactions with Asp^{3.32} (Ballesteros Weinstein numbering). In LA₁R, the ion occupies the orthosteric site by interacting with Glu^{7.43} and Tyr^{2.57} in S₁P. Finally, the ion reaches the allosteric site and establishes stable interaction with Asp^{2.50}. The free energy profiles of sodium ion binding have provided the various degrees of ion-binding specificity across the GPCR family. M₃ and A_{2A}AR bind more specifically to sodium ions compared to other GPCRs. The ion forms a stable, extended interaction at various intermediate states for β₂AR. The receptors H₁, 5HT_{1B}, 5HT_{2B}, and LA₁R exhibit very similar energy profiles. S₁P has less specificity to sodium ions and binds loosely compared to other GPCRs.

We determined the free energy profile of sodium ion binding for the first time using MD simulation. The sodium binding site can be used as a potential allosteric drug-binding site to lock the receptor in the GPCR inactive state to restrict the downstream signaling mechanisms. Our results provide molecular-level details of the specificity of ion binding to various GPCRs.

WHY BLUE WATERS

The understanding of long-range network interaction in proteins requires several hundred microseconds-long simulations. Blue Waters provides the necessary computer architecture needed to carry out these computational studies. The current GPU and CPU framework allows us to run hundreds of parallel simulations. Without Blue Waters, the current work would not be possible.

METHODS & CODES

To investigate the allosteric effects of CDK and GPCRs, we performed extensive molecular simulations of both proteins using Blue Waters. The simulation data were clustered based on the kinetically relevant state, and a Markov state model (MSM) for each system was constructed.

RESULTS & IMPACT

Our MSM reveals that kinase activation is a highly complicated dynamic process that involves multiple intra- and intermolecular switches that regulate the kinase conformational preferences. We hypothesized that in the absence of cyclin, at least one regulatory switch in CDK2 is “off” while it is “on” in the CMGI. To find the suitable regulatory switches, all available crystal structures of CDK2 have been compared. Four possible molecular switches have been identified including the conserved K-E bond, alignment of the R spine residues, the helical twist in the A-loop, and availability of DFG-loop aspartate residue for substrate interaction.

Formation of a helical region in the beginning of the activation loop is an important characteristic of inactive CDK structures, which prevents binding of the substrate protein (PDB ID: 3XPR and 3PXF). The helical turn pushes the αC-helix out, thereby acting as a molecular switch that could alter the cyclin dependence of CDK kinases. This auto-inhibitory mechanism is observed in several kinases such as CDKs, Src and Abl [2]. The crystal structural analysis shows high degrees of correlation between the presence of the helical turn and the K-E bond. The simulation results mapped into a two-dimensional conformational landscape of K-E distance versus the root mean squared deviation of the helical turn with respect to the inactive structure (PDB ID: 3XPR) shows a barrier of ~6 kcal/mol for unfolding of the helical turn in CDK2, whereas the barrier is not observed due to a stable intermediate state in CMGI (Fig.1 minima B and β respectively). The helical secondary structure moves from the beginning of the A-loop toward its end in the β intermediate state in CMGI (Fig. 1 b), which helps the

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

Allocation: NSF PRAC/707.6 Knh
PI: Ivan Soltesz¹
Collaborator: Michael Hines²

¹Stanford University
²Yale University

EXECUTIVE SUMMARY

We have extensively validated a first-of-its-kind 1:1 scale, strictly biological data-driven computational network model of the CA1 region of the rodent hippocampus. The model spontaneously generates theta oscillations, which occur in the brain during locomotion and spatial navigation. Theta oscillations are critical to the formation of episodic memory, which associates events with places. We have simulated a large number of perturbed network configurations in order to determine the range of parameters and conditions that are necessary and sufficient for the emergence of the theta oscillations. In addition, we have conducted initial simulations that mimic realistic spatial information input from the entorhinal cortex, which is the principal source of information for the hippocampus. These results confirm that our computational models are capable of replicating major dynamic phenomena in the hippocampus and provide a stepping stone to modeling cognitive processes associated with the hippocampus such as episodic memory.

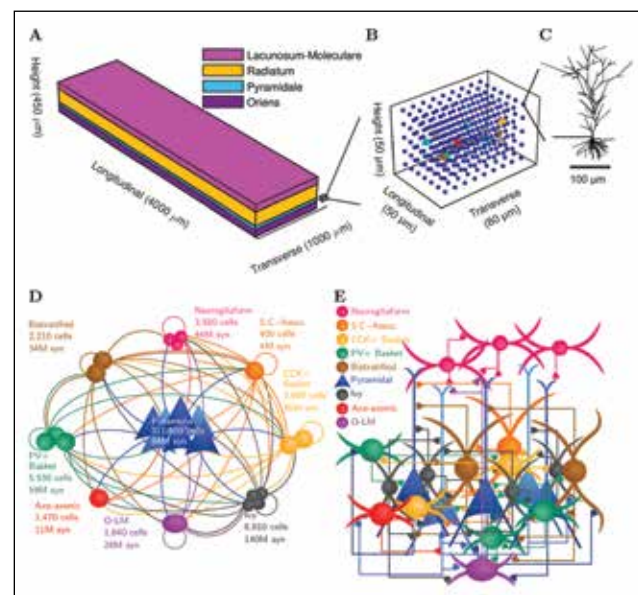


Figure 1: Structure of the CA1 network model. (A) The model network has dimensions identical to the rat hippocampus, (B) model cell bodies laid out in their respective layers, (C) detailed morphology and synapse placement for each model neuron, (D) illustration of the connectivity types, (E) the characteristics of each connection are constrained by experimental data.

RESEARCH CHALLENGE

The main goal of the project is to study the mechanisms that govern the emergence of characteristic oscillatory behavior in the hippocampal network and its implications for information processing using 1:1 scale biophysical computational models that are closely based on electrophysiological, morphological, and imaging data. The hippocampal circuits that store and recall information are comprised of diverse cell types, each exhibiting distinct dynamics and complex patterns of synaptic connectivity. Thus, even highly specific experimental perturbations of a single component of these neuronal circuits can have highly nonlinear and counterintuitive effects on their internal dynamics and output. Our computational models offer a framework to integrate knowledge and quantitatively predict how each element of a neuronal network is expected to respond to specific perturbations. Our computational models completed so far represent a major milestone in the development of large-scale, anatomically and biophysically realistic models of the brain, which allow for the generation and testing of hypotheses concerning synaptic and network mechanisms of behaviorally relevant oscillations with unparalleled biological realism and precision.

METHODS & CODES

Our principal simulation environment is NEURON 7.4 [1]. NEURON is designed to simulate neuronal models that are described in terms of the membrane properties and geometric structure of neurons [2], and supports computationally efficient representation of connections among neurons in a network [3]. NEURON is formulated around the notion of continuous cable “sections” that can be connected together to form any kind of branched cable. A section can be assigned properties that vary continuously with position along its length. User-defined biophysical properties of membranes, such as ion channel dynamics, are described in terms of differential equations, kinetic schemes, and sets of simultaneous equations. These model descriptions are compiled to C, so that membrane voltage and gating states can be computed efficiently using an implicit integration method optimized for branched structures [2]. NEURON is very well supported and has been used in more than 1,532 publications as of September 2014 and can be fully parallelized [2]. NEURON includes an introspective interpreter for simulation code (HOC

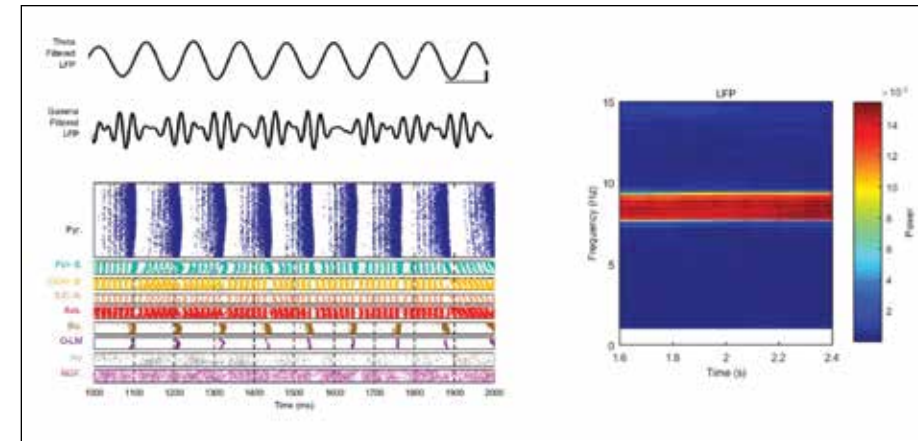


Figure 2: Oscillatory network activity of the CA1 model. (A–B) Local field potentials (LFP) generated by the model, filtered at (A) theta (5–10 Hz) or (B) gamma range (25–40 Hz) show clear oscillatory activity. (C) Neuronal spike times generated by the model. (D) Spectrogram of LFP shows the stability and intensity of the theta oscillation.

or Python) that is flexible and supports a wide range of model description paradigms.

RESULTS & IMPACT

We have made major advances toward achieving our aims. Specifically, we have extensively validated a first-of-its-kind 1:1 scale, strictly biological data-driven computational network model of the CA1 region of the rodent hippocampus. The model spontaneously generates theta oscillations, which occur in the brain during locomotion and spatial navigation. In addition, we have studied large numbers of perturbed network configurations in order to determine the range of parameters and conditions that are necessary and sufficient for the emergence of the theta oscillations. Furthermore, we have advanced our technique for generating biologically realistic dendritic trees and applied it toward constructing a 1:1 scale model of the dentate gyrus (DG), another major hippocampal region and a key component of our modeling efforts. We have conducted a number of control simulations to validate the dynamic behavior of this model and tuned its synaptic parameters so that its spatial coding properties are consistent with recently published experimental results. These advances, which have never been achieved in computational neuroscience before, are critical for ensuring that our computational models have capabilities for processing spatial information that are comparable with those of the rodent hippocampus.

The CA1 model that we have completed during the project period is capable of spontaneous generation of theta and gamma oscillations, replicating published studies. In addition, we have analyzed the firing patterns of each neuron type in the network. Next, we have conducted extensive work on perturbing various network and neuronal parameters to determine the relative contributions of each component to the generation of biophysically relevant oscillations. The model neurons exhibited firing patterns relative to the global network oscillations that are consistent with recently published experimental data. This is a major advance that indicates that the characteristic firing patterns of hippocampal neurons are possibly a function of the intrinsic network wiring.

Furthermore, we have tested whether theta rhythm was differentially sensitive to the contribution of each inhibitory neuron type, as well as the intrinsic properties of each interneuronal type by respectively muting the output of different neuron classes or configuring all neuron models with the same electrophysiological profile. Theta oscillations were not apparent in any of these perturbed configurations and, therefore, these results indicate, for the first time, that interneuronal diversity itself is an important factor in the emergence of biophysical oscillations in the CA1 network.

WHY BLUE WATERS

In order to simulate realistic spatial navigation and formation of place cells, the simulations must be long enough to reflect the brain’s changing representation of space during locomotion. A typical behavioral experiment with animals running on a linear track has a duration of tens of seconds, and therefore our simulations must have a minimum duration of 10 seconds in order to be comparable to behavioral experiments. Simulations of our CA1 model of 10 seconds of physical time took 14 hours to run on 1,024 Blue Waters nodes. It would not have been practical or affordable to conduct simulations of such scale on other publicly or commercially available computational platforms. Our research plans require the ability to run even longer simulations of the combined hippocampal model, and for those only Blue Waters can provide the necessary computational capacity.

PUBLICATIONS AND DATA SETS

Bezaire, M., et al., Interneuronal mechanisms of hippocampal theta oscillations in a full-scale model of the rodent CA1 circuit. *eLife*, 5:e18566 (2016), DOI: 10.7554/eLife.18566.

Bezaire, M., I. Raikov, K. Burk, D. Vyas, and I. Soltesz, Computational model of CA1 network on ModelDB, accession number: 187604.

Bezaire, M., I. Raikov, K. Burk, D. Vyas, and I. Soltesz, Simulation results from full scale and rationally reduced network models of the isolated hippocampal CA1 subfield in rat on CRCNS.

STRETCHING THE CADHERIN MOLECULAR VELCRO® OF CELL-CELL JUNCTIONS

Allocation: GLCPC/350 Knh
PI: Marcos Sotomayor¹

¹The Ohio State University

EXECUTIVE SUMMARY

The extracellular domains of classical cadherins form a Velcro®-like surface that glues cells together in the presence of calcium. This is essential for cell-cell adhesion and multicellular life, but the mechanics of multi-cadherin complexes is poorly understood. Using VMD (visual molecular dynamics) and NAMD (nanoscale molecular dynamics) on Blue Waters we were able to perform atomistic simulations of realistic adhesive systems that included up to twelve cadherin-cadherin bonds with 3.7 million atoms. These simulations revealed how cadherins respond to mechanical stimulation that mimics physiological forces, such as those experienced by tissue stretched by blood pressure, muscle movement, or impact with a foreign object. In addition, we were able to determine the collective behavior of cadherins in the absence of calcium and the relaxation of partially ruptured cadherin lattices as it may occur during wound healing. These simulations shed light on the basic molecular mechanisms that underlie tissue development, mechanics, and repair.

RESEARCH CHALLENGE

Selective and robust adhesion between cells is essential for multicellular life. Classical cadherin proteins are found on the surface of cells and act as molecular Velcro® that glues adhesive cells together [1,2]. The adhesion mediated by cadherins depends on calcium ions and is the basis for the formation of organs and the maintenance of tissue integrity in humans and in multiple species across the animal kingdom. The molecular architecture of a key cadherin protein named E-cad revealed how its extracellular domain protrudes from the cell surface as a hook that engages with another E-cad from an adjacent cell [3]. This E-cad/E-cad complex effectively forms a bond that links adjacent cells together in cell-cell junctions. E-cad also has a transmembrane helix and a cytoplasmic domain that anchors the protein to the actin cytoskeleton, thus restricting its motion in the membrane plane. A single E-cad/E-cad bond is weak, so the strength of E-cad-mediated adhesion derives from multiple E-cad/E-cad bonds arranged in a large and robust lattice (see Fig. 1). Previous studies have provided a detailed description of how single E-cad molecules and bonds behave and respond to force in the presence and the absence of calcium [4–7]. However, very little is known about the collective behavior of a “society” of E-cad molecules in a lattice. Understanding how lattices of E-cad proteins respond to forces is essential to determine how tissue is maintained and broken, how wound healing occurs

at the molecular level, and how mutations that change E-cad mechanics lead to disease.

METHODS & CODES

To test how an E-cad lattice responds to forces that mimic the effect of cells being stretched by blood pressure, muscle movement, or impact with a foreign object, we generated atomistic models with lattices containing one to twelve E-cad/E-cad bonds [3]. Mechanical properties of cadherin proteins are best studied in simulations where all atoms in the system, including critical calcium ions, are explicitly modeled, and where the system is hydrated with water molecules and ions that mimic the native physiological environment of cell-cell junctions. Such systems were built and simulated with VMD and NAMD [8,9], a pair of programs for molecular visualization and dynamics simulation created by the Theoretical and Computational Biophysics Group at the University of Illinois at Urbana-Champaign. This software suite can handle large atomistic systems, which in our case encompassed up to 3.7 million atoms. NAMD can efficiently use thousands of cores in Blue Waters, which allowed us to simulate these large systems and use steered molecular dynamics (SMD), a technique in which forces are applied to proteins to test their mechanical response *in silico* [10].

RESULTS & IMPACT

Systems including one, four, and twelve E-cad/E-cad bonds in a lattice arrangement were simulated for twenty or more nanoseconds, in equilibrium and under tension. To make the systems realistic, we also incorporated constraints that mimic cytoplasmic attachment to the cytoskeleton. To the best of our knowledge, these are the first all-atom SMD simulations revealing detailed dynamics of a complete and realistic cadherin lattice. We were able to observe how the initially curved E-cad/E-cad bonds became straight upon application of force, suggesting that cadherins at the cell-cell junction can act as molecular shock absorbers that extend without the E-cad/E-cad bond breaking at low force. As tension increased, we observed E-cad/E-cad bond rupture without any unfolding of the individual E-cad molecules, despite lateral interactions that stabilized the lattice. To test reversibility, we released forces and observed the recovery of curvature for individual E-cad molecules, hinting at the steps required to reestablish cell-cell adhesion after rupture. We also explored the behavior of the lattice in the absence of calcium, which resulted in disordered and floppy E-cad chains. Additional




Figure 1: Cell-cell adhesion junction with twelve E-cad/E-cad bonds involving the extracellular domains of 24 cadherin molecules in a lattice arrangement with 12 cadherin bonds (3.7 million atoms). The system is shown after equilibration (top left) and after stretching (center). The membrane planes of adjacent cells are shown in light blue. Molecular image prepared by B. L. Neel using VMD [8].

simulations will explore the elastic response of the calcium-free lattice as well as of lattices that are not anchored to the cytoskeleton or that are subjected to shearing stress. Overall, NAMD simulations of cadherin lattices are providing an unprecedented atomistic view of the mechanics of cellular adhesion.

WHY BLUE WATERS

Molecular dynamics simulations of large atomistic systems are extremely computationally demanding and cannot be divided into smaller independent simulations to be distributed among poorly networked computational resources. Only a fast networked and massively parallel system like Blue Waters can be used to achieve the multi-nanosecond time scales relevant to study the dynamics and elasticity of cadherin complexes.

SIMULATION OF VIRAL INFECTION PROPAGATION THROUGH AIR TRAVEL

Allocation: NSF PRAC/150 Knh
PI: Ashok Srinivasan¹
Co-PIs: Sirish Namilae², Anuj Mubayi³, Matthew Scotch³, Robert Pahle³
Collaborator: C.D. Sudheer⁴

¹Florida State University
²Embry-Riddle Aeronautical University
³Arizona State University
⁴Argonne National Lab

EXECUTIVE SUMMARY

There is direct evidence that commercial air travel spreads common infectious diseases including influenza, SARS, tuberculosis, and measles. This has motivated calls for restrictions on air travel, for example, during the 2014 Ebola outbreak. Such restrictions, however, carry considerable economic and human costs. Ideally, decision-makers ought to take steps to mitigate the likelihood of an epidemic without imposing the above costs. Science-based policy analysis can yield useful insight to decision-makers.

The effectiveness of a policy depends on human response to it. Given inherent uncertainties in human behavior, we simulate a variety of scenarios and identify the vulnerability of policies under these potential scenarios. Supercomputing is used to deal with the large number of scenarios and the need for a short response time in case of national emergencies. Our results identify new boarding procedures that can result in substantial reduction in the risk of the spread of Ebola and SARS.

RESEARCH CHALLENGE

Our goal is to develop models and a novel methodology that can provide insight to decision-makers on policies and procedures that will reduce the likelihood of infection spread during air travel. In addition, our research contributions promise major advances in the disciplinary areas of our expertise—pedestrian movement modeling, mathematics, epidemic modeling, computer science, and bioinformatics—with a consequent transformative effect on transportation infrastructure and management.

METHODS & CODES

We modeled pedestrian movement during air travel as particles based on a force-field approach proposed by Helbing, et al. [1]. Both pedestrian density and speed of immediate neighbor in a pedestrian line determine pedestrian speed and trajectory [2–3]. Our modifications incorporate these aspects into the pedestrian movement model. The pedestrian trajectory information is then integrated with a discrete-time stochastic Susceptible–Infected (SE) model for infection transmission. When i_0^i infectives come into contact with m_i susceptibles estimated by the pedestrian movement model, the newly infected and the probability of their infection

can be estimated using a binomial distribution approximated as Poisson distribution. This accounts for demographic stochasticity and variations in susceptibility of the population. This approach (see Fig. 1) provides insight into the consequences of policy choices that change a passenger’s behavior. We input this information to a global phylogeography model to assess the impact of these policies at global scale. Several airports (Phoenix, Tallahassee, and Daytona Beach) have expressed interest in collaborating with us on developing feasible policies.

Since inherent uncertainties in human behavior and insufficient data during the initial stages of an epidemic make prediction difficult, we parameterize the sources of uncertainty and evaluate vulnerability under different possible scenarios. Even if individual parameters cannot be reliably estimated due to identifiability issues, it may still be possible to precisely estimate a compound quantity of interest, such as the reproductive numbers. We use Blue Waters to deal with the computational load that arises from a large parameter-space.

We link the results of this fine-scale model with a phylogeography model, which uses genetic mutation information and geographic locations of viruses to model the spread of epidemics across large

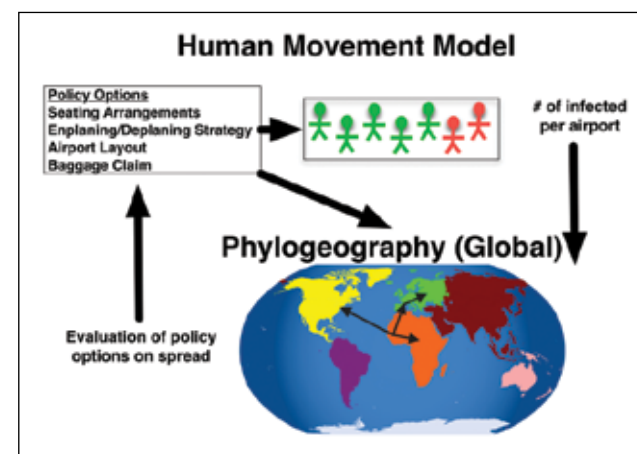


Figure 1: Schematic for combining fine-scale model at airport level with global level phylogeography model. The impact of airport-level public health policies on global disease spread can be evaluated through this approach.




Figure 2: Infection distribution profile at different boarding strategies for: (a) Boeing 757–200 capable of 182 passengers, (b) Airbus A320 capable of 144 passengers. The pictures on the bottom show the corresponding aircraft seating configurations with seats (blue dots) and pedestrians (green dots).

geographic scales. For the phylogeography work, we used Blue Waters to analyze 264 full-genome Ebola sequences from Guinea, Liberia, Sierra Leone, Italy, the United Kingdom, and the U.S., then used the BEAST software on Blue Waters to implement the phylogeography model. As part of this, we specified a generalized linear model (GLM) to study predictors of spread between discrete locations. We used five predictors: worst-case (infection) scenario, best-case scenario, airplane passenger flux between airports, and sample size from the origin and destination. We ran multiple Bayesian Markov-chain Monte Carlo simulations to produce posterior estimates, and also implemented a traditional discrete phylogeographic run without a GLM that evaluated the transmission rate between locations and estimated the routes with strongest support

RESULTS & IMPACT

In prior work, we used the above approach with Ebola, studying the impact of different procedures for boarding, disembarkation, and seat assignment on infection spread. Our results show promise for substantial impact from such choices. For example, we showed that on a 182-passenger Boeing 757 airplane, random boarding can lead to substantial reduction in infection transmission compared with current zone-wise boarding (see Fig. 2). We obtained similar results showing the potential for changes in in-plane movement, deplaning procedure, seating arrangement, and plane sizes in reducing the likelihood of infection transmission. The improvements obtained for individual flights by these policy changes can be of substantial benefit over the course of an epidemic. Based on the transportation data from 2013, if unrestricted air travel were to have occurred during the 2014 Ebola epidemic, then the probability of generating 20 infectives per month from air travel could have been reduced from 67% to 40% using better pedestrian movement strategies. This could further be reduced to 13% by exclusively using smaller 50-seat airplanes. The approach is generic and can be used for other directly transmitted diseases such as SARS.

WHY BLUE WATERS

In a new emergency, due to lack of data, one usually needs to model for a variety of scenarios. This leads to a large parameter space of uncertainties, which requires a large computational effort. In addition, the models typically need fine-tuning, which leads to an iterative process where the model is repeatedly tuned based on results from its previous validation steps. Consequently, rapid turnaround time is critical, which requires massive parallelism. Such parallelism becomes even more crucial during the course of a decision meeting, where results are typically needed in the span of a couple of minutes.

Thus cloud computing is not effective in national emergencies because of the large parameter space and need for quick runtimes during policy decision meetings. After extensive optimization (with support from the Blue Waters team) the simulation time per run is currently around 20 minutes. This can be reduced further only by parallelizing an individual simulation, which is communication-intensive. Again, a cloud environment is not useful.

Among supercomputers available to us, only Blue Waters and Stampede provide the necessary computing power. However, Stampede’s queuing policy prevents large-scale parallelism except in an emergency. An optimized scalable infrastructure cannot be set up unless there is access to large-scale parallelism ahead of the emergency.

PUBLICATIONS AND DATA SETS

Srinivasan, A., C.D. Sudheer, and S. Namilae, Optimizing Massively Parallel Simulations of Infection Spread through Air-Travel for Policy Analysis. *Proceedings of the 16th IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing (IEEE Computer Society, Cartagena, Colombia, May 16–19, 2016)*, DOI: 10.1109/CCGrid.2016.23.

Namilae, S., et al., Self-propelled pedestrian dynamics model: Application to passenger movement and infection propagation in airplanes. *Physica A*, 465 (2017), pp. 248–260.

Namilae, S., A. Srinivasan, C. Sudheer, A. Mubayi, R. Pahle, and M. Scotch, Self-Propelled Pedestrian Dynamics Model for Studying Infectious Disease Propagation during Air-Travel. *Journal of Transport & Health*, 3:2 (2016), S40.

Namilae, S., et al., Multiscale Model For Infection Dynamics During Air Travel. *Physical review E*, 95:5 (2017), DOI: 10.1103/PhysRevE.95.052320.

Derjany, P., S. Namilae, A. Mubayi, M. Scotch, and A. Srinivasan, Effect of pedestrian movement on infection transmission during air travel: A modeling study. *Transportation Research Forum Proceedings*, April (2017).

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

PI: Brad Sutton¹

Collaborators: Curtis Johnson², Alex Cerjanic¹, Aaron Anderson¹, Joseph Holtrop³

¹University of Illinois at Urbana-Champaign

²University of Delaware

³St. Jude Children's Research Hospital

EXECUTIVE SUMMARY

There is a large and growing body of evidence to support the interplay of mechanical properties of the human brain being affected by and affecting the function of the human brain. The ability to quantitatively assess the structural integrity of the brain using noninvasive techniques such as magnetic resonance imaging (MRI) may provide the necessary insight to understand the status and function of the living brain. Clinical MRI systems are capable of advanced quantitative imaging, both in high resolution and unique biomarkers, but they require significant technological development, tissue modeling, and computational resources. The Blue Waters environment has allowed for more rapid progress in all three areas through analysis of the computationally intensive inversion algorithms and validation of large experimental datasets. Two studies have been completed showing correlations between memory performance and hippocampal mechanical properties. Additionally, two ongoing studies look to characterize the structural–functional properties of the normal aging human brain.

RESEARCH CHALLENGE

The human brain changes throughout its lifetime, both in subtle and drastic ways. Subtle changes to the brain microstructure occur with lifestyle choices and normal aging. More drastic changes can come from traumatic injuries and neurodegenerative diseases. Quantitative MR imaging techniques provide the opportunity to interrogate the tissue microstructure for these changes, but traditional approaches require resolution on the order of the disruption of the microstructure, which is much smaller than

the achievable spatial resolution of MRI. Magnetic resonance elastography (MRE) is a quantitative imaging method that is able to accurately estimate the viscoelastic material properties of the brain tissue, from the small grey matter regions (e.g., hippocampus) to large white matter (WM) regions. MRE relies on advanced imaging methods and a nonlinear inversion (NLI) to estimate the material properties. NLI is technically and computationally challenging but has shown it is sensitive enough to detect subtle changes in human brain structures during disease and even during exercise interventions in both healthy and diseased brains. Each brain requires the computationally expensive step of NLI, resulting in the need for significant resources for analysis of a single brain. Blue Waters enables running full studies of 10–50 subjects to answer specific questions about the relationship between brain material properties and function.

METHODS & CODES

MRE is a noninvasive quantitative technique for measuring the mechanical properties via the response of the tissue to externally applied shear waves. The full 3D mechanical waves are captured using synchronized motion-encoding imaging, then the material properties are estimated using an iterative, finite-element-based nonlinear inversion (NLI) of the heterogeneous, steady-state Navier equation [1–2]. The size of the computational problem initially required dividing the brain into smaller regions (called “subzones”) for FEM meshing and optimization of properties at a local level, then reconstituting as a global solution. Blue Waters’ combination of a large number of nodes, high-speed interconnect, and large memory systems enabled us to perform a parameter study of a wide range of subzone sizes to determine the dependence of the estimated mechanical properties on these parameters.

RESULTS & IMPACT

Four separate human studies were completed, or are underway, and focus on two areas of brain health: understanding the links between tissue mechanical properties and brain performance [3–4] and effects of normal aging on the microstructure of the human brain. We have built on previous findings with MRE in studies run at the Beckman Institute exploring the relationship between tissue viscoelasticity and memory to include an analysis of fitness effects, which are known to improve brain health. The positive relationship between hippocampal structure, aerobic




Figure 2: Hippocampal viscoelasticity is improved by exercise training in persons with multiple sclerosis. Exercise therapy improves brain health and cognition, and the MRE measure is used to evaluate the recovery of hippocampal integrity following intervention. It can be used to better design and deliver rehabilitative treatments in the future.

fitness, and memory performance has been demonstrated at either end of the lifespan in children and older adults; evidence of this relationship among young adults, for whom the hippocampus is neither developing nor atrophying, is less consistent. Studies have typically relied on hippocampal volumetry (a gross proxy of tissue composition) to assess individual differences in hippocampal structure. Microstructural differences in hippocampal integrity may exist even among healthy young adults when volumetric differences are either absent or not diagnostic of tissue health or cognitive function.

In the current study, we investigated the relationship between hippocampal viscoelasticity from MRE and cardiovascular health, and their mutual effect on relational memory in a group of healthy young adults. We replicated our previous finding that hippocampal viscoelasticity correlates with relational memory performance. We extended this work by demonstrating better aerobic fitness, as measured by VO₂max, was associated with hippocampal viscoelasticity that mediated the benefits of fitness on memory function. In contrast, hippocampal volume did not account for individual differences in memory. Therefore, these data suggest hippocampal viscoelasticity may provide a more sensitive measure of microstructural tissue organization and its consequences on brain performance among healthy young adults.

We followed up on the findings of relationships among hippocampal viscoelasticity, memory performance, and fitness by using MRE in a study of exercise training in persons with multiple sclerosis (MS). Five subjects with MS completed three-month exercise training, and we observed this intervention significantly improved their hippocampal viscoelasticity. Further, we observed the change in viscoelasticity correlated with the improvement in memory performance with training. These results highlight the sensitivity of the MRE measures of brain health as they relate to cognitive performance. By using a highly sensitive neuroimaging battery that includes MRE to monitor therapy, we can better

understand the mechanisms contributing to brain health and recovery.

Finally, Blue Waters enabled a parameter study of NLI to further develop the NLI algorithms to improve the capabilities of quantitative imaging via a careful characterization of the effects of inversion parameters on the reconstructed material properties. Using the capabilities of Blue Waters’ network and local memory, the size of the subzones was increased up to 15 times the previous subzone size, previously impossible on other systems. Additionally, processor speed afforded an investigation into increasing the number of conjugate gradient estimations at the subzone level to ensure proper convergence of the material property estimates. This study was the first of its kind in the breadth and depth of inversion parameter characterization.

WHY BLUE WATERS

The Blue Waters system provides a unique resource for MRI because it allows for rapid transfer of imaging data from the scanners and the low-latency memory transfers necessary for highly parallel computations. This is especially critical for quantitative, high-resolution MRI that inherently involves very large datasets and extensive computation. Blue Waters has enabled the refinement of these techniques and their translation in research by significantly reducing computational burden.

PUBLICATIONS AND DATA SETS

Schwarb, H., et al., Aerobic fitness, hippocampal viscoelasticity, and relational memory performance. *NeuroImage*, 153 (2017), pp. 179–188, DOI: 10.1016/j.neuroimage.2017.03.061.

Sandhoff, B. M., C. L. Johnson, and R. W. Motl, Exercise training effects on memory and hippocampal viscoelasticity in multiple sclerosis: a novel application of magnetic resonance elastography. *Neuroradiology*, 59:1 (2017), pp. 61–67, DOI: 10.1007/s00234-016-1767-x.


Anderson, A.T., et al., Inversion Parameters based on Convergence and Error Metrics for Nonlinear Inversion MR Elastography. *25th Annual Meeting of the International Society for Magnetic Resonance in Medicine*, Honolulu, Hawaii, April 22–27, 2017.

Johnson, C.L., et al., Multiple Sclerosis Lesions are Softer than Surrounding White Matter: An MR Elastography Study. *25th Annual Meeting of the International Society for Magnetic Resonance in Medicine*, Honolulu, Hawaii, April 22–27, 2017.

H Schwarb, H., et al., Double Dissociation of Structure-Function Relationships Between Memory and Fluid Intelligence Using Magnetic Resonance Elastography, *24th Annual Meeting of the Cognitive Neuroscience Society*, San Francisco, Calif., March 25–28, 2017.

Anderson, A.T., et al., Study of Nonlinear Inversion Parameters for MR Elastography on Phantoms and Human Brain. *15th International Tissue Elasticity Conference*, Lake Morey, Vt., Oct 16–19, 2016.

Figure 1: Hippocampal viscoelasticity is correlated with memory performance in young adults. This measure of hippocampal tissue health, unique to high-resolution MRE, provides a novel method to evaluate the brain, as illustrated by its improved sensitivity over hippocampal volume, a classical brain-health measure.



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

Allocation: NSF PRAC/300 Knh
 PI: Ilias Tagkopoulos¹

¹University of California, Davis

EXECUTIVE SUMMARY

The democratization of mass sequencing and profiling technologies has resulted in a plethora of data that can reveal how organisms are organized and function on a cellular level. The goal of this project is to integrate the millions of data points collected for a model bacterium, *Escherichia coli*, so we can build predictive models of its behavior in novel, untested environments. We used Blue Waters for two major tasks. The first was to apply novel deep-learning algorithms for making sense of proteomics samples to find what proteins are present in given experimental settings. The second was to run massively parallel simulations with genome-scale, integrative models that predict the omics expression and bacterial behavior in novel environments. This work led to novel algorithms for omics data processing, multi-omics modeling, and the most integrative predictive model of bacterial behavior, with forward predictions validated experimentally. Parallelization

of these processes on the Blue Waters supercomputer enabled a precise data-processing pipeline, search of optimal model architecture, and hyper-parameter optimization.

RESEARCH CHALLENGE

One of the grand challenges of systems and synthetic biology is the development of an accurate predictive model of any organism. If we had such a tool at our disposal, we could interrogate the cellular states and molecular abundances that exist in general under a specific selection pressure or environmental setting. In the realm of synthetic biology, it would allow the in silico testing of synthetic circuits and the genetic engineering of the chassis cells so that they exhibit optimal behavior. However, there are many obstacles to achieving this vision, including the large complexity of the cellular organization and machinery, the lack of cohesive datasets, and models that are capable of integrating them into one system that is more than the sum of its parts.

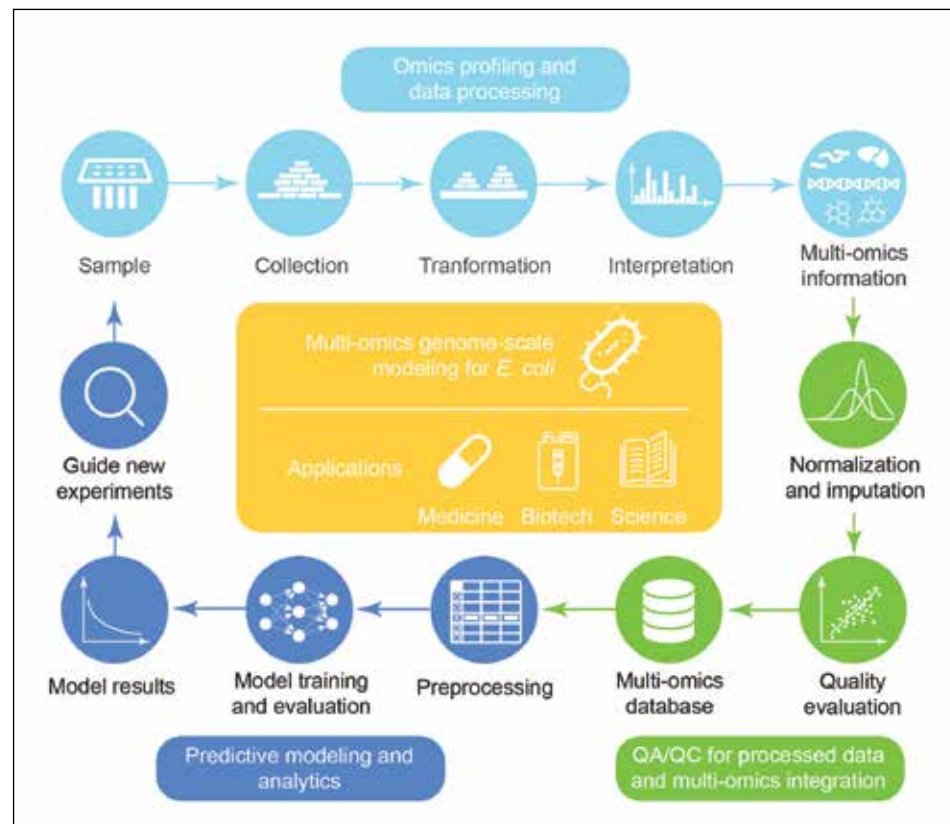


Figure 1: Overview of the processing pipeline and genome-scale modeling for *E. coli*. There are three major steps: 1) omics profiling and data processing, 2) QA/QC for processed data and multi-omics integration, and 3) predictive modeling and analysis.

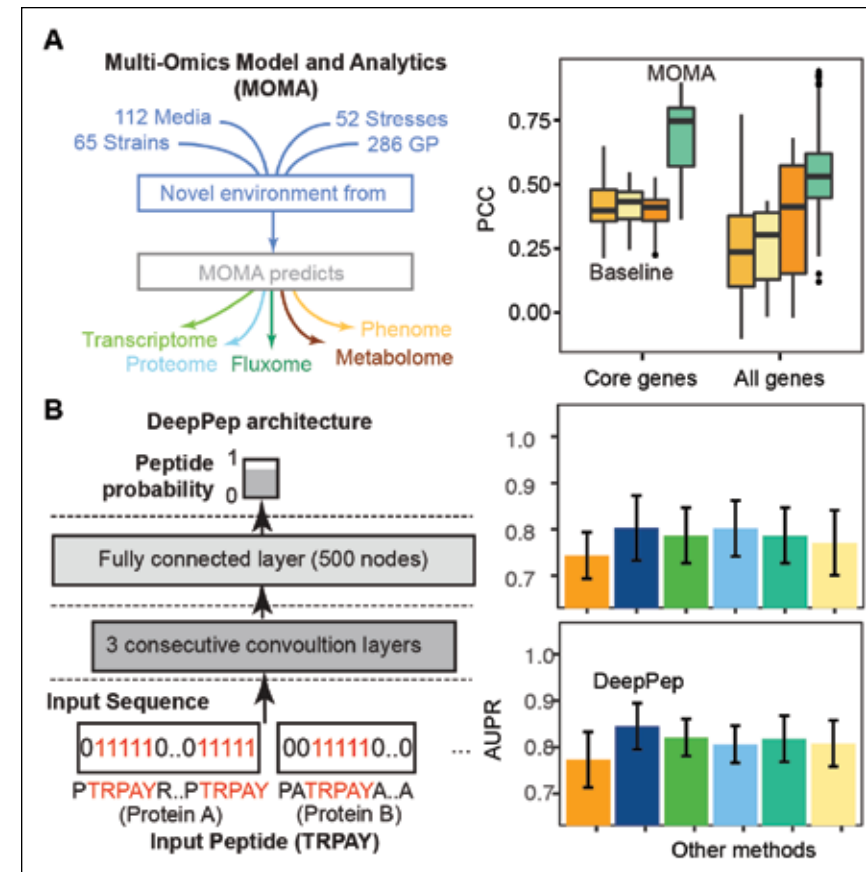


Figure 2: Performance evaluation of the proposed processing method (DeepPep) and the genome-scale model for *E. coli*. (A) DeepPep outperforms existing methods for protein inference from peptide-level proteome data. The multi-omics genome-scale model for *E. coli* shows predictive capacity in predicting molecular responses in transcriptome layer (B) and proteome layer (C) compared to multiple baseline methods.

METHODS & CODES

The proposed protein inference method (DeepPep) in the omics data-processing pipeline is based on convolutional neural networks with deep-learning techniques to efficiently train the complex model. DeepPep is written in torch/python, and the method was validated with the six external data sets (Sigma49 [1], UPS2 [2], 18Mix [3], Yeast [4], DME [5], HumanMD [6], and HumanEKC [7]). The genome-scale model employs a recurrent neural network and a constrained regression altogether to predict genome-wide responses layer by layer.

RESULTS & IMPACT

The predictive performance of the multi-scale model has shown substantial improvement over prior work as well as the ability to predict not only phenotypic information (growth, traits, etc.) but also genome-wide gene, protein, and metabolite expression. This brings us a step closer to having a crystal ball for prediction of bacterial states and behavior in novel environments. Having such a tool in our arsenal will allow the fast and inexpensive testing of experimental settings, which in turn will allow us to navigate the vast experimental space in search of suitable experimental environments of phenotypes. From biotechnology to medicine, such capability has several applications and can be transformational if applied at an industrial scale.

WHY BLUE WATERS

Our large-scale simulation for multi-scale modelling as well as the size and complexity of our datasets (30 million data points from several platforms and molecular species) necessitate the use of high-performance computing for this project. Furthermore, the effectiveness of GPU capabilities in Blue Waters for our application of deep neural network models both in multi-omics- and proteomics-specific projects allowed us to train large and complex models.

PUBLICATIONS AND DATA SETS

Kim, M., N. Rai, V. Zorraquino, and I. Tagkopoulos, Multi-omics integration accurately predicts cellular state in unexplored conditions for *Escherichia coli*. *Nat. Commun.* (2016), DOI: 10.1038/ncomms13090.

Kim, M., A. Eetemadi, and I. Tagkopoulos, DeepPep: deep proteome inference from peptide profiling. *PLoS Computational Biology*, accepted (2017).

<https://deeppep.github.io/DeepPep/>
<http://www.prokaryomics.com>

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

Allocation: NSF PRAC/1,200 Knh

PI: Gregory A. Voth¹

Co-PI: John M.A. Grime¹

Collaborators: Prabhuddha Sengupta^{2,3}, Antony K. Chen^{3,4}, Jennifer Lippincott-Schwartz^{3,4}, Anand Srivastava⁵, Mark Yeager^{6,7}, John A. G. Briggs^{8,9}

¹University of Chicago

²Howard Hughes Medical Institute

³National Institutes of Health

⁴Peking University

⁵Indian Institute of Science

⁶University of Virginia School of Medicine

⁷University of Virginia Health System

⁸Cambridge Biomedical Campus, UK

⁹European Molecular Biology Laboratory, Heidelberg

EXECUTIVE SUMMARY

A key step in the lifecycle of human immunodeficiency virus type-1 (HIV-1) is the production of enveloped particles containing viral proteins and genomes from infected cells. Experimental data suggest that the HIV-1 Gag polyprotein orchestrates this process through subdomains that specifically interact with the cell membrane, viral nucleic acid, or other Gag polyproteins. Nonetheless, the molecular details of this process have remained elusive. Using large-scale coarse-grained molecular simulations, enabled by the use of Blue Waters, we investigate a network of critical interactions that regulate the early stages of HIV-1 assembly, packaging, and budding (Fig. 1).

RESEARCH CHALLENGE

The proliferation of HIV-1 requires the aggregation of viral proteins and genetic material at the membrane of an infected cell, leading to the release of viral particles that spread the infection. Specifically, thousands of copies of the HIV-1 Gag polyprotein self-assemble into the so-called immature lattice at the cell membrane in the presence of viral RNA [1]. Disruption of this highly dynamical process is therefore a potential therapeutic target, and offers a blueprint for treating a range of viral infections. However, we lack a detailed molecular understanding of the factors that regulate this process due to limitations in conventional experimental techniques. Controlled study of the aggregation and interactions of large numbers of biomolecules at cell membranes is a challenging problem but can provide significant benefits for biomedical research and advance fundamental biophysical knowledge.




Figure 1: Molecular snapshot depicting early assembly of Gag polyprotein (blue, yellow, and grey tubes) as catalyzed by RNA (red chain) at the site of a deforming membrane puncta (translucent green sheet).

METHODS & CODES

To investigate the self-assembly of the immature protein lattice of HIV-1, we created large-scale coarse-grained (CG) molecular models of the relevant viral components and the cell membrane based on experimental data from collaborators [2,3]. CG models are computationally efficient representations, enabling computer simulations to be performed at time- and length-scales that are otherwise impossible. We performed molecular dynamics simulations using the LAMMPS software package (Sandia National Laboratories, USA) to examine the self-assembly of viral proteins, and the influence of the cell membrane and viral genome on the “budding” of viral particles. In collaboration with experimental colleagues from around the globe, we combined our simulation results with experimental data to better understand key aspects of the early stages of the HIV-1 viral lifecycle.

RESULTS & IMPACT

By combining computer simulations with fluorescence localization experiments [4], we elucidate the interactions that regulate HIV-1 viral assembly dynamics. Specifically, analysis of our results reveals the influence of nucleic acids and the cell membrane in promoting the aggregation of HIV-1 Gag polyprotein via a multi-step, self-correcting nucleation process. We also illustrate the functional importance of the N-terminal, C-terminal, and spacer peptide 1 (SP1) protein domains, which are each responsible for regulating different lattice qualities. These aspects are difficult or impossible to control precisely in conventional experimental approaches but are tractable in the context of our large-scale CG models. The success of our simulations, and the methodology we use to generate CG models of complicated biomolecules, suggest a simple and robust approach to the direct incorporation of experimental data into CG model generation. We envisage that our general methodology will be transferable to the study of other biomolecular processes and open up new avenues of research into phenomena that occur at physiologically relevant time and length scales. The insight gained by large-scale CG analysis can then be used to assist in the design of new therapeutic approaches to viral infection.

WHY BLUE WATERS

Despite the relative efficiency of CG models, our simulations require large numbers of individual molecules interacting over relatively long molecular time scales. Furthermore, in order to study the influence of viral RNA and the cell membrane on protein self-assembly, we must probe the behaviors of this system under a wide range of biologically relevant conditions. It was therefore crucial to the success of our project that we had access to very large amounts of computational power, and that power needed to be combined with superior network performance: the tightly coupled parallel nature of our simulations require latency and bandwidth characteristics that are simply not available in cloud computing. The Blue Waters computing platform thus presented

a natural choice for our work; the combination of leadership-class computing capabilities with cutting-edge network hardware allowed us to successfully investigate a system of significant biomedical interest. Our large-scale CG models are generated using techniques developed for previous work on the Blue Waters platform [5], and existing relationships with Blue Waters technical project staff greatly assisted in the deployment of novel simulations on Blue Waters.

PUBLICATIONS AND DATA SETS

Grime, J., et al., Coarse-grained simulation reveals key features of HIV-1 capsid self-assembly. *Nat. Commun.*, 7 (2016), p. 11568.

SOCIAL SCIENCE, ECONOMICS, & HUMANITIES

ECONOMICS

SOCIAL SCIENCE

252 *Policy Responses to Climate Change in a Dynamic Stochastic Economy*

254 *Enabling Redistricting Reform: A Computational Study of Zoning Optimization*

POLICY RESPONSES TO CLIMATE CHANGE IN A DYNAMIC STOCHASTIC ECONOMY

Allocation: GLCPC/250 Knh

PIs: Lars Hansen¹, Yongyang Cai²

Co-PIs: Kenneth Judd³, William Brock⁴, Thomas Hertel⁵

Collaborators: Simon Scheidegger⁶, T. J. Canann⁷, Carlos Rangel³

¹University of Chicago

²The Ohio State University

³Hoover Institution

⁴University of Wisconsin

⁵Purdue University

⁶University of Zurich

⁷University of Minnesota

EXECUTIVE SUMMARY

We extended our DSICE (Dynamic Stochastic Integration of Climate and the Economy) framework for evaluating policy responses to future climate change, to control the global average temperature anomaly so as not to exceed 2°C, with a stochastic economic production process [1]. We find that the social cost of carbon should be significantly greater in order to meet the 2°C target with a high probability [2].

We are extending DSICE to incorporate spatial structure of temperature and economy, sea level rise, thawing of permafrost, partial competition and collaboration, and adaptation. Our preliminary results show that rich countries in the region above 30 degrees north latitude should enact higher carbon taxes than poor countries in the tropic region [3]. Another extension of DSICE is to incorporate dirty/clean energy sectors. We find that we should significantly increase investment in clean energy and reduce investment in dirty energy, and very soon [4].

RESEARCH CHALLENGE

There are significant uncertainties in the climate and economic systems. Integrated Assessment Models (IAMs) of climate and the economy aim to analyze the impact and efficacies of policy responses to climate change. DICE (Dynamic Integrated model of Climate and Economy) is the benchmark IAM model developed over the past 20 years [5] and used frequently in the literature, e.g., by the United States Interagency Working Group on Social Cost of Carbon [6]. It is a simple perfect foresight forward-looking model that assumes we know all future information and that there is no heterogeneity. Ignoring heterogeneity in the economy as well as uncertainty in economic and climate conditions, or giving a relatively primitive treatment of uncertainty and heterogeneity, is often excused on the grounds that computational limitations make it impossible to do better. Our work from recent and previous Blue Waters projects clearly shows otherwise. We develop and solve new computational IAMs that merge the basic elements necessary, such as the spatial temperature system, climate tipping points, economic risks, clean energy usage, regional economic activities,

and so on. We then use their solutions to do economic analysis about the optimal climate policy under uncertainty and risks and how such a policy will impact economic activities.

METHODS & CODES

In the past three years, Cai, Judd, and Lontzek used GLCPC allocations to build a prototype computational framework, called DSICE [1], to address these issues. DSICE allows shocks to both the economic and climate systems, and uses recursive preferences that isolate the intertemporal elasticity of substitution and risk aversion. Cai, Brock, and Xepapadeas are working on a spatial DSICE model [3], based on DSICE and the energy balance model developed by Brock and his collaborators. The spatial DSICE model incorporates heat and moisture transport between poles and tropics, sea level rise, and regional economic entities with interaction, as well as climate adaptation. Cai and his collaborators are also working on clean energy usage and Bayesian learning of uncertain critical parameters in DSICE.

We have developed three general parallel code packages for stochastic dynamic programming (DP), nonlinear certainty equivalent approximation (NLCEQ), and supergames, respectively, with high parallel efficiency. For stochastic DP models, we use the value/policy function iteration method (backward iteration) and the master-worker structure; the master assigns N tasks for workers to solve in parallel, and then gathers the results of these tasks from workers, and repeats this process until our iteration stops. For NLCEQ, it can be implemented in parallel naturally, and we have used Blue Waters allocations to solve huge-dimensional (up to 400 state variables) dynamic stochastic problems [7]. While tasks in the DP method are of small or moderate size, the tasks in NLCEQ are solving large-scale optimization problems that could take hours for each solution. Moreover, the number of tasks could also be large, but they are naturally parallelizable, so parallelization over these tasks is efficient in Blue Waters. The algorithm for solving supergames uses iteration methods until convergence, and each iteration consists of solving an enormous number of linear programming problems (of moderate size) [8].

RESULTS & IMPACT

In the past year, we published three papers based on the support of Blue Waters. In the first paper, published in *Nature Climate Change*, we extended DSICE to study the impact of multiple interacting tipping points. It is a large-size stochastic dynamic programming model with up to 10 continuous-state variables and five discrete-state variables. The paper showed that the five multiple interacting tipping points increases the social cost of carbon nearly eightfold.

In the second paper, published in *Quantitative Economics*, we used Blue Waters to develop an NLCEQ method to solve huge-dimensional dynamic stochastic problems efficiently and in parallel. NLCEQ has also been implemented to analyze the effect of climate and technological uncertainty in crop yields on the optimal path of global land use by Cai, Judd, Hertel and their collaborators [7].

In the third paper, published in *Operations Research*, we used Blue Waters to develop a parallel algorithm that can solve super games with states, which model strategic interactions among multiple players.

Economics, particularly when integrated with climate change science, is a field with many problems that have the size and complexity that justify the use of massively parallel computer systems. Our work using Blue Waters has allowed some economists to attack those computationally intensive economics problems.

WHY BLUE WATERS

Our parallel algorithms use the master-worker structure, and the communications among the master and workers are frequent and of small or moderate sizes, so the high latency inherent in a commodity cluster limited the ability to solve large problems in a reasonable amount of time. Moreover, the number of tasks could be huge. For example, the largest DP problem (for DSICE) we solved using Blue Waters had 372 billion of such tasks. Blue Waters allows us to use MPI (Message-Passing Interface) and solve far larger problems efficiently, as we have already shown in our previous work using Blue Waters.

PUBLICATIONS AND DATA SETS

Cai, Y., T.M. Lenton, and T.S. Lontzek, Risk of multiple climate tipping points should trigger a rapid reduction in CO₂ emissions. *Nature Climate Change* 6 (2016), pp. 520–525.

Cai, Y., K.L. Judd, and J. Steinbuks, A nonlinear certainty equivalent approximation method for stochastic dynamic problems. *Quantitative Economics*, 8:1 (2017), pp. 117–147.

Yeltekin, S., Y. Cai, and K. L. Judd, Computing equilibria of dynamic games. *Operations Research*, 65:2 (2017), pp. 337–356.

ENABLING REDISTRICTING REFORM: A COMPUTATIONAL STUDY OF ZONING OPTIMIZATION

Allocation: Illinois/400 Knh

PI: Wendy K. Tam Cho¹

Co-PI: Yan Liu¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

We have developed a scalable computational tool for redistricting that synthesizes and organizes massive amounts of computation and data to evaluate redistricting schemes. The tool allows one to create high-quality maps and tailor them to notions of “fairness” and democratic rule. It can also be used as an evaluation tool by courts, advocates, and the public to ensure nondiscriminatory representation. Specifically, we developed a scalable, parallel, evolutionary algorithm for redistricting that includes a set of spatial evolutionary algorithm operators to handle the costly spatial configuration of redistricting maps. These maps provide the basis for additional statistical analyses.

RESEARCH CHALLENGE

In the United States, political redistricting occurs at the national level every 10 years following the decennial census. It is intended to provide fair representation in Congress to all communities and interest groups. Gerrymandering occurs when districts are drawn in a manner that discriminates against a partisan or racial group. Both partisan and racial gerrymanders are commonly alleged. Despite broad disdain for the practice of gerrymandering, the Supreme Court has found it difficult to identify a workable standard by which we might regulate gerrymandering. We lack sufficient tools to analyze and synthesize redistricting data, in part, because the requisite computation is massive. Without the tools to quantify the effect of electoral maps, the court is left without the ability to issue legal and consistent judgments. As a result, despite the five decades since the Supreme Court declared gerrymandering



Figure 1: The algorithm is able to produce billions of high-quality, viable redistricting maps that may be used by the Courts in adjudicating partisan gerrymandering cases. Image credit: Carlos Waters/Vox.

to be capable of being decided by legal principles or by a court, the court has yet to identify manageable standards under which one could measure and thus declare a partisan gerrymander. The failure of the legal system in this political realm has significant ramifications for our democratic system of governance.

Using the Supreme Court’s articulated legal reasoning and mandates, we have developed a computational redistricting tool. The redistricting problem can be formulated as a combinatorial optimization problem, with objectives and constraints defined to meet legal requirements. Drawing electoral maps amounts to arranging a finite number of indivisible geographic units of a region into a small number of districts. Since every unit must belong to exactly one district, a districting map is a partition of the set of all units into a preestablished number of nonempty districts. The redistricting problem is an application of the set-partitioning problem that is known to be NP-complete and, thus, the time required to solve the problem increases very quickly as the size of the problem grows. We have developed a scalable, evolutionary computational approach utilizing massively parallel high-performance computing for redistricting optimization and analysis at fine levels of granularity.

METHODS & CODES

Our algorithm, PEAR, or Parallel Evolutionary Algorithm for Redistricting, is implemented in ANSI C. It can be compiled on Linux, OS X, and Windows as a standard *makefile* project. PEAR uses MPI nonblocking functions for asynchronous migration for load balancing and efficiency. It uses the C SPRNG 2.0 library to provide a unique random number sequence for each MPI process, which is necessary for running a large number of evolutionary algorithm (EA) iterations.

Since the spatial configuration plays a critical role in the effectiveness and numerical efficiency of redistricting algorithms, we designed spatial EA operators that incorporate spatial characteristics to effectively search the solution space. Our parallelization of the algorithm further harnesses massive parallel computing power via the coupling of EA search processes and a highly scalable message-passing model that maximizes the overlapping of computing and communication at runtime.

RESULTS & IMPACT

Our approach is designed to identify redistricting maps that satisfy a set of user-defined criteria with a particular focus on

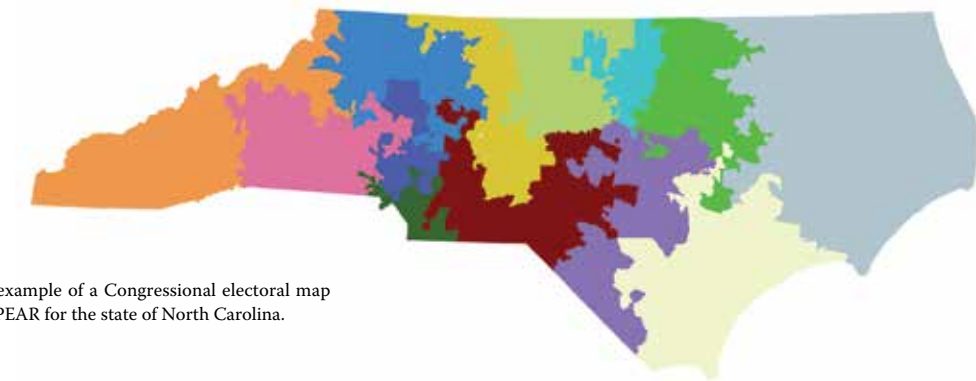


Figure 2: An example of a Congressional electoral map generated by PEAR for the state of North Carolina.

addressing fine levels of spatial granularity. We leveraged and enhanced a scalable Parallel Genetic Algorithm (PGA) library to develop PEAR for the computationally intensive redistricting problem. By incorporating a set of spatial configuration operators and spatial EA operators to handle spatial characteristics and the associated computational challenges, and harnessing massive computing power, PEAR provides a powerful and computationally scalable redistricting tool that has never existed before. In the substantive realm, our computational model 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.

The discrete optimization framework identifies large sets of quality electoral maps. This set of plans is useful at the districting drawing stage as well as for judges who are adjudicating the constitutionality of a redistricting plan. At the drawing stage, this information suitably drives an iterative bargaining process whether that process involves partisan legislators or members of an independent redistricting commission. When this bounded set of plans can be identified, the redistricting process is imbued with valuable structure that is otherwise nonexistent. That structure alone makes the redistricting process more transparent and may serve to reduce legal challenges. When a legal challenge is mounted, this large set of reasonably imperfect plans produces the relevant background and allows one to place and understand the proposed plan in context. The purpose of our optimization strategy, then, is to search, synthesize, and organize massive amounts of information that will supply a common base of knowledge to guide informed and intelligent debate and decisions.

The impact of our project is yet to be determined, but the project has thus far been well received. Our proposal for a new partisan gerrymandering standard, based on our computational tool, won the first place prize in Common Cause’s “Gerrymander Standard” writing competition that was judged by law school deans, law professors, and lawyers. The project has also received some media coverage.

WHY BLUE WATERS

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

PUBLICATIONS AND DATA SETS

Cho, W. K. T., and Y. Y. Liu, Toward a Talismanic Redistricting Tool: A Computational Method for Identifying Extreme Redistricting Plans. *Election Law Journal* 15:4 (2016), pp. 351–366.

Cain, B. E., W.K. T. Cho, Y. Y. Liu, and E. Zhang, A Reasonable Bias Method for Redistricting: A New Tool for an Old Problem. *William & Mary Law Review*, 59 (2017).

Cho, W. K. T., and Y. Y. Liu, A Scalable Evolutionary Algorithm with Intensification and Diversification Protocols Designed for Statistical Models. *SC16: The International Conference for High Performance Computing, Networking, Storage and Analysis*, Salt Lake City, Utah, November 13–18, 2016 (extended abstract).

Liu, Y. Y., W. K. T. Cho, and S. Wang, PEAR: A Massively Parallel Evolutionary Computation Approach for Political Redistricting Optimization and Analysis. *Swarm and Evolutionary Computation*, 30 (2016), pp. 78–92.

Cho, W. K. T., and Y. Y. Liu, A High-Performance Approach for Solution Space Traversal in Combinatorial Optimization. *SC15: The International Conference for High Performance Computing, Networking, Storage and Analysis*, Austin, Texas, November 16–19, 2015 (extended abstract).

Liu, Y. Y., W. K. T. Cho, and S. Wang, A Scalable Computational Approach to Political Redistricting Optimization. in *Proceedings of the 2015 Annual Conference on Extreme Science and Engineering Discovery Environment*. (XSEDE’15: Scientific Advancements Enabled by Enhanced Cyberinfrastructure, St. Louis, Mo., July 26–30, 2015), DOI: 10.1145/2792745.2792751.

GRADUATE FELLOWS

- 258** *Resolving Plant Functional Biodiversity to Quantify Forest Drought Resistance in the Amazon*
- 260** *Modeling Nonlinear Physical-Biological Interactions: Eddies and Sargassum in the North Atlantic*
- 262** *The Role of Cosmic Rays in Isolated Disk Galaxies*
- 264** *Analyzing the Propagation of Soft Error Corruption in HPC Applications*
- 266** *Toward Developing a Thermodynamic Model of Binding-Induced Conformational Transitions in Short, Disordered Protein Regions*
- 268** *Genomic Perspectives on the Amphibian Tree of Life*
- 270** *Multiscale Simulations of Complex Fluid Rheology*
- 272** *High Accuracy Radiative Transfer in Cloudy Atmospheres*
- 274** *High-Fidelity Blade-Resolved Wind Farm Simulations*
- 276** *Reducing the Computational Cost of Coupled Clustery Theory*
- 278** *Simulated Effects of Urban Environments on the Dynamics of a Supercell Thunderstorm*
- 280** *Monte Carlo Neutrino Closures in 3D GRMHD Simulations of Core-Collapse Supernovae and Neutron Star Mergers*
- 282** *Understanding the Role of Hydrodynamic Fluctuations in Biomacromolecular Dynamics through the Development of Hybrid Atomistic-Continuum Simulation*
- 284** *The Impacts of Hydrometeor Centrifuging on Tornado Dynamics*
- 286** *Glassy Dynamics and Identity Crises in Hard-Particle Systems*
- 288** *Magnetic Reconnection in Laser-Driven Plasmas: from Astrophysics to the Laboratory in silico*

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

Elizabeth Agee, University of Michigan
2016–2017 Graduate Fellow

EXECUTIVE SUMMARY

Cases of heat- and drought-induced mortality have been documented in every biome of the world, indicating that changes in global temperatures and precipitation patterns are pushing the world's forests beyond current thresholds of stress resilience. The increased frequency and severity of droughts and their regional consequences have highlighted the potential vulnerability of the Amazon Basin region to heat- and drought-induced stress. To adequately capture the response of tropical rainforests to water limitation, mechanistic models that incorporate three-dimensional plant morphology and traits are needed. Three-dimensional root water uptake is modeled for a forest plot in the Tapajós National Forest for the 2015–2016 El Niño drought. Results confirm the model's ability to capture differential response of individual species to standard diurnal cycles and water limitation. Using Blue Waters, future work will model whole-forest response to water limitation and the contribution of root systems to drought resilience.

RESEARCH CHALLENGE

Cases of heat- and drought-induced mortality have been documented in every biome of the world, indicating that changes in global temperatures and precipitation patterns are pushing the world's forests beyond current thresholds of stress resilience [1]. The Amazon Basin region, home to the world's largest area of undisturbed tropical biomass, is critical to global energy, water, and carbon cycles. Over the past two decades, the region has been hit with multiple drought events triggered by strong shifts in sea surface temperature cause by the El Niño–Southern Oscillation (ENSO). The increased frequency and severity of droughts and their regional consequences have highlighted the potential vulnerability of the Amazon to heat- and drought-induced stress [2]. To adequately capture the response of tropical rainforests to water limitation, mechanistic models that incorporate diverse plant morphology and hydraulic function are needed.

METHODS & CODES

The computational complexity of single plant models has previously limited incorporation into hydrological models at the forest plot or ecosystem scale, but recent developments in microscale hybridization of root hydraulic architecture has opened the door to coupled models of three-dimensional root water uptake and soil water physics. Root architectures that represent the structural and spatial distribution of roots were modeled using the open source RootBox model [3]. Each tree system is assigned hydraulic parameterization (e.g., root hydraulic conductivity, water potential thresholds) based on statistically generated water-usage strategies. These strategies may range from risky, which favor carbon assimilation over hydraulic integrity, to conservative, which will limit carbon assimilation and, therefore, water uptake to protect hydraulic pathways from damage. Root water uptake has been coupled with the massively parallel flow and transport model, PFLOTRAN [4], using hybridization techniques from [5].

Using these tools, we are exploring how tree roots contribute to forest drought resilience in areas of the Amazon rainforest during the recent 2015–2016 El Niño drought event. To tease apart the contributions of various ecophysiological properties, ensemble modeling approaches are employed that test a multitude of risk configurations and root distributions. Each of these approaches uses spatial distributions from and is validated with data collected from our field site in the Tapajós National Forest (K67), located in the eastern Amazon River Basin.

RESULTS & IMPACT

Preliminary simulations focused on canopy-dominant species. Representing the largest individuals in the forest, canopy trees

have the greatest water demand and are thus susceptible to stress or damage during water limitation. Initial results confirm the model's capacity to model water uptake for individual trees at the forest scale. Individuals show differential response to water limitation based on both their size and physical traits such as tissue conductivity and stomatal response. Furthermore, intrinsic properties are not the only driver of an individual tree's response. Social position, the relative location of a tree to its peers, is an important factor in determining where in the soil water column uptake occurs. As soils continue to dry, uptake will be shifted away from the center of mass and into deeper soil layers.

Future simulations will include all canopy layers—upwards of 1500 trees per hectare—comprising the largest simulations of three-dimensional root water uptake ever attempted. It is expected that the additional root density and root-zone overlap will further tune individual response as competition increases in soil-water reservoirs. Individual and community drought resilience will be assessed, elucidating the contributions of root systems to forest drought resilience in the Amazon rain forest.

WHY BLUE WATERS

Blue Waters is critical to the ongoing success of this project. Simulations of this complexity and scale require the computational power of this system to make meaningful analyses. Not only are the simulation domains complex, multiple simulations are needed to account for system uncertainty. The enormous biodiversity and harsh environmental conditions of tropical forests hinder data collection needed for model parameterization. Scalable, physically based models provide a necessary tool with which to explore modes of uncertainty and help target data collection efforts.

Elizabeth Agee is a fourth-year Ph.D. student in environmental engineering at the University of Michigan. She is working under the supervision of Valeriy Ivanov, and hopes to graduate in April 2019.




Figure 1: Root water uptake was modeled for individuals within the canopy. This representation shows the spatial distribution of trees and their relative heights. Grey circles indicate the lateral spread of root systems. Root systems overlap (inset), pulling water from shared soil water reserves.

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

Maureen T. Brooks, University of Maryland Center for Environmental Science
2016-2017 Graduate Fellow

EXECUTIVE SUMMARY

The macroalgae *Sargassum*, commonly known as “gulf weed,” inhabits the Atlantic Ocean and poses natural resource management challenges for coastal communities in its range. This study seeks to highlight interactions among *Sargassum* and mesoscale eddies and fronts to better predict *Sargassum* dispersal and growth. To this end, I developed a coupled model system that spans scales from individual organisms up through

the basinwide circulation of the Atlantic. The resources of Blue Waters facilitated this model development and allowed for implementation at high resolution. Model results suggest that the Gulf of Mexico (GoM) and Western Tropical Atlantic play a key role in determining *Sargassum* distribution, and highlight the need for better understanding of the reproductive strategy of this organism. The eddy field in the GoM in particular appears to both influence *Sargassum* dispersal into the greater Atlantic and alter the growth conditions experienced by *Sargassum* in the region.

RESEARCH CHALLENGE

Pelagic *Sargassum* is comprised of two species: *Sargassum fluitans* and *Sargassum natans*. These are the only species of macroalgae in the world that spend their entire life cycle floating on the ocean surface. They serve as keystone species in the Sargasso Sea and throughout their range, supporting a thriving ecosystem, from invertebrates to commercial fish, in low-nutrient “ocean desert” regions [1]. Recently, increased reports of *Sargassum* wash-ups have highlighted its negative impacts. These events are associated with lost fishing and tourism revenue, as well as large clean-up costs for coastal communities on both sides of the Atlantic [2].

Sargassum has many air bladders which provide buoyancy, and this causes colonies to accumulate along convergent eddies and fronts. These features not only affect the regional surface flow, but also can induce vertical velocities that can alter the nutrient availability in the surface waters where *Sargassum* grows. Mesoscale eddies can persist at monthly time scales, long enough to impact *Sargassum* biomass based on measured growth rates [3]. This study examines the physical and physiological impacts of mesoscale features on pelagic *Sargassum*.

METHODS & CODES

This research uses a system of four coupled models to simulate *Sargassum* growth and interactions with ocean circulation features. A Hybrid Coordinate Ocean Model (HYCOM) [4] domain was implemented at $1/12^\circ$ ($< 10\text{m}$) resolution over a domain that encompasses the known *Sargassum* distribution, from 15°S to 64°N and 100°W to 15°E . This model has 28 hybrid vertical layers which capture the 3-D ocean circulation and their resolution is concentrated in the upper 200m to better simulate vertical velocities associated with surface eddies. Coupled to the HYCOM model is a biogeochemical model adapted from the work of Fennel [5], which includes nitrogen and phosphorus, as well as phytoplankton, zooplankton, and detritus to effectively

capture the dynamics of biologically-mediated nutrient cycling in the upper ocean.

Sargassum colonies are modeled using a combination of Lagrangian particle and individual-based physiology models. I have modified the HYCOM Lagrangian particle code to allow for particle buoyancy to better simulate *Sargassum*. The particle code has also been improved to allow both forward- and backward-time integration as well as inertial effects, and interpolates the physical and biogeochemical conditions from the first two models along each particle trajectory. The *Sargassum* physiology model was developed for this study and is run within every individual Lagrangian particle. It uses light, temperature, and nutrient availability to determine growth rate, and because each colony is tracked it can also account for age and reproductive strategy. Vegetative propagation is simulated by initializing a new *Sargassum* propagule in place when a colony dies to age-related causes in regions where conditions are otherwise favorable. The *Sargassum* biomass distributions generated by the 4-model suite were validated against monthly satellite climatologies derived from observations over a 10-year period [6].

RESULTS & IMPACT

This multi-scale modeling effort gives an unprecedented view of basin-wide *Sargassum* biomass and mesoscale biological-physical interactions. Particle seeding and Lagrangian trajectory analyses examined a total of 17 sub-regions of the model domain, with an average area of $1.3 \times 10^6 \text{ km}^2$, and found evidence for two potential “seed” regions that exert disproportionate influence on the *Sargassum* seasonal cycle in the Atlantic. When *Sargassum* particles are seeded in the Gulf of Mexico and the Western Tropical Atlantic near the mouth of the Amazon River the

seasonal distribution of biomass has a 30% reduction in RMS error as measured against the satellite observations. The inclusion of *Sargassum*’s vegetative reproductive strategy also improved model accuracy, reducing model bias to within 1.5% of the mean observed biomass.

Lagrangian Coherent Structure (LCS) analysis was used to accurately determine the boundaries of mesoscale eddies and fronts via the finite-time Lyapunov exponent field (Fig. 1a). *Sargassum* particles aggregate along attracting LCS as expected, and biomass is influenced by these structures (Fig. 1b). Because the colonies tend to stay at eddy boundaries, they are spared the low-nutrient conditions in the interior of the large, convergent eddies that pinch off from the Loop Current. Buoyant particles tend not to cross LCS lines, and the differential eddy activity in the western and eastern Gulf of Mexico also helps maintain the *Sargassum* population in the western GoM where the potential for growth is high and biomass can accumulate. This improved understanding should help us better predict the potential for *Sargassum* growth based on local oceanographic conditions.

WHY BLUE WATERS

The resources of Blue Waters have made the scale and scope of this project possible. High-resolution ocean circulation modeling alone has a high computational cost. By utilizing Blue Waters, I accomplished the high-resolution ocean circulation modeling as well as coupling it with ocean biogeochemistry, Lagrangian particles, and individual organism physiology at temporal and spatial scales that span orders of magnitude. The professionalism of the NCSA staff has also been key to the success of this project. Their responsiveness and expertise made implementing and running this code on Blue Waters as straightforward as possible.

Maureen T. Brooks is in the fourth year of her Marine-Estuarine-Environmental Sciences Ph.D. program, working under the direction of Victoria Coles at the University of Maryland Center for Environmental Science Horn Point Laboratory. She plans to complete her degree in 2018.




Figure 1: (a) Attracting Lagrangian coherent structure field for the Gulf of Mexico for November 6, 1993, as calculated by finite-time Lyapunov exponent from a 28-day backwards-time particle integration. (b) Normalized biomass of *Sargassum* particles, overlaid on the LCS field for the same date. Note higher biomass in the western GoM.

THE ROLE OF COSMIC RAYS IN ISOLATED DISK GALAXIES

Iryna Butsky, University of Washington
2016-2017 Graduate Fellow

EXECUTIVE SUMMARY

Cosmic ray feedback in cosmological simulations is integral to creating robust models of stellar feedback and to reproducing galactic structure that is consistent with observations. Unlike thermal gas, relativistic cosmic rays retain more pressure under adiabatic expansion and cool on significantly longer time scales. Cosmic rays stream along magnetic field lines and drive instabilities that initiate strong outflows. The goal of this project is to use simulations of isolated disk galaxies for which cosmic rays are dynamically important to study how cosmic ray-driven outflows shape the structure and kinematics of the circumgalactic medium. In order to model realistic cosmic ray behavior, it is necessary to employ a fully anisotropic, magnetohydrodynamic (MHD) treatment, which I have implemented in Enzo, a publicly available cosmological simulation code.

RESEARCH CHALLENGE

The circumgalactic medium (CGM) is a dynamically complex and diffuse multiphase gas that extends to the outer edges of galactic halos and contains the majority of a galaxy's baryons. The interplay in the CGM between outflows from the star-forming disk and inflows from the intergalactic medium provides constraints to theories of galaxy formation and evolution. Recent observations, such as those from the COS-Halos survey on the Hubble Space Telescope (HST), show that the CGM of distant galaxies contains a substantial amount of metal-enriched, cool gas [1]. Cooling times of such a gas are very short compared to galactic timescales, and it is unclear how this material survives in such abundance. The data seem to imply an additional unknown source of nonthermal pressure—for which cosmic rays are a candidate—that supports the cool gas against evaporation. Furthermore, although metals are produced within galactic disks, galaxies retain only 20–25% of these metals in their stars and interstellar medium [2]. Data from the Sloan Digital Sky Survey suggest that metals have been lost to outflows [3], which strongly affect the mass structure and kinematics of the CGM.

Recent studies have shown that including cosmic ray stellar feedback can drive stronger and more mass-loaded galactic winds than thermal stellar feedback alone [4]. Cosmic rays are relativistic charged particles accelerated by shocks (such as supernovae). Unlike thermal gas, cosmic rays retain pressure under adiabatic expansion better and cool on significantly longer time scales. Because cosmic rays are charged particles, they do not move freely through space but propagate along magnetic field lines. This streaming drives instabilities in the magnetic field lines, which contribute to the exponential growth of the galactic magnetic field and drives strong outflows. Because magnetic fields vary greatly in strength and shape between the galactic disk and the CGM, implementing anisotropic cosmic ray transport is crucial to improved accuracy and predictive power within simulations. Therefore, the goal of this project is to employ realistic cosmic ray physics in simulations of isolated disk galaxies in order to study galactic outflows and their effects on the CGM structure and composition.

METHODS & CODES

The challenge of this project was in implementing a new cosmic ray fluid that is compatible with the Riemann solvers in Enzo, a multi-physics astrophysical simulation code. In addition to cosmic ray advection, I implemented several different modes of cosmic ray propagation along magnetic field lines, including anisotropic diffusion, streaming, and gas heating. The new cosmic ray physics will soon be publicly available to all Enzo users.

After extensive testing, I've begun running a suite of isolated disk galaxies following initial conditions prescribed by the AGORA [5] project. In these galaxy models, supernovae inject thermal, magnetic, and cosmic ray energy into the surrounding gas. I will systematically compare models of varying resolution and cosmic ray transport to the control model with no cosmic rays to isolate the effects of the different propagation modes.

RESULTS & IMPACT

Initial results show evidence of strong, mass-loaded outflows that enrich the CGM out to the virial radius of the halo when cosmic rays are present. Fig. 1 shows an edge-on view of the metal-enriched CGM after only 1.5 Gyr (gigayears). The contours of the magnetic field lines in black show the turbulent nature of the outflowing gas. Follow-up work will focus on comparing the simulated CGM structures to existing COS-Halos observations.

I will use the new Trident tool to create synthetic spectra that are tuned to the specifications of the COS-Halos instrument. Directly comparing the simulations to existing data is a state-of-the-art approach that will place better constraints on theories pertaining to the CGM and make predictions for the structure and metallicity distribution of the CGM for future COS-Halos observations.

WHY BLUE WATERS

The universal challenge faced by galaxy simulations is reconciling the great dynamic range in physical and temporal scales. Although a galaxy's CGM scales hundreds of kiloparsecs (1kpc = 3.09e16 km), its evolution is dictated by stellar feedback, which happens on subparsec scales. Even with adaptive mesh refinement, this is a formidable task, and galaxy simulations have yet to resolve the scale of individual stars. Furthermore, each cell in the simulation contains many different fluid properties and follows a complicated set of interaction rules in order to capture the equally complicated physics governing galaxy evolution. Therefore, galaxy simulations require the use of massively parallel, high-performance supercomputers such as Blue Waters. In addition to Blue Waters' computational resources, I have benefited greatly from the support of its staff, who are admirably dedicated to resolving issues in a timely manner.

Iryna Butsky is in the second year of a Ph.D. program in astronomy at the University of Washington (UW). Her advisor is Tom Quinn, also of UW; she plans to graduate in June 2021.




Figure 1: An edge-on view of outflows from an isolated disk galaxy after 1.5 Gyr. The color depicts the metallicity (metal enrichment relative to solar abundances) of the gas, and the streamlines follow the topology of the magnetic field. The dimensions of the image are 90 kpc x 150 kpc.

ANALYZING THE PROPAGATION OF SOFT ERROR CORRUPTION IN HPC APPLICATIONS

Jon Calhoun, University of Illinois at Urbana-Champaign
2016-2017 Graduate Fellow

EXECUTIVE SUMMARY

Because the rate of radiation-induced soft errors impacting application data are expected to increase on future high-performance computing (HPC) systems, analyzing how corruption due to soft errors propagates inside applications becomes critical for developing efficient detection and recovery schemes. Quantifying the latency (the number of instructions or iterations) of common symptoms of soft errors—e.g., crash or detection—allows for measuring the effectiveness of soft error detectors at containing corruption. Containing corruption allows for low-cost localized recovery instead of the high-cost global roll-back recovery checkpoint-restart. To analyze corruption propagation in HPC applications, an LLVM-based compiler tool instruments an application, allowing for tracking of corruption at an instruction and application variable level. Results show that the latency of common symptoms of soft errors along with speed of propagation to other processes varies dramatically. Blue Waters provided us with a platform to run thousands of fault-injection experiments and analyze hundreds of gigabytes of data efficiently.

RESEARCH CHALLENGE

Design constraints, such as cost of procurement and power budget, may make future HPC systems more susceptible to transient radiation-induced soft errors [1]. Soft errors commonly result in bit inversion and are the reason that all HPC main memories uses protection mechanisms such as error-correcting codes (ECC). On future systems, full protection from the effects of soft errors will be prohibitively expensive. Thus, HPC applications running for long durations and at large scales may experience data

corruption during execution [2]. Furthermore, global rollback-recovery becomes increasingly expensive at larger scales. Analyzing the effectiveness of current software-based detection schemes at containing data corruption is central in facilitating low-cost localized recovery and ensuring correct simulation results.

METHODS & CODES

To track corruption propagation inside HPC applications, we constructed an LLVM compiler pass to create an executable file that emulates lockstep execution between a nonfaulty version of the application called *Gold* and a version that experiences a single bit-flip fault injection called *Faulty*. *Gold* produces a reference set of loads, stores, and correct program behavior that *Faulty* is judged against. The compiler pass completes the following steps when instrumenting the code: (1) it duplicates compiled code, forming two applications—*Gold* and *Faulty*, (2) it interleaves instructions from *Gold* and *Faulty*, emulating lockstep execution, (3) it instruments all loads and stores in *Faulty* with a function call to log deviations, (4) it instruments *Faulty* branches to check for control flow divergence, (5) it instruments faulty code for fault injection with FlipIt [3], and (6) it logs deviation in application-level variables at the end of iterations. This tool is used to investigate the latency (in number of LLVM instructions executed) of common symptoms of soft errors, the speed and extent of propagation between variables and parallel processes, and the impact of local problem size and compiler optimizations on propagation. The applications we test include: Jacobi, CoMD, and HPCCG. The latter two are taken from the Mantevo mini-app collection [4]. Each of these applications is modified with a lightweight soft-error detection scheme to investigate the latency of soft-error detection.

RESULTS & IMPACT

This project tracks corruption propagation at both the instruction level and the application variable level. Investigation at the instruction level shows the latency (in number of LLVM instructions) of common symptoms of soft errors: crash, control flow divergence, and soft-error detection. Application crashes have the shortest latency of the symptoms considered. Results show that 90% of crashes occur within four instructions. During those four instructions, there is little chance of corruption propagation. The latency of control-flow divergence and soft-error detection are noticeably larger. The larger latency allows for corruption to propagate to other variables and processes. To measure the extent and magnitude of corruption propagation in variables and processes, we track propagation at the application variable level.

Tracking propagation in application variables highlights dependencies between variables and processes. The application HPCCG uses the conjugate gradient linear solver to solve the sparse system of $Ax=b$. HPCCG is run 1,500 times with a single bit-flip fault injected on process rank 3 during every run. Fig. 1 shows the average percentage of elements deviated by more than $1e-10$ (color) in the solution vector, x , for the application HPCCG across all MPI processes (y-axis) for subsequent iterations after injection (x-axis). As corruption in the solution vector propagates locally, the horizontal color for that row grows darker. As corruption is removed, the color lightens. Propagation between processes can be seen by looking at the color progression of columns at each iteration. Due to an `MPI_Allreduce` in the algorithm, corruption is present on all processes within one iteration after injection. However, the magnitude of the corruption is below our threshold deviation tolerance of $1e-10$ and does not appear in Fig. 1. As HPCCG continues to iterate, corruption propagates in all processes. Over time, corruption is removed as HPCCG refines the solution to $Ax=b$. The largest in magnitude and most difficult to remove corruption remains on the process that experienced the fault, rank 3.

Analyzing HPC applications and how they propagate corruption is useful in determining what types of soft-error detectors to use and where they are best placed. Furthermore, knowing which variables are corrupted and on which processes corruption resides

when a soft error is detected enables low-cost localized recovery instead of an expensive global rollback recovery, checkpoint-restart. Going forward, as HPC design constraints increase the likelihood of soft errors that can impact HPC application data, efficient detection and recovery schemes are needed to ensure applications obtain correct results.

WHY BLUE WATERS

The Blue Waters system allows us to perform fault injection and track propagation in thousands of HPC application runs and analyze the data generated more efficiently than other available systems. The fast turnaround time in obtaining and analyzing results greatly increased the speed and quality of this project.

PUBLICATIONS AND DATA SETS

Calhoun, J., M. Snir, L. Olson, and M. Garzaran, Understanding the Propagation of Error Due to a Silent Data Corruption in a Sparse Matrix Vector Multiply. *Proceedings of the 2015 IEEE International Conference on Cluster Computing*, (IEEE Computer Society, Chicago, Ill., September 8–11, 2015), pp. 541–542.

Calhoun, J., M. Snir, L. Olson, and W.D. Gropp, Towards a More Complete Understanding of SDC Propagation. To appear in *Proceedings of the 26th International Conference on High Performance and Distributed Computing*, (ACM, Washington, D.C., June 26–30, 2017).

Jon Calhoun was a fifth-year Ph.D. student in computer science working under the direction of Luke N. Olson and Marc Snir at the University of Illinois at Urbana-Champaign when this work was completed. He graduated in August 2017.

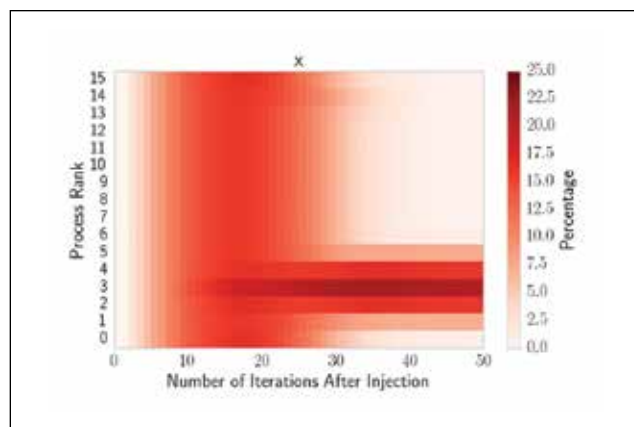


Figure 1: Propagation of corruption (percentage of elements deviated by at least $1e-10$) in the solution variable of x of the application HPCCG.

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

Justin Drake, University of Texas Medical Branch
2016–2017 Graduate Fellow

EXECUTIVE SUMMARY

Intrinsically disordered regions (IDRs) in proteins are highly dynamic and often undergo conformational transitions upon binding a protein target. The thermodynamics underlying these transitions, particularly conformational entropy and its contribution to protein binding, remain elusive. In this work, we use molecular dynamics simulations to sample the conformational states of a model disordered polypeptide as a function of chain length. Backbone dihedral angles for each polypeptide were measured from the simulation trajectories and used to calculate the conformational entropy using the Quasi-Harmonic Analysis (QHA), Boltzmann Quasi-Harmonic (BQH) Analysis, and two variants of the Mutual Information Expansion (MIE) method. Our results suggest that short IDRs may provide a significant source of free energy that proteins may tap through order-disorder transitions to modulate or regulate protein binding. We are currently working to develop a model that more fully describes the energetics of these conformational transitions and the extent to which they contribute to protein binding processes.

RESEARCH CHALLENGE

For much of the twentieth century, it was widely held that the three-dimensional structure of a protein dictates its cellular function. While true for a large group of proteins, at the turn of the twenty-first century, seminal papers (1, 2) opened the door to an entirely new class of proteins that rely on highly dynamic, flexible (i.e., disordered) regions to carry out their functions. These




Figure 1: Conformational entropy calculated from the dihedral angles of successively longer oligoglycine polypeptides using the QHA, BQH, and two variants of the MIE method.

intrinsically disordered regions (IDRs) found within proteins play an important role in protein-signaling networks and regulation and have been implicated in a number of diseases (1, 2). Upon binding a target protein, IDRs often undergo conformational transitions to a more ordered or disordered state and may directly interface with the target or be located distally to the binding interface (3, 4). These conformational transitions involve relatively short IDRs on the order of ~10 residues (4). To successfully target drugs to IDRs or genetically engineer IDRs with certain therapeutic properties to treat various diseases, we need a complete understanding of the thermodynamics associated with binding-induced conformational transitions of short IDRs.

METHODS & CODES

Conformational entropy (i.e., a proxy for disorder) of IDRs associated with order-disorder transitions is believed to play an important role in protein-binding processes (3). However, it is notoriously difficult to measure or approximate conformational entropy using traditional solution biophysics techniques. In this work, we use all-atom molecular dynamics simulations to sample the conformational states of successively longer oligoglycine polypeptides (Gly_N where N = 3, 4, 5, 10, and 15 residues), and use a number of methods to calculate the dihedral angle contribution to the conformational entropy. Oligoglycine is an ideal protein backbone model and is found in varying lengths in a number of IDRs. Simulations were performed with the NAMD (5) molecular dynamics (MD) package and Amber ff12SB (6) force field. Backbone dihedral angles for each model were measured from the simulation trajectories and used to calculate the conformational entropy using QHA (7), BQH (8), and two variants of the MIE (9) method.

RESULTS & IMPACT

The structural and thermodynamic properties of IDRs depend on chain length. In this work, we have calculated the conformational entropy (S) of an oligoglycine model as a function of chain length using four methods that account in different ways for the reduction in entropy due to correlated dihedral motions. Fig. 1 shows that conformational entropy scales remarkably linearly with chain length, with slopes ranging between 4.1 and 5.6 cal/mol/K per residue. This is consistent with the experimentally measured loss of backbone entropy (4.5 cal/mol/K per residue) upon protein folding (10). S^{QHA} is a strict upper bound on the conformational entropy. S^{MIE}_{res} and S^{MIE}_2 differ in that the former accounts only for

dihedral angle correlations within each residue, whereas the latter considers all possible pairs of dihedral angles along the backbone, yielding a slight reduction in the slope of S^{MIE}_2 (4.1 cal/mol/K/residue) compared to S^{MIE}_{res} (4.7 cal/mol/K/residue). We note that S^{MIE}_{res} is the simplest and most computationally efficient method of the four. Coupling of the dihedral angles within each residue appears to be the most significant source of correlations affecting the conformational entropy (Fig. 2). Preliminary data also suggest that oligoglycine conformational entropy is relatively insensitive to constraints on end-to-end distance, radius of gyration, and solvent-accessible surface area.

At a temperature of 300°K, the conformational entropy of these short IDR models (e.g., ~21 kcal/mol for Gly¹⁵) provides a significant free energy reservoir that proteins may tap through order-disorder transitions to modulate or regulate protein binding. We will incorporate our current results with those reported in last year's Blue Waters Annual Report on the solvation thermodynamics of oligoglycines to create a more complete model of the thermodynamic contribution to protein binding energetics associated with these conformational transitions. We will then evaluate this current model in the context of biologically relevant examples.

WHY BLUE WATERS

Statistically converged estimates of conformational entropy require extensive sampling of the conformational states of a protein or polypeptide during molecular dynamics (MD) simulations. Access to GPU-optimized MD packages such as NAMD for use with Blue Waters XK nodes, assistance from the project staff to develop an efficient workflow, and the high-throughput of Blue Waters was critical to performing our research in a tractable amount of time. Additionally, the Blue Waters Graduate Fellowship provided me the opportunity to successfully propose, manage, and partially fund my doctoral research, and expanded my understanding of high-performance computing. These experiences are instrumental in my pursuit of becoming an independent researcher.

PUBLICATIONS AND DATA SETS

Ou, S.C., J. A. Drake, and B. M. Pettitt, Nonpolar Solvation Free Energy from Proximal Distribution Functions. *J. Phys. Chem. B*, 121 (2017), pp. 3555–3564.

Drake, J.A., R. C. Harris, and B. M. Pettitt, Solvation Thermodynamics of Oligoglycine with Respect to Chain Length and Flexibility. *Biophysical Journal*, 111 (2016), pp. 756–767.




Figure 2: Representative free energy map of the two backbone dihedral angles (ϕ , ψ) in a glycine residue highlights correlations between the two. The free energy is color coded and relative to the most populated dihedral pair state. The underlying pair probability distributions of the map are used in MIE entropy calculations.

Drake, J.A., and B.M. Pettitt, Solvation thermodynamics of oligoglycine: implications for aggregation and collapse. *The 4th Annual NCSA Blue Waters Symposium* (Sunriver, Ore., June 12–15, 2016).

Drake, J.A., and B.M. Pettitt, Solvation thermodynamics of oligoglycine with respect to chain length and flexibility: implications for aggregation and collapse. *72nd Annual American Chemical Society Southwest Regional Meeting: Innovations in Energy and Medicine* (Galveston, Tex., November 10–13, 2016).

Drake, J.A., and B.M. Pettitt, The protein backbone: how structural and thermodynamic properties scale with length. *251st American Chemical Society National Meeting and Exposition* (San Diego, Calif., March 13–17, 2016).

Drake, J.A., and B.M. Pettitt, The protein backbone: how structural and thermodynamic properties scale with length. *21st Annual Structural Biology Symposium* (Sealy Center for Structural Biology and Molecular Biophysics, Galveston, Tex., April 23, 2016).

Drake, J.A., and B.M. Pettitt, Protein Backbone Thermodynamics: Towards Estimating Conformational Entropy. *21st Annual Structural Biology Symposium* (Sealy Center for Structural Biology and Molecular Biophysics, Galveston, Tex., April 23, 2016).

Drake, J.A., and B.M. Pettitt, Protein Backbone Thermodynamics: Towards Estimating Conformational Entropy. *20th Annual Structural Biology Symposium* (Sealy Center for Structural Biology and Molecular Biophysics, Galveston, Tex., May 2, 2015).

A sixth-year doctoral student in biochemistry and molecular biology at the University of Texas Medical Branch at Galveston, Justin Drake is working under the direction of B. Montgomery Pettitt. He expects to graduate in December 2017.

GENOMIC PERSPECTIVES ON THE AMPHIBIAN TREE OF LIFE

Paul Hime, University of Kentucky
2016–2017 Graduate Fellow

EXECUTIVE SUMMARY

The relationships among extant amphibians (frogs, salamanders, and caecilians) has been a longstanding debate in phylogenetics. Previous studies have supported any of the three possible topologies relating these three extant orders of amphibians. This project developed a novel amphibian-specific gene-capture system to target and sequence 220 nuclear genes in a diverse set of 296 amphibian species representing all major amphibian lineages. Using an information-theoretic approach to compare the relative model support for interordinal relationships and to calculate the magnitude of support for these relationships across genes, this project has demonstrated substantial variation in how strongly different genes support any interordinal topology. A few genes supporting each of the three interordinal models resoundingly reject the alternative models (high Δ AIC values in Fig. 1), while most other genes only weakly support any model for relationships among the three amphibian models, suggesting an erosion of phylogenetic signal over these deep evolutionary timescales.

Phylogenetic information content in many genes in the genome appears to have eroded in the nearly 300 million years since the three amphibian orders diverged, yet our results support a model of frogs and salamanders sharing a common ancestor despite substantial variation in phylogenetic signal across different genes. Our results also suggest a revised hypothesis for the relationships among the extant families of amphibians and support a more recent origin for many of the hyperdiverse lineages of frogs. Overall, this project highlights the power of phylogenomics and




Figure 1: Support for competing models of interordinal amphibian relationships—194 genes are binned on the horizontal axis by support for frogs+caecilians (Acauda hypothesis in green), frogs+salamanders (Batrachia hypothesis in blue), or caecilians+salamanders (Procera hypothesis in red). The upper panel depicts the magnitude of support *against* competing hypotheses measured by the AIC.

a model-based theoretical framework for testing phylogenetic hypotheses in the era of genome-scale evolutionary biology.

RESEARCH CHALLENGE

All organisms trace their ancestry back to a single common ancestor nearly four billion years ago [1]. Yet today, life has diversified into tens of millions of species. Reconstructing these evolutionary relationships is a key aim of the field of phylogenetics, and such insights may inform nearly all aspects of modern biology. The recent advent of genome sequencing technologies has ushered in a new era of phylogenomics in which genetic data generation is no longer a rate-limiting step toward testing phylogenetic hypotheses. Yet, it is becoming increasingly clear that different regions of the genome can support conflicting phylogenetic hypotheses, and reconciling these discordant gene genealogies is a key problem facing evolutionary biologists today. Information-theoretic approaches have great promise for parsing signal from noise in large phylogenomic data sets.

METHODS & CODES

Three possible topologies exist for relationships among frogs, caecilians, and salamanders. The Procera hypothesis supports caecilians+salamanders, the Acauda hypothesis supports frogs+caecilians, and the (canonical) Batrachia hypothesis supports frogs+salamanders. To evaluate support for each of these three hypotheses across the genome, we conducted gene-by-gene tests of constrained topology, comparing the maximum likelihood estimates of gene trees for 194 nuclear genes between constraints for the three possible interordinal topologies. Gene trees were estimated in RAxML [2]. The Akaike information criterion (AIC) [3] was used to quantify the direction and magnitude of support across genes. Gene trees were then reconciled into an estimate of the species tree using Astral [4]. This species tree topology was then used in concert with a set of 25 fossil calibrations to estimate divergence times across Amphibia in MCMCTree [5].

RESULTS & IMPACT

This project has demonstrated that different genes in the amphibian genome support conflicting topologies for the relationships among the three amphibian orders. Roughly half of the genes examined support frogs and salamanders as each other's closest relatives (sister taxa). Roughly one-quarter of the genes support frogs and caecilians as sister taxa, while another quarter support caecilians and salamanders as sister taxa. Significant numbers of genes are found to support each of these three competing topologies, suggesting that either there has been an erosion of phylogenetic signal over deep time and/or




Figure 2: Estimated divergence times across Amphibia—Amniotes, caecilians, salamanders, and frogs are depicted by magenta, red, blue, and green branches, respectively. The dashed red line denotes the Cretaceous–Tertiary boundary, a period that witnessed several independent increases in the net rate of diversification in multiple amphibian lineages.

that population-level processes (large effective population sizes and rapid divergences) have led to incomplete lineage sorting and gene tree, or species tree discordance. Results from our AIC-based topology testing support the Batrachia hypothesis (frogs and salamanders are each other's closest relatives) at the level of the species tree and suggest that large numbers of genes may need to be analyzed in order to overcome stochastic phylogenetic noise at deep timescales. Our divergence time analyses also provide a revised timescale for amphibian diversification through time and indicate that a rapid increase in net rates of species diversification occurred at the Cretaceous–Tertiary boundary. Analyzing hundreds of genes for hundreds of species has traditionally been computationally intractable for empirical data sets, and this study is one of the first to use an information-theoretic framework to address not only the direction of support for phylogenetic hypotheses across the genome but also the magnitude of that support.

WHY BLUE WATERS

The Blue Waters fellowship has provided essential support to enable me to make progress on this project and to successfully write and defend my Ph.D. dissertation. Although many of the analyses I intended to execute on Blue Waters have been delayed due to unforeseen issues with our data set, I am hopeful that with a time extension to this allocation, I will be able to leverage the massively parallel resources on Blue Waters to perform topological testing in Neobatrachian frogs at a level of gene- and species-sampling that has not been attempted in any other empirical data set.

Paul Hime—a sixth-year Ph.D. student in evolutionary biology at the University of Kentucky—successfully defended his dissertation in April 2017. His work was directed by David W. Weisrock.

MULTISCALE SIMULATIONS OF COMPLEX FLUID RHEOLOGY

Michael P. Howard, Princeton University
2016-2017 Graduate Fellow

EXECUTIVE SUMMARY

We developed the first massively parallel, open-source implementation of the multiparticle collision dynamics (MPCD) algorithm for graphics processing units (GPUs). MPCD is a mesoscale particle-based simulation method for hydrodynamics, and is particularly useful for modeling complex fluids and soft matter. Our implementation of the MPCD algorithm scales up to 1024 nodes on Blue Waters, and GPU acceleration gives a 3x speedup for an XK node compared to an XE node. The developed software will enable studies of complex fluids at length and time scales that would be otherwise inaccessible.

RESEARCH CHALLENGE

Complex fluids, readily encountered in biology, consumer products, and industrial processing, are multicomponent mixtures that exhibit a rich variety of flow behaviors. A classic example is the cornstarch-water “oobleck” mixture, which acts like a liquid when pressed slowly but can thicken to support the weight of a person

when struck quickly. Such peculiar macroscopic flow properties of complex fluids are fundamentally controlled by microscopic molecular structures and interactions. Computer simulations are ideal tools for studying this nontrivial and difficult-to-predict relationship; however, performing simulations of complex fluids at physically relevant scales presents a considerable challenge.

Many complex fluids consist of a mixture of solute macromolecules (polymers) and hard spherical particles (colloids) suspended in a liquid solvent such as water. To obtain the correct dynamics, it is essential to resolve both the direct solute-solute interactions as well as the solvent-mediated interactions. The latter dominate the computational cost for a molecular model of such a mixture; however, a molecular-level description of the solvent itself is often not of interest. A multiscale approach that simplifies the solvent model while preserving its most important interactions is required to study complex fluids at relevant length and time scales.

METHODS & CODES

Multiparticle collision dynamics (MPCD) [1] is a mesoscale simulation method that combines classical molecular dynamics (MD) simulations of microscopic solutes with a coarse-grained solvent representation that faithfully reproduces long-range hydrodynamics and thermal fluctuations. In MPCD, the solvent is modeled by a set of point particles that alternate between ballistic streaming steps and cell-based, momentum-conserving, stochastic multiparticle collisions. The frequency and nature of the collisions control the transport coefficients of the solvent. Solute molecules propagate using the MD equations of motion, and are coupled to the MPCD solvent during either the streaming or collision step [2, 3]. Significantly larger length and time scales can be accessed in MPCD than with an explicit-solvent MD model because MPCD solvent particles have simple interactions with each other. However, in practice, tens of millions of MPCD particles may be required for a simulation, necessitating a parallel computational approach. We developed the first massively-parallel, open-source implementation of MPCD for graphics processing units (GPUs).

We implemented the MPCD algorithm as part of the open-source simulation package HOOMD-blue [4, 5], which has optimized MD methods designed for NVIDIA GPU architectures. All MPCD data are stored independently from HOOMD-blue’s MD data to ensure high performance. Nearly all computations are performed exclusively on the GPU, which minimizes latency from host-device data transfers. We employ a spatial domain decomposition strategy to extend support to multiple GPUs [5]. One MPI process is assigned per GPU, and we perform runtime autotuning to ensure optimal CUDA kernel launch parameters. Flexible initialization and simulation setup is supported through a scriptable Python user interface.

RESULTS & IMPACT

We performed strong-scaling benchmarks of our MPCD implementation on Blue Waters. We benchmarked two cubic simulation boxes with edge lengths $L = 200a$ and $L = 400a$, where a is the size of an MPCD cell. There were 10 MPCD solvent particles per cell, giving 80 million and 640 million total particles in each simulation box, respectively. We benchmarked a CPU-only implementation of our code using 16 processes per XE node and our GPU implementation using 1 process per XK node. The CPU-only code showed excellent strong scaling up to 1024 nodes. The GPU code also showed good scaling to 1024 nodes, with some performance lost at the highest node counts owing to communication latency for the MPCD cells. GPU acceleration on the XK nodes gave a roughly 3x speedup compared to the XE nodes. The complex fluid and soft matter research communities will significantly benefit from the developed MPCD software, especially the GPU implementation, which will permit studying processes at scales that would otherwise be inaccessible.

WHY BLUE WATERS

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




Figure 1: Strong-scaling benchmarks of MPCD software on Blue Waters for CPU-only (XE nodes) and GPU (XK nodes) implementations.

Michael P. Howard is currently completing the fourth year of a Ph.D. program in chemical engineering at Princeton University, under the direction of Athanassios Z. Panagiotopoulos. He expects to graduate in 2018.

HIGH ACCURACY RADIATIVE TRANSFER IN CLOUDY ATMOSPHERES

Alexandra L. Jones, Princeton University
2016-2017 Graduate Fellow

EXECUTIVE SUMMARY

One of the most important roles clouds play in the atmosphere is in redistributing radiative energy from the sun and that which 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.

RESEARCH CHALLENGE

A model that treats broadband integration and 3D radiative transfer in a highly accurate and unbiased way is needed to quantify the bias in the simpler models ubiquitously used. This model will serve as a previously nonexistent standard of comparison for other similar models and provide accuracy bounds for simpler models and parameterizations attempting to capture 3D effects at lower computational cost. Such a model was not publicly available prior to this project. So, one was developed that uses Monte Carlo methods to capture the 3D transfer of radiation and to sample at high resolution the broad range of the electromagnetic spectrum.

METHODS & CODES

Unlike the direct approach to solving the radiative transfer equation, the Monte Carlo approach has the potential to be embarrassingly parallel, since the random samples are independent from one another. The figure shows weak and strong scaling tests for both solar and internal sources of photons for the monochromatic model, “IMC+emission,” in terms of throughput, or number of photons simulated per minute, and throughput efficiency. For both weak and strong scaling experiments, the solar source of photons exhibits a faster drop-off in efficiency but has an overall higher throughput than the internal source of photons. The “IMC+emission” model shows better weak scaling efficiency than strong scaling efficiency for both sources of photons over the range of processes tested.

RESULTS & IMPACT

The overarching goal of this project is to make publicly available to the radiative transfer community the models, tools, data, and products developed to aid in faster and more robust progress in addressing scientific questions about the interactions of clouds and realistic radiative transfer. An existing monochromatic 3D Monte Carlo community solar radiative transfer model was further developed to include terrestrial emission in addition to solar sources of radiation. That model was then further developed to include integration over the electromagnetic spectrum to produce the broadband 3D model discussed above. In addition to the development of these two community models, several other products have resulted so far and will be made available to the community. These include databases of high spectral resolution radiative properties of earth’s gaseous atmosphere and liquid

water clouds, the largest and highest resolution publicly available databases of their kind. The tools and workflow to create and subset them will also be made available. These data can be mined to update the decades-old broadband parameterizations of cloud radiative properties that are still in wide use today, for example. Each product has been thoroughly vetted for accuracy. The results of these tests will be made available for reproduction by other scientists to test these models or their own. Finally, the first few idealized experiments with long heritage in the literature have been conducted to provide the first set of benchmark simulation results that can be used to evaluate other models.

WHY BLUE WATERS

Access to debugging and profiling tools such as CrayPat and DDT allowed me to streamline the development process. Having access to a point of contact on the SEAS staff helped me think through issues and find tailored solutions for my problems that would have otherwise delayed progress for weeks. The responsiveness of the Blue Waters staff through the JIRA ticket system allowed for limited interruption in progress when small issues or questions arose. My experience as a Blue Waters graduate fellow has been invaluable to my professional development. I hope to make use of Blue Waters for the rest of its lifetime.

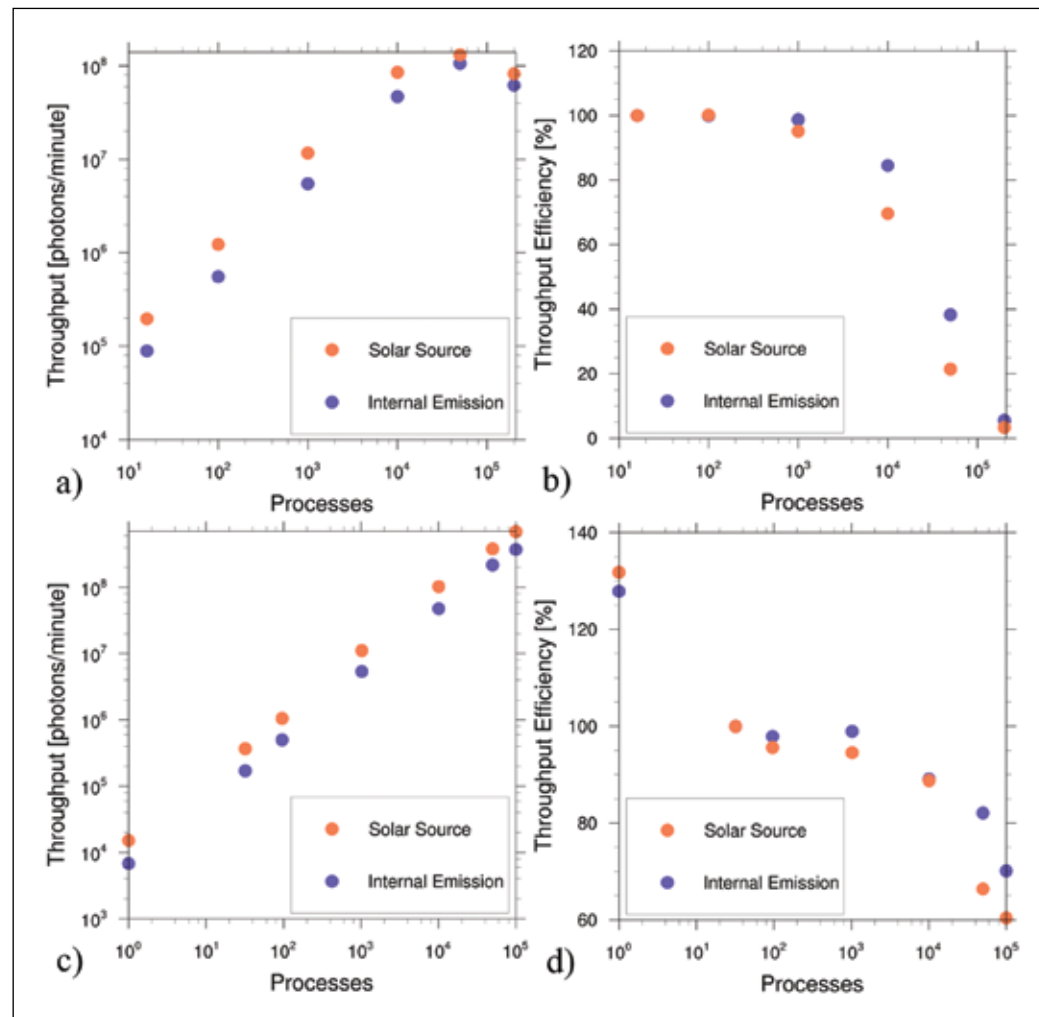
PUBLICATIONS AND DATA SETS

Jones, A. L., Development of an Accurate 3D Monte Carlo Broadband Atmospheric Radiative Transfer Model (doctoral dissertation—2016), University of Illinois, Urbana, Ill.

Jones, A. L., Alexandraljones/imc-emission: Code base plus select benchmark results (2017), DOI:10.5281/zenodo.574872.

Alexandra L. Jones received her Ph.D. in May 2016 in atmospheric science from the University of Illinois at Urbana-Champaign. She currently is a postdoctoral scholar at the Cooperative Institute for Climate Science, which is a collaboration between Princeton University and the National Oceanographic and Atmospheric Administration’s Geophysical Fluid Dynamics Laboratory.

Figure 1: Performance of the “IMC+emission” model runs in both internal emission mode and solar source mode for strong scaling experiments (a-b) and weak scaling experiments (c-d). Throughput (a, c), defined as the number of photons simulated per minute of wall-clock time, and throughput efficiency [%] (b, d) are shown as a function of number of processes. Note that logarithmic spacing is used on the horizontal axis in all panels and on the vertical axis for panels a and c.



HIGH-FIDELITY BLADE-RESOLVED WIND FARM SIMULATIONS

Andrew Kirby, University of Wyoming
2016-2017 Graduate Fellow

EXECUTIVE SUMMARY

Wind energy is becoming an emergent renewable energy source throughout the world. Costs have dropped dramatically over the past two decades, making wind energy a desirable alternative to fossil fuels. Improvement in wind energy application simulation technologies may have a profound economic impact through improved wind plant efficiency. The goal of this work is to develop state-of-the-art aerodynamics modeling techniques using high-fidelity blade-resolved turbine models to simulate complete wind farms. The numerical methods in this research utilize multiple mesh and multiple computational fluid dynamics flow solvers coupled in an overset framework.

RESEARCH CHALLENGE

High-fidelity numerical simulation of wind energy applications is becoming a precedent for future technologies, not only for the wind energy sector but also for lower-fidelity modeling as well. The need for high-fidelity simulation using a complete geometric description of wind turbines with tower and nacelle is essential for capturing the true aerodynamic nature of the flow, which is highly turbulent and chaotic. This is especially important in the study of downstream wake effects on wind turbines that cause

the primary decrease in wind plant power production efficiency. Wind farm simulation is a truly multiscale problem; the smallest turbulent structures are on the order of microns and the largest scales present from atmospheric turbulent inflow are on the order of tens of kilometers. To simulate ten orders of spatial magnitude, tens to hundreds of thousands of computing cores are required to approach this grand challenge problem.

METHODS & CODES

The numerical methods developed in this research utilize an overset grid paradigm where multiple meshes and multiple flow solvers are used in a coupled manner to enable efficient simulation of this truly multiscale problem. Nearly all software used on our framework is developed in-house at the University of Wyoming with the exception of the p4est adaptive mesh refinement framework developed by Carsten Burstedde, et al. [1]. The two flow solvers developed at the University of Wyoming are NSU3D, an unstructured 3D finite-volume solver, and dg4est, a high-order discontinuous Galerkin finite-element solver. The overset solver used is TIOGA, developed by Jay Sitaraman of Parallel Geometric Algorithms, LLC. The complete framework is known as the Wyoming Wind and Aerodynamics Applications Komputation Environment (WwAaKE3D).

RESULTS & IMPACT

This research enabled the largest complete wind farm simulations to date using a high-fidelity blade-resolved turbine model. This demonstrates the capability for large-scale simulations of entire wind farms and serves as a milestone for high-fidelity methods in a multi-scale problem spanning 10 orders of spatial magnitude.

WHY BLUE WATERS

Blue Waters offers a unique environment not only as a computational resource but also for its expert project staff. The design of Blue Waters makes it an excellent machine geared toward scientific output rather than just its flop rate. Blue Waters allowed us to perform large-scale wind farm simulations using tens of thousands of compute cores. In addition, the project staff provided excellent insight for optimization and maximizing throughput.

PUBLICATIONS AND DATA SETS

Kirby, A., et al., Wind Farm Simulations Using an Overset *hp*-Approach with Blade-Resolved Turbine Models. *Proceedings of the 23rd AIAA Computational Fluid Dynamics Conference* (AIAA Aviation, Denver, Colo., June 5–9, 2017), DOI: 10.2514/6.2017-3958.

Brazell, M., A. Kirby, and D. Mavriplis, A high-order discontinuous-Galerkin octree-based AMR solver for overset simulations. *Proceedings of the 23rd AIAA Computational Fluid Dynamics Conference* (AIAA Aviation, Denver, Colo., June 5–9, 2017), DOI: 10.2514/6.2017-3944.

Andrew Kirby is a fifth-year Ph.D. student in mechanical engineering at the University of Wyoming. He is working under Dimitri Mavriplis and expects to graduate in December 2017.




Figure 1: Volume rendering of vertical wake structures generated downstream of a Siemens SWT-2.3-93 wind turbine with uniform inflow.

REDUCING THE COMPUTATIONAL COST OF COUPLED CLUSTER THEORY

Sara Kokkila Schumacher, Stanford University
2016-2017 Graduate Fellow

EXECUTIVE SUMMARY

The retinal chromophore, found in the photoactive retinylidene proteins, is a challenging system for many of the excited state methods that are currently available. This chromophore plays an important role in vision in humans and other species. In order to accurately and efficiently simulate the absorption spectra of retinal models, our project aims to develop a parallel implementation of the tensor hypercontraction equation-of-motion second-order approximate coupled cluster singles and doubles (THC-EOM-CC2) method. We implemented an MPI (message passing interface) version of THC-EOM-CC2 that is accelerated by the graphics processing units (GPUs) available on Blue Waters. We successfully applied our new approach to several retinal models. Our initial findings indicate that THC-EOM-CC2 performs better than other available excited state methods that are capable of treating this system size.

RESEARCH CHALLENGE

Retinylidene proteins are photosensitive proteins that serve a variety of biological functions and have been found in all domains of life. [1] For example, the rhodopsin protein and the visual pigments found in human eyes enable vision. [2] Other retinylidene proteins can induce movement in certain types of bacteria or act as light-induced pumps. Retinylidene proteins consist of an opsin protein and the retinal chromophore. Through the use of Blue Waters, we aim to simulate the absorption spectra of different retinal chromophore models at an unprecedented level of accuracy.

Simulating the excited states of the retinal chromophore necessitates the use of electronic structure methods since simulation requires an accurate description of the electronic wave function. Unfortunately, the accuracy and computational

demands of an electronic structure method are often at odds with each other. The second-order approximate coupled cluster singles and doubles (CC2) method can be extended to excited states through the equation-of-motion formalism (EOM). [3,4] The EOM-CC2 method is robust against many of the issues found in inexpensive excited state alternatives such as configuration interaction singles or time-dependent density functional theory. [5] However, studying the retinal chromophore with canonical EOM-CC2 is infeasible with typical computational resources. This is because the runtime of canonical EOM-CC2 increases formally on the order of $O(N^5)$ where N is related to the size of the chemical system.

METHODS & CODES

In an effort to reduce the computational cost of EOM-CC2 and other quantum chemistry methods, we developed the tensor hypercontraction (THC) approximation. [6] We applied the THC approximation to EOM-CC2 (THC-EOM-CC2) and showed that this lowered the cost of the calculation from an order of $O(N^5)$ to $O(N^4)$. [7,8] To further improve the efficiency of THC-EOM-CC2, we redesigned the algorithms to take advantage of parallelism. This included developing a code that was parallelized at a higher level to take advantage of multiple compute nodes and at a finer level to take advantage of acceleration with GPUs. The parallel THC-EOM-CC2 approach was developed in TeraChem. [9]

We are working with the grid-based variant of the THC approximation, which allows the new algorithm to block over grid point indices. Additionally, we use a Laplace transformation to express certain terms by numerical quadrature in this method. In the parallel implementation, each MPI task is assigned a quadrature point and set of grid points. The THC-EOM-CC2

approach is memory-intensive, and by blocking over grid points on an MPI task, we can exploit this level of parallelism while ensuring that the computations fit within the limited memory available on the GPU. This new approach allows us to extend the THC-EOM-CC2 method to system sizes that are challenging for canonical EOM-CC2.

RESULTS & IMPACT

Our approximation, combined with the parallelization approach developed in this work, offers a new way to parallelize and improve the efficiency of quantum chemistry methods. In addition to its use in CC2 and EOM-CC2, the THC approximation has been shown to reduce the computational complexity of other electronic structure methods. [6] While this work was focused on parallelizing THC-CC2 and THC-EOM-CC2, the general parallelization design outlined can be applied to other quantum chemistry methods using the THC approximation. We expect this new approach will be used to improve the efficiency of many other electronic structure methods in different quantum chemistry software.

The development of the MPI-enabled and GPU-accelerated THC-EOM-CC2 code and the use of Blue Waters allow us to study the absorption spectra of several model retinal chromophores. We find that for a fixed problem size, the parallel implementation of THC-EOM-CC2 scales superlinearly in the regime where we parallelize only Laplace quadrature points across compute nodes. In the regime where we distribute both quadrature points and grid point blocks across compute nodes, we see a computational speedup that levels off as the number of compute nodes increases.

We are currently comparing the results of THC-EOM-CC2 to other excited state methods. Our initial findings indicate that the THC-EOM-CC2 method more accurately describes the excited states of the retinal models compared to the results from time-dependent density functional theory. The THC-EOM-CC2 method will be a useful tool for studying the excited states of chemical systems that are otherwise challenging for the available excited state methods.

WHY BLUE WATERS

Access to Blue Waters allowed us to develop the MPI-enabled and GPU-accelerated THC-EOM-CC2 method. Blue Water's code development environment and tools helped us test and debug different implementations more rapidly. This work represents the first time we have tried to combine MPI and GPU acceleration for a method within TeraChem. The Blue Waters project staff offered valuable insight into code development for high-performance computing systems. Additionally, studying the absorption spectra of the retinal models requires sampling different conformations of each retinal model. Blue Waters enabled us to calculate the ground state and multiple excited states of many different configurations.

PUBLICATIONS AND DATA SETS

Kokkila Schumacher, S.I.L., E. G. Hohenstein, R. M. Parrish, L.-P. Wang, and T. J. Martínez, Tensor Hypercontraction Second-Order Møller–Plesset Perturbation Theory: Grid Optimization and Reaction Energies. *J. Chem. Theory Comput.*, 11 (2015), pp. 3042–3052.

Sara Kokkila Schumacher received a Ph.D. in chemistry from Stanford University in January 2017. She currently is a postdoctoral researcher in high-performance computing at IBM.




Figure 1: One of the retinal models used in this study.

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

Larisa Reames, The University of Oklahoma
2016–2017 Graduate Fellow

EXECUTIVE SUMMARY

This study used real-data simulations to quantify the impacts of a large Great Plains (U.S.) urban area on the evolution and strength of a supercell thunderstorm. Simulations with urban areas are compared to those without, with the aid of hierarchical clustering analysis, to form statistically similar groups of simulations. We investigated the effects of the storm with various city-relative paths as well as the storm lifecycle stage during urban interactions. These comparisons concentrate on differences in boundary layer characteristics prior to storm formation as well as changes in supercell structure, dynamics, and evolution. Additionally, we performed a factor separation analysis to determine which aspects of the urban area (i.e., roughness or thermal differences) have the most effect. The results suggest that the urban area, particularly surface roughness, can have a significant impact on storm strength, and that these modifications change based on the city-relative path of the storm.

RESEARCH CHALLENGE

Earth's population is increasingly concentrated in urban areas, with nearly two-thirds of the world's population expected to live in urban areas by 2050. As the number of people within cities grows, it is becoming more important to understand and to correctly predict the interactions between urban environments and the atmosphere. As such, many studies have investigated the effect of urban areas on weakly forced precipitation systems. However, interactions

between urban areas and synoptically-active convection, such as supercells, remain relatively unexamined.

METHODS AND CODES

Using the Weather Research and Forecasting (WRF) [1] model—a community mesoscale numerical weather prediction model—we investigated a total of 134 simulations of a supercell thunderstorm to quantify the impacts of a large Great Plains urban area on the evolution and strength of a supercell thunderstorm. In order to properly resolve complex urban structure, all simulations were run on a 500-m horizontal grid over a 250-km x 250-km grid. In addition, to resolve the atmospheric boundary layer well, 120 vertical grid points were used, with 20 in the lowest 1.5 km above ground. In all, over 29.7 million points were integrated over 75,600 time steps for each simulation. Ten of the simulations contained homogeneous land use (CTRL) to serve as a comparison point for simulations with urban areas. An urban area simulated to have both increased surface roughness and thermal properties characteristic of manmade surfaces (i.e., “full physics”) was placed in 108 gridded locations (WestF and EastF) throughout the domain to determine effects of the city-relative path of the storm. At eight of these city locations (four near the middle (East) and four near the beginning (West) of the storm's track), two additional simulations were performed with only either thermodynamic (WestT and EastT) or roughness (WestR or EastR) properties of the urban area simulated, providing an additional 16 simulations.

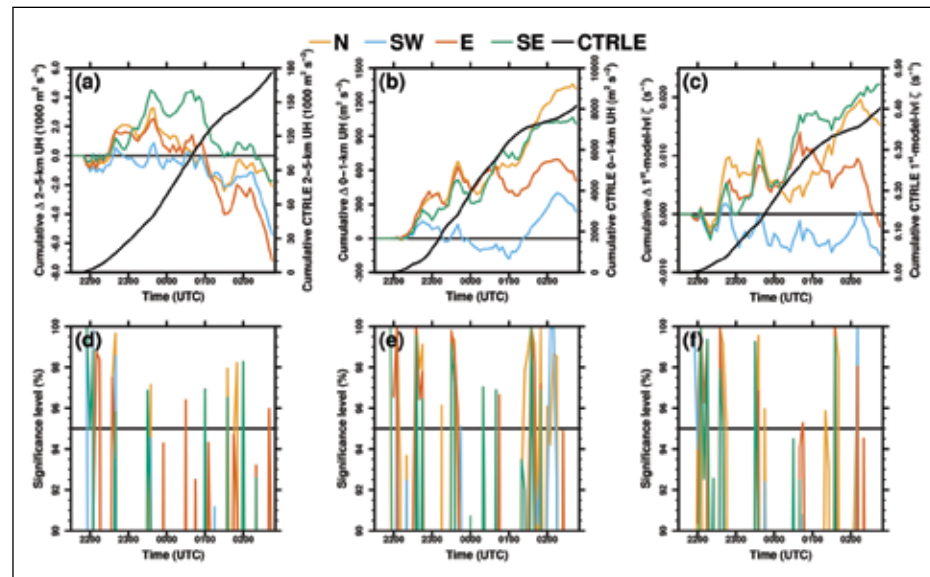


Figure 1: Cumulative CTRL group average (black) and accumulated difference from CTRL of N (yellow), SW (blue), E (red), and SE (green) groups of (a) mid-level (2- to 5-km) storm strength, (b) low-level (0 to 1-km) storm strength, and (c) near-surface circulation strength as a function of time.

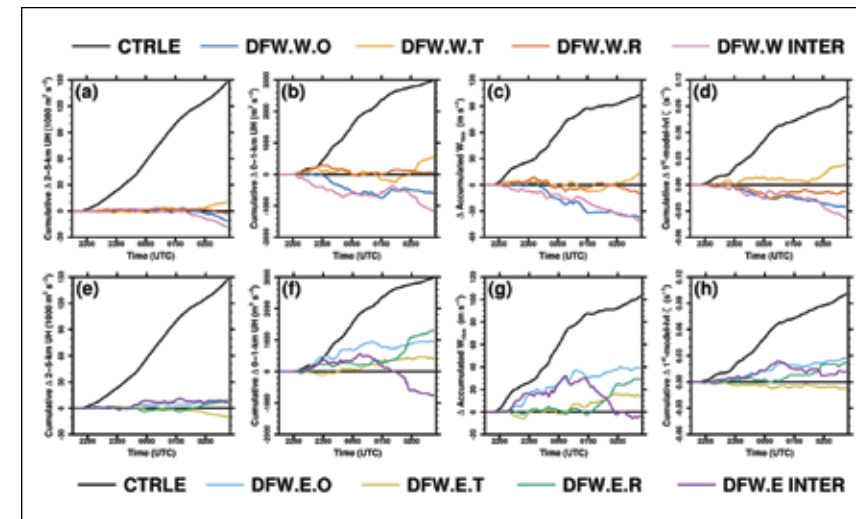


Figure 2: Cumulative CTRL group average (black) and accumulated difference from CTRL of urban simulations (colored lines) of mesocyclone-averaged [(a) and (e)] mid-level (2- to 5-km) storm strength, [(b) and (f)] low-level (0 to 1-km) storm strength, [(c) and (g)] maximum updraft speed, and [(d) and (h)] near-surface circulation strength.

RESULTS AND IMPACT

We compared full-physics urban simulations to CTRL with the aid of hierarchical clustering analysis (HCA) to form statistically similar groups of simulations. In this analysis, we investigated the effects of the storm having various city-relative paths, as well as the storm lifecycle stage during urban interactions. These comparisons concentrate on differences in boundary-layer characteristics prior to storm formation as well as changes in supercell structure, dynamics, and evolution. The results (Fig. 1), suggest that when the storm passes to the north of or directly over the city center late in its life cycle, low- and mid-level storm rotation strength increases, and the center of rotation tracks farther south. Although differences in mid-level rotation are minimal when the storm tracks south of the city, low-level rotation increases, especially late in the storm's lifetime.

Using the 16 individual physical models (i.e., surface roughness or thermodynamic urban properties only), we undertook a factor separation approach to determine the relative importance of the roughness and thermal characteristics of urban areas on storm modification. Results (Fig. 2) generally suggest that surface roughness and its interactions among thermodynamic properties are the dominant contributors to urban-induced effects on storm strength and evolution. Additionally, the amplitude of interactions between shear and thermodynamic modifications is often similar in magnitude to either effect individually.

Most investigations of near-surface urban properties and urban-storm interactions have focused on the impact of the thermal effects of cities, but the results of this investigation argue for a greater focus on the implications of urban surface roughness. Unlike the full-physics simulations, those simulations with the city parameterized only as a roughness element, and with a long urban-storm interaction, resulted in the largest differences from nonurban simulations of any of the factor-separation analyses.

WHY BLUE WATERS

While HCA has been used previously for attribution of variations in synoptic and mesoscale fields to various factors, this is the first time it has been used to analyze storm-scale modifications. Given their large scale of motion, synoptic [O (1000 km)] and mesoscale [O (100 km)] phenomena are generally more predictable than severe storms [O (10 km)], thus few simulations are required to attribute large-scale field variations to modifications in boundary conditions and parameterization options. However, to perform attribution of small-scale effects to various factors in a real-data (i.e., nonidealized) simulation, many simulations are required to ensure that the simulated changes are significant. The general hindrance to such an analysis is the large computational requirement; hence, the resources made available on Blue Waters were vital to this work. While each simulation was relatively small, the large quantity of simulations needed to produce significant results required the large computational and data storage capacities of Blue Waters.

Larisa Reames graduated from the University of Oklahoma in May 2017 with a Ph.D. in meteorology. There, her work was directed by David Stensrud at Pennsylvania State University. She currently is a postdoctoral research associate at the National Oceanic and Atmospheric Administration's National Severe Storms Laboratory.

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

Sherwood Richers, California Institute of Technology
2016-2017 Graduate Fellow

EXECUTIVE SUMMARY

The transport of neutrinos plays a dominant role in engine-driving core-collapse supernovae, but results from simulations of neutrino transport depend sensitively on the method and implementation details of neutrino transport codes. We have enhanced the open-source Monte Carlo neutrino transport code Sedonu and compared the results against a discrete ordinates transport code to verify for the first time multidimensional Boltzmann neutrino transport algorithms. We used these results to assess the accuracy of approximate two-moment transport schemes. We then motivate the use of time-independent Monte Carlo transport calculations in future large-scale time-dependent supernova simulations. In an orthogonal study, we also analyzed how gravitational waves—the only signal other than neutrinos that directly reflects conditions deep within the explosion—could inform our understanding of dynamics in exotic supernovae.

RESEARCH CHALLENGE

Core-collapse supernovae are immense explosions following the collapse of massive stars; they release around 10^{53} ergs in neutrinos and result in explosions with kinetic energies of around 10^{51} ergs [1]. These supernovae are the origin of most of the elements of which we and the world around us are made. The extreme conditions during the explosions are also wonderful laboratories for a wide range of physics, including nuclear structure, general relativity, neutrino interactions, and magnetohydrodynamical instabilities. Though we know *what* causes these explosions, understanding *how* they happen is a difficult and computationally intense problem that has been an active area of research for over 50 years.

In the canonical theory, neutrinos emitted from the dense inner core heat and drive turbulence in matter that pushes a shock wave through the collapsing star, resulting in explosion. Due to the complexity of the involved physics and dynamics, computation has become the primary tool for understanding this mechanism. In particular, the transport of neutrinos in these systems is modeled by the six-dimensional (plus time) integro-differential Boltzmann equation, which cannot currently be adequately simulated for the duration of the explosion. The neutrino transport component of simulations is treated by many different methods, all of which include some level of approximation, even in the largest and most sophisticated simulations. Small differences in how neutrino transport is implemented can lead to entirely different simulation outcomes. Understanding the mechanism that leads to stellar

explosion thus requires both improved simulation techniques to more accurately treat neutrino transport and a means of ensuring that the techniques and discretization choices adequately simulate the physics they represent. Though verification efforts of Boltzmann-level transport schemes have been performed in spherical symmetry, in this project we make the first jump to verification of Boltzmann transport in multiple spatial dimensions. We also motivate improvements to more approximate transport methods.

METHODS & CODES

We use a variety of simulation codes to approach the problem from different directions. The key component of this project is the open-source Monte Carlo neutrino transport code Sedonu, which computes with high accuracy time-independent effects (e.g., instantaneous heating/cooling and lepton gain/loss rates) of neutrinos on one, two, and three-dimensional stationary fluid snapshots. With a paucity of direct data to validate against, we compare Monte Carlo results to those from a vastly different discrete ordinates method [2] to allow us to quantify how well we understand supernova neutrinos as we work toward a working theory of neutrino-driven core-collapse supernovae. We also assess the accuracy of popular approximate two-moment transport schemes [3] used in multidimensional core-collapse simulations. In addition, we couple Sedonu to the open-source supernova code GR1D [4]. The latter code simulates stellar core collapse in spherical symmetry using an approximate neutrino transport method, while Sedonu informs the transport approximation to bring it closer to the exact solution. Finally, we perform three-dimensional test calculations using snapshots from simulations using the open-source package Zelmani [5,6].

RESULTS & IMPACT

Understanding the core-collapse supernova engine requires not only ever-increasing simulation size and complexity, but an understanding of how well these methods reflect the physics they are meant to simulate. With this in mind, we developed a new version of the open-source Monte Carlo neutrino transport code Sedonu, in which we implemented the random walk approximation. This improvement, along with many other feature additions and optimizations made as part of this project, makes the code capable of calculating highly accurate steady-state neutrino fields and interaction rates through the entire system without resorting to artificial boundary conditions. We performed the




Figure 1: The heating rate in a supernova snapshot computed by the Monte Carlo code Sedonu (bottom left) and the discrete ordinates code of [1] (top left). They agree very well, but the small differences between them (right) could make the difference between an explosion and a dud in a dynamical simulation.

first multi-dimensional comparison of full Boltzmann neutrino transport methods, using Sedonu and the discrete ordinates code of [2]. Given that core collapse simulation results depend sensitively on small details of the neutrino transport implementation, this provides a much-needed measurement of the numerical errors associated with each method, and provides the first multi-dimensional code verification tools to the community.

Observations of neutrinos and gravitational waves from nearby supernovae both give direct, unobscured information about the nature of the processes occurring deep within the collapsing star that lead to explosion. In addition to the above studies of neutrinos in supernovae, we performed a parameter study to determine what gravitational wave signals from rotating core collapse can tell us about the nature of matter at super-nuclear densities. In this study, we performed over 1,800 axisymmetric (2D) core collapse simulations. We used 18 different parameterizations for nuclear matter properties and over 100 rotation profiles, and concluded that differences in the descriptions of nuclear matter allowed by current constraints are not likely discernible by current gravitational wave detectors, even for a galactic supernova. We also established a simple universal relationship between the gravitational wave frequency and properties of the collapsed core. The resulting publicly available data set is an order of magnitude larger than those in previous studies and represents the first thorough exploration of both rotation and nuclear matter parameters.

WHY BLUE WATERS

Without access to the unique Blue Waters environment, the careful method development and verification in this project would not have been possible. The fellowship allocation enabled these calculations independently from other resources, since the calculations are far too large for local clusters and are separate from other resource requests. This project is targeted at working toward large-scale three-dimensional core-collapse simulations with a Monte Carlo-informed approximate method for neutrino transport. The domain-replication parallelism used in this approach requires both a large amount of compute time and a large amount of memory on each node. Blue Waters provides a unique environment with both of these components. In addition, the support staff with both domain and system expertise significantly simplify the task of implementing and optimizing our algorithms on Blue Waters, allowing user time to be spent on science and development rather than code porting.

PUBLICATIONS AND DATA SETS

Richers, S., et al., Equation of state effects on gravitational waves from rotating core collapse. *Physical Review D*, 95:6 (2017), DOI: 10.1103/PhysRevD.95.063019.

Richers, S., et al., Equation of state effects on gravitational waves from rotating core collapse [Data set]. *Zenodo* (2017).

Richers, S., et al., A detailed comparison of multi-dimensional Boltzmann neutrino transport methods in core-collapse supernovae [Data set]. *Zenodo* (2017).

A fifth-year Ph.D. student in physics at California Institute of Technology when he completed this research, Sherwood Richers was working under the direction of Christian Ott. He successfully defended his dissertation in June 2017 and will graduate in June 2018.

UNDERSTANDING THE ROLE OF HYDRODYNAMIC FLUCTUATIONS IN BIOMACROMOLECULAR DYNAMICS THROUGH THE DEVELOPMENT OF HYBRID ATOMISTIC-CONTINUUM SIMULATION

Sean L. Seyler, Arizona State University
2016-2017 Graduate Fellow

EXECUTIVE SUMMARY

Biological macromolecules are nanoscale structures that largely exist in ionic aqueous conditions. Proteins, in particular, are biopolymers that fold into functional 3D structures that—under the right solvent conditions and, possibly, lipid environment—can perform mechanicochemical work, and whose dynamics span femtosecond timescales (i.e., covalent bond oscillations) to beyond the millisecond regime (e.g., glucose transport across a lipid membrane). Explicit-solvent, atomistic molecular dynamics (MD) is necessary to fully capture solute-solvent interactions but is currently limited to microsecond timescales—orders of magnitude shorter than the timescale range of most processes of biophysical interest, including even the fastest conformational transitions (tens of microseconds). We are developing a hybrid atomistic-continuum (HAC) method—with a view toward biomolecular simulation—coupling an MD engine to a novel discontinuous-Galerkin-based fluctuating hydrodynamics (FHD) solver; our current focus is to develop an FHD model capturing transport phenomena in dense fluids (e.g., water) arising at subnanometer scales.

RESEARCH CHALLENGE

There is growing interest in hybrid multi-physics simulations, where MD is used in a restricted subdomain requiring atomistic resolution, while an efficient FHD model replaces some, or all, of the solvent [1,2]. Nearly all HAC methods employ Landau-Lifschitz Navier-Stokes (LLNS), a system of stochastic equations describing mass, momentum, and energy transport subject to thermal fluctuations in the viscous stress tensor and heat flux vector [3]. LLNS assumes linear constitutive relations: Newton's law of viscosity (stress is proportional to velocity gradients) and Fourier's law (heat flux is proportional to the temperature gradient). Substituting the constitutive relations into the momentum and energy equations leads to second-order spatial derivatives—the LLNS system is semiparabolic—implying, for example, an infinite speed of heat conduction; physically, however, we expect a finite transport speed. Though LLNS has proven a powerful approach to modeling many nanoflows, we expect the aforementioned assumptions to break down for dense fluids when simulation grid cells approach subnanometer dimensions (i.e., a water molecule) and hydrodynamic timescales of interest are comparable to collision times.

METHODS & CODES

Our HAC method currently uses the LAMMPS MD engine and is being developed in concert with a novel FHD model that extends LLNS. Development is currently focused on our FHD implementation, called HERMESHD (Hyperbolic Equations and Relaxation Model for Extended Systems of HydroDynamics), which is based on Grad's 13-moment (G13) approximation [4]. G13 includes additional equations describing time-dependent transport of viscous stresses (tensorial, symmetric/traceless—five equations) and heat flux (vectorial—three equations). Fluctuating terms can be incorporated either as stochastic fluxes in the momentum and energy equations, or as stochastic sources in the stress and heat flux equations. HERMESHD uses a discontinuous Galerkin (DG) spatial discretization of G13 with stochastic sampling performed on a per-cell basis (i.e., not individual DG Gaussian quadrature points). For temporal integration, we leverage FG13's (fluctuating G13) hyperbolic structure to form a split-level, semi-implicit scheme: explicit time advance is carried out with a second- or third-order strong-stability preserving (SSP) Runge-Kutta method, while a locally implicit relaxation method is used to step over the time step constraint imposed by stiff source terms [5,6].

RESULTS & IMPACT

The additional equations in G13 generate hysteresis in stress and heat flux transport, leading to viscoelastic behavior on small spatiotemporal scales; in fluctuating G13 (FG13), stress and heat flux fluctuations, modeled as white Gaussian noise, are thus also subject to memory effects and give rise to colored Gaussian noise. For nanoscale FHD and HAC simulations of dense fluids like water, we expect FG13 to produce transport phenomena that are neglected by LLNS, such as finite-speed thermoacoustic waves, which may have consequences for protein dynamics, especially as recent experiments have demonstrated a connection between heat released during enzyme catalysis and enhancement of diffusion [7].

As a first test of G13-based FHD, we augmented a linearized version of G13 (L13), with appropriate stochastic terms, to construct a fluctuating L13 (FL13) system. The numerical FL13 model was examined for qualitative correctness in the large collision frequency limit (i.e., solutions should relax to ordinary

LLNS); Fig. 1 compares simulations with and without the fluctuating terms of an unstable nanoscale hydrodynamic jet for viscous, compressible, isothermal flow. The jet was initialized in a periodic domain with small, random velocity perturbations throughout and simulated until instability occurred. As expected, the jet went unstable in the FHD simulation first, though we have yet to quantify differences in the manner of jet breakup or in the overall flow fields between the two models. This comparison will soon be extended to include FL13 results for moderate collision frequency, and a complete comparison with the full FG13 system is being planned.

HERMESHD has a Python-wrapped library interface to facilitate rapid prototyping, Pythonic data manipulation, and simplified communication with external codes. Our HAC method employs a Python-based driver code to couple HERMESHD to the LAMMPS MD engine, though minimizing MD/FHD code interdependence is a priority. We plan to release the HERMESHD code under the GPLv3 license on Github.

WHY BLUE WATERS

We must verify that our G13-based FHD models reduce to LLNS when appropriate limits are taken and also pass common computational hydrodynamic benchmarks. Blue Waters enables these benchmarks—especially high-resolution 3D simulations that would otherwise create long turnaround times—to be executed quickly without hampering the development process. FHD and HAC model validation requires examining their correspondence to numerical experiments using gold-standard atomistic MD, necessitating large microcanonical ensemble simulations in order to avoid spurious effects introduced by thermo-/baro-stating, to mitigate long-range spatial correlations across periodic boundaries, and to obtain adequate statistical sampling. Furthermore, HAC simulations require matching the continuum region to the atomistic domain through empirical constitutive relations and equations of state, which must be calculated numerically from MD simulations of bulk fluids [8]. The multifaceted nature of our project requires running custom and existing codes as well as post-processing and visualizing diverse data, making access to Blue Waters staff and resources essential.

Sean L. Seyler, a fifth-year Ph.D. student in physics at Arizona State University, is working under the direction of Oliver Beckstein. He expects to receive his Ph.D. in December 2017.




Figure 1: Comparison of rightward horizontal velocities for a simple fluid around water density at STP for: Navier-Stokes (top) and fluctuating hydrodynamics (bottom) over 20 nanoseconds. Blue (red) corresponds to small (large) velocities (0–100 nm/ns). Simulation measures 300 nm per side (30² cells) using four-element linear quadrature and second-order Runge-Kutta integration.

THE IMPACTS OF HYDROMETEOR CENTRIFUGING ON TORNADO DYNAMICS

Ronald Stenz, University of North Dakota
2016-2017 Graduate Fellow

EXECUTIVE SUMMARY

Improving the understanding of tornadoes becomes more important as the population in regions prone to these violent weather phenomena increases. This research aims to advance our understanding of tornadoes by making the simulations used to study these destructive and dangerous storms more physically realistic. For the first time, we are quantifying the impacts that centrifuging of precipitation have on the vorticity budgets of these numerically simulated tornadoes. Preliminary findings have removed an unrealistic build-up of precipitation in the vortex center (widely seen in tornado simulations) for both idealized vortices and simulations of an entire storm and the tornado it produces. Ongoing work will examine a large number of tornado simulations to evaluate the significance of the inclusion of precipitation centrifuging in tornado dynamics, as well as to more generally study how a tornado acquires its vorticity, or spin, in different environmental conditions.

RESEARCH CHALLENGE

The primary research challenge being addressed is the lack of precipitation centrifuging in numerical simulations of tornadoes. In current simulations, precipitation follows the air flow, which creates an unrealistic build-up of precipitation in the vortex center. This, in turn, creates a source of negative buoyancy that

potentially limits the stretching of vorticity in these simulated tornadoes. In nature, as precipitation moves around a circulation, such as a tornado, there is no force strong enough to keep the precipitation from moving outward, or being centrifuged, away from the circulation center. Observed tornadoes have a minimum of precipitation in the vortex center, while simulated tornadoes often have a relative maximum of precipitation in the vortex center.

With millions and sometimes billions of dollars of damage caused by tornadoes every year, along with the risk of fatalities or serious injuries, a better understanding of these destructive weather events is needed to improve forecasting, preparedness, and mitigation of their impacts. By including the centrifuging of precipitation into the model we use to learn about tornadoes, our simulations become more consistent with what is observed in nature, facilitating the improvement of our understanding. Research findings have and continue to shape forecasting methods and plans for preparedness and damage mitigation; therefore, continued improvement of our understanding of tornadoes will provide results that can be used in operational settings, ultimately aiding those living in regions prone to tornadoes.

METHODS & CODES

Cloud Model 1 (CM1), which was designed for studying small-scale atmospheric phenomena such as thunderstorms [1], can run

efficiently on supercomputers such as Blue Waters and was used for our simulations in this research. To quantify the impact that the inclusion of centrifuging has on tornado dynamics, simulations were first run without centrifuging. Just prior to the formation of a tornado, a checkpoint is used, allowing the model to be run both with and without centrifuging from this point to determine what impacts the centrifuging of precipitation had on the tornado dynamics. To determine the magnitude of the centrifuging occurring, a centrifuging algorithm based on [2] uses trajectories released within the simulation to calculate the curvature of the flow and ultimately how quickly precipitation will be centrifuged, or moved outward, from the tornadic circulation. To quantify these impacts over a large sample size, atmospheric profiles of temperature, moisture, and wind from atmospheric soundings that were in close proximity to observed supercells [3] were used as the environmental conditions for our simulations of storms and their resulting tornadoes. A subset of these environments known to produce simulated tornadoes in previous research has been used for this study.

RESULTS & IMPACT

Both idealized simulations and a simulation of a full-scale storm with a resulting tornado have been completed with and without centrifuging. In simulations without centrifuging, the unrealistic

maximum of precipitation develops with the vortex core; however, after turning centrifuging on the precipitation in the vortex center is removed and a physically realistic precipitation minimum forms in the vortex center for both the idealized and full-scale tornado simulations. The removal of the unrealistic precipitation in the vortex center is completed within several minutes in both types of simulations. Work is underway to optimize and improve this centrifuging algorithm further. We will then share these findings and, eventually, the centrifuging code to allow future research to benefit from the improved realism of the tornado simulations. Findings from our study on both the importance of centrifuging, and also more general findings about how tornadoes work, have the potential to improve future forecasting and also to facilitate further research into understanding these deadly and destructive storms.

WHY BLUE WATERS

Blue Waters was critical to this project because tornado simulations require thousands of computing cores and produce large amounts of data that must be stored and analyzed. The computing power of Blue Waters along with the available storage for our data was a perfect match for our project. Additionally, the technical and visualization support available with Blue Waters greatly facilitated accomplishment of our research goals.

A fourth-year Ph.D. student in atmospheric sciences at the University of North Dakota, Ronald Stenz is working under the supervision of Matthew Gilmore. He expects to graduate in 2018.

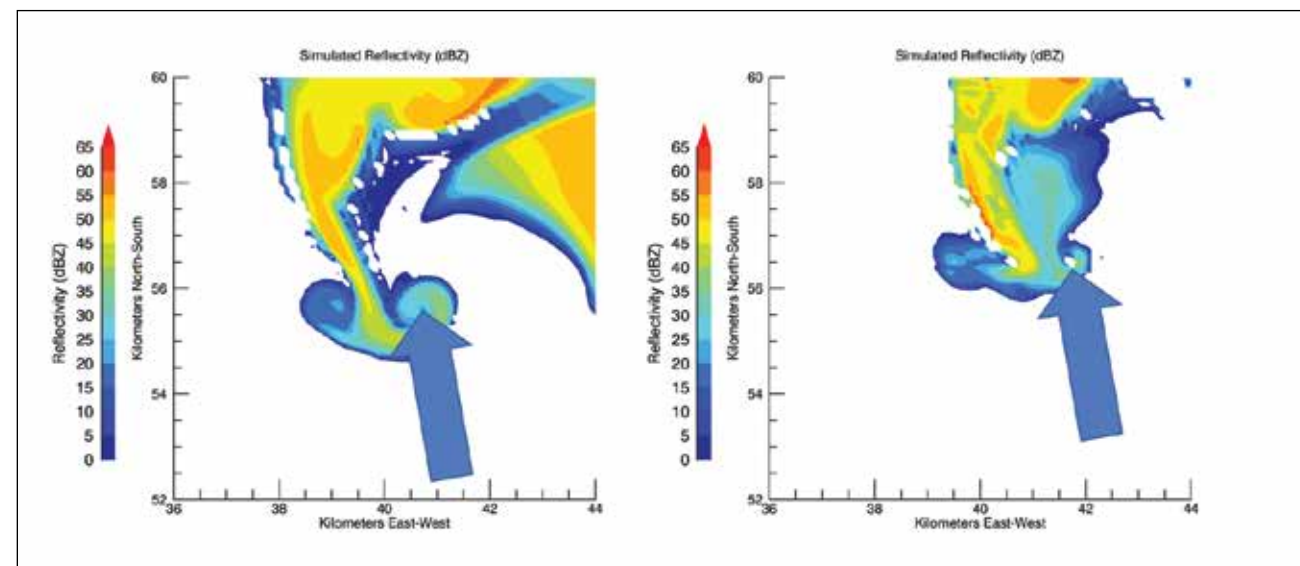


Figure 1: Rain is beginning to accumulate in the simulated tornado center (left). With centrifuging turned on (right), there is an eye-like feature in the simulated tornado center consistent with radar observations of actual tornadoes.

GLASSY DYNAMICS AND IDENTITY CRISES IN HARD-PARTICLE SYSTEMS

Erin Teich, University of Michigan
2016-2017 Graduate Fellow

EXECUTIVE SUMMARY

Glass formation is a well-known outstanding mystery in the physical sciences: It is a phenomenon with no canonical thermodynamic description, yet one with wide-ranging industrial applications, occurring in materials as disparate as superconductors and sand. We utilized our Blue Waters allocation to help shed light on this murky topic, by simulating and investigating the dynamics of several glass-forming soft matter systems on time scales spanning about six orders of magnitude. We used our open-source simulation toolkit HOOMD-blue [1], scalable software built to run in parallel using domain decomposition on both CPUs and GPUs. We found that all systems studied displayed dynamical signatures indicative of cooperative rearrangement and consequent dynamic heterogeneity, both hallmarks of glass-forming behavior [2,3]. We examined local structure in these systems on other machines, and found a link between glass formation and a structural identity crisis, or competition between various local structural motifs.

RESEARCH CHALLENGE

Soft matter systems, in which thermal fluctuations are strong enough to drive particle rearrangements, are capable of self-assembling into a staggering variety of simple and complex crystalline structures. Often, however, no such assembly occurs, and the system remains disordered, displaying instead dynamical signatures characteristic of glass-forming materials. To develop robust methods of self-assembly, then, an understanding of glass formation in soft matter systems is needed. Beyond the realm of soft matter, glasses have proven useful materials to humans for over four millennia, for applications ranging from dinnerware to phase-change memory devices [3].

Despite this ubiquity, a clear picture of the thermodynamics and structural change undergirding glass formation remains unknown. This is due in large part to the significant slowing down of any system as it approaches the glass transition [2,3], requiring investigations of glass formation to resolve, dauntingly, system dynamics on time scales that vary by orders of magnitude. We were able to tackle this problem using our highly parallel simulation package, HOOMD-blue, and by accessing both the Cray XK7 nodes hosted on Blue Waters and its very generous storage capabilities.

METHODS & CODES

To help shed light on the dynamics and structural underpinnings of glass formation in soft matter systems, we chose to simulate and

analyze nonassembling monatomic systems of hard polyhedral particles, with no interactions aside from those of excluded volume. We chose particle shapes consisting of tetrahedra with varying degrees of edge and vertex truncation. These are simple systems mediated solely by entropy maximization. Entropy maximization is known to drive phase transitions in a variety of systems, from those composed of hard spheres [4] to hard rods [5] to colloids and polymer depletants [6]. Recently, it was shown that entropy maximization drives systems of hard polyhedral particles toward a vast array of ordered structures upon compression [7], and a mechanism involving emergent, directional entropic forces was proposed to explain this ordering [8]. In [7], however, a number of systems displayed no order at all, instead exhibiting the signature slow dynamics of glassiness. In these cases, directional entropic forces drove systems toward glassy disorder rather than crystallization. Our project further explores these entropic effects.

We used the hard particle Monte Carlo (HPMC) simulation method [9] included in HOOMD-blue to sample our systems in the isochoric ensemble in equilibrium. In brief, this classic simulation method proposes trial moves for particles, and accepts or rejects them based on an acceptance criterion that dictates rejection if there are any overlaps created in the system, and acceptance otherwise. We simulated systems of 4096 particles on a single GPU at a variety of densities, and subsequently measured structural and dynamical information. To gather trajectory information at a wide range of time scales for later analysis, we wrote our system trajectories to disk remarkably frequently—every 10 Monte Carlo (MC) sweeps—and collected data for the far longer period of 100 million MC sweeps. We produced trajectories as large as 3.25 terabyte (TB) MC sweeps per simulation, each containing about 10 million simulation frames.

RESULTS & IMPACT

For each trajectory, we measured several dynamical parameters known in the glass community to be indicators of glass-forming behavior, and confirmed that our hard particle systems are indeed glass-formers. These were the mean-squared displacement of all particles in the system, the van Hove autocorrelation function, the non-Gaussian parameter characterizing the particle displacement distribution, and the self-part of the intermediate scattering function [2,3]. The mean-squared displacement and self-intermediate scattering function for a sample system are shown in Fig. 1. Plateaus in both parameters at intermediate times and increasing densities indicate caging behavior typical of glass-formers; at these times, particles are trapped on average




Figure 1: A glass-forming system of hard polyhedral particles. (A) A system snapshot at a volume density of 0.62. (B) Both the mean-squared displacement and the self-intermediate scattering function of all particles in the system over time, for a variety of densities. Plateaus in both parameters are indicative of caging.

by a cage of their surrounding neighbors, and only escape due to cooperative structural rearrangements and consequent dynamic heterogeneity. The mean-squared displacement shows this caged behavior directly, while the plateau in the self-intermediate scattering function (a Fourier transform of the real-space van Hove autocorrelation function) shows a slowing-down of the relaxation of the system.

We have determined via simulations on other platforms that observed glass-forming tendencies in these systems are related to competition between local structural motifs, mediated by entropic forces; future work will strive to make more concrete the connection between local structure and observed heterogeneous dynamics.

WHY BLUE WATERS

HOOMD-blue HPMC's parallel nature on both CPUs and GPUs allowed us to take unique advantage of the computational resources on Blue Waters, and, in particular, its GPU resources. We also produced tens of TB of raw data that we needed to sift through later, and we were able to accomplish this by storing it all on Blue Waters' generous Online file system, analyzing it, and storing resultant analysis results somewhere more permanently. Without the 500 TB quota afforded to us by Blue Waters, we would find it far more difficult to explore such fine-grained dynamical phenomena.

Erin Teich is a fifth-year Ph.D. student in applied physics at the University of Michigan. She is working under the direction of Sharon C. Glotzer, and expects to graduate in June 2018.

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

Samuel Totorica, Stanford University
2016-2017 Graduate Fellow

EXECUTIVE SUMMARY

In this project, we use the state-of-the-art particle-in-cell code OSIRIS to study a fundamental plasma process known as magnetic reconnection, which plays a key role in the evolution of plasmas from astrophysics to the laboratory. Reconnection is a promising candidate for producing the energetic particle distributions associated with explosive astrophysical sources; however, the particle acceleration properties of reconnection are not fully understood. Recently, laser-driven plasma experiments have been used to study reconnection in conditions relevant for astrophysics. By modeling these experiments on Blue Waters we were able to show that for current experimental conditions, electrons can be accelerated by reconnection with sufficient quantity and energy to be detected in the laboratory, which opens the way for new

experimental studies of particle acceleration from reconnection. We are also working on developing a new simulation method called simplex-in-cell, which may improve the accuracy and reduce the expense of plasma simulations.

RESEARCH CHALLENGE

Magnetic reconnection is a fundamental plasma process that converts magnetic field energy into plasma kinetic energy through the breaking and rearrangement of magnetic field lines [1]. It is believed to play a key role in frontier problems in physics including the origin of cosmic rays, and is relevant for applications with societal benefit such as space weather and nuclear fusion energy. In astrophysics, reconnection is currently being studied intensely as a promising candidate for producing the energetic particle distributions associated with explosive astrophysical sources such as gamma-ray bursts and jets from active galactic nuclei. However, the efficiency of reconnection in accelerating nonthermal particles, and how this depends on the plasma conditions, remains poorly understood. It is currently an active area of research to determine whether reconnection can account for the astrophysical observations. As a result of the inertial confinement fusion program, high-energy laser facilities have been developed that can produce extremely hot and dense plasmas that reach a regime where scaling laws allow comparisons with astrophysical systems. The goal of this project is to use simulations to study particles from reconnection in varied plasma conditions, and in particular to investigate whether laser-driven plasma experiments could be used to study the particle acceleration properties of reconnection in the laboratory.

METHODS & CODES

One of the most powerful tools for *ab initio* plasma simulation is the particle-in-cell (PIC) method, which treats the plasma as a collection of discrete simulation particles that interact via self-consistent electromagnetic forces. The simulations for this project were run using the state-of-the-art, massively parallel, and fully relativistic PIC code OSIRIS [2] and match the experimental conditions produced by the most energetic laser systems in the world, such as the National Ignition Facility.

RESULTS & IMPACT

From the results of these simulations we were able to show clearly that for current experimental conditions, electrons can be accelerated by reconnection with sufficient quantity and energy to be detected in the laboratory. For the conditions of recent experiments, the nonthermal electrons can be accelerated

to energies more than an order of magnitude larger than the initial thermal energy. The nonthermal electrons are primarily accelerated by the reconnection electric field near the X-points, which establishes a distribution of energies that resembles a power-law spectrum. After being energized, the electrons can also become trapped inside the plasmoids (magnetic islands) that form in the current layer and gain additional energy from the electric field arising from the motion of the plasmoid. By comparing simulations for finite and infinite periodic systems, we were able to demonstrate the importance of particle escape on the shape of the spectrum.

Based on our findings, we derived an analytical estimate of the maximum electron energy and a threshold condition for observing suprathermal electron acceleration in terms of the initial plasma conditions, which can now be tuned in future experiments to optimize the particle acceleration. Through the use of 3D simulations (Fig. 1) we studied the angular distribution of the accelerated particles and constructed synthetic detector spectra to determine experimental signatures. These results provide new insight into the physics of reconnection and particle acceleration, and are now helping to guide several experimental programs in the United States.

Due to limitations such as noise from artificial two-body collisions and the computational expense associated with the large number of particles required to accurately capture the development of nonthermal tails in the particle distribution, multiscale PIC simulations such as those used to study laser-driven reconnection are extremely challenging. It is thus critical to work on improved methods that could reduce the computational expense of these simulations and improve their physical accuracy. To this end, we are also developing a novel method for plasma simulation, which we refer to as simplex-in-cell (SIC). The foundation of SIC is an interpretation of the simulation particles as the vertices of an unstructured mesh that traces the evolution of the plasma distribution function in phase space [3]. This enables a new discretization using deformable phase space volume elements rather than fixed-shape, localized particles. We are using the SIC interpretation of the simulation particles for data analysis and visualization of standard PIC simulations performed using OSIRIS, and have been able to show that in certain regimes SIC can reach a given noise level using one thousand times fewer simulation particles than standard methods (Fig. 2). Future work will involve implementing SIC directly in the simulations to reduce noise and unphysical artifacts.




Figure 2: Momentum space for simulations of the Weibel instability with 16 particles-per-cell (PPC) (top) and 65,536 PPC (bottom). The SIC (simplex in cell) deposit (right) captures the tails and shape of the distribution at low PPC much better than a high-order (quartic) particle deposit (left).

WHY BLUE WATERS

This project required the use of large-scale 2D and 3D simulations with sufficient size and resolution to bridge the multiscale physics, from fluid dynamics to the kinetic microscopic processes. These computationally demanding simulations can require billions of simulation particles, and demand the cores, memory, and communication performance available on Blue Waters. The quick support from the NCSA staff on technical issues helped me to maximize my productivity on the machine.

PUBLICATIONS AND DATA SETS

Totorica, S. R., T. Abel, and F. Fiuza, Nonthermal electron energization from magnetic reconnection in laser-driven plasmas. *Physical Review Letters*, 116 (2016), DOI:10.1103/PhysRevLett.116.095003.

Totorica, S. R., T. Abel, and F. Fiuza, Particle acceleration in laser-driven magnetic reconnection. *Physics of Plasmas*, 24 (2017), DOI:10.1063/1.4978627.

Samuel Totorica, a fifth-year Ph.D. student in physics at Stanford University, is working under the direction of Tom Abel and Federico Fiuza. He expects to graduate in 2018.




Figure 1: Cover of the April 2017 edition of *Physics of Plasmas*, showing a 3D visualization of a simulation of laser-driven magnetic reconnection. The trajectories of the energetic electrons are colored by their energy and plotted over an isocontour of the magnetic field (gray).

TRANSFORMATIVE SCIENCE



The Blue Waters project provides the national open-source research community with the computational power and data analysis capabilities that have become increasingly vital to virtually all areas of inquiry and discovery. With the power of integrated computing and the data analysis capabilities of the Blue Waters system, researchers are pushing the boundaries of their disciplines and achieving breakthrough results on a wide range of transformative problems, including the fields of molecular bioscience, astronomy and astrophysics, geoscience, social science, and materials research. Fig. 1 shows the diverse disciplines using the Blue Waters system from June 1, 2016–May 31, 2017 (Project Year 4–PY4). In PY4, 1,047 researchers on 352 projects ran over 1.6 million jobs on Blue Waters, consuming over five billion core hour equivalents. Indeed, science teams in PY4 utilized 8.5 percent more time than in PY3, easily making it the most productive year since Blue Waters entered service.

Because Blue Waters has different types of nodes (XE and XK), we use the term “core hour equivalent” to help compare node

hours to core hours, which may be more familiar to some. There are two ways to use the AMD Interlagos process; one mode is as 16 floating-point cores, each with one integer core. The other mode is as 32 integer cores sharing 16 floating-point cores. We define a “core hour equivalent” to be the number of node hours multiplied by 32. Using this relationship, the median job size for the past year, based on node hours consumed, utilized 256 nodes, or over 8,000 core equivalents. This is a lower median than the past years primarily because, as we move to support a workload with more data analytics, and we have improved the ability to increase node utilization with backfilling jobs while not delaying any large or very large jobs, the median job size decreases without perturbing service to the very scalable jobs. In addition to the compute power of Blue Waters there is an incredible data analysis capability evidenced by the ability of science teams to read and write in excess of 5 petabytes of data in 24-hour periods multiple times during the year.




Figure 1: Distribution of annual actual usage by discipline June 2016–May 2017

From the beginning of the project, the primary goal of Blue Waters has been to enable transformative science and engineering research by giving research teams access to a large number of powerful compute nodes and large data capacity. The target is to have at least 40 percent of the system’s capacity dedicated to research challenges that cannot be addressed on smaller-scale systems. These “capability” jobs are defined as those that are either *large*, using 512 to 4,528 XE nodes (i.e., 16,384 to 144,896 core equivalents) or 64 to 845 XK nodes, or *very large*, using more than 4,528 XE nodes (i.e., greater than 144,896 core equivalents), or more than 845 XK nodes. Large jobs are equal to the largest computations possible on many other compute resources, and very few systems in the world other than Blue Waters can support very large jobs. The project generally surpassed this 40 percent benchmark over the past year, though the average is lower than previous years due to more data-focused, “big data” single node workloads being added to the mix of projects awarded time on Blue Waters. Blue Waters continues to prioritize the “capability” jobs that allow researchers to push the boundaries of their disciplines and to achieve results that would not be possible without Blue Waters (see Table 2). Several science teams are regularly running jobs that use 2,048; 4,096; and 8,192 nodes on Blue Waters.

Many researchers who use Blue Waters have told us that their work would simply not be possible without this Leadership-Class resource:

“The simulations required to conduct our research simply cannot be accomplished in a timely manner on less powerful machines. Without Blue Waters, iNEMO research is, at best, hampered, and, at worst, impossible. Blue Waters allows us to get results in a day, rather than weeks as on other computer systems.”

Gerhard Klimeck, Purdue University

“Without Blue Waters...neither the HIV nor the chromatophore project would be possible. ...These projects are examples of how Blue Waters enables bold new projects that push the limits of what can be done with scientific computing.”

The late Klaus Schulten, University of Illinois at Urbana-Champaign

“On Blue Waters, what would be a tremendous calculation extending over months or an entire year on another system took a week. Blue Waters entirely changed our concept of what is practical and what important research is sensible to undertake.”

Paul Woodward, University of Minnesota

Support from the Blue Waters team is essential in enabling research partners to perform extreme-scale research and to achieve breakthrough results. Key features of the Blue Waters partner support include:

- Traditional Partner Consulting. Standard requests for assistance with porting, debugging, allocation issues, and software requests are handled through the Jira service request ticket system at help+bw@ncsa.illinois.edu.
- Advanced Application Support. Requests beyond the traditional consulting can be made via the service request ticket system. These requests are reviewed and evaluated by the Blue Waters Project Office (BWPO) for breadth, reach, and impact for the project and the community at large. Generally they are fulfilled once a work plan has been agreed upon with the research team.
- Major research teams (for example, those receiving allocations through the NSF PRAC program) are assigned an individual Point of Contact (PoC) from the Blue Waters Science and Engineering Application Support (SEAS) group for the duration of their project time. Initial contact with a science team starts with a questionnaire, followed by meetings to gather information on project goals, approaches, status, challenges, requirements, and concerns. The assigned PoC works with the science team on issues such as tuning, modeling, I/O, and optimizing application codes as well as requests for special resources or priorities.
- Advanced Application Support activities for these major science teams have PoCs participate in code restructuring, re-engineering, or redesign, such as implementing GPU functionality via OpenACC, or alternatives to MPI collective operations. Such work is tracked via a coordinated work plan that is developed between the PoC and the science team to clearly indicate the scope and scale of the work involved, including milestones and deliverables. Work plans are reviewed and approved by the BWPO.

METRIC	TARGET	6/1/2016-8/31/2016	9/1/2016-11/30/2016	12/1/2016-2/29/2017	3/1/2017-5/31/2017	PY3 6/1/2016-5/31/2017
The majority of computational time goes to capability jobs	40% of all computational time will go to capability jobs	51.3% Exceeds Expectations	45.8% Meets Expectations	38.3% Below Expectations	47.6% Meets Expectations	45.7% Meets Expectations

Table 1: Blue Waters exceeded its metric in regard to enabling large and very large “capability” jobs.

• Support for workflows, data movement, and visualization are provided by the appropriate Blue Waters support teams and can be requested by sending email to help+bw@ncsa.illinois.edu or entering a Service Request.

The leadership of the Blue Waters project monitors the time required to respond to and resolve service requests, which include traditional “trouble tickets” as well as requests for basic issues, expanded assistance, added services, and suggestions. Table 3 demonstrates how Blue Waters staff met or exceeded expectations for key service metrics over the past year.

Many users have mentioned the crucial role of Blue Waters support staff in enabling their work:

“One key aspect of the Blue Waters facility was the availability of the Blue Waters staff, who provided invaluable technical insight and frequent advice in order to tune the performance of our software.”

Greg Voth, University of Chicago

“The Blue Waters Fellowship has opened new doors for me by providing me first-hand experience with the high-performance computing resources necessary for cutting-edge science in oceanography. The expert staff at NCSA have been extremely responsive to my questions and needs and have helped me grow from a novice into a proficient user with plans to continue using high-performance computing to study physical–biological interactions across marine ecosystems.”

Maureen Brooks, Blue Waters Graduate Fellow, University of Maryland College Park

The Petascale Application Improvement Discovery (PAID) program is another way the Blue Waters project supports advanced application development and improvement work. PAID forges collaborations between Improvement Method Enablers (IMEs) and science and engineering research teams, helping them to create and implement new technologies that improve application performance. (A detailed report on the PAID program can be found elsewhere in this volume.)

Additionally, the annual Blue Waters Symposium provides an excellent opportunity for cross-pollination among research teams and across disciplines. Attendees report that the facilitation of personal networking, interaction, and discussion is a particular strength of the symposium and that the event provides useful resources and information.

The Blue Waters project also strives to develop the next generation of computational researchers by supporting education and training activities for undergraduate students, graduate students, and professional staff. Each year, an elite group of Ph.D. students from a diverse range of domains and from institutions across the country are awarded Blue Waters Graduate Fellowships. These competitive fellowships include both financial support and

access to the Blue Waters system to substantially advance the fellows’ computational research. To date, 36 young scholars have participated in the program, which over three years will award more than \$1.3 million in direct financial support and over 60 million core equivalent hours to support graduate research. The 36 fellows are from 24 states and 26 institutions, including eight institutions from EPSCoR (Established Program to Stimulate Competitive Research) states and one that is a minority-serving institution.

In order to ensure that the fellows are able to effectively utilize the power of the Blue Waters system, each fellow is paired with a member of the Blue Waters SEAS group; these “points of contact” assist the fellows with essential activities such as porting and tuning codes and with any support requests. In the December 2015 review of the Blue Waters project by the National Science Foundation, the review panel pointed to this in-depth PoC model as a “best practice” for fellowship and education programs.

Some of the impressive research results achieved by the 2016–2017 Blue Waters Graduate Fellows are showcased in this annual report. Many fellows highlighted the impact that Blue Waters and their fellowship experience will have on their research and their future career path:

“The Blue Waters Graduate Fellowship has not only played a pivotal role in helping me achieve my doctorate, but it has also exposed me to a variety of fields and topics related to high-performance computing that I would have otherwise not experienced as part of my degree plan. I believe this breadth of knowledge will give me a competitive edge as I continue to the next stages of my career.”

Justin Drake, University of Texas Medical Branch

Approximately 20 students are selected each year for the Blue Waters Undergraduate Internship program, which includes an intensive two-week workshop covering the fundamentals of high-performance computing (HPC) and engages undergraduate students in year-long projects in which they apply HPC to problems in science, mathematics, and engineering with guidance from faculty mentors. To date, the Blue Waters project has funded 120 interns from 64 unique institutions in 27 states, the District of Columbia, and Puerto Rico. Of the 99 interns, 52 are women and/or minorities. Thirty-five of the 69 institutions were in EPSCoR regions, and 17 were classified as minority-serving institutions.

Blue Waters has directly benefitted people in 49 states (Hawaii is the outlier) as well as Puerto Rico. Blue Waters activities have also benefited individuals in more than 20 countries working as part of projects led by principal investigators in the United States.

External evaluation has found that both interns and their advisors find the program beneficial. Some participants have reported altering their educational and career paths or embarking on a path with greater confidence thanks to the internship experience.

METRIC	TARGET	6/1/2016-8/31/2016	9/1/2016-11/30/2016	12/1/2016-2/29/2017	3/1/2017-5/31/2017	PY3 6/1/2016-5/31/2017
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	96.6% of partner service tickets had a human response within four business hours Meets Expectations	97.8% of partner service tickets had a human response within four business hours Exceeds Expectations	97.3% of partner service tickets had a human response within four business hours Meets Expectations	98.2% of partner service tickets had a human response within four business hours Exceeds Expectations	97.4% 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)	83.6% of partner service requests were resolved within three business days Meets Expectations	82.4% of partner service requests were resolved within three business days Meets Expectations	85.1% of partner service requests were resolved within three business days Exceeds Expectations	81.5% of partner service requests were resolved within three business days Meets Expectations	83.2% of partner service requests were resolved within three business days Meets Expectations

Table 2: Metrics for service request response and resolution time show the Blue Waters support staff are meeting and often exceeding the metric expectations throughout the year.

Another aspect of the Blue Waters engagement with students and early career workers is the fact that from the beginning of the project, 56.4 percent of the people using Blue Waters were students and/or postdocs. This includes 881 graduate students, 125 undergraduate students, 17 high school students, and 350 postdocs.

The project also funds, in partnership with the XSEDE project, the content in HPC University (<http://hpcuniversity.org>), which is focused on developing educational materials for and about petascale computing and analysis. This year, the HPC University had 67,533 petascale learning modules and training materials downloaded.

In partnership with the University of Buffalo, Blue Waters published a extensive study of the work that is using Blue Waters over the first three and a half years of service. This study shows many of the characteristics of how the applications use the different architectural features of Blue Waters such as the amount

of memory used and the I/O characteristics of the workload. The full report is able to provide much more detail and is available at <https://arxiv.org/ftp/arxiv/papers/1703/1703.00924.pdf>

Because of excellent overall project management, the partners, leadership and staff of the Blue Waters project is pleased that the National Science Foundation has approved continued operations of the Blue Waters project through at least March 2019. With this one-year-plus extension of full-service operations, the project will continue to serve the research needs of the scientific community as well as undergraduate and graduate education activities. The dedicated Blue Waters team at NCSA looks forward to continuing productive collaborations with our science and engineering partners as they continue to innovate and discover.

SCIENCE AND ENGINEERING TEAM ADVISORY COMMITTEE

The Science and Engineering Team Advisory Committee (SETAC) brings together a diverse group of scientists and engineers who represent the various Blue Waters the science and engineering research teams using Blue Waters and the breadth of research powered by the Blue Waters system. The committee provides guidance and assessment to help the Blue Waters project deliver the best possible performance and services that will in turn assist research teams in achieving breakthrough results.

The SETAC makes recommendations on technical directions, strategies, and management while identifying potential challenges for petascale applications. As users themselves, the SETAC members also provide advice for solving common issues that arise from moving applications to Blue Waters and from system software at scale.

The SETAC convenes three to four times per year and is available to the Blue Waters Project to provide guidance and advice as needed throughout the year. Members are nominated and typically serve two-year terms.

SETAC Members

PETASCALE COMPUTING RESOURCE TEAMS

- **Paul Woodward**, Physics and Astrophysics, University of Minnesota (SETAC Chair)
- **Susan Bates**, Oceanography & Climate Dynamics, NCAR
- **Manuela Campanelli**, Physics and Astronomy, Rochester Institute of Technology
- **Tom Cheatham**, Chemistry, University of Utah
- **David Ceperley**, Physics and Material Science, University of Illinois at Urbana-Champaign
- **Tiziana Di Matteo**, Physics and Cosmology, Carnegie Mellon University
- **Said Elghobashi**, Mechanical and Aerospace Engineering, University of California, Irvine
- **Tom Jordan and Philip J. Maechling**, Seismology, University of Southern California
- **Paul Morin**, Earth Sciences, University of Minnesota
- **Brian O'Shea**, Physics and Astronomy, Michigan State University
- **Nikolai Pogorelov**, Space Science, University of Alabama, Huntsville

GREAT LAKES CONSORTIUM FOR PETASCALE COMPUTATION TEAMS

- **Ben Rogers**, Information Technology Research Services, University of Iowa

INDUSTRY TEAMS

- **Rick Arthur**, General Electric Global Research, Advanced Computing

UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN TEAMS

- **Athol Kemball**, Astronomy, University of Illinois at Urbana-Champaign

OTHER BLUE WATERS PROJECTS

The projects listed here had a Blue Waters allocation during this reporting period but did not submit a report for the project.

QC HathiTrust Digital Text Archive
Scott Althaus

Spot Scanning Proton Therapy Optimiz.
Chris Beltran

EpiSimdemics Simulations of Influenza
Abhinav Bhatle

SC16 SciNet Cosmology Experiment
Franck Cappello

Simulations of Large-N Star Clusters
Sourav Chatterjee

HathiTrust Public Research Dataset
J. Stephen Downie

Reservoir Simulation on Blue Waters
Vadim Dyadechko

Mechanics of Biofilms
Ahmed Elbanna

Co-Evolution of Galaxies and Black Holes
Claude-Andre Faucher-Giguere

Simulation and Scaling of MultiPhysics
Jonathan Freund

Scalable Social Simulations
Leslie Gasser

Protein Sequence Similarity Networks
John Gerlt

Forecasting Volcanic Activity
Patricia Gregg

Global crop forecasting from satellite
Kaiyu Guan

Cloud Computing Systems at Scale
Indranil Gupta

Tilted Disks Around Black Holes
John Hawley

Image Uncertainty Quantification
Athol Kemball

Paracellular Transport Mechanism
Fateme Khalili-Araghi

Improving Performance of an Implicit FEA
Seid Koric

Blue Waters Workload Analysis
William T.C. Kramer

Ecohydrologic dynamics-climate change
Praveen Kumar

LES of a Film Cooled Turbine Vane
Gregory Michael Laskowski

Petascale Simulation of Two-Phase Flows
Vincent Le Chenadec

Numerical simulation of wing-body-nacell
Jong Lee

Numerical study of many body transition
David Luitz

Genotyping for very large cohorts
Liudmila Mainzer

Enabling Discoveries at the LHC
Mark Neubauer

Deep learning for structure prediction
Jian Peng

Solvers for Turbulent Combustion
Venkat Raman

AGN Feedback in Galaxy Clusters
Paul Ricker

Evaluating Theft in Smart Grids
William Sanders

High-Speed Link Simulation
Jose Schutt-Aine

Large-scale Deep Learning
Justin Sirignano

Atmospheric Response to a Wildland Fire
Marshall Stageberg

Variability of halo properties in Enzo
Victoria Stodden

The Computational Microscope
Emad Tajkhorshid

Chemistry in the Early Universe
Matthew Turk

REFERENCES

SPACE SCIENCE

Balsara, Dinshaw

- [1] Tilley, D.A. and D.S. Balsara, Direct Evidence for Two-Fluid Effects in Molecular Clouds. *Monthly Notices of the Royal Astronomical Society*, 406 (2010), pp. 1201–1207.
- [2] Meyer, C., D.S. Balsara, B. Burkhardt, and A. Lazarian, Observational Diagnostics for Two-Fluid Turbulence in Molecular Clouds as Suggested by Simulations. *Monthly Notices of the Royal Astronomical Society*, 439 (2014), pp. 2197–2210.
- [3] Burkhardt, B., et al., Alfvénic Turbulence Beyond the Ambipolar Diffusion Scale. *The Astrophysical Journal*, 805 (2015), pp. 118–126.
- [4] Li, H.-B. and M. Houde, Probing the Turbulence Dissipation Range and Magnetic Field Strengths in Molecular Clouds. *The Astrophysical Journal*, 677 (2008), p. 1151.
- [5] Balsara, D.S., Wave Propagation in Molecular Clouds. *The Astrophysical Journal*, 465 (1996), pp. 775–794.
- [6] Tilley, D.A. and D.S. Balsara, A Two-fluid Method for Ambipolar Diffusion. *Monthly Notices of the Royal Astronomical Society*, 389 (2008), p. 1058.
- [7] Tilley, D.A. and D.S. Balsara, Two-Fluid Ambipolar Diffusion for Molecular Clouds with Realistic Heating and Cooling. *Monthly Notices of the Royal Astronomical Society*, 415 (2011), pp. 3681–3692.
- [8] Tilley, D.A., D.S. Balsara, and C. Meyer, A Numerical Scheme and Benchmark Tests for Two-fluid Ambipolar Diffusion. *New Astronomy*, 17 (2012), pp. 368–376.

Holder, Gilbert

- [1] Omori, Y., et al., A 2500 square-degree CMB lensing map from combined South Pole Telescope and Planck data. *The Astrophysical Journal*, submitted May 2017, arXiv:1705.00743.
- [2] Hezaveh, Y., et al., Detection of Lensing Substructure Using ALMA Observations of the Dusty Galaxy SDP.81. *The Astrophysical Journal*, 823:37 (2016), DOI:10.3847/0004-637X/823/1/37.

Huerta, Eliu (1)

- [1] Huerta, E.A., et al., Complete waveform model for compact binaries on eccentric orbits. *Physical Review D*, 95 (2017), DOI:10.1103/PhysRevD.95.024038.
- [2] Harry, I., et al., Searching for gravitational waves from compact binaries with precessing spins. *Physical Review D*, 94 (2016), DOI:10.1103/PhysRevD.94.024012.

- [3] Veitch, J., et al., The formation of eccentric compact binary inspirals and the role of gravitational wave emission in binary-single stellar encounters. *The Astrophysical Journal*, 784:71 (2014), DOI:10.1088/0004-637X/784/1/71.
- [4] Blackman, J., et al., Fast and accurate prediction of numerical relativity waveforms from binary black hole coalescences using surrogate models. *Physical Review Letters*, 115 (2015), DOI:10.1103/PhysRevLett.115.121102.
- [5] Vaibhav, T., et al., A Proposed Search for the Detection of Gravitational Waves from Eccentric Binary Black Holes. *Physical Review D*, 93 (2016), DOI:10.1103/PhysRevD.93.043007.

Huerta, Eliu (2)

- [1] Abbott, B.P., et al., Observation of Gravitational Waves from a Binary Black Hole Merger. *Physical Review Letters*, 116 (2016), DOI: 10.1103/PhysRevLett.116.061102.
- [2] Abbott, B.P., et al., GW151226: Observation of Gravitational Waves from a 22-Solar-Mass Binary Black Hole Coalescence. *Physical Review Letters*, 116 (2016), DOI: 10.1103/PhysRevLett.116.241103.
- [3] Samsing, J., et al., The formation of eccentric compact binary inspirals and the role of gravitational wave emission in binary-single stellar encounters. *The Astrophysical Journal*, 784:71 (2014), DOI: 10.1088/0004-637X/784/1/71.
- [4] Strader, J., et al., Two stellar mass black holes in the globular cluster M22. *Nature*, 490:71 (2012), DOI: 10.1038/nature11490.
- [5] Antonini, F., et al., Black hole mergers and blue stragglers from hierarchical triples formed in globular clusters. *The Astrophysical Journal*, 816 (2016), DOI: 10.3847/0004-637X/816/2/65.
- [6] Miller, M.C., Implications of the Gravitational Wave Event GW150914. *General Relativity and Gravitation*, 48 (2016), DOI: 10.1007/s10714-016-2088-4.
- [7] Lehner, L., and F. Pretorius, Numerical Relativity and Astrophysics. *Annual Review of Astronomy and Astrophysics*, 52 (2014), pp. 661–694.

Lentz, Eric

- [1] Bruenn, S. W., et al., The Development of Explosions in Axisymmetric *Ab Initio* Core-collapse Supernova Simulations of 12–25 M_{\odot} Stars. *Astrophys. J.*, 818 (2016), p. 123.
- [2] Bruenn, S. W., et al., Chimera: A massively parallel code for core-collapse supernova simulation, in preparation.
- [3] Cassanova, J., et al., Resolution effects in high resolution core-collapse supernova models, in preparation.
- [4] Radice, D., et al. Neutrino-driven Convection in Core-collapse Supernovae: High-resolution Simulations. *Astrophys. J.*, 820 (2016), p. 76.
- [5] Lentz, E. J., et al., Three-dimensional Core-collapse Supernova Simulated Using a 15 M_{\odot} Progenitor. *Astrophys. J. Lett.*, 807 (2015), p.31.

- [6] Melson, T., et al., Neutrino-driven Explosion of a 20 Solar-mass Star in Three Dimensions Enabled by Strange-quark Contributions to Neutrino-Nucleon Scattering. *Astrophys. J. Lett.*, 808 (2015), p. 42.
- [7] Lentz, E. J., et al., Explosion and ejecta of a light iron-core supernova in three-dimensions, in preparation.
- [8] Wanajo, S., et al., Nucleosynthesis in the Innermost Ejecta of Neutrino-Drive Supernova Explosions in Two Dimensions. *Astrophys. J.* (submitted).

Levin, Deborah

- [1] Jambunathan, R., and D. A. Levin, Advanced Parallelization Strategies Using Hybrid MPI-CUDA Octree DSMC Method for Modeling Flow Through Porous Media. *Computers and Fluids*, 149 (2017), DOI: 10.1016/j.compfluid.2017.02.020.
- [2] Jambunathan, R., and D. A. Levin, Modeling of Ion Thruster Plasma Plume using CHAOS PIC-DSMC solver with multiple GPUs. 44th International Conference on Plasma Science (IEEE Computer Society, Atlantic City, N.J., May 21–25, 2017).

Noble, Scott

- [1] Noble, S., et al., Circumbinary Magnetohydrodynamic Accretion into Inspiring Binary Black Holes. *Astrophysical Journal*, 755:1 (2012), p. 51.
- [2] Zilhao, M., S. Noble, M. Campanelli, and Y. Zlochower, Resolving the relative influence of strong field spacetime dynamics and MHD on circumbinary disk physics. *Physical Review D*, 91:2 (2015), 024034.
- [3] Bowen, D., et al., Relativistic Dynamics and Mass Exchange in Inspiring Binary Black Hole Mini-Disks. *Astrophysical Journal*, 838 (2017), p. 42.
- [4] Noble, S., J. Krolik, and J. Hawley, Direct calculation of the radiative efficiency of an accretion disk around a black hole. *Astrophysical Journal*, 692 (2009), pp. 411–421.
- [5] Zilhao, M., and S. Noble, Dynamic fisheye grids for binary black hole simulations. *Classical Quantum Gravity*, 31:6 (2014), 065013.
- [6] Shiokawa, H., R. Cheng, S. Noble, and J. Krolik, PATCHWORK: A Multipatch Infrastructure for Multiphysics/Multiscale/Multiframe Fluid Simulations. *Astrophysical Journal*, arXiv:170105610, under review (2017).

O'Shea, Brian

- [1] Bryan, G. et al., Enzo: An Adaptive Mesh Refinement Code For Astrophysics. *The Astrophysical Journal*, 211:2 (2014), DOI: 10.1088/0067-0049/211/2/19.
- [2] Turk, M., et al., yt: A Multi-code Analysis Toolkit for Astrophysical Simulation Data. *The Astrophysical Journal Supplements*, 192:1 (2011), DOI: 10.1088/0067-0049/192/1/9.

Pogorelov, Nikolai

- [1] Burlaga, L.F., and N.F. Ness, Observations of the Interstellar Magnetic Field in the Outer Heliosheath: Voyager 1. *Astrophys. J. Lett.*, 829:134 (2016), DOI: 10.3847/0004-637X/829/2/134.
- [2] Gurnett, D.A., et al., Precursors to Interstellar Shocks of Solar Origin. *Astrophys. J.*, 809:121 (2015), DOI: 10.1088/0004-637X/809/2/121.
- [3] McComas, D.J., et al., Seven Years of Imaging the Global Heliosphere with IBEX. *Astrophys. J. Suppl.*, 229:41 (2017), DOI: 10.3847/1538-4365/aa66d8.

Quinn, Thomas

- [1] Kale, L. V., and S. Krishnan, Charm++: Parallel Programming with Message-Driven Objects. in *Parallel Programming using C++*, MIT Press (1996), pp. 175–213.
- [2] Menon, H., et al., Adaptive techniques for clustered N-body cosmological simulations. *Computational Astrophysics and Cosmology*, 2 (2015), DOI: 10.1186/s40668-015-0007-9.
- [3] Tremmel, M., E. Governato, M. Volonteri, and T. Quinn, Off the beaten path: a new approach to realistically model the orbital decay of supermassive black holes in galaxy formation simulations. *Monthly Notices of the Royal Astronomical Society*, 451 (2015), pp. 1868–1874.
- [4] Tremmel, M., et al., The Romulus cosmological simulations: a physical approach to the formation, dynamics and accretion models of SMBHs. *Monthly Notices of the Royal Astronomical Society*, 470 (2017), pp. 1121–1139.
- [5] Anderson, L., et al., The little galaxies that could (reionize the universe): predicting faint end slopes and escape fractions at $z > 4$. *Monthly Notices of the Royal Astronomical Society*, 468 (2017), pp. 4077–4092.

Roytershteyn, Vadim

- [1] Podesta, J. J., and V. Roytershteyn, The most intense electrical currents in the solar wind: Comparisons between single spacecraft measurements and plasma turbulence simulations. *J. Geophys. Res. Sp. Phys.* (2017), DOI: 10.1002/2017JA024074.

Schive, Hsi-Yu

- [1] ZuHone, J.A., A parameter space exploration of galaxy cluster mergers I: Gas mixing and the generation of cluster entropy. *The Astrophysical Journal*, 728 (2011), DOI: 10.1088/0004-637X/728/1/54.
- [2] Fryxell, B., et al., FLASH: An Adaptive Mesh Hydrodynamics Code for Modeling Astrophysical Thermonuclear Flashes. *The Astrophysical Journal Supplement Series*, 131:1 (1999), DOI: 10.1086/317361.
- [3] Schive, H.-Y., Y.-C. Tsai, and T. Chiueh, GAMER: A graphic processing unit-accelerated adaptive-mesh-refinement code for astrophysics. *The Astrophysical Journal Supplement Series*, 186 (2010), pp. 457–484.

- [4] Turk, M.J., et al., *yt*: A multi-code analysis toolkit for astrophysical simulation data. *The Astrophysical Journal Supplement Series*, 192 (2011), DOI: 10.1088/0067-0049/192/1/9.

Shapiro, Stuart

- [1] Schödel, R., et al., Closest Star Seen Orbiting the Supermassive Black Hole at the Centre of the Milky Way. *Nature*, 419 (2002), pp. 694–696.
- [2] Madau P., and M. Rees, Massive Black Holes as Population III Remnants. *Astrophys. J.*, 551 (2001), pp. L24–L30.
- [3] Heger, C. L., and S.E. Woosley, The Nucleosynthetic Signature of Population III. *Astrophys. J.*, 567 (2002), pp. 532–543.
- [4] Regan, J., P. Johansson, and J. Wise, The Direct Collapse of a Massive Black Hole Seed Under the Influence of an Anisotropic Lyman–Werner Source. *Astrophys. J.*, 795 (2014), pp. 137–154.
- [5] Sun, L., V. Paschalidis, M. Ruiz, and S. Shapiro, Magnetorotational Collapse of Supermassive Stars: Black Hole Formation, Gravitational Waves and Jets. ArXiv:1704.04502.
- [6] Etienne, Z.B., Y. Liu, and S. Shapiro, Relativistic magnetohydrodynamics in dynamical spacetimes: A new AMR implementation. *Phys. Rev. D*, 82 (2010), pp. 084031–084052.
- [7] Liu, Y. T., S. Shapiro, and B. Stephens, Magnetorotational collapse of very massive stars to black holes in full general relativity. *Phys. Rev. D*, 76 (2007), pp. 084017–084034.
- [8] Shibata, M., and S. Shapiro, Collapse of a Rotating Supermassive Star to a Supermassive Black Hole: Fully Relativistic Simulations. *Astrophys. J.*, 572 (2002), pp. L39–L44.
- [9] Bosnjak, Z., et al., The spectral catalogue of INTEGRAL gamma-ray bursts: results of the joint IBIS/SPI spectral analysis. *Astron. Astrophys.*, 561 (2014), A25–A40.

Tchekhovskoy, Alexander

- [1] Shiokawa, H., et al., General Relativistic Hydrodynamic Simulation of Accretion Flow from a Stellar Tidal Disruption. *The Astrophysical Journal*, 804 (2015), p. 85.
- [2] Sadowski, A., et al., Magnetohydrodynamical simulations of a deep tidal disruption in general relativity. *Monthly Notices of the Royal Astronomical Society*, 458 (2016), p. 4250.
- [3] Fragile, P. C., O.M. Blaes, P. Anninos, and J.D. Salmonson, Global General Relativistic Magnetohydrodynamic Simulation of a Tilted Black Hole Accretion Disk. *The Astrophysical Journal*, 668 (2007), p. 417.

Woodward, Paul R.

- [1] <https://www.youtube.com/watch?v=i8xX3tHkicE>, talk by Paul Woodward at the 2017 Blue Waters Symposium on the motivation for and design of our new PPMstar code, focusing on the Level 3 AMR and dynamic load balancing strategy and touching on GPU acceleration of the code.

- [2] www.lcse.umn.edu/PAID, slides from the talk by Paul Woodward at the 2017 Blue Waters Symposium on the techniques for implementing our new PPMstar code to run efficiently on GPU accelerators.
- [3] Lin, P.-H., and P. R. Woodward, Transforming the Multifluid PPM Algorithm to Run on GPUs. *Journal of Parallel and Distributed Computing*, 93 (2016), pp. 56–65.
- [4] Woodward, P. R., F. Herwig, and P.-H. Lin, Hydrodynamic Simulations of H Entrainment at the Top of He-Shell Flash Convection. *Astrophys. J.*, 798 (2015), p. 49.
- [5] Herwig, F., et al., Global Non-Spherical Oscillations in 3-D 4 π Simulations of the H-Instability Flash. *Astrophys. J. Letters*, 792 (2014), p. L3.
- [6] Jones, S., et al., Idealized hydrodynamic simulations of turbulent oxygen-burning shell convection in 4 π geometry. *Monthly Notices Royal Astron. Soc.*, 465 (2017), pp. 2991–3010.
- [7] Ritter, C., et al., Convective-reactive nucleosynthesis of K, Sc, Cl, and p-process isotopes in O-C shell mergers, submitted to *Monthly Notices of the Royal Astronomical Society* (2017).
- [8] Woodward, P. R., et al., Scaling the Multifluid PPM Code on Blue Waters and Intel MIC. *Proc. Extreme Scaling 2013 Workshop* (Boulder, Colo., August 15–16, 2013).

GEOSCIENCE

DiGirolamo, Larry

- [1] National Research Council, *Earth Science and Applications from Space: National Imperatives for the Next Decade and Beyond* (2007).
- [2] Wolfe, R., and H. Ramapriyan, Scaling the pipe: NASA EOS Terra data systems at 10. *Geoscience and Remote Sensing Symposium (IGARSS), 2010 IEEE International* (IEEE, Honolulu, Hawaii, July 25–30, 2010) pp. 1300–1303.
- [3] Di Girolamo, L., et al., The Terra Data Fusion Project. *MISR Data Users' Symposium*, Pasadena, Calif., Dec. 8–9, 2016.
- [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.*, 120:15 (2015), DOI: 10.1002/2015JD023256.
- [5] Zhao, G., et al., Regional changes in Earth's color and texture as observed from space over a 15-year period. *IEEE Trans. Geosci. Remote Sens.*, 54:7 (2016), DOI: 10.1109/TGRS.2016.2538723.
- [6] Hioki, S., et al., Retrieving degree of ice particle surface roughness in clouds from multi-angle measurements. *MISR Data Users' Symposium*, Pasadena, Calif., Dec. 8–9, 2016.
- [7] Bell, A., S. Hioki, Y. Wang, P. Yang, and L. Di Girolamo, Sensitivity analysis of observed reflectivity to ice particle surface roughness using MISR satellite observations, *American Geophysical Union 2016 Fall Meeting*, San Francisco, Calif., December 12–16, 2016.

- [8] Fu, D., L. Di Girolamo, L. Liang, and G. Zhao, The effect of cloud heterogeneity on microphysical properties through MISR-MODIS data fusion. *MISR Data Users' Symposium*, Pasadena, Calif., Dec. 8–9, 2016.

Garcia, Marcelo

- [1] Bulle, H., Untersuchungen ber die geschiebeableitung bei der spaltung von wasserlufen (in German). *V.D.I Verlag, Berlin, Germany, Tech. Rep.* (1926).
- [2] Dutta, S., D. Wang, P. Tassi, and M.H. Garcia, Three-Dimensional Numerical Modeling of the Bulle-Effect: the non-linear distribution of near-bed sediment at fluvial diversions. *Earth Surface Processes and Landforms* (2017), DOI: 10.1002/esp.4186.
- [2] Giosan, L., J. Syvitski, S. Constantinescu, and J. Day, Protect the world's deltas. *Nature*, 516 (2014), pp. 31–33.
- [3] Gaweesh, A., and E. Meselhe, Evaluation of Sediment Diversion Design Attributes and their Impact on the Capture Efficiency. *J. Hydraulic Eng.*, 142:5 (2016), pp. 1–15.
- [4] Dutta, S., et al., Application of computational fluid dynamic modelling to improve flow and grit transport in Terence J. O'Brien water reclamation plant, Chicago, Illinois. *J. Hydrul. Res.*, 52:6 (2014), pp. 759–774.
- [5] Lee, S., et al., Direct numerical simulation of transitional flow in stenosed carotid bifurcation. *J. Biomech.*, 41 (2008), pp. 2551–2561.
- [6] Fischer, P.F., J.W. Lottes, and S.G. Kerkemeier, Nek5000 webpage (2008), <http://nek5000.mcs.anl.gov>.
- [7] Deville, M., P. Fischer, and E. Mund, *High-Order Methods for Incompressible Fluid Flow*. Cambridge University Press (Cambridge, England, 2002).

Lasher-Trapp, Sonia

- [1] Bryan, G. H., and H. Morrison, Sensitivity of a Simulated Squall Line to Horizontal Resolution and Parameterization of Microphysics. *Mon. Wea. Rev.*, 140 (2012), pp. 202–225.
- [2] Mansell, E. R., C. L. Ziegler, and E. C. Bruning, Simulated Electrification of a Small Thunderstorm with Two-Moment Bulk Microphysics. *J. Atmos. Sci.*, 67 (2010), pp. 171–194.
- [3] Moser, D. H., and S. Lasher-Trapp, The Influence of Successive Thermals on Entrainment and Dilution in a Simulated Cumulus Congestus. *J. Atmos. Sci.*, 74 (2017), pp. 375–392.

Liu, Lijun

- [1] Zhong, S., et al., A benchmark study on mantle convection in a 3-D spherical shell using CitcomS. *Geochem. Geophys. Geosyst.*, 9:10 (2008), DOI: 10.1029/2008GC002048.
- [2] Leonard, T., and L. Liu, The Role of a Mantle Plume in the Formation of Yellowstone Volcanism. *Geophys. Res. Lett.*, 43 (2016), DOI: 10.1002/2015GL067131.

- [3] Hu, J., L. Liu, A. Hermosillo, and Q. Zhou, Simulation of Late Cenozoic South American Flat-Slab Subduction Using Geodynamic Models with Data Assimilation. *Earth Planet. Sci. Lett.*, 438 (2016), pp. 1–13.
- [4] Hu, J., and L. Liu, Abnormal Seismological and Magmatic Processes Controlled by the Tearing South American Flat Slabs. *Earth Planet. Sci. Lett.*, 450 (2016), pp. 40–51.
- [5] Hu, J., M. Faccenda, and L. Liu, Subduction-controlled Mantle Flow and Seismic Anisotropy in South America. *Earth Planet. Sci. Lett.*, 470 (2017), pp. 13–24.
- [6] Liu, L., and D. Hasterok, High-Resolution Lithosphere Viscosity and Dynamics Revealed by Magnetotelluric Tomography, *Science*, 353 (2016), pp. 1515–1519.

Maechling, Philip

- [1] Cornell, C.A., Engineering seismic risk analysis. *Bull. Seismol. Soc. Am.*, 58 (1968), pp. 1583–1606.
- [2] Petersen, M.D., et al., Documentation for the 2014 update of the United States' national seismic hazard maps. *U.S. Geological Survey Open-File Report 2014–1091* (2014), p. 243, DOI: 10.3133/ofr20141091.
- [3] Withers, K.B., K.B. Olsen, and S.M. Day, Memory Efficient Simulation of Frequency Dependent Q. *Bull. Seis. Soc. Am.*, 105 (2015), pp. 3129–3142.
- [4] 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.
- [5] Roten, D., K.B. Olsen, S.M. Day, and Y. Cui, Quantification of Fault-Zone Plasticity Effects with Spontaneous Rupture Simulations. *Pure Appl. Geophys.*, (2017), DOI: 10.1007/s00024-017-1466-5.
- [6] Roten, D., et al., High-frequency nonlinear earthquake simulations on petascale heterogeneous supercomputers. *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis* (IEEE Press, Piscataway, N.J., November 13–18, 2016), pp. 957–968.
- [7] Tabor, R., N. Khoshnevis, S. Azizzadeh-Roodpish, and M. Huda, Influence of the source, seismic velocity, and attenuation models on the validation of ground motion simulations. Poster presentation at *World Conference on Earthquake Engineering*, Santiago, Chile, January 9–13, 2017, number 4574.
- [8] Graves, R., et al., CyberShake: A physics-based seismic hazard model for Southern California. *Pure and Applied Geophysics*, 168:3 (2011), pp. 367–381.
- [9] Breuer, A., A. Heinecke, and Y. Cui, EDGE: Extreme Scale Fused Seismic Simulations with the Discontinuous Galerkin Method. *ISC High Performance'17*, Frankfurt, Germany, June 18–22, 2017.
- [10] National Science Foundation, *Investing in Science, Engineering, and Education for the Nation's Future Strategic Plan for 2014–2018*, March 2014.

Morin, Paul

- [1] Noh, M.-J., and I. M. Howat, Automated stereo-photogrammetric DEM generation at high latitudes: Surface Extraction with TIN-based Search-space Minimization (SETSM) validation and demonstration over glaciated regions. *GIScience Remote Sens.*, 52:2 (2015), pp. 1–20.
- [2] U.S. White House, Using new elevation data to explore the Arctic. Accessed from: <https://www.whitehouse.gov/blog/2015/09/02/using-new-elevation-data-explore-arctic>.

Rauber, Robert

- [1] Zarzycki, C. M., et al., Impact of ocean coupling strategy on extremes in high-resolution atmospheric simulations. *Geosciences Model Development*, 9 (2016), pp. 779–788, DOI: 10.5194/gmd-9-779-2016.
- [2] Zhao, M., I. M. Held, S.-J. Lin, and G. A. Vecchi, Simulations of global hurricane climatology, inter-annual variability, and response to global warming using a 50km resolution GCM. *Journal of Climate*, 22:24 (2009), DOI: 10.1175/2009JCLI3049.1.

Simpson, Jamesina

- [1] 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.
- [2] Taflove, A., and S. C. Hagness, *Computational Electrodynamics: The Finite-Difference Time-Domain (FDTD) Method*, Artech House, Third Edition (2005).
- [3] Pokhrel, S., et al., FDTD modeling of space weather hazards to electric power grids near ocean-continent boundaries, *JGR—Space Physics*, (submitted).

Srifer, Ryan

- [1] Wehner, M. F., et al., The effect of horizontal resolution on simulation quality in the Community Atmospheric Model, CAM5.1. *Journal of Advances in Modeling Earth Systems*, 6:4 (2014), DOI: 10.1002/2013MS000276.
- [2] Small, J. R., et al., A new synoptic scale resolving global climate simulation using the Community Earth System Model. *Journal of Advances in Modeling Earth Systems*, 6 (2014), DOI: 10.1002/2014MS000363.
- [3] Li, H., R. L. Srifer, and M. Goes, Modeled sensitivity of the Northwestern Pacific upper-ocean response to tropical cyclones in a fully coupled climate model with varying ocean grid resolution: Ocean response to TCs in coupled model. *Journal of Geophysical Research—Oceans*, 121 (2016), DOI: 10.1002/2015JC011226.
- [4] Li, H., and R. L. Srifer, Effects of ocean grid resolution on tropical cyclone-induced upper ocean responses using a global ocean general circulation model. *Journal of Geophysical Research—Oceans*, 121 (2016), DOI: 10.1002/2016JC011951.

- [5] Hurrell, J. W., et al., The Community Earth System Model: A Framework for Collaborative Research. *Bulletin of the American Meteorological Society*, 94:9, (2013), DOI: 10.1175/BAMS-D-12-00121.1.

West, Matthew

- [1] Dockery, D., and A. Pope, Epidemiology of acute health effects: summary of time-series studies. *Particles in our air: Concentrations and health effects* (1996), pp. 123–147.
- [2] Stocker, T. F., et al., Climate Change 2013: The Physical Science Basis (Technical Summary), *Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press (2013), pp. 33–115.
- [3] Riemer, N., M. West, R. A. Zaveri, and R. C. Easter, Simulating the evolution of soot mixing state with a particle-resolved aerosol model. *Journal of Geophysical Research: Atmospheres*, 114 (2009), DOI: 10.1029/2008JD011073.
- [4] Skamarock, W., et al., A Description of the Advanced Research WRF Version 3. NCAR Tech. Note NCAR/TN-475+STR (2008).
- [5] Environmental Protection Agency, *Emissions Modeling for the Final Mercury and Air Toxics Standards*. Technical Support Document, EPA-454/R-11-011 (2011).

Wuebbles, Donald J.

- [1] Lamarque, J.F., et al., CAM-chem: Description and evaluation of interactive atmospheric chemistry in the Community Earth System Model. *Geoscientific Model Development*, 5:2 (2012), p.369.
- [2] Yu, H., et al., Aerosols from overseas rival domestic emissions over North America. *Science*, 337:6094 (2012), pp. 566–569.
- [3] Lin, J., et al., China's international trade and air pollution in the United States. *Proceedings of the National Academy of Sciences of the U.S.A.*, 111:5 (2014), pp.1736–1741.

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 (2011), pp. 2234–2247.
- [2] 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.
- [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.

Valocchi, Albert J.

- [1] Liu, H., et al., Multiphase lattice Boltzmann simulations for porous media applications. *Computational Geosciences*, 20:4 (2014), pp. 1–29.

- [2] Tölke, J., S. Freudiger, and M. Krafczyk, An adaptive scheme using hierarchical grids for lattice Boltzmann multi-phase flow simulations. *Computers & Fluids*, 35:8–9 (2006), pp. 820–830.

PHYSICS & ENGINEERING**Ackerman, David**

- [1] Fredrickson, G.H. *The Equilibrium Theory of Inhomogeneous Polymers*. Oxford University Press (2006).
- [2] Balay, S., et al., *PETSc Users Manual*. Argonne National Laboratory (2015).
- [3] Feyel, F., A multilevel finite element method (FE2) to describe the response of highly non-linear structures using generalized continua. *Computer Methods in Applied Mechanics and Engineering*, 192 (2003), pp. 3233–3244.

Akono, Ange-Therese

- [1] Davidovits, J., Properties of geopolimer cements. *Proceedings of the First International Conference on Alkaline Cements and Concretes* (Kiev, Ukraine, October 11–14 1994), pp. 131–149.
- [2] Barbosa, V. F., and K. J. MacKenzie, Synthesis and thermal behaviour of potassium sialate geopolymers. *Mater. Lett.*, 57:9 (2003), pp. 1477–1482.
- [3] McLellan, B. C., et al., Costs and carbon emissions for geopolimer pastes in comparison to ordinary Portland cement. *Journal Cleaner Product.*, 19:9 (2011), pp. 1080–1090.
- [4] Moon, J., et al., Characterization of natural pozzolan-based geopolimeric binders, *Cem. Concr. Comp.*, 53 (2014), pp. 97–104.
- [5] Sukontasukkul, P., et al., Use of phase change material to improve thermal properties of lightweight geopolimer panel. *Mater. Struct.*, 49:11 (2016), pp. 4637–4645.
- [6] Davidovits, J., Environmentally driven geopolimer cement applications. In *Proceedings of 2002 Geopolimer Conference*, Melbourne, Australia, October, 2002.
- [7] Sakulich, A. R., Reinforced geopolimer composites for enhanced material greenness and durability, *Sustain. Cit. Soc.*, 1:4 (2011), pp. 195–210.
- [8] Davidovits, J., Geopolimer cement in *Geopolimer Institute, Technical Papers*, 21 (2013), pp. 1–11.

Allain, Jean Paul

- [1] Facsko, S. et al., Formation of Ordered Nanoscale Semiconductor Dots by Ion Sputtering. *Science*, 285:5433 (1999), pp. 1551–1553.
- [2] Lively, M. A., B. Holybee, M. Toriyama, and J. P. Allain, Massive-scale molecular dynamics of ion-irradiated III–V compound semiconductors at the onset of nanopatterning. *Nucl. Inst. Methods Phys. Res. B*, in press (2017).

- [3] El-Atwani, O. et al., Ion beam nanopatterning of III–V semiconductors: consistency of experimental and simulation trends within a chemistry-driven theory. *Sci. Rep.*, 5:February (2015), DOI:10.1038/srep18207
- [4] Holybee, B. J., Novel analytical method for investigating compositional driven ion-induced nanopatterning of multicomponent materials. University of Illinois at Urbana-Champaign (2016).
- [5] Plimpton, S., Fast Parallel Algorithms for Short-Range Molecular Dynamics. *J. Comput. Phys.*, 117:June 1994 (1995), pp. 1–19.
- [6] Bracht, H. et al., Large disparity between gallium and antimony self-diffusion in gallium antimonide. *Nature*, 408:6808 (2000), pp. 69–72.

Aluru, Narayana

- [1] Wells, D.B., M. Belkin, J. Comer, and A. Aksimentiev, Assessing Graphene Nanopores for Sequencing DNA. *Nano Letters*, 12 (2012), pp. 4117–4123.
- [2] Branton, D., et al., The Potential and Challenges of Nanopore Sequencing. *Nat. Biotech.*, 26 (2008), pp. 1146–1153.

Ceperley, David

- [1] McMahan, J. M., M. A. Morales, C. Pierleoni, and D. M. Ceperley, *Rev. Mod. Phys.*, 84 (2012), pp. 1607–1653 (2012).
- [2] Morales, M. A., C. Pierleoni, E. Schwegler, and D. M. Ceperley, *Proc. Nat. Acad. Science*, 107 (2010), pp. 12799–12803.
- [3] Knudson, M. D. et al., *Science*, 348 (2015), p. 1455.
- [4] Zaghoo, M. A. Salamat, and I. Silvera, *Phys. Rev. B*, 93 (2016), p. 155128.
- [5] Pierleoni, C., et al., Liquid-liquid phase transition in hydrogen by coupled electron-ion Monte Carlo simulations, *Proc. Nat. Acad. Science (US)*, 113 (2016), pp. 4953–4957.
- [6] Dias, R.P., and I. F. Silvera, *Science*, 355 (2017), pp. 715–718.
- [7] Rillo, G., M. A. Morales, D. M. Ceperley, and C. Pierleoni, Coupled Electron-Ion Monte Carlo simulations of hydrogen molecular crystals. *J. Chem. Phys.*, (submitted 2017).

Draayer, Jerry

- [1] Dytrych, T., et al., Collective Modes in Light Nuclei from First Principles. *Phys. Rev. Lett.*, 111:25 (2013), 252501.
- [2] Launey, K. D., et al., Symmetry-guided large-scale shell-model theory. *Prog. Part. Nucl. Phys.*, 89 (2016), pp. 101, DOI: 10.1016/j.pnpnp.2016.02.001.
- [3] *LSU3shell* code, Louisiana State University, available under the GNU *General Public License* at the *git* repository, <http://sourceforge.net/projects/lsu3shell>.
- [4] Dytrych, T., et al., Efficacy of the SU(3) scheme for *ab initio* large-scale calculations beyond the lightest nuclei. *Comput. Phys. Commun.*, 207:202 (2016).

- [5] Draayer, J. P., et al., Symmetry-adapted no-core shell model—the ab initio picture of nuclear collectivity. in *Emergent Phenomena in Atomic Nuclei from Large-scale Modeling: A Symmetry-guided Perspective*, in press (2017); ISBN: 978-981-3146-044.

Duan, Lian

- [1] Laufer, J., Some statistical properties of the pressure field radiated by a turbulent boundary layer. *Physics of Fluids*, 7:8 (1964), pp. 1191–1197.
- [2] Martín M. P., et al., A bandwidth-optimized weno scheme for the direct numerical simulation of compressible turbulence. *Journal of Computational Physics*, 220:1 (2006), pp. 270–289.
- [3] Taylor E. M., M. Wu, and M. P. Martín, Optimization of nonlinear error sources for weighted non-oscillatory methods in direct numerical simulations of compressible turbulence. *Journal of Computational Physics*, 223:1 (2006), pp. 384–397.
- [4] Williamson, J., Low-storage Runge-Kutta schemes. *Journal of Computational Physics*, 35:1 (1980), pp. 48–56.
- [5] Xu, S. and M. P. Martín, Assessment of inflow boundary conditions for compressible turbulent boundary layers. *Physics of Fluids*, 16:7 (2004), pp. 2623–2639.
- [6] E. Toubert and N. D. Sandham, “Oblique shock impinging on a turbulent boundary layer: Low-frequency mechanisms,” AIAA Paper 2008–4170, *38th Fluid Dynamics Conference and Exhibit*, Seattle, Wash., June 23–26, 2008.

Ertekin, Elif

- [1] E. Ertekin, Autogen V2, (2017), GitHub repository, <https://github.com/elifleaf/autogenv2>.
- [2] Dovesi, R., et al., CRYSTAL14: A program for the ab initio investigation of crystalline solids. *International Journal of Quantum Chemistry*, 114 (2014), pp. 1287–1317.
- [3] Wagner, L. K., M. Bajdich, and L. Mitas, QWalk: A quantum Monte Carlo program for electronic structure. *Journal of Computational Physics*, 228, 3390 (2009).

Fischer, Paul

- [1] Roche, K. R., et al., Benthic biofilm controls on fine particle dynamics in streams. *Water Resources Research*, 53:1 (2017), pp. 222–236.
- [2] Bulle, H., Untersuchungen über die geschiebeableitung bei der spaltung von wasserlufen (in German), V.D.I Verlag, Berlin, Germany, Tech. Rep., (1926).
- [3] Dutta, S., D. Wang, P. Tassi, and M.H. Garcia, Three-Dimensional Numerical Modeling of the Bulle-Effect: the non-linear distribution of near-bed sediment at fluvial diversions. *Earth Surface Processes and Landforms*, (2017), DOI:10.1002/esp.4186.
- [4] Tinoco, R., G. Coco, A laboratory study on sediment resuspension within arrays of rigid cylinders. *Adv. Water Res.*, 92 (2016), pp. 1–9.

- [4] Fischer, P.F., J.W. Lottes, and S.G. Kerkemeier, Nek5000 webpage, (2008), <http://nek5000.mcs.anl.gov>.
- [5] Deville, M., P. Fischer, and E. Mund, *High-Order Methods for Incompressible Fluid Flow*. Cambridge University Press, Cambridge, U.K., (2002).
- [6] Dutta, S., P. Fischer, M.H. Garcia, A fast semi-implicit Lagrangian particle tracking model for efficient computation of poly-disperse particle transport. In preparation for *Journal of Scientific Computing*.
- [7] Mittal, K., Spectral element mesh generation and improvement methods, M.S. thesis, University of Illinois at Urbana-Champaign, <http://hdl.handle.net/2142/95509>.

Garcia, Marcelo

- [1] Mier, J. M., Experimental analysis of flow and turbulence characteristics in oscillatory boundary layers from LDV measurements. Ph.D. Thesis, University of Illinois at Urbana-Champaign (2014).
- [2] Fischer, P. F., J. W. Lottes, and S. G. Kerkemeier, nek5000 web page, <http://nek5000.mcs.anl.gov>.
- [3] Cantero, M. I., S. Balachandar, and M. H. Garcia, An Eulerian–Eulerian model for gravity currents driven by inertial particles. *Intl J. Multiph. Flow*, 34:5 (2008), pp. 484–501.
- [4] Garcia, M.H., ed., Chapter 2: *Sediment Transport and Morphodynamics. In Sedimentation Engineering, ASCE Manuals and Reports on Engineering Practice No. 110*, ASCE/EWRI, Reston, Virginia (2008).
- [5] Pedocchi, F., M. I. Cantero, and M. H. Garcia, Turbulent kinetic energy balance of an oscillatory boundary layer in the transition to the fully turbulent regime. *Journal of Turbulence*, 12:32 (2011), pp. 1–27.
- [6] Carstensen, S., B. M. Sumer, and J. Fredsøe, Coherent structures in wave boundary layers, Part 1: Oscillatory motion. *J. Fluid. Mech.*, 646 (2010), pp. 169–206.
- [7] Carstensen, S., B. M. Sumer, and J. Fredsøe, A note on turbulent spots over a rough bed in wave boundary layer. *Physics of Fluids*, 24 (2010), DOI: 10.1063/1.4767536.
- [8] Pedocchi, F., and M. H. García, Friction coefficient for oscillatory flow: the rough-smooth turbulence transition. *J. of Hydraulic Research*, 47:4 (2009), pp. 438–444.
- [9] Niño, Y., and M. H. García, Experiments on particle-turbulence interactions in the near-wall region of an open channel flow: implications for sediment transport. *J. Fluid Mech.*, 326 (1996), pp. 85–319.

Gardoni, Paolo

- [1] Mazziari, I., M. Stupazzini, R. Guidotti, and C. Smerzini, SPEED: SPectral Elements in elastodynamics with discontinuous Galerkin: a non-conforming approach for 3D multi-scale problems. *International Journal for Numerical Methods in Engineering*, 95:12 (2013), pp. 991–1010.

- [2] Guidotti, R., et al., Numerical Study on the Role of Basin Geometry and Kinematic Seismic Source in 3D Ground Motion Simulation of the 22 February 2011 Mw 6.2 Christchurch Earthquake. *Seismological Research Letters*, 82:6 (2011), pp. 767–782.

Gazzola, Mattia

- [1] Gazzola, M., P. Chatelain, W. van Rees, and P. Koumoutsakos, Simulations of single and multiple swimmers with non-divergence free deforming geometries. *Journal of Computational Physics*, 230 (2011), pp. 7093–7114.
- [2] Gazzola, M., B. Hejazialhosseini, and P. Koumoutsakos, Reinforcement learning and wavelet adapted vortex methods for simulations of self-propelled swimmers. *SIAM Journal on Scientific Computing*, 36 (2014), pp. B622–B639.
- [3] D. Rossinelli, et al., MRAG-I2D: Multi-resolution adapted grids for remeshed vortex methods on multicore architectures. *Journal of Computational Physics*, 288 (2015), pp.1–18.
- [4] Sbalzarini, I., et al., PPM a highly efficient parallel particle–mesh library for the simulation of continuum systems. *Journal of Computational Physics*, 215 (2006), pp.566–588.
- [5] Gazzola, M., W. van Rees, and P. Koumoutsakos, C-start: optimal start of larval fish. *Journal of Fluid Mechanics*, 698 (2012), pp. 5–18.
- [6] van Rees, W., M. Gazzola, and P. Koumoutsakos, Optimal shapes for anguilliform swimmers at intermediate Reynolds numbers. *Journal of Fluid Mechanics*, 722 (2013), p. R3.
- [7] Gazzola, M., et al., Learning to school in the presence of hydrodynamic interactions. *Journal of Fluid Mechanics*, 789 (2015), pp. 726–749.

Huff, Kathryn

- [1] Lindsay, A., K. Huff, and A. Rykhlevskii, *ARFC/Moltres: Initial Moltres release*, Zenodo (2017), DOI: 10.5281/zenodo.801823.
- [2] Gaston, D., C. Newman, G. Hansen, and D. Lebrun-Grandie, MOOSE: A parallel computational framework for coupled systems of nonlinear equations. *Nuclear Engineering and Design*, 239 (2009), pp. 1768–1778.
- [3] Kirk, B. S., J.W. Peterson, R.H. Stogner, and G.F. Carey, libMesh: a C++ library for parallel adaptive mesh refinement/coarsening simulations. *Engineering with Computers*, 22 (2006), 237–254.
- [4] Satish Balay, et al., *PETSc Users Manual* (2015).
- [5] Turk, M. J., et al., yt: A Multi-Code Analysis Toolkit for Astrophysical Simulation Data. *The Astrophysical Journal Supplement Series*, 192: 9 (2011).

Ismail-Beigi, Sohrab

- [1] Jain, N., et al., OpenAtom: Scalable *Ab-Initio* Molecular Dynamics with Diverse Capabilities. In *High Performance Computing* (2016), pp. 139–158.

Jain , Prashant

- [1] Giannozzi, P., et al., QUANTUM ESPRESSO: A Modular and Open-Source Software Project for Quantum Simulations of Materials. *J. Phys. Condens. Matter*, 21:39 (2009), 395502.
- [2] Pamidighantam, S., et al., Community Science Exemplars in SEAGrid Science Gateway: Apache Airavata Based Implementation of Advanced Infrastructure. *Procedia Computer Science*, 80 (2016), pp. 1927–1939.

Johnsen, Eric

- [1] Brennen, C. E., *Cavitation and Bubble Dynamics*, Cambridge University Press (2013).
- [2] Arndt, R. E., Cavitation in fluid machinery and hydraulic structures, *Annu. Rev. Fluid Mech.*, 34 (1991), pp. 273–328.
- [3] Deplancke, T., et al., Outstanding cavitation erosion resistance of Ultra High Molecular Weight Polyethylene (UHMWPE) coatings. *Wear*, 328–329 (2015), pp. 301–308.
- [4] Franc, J. P., M. Riondet, A. Karimi, and G. L. Chahine, Material and velocity effects on cavitation erosion pitting. *Wear*, 274–275 (2012), pp. 248–259.
- [5] Tomita, Y., and A. Shima, Mechanisms of impulsive pressure generation and damage pit formation by bubble collapse. *J. Fluid Mech.*, 169 (1986), pp. 535–564.
- [6] Johnsen, E., and T. Colonius, Numerical simulations of non-spherical bubble collapse. *J. Fluid Mech.*, 29 (2009), pp. 231–261.
- [7] Alahyari Beig, S., and E. Johnsen, Maintaining interface equilibrium conditions in compressible multiphase flows using interface capturing. *J. Comput. Phys.*, 302 (2015), pp. 548–566.
- [8] Gottlieb, S., and C. W. Shu, Total variation diminishing Runge-Kutta schemes. *Math. Comput.*, 67 (1996), pp. 73–85.
- [9] Fuster, D., and T. Colonius, Modelling bubble clusters in compressible liquids. *J. Fluid Mech.*, 688 (2011), pp. 352–389.

Klimeck, Gerhard

- [1] World semiconductor trade statistics, accessed at <https://www.wsts.org/Teaser-Left/Historical-Billings-Report>.
- [2] Steiger, S., et al., NEMO5: A Parallel Multiscale Nanoelectronics Modeling Tool. *IEEE Transactions on Nanotechnology*, 10:6 (2011), pp. 1464–1474.
- [3] Fonseca, J. E., et al., Efficient and realistic device modeling from atomic detail to the nanoscale. *Journal of Comp. Electronics*, 12:4 (2013), pp. 592–600.
- [4] U.S. Patent and Trademark Office, accessed at <https://www.uspto.gov/web/offices/ac/ido/oeip/taf/data/semicon.htm>.

Mackenzie, Paul

- [1] Komijani, J., et. al, Decay constants f_B and f_{B_s} and quark masses m_b and m_c from HISQ simulations. in *Proceedings of the 34th International Symposium on Lattice Field Theory* (Lattice, Southampton, UK, July 24–30, 2016).

Militzer, Burkhard

- [1] Betti, R., editor. *Advancing the Science of High Energy Density Laboratory Plasmas* (2009).
- [2] Cook, L., Materials research for fusion energy. *Nat. Mater.*, 5 (2006), DOI: 10.1038/nmat1584.
- [3] Cauble, R., et al., Experiments Using Laser driven Shock-waves for EOS and Transport Measurements. *Contrib. Plasma Physics*, 41:2 (2001), pp. 239–242.
- [4] Van Horn, H. M., Dense astrophysical plasmas. *Science*, 252 (1991), pp. 384–389.
- [5] Militzer, B., Path Integral Monte Carlo Simulations of Hot Dense Hydrogen. Ph.D. thesis, University of Illinois at Urbana-Champaign (2000).
- [6] Ceperley, D.M., Path integral in the theory of condensed helium. *Rev. Mod. Phys.*, 67:2 (1995), DOI: 10.1103/RevModPhys.67.279.
- [7] Militzer, B., F. Soubiran, S. M. Wahl, and W. Hubbard, Understanding Jupiter's Interior. *J. Geophysical Research*, 121:9 (2016), DOI: 10.1002/2016JE005080.
- [8] Kresse, G., and J. Furthmüller, Efficient iterative schemes for *ab initio* total-energy calculations using a plane-wave basis set. *Phys. Rev. B*, 54:16 (1996), DOI: 10.1103/PhysRevB.54.11169.
- [9] Zhang, S., K. P. Driver, F. Soubiran, and B. Militzer, First-principles equation of state and shock compression predictions of warm dense hydrocarbons. *Phys. Rev. E*, 96:1 (2017), DOI:10.1103/PhysRevE.96.013204.

Rani, Sarma

- [1] Rani, S.L., R. Dhariwal, and D.L. Koch, A stochastic model for the relative motion of high Stokes number particles in isotropic turbulence. *Journal of Fluid Mechanics*, 756 (2014), pp. 870–902.
- [2] D. Pekurovsky, P3DFFT: A framework for parallel computations of Fourier transforms in three dimensions. *SIAM J. Sci. Comput.*, 34:4 (2012), pp. C192–C209.
- [3] Février, P., O. Simonin, and D. Legendre, Particle dispersion and preferential concentration dependence on turbulent Reynolds number from direct and large-eddy simulations of isotropic homogeneous turbulence. *Proceedings of the Fourth International Conference on Multiphase Flow* (New Orleans, La., May 27–June 1, 2001).
- [4] Zaichik, L. I., O. Simonin, and V.M. Alipchenkov, Two statistical models for predicting collision rates of inertial particles in homogeneous isotropic turbulence. *Physics of Fluids*, 15:10 (2003), pp. 2995–3005.

- [5] Dhariwal, R., S.L. Rani, and D.L. Koch, Stochastic Theory and Direct Numerical Simulations of the Relative Motion of High-Inertia Particle Pairs in Isotropic Turbulence. *Journal of Fluid Mechanics*, 813 (2017), pp. 205–249.

Riedl, Caroline

- [1] The DOE/NSF Nuclear Science Advisory Committee Working Group, *The Frontiers of Nuclear Science: A Long-Range Plan* (2007), arXiv:0809.3137.
- [2] Deshpande, A., et al., Electron Ion Collider: The Next QCD Frontier—Understanding the glue that binds us all. *European Physical Journal*, A:52 (2016), p. 268.
- [3] Riahi H., et al., FTS3: Quantitative Monitoring. *Journal of Physics: Conference Series* 664 (2015), DOI: 10.1088/1742-6596/664/6/062051.
- [4] Allison J., et al., Facilities and Methods: Geant4—A Simulation Toolkit. *Nuclear Instruments and Methods in Physics Research*, A:506 (2003), pp. 250–303.
- [5] Meyer, M., Blue Waters, A Petascale Computer Facility for the Reconstruction of CERN COM-PASS-II data. *22nd International SPIN Symposium*, Urbana, Ill., September 26–30, 2016.
- [6] Sjostrand, T., et al., An Introduction to PYTHIA 8.2. *Computer Physics Communications*, 191 (2015), pp. 159–177.

Schleife, Andre

- [1] Gajdos, M., et al., Linear optical properties in the projector-augmented wave methodology. *Phys. Rev. B*, 73:4 (2006), DOI: 10.1103/PhysRevB.73.045112.
- [2] Kresse, G., and D. Joubert, From ultrasoft pseudopotentials to the projector augmented-wave method. *Phys. Rev. B*, 59:3 (1999), DOI: 10.1103/PhysRevB.59.1758.
- [3] Kresse, G., and J. Furthmüller, Efficient iterative schemes for *ab initio* total-energy calculations using a plane-wave basis set. *Phys. Rev. B*, 54:16 (1996), DOI: 10.1103/PhysRevB.54.11169.
- [4] Rödl, C., F. Fuchs, J. Furthmüller, and F. Bechstedt, *Ab initio* theory of excitons and optical properties for spin-polarized systems: Application to antiferromagnetic MnO. *Phys. Rev. B*, 77:18 (2008), DOI: 10.1103/PhysRevB.77.184408.
- [5] Fuchs, F., C. Rödl, A. Schleife, and F. Bechstedt, Efficient $O(N^2)$ approach to solve the Bethe–Salpeter equation for excitonic bound states. *Phys. Rev. B*, 78:8 (2008), DOI: 10.1103/PhysRevB.78.085103.
- [6] Jun Lim, S., A. Schleife, and A. M. Smith, Optical determination of crystal phase in semiconductor nanocrystals. *Nat. Comm.*, 8 (2017), DOI: 10.1038/ncomms14849.

Taha, Ahmed

- [1] Idkaidek, A., V. Agarwal, and I. Jasiuk, Finite element simulation of Reference Point Indentation on bone. *Journal of the Mechanical Behavior of Biomedical Materials*, 65 (2017), pp. 574–583.

- [2] Idkaidek, A., and I. Jasiuk, Cortical bone fracture analysis using XFEM—case study. *International Journal for Numerical Methods in Biomedical Engineering*, 33: e2809 (2016), DOI: 0.1002/cnm.2809.

Thomas, Brian

- [1] *Steel Statistical Yearbook 2014* (World Steel Association, Brussels, Belgium, 2014).
- [2] Cho, S.-M., et al., Effect of Nozzle Clogging on Surface Flow and Vortex Formation in the Continuous Casting Mold. *Iron and Steel Technology*, 9 (2012), pp. 85–95.
- [3] Rackers, K., and B. G. Thomas, Clogging in Continuous Casting Nozzles. *Proc. of 78th Steelmaking Conference* (Iron and Steel Society, Warrendale, Pa., 1995), pp. 723–734.
- [4] Cho, S.-M., B. G. Thomas, H.-J. Lee, and S.-H. Kim, Effect of Nozzle Port Angle on Mold Surface Flow in Steel Slab Casting. *Iron and Steel Technology*, 3 (2017), pp. 76–84, reprinted from *Proceedings of AISTECH2016*.
- [5] Jin, K., B. G. Thomas, and X. Ruan, Modeling and Measurements of Multiphase Flow and Bubble Entrapment in Steel Continuous Casting. *Metallurgical and Materials Transactions B*, 47:1 (2016), pp. 548–565, DOI: 10.1007/s11663-015-0525-5.
- [6] Cho, S.-M., B. G. Thomas, and S.-H. Kim, Transient Two-Phase Flow in Slide-Gate Nozzle and Mold of Continuous Steel Slab Casting with and without Double-Ruler Electro-Magnetic Braking. *Metallurgical and Materials Transactions B*, 47:5 (2016), pp. 3080–3098, DOI: 10.1007/s11663-016-0752-4
- [7] Jin, K., S. P. Vanka, B. G. Thomas, and X. Ruan, Large Eddy Simulations of the Effects of Double-Ruler Electromagnetic Braking and Nozzle Submergence Depth on Molten Steel Flow in a Commercial Continuous Casting Mold. *Proc. TMS Annual Meeting 2016* (The Minerals, Metals and Materials Society, Nashville, Tenn., February 14–18, 2016), pp. 159–166.
- [8] 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. *Physics of Fluids*, 28:9 (2016), DOI: 10.1063/1.4961561.

Tinoco Lopez, Rafael

- [1] Koch, E. W., Beyond light: physical, geological, and geochemical parameters as possible submersed aquatic vegetation habitat requirements. *Estuaries*, 24:1 (2001), pp. 1–17.
- [2] Nepf, H. M., Hydrodynamics of vegetated channels. *Journal of Hydraulic Research*, 50:3 (2012), pp. 262–79.
- [3] Nepf, H. M., Flow and transport in regions with aquatic vegetation. *Ann. Rev. Fluid Mech.*, 44 (2012), pp. 123–142.
- [4] Tinoco, R., G. Coco, A laboratory study on sediment resuspension within arrays of rigid cylinders. *Adv. Water Res.*, 92 (2016), pp. 1–9.

- [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. Hydr. Res.*, 52:6 (2014), pp. 759–774.
- [6] Fischer, P. F., J. W. Lottes, and S. G. Kerkemeier, *Nek5000* webpage (2008), <http://nek5000.mcs.anl.gov>.
- [7] Deville, M., P. Fischer, and E. Mund, *High-Order Methods for Incompressible Fluid Flow*. (Cambridge University Press, Cambridge, England, 2002.)

Wagner, Lucas

- [1] Wagner, L.K., M. Bajdich, and L. Mitas, QWalk: A quantum Monte Carlo program for electronic structure. *Journal of Computational Physics*, 228 (2009), pp. 3390–3404.
- [2] Zheng, H., and L.K. Wagner, Computation of the correlated metal-insulator transition in vanadium dioxide from first principles. *Physical Review Letters*, 114 (2015), 176401.
- [3] Cao, J., et al., Strain engineering and one-dimensional organization of metal–insulator domains in single-crystal vanadium dioxide beams. *Nat Nano*, 4 (2009), pp. 732–737.

Yeung, Pui-kuen

- [1] Frisch, U. *Turbulence: The Legacy of A.N. Kolmogorov*. Cambridge University Press (1995).
- [2] Yeung, P. K., X. M. Zhai, and K. R. Sreenivasan, Extreme events in computational turbulence. *Proceedings of the National Academy of Sciences of the U.S.A.*, 112:41 (2015), pp. 12633–12638.
- [3] Gotoh, T., S. Hatanaka, and H. Miura, Spectral compact difference hybrid computation of passive scalar in isotropic turbulence. *Journal of Computational Physics*, 231:21 (2012), pp. 7398–7414.
- [4] Clay, M. P., D. Buaria, T. Gotoh, and P. K. Yeung, A dual communicator dual grid-resolution algorithm for petascale simulations of turbulent mixing at high Schmidt number. *Computer Physics Communications*, in press (2017), DOI: 10.1016/j.cpc.2017.06.009.

COMPUTER SCIENCE & ENGINEERING**Cox, Donna**

- [1] O'Shea, B., et al., Probing the Ultraviolet Luminosity Function of the Earliest Galaxies with the Renaissance Simulations. *The Astrophysical Journal Letters*, 807:1 (2015), DOI:10.1088/2041-8205/807/1/L12.
- [2] Chen, P., Norman, L., Xu, H., and Wise, J. Fully Coupled Simulation of Cosmic Reionization. III. Stochastic Early Reionization by the Smallest Galaxies. Submitted to *The Astrophysical Journal*. Preprint at arXiv:1705.00026v1 [astro-ph.CO].

- [3] Smith, B., et al., The First Population II Stars Formed in Externally Enriched Mini-Halos. *Monthly Notices of the Royal Astronomical Society*, 452:3 (2015), DOI:10.1093/mnras/stv1509.
- [4] Orf, L., et al., Evolution of a Long-Track Violent Tornado within a Simulated Supercell. *Bulletin of the American Meteorological Society*, 98 (2017), DOI:10.1175/BAMS-D-15-00073.1.
- [5] Canup, R., Forming a Moon with an earth-like composition via a giant impact. *Science* 338 (2012), pp. 1052–1055, DOI:10.1126/science.1226073.
- [6] Large Synoptic Survey Telescope Model provided by LSST project office/NSF/AURA.

Gropp, Bill

- [1] Eller, P. R., and W. Gropp, Scalable nonblocking preconditioned conjugate gradient methods. *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis* (IEEE Press, Salt Lake City, Utah, November 13–18, 2016), pp. 1–12.
- [2] Gropp, W., et al., Modeling MPI Communication Performance on SMP Nodes: Is it Time to Retire the Ping Pong Test? *Proceedings of the 23rd European MPI Users' Group Meeting* (ACM Digital Library, Edinburgh, United Kingdom, September 25–28, 2016), pp. 41–50.
- [3] Bienz, A., et al., Reducing Parallel Communication in Algebraic Multigrid through Sparsification. *SIAM Journal on Scientific Computing*, 38:5 (2016), pp. S332–S357.

Gürel, Levent

- [1] Ergül, Ö., and L. Gürel, A Hierarchical Partitioning Strategy for an Efficient Parallelization of the Multilevel Fast Multipole Algorithm. *IEEE Trans. Antennas Propagation*, 57:6 (2009), pp. 1740–1750.
- [2] Gürel, L. and Ö. Ergül, Hierarchical Parallelization of the Multilevel Fast Multipole Algorithm (MLFMA). *Proceedings of IEEE*, 101:2 (2013), pp. 332–341.
- [3] Hidayetoğlu, M., C. Pearson, W. C. Chew, L. Gürel, and W.-M. Hwu, Large Inverse-Scattering Solutions with DBIM on GPU-Enabled Supercomputers. *International Applied Computational Electromagnetics Society Symposium* (ACES, Florence, Italy, March 26–30, 2017).
- [4] Hidayetoğlu, M., and L. Gürel, Parallel Out-of-Core MLFMA on Distributed-Memory Computer Architectures. *Computational Electromagnetics International Workshop* (CEM'15, Izmir, Turkey, July 1–4, 2015), pp. 18–19.
- [5] Hwu, W.-M., et al., Thoughts on Massively-Parallel Heterogeneous Computing for Solving Large Problems. *CEM'17 Computing and Electromagnetics International Workshop* (CEM'17, Barcelona, Spain, June 21–24, 2017), submitted.
- [6] Pearson, C., et al., Comparative Performance Evaluation of Multi-GPU MLFMA Implementation for 2-D VIE Problems. *CEM'17 Computing and Electromagnetics International Workshop* (CEM'17, Barcelona, Spain, June 21–24, 2017), submitted.

- [7] Hidayetoğlu, M., C. Pearson, L. Gürel, W. C. Chew, and W.-M. Hwu, Scalable Parallel DBIM Solutions of Inverse-Scattering Problems. *CEM'17 Computing and Electromagnetics International Workshop* (CEM'17, Barcelona, Spain, June 21–24, 2017), submitted.

Nagi, Rakesh

- [1] Koopmans, T., and M. Beckmann, Assignment problems and the location of economic activities. *Econometrica: Journal of the Econometric Society*, 25:1 (1957), pp. 53–76.
- [2] Adams, W., et al., A level-2 reformulation-linearization technique bound for the Quadratic Assignment Problem. *European Journal of Operational Research*, 180:3 (2007), pp. 983–996.
- [3] Date, K., R. Nagi, GPU-accelerated Hungarian algorithms for the Linear Assignment Problem. *Parallel Computing*, 57 (2016), pp. 52–72, DOI: 10.1016/j.parco.2016.05.012.
- [4] Burkard, R., S. Karisch, and F. Rendl, QAPLIB - A Quadratic Assignment Problem Library. *Journal of Global Optimization*, 10:4 (1997), pp. 391–403.

Olson, Kevin

- [1] Reguly I.Z., et al., The OPS domain specific abstraction for multi-block structured grid computations. *Proceedings of the 4th International Workshop on Domain-Specific Languages and High-Level Frameworks for High Performance Computing* (IEEE Press, Piscataway, N.J., November 16–21, 2014), pp. 58–67.
- [2] G. R. Mudalige, G.R., et al., OP2: An active library framework for solving unstructured mesh-based applications on multi-core and many-core architectures. *Proceedings of the 2012 Innovative Parallel Computing (InPar)* (IEEE Press, Piscataway, N.J., May 13–14, 2012), pp. 1–12.

Warnow, Tandy

- [1] Nguyen, N., M. Nute, S. Mirarab, and T. Warnow, HIPPI: Highly accurate protein family classification with ensembles of HMMs. *BMC Genomics*, 17 (Supplement 10), 765 (2016), DOI: 10.1186/s12864-016-3097-0.
- [2] Altschul, S.F., et al., Basic local alignment search tool. *J. Mol. Biol.*, 215:3 (1990), pp. 403–10.
- [3] Eddy, S.R., A new generation of homology search tools based on probabilistic inference. *Genome Inform*, 23 (2009), pp. 205–211.
- [4] Bateman, A., et al., The Pfam protein families database. *Nucleic Acids Research*, 30 (2002), pp. 276–280.
- [5] Söding, J., Protein homology detection by HMM-HMM comparison. *Bioinformatics*, 21:7 (2005), pp. 951–960.

Xie, Tao

- [1] Hochreiter, S., and J. Schmidhuber, Long short-term memory. *Neural computation*, 9:8 (1997), pp. 1735–1780.

- [2] Krizhevsky, A., I. Sutskever, and G. E. Hinton, Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25 (2012), pp. 1097–1105.

BIOLOGY, CHEMISTRY, & HEALTH

Aksimentiev, Aleksei (1)

- [1] Göpfrich, K., et al., Ion channels made from a single membrane-spanning DNA duplex. *Nano Letters* 16:7 (2016), pp. 4665–4669.
- [2] Göpfrich, K. and Li, C., et al., Large-conductance transmembrane porin made from DNA origami. *ACS Nano* 10:9 (2016), pp. 8207–8214.
- [3] Langecker, M., et al., Synthetic lipid membrane channels formed by designed DNA nanostructures. *Science*, 338:6109 (2012), pp. 932–936.
- [4] Burns, J., E. Stulz, and S. Howorka, Self-assembled DNA nanopores that span lipid bilayers. *Nano Letters* 13:6 (2013), pp. 2351–2356.
- [5] Burns, J., et al., Lipid-Bilayer-Spanning DNA Nanopores with a Bifunctional Porphyrin Anchor. *Angewandte Chemie International Edition* 52:46 (2013), pp. 12069–12072.
- [6] Seifert, A., et al., Bilayer-spanning DNA nanopores with voltage-switching between open and closed state. *ACS Nano* 9:2 (2014), pp. 1117–1126.
- [7] Göpfrich, K., et al., DNA-tile structures induce ionic currents through lipid membranes. *Nano Letters* 15:5 (2015), pp. 3134–3138.
- [8] Burns, J., et al., A biomimetic DNA-based channel for the ligand-controlled transport of charged molecular cargo across a biological membrane. *Nature Nanotechnology* 11:2 (2016), pp. 152–156.
- [9] Kalé, L., et al., NAMD2: greater scalability for parallel molecular dynamics. *Journal of Computational Physics* 151:1 (1999), pp. 283–312.
- [10] Phillips, J., et al., Scalable molecular dynamics with NAMD. *Journal of Computational Chemistry* 26:16 (2005), pp. 1781–1802.
- [11] Yoo, J., and A. Aksimentiev, *In situ* structure and dynamics of DNA origami determined through molecular dynamics simulations. *Proceedings of the National Academy of Sciences* 110:50 (2013), pp. 20099–20104.

Aksimentiev, Aleksei (2)

- [1] Jovanovic-Talman, T., et al., Artificial nanopores that mimic the transport selectivity of the nuclear pore complex. *Nature* 457:7232 (2009), pp. 1023–1027.
- [2] Kowalczyk, S., et al., Single-molecule transport across an individual biomimetic nuclear pore complex. *Nature Nanotechnology* 6:7 (2011), pp. 433–438.
- [3] Bui, K., et al., Integrated structural analysis of the human nuclear pore complex scaffold. *Cell* 155:6 (2013), pp. 1233–1243.

- [4] Lin, D., et al., Architecture of the symmetric core of the nuclear pore. *Science* 352:6283 (2016), p. aaf1015.
- [5] Kale, L., et al., NAMD2: Greater scalability for parallel molecular dynamics. *Journal of Computational Physics* 151:1 (1999), pp. 283–312.
- [6] Phillips, J. C., et al., Scalable molecular dynamics with NAMD. *Journal of Computational Chemistry* 26:16 (2005), pp. 1781–1802.
- [7] Le Sage, V., and A. Mouland, Viral subversion of the nuclear pore complex. *Viruses* 5:8 (2013), pp. 2019–2042.
- [8] Köhler, A., and E. Hurt, Gene regulation by nucleoporins and links to cancer. *Molecular Cell* 38:1 (2010), pp. 6–15.
- [9] Jovičić, A., et al., Modifiers of C9orf72 dipeptide repeat toxicity connect nucleocytoplasmic transport defects to FTD/ALS. *Nature Neuroscience* 18:9 (2015), pp. 1226–1229.

Bauer, Gregory

- [1] Anisimov, V. M., et al., Optimization of the Coupled Cluster Implementation in NWChem on Petascale Parallel Architectures. *J. Chem. Theory Comput.*, 10:10 (2014), pp. 4307–4316.
- [2] Valiev, M., et al., NWChem: a comprehensive and scalable open-source solution for large scale molecular simulations. *Comput. Phys. Commun.*, 181:9 (2010), pp.1477–1489.

Caetano-Anolles, Gustavo

- [1] Feller, S. M., and M. Lewitzky, What's in a loop? *Cell Communication and Signaling*, 10:1 (2012), 31, DOI: 10.1186/1478-811X-10-31.
- [2] Wilke, C. O., Bringing molecules back into molecular evolution. *PLoS Computational Biology*, 8:6 (2012), e1002572, DOI:10.1371/journal.pcbi.1002572.
- [3] Liberles, D., S. A. Teichmann, I. Bahar, U. Bastolla, et al. The interface of protein structure, protein biophysics, and molecular evolution. *Protein Science: A Publication of the Protein Society*, 21:6 (2012), pp. 769–85. doi:10.1002/pro.2071.
- [4] 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.
- [5] Hensen, U., T. Meyer, J. Haas, R. Rex, G. Vriend., & H. Grubmüller. Exploring protein dynamics space: the dynasome as the missing link between protein structure and function. *PloS One*, 7:5 (2012), e33931.
- [6] Solé, R., and S. Valverde. Information theory of complex networks: on evolution and architectural constraints. In *Complex networks*, Lecture Notes in Physics book series volume 650 (2004), pp. 189–207.
- [7] Ravasz, E., et al., Hierarchical organization of modularity in metabolic networks. *Science*, 297 (2002), pp. 1551–1555, DOI: 10.1126/science.1073374.

- [8] Zooko's triangle (n.d.). Retrieved June 26, 2017, from https://en.wikipedia.org/wiki/Zooko%27s_triangle.
- [9] Mendes, C. L., et al., Deploying a Large Petascale System: The Blue Waters Experience. *Procedia Computer Science*, 29 (2014), pp. 198–209.

Cann, Isaac

- [1] Wu, M., et al., Genetic determinants of *in vivo* fitness and diet responsiveness in multiple human gut *Bacteroides*. *Science*, 350:6256 (2015), aac5992, pp. 1–8.
- [2] Dodd, D., and I.K.O. Cann, Enzymatic deconstruction of xylan for biofuel production. *Glob. Change Biol. Bioenergy*, 1 (2009), pp. 2–17.
- [3] Fontes, C. M. G. and H.J. Gilbert, Cellulosomes: highly efficient nanomachines designed to deconstruct plant cell wall complex carbohydrates. *Annu. Rev. Biochem.*, 79 (2010), pp. 655–681.
- [4] Borne, R., et al., Unraveling enzyme discrimination during cellulosome assembly independent of cohesin–dockerin affinity. *FEBS J.*, 280 (2013), pp. 5764–5779.
- [5] Ribeiro, J. V., et al., QwikMD—Integrative Molecular Dynamics Toolkit for Novices and Experts. *Sci. Rep.*, 6 (2016), p. 26536.
- [6] Bernardi, R. C., M.C.R. Melo, and K. Schulten, Enhanced sampling techniques in molecular dynamics simulations of biological systems. *Biochim. Biophys. Acta—Gen. Subj.*, 1850 (2015), pp. 872–877.
- [7] Melo, M., et al., GSAFold: a new application of GSA to protein structure prediction. *Proteins*, 80 (2012), pp. 2305–2310.

Carnevale, Vincenzo

- [1] Cao, E., M. Liao, Y. Cheng, D. Julius, TRPV1 structures in distinct conformations reveal activation mechanisms. *Nature*, 504 (2013), pp. 113–118.
- [2] Kaufmann, K.W., et al., Practically useful: what the Rosetta protein modeling suite can do for you. *Biochemistry (Moscow)*, 49:14 (2010), pp. 2987–2998.
- [3] Elokely, K., et al., Understanding TRPV1 activation by ligands: Insights from the binding modes of capsaicin and resiniferatoxin. *Proc Natl Acad Sci USA*, 113:2 (2016), pp. E137–E145.
- [4] Mackerell, A.D. Jr., M. Feig, and C.L. Brooks 3rd, Extending the treatment of backbone energetics in protein force fields: limitations of gas-phase quantum mechanics in reproducing protein conformational distributions in molecular dynamics simulations. *J Comput Chem.*, 25:11 (2004), pp. 1400–1415.
- [5] Jorgensen, W.L., et al., Comparison of simple potential functions for simulating liquid water. *J Chem Phys.*, 79:2 (1983), p. 926.
- [6] Liao, M., E. Cao, D. Julius, and Y. Cheng, Structure of the TRPV1 ion channel determined by electron cryo-microscopy. *Nature*, 504:7478 (2013), pp. 107–112.
- [7] Phillips, J.C., et al., Scalable molecular dynamics with NAMD. *J Comput Chem.*, 26:16 (2005), pp. 1781–1802.

- [8] Darden, T., D. York, and L. Pedersen, Particle mesh Ewald: An Nlog(N) method for Ewald sums in large systems. *J Chem Phys.*, 98:12 (1993), pp. 10089–10092.
- [9] Laio, A., et al., Assessing the Accuracy of Metadynamics. *J Phys Chem B*, 109:14 (2005), pp. 6714–6721.
- [10] Fiorin, G., M.L. Klein, and J. Hénin, Using collective variables to drive molecular dynamics simulations. *Mol Phys.*, 111:22–23 (2013), pp. 3345–3362.
- [11] Cooley, R.B., D.J. Arp, and P.A. Karplus, Evolutionary origin of a secondary structure: π -helices as cryptic but widespread insertional variations of α -helices that enhance protein functionality. *J Mol Biol.*, 404:2 (2010), pp. 232–246.
- [12] Susankova, K., et al., Contribution of the putative inner-pore region to the gating of the transient receptor potential vanilloid subtype 1 channel (TRPV1). *J Soc Neurosci.*, 27:28 (2007), pp. 7578–7585.
- [13] Salazar, H., et al., Structural determinants of gating in the TRPV1 channel. *Nat Struct Mol Biol.*, 16:7 (2009), pp. 704–710.
- [14] Dong, H., et al., Pore waters regulate ion permeation in a calcium release-activated calcium channel. *Proc Natl Acad Sci. USA*, 10:43 (2013), pp. 17332–17337.

Cheng, Chi-Hing

- [1] Siezen, R. J., J. A. Thomson, E. D. Kaplan, and G. B. Benedek, Human lens gamma-crystallins: isolation, identification, and characterization of the expressed gene products. *Proc Natl Acad Sci USA*, 84:17 (1987), pp. 6088–6092.
- [2] Papaleo, E., et al., Flexibility and enzymatic cold-adaptation: a comparative molecular dynamics investigation of the elastase family. *Biochim Biophys Acta*, 1764:8 (2006), pp. 1397–1406.
- [3] Grünberg, R., M. Nilges, and J. Leckner, Flexibility and conformational entropy in protein–protein binding. *Structure*, 14:4 (2006), pp. 683–693.
- [4] Swapna, L.S., R.M. Bhaskara, J. Sharma, and N. Srinivasan, Roles of residues in the interface of transient protein–protein complexes before complexation. *Sci Rep*, 2:334 (2012), DOI: 10.1038/srep00334.
- [5] Broide, M. L., et al., Binary-liquid phase separation of lens protein solutions. *Proc Natl Acad Sci USA*, 88:13 (1991), pp. 5660–5664.
- [6] Kiss, A.J., and C.H. Cheng, Molecular diversity and genomic organisation of the alpha, beta and gamma eye lens crystallins from the Antarctic toothfish *Dissostichus mawsoni*. *Comp Biochem Physiol Part D Genomics Proteomics*, 3:2 (2008) pp. 155–171.

Dill, Ken

- [1] Perez, A., J.L. MacCallum, and K. Dill, Accelerating molecular simulations of proteins using Bayesian inference on weak information. *Proc Natl Acad Sci USA*, 112 (2015), pp. 11846–11851.

- [2] MacCallum, J. L., A. Perez, and K. Dill, Determining protein structures by combining semireliable data with atomistic physical models by Bayesian inference. *Proc Natl Acad Sci USA*, 112 (2015), pp. 6985–6990.
- [3] Eastman, P., et al., OpenMM 4: A reusable, extensible, hardware independent library for high performance molecular simulation. *J. Chem. Theory Comput.*, 9 (2013), pp. 461–469.
- [4] Case, D. A., et al., *Amber12*, University of California, San Francisco (2012).
- [5] Skolnick, J. and H. Zhou, Why Is There a Glass Ceiling for Threading Based Protein Structure Prediction Methods? *J. Phys. Chem. B*, 121:15 (2016), DOI:10.1021/acs.jpcc.6b09517.

Freddolino, Peter

- [1] Berger, M. F., et al., Compact, universal DNA microarrays to comprehensively determine transcription-factor binding site specificities. *Nature Biotechnology*, 24 (2006), pp. 1429–1435.
- [2] Jolma, A., et al., DNA-binding specificities of human transcription factors. *Cell*, 152 (2013), pp. 327–339.
- [3] Liu, L. A. and P. Bradley, Atomistic modeling of protein-DNA interaction specificity: progress and applications. *Curr Opin Struct Biol*, 22 (2012), pp. 397–405.
- [4] Seeliger, D., et al., Towards computational specificity screening of DNA-binding proteins. *Nucleic Acids Research*, 39 (2011), pp. 8281–8290.
- [5] Khabiri, M. and P. L. Freddolino, Deficiencies in Molecular Dynamics Simulation-Based Prediction of Protein-DNA Binding Free Energy Landscapes. *J. Phys. Chem. B*, 121:20 (2017), pp. 5151–5161.
- [6] Goette, M. and H. Grubmüller, Accuracy and convergence of free energy differences calculated from nonequilibrium switching processes. *Journal of Computational Chemistry*, 30 (2009), pp. 447–456.

Hammes-Schiffer, Sharon

- [1] Kennis, J. T. M., and T. Mathes, Molecular eyes: Proteins that transform light into biological information. *Interface Focus*, 3:5 (2013), DOI: 10.1098/rsfs.2013.0005.
- [2] Webb, S. P., T. Iordanov, and S. Hammes-Schiffer, Multiconfigurational nuclear-electronic orbital approach: Incorporation of nuclear quantum effects in electronic structure calculations. *J. Chem. Phys.* 177:9 (2002), pp. 4106–4118.
- [3] Sirjoosingh, A., M. V. Pak, K. R. Brorsen, and S. Hammes-Schiffer, Quantum treatment of protons with the reduced explicitly correlated Hartree-Fock approach. *J. Chem. Phys.* 144:21 (2015), pp. 214107, 1–16.
- [4] Dickson, B. M., H. Huang, and C. B. Post, Unrestrained Computation of Free Energy along a Path. *J. Phys. Chem. B*, 116:36 (2012), pp. 11046–11055.
- [5] Torrie, G. M., and J. P. Valleau, Nonphysical sampling distributions in Monte Carlo free-energy estimation: Umbrella sampling. *J. Comput. Phys.* 23:2 (1977), pp. 187–199.

- [6] Brooks, B. R., et al., CHARMM: The biomolecular simulation program. *J. Comput. Chem.* 30:10 (2009), pp. 1545–1614.
- [7] Casida, M. E., and M. Huix-Rotllant, Progress in Time-Dependent Density-Functional Theory. *Annu. Rev. Phys. Chem.* 63 (2012), pp. 287–323.
- [8] Frisch, M. J., et al., Gaussian 09, Revision E.01, Gaussian, Inc., Wallingford, Conn. (2009).
- [9] Goyal, P., and S. Hammes-Schiffer, Role of active site conformational changes in photocycle activation of the AppA BLUF photoreceptor. *Proc. Natl. Acad. Sci. USA*, 114:7 (2017), pp. 1480–1485.

Hirata, So

- [1] Hirata, S., X. He, M. R. Hermes, and S. Y. Willow, Second-order many-body perturbation theory: An eternal frontier. *J. Phys. Chem. A*, 118 (2014), pp. 655–672.
- [2] Willow, S.Y., K. S. Kim, and S. Hirata, Stochastic evaluation of second-order many-body perturbation energies. *J. Chem. Phys.*, 137 (2012), p. 204122.
- [3] Willow, S. Y., K. S. Kim, and S. Hirata, Stochastic evaluation of second-order Dyson self-energies. *J. Chem. Phys.*, 138 (2013), p. 164111.
- [4] Willow, S., J. Zhang, E. F. Valeev, and S. Hirata, Communication: Stochastic evaluation of explicitly correlated second-order many-body perturbation energy. *J. Chem. Phys.*, 140 (2014), p. 031101.
- [5] Johnson, C. M., et al., Monte Carlo explicitly correlated second-order many-body perturbation theory. *J. Chem. Phys.*, 145 (2016), p. 154115.
- [6] Johnson, C. M., A. E. Doran, S. Ten-no, and S. Hirata, Monte Carlo explicitly correlated second-order many-body Green's function theory. In preparation (2017).
- [7] Doran, A. E. and S. Hirata, Monte Carlo MP2 on many graphical processing units. *J. Chem. Theory Comput.*, 12 (2016), pp. 4821–4832.
- [8] Willow, S. Y., K. S. Kim, and S. Hirata, Brueckner–Goldstone quantum Monte Carlo for correlation energies and quasiparticle energy bands of one-dimensional solids. *Phys. Rev. B*, 90 (2014), p. 201110(R).
- [9] Willow, S. Y. and S. Hirata, Stochastic, real-space, imaginary-time evaluation of third-order Feynman–Goldstone diagrams. *J. Chem. Phys.*, 140 (2014), p. 024111.
- [10] Willow, S. Y., M. R. Hermes, K. S. Kim, and S. Hirata, Convergence acceleration of parallel Monte Carlo second-order many-body perturbation calculations using redundant walkers. *J. Chem. Theory Comput.*, 9 (2013), pp. 4396–4402.

Luthey-Schulten, Zaida

- [1] Gillespie, D., Exact stochastic simulation of coupled chemical reactions. *J. Phys. Chem.*, 81 (1977), pp. 2340–2361.
- [2] Alfonsi, A., et al., Adaptive simulation of hybrid stochastic and deterministic models for biochemical systems. *ESAIM Proceeding*, 14 (2005), pp. 1–13.

- [3] Cao, Y., D.T. Gillespie, and L.R. Petzold, The slow-scale stochastic simulation algorithm. *J. Chem. Phys.*, 122:1 (2005), p. 014116.
- [4] Jahnke, T., and M. Kreim, *SIAM Multiscale Model. Simul.*, 10 (2012), pp. 1119–1131.
- [5] Roberts, E., J.E. Stone, and Z. Luthey-Schulten, Lattice-Microbes: high-performance stochastic simulations of the reaction-diffusion master equation. *J. Comput. Chem.*, 34:3 (2013), pp. 245–255.
- [6] Peterson, J.R., M.J. Hallock, J.A. Cole, and Z. Luthey-Schulten, A Problem Solving Environment for Stochastic Biological Simulations. *PyHPC'13: Proceedings of the 3rd Workshop on Python for High-Performance and Scientific Computing* (2013), DOI: 10.13140/2.1.3207.7440.

Masud, Arif

- [1] Masud, A., and J. Kwack, A stabilized mixed finite element method for the incompressible shear-rate dependent non-Newtonian fluids: Variational Multiscale framework and consistent linearization. *Computer Methods in Applied Mechanics and Engineering*, 200 (2011), pp. 577–596.
- [2] Kwack, J., and A. Masud, A stabilized mixed finite element method for shear-rate dependent non-Newtonian fluids: 3D benchmark problems and application to blood flow in bifurcating arteries. *Computational Mechanics*, 53:4 (2014), pp. 751–776.

Moradi, Mahmoud

- [1] Paulsen, I. T., and R. A. Skurray, The POT family of transport proteins. *Trends Biochem. Sci.*, 19:10 (1994), p. 404.
- [2] Ito, K., et al., Analysing the substrate multispecificity of a proton-coupled oligopeptide transporter using a dipeptide library. *Nat. Commun.*, 4 (2013), p. 2502.
- [3] Newstead, S., Molecular insights into proton coupled peptide transport in the PTR family of oligopeptide transporters. *Biochim. Biophys. Acta, Gen. Subj.*, 1850 (2015), pp. 488–499.
- [4] Smith, D. E., B. Clemençon, and M. A. Hediger, Proton-coupled oligopeptide transporter family SLC15: Physiological, pharmacological and pathological implications. *Mol. Aspects Med.*, 34 (2013), pp. 323–336.
- [5] Rautio, J., et al., Prodrugs: design and clinical applications. *Nat. Rev. Drug Discovery*, 7 (2008), pp. 255–270.
- [6] Solcan, N., et al., Alternating access mechanism in the POT family of oligopeptide transporters. *EMBO J.*, 31 (2012), pp. 3411–3421.
- [7] Newstead, S., et al., Crystal structure of a prokaryotic homologue of the mammalian oligopeptide-proton symporters, PepT1 and PepT2. *EMBO J.*, 30 (2011), pp. 417–426.
- [8] Guettou, F., et al., Structural insights into substrate recognition in proton-dependent oligopeptide transporters. *EMBO Rep.*, 14 (2013), pp. 804–810.

- [9] Lyons, J. A., et al., Structural basis for polyspecificity in the POT family of proton-coupled oligopeptide transporters. *EMBO Rep.*, 15 (2014), pp. 886–893.
- [10] Doki, S., et al., Structural basis for dynamic mechanism of proton-coupled symport by the peptide transporter POT. *Proc. Natl. Acad. Sci. U. S. A.*, 110 (2013), pp. 11343–11348.
- [11] Immadisetty, K., J. Hettige, and M. Moradi, What can and cannot be learned from molecular dynamics simulations of bacterial proton-coupled oligopeptide transporter GkPOT? *J. Phys. Chem. B.*, 121:15 (2017), pp. 3644–3656.
- [12] 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.
- [13] Moradi, M., and E. Tajkhorshid, Mechanistic picture for conformational transition of a membrane transporter at atomic resolution. *Proc. Natl. Acad. Sci. U. S. A.*, 110 (2013), pp. 18916–18921.
- [14] Moradi, M., G. Enkavi, and E. Tajkhorshid, Atomic-level characterization of transport cycle thermodynamics in the glycerol-3-phosphate:phosphate antiporter. *Nat. Commun.*, 6 (2015), p. 8393.
- [15] Fakharzadeh, A., and M. Moradi, Effective Riemannian diffusion model for conformational dynamics of biomolecular systems. *J. Phys. Chem. Lett.*, 7 (2016), pp. 4980–4987.

Pande, Vijay

- [1] Nelson, L. S. and J. Perrone, Curbing the opioid epidemic in the United States: the risk evaluation and mitigation strategy (REMS). *JAMA*, 308 (2012), pp. 457–458.
- [2] Weiss, D. R., et al., Conformation guides molecular efficacy in docking screens of activated β -2 adrenergic G protein coupled receptor. *ACS Chem. Biol.*, 8 (2013), pp. 1018–1026.
- [3] Rasmussen, S. G. F., et al., Structure of a nanobody-stabilized active state of the β -2 adrenoceptor. *Nature*, 469 (2011), pp. 175–180.
- [4] Manglik, A., et al., Crystal structure of the μ -opioid receptor bound to a morphinan antagonist. *Nature*, 485 (2012), pp. 321–326.
- [5] Huang, W., et al., Structural insights into μ -opioid receptor activation. *Nature*, 524 (2015), pp. 315–21.
- [6] Manglik, A., et al., Structure-based discovery of opioid analgesics with reduced side effects. *Nature*, 537:7619 (2016), pp. 185–190.

Roux, Benoit

- [1] Ostmeyer, J., et al., Recovery from Slow Inactivation in K⁺ Channels Controlled by Water Molecules. *Nature*, 501 (2013), pp. 121–124.
- [2] Weingarth, M., et al., Quantitative Analysis of the Water Occupancy around the Selectivity Filter of a K⁽⁺⁾ Channel in Different Gating Modes. *J Am Chem Soc.*, 136:5 (2014), pp. 2000–2007.

- [3] Lueck, J.D., et al., Atomic mutagenesis in ion channels with engineered stoichiometry. *eLife*, 5 (2016), DOI: 10.7554/eLife.18976.

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, DOI: 10.1038/nature12162.
- [2] Perilla, J.R., et al., All-atom molecular dynamics of virus capsids as drug targets. *J. Phys. Chem. Lett.*, 7:10 (2016), pp. 1836–1844, DOI: 10.1021/acs.jpcclett.6b00517.
- [3] Phillips, J.C., et al., Scalable molecular dynamics with NAMD. *J. Comp. Chem.*, 26:16 (2005), pp. 1781–1802, DOI: 10.1002/jcc.20289.

Soltész, Ivan

- [1] Carnevale, N., and M. Hines, *The NEURON Book*. Cambridge University Press (2006).
- [2] Hines, M., and N. Carnevale, Computer modeling methods for neurons. in *The Handbook of Brain Theory and Neural Networks* (1995), p. 226–230.
- [3] Hines, M., and N. Carnevale, Translating network models to parallel hardware in NEURON. *J. Neurosci. Methods*, 169 (2008), pp. 425–455.

Sotomayor, Marcos

- [1] Takeichi, M., Cadherins: a molecular family important in selective cell-cell adhesion. *Annu. Rev. Biochem.*, 59 (1990), pp. 237–252.
- [2] Brasch, J., O.J. Harrison, B. Honig, and L. Shapiro, Thinking outside the cell: how cadherins drive adhesion. *Trends Cell Biol.*, 22 (2012), pp. 299–310.
- [3] Harrison, O. J., et al., The extracellular architecture of adherens junctions revealed by crystal structures of type I cadherins. *Structure*, 19 (2011), pp. 244–256.
- [4] Cailliez, F. and R. Lavery, Dynamics and stability of E-cadherin dimers. *Biophys. J.*, 91 (2006), pp. 3964–3971.
- [5] Sotomayor, M., and K. Schulten, The allosteric role of the Ca²⁺ switch in adhesion and elasticity of C-cadherin. *Biophys. J.*, 94 (2008), pp. 4621–4633.
- [6] Oroz, J., et al., Nanomechanics of the cadherin ectodomain: “canalization” by Ca²⁺ binding results in a new mechanical element. *J. Biol. Chem.*, 286 (2011), pp. 9405–9418.
- [7] Manibog, K., H. Li, S. Rakshit, and S. Sivasankar, Resolving the molecular mechanism of cadherin catch bond formation. *Nat. Commun.*, 5 (2014), 3941.
- [8] Humphrey, W., A. Dalke, and K. Schulten, VMD: Visual molecular dynamics. *J. Mol. Graph.*, 14 (1996), pp. 33–38.
- [9] Phillips, J. C., et al., Scalable molecular dynamics with NAMD. *J. Comput. Chem.*, 26 (2005), pp. 1781–1802.
- [10] Sotomayor, M. and K. Schulten, Single-molecule experiments *in vitro* and *in silico*. *Science*, 316 (2007), pp. 1144–1148.

Srinivasan, Ashok

- [1] Helbing, D., and P. Molnar, Social force model for pedestrian dynamics. *Physical review E*, 51:5 (1995), DOI: 10.1103/PhysRevE.51.4282.
- [2] Nikolić, M., M. Bierlaire, B. Farooq, and M. de Laparent, Probabilistic speed–density relationship for pedestrian traffic. *Transportation Research Part B: Methodological*, 89 (2016), pp. 58–81.
- [3] Namilae S., et al., Self-propelled pedestrian dynamics model: Application to passenger movement and infection propagation in airplanes, *Physica A*, 465 (2017) pp. 248–260.

Sutton, Brad

- [1] McGarry, M. D., et al., Multiresolution MR Elastography Using Nonlinear Inversion. *Med Phys*, 39:10 (2012), pp. 6388–6396, DOI: 10.1118/1.4754649.
- [2] McGarry, M. D., et al., Including Spatial Information in Nonlinear Inversion MR Elastography Using Soft Prior Regularization. *IEEE T Med Imaging*, 32:10 (2013), pp. 1901–1909, DOI: 10.1109/TMI.2013.2268978.
- [3] Schwardt, H., et al., Aerobic fitness, hippocampal viscoelasticity, and relational memory performance. *NeuroImage*, 153 (2017), pp. 179–188, DOI: 10.1016/j.neuroimage.2017.03.061.
- [4] Sandroff, B. M., C. L. Johnson, and R. W. Motl, Exercise training effects on memory and hippocampal viscoelasticity in multiple sclerosis: a novel application of magnetic resonance elastography. *Neuroradiology*, 59:1 (2017), pp. 61–67, DOI: 10.1007/s00234-016-1767-x.

Tagkopoulos, Ilias

- [1] Huang, T., H. Gong, C. Yang, and Z. He, Protein-Lasso: A Lasso regression approach to protein inference problem in shotgun proteomics. *Computational Biology and Chemistry*, 43 (2013), pp. 46–54.
- [2] Ahrné, E., L. Molzahn, T. Glatter, and A. Schmidt, Critical assessment of proteome-wide label-free absolute abundance estimation strategies. *Proteomics*, 13:17 (2013), pp. 2567–2578.
- [3] Klimek, J., et al., The standard protein mix database: a diverse data set to assist in the production of improved peptide and protein identification software tools. *The Journal of Proteome Research*, 7:1 (2007), pp. 96–103.
- [4] Ramakrishnan, S., and C. Vogel, Gold Standard of Protein Expression in Yeast (data set), 2009, http://www.marcottelab.org/MSdata/gold_yeast.
- [5] Brunner, E., et al., A high-quality catalog of the *Drosophila melanogaster* proteome. *Nature Biotechnology*, 25:5 (2007), pp. 576–583.
- [6] Penalva, L., and C. Vogel, Human–Orbitrap–Daoy medulloblastoma wildtype, cell lysate (data set), 2009, http://www.marcottelab.org/MSdata/Data_05/.
- [7] Penalva, L., and C. Vogel, Human–Orbitrap–T293 embryonic kidney cells, overexpressing GFP, cell lysate and pellet (data set), 2009, http://www.marcottelab.org/MSdata/Data_07/.

Voth, Gregory

- [1] Briggs, J. A. G., et al., Structure and assembly of immature HIV. *PNAS*, 106 (2009), pp. 11090–11095.
- [2] Schur, F. K. M., et al., An atomic model of HIV-1 capsid-SP1 reveals structures regulating assembly and maturation. *Science*, 353 (2016), pp. 506–508.
- [3] Wagner, J. M., et al., Crystal structure of an HIV assembly and maturation switch. *eLife*, 5 (2016), e17063, DOI: 10.7554/eLife.17063.002.
- [4] Manley, S., et al., High-density mapping of single-molecule trajectories with photoactivated localization microscopy. *Nat. Methods*, 5 (2008), pp. 155–157.
- [5] Grime, J., et al., Coarse-grained simulation reveals key features of HIV-1 capsid self-assembly. *Nat. Commun.*, 7 (2016), p. 11568.

SOCIAL SCIENCE, ECONOMICS, & HUMANITIES**Cai, Yongyang**

- [1] Cai, Y., K. L. Judd, and T. S. Lontzek, The social cost of carbon with economic and climate risks. R&R in *Journal of Political Economy* (2015), arXiv preprint arXiv:1504.06909.
- [2] Cai, Y., K. L. Judd, and T. S. Lontzek, Climate policies for climate targets (2017), working paper.
- [3] Cai, Y., W. Brock, and A. Xepapadeas, Climate change economics and heat transport across the globe: Spatial-DSICE. *ASSA 2017*, Chicago, Ill., January 6–8, 2017.
- [4] Baldwin, E. C., Y. Cai, and K. Z. Kuralbayeva, Build now, regret later? Infrastructure and Climate Policy (2017), working paper.
- [5] Nordhaus, W. D., *A Question of Balance: Weighing the Options on Global Warming Policies*. Yale University Press, 2008.
- [6] IWG, *Social Cost of Carbon for Regulatory Impact Analysis under Executive Order 12866*. United States Government (2010).
- [7] Cai, Y., et al., The effect of climate and technological uncertainty in crop yields on the optimal path of global land use (2016), working paper.

GRADUATE FELLOWS**Agee, Elizabeth**

- [1] Allen, C. D., et al., A global overview of drought and heat-induced tree mortality reveals emerging climate change risks for forests. *Forest Ecology and Management*, 259:4 (2010), pp. 660–684.
- [2] Phillips, O. L., et al., Drought Sensitivity of the Amazon Rainforest. *Science*, 323:5919 (2009), pp. 1344–1347.
- [3] Leitner, D., et al., A dynamic root system growth model based on L-Systems, *Plant Soil*, 332:1–2 (2010), pp. 177–192.
- [4] Hammond, G. E., et al., Evaluating the performance of parallel subsurface simulators: An illustrative example with PFLOT-TRAN: Evaluating the Parallel Performance of PFLOT-TRAN. *Water Resources Research*, 50:1 (2014), pp. 208–228.
- [5] Couvreur, V., et al., A simple three-dimensional macroscopic root water uptake model based on the hydraulic architecture approach. *Hydrology and Earth System Sciences*, 16:8 (2012), pp. 2957–2971.

Brooks, Maureen

- [1] Laffoley, D., et al. *The Protection and Management of the Sargasso Sea: The golden floating rainforest of the Atlantic Ocean: Summary Science and Supporting Evidence Case*, Sargasso Sea Alliance (2011), pp. 1–44.
- [2] Smetacek, V., and A. Zingone, Green and golden seaweed tides on the rise., *Nature* 504 (2013), pp. 84–88.
- [3] Hanisak, M. D., and M. A. Samuel, Growth rates in culture of several species of *Sargassum* from Florida, USA. *Hydrobiologia*, 151–152 (1987), pp. 399–404.
- [4] Chassignet E., et al., US GODAE: Global Ocean Prediction with the HYbrid Coordinate Ocean Model (HYCOM). *Oceanography* 22:2 (2009), pp. 64–75.
- [5] Fennel, K., R. Hetland, Y. Feng, and S. DiMarco, A coupled physical-biological model of the Northern Gulf of Mexico shelf: model description, validation and analysis of phytoplankton variability. *Biogeosciences*, 8:7 (2011), pp. 1881–1899.
- [6] Gower, J. F. R., and S. A. King, Distribution of floating *Sargassum* in the Gulf of Mexico and the Atlantic Ocean mapped using MERIS. *Int. J. Remote Sens.*, 32:7 (2011), pp. 1917–1929.

Butsky, Iryna

- [1] Werk, J. K., et al., The COS-Halos Survey: An Empirical Description of Metal-line Absorption in the Low-redshift Circumgalactic Medium. *The Astrophysical Journal*, 204:17 (2013), DOI: 10.1088/0067-0049/204/2/17.
- [2] Peebles, M. S., et al., A Budget and Accounting of Metals at $z \sim 0$: Results from the COS-Halos Survey. *The Astrophysical Journal*, 786:54 (2014), DOI:10.1088/0004-637X/786/1/54.
- [3] Tremonti, C. A., et al., The Origin of the Mass-Metallicity Relation: Insights from 53,000 Star-forming Galaxies in the Sloan Digital Sky Survey. *The Astrophysical Journal*, 613:898 (2004), pp. 898–913.
- [4] Enßlin, T. A., C. Pfrommer, V. Springel, and M. Jubelgas, Cosmic ray physics in calculations of cosmological structure formation. *Astronomy and Astrophysics*, 473:41 (2007), pp. 41–57.
- [5] Kim, J.-H., et al., The AGORA High-resolution Galaxy Simulations Comparison Project. *The Astrophysical Journal*, 210:14 (2014), DOI:10.3847/1538-4357/833/2/202.

Calhoun, Jon

- [1] Snir, M., et al., Addressing failures in exascale computing. *International Journal of High Performance Computing Applications*, 28:2 (2014), pp.127–171.
- [2] Sridharan, V., et al., Memory errors in modern systems: The good, the bad, and the ugly. *Proceedings of the Twentieth International Conference on Architectural Support for Programming Languages and Operating Systems* (ACM, Istanbul, Turkey, March 14–18, 2015), pp. 297–310.
- [3] Calhoun, J., L. Olson, and M. Snir, FlipIt: An LLVM Based Fault Injector for HPC. *Euro-Par 2014: Parallel Processing Workshops* (Springer International Publishing, Porto, Portugal, August 25–29, 2014), pp. 547–558.
- [4] <https://mantevo.org>.

Drake, Justin

- [1] Wright, P.E., and H. J. Dyson, Intrinsically unstructured proteins: re-assessing the protein structure-function paradigm. *J. Mol. Biol.*, 293 (1999), pp. 321–331.

- [2] Dunker, A.K., et al., Intrinsic protein disorder in complete genomes. *Genome Inform. Workshop Genome Inform.*, 11 (2000), pp. 161–171.
- [3] Flock, T., R. J. Weatheritt, N. S. Latysheva, and M. M. Babu, Controlling entropy to tune the functions of intrinsically disordered regions. *Curr. Opin. Struct. Biol.* 26 (2014), pp. 62–72.
- [4] Jakob, U., R. Kriwacki, and V. N. Uversky, Conditionally and Transiently Disordered Proteins: Awakening Cryptic Disorder to Regulate Protein Function. *Chem. Rev.*, 114 (2014), pp. 6779–6805.
- [5] Phillips, J.C., et al., Scalable molecular dynamics with NAMD. *J. Comput. Chem.*, 26 (2005), pp. 1781–1802.
- [6] Case, D.A., et al., Amber 12 (2012).
- [7] Levy, R.M., A. R. Srinivasan, W. K. Olson, and J. A. McCammon, Quasi-harmonic method for studying very low frequency modes in proteins. *Biopolymers*, 23 (1984), pp. 1099–1112.
- [8] Harpole, K.W., and K. A. Sharp, Calculation of Configurational Entropy with a Boltzmann–Quasiharmonic Model: The Origin of High-Affinity Protein–Ligand Binding. *J. Phys. Chem. B*, 115 (2011), pp. 9461–9472.
- [9] Killian, B.J., J. Y. Kravitz, and M. K. Gilson, Extraction of Configurational Entropy from Molecular Simulations via an Expansion Approximation. *J. Chem. Phys.*, 127 (2007), DOI:10.1063/1.2746329.
- [10] Thompson, J.B., H. G. Hansma, P. K. Hansma, and K. W. Plaxco, The Backbone Conformational Entropy of Protein Folding: Experimental Measures from Atomic Force Microscopy. *J. Mol. Biol.*, 322 (2002), pp. 645–652.

Guang, August

- [1] Eddy, S. R., HMMER: Profile hidden Markov models for biological sequence analysis (2001).
- [2] Jombart, T., et al., Bayesian reconstruction of disease outbreaks by combining epidemiologic and genomic data. *PLoS Comput. Biol.* 10 (2014), DOI:10.1371/journal.pcbi.1003457.
- [3] Mallo, D., L. De Oliveira Martins, and D. Posada, SimPhy: Phylogenomic Simulation of Gene, Locus, and Species Trees. *Syst. Biol.* 65 (2016), pp. 334–344.
- [4] Spielman, S. J. and C.O. Wilke, Pyvolve: A Flexible Python Module for Simulating Sequences along Phylogenies. *PLoS One* 10 (2015), DOI:10.1371/journal.pone.0139047.
- [5] Huang, W., L. Li, J.R. Myers, and G.T. Marth, ART: a next-generation sequencing read simulator. *Bioinformatics* 28 (2012), pp. 593–594.

Hime, Paul

- [1] Baum, D.A., et al., Statistical evidence for common ancestry: Application to primates. *Evolution*, 70:6 (2016), pp. 1354–1363.
- [2] Stamatakis, A., RAxML version 8: a tool for phylogenetic analysis and post-analysis of large phylogenies. *Bioinformatics*, 30:9 (2014), pp. 1312–1313.
- [3] Akaike, H., A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19:6 (1974), pp. 716–723.
- [4] Mirarab, S., and T. Warnow, ASTRAL-II: coalescent-based species tree estimation with many hundreds of taxa and thousands of genes. *Bioinformatics*, 31:12 (2015), pp. i44–i52.
- [5] Yang, Z., PAML 4: phylogenetic analysis by maximum likelihood. *Molecular Biology and Evolution*, 24:8 (2007), pp. 1586–1591.

Howard, Michael

- [1] Malevanets, A. and R. Kapral, Mesoscopic model for solvent dynamics. *J. Chem. Phys.*, 110:17 (1999), pp. 8605–8613.
- [2] Malevanets, A. and J. M. Yeomans, Dynamics of short polymer chains in solution. *Europhys. Lett.*, 52:2 (2000), pp. 231–237.
- [3] Gompper, G., T. Ihle, D. M. Kroll, and R. G. Winkler, Multi-Particle Collision Dynamics: A Particle-Based Mesoscale Simulation Approach to the Hydrodynamics of Complex Fluids. in *Advanced Computer Simulation Approaches for Soft Matter Sciences III*, Advances in Polymer Science, 221 (2009), pp. 1–87.
- [4] Anderson, J. A., C. D. Lorenz, and A. Travesset, General purpose molecular dynamics simulations fully implemented on graphics processing units. *J. Comput. Phys.*, 227 (2008), pp. 5342–5359.
- [5] Glaser, J., et al., Strong scaling of general-purpose molecular dynamics simulations on GPUs. *Comput. Phys. Commun.*, 192 (2015), pp. 97–107.

Kirby, Andrew

- [1] C. Burstedde, L.C. Wilcox, and O. Ghattas, p4est: Scalable Algorithms for Parallel Adaptive Mesh Refinement on Forests of Octrees. *SIAM Journal on Scientific Computing* 33:3 (2011), pp. 1103–1133.

Kokkila Schumacher, Sara

- [1] Spudich, J.L., C.-S. Yang, K.-H. Jung, and E. N. Spudich, Retinylidene Proteins: Structures and Functions from Archaea to Humans. *Annu. Rev. Cell. Dev. Biol.* 16 (2000), pp. 365–392.
- [2] Nielsen, I. B., et al., Absorption Studies of Neutral Retinal Schiff Base Chromophores. *J. Phys. Chem. A*, 110:46 (2006), pp. 12592–12596.
- [3] Geertsen, J., M. Rittby, and R. J. Bartlett, The equation-of-motion coupled-cluster method: Excitation energies of Be and CO. *Chem. Phys. Lett.*, 164 (1989), pp.57–62.
- [4] O. Christiansen, H. Koch, and P. Jørgensen, The Second-order Approximate Coupled Cluster Singles and Doubles Model CC2. *Chem. Phys. Lett.*, 243 (1995), pp. 409–418.
- [5] Plötner, J., D. J. Tozer, and A. Dreuw, Dependence of Excited State Potential Energy Surfaces on the Spatial Overlap of the Kohn–Sham Orbitals and the Amount of Nonlocal Hartree–Fock Exchange in Time-Dependent Density Functional Theory. *J. Chem. Theory Comput.*, 6:8 (2010), pp. 2315–2324.
- [6] Hohenstein, E. G., R. M. Parrish, and T. J. Martínez, Tensor hypercontraction density fitting. I. Quartic Scaling Second- and Third-order Møller–Plesset Perturbation Theory. *J. Chem. Phys.*, 137 (2012), pp. 044103:1–10.
- [7] Hohenstein, E. G., S. I. L. Kokkila, R. M. Parrish, and T. J. Martínez, Quartic Scaling Second-order Approximate Coupled Cluster Singles and Doubles via Tensor Hypercontraction: THC-CC2. *J. Chem. Phys.*, 138 (2013), pp. 124111:1–10.
- [8] Hohenstein, E. G., S. I. L. Kokkila, R. M. Parrish, and T. J. Martínez, Tensor Hypercontraction Equation-of-Motion Second-order Approximate Coupled Cluster: Electronic Excitation Energies in $O(N^4)$ Time. *J. Phys. Chem. B*, 117 (2013), pp. 12972–12978.
- [9] Ufimtsev, I. S. and T. J. Martínez, Quantum Chemistry on Graphical Processing Units. 3. Analytical Energy Gradients and First Principles Molecular Dynamics. *J. Chem. Theo Comp.*, 5:10 (2009), pp. 2619–2628.

Richers, Sherwood

- [1] Bethe, H., Supernova mechanisms. *Review of Modern Physics*, 62:4 (1990), pp. 801–866.
- [2] Nagakura, H., K. Sumiyoshi, and S. Yamada, Three-dimensional Boltzmann Hydro Code for Core Collapse in Massive Stars. I. Special Relativistic Treatments. *The Astrophysical Journal Supplement Series*, 214:2 (2014), 16–35.
- [3] Smit, J. M., L.J. van den Horn, and S.A. Bludman, Closure in flux-limited neutrino diffusion and two-moment transport. *Astronomy and Astrophysics*, 356 (2000), pp.559–569.
- [4] O'Connor, Evan, An Open-source Neutrino Radiation Hydrodynamics Code for Core-collapse Supernovae. *The Astrophysical Journal Supplement Series*, 219:2 (2015), 24–44.
- [5] Ott, C., et al., New open-source approaches to the modeling of stellar collapse and the formation of black holes. *Astrophysics and Space Science*, 336:1 (2011), pp.151–156.
- [6] Moesta, P., et al., GRHydro: a new open-source general-relativistic magnetohydrodynamics code for the Einstein Toolkit. *Classical and Quantum Gravity*, 31:1 (2015), DOI: 10.1088/0264-9381/31/1/015005.

Seyler, Sean

- [1] Wijesinghe, H. S., and N.G. Hadjiconstantinou, Discussion of Hybrid Atomistic–Continuum Methods for Multiscale Hydrodynamics. *Int. J. Multiscale Com.*, 2 (2004), pp. 189–202.
- [2] Mohamed, K.M., and A.A. Mohamad, A review of the development of hybrid atomistic–continuum methods for dense fluids. *Nanofluidics*, 8 (2010), pp. 283–302.
- [3] Landau, L.D., and E.M. Lifshitz, *Fluid Mechanics*. Vol. 6 (1959), pp. 523–529.
- [4] Grad, H., On the kinetic theory of rarefied gases. *Comm. Pure Appl. Math.*, 2 (1949), pp. 331–407.
- [5] Seyler, C.E., and M.R. Martin, Relaxation model for extended magnetohydrodynamics: Comparison to magnetohydrodynamics for dense Z-pinches. *Phys. Plasmas*, 18 (2011), DOI:10.1063/1.3543799.
- [6] Zhao, X., Y. Yang, and C.E. Seyler, A positivity-preserving semi-implicit discontinuous Galerkin scheme for solving extended magnetohydrodynamics equations. *J. Comput. Phys.*, 278 (2014), pp. 400–415.
- [7] Riedel, C., et al., The heat released during catalytic turnover enhances the diffusion of an enzyme. *Nature*, 517 (1999), pp. 227–230.
- [8] De Fabritiis, G., M. Serrano, R. Delgado-Buscalioni, and P.V. Coveney, Fluctuating hydrodynamic modeling of fluids at the nanoscale. *Phys. Rev. E*, 75 (2007): DOI:10.1103/PhysRevE.75.026307.

Stenz, Ronald

- [1] Bryan, G. H., and J. M. Fritsch, A benchmark simulation for moist nonhydrostatic numerical models. *Mon. Wea. Rev.*, 130 (2002), pp. 2917–2928.
- [2] Gilmore M. S., R. Davies-Jones, J. Straka, E. Rasmussen, and R. Wilhelmson, Centrifugal precipitation transport in tornadic supercells: An algorithm consistent for use with bulk microphysics schemes. *14th Intl. Conf. on Clouds and Precipitation* (Bologna, Italy), Intl. Commission Clouds Precip. (2004), pp. 1654–1656.

- [3] Thompson, R. L., R. Edwards, J. A. Hart, K. L. Elmore, and P. M. Markowski, Close proximity soundings within supercell environments obtained from the rapid update cycle. *Wea. Forecasting*, 18 (2003), pp. 1243–1261.

Teich, Erin

- [1] Available at <https://glotzerlab.engin.umich.edu/hoomd-blue>.
- [2] Berthier, L., and G. Biroli, Theoretical perspective on the glass transition and amorphous materials. *Rev. Mod. Phys.*, 83 (2011), pp. 587–645.
- [3] Royall, C. P., and S. R. Williams, The role of local structure in dynamical arrest. *Phys. Reports*, 560 (2015), pp. 1–75.
- [4] Alder, B. J., and T.E. Wainwright, Phase Transition for a Hard Sphere System. *J. Chem. Phys.*, 27 (1957), pp. 1208–1209.
- [5] Onsager, L., The Effects of Shape on the Interaction of Colloidal Particles. *Ann. N.Y. Acad. Sci.*, 51 (1949), pp. 627–659.
- [6] Asakura, S., and F. Oosawa, Interaction between particles suspended in solutions of macromolecules. *J. Polym. Sci.*, 33 (1958), pp. 183–192.
- [7] Damasceno, P.F., M. Engel, and S.C. Glotzer, Predictive Self-Assembly of Polyhedra into Complex Structures. *Science*, 337 (2012), pp. 453–457.
- [8] van Anders, G., et al., Understanding shape entropy through local dense packing. *Proc. Natl. Acad. Sci. U.S.A.*, 111 (2014), pp. E4812–E4821.
- [9] Anderson, J. A., M. E. Irrgang, and S. C. Glotzer, Scalable Metropolis Monte Carlo for simulation of hard shapes. *Comput. Phys. Commun.*, 204 (2016), pp. 21–30.

Totorica, Sam

- [1] Zweibel, E. G., and M. Yamada, Magnetic Reconnection in Astrophysical and Laboratory Plasmas. *Annu. Rev. Astron. Astrophys.*, 47 (2009), 291–332.
- [2] Fonseca, R. A., et al., Exploiting multi-scale parallelism for large scale numerical modelling of laser wakefield accelerators. *Plasma Phys. Controlled Fusion*, 55 (2013), DOI:10.1088/0741-3335/55/12/124011.
- [3] Abel, T., O. Hahn, and R. Kaehler, Tracing the dark matter sheet in phase space. *Mon. Not. R. Astron. Soc.*, 427 (2012), 61–76.

INDEX

A

Ackerman, David 102
Agee, Elizabeth 258
Akono, Ange-Therese 104
Aksimentiev, Aleksei 192, 194
Allain, Jean Paul 106
Aluru, Narayana 108

B

Balsara, Dinshaw 18
Bauer, Gregory 196
Bernholc, Jerzy 110
Bodony, Daniel 112
Brooks, Maureen 260
Butsky, Iryna 262

C

Caetano-Anolles, Gustavo 198
Cai, Yongyang 252
Calhoun, Jon 264
Cann, Isaac 200
Carnevale, Vincenzo 202
Ceperley, David 114
Cheatham, Thomas 204
Cheng, Chi-Hing 206
Chew, Huck Beng 116
Cho, Wendy 254
Clark, Marin 64
Corcoran, Jennifer 66
Cox, Donna 170
Curreli, Davide 118

D

Di Girolamo, Larry 68
Di Matteo, Tiziana 20
Dill, Ken 208
Draayer, Jerry 120
Drake, Justin 266
Duan, Lian 122

E

Elbanna, Ahmed 210
Elghobashi, Said 124
Ertekin, Elif 126

F

Fischer, Paul 128
Freddolino, Peter 212

G

Garcia, Marcelo 70, 130
Gardoni, Paolo 132
Gazzola, Mattia 134
Gropp, Bill 172
Guang, August 68
Gürel, Levent 174

H

Hammes-Schiffer, Sharon 214
Hansen, Lars 252
Hime, Paul 268
Hirata, So 216
Holder, Gilbert 22
Howard, Michael 270
Howat, Ian 72
Huerta, Eliu 24, 26
Huff, Kathryn 136

I

Ismail-Beigi, Sohrab 138
Iyer, Ravishankar 176

J

Jain, Prashant 140
Johnsen, Eric 142
Jones, Alexandra 272
Jones, Thomas 28
Jordan, Thomas 78

K

Kasson, Peter 218
Kirby, Andrew 274
Klimeck, Gerhard 144
Kokkila Schumacher, Sara 276

L

Lasher-Trapp, Sonia 74
Lentz, Eric 30

Levin, Deborah 32, 146
Liu, Lijun 76
Liu, Yi-Hsin 34
Luthey-Schulten, Zaida 220

M

Mackenzie, Paul 148
Maechling, Phil 78
Makri, Nancy 222
Masud, Arif 224
Mazzuca, James 226
Militzer, Burkhard 150
Moradi, Mahmoud 228
Mori, Warren 36
Morin, Paul 80

N

Nagi, Rakesh 178
Noble, Scott 38
Norman, Michael 40

O

O'Shea, Brian 42
Olson, Luke 180
Olson, Kevin 188
Orf, Leigh G. 82
Ott, Christian 44

P

Pande, Vijay 230
Pogorelov, Nikolai 46

Q

Quinn, Thomas 48

R

Rani, Sarma 152
Rauber, Robert 84
Reams, Larissa 278
Richers, Sherwood 280
Riedl, Caroline 154
Roux, Benoit 232
Roytershteyn, Vadim 50

S

Schive, Hsi-Yu 52
Schleife, Andre 156
Schulten, Klaus 234

Seyler, Sean 282
Shapiro, Stuart 54
Shukla, Diwakar 236
Simpson, Jamesina 86
Solomonik, Edgar 182
Soltesz, Ivan 238
Sotomayor, Marcos 240
Srinivasan, Ashok 242
Sriver, Ryan 88
Stenz, Ronald 284
Sutton, Brad 244

T

Tagkopoulos, Ilias 246
Taha, Ahmed 158
Tajkhorshid, Emad 234
Tchekhovskoy, Alexander 56
Teich, Erin 286
Thomas, Brian 160
Tinoco Lopez, Rafael 163
Toth, Gabor 58
Totorica, Sam 288
Trapp, Robert 90

U

Um, Junshik 92

V

Valocchi, Albert J. 94
Voth, Gregory 248

W


Wagner, Lucas 164
Warnow, Tandy 184
West, Matthew 96
Wilhelmson, Robert 82
Woodward, Paul R. 60
Wuebbles, Donald J. 84, 98

X

Xie, Tao 186

Y

Yeung, Pui-kuen 166





NCSA

I ILLINOIS

Blue Waters is supported by the
National Science Foundation

